LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-701390

Advanced Simulation and Computing FY17 Implementation Plan, Version 0

D. Wade, T. Hoang, M. McCoy, B. Archer, B. Hendrickson

August 26, 2016

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
FY17 IMPLEMENTATION PLAN

Version 0

August 29, 2016

NA-ASC-127R-16-Rev.0-IP

LLNL-TR-701390

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*	*8/29/16*	*Initial release*

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.

ASC is now 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 counter nuclear 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 (NPR) Report* 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 counter nuclear 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 FY17. The Initial IP effective October 1, 2016, should be consistent with: 1) the Department's Base Table when operating under a

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

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	FY2017 Target	Endpoint Target
Advanced Simulation and Computing Program	Reduced Reliance on Calibration	60%	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 FY17 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 FY17 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

[2] U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2017 Stockpile Stewardship and Management Plan—Biennial Plan Summary*, Report to Congress, March 2016.

NNSA's Science Program, which provides the experimental data that informs development of new models used in simulation codes.

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. The efficiency of the integrated design codes (IDCs) 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 three focus areas for investment. The Next-Generation Code Development and Application (CDA) product 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 (ASD) product 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 Interagency Co-Design (ICD) product will support the U.S. National Strategic Computing Initiative (NSCI) in increasing the capacity and capability of an enduring national HPC ecosystem via inter-agency collaborations with other U.S. federal agencies. These projects will utilize the HPC systems, software, and applications to address the sponsor agencies' mission needs.

The ATDM subprogram tackles the most critical subset of issues occurring during the upcoming period of disruptive change in HPC architectures in order to continue the current level of support to the DSW mission. It aligns and coordinates activities with the DOE Office of Science to achieve the goals of the department's Exascale Computing Initiative (ECI).

The **Computation Systems and Software Environment (CSSE)** subprogram procures and integrates the computing systems needed for weapons simulations. Since requirements of the ASC codes drive the need to achieve its predictive capability goals, the ASC Program must continue to invest in and consequently influence the evolution of computational environment. 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. In this subprogram, ASC will embark on research investigations of Beyond Moore's Law to include quantum, neuromorphic, and non-complementary metal-oxide-semiconductor (CMOS)-based computing techniques.

The **Facility Operations and User Support (FOUS)** subprogram provides the facilities and services required to run nuclear weapons simulations. Facility operations include 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 FY17 major activities for the ASC Program include:

- Complete work on defining early initial conditions for boost; begin updating the IDC with results.
- Bring to Production Readiness the CT systems and complete the Trinity-Knights Landing (Phase 2) system for the tri-labs' production computing environment in early FY17 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 and future 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, under the ECI aegis, to advance technologies that will eventually enable procurement and utility of an exascale-class HPC platform.

The ASC budgetary increase in FY17 provides the additional funding required to facilitate transitioning the integrated codes to work efficiently on emerging high-performance computers and develop next-generation codes. 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 IDCs will open design options for subsystem components for future LEPs.

The ASC computing capabilities are the key for integrating mechanisms across the nuclear weapons program through the IDCs. The assessment of the nation's stockpile requires high-fidelity physical models. The IDCs support design studies, maintenance analyses, the annual assessment reports, LEPs, SFIs, and weapons dismantlement activities. The IDCs 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 IDCs allow detailed simulations of nuclear weapons performance assessment, without the need for underground nuclear testing. Since the 1992 nuclear weapons testing moratorium, the IDCs 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 IDCs 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 IDCs 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 exascale computing capabilities in the future. The ASC Program's approach to advancing HPC technologies is scoped to contribute to the foundation for an exascale computing capability for the nation. The relatively new ATDM subprogram consolidates the investments Congress directed in FY14 for exascale into a unified effort that ASC will manage in order to tackle the technical challenges as it carries out its stockpile stewardship mission. Since the technical problems facing the program today are similar but at lower scale, investments in ATDM will 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 Work Authorization s (WAs) and obligated via formal contract modification.

Program/Operational Control Level	President's FY2017 Budget Request	FY2017 CR Operating Target	FY2017 Enacted/Full-Year CR	Difference between Request versus Enacted
ASC	$663,184M	$	$	

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 FY17. A more comprehensive list (including milestone description and evaluation/exit criteria) is included in the individual subprogram detail in the appendices.

Table V-1. ASC Level 2 Milestone for FY17[3]

Sub-Program	ID#	Milestone Title	Complete Date	Site
IC	TBD	Evaluate Full-System Methodologies	6/30/17	LLNL
IC	TBD	Evaluate Three-Dimensional Dendrites in an ASC Code	9/30/17	LLNL
IC	TBD	Initial Adaptive Mesh Refinement Demonstrations	9/30/17	LLNL
IC	TBD	Monte Carlo Applications Toolkit Physics and Performance Upgrades	9/30/17	LANL
IC	TBD	Eulerian Applications Project Code Upgrades for Modeling Inertial Confinement Fusion Experiments	6/30/17	LANL

[3] Factors such as FY17 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	Demonstrate Thread Scalability within ARIA on Both Sides of Trinity	9/30/17	SNL
IC	TBD	Demonstration of Effective Coupling between the SIERRA Low-Mach Module FUEGO and the RAMSES Boltzmann Transport (Particle and Radiation) Code SCEPTRE	9/30/17	SNL
PEM	TBD	Assessment of Tantalum Strength Models in Light of Experimental Results Collected under the Science Campaigns	9/30/17	LLNL
PEM	TBD	Delivery of Improved Physics Models Supporting FY18 Primary Performance Predictive Capability Framework Peg Post	9/30/17	LANL
PEM	TBD	Multiscale Analysis and Comparisons of Cavity Collapse in High Explosives	3/30/17	SNL
PEM	TBD	Coming Full Circle: A Multiphysics Workflow for the Lifecycle of a Gas Transfer System Reservoir	9/30/17	SNL
V&V	TBD	Expand Secondary Computational Assessment and Metrics Project Suite and Optimize Agreement with Underground Tests	9/30/17	LLNL
V&V	TBD	Deliver a Capability for V&V Assessments of Code Implementations of Physics Models and Numerical Algorithms in Support of Future Predictive Capability Framework Peg Posts	9/30/17	LANL
V&V	TBD	Safety in Abnormal Thermal Environments	12/31/16	SNL
V&V	TBD	Fully Coupled Single-Cell Thermal Battery Performance Model	9/30/17	SNL

Sub-Program	ID#	Milestone Title	Complete Date	Site
ATDM	TBD	Demonstrate at Least One Modular Transport Code Coupled to a Multiphysics Application Using the Computer Science Toolkit Linking Capability*	9/30/17	LLNL
ATDM	TBD	An Integration and Deployment Plan for Next-Generation Software Environment	9/30/17	LLNL
ATDM	TBD	Next-Generation Code Project: Release Version 1.0 of a Production Toolkit*	9/30/17	LANL
ATDM	TBD	Asynchronous Many-Task Software Stack Demonstration	9/30/17	SNL
ATDM, IC, CSSE	TBD	Tri-Lab Co-Design Milestone: In-Depth Performance Portability Analysis of Improved Integrated Codes on Advanced Architectures	9/30/17	LLNL, LANL, SNL
CSSE	TBD	Applications Development Environment for Sierra Early Arrival System	6/30/17	LLNL
CSSE	TBD	File System Selected for Sierra	9/30/17	LLNL
CSSE	TBD	Trinity System Production Readiness	3/31/17	LANL, SNL
CSSE	TBD	Advanced Power Management on Trinity Platform	9/30/17	SNL
FOUS	TBD	Sierra “Early Delivery System” Deployed on Secret Restricted Network	6/30/17	LLNL
FOUS	TBD	Complete Sierra Facilities Electrical and Mechanical Design	9/30/17	LLNL
FOUS	TBD	Next-Generation Infrastructure Plan	6/30/17	LANL
FOUS	TBD	Submit Exascale-Class Computer Cooling Equipment Critical Decision 2/3 Package	9/30/17	LANL
FOUS	TBD	Application of Performance Analysis Tools on SNL ASC Codes	9/30/17	SNL

* 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**. 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)
- *ASC Co-Design Strategy* (February 2016)
- 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* (November 2015, this document is available upon request from the ASC Program Office)
- 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.

Major risks and mitigations associated with the DOE ECI are captured in the risk registry of the planned Exascale Computing Project, which is available upon request from the ASC Program Office.

IX. Points of Contact

Headquarters Programmatic Points of Contact

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

Headquarters Implementation Plan 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

WBS	Title	Contact
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
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 Curt Nilsen, SNL, 925-294-1407, canilse@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 Curt Nilsen, SNL, 925-294-1407, canilse@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 Curt Nilsen, SNL, 925-294-1407, canilse@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 Curt Nilsen, SNL, 925-294-1407, canilse@sandia.gov
1.2.3.4	Advanced Technology Development and Mitigation	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@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 Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov

WBS	Title	Contact
1.2.3.4.2	Next-Generation Architecture and Software Development	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov
1.2.3.4.3	Inter-Agency Co-Design	Chris Clouse, LLNL, 925-422-4576, clouse1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@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 Michael 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 Jim Lujan, LANL, 505-665-0718, jewel@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 Jim Lujan, LANL, 505-665-0718, jewel@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 Michael 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 Michael 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 Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.2.3.5.6	Beyond Moore's Law	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Ken Alvin, SNL, 505-844-9329, kfalvin@sandia.gov

WBS	Title	Contact
1.2.3.6	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@lanl.gov Tom Klitsner, SNL, 505-844-1901, tklitsn@sandia.gov
1.2.3.6.1	Collaborations	Pam Hamilton, LLNL, 925-423-1332, hamilton5@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@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 Jason Hick, LANL, 505-667-4477, jhick@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 Jason Hick, LANL, 505-667-4477, jhick@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 Jason Hick, LANL, 505-667-4477, jhick@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

WBS	Title	Contact
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
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 Ryan McDaniel (Project Administrator) 979-862-5969 ryanmcdaniel@tamu.edu
University of Notre Dame, Center for Shock Wave Processing of Advanced Reactive Materials		Prof. Karel Matous 1376kmatous@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 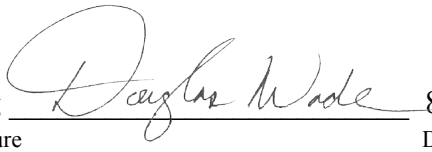 8/24/16
Signature Date

Concurrence:

Michel McCoy, LLNL ASC Executive	*Concurrence received*	8/15/16
Bill Archer, LANL ASC Executive	*Concurrence received*	8/19/16
Bruce A. Hendrickson, SNL ASC Executive	*Concurrence received*	8/19/16

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
APM	Advanced Power Management
ARM	Advanced RISC Machines
ART	Application Regression Test Bed
ASC	Advanced Simulation and Computing (formerly ASCI)
ASCI	Accelerated Strategic Computing Initiative
ASCR	Office of Science's Advanced Scientific Computing Research
ASD	Next-Generation Architecture and Software Development
AT	Advanced Technology
ATCC	Advanced Technology Computing Campaign
ATDM	Advanced Technology Development and Mitigation
B&R	Budget and Reporting

BMGPC	Beyond-Moore General-Purpose Computing
BML	Beyond Moore's Law
CBTF	Component-Based Tool Framework
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CDA	Next-Generation Code Development and Applications
CHAI	Copy Hiding Application Interface
CMOS	Complementary Metal-Oxide-Semiconductor
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
CZ	Collaboration Zone
DisCom	Distance Computing
DOE	Department of Energy
DRAM	Dynamic Random Access Memory
DSW	Directed Stockpile Work
EAP	Eulerian Application Project
ECCCE	Exascale-Class Computer Cooling Equipment
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
ECP-ST	Exascale Computing Project-Software Technologies
ESSIO	Extreme-Scale Storage and Input/Output
F-SEFI	Fine-Grained Soft Error Fault Injection Tool
FleCSI	Flexible Computation Science Infrastructure for Multiphysics
FOUS	Facility Operations and User Support
GPFS	Global Parallel File System

GPGPU	General-Purpose Graphics Processing Unit
GPU	Graphics Processing Unit
HIO	High Performance Input/Output
HPC	High Performance Computing
HPSS	High Performance Storage System
HQ	Headquarters
I/O	Input/Output
IB	Infiniband
IC	Integrated Codes
ICD	Interagency Co-Design
IDC	Integrated Design Code
IP	Implementation Plan
IRB	Institutional Review Board
ITSM	Information Technology Service Management
JBOD	Just a Bunch of Disk
KNL	Knights Landing processors
KVS	Key Value Store
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LAP	Lagrangian Applications Project
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
M&O	Management and Operating
MDHIM	Multidimensional Hashing Indexing Middleware
ML	Machine Learning
MOU	Memorandum of Understanding

MPI	Message Passing Interface
MST	Multisite Target
NAS	Network-Attached Storage
NCI	National Cancer Institute
NCSI	National Strategic Computing Initiative
NFS	Network File System
NGBB	Next-Generation Backbone
NGCE	Next-Generation Computing Enablement
NIC	Neural-Inspired Computing
NNSA	National Nuclear Security Administration
NNTI	Nessie Network Transport Interface
NPR	Nuclear Posture Review
NRE	Non-Recurring Engineering
NSCC	National Security Computing Center
NSCI	National Strategic Computing Initiative
NVRAM	Non-Volatile Random Access Memory
nWBS	National Work Breakdown Structure
O|SS	Open|SpeedShop
OCF	Open Computing Facility
ORNL	Oak Ridge National Laboratory
OS	Operating System
OS/R	Operating System/Runtime
PCF	Predictive Capability Framework
PDX	Patient-Derived Xenograft
PECASE	Presidential Early Career Award for Scientists and Engineers
PEM	Physics and Engineering Models
PI	Principal Investigator
PIC	Particle in Cell
PIDX	Portable Information Data eXchange
PSAAP	Predictive Science Academic Alliance Program
PSEP	Predictive Science Engineering Panel

PSP	Predictive Science Panel
QC	Quantum Computing
QIP	Quantum Information Processing
QMU	Quantification of Margins and Uncertainties
R&D	Research and Development
RAS	Reliable, Available, Secure
RDMA	Remote Direct Memory Access
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
RISC	Reduced Instruction Set Computer
RTS	Runtime System
RZ	Restricted Zone
SAN	Storage Area Network
SAP	Scalable Applications Preparation
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 (in ASCR)
SCN	Sandia Classified Network
SCR	Scalable Checkpoint Restart
SDS	Software-Defined Storage
SFI	Significant Finding Investigation
SIMD	Single Instruction, Multiple Data
SIO	Scalable Input/Output
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SPOT	System for Performance Optimization Tracking
SRN	Sandia Restricted Network
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit

STAT	Stack Trace Analysis Tool
SU	Scalable Unit
SUPER	SUstained Performance, Energy & Resilience Institutes (SCiDAC)
TLCC	Tri-Lab Linux Capacity Cluster
TOSS	Tri-Lab Operating System Stack
UGT	Underground Test
UQ	Uncertainty Quantification
V&V	Verification and Validation
VM	Virtual Machine
VTK	Visualization Toolkit
WA	Work Authorization
WAN	Wide Area Network
WBS	Work Breakdown Structure
WCI	Weapons & Complex Integration (at LLNL)
XTD	X-Theoretical Design Division (at LANL)

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)

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. 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 next-generation HPC systems. The ATDM subprogram is the NNSA implementation of the DOE Exascale Computing Project (ECP) and aligned with the NSCI Executive Order. Performing the ATDM work within the scope of the DOE ECP 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.

Note: This Appendix only includes activities within the ASD and ICD products. Content for the ATDM CDA product is available upon request from the ASC Program Office.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2015, and through quarter 3, fiscal year 2016, are reflected below for the ATDM subprogram.

- Delivered strategic plan and initial scalability assessment by Advanced Architecture and Portability Specialists (AAPS) Team. (LLNL)
- Formed, staffed, and completed initial tasks for four new LLNL teams: AAPS, Production-Quality Tools Development (ProTools), User Workflow and Modernization, and HPC Development Ecosystem. (LLNL)
- Worked with IC code team to develop the Copy Hiding Application Interface (CHAI), a portable data abstraction that enables easier programming for heterogeneous architectures without having to manage locality or movement. (LLNL)
- Completed ProTools work to support IC users, including porting the Spack package manager and Scalable Checkpoint Restart (SCR) to Cray systems (specifically Trinity) and building an initial prototype of SPOT, a continuous performance analysis tool. (LLNL)
- Successfully advanced the state of tools for debugging at scale through research, development, and co-design (for example, debugging for OpenMP, GPU debugging on Sierra with TotalView, and debugging research on scalable record and

deterministic reply for MPI, as well as non-deterministic problems due to message races). (LLNL)

- Developed a Flexible Computation Science Infrastructure for Multiphysics (FLECSI), a mesh topology abstraction on which the new code effort is based. Began initial development of new generation codes on the infrastructure. (LANL)
- Developed Legion storage interface with Intel and HDF5 group. (LANL)
- Completed 2015 Level 2 milestone in co-design development of SNAP with KOKKOS and Legion. Began FY16 Level 2 milestone of moving this work into Partisn. (LANL)
- Finalized the workflow taxonomy to represent multiphysics application flows. (LANL)
- Included in the ISO/C++17 standard first-vote in March 2016 the SNL-led fix to the integration of lambda and parallel dispatch mechanisms. The SNL-led proposal for new multidimensional array mechanism modeled on KOKKOS' View is advancing in the ISO/C++ Standards committee and on-track for ISO/C++20. (SNL)
- Developed first version of the DARMA interface specification to make expression of application intent relatively simple, while insulating the applications developers from the idiosyncrasies of underlying system runtimes. (SNL)
- Completed initial prototype of the core DataWarehouse components, including the key-value based management layer, ports of the remote direct memory access (RDMA) abstraction layer to several HPC interconnects, and support for KOKKOS memory views. (SNL)
- Completed the Level 2 milestone (#5674) by incorporating transient sensitivity analysis capabilities using Sacado automatic differentiation into EMPIRE and SPARC, demonstrating substantially improved performance compared to traditional finite differencing. (SNL)
- Developed and deployed an automated build-system for third party libraries; deployed 32-node Intel Knights Landing (KNL) test bed on the SNL restricted network; and installed eight Power8+/NVIDIA general purpose GPU (GPGPU) nodes on both the SNL open and restricted network environments. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): An Integration and Deployment Plan for Next-Generation Software Environment		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: ATDM		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Building on the existing LLNL software environment and plans for individual next-generation software tools, interfaces, and standards, this milestone addresses and will issue a plan for an integrated LLNL next-generation software environment. The plan includes systems software, the code development environment (including programming models and runtimes), and tools used directly by applications end users. The integrated environment will enable developers and users to leverage the next-generation computer science scalability research as well as systems software innovations.		
Completion Criteria: A report covering the strategy for continued innovation and a description of the LLNL next-generation software environment.		
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: ATDM and CSSE staff.		

Milestone (ID#TBD): Asynchronous Many-Task Software Stack Demonstration		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: ATDM		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: The asynchronous, many-task (AMT) programming models are emerging as a candidate model for addressing the challenges of next-generation computing platforms. However, there are many variants of this model emerging, and SNL's FY15 ATDM Level 2 milestone (#5325) highlighted requirements gaps in these leading AMT runtimes for ASC applications. DARMA is an AMT portability layer and specification under current development within ATDM to address these gaps for next-generation ASC applications. This milestone will evaluate a DARMA-compliant AMT runtime software stack comprising ATDM-ASD software components (DARMA, Data Warehouse, KOKKOS, and Resource Manager) and existing community AMT runtime technologies. SNL will assess the performance and productivity of this software stack on representative ATDM applications and proxies. In order to assess the strengths and weaknesses of the AMT model, SNL will perform computational experiments on test bed machines and one or more AT systems. The outcomes of this effort will include 1) an initial DARMA-compliant AMT software stack, 2) a clear understanding of the strength and limitations of the AMT API and runtime implementation in the context of SNL's ATDM codes, 3) and information to guide future R&D in this area (for example, system performance issues, portability issues, and usability issues on realistic apps).		
Completion Criteria: Completion of milestone review and a report is provided summarizing the key results and recommendations as regards ASC/ATDM code preparation for next-generation platforms.		
Customer: ATDM-CDA/IC/ECP		
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 applications (or proxies) and software components, expertise and feedback from SNL ASC/ATDM code developers, advanced architecture test beds, and Trinity AT platform.		

Milestone (ID#TBD): Tri-Lab Co-Design Milestone: In-Depth Performance Portability Analysis of Improved Integrated Codes on Advanced Architectures		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: ATDM, IC, CSSE		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: This milestone will exercise the IC or ATDM codes using new and emerging programming models, to demonstrate portability and analyze performance characteristics on the latest advanced architectures expected to become widely available in FY17. Performance analysis tools will be used to do an in-depth analysis of those same or similar production codes on available ASC AT platforms and test beds, represented by the significantly disparate architectures of ATS-1/Trinity and ATS-2/Sierra. This will advance ASC co-design efforts by impacting one or more of the following: • Determining performance bottlenecks in new architectures • Characterizing ASC applications (for example, memory footprints, network utilization, instruction mixes) so as to better inform proxy application development and hardware design space evaluations • Evaluating the portability and performance of different data structures and data layouts for important physics motifs across multiple architectures • Evaluating the effectiveness of hardware and/or software-based memory and data management strategies to deal with automating movement between levels of system memory and I/O hierarchies • Evaluating the performance analysis tools in these environments to determine their strengths and weaknesses, which will be invaluable in feedback to the vendors		
Completion Criteria: Evaluate the performance of at least one code per lab on the Trinity/KNL partition and the latest Sierra test beds using the performance-portability abstractions introduced. Develop a database of application characteristics to be used to both improve the representativeness of proxy applications and to better address co-design questions with the vendor community. Report on the strengths and weaknesses of the performance analysis tools available in these environments and recommendations for improvements.		
Customer: IC/ATDM		
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.		

Milestone (ID#TBD): Tri-Lab Co-Design Milestone: In-Depth Performance Portability Analysis of Improved Integrated Codes on Advanced Architectures
Supporting Resources: ATDM, IC, CSSE staff.

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 (NCGE) and Co-Design project efforts are preparing ASC for the next generation of advanced computing technologies. This project coordinates next-generation activities both within the ASC program and externally. This effort includes interactions with the DOE's ASCR, vendors, and academia, including planning and technical coordination for vendor and academic contracts and CORAL working group efforts. It includes standards work, performing feasibility and prototyping studies, and projects in areas that include system level software, power-aware computing, development environment tools, data analysis tools, and programming models. Efforts on full-system optimization covers work across all vertical elements of the software stack. Team members will carry out investigations and co-design activities using test beds and early availability machines, in addition to current technology.

Accomplishments in FY16:

- Participated in vendor and Center of Excellence (COE) activities, as well as DesignForward efforts.
- Participated in standardization activities for both MPI and OpenMP in leading roles, as well as new initiatives within these standardization efforts.
- Participated in investigations of the SNL Level 2 milestone on programming models, in particular in the area of tools.
- Advanced the debugging and code-correctness tool suite to realize LLNL's post-petascale plan through co-design, research, development, and deployment.
- Conducted co-design activities with vendors (for example, prototyping the OpenMP debug interface—OMPD—and working with the vendor to integrate it into TotalView, as part of Sierra NRE). Transitioned the hardening of the Caliper instrumentation library and the SCR checkpointing system into the ProTools team, while continuing active R&D targeting future systems.

Planned Activities in FY17:

- Continue active involvements in standardization effort for MPI and OpenMP; continue and expand investigation of new programming models and especially task-based models in close collaboration with code teams and portability specialists.
- R&D tools and techniques to effectively utilize emerging next-generation architectures (for example, network-aware scheduling; tools for measuring, visualizing, and analyzing the impact of applications on power, energy, and performance; performance modeling; and next-generation debugging).
- Continue development of Caliper coupled with initial integration work of Caliper annotations into LLNL code base and work with ProTools team on continued hardening.
- Exercise applications on test beds to identify gaps and enhancements needed for next-generation simulation codes.
- Investigate power-aware HPC and its impact on LLNL systems and applications as well as coordinate with vendor efforts.
- Continue participation in vendor, COE, and co-design activities, as well as technical management of academic contracts and other external interactions.
- Continue development of the Flux resource management framework, including investigations into novel scheduling algorithms and tools that leverage the Flux core services.
- Extend memory profiling capabilities, including investigations of on-node memory architecture for GPUs.
- Extend CTS-1 network investigation to understand the impact of factors such as congestion, multirail networking, and topology-aware placement.
- Continue work on advanced and scalable debugging with a focus on reproducibility and race detection/avoidance for both MPI and OpenMP applications.

Advanced Architectures and Portability Specialists Project (LLNL)

The LLNL ATDM strategy for supporting the development of next-generation simulation capabilities includes the work 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, the ASC code developers have access to an agile labor pool of computational scientists and computer scientists skilled in scaling applications on new, cutting-edge hardware. The team shares knowledge through a code repository and wiki as well as seminars, hackathons, and publications. The team includes specialists in key areas such as GPGPU programming, many-core programming, I/O, and parallel application development.

Accomplishments in FY16:

- Formed and staffed AAPS team and completed a Level 2 strategic planning milestone.
- Created and populated a shared wiki and code repository that stores descriptions of best practices and lessons learned, including code examples.
- Worked with ALE3D team to develop the CHAI, a portable data abstraction that allows developers to write code for heterogeneous architectures without needing to manage locality or movement.
- Developed simplified prototype application, Cheep, for the Cheetah thermochemical physics package and used it to test strategies for accelerating the underlying solves in batch mode; for example, ordering the solver calls to take advantage of correlations in neighboring phase space points, which resulted in a threefold improvement. Optimized kernels to make better use of single instruction, multiple data (SIMD) hardware.

Planned Activities in FY17:

- Work with ARES code team to conduct performance assessment to determine potential resource bottlenecks on next-generation architectures and help determine optimal porting strategy. Introduce data movement and loop abstractions such as CHAI and RAJA to enable GPU acceleration and ensure performance portability.
- Work with deterministic transport code team to develop porting strategy for existing codes, ARDRA and TETON. Assist with code refactoring of ARDRA code to utilize newly-developed nested loop RAJA constructs.
- Continue development of programming abstractions such as CHAI and RAJA to improve the performance portability of ASC applications.
- Continue leadership in exascale co-design by working with vendors and compiler developers; continue to investigate new and updated programming models within established set of proxy applications.

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 visualization and 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 ProTools development team takes on the responsibilities to ensure these goals are met, working closely with the R&D efforts on one side and user support the AAPS team on the other.

Accomplishments in FY16:

- Developed a strategic plan describing the production path for tools, including testing and documentation and release strategies.
- Identified the initial code teams for ProTools collaboration and used the code teams' input to choose a set of research tools to target.
- Identified the initial set of tools targeted for ProTools productization: SCR, Spack, Stack Analysis Tool (STAT), Caliper, SPOT, DI-MMAP, and mpiIOTools; began work on productizing and providing support for these tools.
- Ported Spack and SCR to Cray systems, specifically to the Trinity platform.
- Built the initial prototype of the System for Performance Optimization Tracking, or SPOT, tool (performance tracking of nightly tests).

Planned Activities in FY17:

- Deliver a working SPOT implementation that does continuous performance analysis for the ALE3D code; integrate with Caliper application monitoring efforts; begin working with other code teams interested in this infrastructure; extend SPOT's visualization capabilities.
- Continue the ongoing effort of providing Spack support to code teams.
- Port STAT, DI-MMAP, and SCR to Sierra.
- Continue the ongoing effort of identifying the next generation of research tools; continue working with code teams, AAPS, and management on ProTools priorities

User Workflow and Modernization (LLNL)

The Workflow project's mission is to build infrastructure and applications that enhance the end-to-end productivity of ASC HPC assets. The Workflow project is focused on building solutions to concrete user problems in three areas prioritized by the user community: problem setup, simulation management, and post-processing. Its approach is to develop applications based on re-usable components and web technologies. In this way, the team mitigates risk by providing the user community with a set of components that encompasses best practices and common operations, allowing the user community to build custom workflows where needed while reducing duplicated work.

FY16 Accomplishments:

- Created a proof-of-concept web application, "Contour Central," for creating and editing collections of LANL neutral-format contours and storing them in an SQL database; created a version of the nuclear data viewer tool that is used to understand simulation inputs, which is the first step towards refreshing this tool.
- Developed VV4ALE3D, a prototype application for configuring and managing groups of ALE3D simulations that incorporate a data management strategy and

workflow enhancements in re-usable Python packages; began leveraging this work in collaboration with SNL and LANL to define a shared requirements document for data management.

- Supported integration of the ZFP compression library into Silo and began initial evaluation of incorporation into VisIt to reduce memory overheads; began developing a GPU accelerated ray-tracing infrastructure targeting both VisIt and HADES (simulated radiography).

FY17 Deliverables:

- Develop requirements for a geometry data representation that enables existing problems to be described using the LANL contour representation and extracted from a back-end database; develop a prototype implementation and demonstrate on a problem of significant complexity.
- Develop a requirements document for simulation data management with SNL and LANL; if successful, begin development of shared software components to support data management requirements of workflows; begin deployment of python workflow management and parameter study tools such as VV4ALE3D and associated components.
- Complete initial integration of GPU accelerated ray-tracing infrastructure in VisIt, and explore integration in HADES for simulated radiography applications.

High Performance Computing Development Ecosystem (LLNL)

LLNL has created a new HPC Development Ecosystem team within ATDM ASD. The mission of the Ecosystem team is to build infrastructure, tools, and secure services that multiply the productivity of developers at the computing center. Development Ecosystem focuses on areas where the mission needs of simulation developers overlap with development needs of center staff, reducing duplication of effort and creating a world-class software development platform.

Development teams require increasingly sophisticated tools and closer integration with the deployment environment. Moreover, the security and maintenance costs are too high for each team to maintain its own tools, and the needs of staff and code teams increasingly overlap. The HPC Ecosystem project bridges this gap by identifying the shared needs of the computing center and its users, and by deploying tools and services that boost productivity and tighten integration of developers with the compute environment. This includes developing monitoring solutions that expose more LC performance information to users, standardizing software deployment methodologies and tools, deploying common systems for build and test automation, and ensuring services and automation tools can run securely within LC's environment.

The HPC Development Ecosystem team takes a holistic view of the center as a platform for developers, and it works closely with other ATDM teams to achieve its vision.

Accomplishments in FY16:

- Assembled initial team, producing team goals and mission statements.
- Developed production monitoring cluster, Sonar; deployed monitoring tools (Lightweight Distributed Metric Service (LDMS), Flume) and database (Cassandra) in the collaboration zone (CZ) and restricted zone (RZ).
- Began implementing security for Cassandra so all users can access monitoring data.
- Popularized Spack package manager among Weapons & Complex Integration (WCI) code teams; grew external Spack contributor community to include NNSA, SC labs, the Ecole Polytechnique Fédérale de Lausanne, the Fermi National Accelerator Laboratory, CERN, NASA, and others.
- Deployed Bamboo continuous integration for WCI and Livermore Computing (LC) on the unclassified networks (CZ and RZ).

Planned Activities in FY17:

- Expand monitoring capabilities to more machines and to the secure computing facility (SCF).
- Develop actionable analytics techniques to tune I/O, networks, and applications.
- Begin to harden and deploy ScrubJay data analytics tool, which allows users to easily query data from the Cassandra database deployed on Sonar.
- Develop a scalable maintenance model for Spack, in close coordination with the internal efforts of ProTools, to integrate external contributions with internal advances and to consistently deploy software across multiple DOE facilities.
- Improve Bamboo build/continuous integration by allowing build agents to run as particular users on demand; automate the agent deployment process.

Storage and Input/Output Research and Development (LLNL)

The storage and I/O R&D project focuses on forward-looking I/O investigations, research projects, and NRE for exascale and pre-exascale timeframes. This work includes participation in CORAL file system, function shipping and burst buffer NRE efforts that develop and scale CORAL file system, and node-local burst buffer software intended to greatly increase application I/O performance while reducing off-machine storage requirements and jitter. Other project efforts include participation in the FastForward2 Extreme-Scale Storage and I/O (ESSIO) R&D project, which focuses on continued development and application porting to the ESSIO stack, and participation in the LLNL file systems futures investigation.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Participate in CORAL NRE File System and Burst Buffer working groups in preparation for the Sierra machine.
- Continue LLNL's file system futures effort investigating I/O technologies, software, and methods for the exascale timeframe.
- Participate in the FastForward 2 ESSIO R&D effort to further develop the core functionality of the ESSIO stack and demonstrate DOE software applications operating on this next-generation I/O software.

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 (ANL, Lawrence Berkeley (LBNL), LLNL, LANL, 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 FY16:

- Provided technical coordination and contractual management for FastForward contracts.

Planned Activities in FY17:

- Provide technical coordination and contractual management for FastForward contracts.

Exascale PathForward Project (LLNL)

The PathForward project is the central element of the DOE ECP Hardware Technology effort to accelerate the R&D of critical technologies needed for exascale computing. The selected vendor contracts will be administered by Lawrence Livermore National Security, LLC, on behalf of ECP and under technical project management of the ECT HT team and personnel of six DOE "partner" national laboratories (ANL, LBNL, LLNL, LANL, ORNL, and SNL). PathForward is the follow-on to FastForward2 and DesignForward2 programs, which will be ending in FY17, and runs through 2019. PathForward seeks solutions that will improve application performance and developer productivity while maximizing energy efficiency and reliability of an exascale system. PathForward activities will include R&D to:

- Substantially improve the competitiveness of the HPC exascale system proposals in 2019, where application performance figures of merit will be the most important criteria.
- Improve the Offeror's confidence in the value and feasibility of aggressive advanced technology options that they are willing to propose for the HPC exascale system acquisitions.
- Identify the most promising technology options that would be included in 2019 proposals for the HPC exascale systems.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Conduct contract negotiations and issue awards.
- Provide technical coordination and contractual management for PathForward contracts.

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 base projects to leverage and extend the expertise, focusing on enabling the ATDM-CDA project and extreme-scale environment.

Accomplishments in FY16:

- Parallelized the Fine-Grained Soft Error Fault Injection (F-SEFI) Tool fault injector.
- Integrated initial version of high performance I/O (HIO) burst buffer library into HDF5 as a virtual file system.

Planned Activities in FY17:

- Continue the initial investigation of machine learning (ML) applied to performance counters and benchmarking of existing and emerging ATDM codes and proxies with an end goal of predicting application performance.
- Collect, analyze, and share storage system diagnostic and tracing data, as part of an effort to share data throughout the DOE complex.
- Implement feature extraction algorithms that automate visualization and analysis processes and that identify, quantify, and track key events.
- Continue the workflow model development for existing workflows, moving toward representing future workflows on possible future systems.
- Further deploy, test, and deploy MarFS (campaign storage—an intermediate storage tier between the parallel file system and archive storage).
- Continue to develop the HIO library with the focus on supporting the Sierra architecture. Continue the development of the HDF5 data model for ASC applications as they apply to HIO and the virtual file interface to HDF5.
- Develop storage capabilities for LANL's ATDM project, including leveraging existing scientific file formats, key value stores, and non-volatile memory (NVM)
- Harden the FSEFI Fault Injector (which uses the QEMU virtual machine interface to inject faults in applications or the OS/Runtime) for internal and external work as a production-quality tool.

Co-Design and Programming Model Research (LANL)

This project contains the forward-looking research for advance computing technologies at extreme-scale. Co-design 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.

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. 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.

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.

Accomplishments in FY16:

- Successfully improved performance of large multiphysics production run on Trinity by using the QUO library that manages OpenMP thread and MPI rank interaction and layout.
- Began DOE Storage FastForward2 Phase 2 project with Intel and HDF-5 group. The Legion runtime will be the focus of this effort.
- Completed the initial design and prototype of FleCSI, a mesh topology abstraction for ATDM.

Planned Activities in FY17:

- Address hardening and control replication issues for the Legion runtime.
- Support multiple back ends and additional mesh topology for FleSCI (node-level, for example, C++17 PAL/Thrust/Pinion/KOKKOS, and/or distributed, for example, MPI).
- Develop an open source (non-export controlled) multiphysics proxy application.
- Implement Legion + X interoperability; investigate and develop methods to combine Legion with other node-level parallelism runtimes.
- Continue co-design summer school, a highly successful vehicle delivering on two main objectives: 1) recruiting staff to LANL as postdocs and staff, and 2) investigating programming models and algorithms in a multidiscipline environment; choose a new programmatically relevant topic on which the student teams will focus.
- Continue development of Gladius asynchronous and scalable tool interfaces; support QUO for MPI/OpenMP multiphysics applications; extend the BYFL productivity tool to support analysis with threaded applications.
- Expand and develop parallel constructors for low-level virtual machine (LLVM).

Future-Generation Computing Technologies (LANL)

This project includes high-risk, high-reward research for future systems, including research on advanced programming models, virtual environments, alternative hierarchical storage uses for advanced in-situ.

Data parallel research includes PINION/PISTON and VTK-m. 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.

Next generation codes, PISTON and vtk-m, are focused on graphics algorithms for large multicore machines.

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.

Additionally, this project investigates virtual environment technology (containers, virtual machines) that provides data-rich turn-key environments for applications. It also investigates virtual environments for automated regression testing of CDA codes and instrumentation of application workflows. Finally, basic workflow taxonomy development is included in this work.

Accomplishments in FY16:

- Packaged multiphysics applications in a virtual container environment and investigated the overhead trade-off between virtualized and container environments.
- Completed initial burst-buffer investigations for streamline computation.
- Completed and reported the workflow taxonomy characterizations for the Capability Computing Campaign 9 Science Campaigns.

Planned Activities in FY17:

- Develop a framework that utilizes both VMs and software containers to create a workflow system that establishes a uniform build and execution environment beyond the capabilities of current systems.
- Support ATDM-CDA investigation in data-parallel programming and algorithm development for the vtk-m library.

Programming Models and Abstractions (SNL)

This project explores programing models and abstractions that allow new ASC 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 AMT programming models (DARMA) with an associated data-warehouse capability.

Accomplishments in FY16:

- Conducted five hands-on KOKKOS tutorials:.
- Confirmed quantitatively (in the tri-lab co-design Level 2 milestone #5498) that porting a C++ application to KOKKOS is an equivalent level of effort as porting that application to OpenMP.
- Fixed the C++ lambda feature in C++17 to correctly interact with parallel and concurrent execution.
- Redesigned and re-implemented KOKKOS' multidimensional array capability to support non-trivial array-member types.
- Developed first version of the DARMA interface specification.
- Completed initial prototype of the core DataWarehouse components.
- Delivered performance-portable graph-coloring algorithm and symmetric Gauss-Seidel preconditioner in KOKKOSKernels.
- Developed Qthreads instrumentation capabilities to expose timing of task and thread executions and communication/computation.

Planned Activities in FY17:

- Deliver updated KOKKOS and release KOKKOS Version 3 with on-node resource management and fine-grained hierarchical task-data parallelism.
- Deliver performance-portable kernels for block matrices and batched (contraction) kernels as required by ATDM applications and integrate them into TRILINOS.
- Integrate DARMA, DataWarehouse, and KOKKOS into an AMT software stack capable of running a representative ATDM code on an ATS platform for FY17 Level 2 milestone.
- Design and prototype support for UQ, provenance, and resilience mechanisms for single-node failures within the DataWarehouse.
- Deliver initial implementation of an integrated AMT software stack on advanced

- Extend and coordinate current efforts, as appropriate and without diminishing ASC impact, in KOKKOS, Darma, DataWarehouse, and Qthreads to meet the needs of the broader the ECP applications.

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. This project focuses on the lower level computer-science components and their APIs that are written at a degree of abstraction that can be used across a wide range of ATDM applications.

Accomplishments in FY16:

- Extended the mesh database to import data directly into KOKKOS performance portability data structures.
- Integrated support for serial and parallel unstructured meshes in the mesh database project.
- Designed the new APIs for a next-generation time integrator based upon a thorough requirements gathering process.
- Met completion criteria for the Level 2 milestone by incorporating embedded analysis capabilities using Sacado automatic differentiation into EMPIRE and SPARC, demonstrating substantially improved performance compared to traditional finite differencing.
- Developed a unified framework that integrated existing contact proximity search products and incorporated KOKKOS for portable GPU and threaded proximity search.

Planned Activities in FY17:

- Integrate mesh database component into EMPIRE and SPARC. Demonstrate hybrid structured/unstructured mesh support on an ATDM problem of interest.
- Develop and demonstrate implicit-explicit time integration capability and demonstrate and document performance on problem of interest for EMPIRE.
- Incorporate tools and APIs for advanced embedded optimization and UQ algorithms, including second derivatives, ensembles, and transient adjoint sensitivities into ATDM applications.
- Collaborate with ANL as part of ECP to develop source transformation-based embedded sensitivity/ensemble propagation techniques for many-core architectures (pending support from ECP).
- Evaluate KOKKOS performance portability and memory hierarchy directives for global search that is a key component for contact.

- Extend and coordinate current efforts, as appropriate and without diminishing ASC impact, in AgileComponents to meet the needs of the broader ECP applications.

Next-Generation Software Technologies (SNL)

Starting in FY17, software technologies originally under development within SNL's CSSE program will be merged into ATDM and refocused to target the needs of next-generation codes. This includes performance analysis tools, operating systems (OS)/runtimes, scalable visualization, scalable input/output (SIO), and scalable solvers. Many of these technologies are foundational so that work can support both the needs of next-generation ASC codes while also benefiting the broader needs of the ECP.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Complete performance study of linear solvers within the context of EMPIRE and SPARC end-to-end efforts.
- Complete assessment of solver API designs to support asynchronous many-task invocation of solvers.
- Develop basic Fortran API to key TRILINOS solver capabilities. Demonstrate use in LANL candidate code and coordinate with broader ECP activities.
- Evaluate and document the performance portability of the Nessie Network Transport Interface (NNTI) on various network abstraction layers as well as modifications to support AMT.
- Evaluate SNL SIO component portability and performance using burst buffer on Trinity, including an evaluation of both HIO and Cray's DataWarp, to inform design for next-generation codes.
- Deliver design of fine-grained networking in Qthreads.
- Deliver design of on-node resource management interfaces, mechanisms, and policies.
- Evaluate lightweight operating system/runtime (OS/R) stacks using Trinity KNL test bed; document results and implications for next-generation codes.
- Work with vendor and ECP community to address OS and runtime system software limitations with specific focus on meeting ATDM and ASC application needs.
- Develop the framework for multi/many core data analysis and visualization algorithms within Visualization Toolkit (VTK)-m and coordinate with broader ECP activities.

- Demonstrate prototype for multi/many core algorithms into ParaView and Catalyst, including integration of VTK-m into VTK and a ParaView plugin and utilize the new capabilities.

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 FY16:

- Developed and deployed an automated build system for third-party libraries and version-of-the-day application components for use on development and test bed environments.
- Deployed 32-node Intel KNL test bed on the SNL restricted network and eight Power8+/NVIDIA GP-GPU nodes on both the SNL open and restricted network environments.

Planned Activities in FY17:

- Develop a simplified, flexible, coordinated build and integration system for all changing software libraries and application codes, automated handling of configuration on various test bed platforms, and process to maintain working libraries and application codes for changes in the libraries, applications, and test bed development environments.

- Target improvements in ATDM hardware infrastructure as necessary, including Intel Phi, IBM Power/NVIDIA GP-GPU, and/or Advanced Reduced Instruction Set Computer (RISC) Machines (ARM) architectures as demanded by ATDM developers.
- Purchase software licenses and maintain test bed software environment in support of ATDM test and development environment.

Projects for the Inter-Agency Co-Design Product (WBS 1.2.3.4.3)

The Inter-Agency Co-Design product will seek to increase the capacity and capability of an enduring national HPC ecosystem via inter-agency collaborations with other U.S. Federal agencies. The projects will leverage NNSA HPC advanced architecture activities and software technologies to address the sponsor agencies' mission needs. Agencies will also have the opportunity to participate in co-design activities with vendors and academia, in addition to workforce development and training opportunities.

ASC Interagency Co-Design (LLNL)

DOE is partnering with the National Cancer Institute (NCI) of the National Institutes of Health in the development of exascale ready tools, algorithms, and capabilities. DOE's efforts will focus on co-design research that will be coordinated with parallel efforts by the NCI to develop the field of predictive oncology as part of the President's Precision Medicine Initiative. In modeling and simulation, DOE efforts will focus on multiscale, multiphysics code frameworks suitable to exascale architectures and with quantified predictive capacity. In data sciences, efforts will include work on scalable data structures adaptable to the exascale based heterogeneous architectures and data analytics. The ASC Interagency Co-Design Project is comprised of three pilots that aim to improve understanding of cancer biology and its application to more effective therapies through targeting of drug therapy for cancer patients (Pilot 1), producing an unprecedented scale of dynamic simulation for the reliable, available, secure (RAS) oncogene to facilitate new drug discovery and development (Pilot 2), and to apply real-time, big data analytics to the a suite of cancer databases to optimize cancer therapies (Pilot 3).

LLNL's goal is to provide a rationale and to accelerate the development of a novel diagnostic and/or targeted therapy based on the ability to predict the fundamental mechanism of RAS-driven cancer initiation and growth in various tissue types by creating an unsupervised machine learning ecosystem that analyzes simulation data, recognizes biologically relevant models, creates new hypothesis and abstracts, creates new systems for finer-scale simulations, and returns newly learned results to higher-scale simulations.

Activities will include R&D to:

- Develop machine learning techniques for Pilot 1.
- Develop multiresolution/multiscale simulation techniques for Pilot 2.

- Develop natural language processing and deep learning techniques for Pilot 3.

Accomplishments in FY16:

- Held Livermore Big Artificial Neural Network Workshop for Pilot 1.
- Hired three staff members for Pilot 2; cross-compared CHARMM, NAMD, AMBER, and GROMACS codes; submitted INCITE proposal.
- Put an Institutional Review Board (IRB) in place to work with shared pathology datasets for Pilot 3; received acceptance of two offers for positions starting early FY17.

Planned Activities in FY17:

- Develop scalalable neural net on LLNL HPC (Surface and Catalyst) for Pilot 1; scale up to LBANN.
- Validate three-way lipid phase diagram (DPPC/DOPC/CHOL) for Pilot 2; define realistic membrane (size, composition); investigate mobility of RAS highly variable region in membrane; develop variable-resolution algorithm.
- Create tensor-flow toolbox for deep learning for Pilot 3; begin application to pathology datasets.

Applying HPC and Machine Learning to Cancer Research (LANL)

LANL will develop a set of models that can predict drug responses across a range of cell lines (starting with the current NCI-60 repository), and over the course of the pilot extend these results to new cell lines and to patient-derived xenograft (PDX) models. This effort will focus on evaluating machine learning approaches through a series of cross comparisons. LANL will evaluate the sets of features (molecular characterizations) to determine optimal features to be used for building classifiers.

LANL will establish the gold standards for RAS protein/membrane systems modeling at the all-atom, explicit solvent, and membrane representation level.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Establish Predictive Modeling Framework.
- Carry out base-line MD simulations of RAS.
- Establish baseline information on natural language processing, statistics, and machine learning for cancer data.

Biological Planetary Protection (SNL, NASA-JPL)

This computational co-design challenge helps address the certification of spacecraft for biological planetary protection requirements. Spacecraft that land or interact with other planetary bodies, including Mars, Europa, and small cometary bodies, must certify low probabilities for both forward and backward biological contamination. This certification must be done for both nominal and failure scenarios and often these conditions cannot be created in experimental test environments such that computational models are the primary means for certification.

One motivating mission for this challenge is the planned Europa lander. Key concerns include failure scenarios for the parts of the lander that melt into the Europan ice, generating an environment in which terrestrial organisms can reproduce. Another motivation for this challenge is future Mars landing missions. A key concern for Mars is a crash scenario in which parts of a lander's radioactive thermal generator power supply become embedded in the Martian regolith. The recent discovery of water-bearing hydrates in the Martian regolith raise the potential for embedded nuclear material to create a habitable zone that can support terrestrial organisms.

This project will co-design computational simulation capabilities leveraging SNL's computational mechanical capabilities within the NNSA/ASC developed SIERRA code suite. In so doing, the project will drive these critical NNSA applications to support new architectures while also providing new capabilities relevant to NASA-JPL missions. The activities are collaborative and the lead organization for each is indicated.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Develop requirements and model/code plan for crash-mechanics, burning/melting of spacecraft, and porous transport of contaminants in heterogeneous planetary materials. (SNL)
- Develop requirements and model/code plan for planetary materials including burning and melting, geochemical processes, and biological processes. (JPL)
- Conduct and document model V&V as well as UQ analysis. (SNL)

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 2015, and through quarter 3, fiscal year 2016, are reflected below for the CSSE subprogram.

- Managed the CTS-1 contract and procured more than 77 CTS-1 scalable units for NNSA labs. (LLNL)
- Took delivery of Sierra early access system and provided technical coordination and contractual management for CORAL NRE and Sierra contracts. (LLNL)
- Deployed Hype Tri-Lab Operating System Software (TOSS) test bed and developed TOSS 3 with a beta release used in vendor CTS-1 testing. (LLNL)
- Completed project to modernize and expand disk cache capacity for archival storage. (LLNL)
- Directed scalability and performance enhancements for the TotalView debugger, and provided enhanced, direct vendor support for debugging issues for ASC applications developers. (LLNL)
- Initially integrated HIO into Eulerian Applications Project (EAP) code, allowing application checkpoint to burst-buffers (FY15 Level 2 Milestone). (LANL)
- Delivered in situ consulting/training for EAP code to expose the user base to the advantages of in situ, specifically, direct 3D imagery and plots of Lagrangian tracers. (LANL)

- Delivered a parallel algorithm to calculate two layers of ghost cells on generalized well-formed unstructured grids in the VTK, the basis for ParaView. Catalyst and VisIt, among other applications, can now simply call the ghost cell generator prior to any parallel filters that require ghost cells. (LANL)
- Populated workflow taxonomy with data requirements for various workflows (from CCC-9) used for the Crossroads procurement, which is the first time workflows have been used for an AT system procurement. (LANL)
- Successfully deployed Haswell partition of Trinity. (LANL)
- Completed acceptance of Trinity Phase I platform and Trinity Open Science. Obtained approval for ATS-3/Crossroads Critical Decision-0 (CD-0) and CD-1/3a package for Crossroads. (LANL/SNL)
- Deployed Bowman, the first rack-scale deployment of Intel KNL outside of Intel or their major system integration partners. (SNL)
- Contributed significantly to a tri-lab co-design milestone (#5498) demonstrating increased performance on proxy applications using KOKKOS. (SNL)
- Performed an early evaluation of the power management capabilities of the Cray XC40 architecture. Released version 1.1 of Power API specification, which incorporates feedback from Cray. (SNL)
- Demonstrated in situ visualization using SIERRA/Catalyst on 4k cores on Cielo and a non-ASC ensemble analysis using Slycat for a problem with more than 50,000 runs worth of data. (SNL)
- Participated in lease-to-own approach to acquisition of CTS-1 systems using LLNL contract to support deliveries to SNL this fiscal year. Two platforms totaling nearly 3 petaflops were delivered and installed at SNL. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Applications Development Environment for Sierra Early Arrival System		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 6/30/17		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The Scalable Applications Preparation (SAP) effort within the Applications Development Environment and Performance Team (ADEPT) will deploy a capable application development environment on Sierra early arrival system that enables ASC code developers to effectively utilize the platform to prepare for the new capabilities and requirements of Sierra. As part of the effort, SAP will extend the knowledge base, documentation, and training necessary to actively support needs of tri-lab code teams, the AAPS team, and COE efforts in porting, tuning, and characterizing codes on the Sierra ID system in preparing for Sierra's arrival. Essential environment elements such as compilers, MPI libraries, GPU runtimes, performance analysis tools, debuggers, and system job launch capabilities will be deployed, integrated, and coordinated. In addition to IBM- and NVIDIA-provided software, third-party vendor and open-source community provided elements will be used to provide capabilities for understanding performance of codes targeting the Power8+ and NVIDIA GPU features. Deployment of development environments tools will include initial testing and validation of the tool suite.		
Completion Criteria: Deployment and documentation of the Sierra early arrival system development environment for the system users.		
Customer: ASC Integrated Codes sub-program		
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: The Sierra early arrival system.		

Milestone (ID#TBD): File System Selected for Sierra		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Complete Sierra/CORAL Clustered File System technical checkpoint and finalize file system selection.		
Completion Criteria: File system has been selected and procurement process has begun.		
Customer: ASC program		
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: ASC technical subject matter experts and procurement team.		

Milestone (ID#TBD): Trinity System Production Readiness		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 3/31/17		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: This milestone is based on meeting the requirements outlined in the Production Capability Readiness Milestone Description Guidelines. These requirements certify, through a review, that Trinity is ready for production capability. These requirements include: machine accessibility and integration, operational support, usage model has been demonstrated, user testing of applications has been done, system reliability has been achieved, application and I/O performance testing has been demonstrated, and a milestone review has been completed.		
Completion Criteria: Follows the ASC Level 2 milestone criteria for AT systems: that all of the topics identified in the description above have been successfully demonstrated for AT-class simulations. These requirements are specifically listed in the usage model for Trinity, which defines that the system has demonstrated an acceptable production user environment with all the associated support, testing, reliability and applications use of the system.		
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: CSSE, FOUS, Trinity		

Milestone (ID#TBD): Advanced Power Management on Trinity Platform		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: In anticipation of practical power consumption limits, the ASC program requires guidance for power management of future platforms and applications. The Trinity program's Advanced Power Management (APM) NRE project is delivering integrated power monitoring and control capabilities in the Trinity platform, building on prior work developing the Power API. This milestone will utilize these capabilities to collect information on the power usage characteristics of the ASC production workload running on Trinity. Methods will be developed to assist with understanding and will be applied to assess the potential impact of power-constraints in future ASC platforms. This milestone will lay groundwork for addressing the long-term goal of determining how to best use and operate future ASC platforms to achieve the greatest benefit subject to a constrained power budget.		
Completion Criteria: This milestone will be completed when the above-mentioned analyses are complete and a report or presentation is provided summarizing the key results and recommendations.		
Customer: NNSA NW program.		
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: Trinity APM project.		

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 FY16:

- Led the CTS-1 contract process and procured more than 77 CT system scalable units for the tri-labs.
- Extensively tested and resolved issues on pre-production CT systems.
- Deployed FY16 CT systems into LLNL simulation environments.

Planned Activities in FY17:

- Support and operate CT systems within LLNL simulation environments.
- Lead and manage the CT systems contract and any FY17 CT system procurements.

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 FY16:

- Coordinated the removal of Tri-Lab Linux Capacity Cluster (TLCC)-1 class system from service (Typhoon).
- Participated in the design and acquisition of CTS-1 systems along with the developing the deployment plan into both the classified and unclassified computing environments.
- Coordinated the design and project execution of the facilities (power and water cooling) to support CTS-1 deliveries.
- Installed initial CTS-1 system named Fire.

Planned Activities in FY17:

- Plan for the integration, acceptance, and deployment of the remaining CTS-1 systems (Ice and Snow) 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 CT 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 CT systems.

Accomplishments in FY16:

- Participated in lease-to-own approach to acquisition of CTS-1 systems using LLNL contract to support deliveries to SNL this fiscal year.
- Purchased and installed at SNL two platforms totaling nearly 3 petaflops.

Planned Activities in FY17:

- Fully deploy production-ready CTS-1 platforms with TOSS-3.
- Reassign existing TLCC-2 resources to other SNL programs and remove from ASC.

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 run best on this machine. This platform is used as a tri-lab resource for tri-lab stockpile stewardship milestones via the AT computing campaign (ATCC) process, which is run every six months when the next suite of codes will be ushered onto the machine.

Accomplishments in FY16:

- Ran two ATCC processes.
- Investigated optimal performance tuning.

Planned Activities in FY17:

- Run two ATCC 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. Under the CORAL procurement, LLNL 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 FY16:

- Provided technical coordination and contractual management for CORAL NRE and Sierra contracts.
- Began deployment of the Sierra early access system, which will help applications prepare for Sierra.
- Continued application preparations for Sierra system through the Sierra COE.
- Participated in activities related to ANL's Aurora system, within the CORAL procurement, providing insight into schedule issues for 10-nm fabrication process.

Planned Activities in FY17:

- Provide technical coordination and contractual management for CORAL NRE and Sierra contracts.
- Complete deployment of the Sierra early access system, which will help applications prepare for Sierra.
- Complete Sierra's clustered file system technical checkpoint and finalize file system selection.
- Conduct Sierra prototype system technical checkpoint.
- Participate in NRE and build subcontract activities for ANL's Aurora system as part of the broader CORAL collaboration activities.

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 FY17, 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 FY16:

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

Planned Activities in FY17:

- 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.

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 takes 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 the New Mexico Alliance for Computing at Extreme Scale, 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.

Accomplishments in FY16:

- Completed the Level 2 milestone, *Trinity-Haswell High Performance Computing System Delivery*.
- Continued to provide the technical coordination and management of the Trinity contract.
- Conducted Open Science runs on the Trinity system.

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

- Transitioned Trinity into the classified environment.
- Supported Intel Dungeon session as well as the Trinity COE meeting in April.
- Integrated application readiness test beds Mutrino and Trinitite to support porting of SNL and LANL codes to Trinity.
- Demonstrated performance improvements of SNL and LANL codes on Trinity.
- Managed SNL and LANL participation in the initial Advanced Technology Computing Campaign (ATCC-1).

Planned Activities in FY17:

- Deliver and install the Xeon Phi, or KNL phase of Trinity.
- Bring Trinity into full production and run the initial ATCC-1.
- Successfully complete Trinity Phase 2 Acceptance and Production Readiness Level 2 milestone.
- Work with tri-lab code teams to support improved code performance, analyst workflows, and proposal process and workload prioritization for simulations on future ATCCs.

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 a memorandum of understanding (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 FY16:

- Supported ongoing Cielo production work for CCC-8 and CCC-9.
- Commenced ATCC-1.

- Retired Cielo and removed from service.
- Retired Muzia, Cielo del Sur, and Cielito, the small proxy systems.

Planned Activities in FY17:

- Retire Cielo and remove from service.
- Retire Muzia, Cielo del Sur, and Cielito, the small proxy systems.

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 AT systems 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 third AT system (ATS-3) in the NNSA Computing Strategy, to be called Crossroads. 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 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 from each of the NNSA labs.

Accomplishments in FY16:

- Completed the CD-0 process in July 2015 to instantiate the project.
- Developed and released Draft Technical Specifications on Nov. 2, 2015, for ATS-3/Crossroads, including workflows and benchmarks.
- Completed a multilab, formal review of the Crossroads technical design and approach.
- Developed the acquisition and source selection process for the request for proposal (RFP).
- Completed Design Review for Draft Technical Specifications and released version 2 in Mar. 2016 based on review feedback as well as vendor feedback.

- Supported NERSC internal program review in Apr. 2016.

Planned Activities in FY17

- Complete the CD-1/3a documentation and obtain concurrence.
- Release the Crossroads RFP for vendor responses.
- Complete negotiations and draft a statement of work for the selected vendor.
- Conduct an independent cost review of the negotiated system.
- Complete the CD-2/3b documentation and obtain concurrence.
- Complete the formal contracts for the Crossroads systems.
- Define and initiate NRE work scope with ATS-3/Crossroads vendor(s).

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 and the associated NRE activities. The project will track HPC industry's 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 PathForward R&D projects.

Accomplishments in FY16:

- Completed final reviews of FastForward and DesignForward projects.

Planned Activities in FY17:

- Support ECP PathForward RFP planning, proposal evaluation, and reviews.

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 Structural Simulation Toolkit (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 FY16:

- Received Hybrid Memory Cube pre-production hardware from Micron and performed initial analysis in support of the Advanced Memory Technology Level 2 milestone.
- Deployed Intel KNL test beds and refreshed to Trinity production processors. Completed initial performance analysis for proxy applications, TRILINOS and Nalu.
- Refreshed Power 8 test bed to incorporate NVLINK 1, which allows SNL to evaluate progress on support for upcoming CORAL architecture.
- Improved online user guide and support for test bed users.
- Installed two small test bed systems sourced from Penguin using an industry standard open cabinet architecture. The initial deployment of the cabinet revealed electrical safety issues that were evaluated with the vendor and mediated by operational processes. These led to design modifications that were implemented in the CTS-1 deliveries.

Planned Activities in FY17:

- Assess impact of advanced architectures on SNL code performance in support of future procurements including CORAL/Sierra, ATS-3/Crossroads, and ATS-4.
- Update advanced architecture test beds with technology refresh; publish report and user guide on capabilities for R&D users.
- Obtain and evaluate larger OmniPath interconnect platform.

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 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 the FOUS section of this document.

Accomplishments in FY16:

- Released updates to TOSS (version TOSS 2.4-3, 2.4-4, and 2.4-7) that included security updates and bug fixes.
- Developed TOSS 3 (based on RHEL 7) with a beta release used for hardware check-out at the CTS-1 vendor site and deployed it on the delivered CTS-1 system.

- Developed enhancements to the IB OpenSM Monitoring Service, which provides an application interface to IB diagnostics, monitoring, management, and control functions.
- Developed a prototype software architecture that can continuously collect and analyze system data produced by the LDMS along with performance measurements from WCI proxy applications.

Planned Activities in FY17:

- 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.
- Develop enhancements to the OpenSM Monitoring Service, to provide an application interface to Intel's Omni-Path Architecture, the next-generation fabric deployed with the CTS-1 systems.
- Utilize prototype software architecture that continuously collects and analyzes system data produced by the LDMS to produce performance measurements from WCI proxy applications.

Applications Development Environment and Performance Team (LLNL)

The 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 FY16:

- Deployed the development environment on CTS-1 systems.
- Directed scalability and performance enhancements for the TotalView debugger, and provided enhanced, direct vendor support for debugging issues for ASC applications developers.
- Conducted co-design activities with vendors (including the development of a prototype of the OpenMP debug interface and working with the RogueWave team to integrate it into Totalview, as part of Sierra NRE).
- Informed the LLNL network procurement for CTS-1 by investigating the impact of a tapered fat tree on the performance of a representative suite of applications.

- Supported tri-lab applications teams with debugging and problem resolution for successful ATCC on Sequoia.
- Organized several tool training workshops held by the respective external tool developers at LLNL specifically targeting LLNL users, covering a wide range of tools, including O|SS, TAU, Scalasca, Vampir, and MUST.

Planned Activities in FY17:

- Continue code development environment support on all LLNL ASC platforms.
- Provide dedicated performance tuning expertise for applications teams on CTS-1 systems.
- Conduct Sierra SAP activities to assist tri-lab IC teams and their users, including efforts on early availability systems.
- Develop tool infrastructures to improve scalability and performance of applications and the code development environment.
- Conduct network investigations to understand the impact of congestion, multirail networking, and topology-aware placement within the new CTS-1 systems.

High Performance Computing Systems Research (LANL)

HPC systems research is a broad project focusing on near- to long-term research of all the components needed to support a rich environment for very large-scale applications and includes a strong effort in system resilience.

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 communication (OpenMPI), power, monitoring, and job launch. 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 and scheduling of HPC systems will be investigated.

Hardware reliability in HPC systems remains a challenge to characterize. Modeling and investigations into system resilience is key to this effort, along with continued work with a virtual machine (VM) fault injection tool to investigate both runtime software and ASC applications. Machine learning and 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 FY16:

- Completed initial integration of the Multidimensional Hashing Indexing Middleware (MDHIM) key value store into LAP code.
- Completed initial prototype of Legion/Realm runtime over libfabric.
- Implemented the initial ML classification of memory errors to provide more information than vendor-provided heuristics.

Planned Activities in FY17:

- Continue optimizing and hardening MPI, including threading issues with OpenMP and MPI, scaling, fault tolerance (sessions), and support for many-core nodes and heterogeneous nodes for Trinity, Sierra, and Crossroads.
- Continue development of system fault models for Trinity and Sierra.

Advanced System Test Beds (LANL)

The Advanced System Test Beds project 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 FY16:

- Integrate and support early access hardware and KNL for Trinity.

Planned Activities in FY17:

- Integrate Early Access Sierra type hardware.
- Expand support for KNL compute nodes and investigate emerging ARM-based processors for computational nodes.

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. An important aspect is providing lightweight services and functionality that does not compromise scalability and therefore performance. The focus will be on enhancing 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.

- SNL will continue to explore the relationship between the runtime system (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.

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.

Accomplishments in FY16:

- Completed negotiations with Cray on scope of Trinity program's APM NRE project. Worked with Cray to successfully deliver APM-2 milestone "Python Interfaces for Power API."
- Placed new NRE contract with Adaptive Computing to develop "Power Aware Job Scheduling" for Trinity.
- Released several new versions of the Power API specification, incorporating feedback received from Cray and the HPC community.
- Deployed Power API reference implementation on several SNL advanced architecture test bed clusters.
- Integrated Power API reference implementation with TOSS software stack. Power API is now part of the standard TOSS software distribution.

Planned Activities in FY17:

- Continue close collaboration with Trinity APM NRE vendors Cray and Adaptive to complete and deliver planned milestones.
- Collect information on the power usage characteristics of the ASC production workload running on Trinity and build tools for understanding (Level 2 milestone).
- Integrate lightweight kernel operating system functionality with Cray Linux environment and evaluate on Trinity Application Regression Test Bed (ART) system.

High Performance Computing Hardware Architecture Simulation (SNL)

The SST is a suite of tools enabling multiscale 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 FY16:

- Released SST version 6.0 in April 2016 with improved modularity and continuous integration testing as well as deployment on Github.
- Integrated Micro and Macro models in SST to utilize a common parallel core for discrete event simulation.
- Analyzed Hybrid Memory Cube and KNL architectures in support of the Advance Memory Architectures Level 2 milestone and ATS-1/Trinity.
- Added Dragonfly adaptive routing to Merlin interconnect module in SST to support evaluation of Trinity Aries interconnect.
- Demonstrated full-scale machine simulation models of key ASC computing platforms, pointing the way to accurate and scale-free performance predictions for important parallel workload algorithms.
- Characterized AMR applications congestion and completed performance statistics.

Planned Activities in FY17:

- Compare and contrast the multifidelity network/interconnect models, including the Merlin and Ember-MPI models.
- Evaluate potential ATS-3/Crossroads network and node designs.
- Deploy Knights Hill model and characterize performance.
- Explore performance characteristics of potential NV-RAM architectures.

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.

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 FY16:

- Provided SNL's Portals communication technology, which is the basis for new interconnection networks being designed and built for next-generation systems by two hardware vendors. Includes updates to Portals, MPI, and OpenSHMEM software support, with SNL OpenSHMEM forming the basis for vendor (Intel) delivered SHMEM implementations.
- Delivered first-of-a-kind benchmarking suite for one-sided data transfer multithreaded/tasking capable MPI software, significantly impacting modern MPI implementations in both functionality and performance.
- Demonstrated the impact of one-sided communication overlap on current systems and provided multiple solutions to reduce impact to negligible levels.
- Developed Qthreads runtime support for AMT parallel programming models, MPIQ.

Planned Activities in FY17:

- Further advance Portals technology to meet future generation networking requirements.
- Develop new lightweight communication methods that are compatible with existing networking software stacks, to improve existing communication performance and enable future communication semantics.

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 FY16:

- Investigated the ability of AMT runtimes (for example, Legion and Charm++) to deliver performance in the presence of the local performance variability common to extreme-scale resilience activities.

- Demonstrated how synchronous resilience activities must be in order to keep overheads low. Also, showed the role of collective inter-arrival time in sensitivity to synchronization.
- Merged the Local Failure Local Recovery (LFLR) idea for traditional MPI programming model into FENIX, a framework developed by Rutgers University. In collaboration with Intel, the SNL-Rutgers team is working on the specification of FENIX APIs to serve the resilience of a wide range of MPI-based applications.
- Began development of the LFLR framework for MPI-X model (LFLR-X). The test code can accommodate task-based recovery for application state.
- Identified aspects of hardware properties such as checkpointing efficiency that could impact broader resilience strategies incorporating silent-error tolerance.
- Examined simple checkpoint-based approach for mitigating silent errors. Began adapting interpolation-based robust conjugate gradient solver to a more complex nonlinear problem for more realistic performance evaluation.
- Began implementation of source-based fault injection representing silent errors in processor floating-point operations.
- Demonstrated more realistic extensions of robust solver techniques for tolerating memory bit flips, including reduction in overhead with increased solver scale and complexity.

Planned Activities in FY17:

- Extend and validate the resilience-activity-as-CPU-detour model developed previously to account for resilience mechanisms found in dynamic and task-based runtimes (for example, Charm++, Legion, and UINTAH).
- Demonstrate scalable message logging framework using LFLR-X.
- Release FENIX 1.1 (with minor fix from 1.0) to support new capabilities based on feedback from collaborators.
- Demonstrate candidate API for introducing mitigation of memory bit flips in a maintainable way, in SIERRA's Nalu fluid code or ExaCT's SMC combustion code.
- Conduct initial testing of parallel performance for Nalu or SMC with mitigation of injected memory errors on at least 10,000 cores.

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 FY16:

- Continued 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.

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

- Deployed HPSS 7.5.1 on pre-production HPSS core server platforms and metadata disk subsystems to fully test in advance of upcoming production deployments.
- Finished repack of T10K Gen1 media to T10K Gen2 media to minimize data loss due to aging media; ejected/destroyed Gen1 media to reclaim slot capacity in libraries.
- Deployed SCF HPSS disk subsystems to increase disk cache capacity and disk file residency; deployed 40Gbe 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 technology refresh, and data stewardship.

Planned Activities in FY17:

- 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 tape mover clusters to modernize remaining HPSS archival subsystems and retire existing five-year-old tape movers.
- Evaluate and potentially select and procure HPSS Core Server platforms and metadata disk subsystems to refresh existing five-year-old systems.
- Evaluate and potentially select and procure enterprise tape drives to leverage technological advances that would increase archive capacity and bandwidth.
- 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 FY16:

- Participated in and led aspects of the continuing community software development efforts of Lustre 2.5 and beyond.
- Enhanced ZFS-based Lustre performance in support of user and purge performance.
- Procured and put into production a Lustre file system with new hardware and software architectures leveraging Just a Bunch of Disk (JBOD) arrays and multiple metadata servers.
- Supported LLNL production Lustre deployments.
- Actively participated in CORAL NRE File System and Burst Buffer working groups in preparation for the Sierra machine; launched file system futures investigation.

Planned Activities in FY17:

- Support the development, testing, deployment, and production use of new Lustre versions in classified and unclassified environments.
- Augment drive monitoring and management in support of JBOD arrays under Lustre.
- Improve ZFS drive resilvering rates in support of enhanced availability of ZFS-based Lustre file systems.
- Continue to make general performance improvements to ZFS-based Lustre.

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 tests 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 FY16:

- Integrated Intel OMNI-Path support into TOSS 3.0.
- Tested Mellanox EDR with collective offload.
- Installed RedHat Power8 Little Endian on bare hardware.
- Performed OpenPower dual socket testing.
- Supported Broadwell for CTS-1.

- Conducted functionality testing of Cavium on RHEL7.

Planned Activities in FY17:

- Conduct functionality testing of latest version of Intel KNL.
- Test Broadcom ARM64 hardware with RedHat 7.
- Test Nvidia Pascal with RedHat 7 on x86_64, Power8, and OpenPower.
- Conduct Intel Skylake testing on RedHat 7 (if available).

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, SIO middleware development and support, SAN development and deployment, and archive.

The file systems element of the project provides end-to-end, high-performance networking and storage 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.

Accomplishments in FY16:

- Application readiness team (ARTeam) supported integrated code work on Sequoia, addressing file system and scaling issues for performance on that platform.
- Initially developed and deployed first production campaign storage system (called MarFS), which was specifically focused on the scaling of the erasure coded object store.

Planned Activities in FY17:

- Continue the hardening and testing of MarFS.
- ARTeam will support integrated code work with porting to the KNL partition of Trinity.
- Demonstrate running the EAP Trinity Demonstration Problem at its full scale using HIO to DataWarp with an N-M IO pattern.
- Release a HDF5-HIO bridge and demonstrate running LAP with this code on Trinity.

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 FY16:

- Developed initial version of FreTT data transfer tool and asserted SNL copyright.
- Tested and installed HPSS version 7.5.1 on operational systems at SNL.
- Supported technology refresh through repack of T10K generation 1 tapes to newer denser T10K generation 3 media. This activity prolongs data lifetimes and reduces demand on tape silo locations for equivalent data volume.
- Completed initial 40GE adaptor testing and data transfer using SNL networking and servers.

Planned Activities in FY17:

- Further develop HPSS beyond 7.5.
- Conduct initial production evaluation tests for FreTT data transfer tool.
- Evaluate potential of advanced technology features such as non-volatile random access memory (NVRAM) or solid state disk in various HPSS modules.
- Perform initial wide area network data transfer testing at 40GE rates using dedicated LANL to SNL link.

Scalable Data Services (SNL)

The Scalable Data Services project provides R&D to investigate and contribute to an open-source framework for software-defined storage (SDS). SDS enables the composition of custom data-management and storage services to address the specific needs of applications and adapts to rapid changes in the memory and storage

architectures. SDS provides an alternative to POSIX that allows, for example, a diverse set of data-model representations; consistency, synchronization, and resilience semantics that match the requirements of the application; a building-block architecture model that allows for the rapid development of custom services; and lightweight, user space, services that are instantiated and torn down in response to application needs.

This project will collaborate with ANL, LANL, and Carnegie Mellon University to develop the building blocks necessary for SDS, and it will extend ongoing and past ASC research projects such as the ATDM/DataWarehouse, CSSE/Sirocco file system, and the lightweight file systems project to develop data services that address the specific needs of ASC applications using the SDS building blocks.

Accomplishments in FY16:

This is a new project for FY17.

Planned Activities in FY17:

- Design a low-level "storage pool" service that supports the integration of diverse types of storage media (for example, memory, NVRAM, disk, and tape).
- Design an analytics-integration service that provides an interface for deploying analytics into the data-management layer. This could be as simple as annotating data on the fly and as complex as deploying transformations/filters into the data flow.

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 FY16:

- Supported large-scale data analysis and visualization activities, including supporting ASC scientists with creation of visuals and movies for presenting and analyzing scientific data.
- 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.
- Procured 20 additional nodes for an expansion of the unclassified visualization server to meet customer needs for more GPUs for training and porting activities.
- Developed a framework for tuning memory-mapped I/O with non-native page sizes for improving the runtime performance of the data-intensive streamline tracing application.
- Developed a visualization workflow using VisIt for rendering training images from Arbitrary Lagrangian-Eulerian (ALE) simulations that can be used for deep reinforcement learning on the TrueNorth neuromorphic chip.

Planned Activities in FY17:

- Support ATCC, Grand Challenge, and other LC users with visualization and data analysis activities, including creation of visuals and movies for presenting and analyzing scientific data.
- Continue to maintain a stable data analysis and visualization environment across LC platforms.
- Provide operational support for all visualization theaters, including demonstrations for high-level visits.
- Pursue research activities in data compression and topological methods.
- Exploit research results in data analysis and visualization for ASC simulations, including data compression, topological methods, and machine learning algorithms.
- Implement and deliver visual analytics capabilities for simulation data and machine learning features using web technologies, such as D3.js, through the Lorenz MyLC 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 FY16:

- Released new versions of Hopper and Chopper that provided a faster metadata scanning option for Lustre, MPI-based file transfers and deletes, reduced resource usage, and improved usability.
- Extended the MyLC dashboard to support a focus-area view, where the user can display just items related to specific types of information and topics.
- Implemented Lorenz components for viewing NFS data and metrics as they relate to specific user jobs on the HPC systems.

Planned Activities in FY17:

- Release new versions of Hopper and Chopper with a focus on performing smart deletes within large datasets, automatically generating debug and diagnostic reports, and expanding directory viewing flexibility and general usability enhancements.
- Extend the MyLC dashboard to support display of information from the Cassandra or equivalent center monitoring repository.
- Investigate use of new frameworks for building re-usable Lorenz-based user interface components that can be easily incorporated into portlets or external applications.

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, data reduction, visualization, and data-driven (Hadoop/MapReduce based) post-processing analysis and visualization.

Accomplishments in FY16:

- Deployed Paraview client to users and investigated the use of Paraview for rendering of IC visualizations via backend compute nodes of HPC clusters.
- Delivered with X-Theoretical Design (XTD) division users an in situ visualization of direct 3D imagery and plots of Lagrangian tracers.
- Demonstrated ParaView + Intel's OpenSWR scaling on Trinity: 1024 nodes, 8192 MPI ranks, 128 spheres per rank, each with ~1M polys results in 1099274649612 polygons/frame (~ 1.1 trillion polygons).

Planned Activities in FY17:

- Expand XDT division user base for in situ targeting workflow optimizations.
- Investigate replacement/virtualization of end-user visualization drops.

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 Trinity and subsequent ACES platforms.

Current tools include scalable data analysis software released open source through ParaView and the 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 FY16:

- Completed Level 2 milestone, *Demonstrate and Evaluate Advanced Analysis, Visualization, and Input/Output Capabilities for the SIERRA Toolkit*.
- Integrated Catalyst in situ analysis into numerous new SIERRA Toolkit-based codes Nalu and Aero. Developed Catalyst adapter for SNL rarified gas dynamics code SPARTA.

Planned Activities in FY17:

- Expand capabilities and use of analysis tools (ParaView, Catalyst, Slycat) into ASC mission codes.
- Integrate Slycat ensemble analysis tool into Sandia Analysis Workbench.
- Continue ParaView and Catalyst releases, with production support.

Projects for Beyond Moore's Law (WBS 1.2.3.5.6)

The Beyond Moore's Law product will evaluate potential NNSA Defense Programs' applications of computing technologies that go beyond Moore's Law scaling and Von Neumann architectures. The ASC program will investigate the application of non-CMOS-based logical devices, as well as quantum and neuromorphic computing algorithms and hardware to NNSA computing needs. The goal is to gain a detailed understanding and investigate the best technical approaches and benefits of these emerging technologies for NNSA applications and a roadmap for their integration into ASC computing platforms.

Beyond Moore's Law Computing (LLNL)

ASC's Beyond Moore's Law (BML) program evaluates potential NNSA applications of computing technologies that go beyond Moore's Law scaling and Von Neumann architectures. The application studies focus on NNSA Defense Programs' needs. The BML program is motivated by the NSCI call for "Coordinated research and a technical path forward regarding an effective post Moore's Law computing architecture." This ASC project will investigate the application of both quantum annealing and neuromorphic computing approaches to NNSA computing needs. The objective of the neuromorphic project is a detailed understanding of the technical approaches and benefits of neuromorphic computing for NNSA applications and a roadmap for their integration into ASC computing platforms. The objective of the Quantum Computing (QC) program is to provide a pathway for exploring QC for ASC applications, including applications work as well as evaluation of emerging hardware. The scope of this project includes research, development, and evaluation of prototype computing systems and algorithms, as well as developing potential industry and academic collaborations.

Accomplishments in FY16:

- Organized and hosted a neuromorphic computing workshop Feb. 23–24, 2016. Conclusions from the workshop show that there are clear areas in which neuromorphic computing can make important contributions to NNSA applications, and ASC should plan and execute a program to evaluate NNSA applications in more detail using test beds and simulations. A 15-page report from this workshop was delivered to NNSA on Mar. 31, 2016.
- Received IBM NS16e TrueNorth board along with associated tools and software support packages (Mar. 31, 2016). Performed acceptance testing on both the hardware as well as the toolkits. (Mar. 11, 2016).
- Accomplished significant progress on the exploration of algorithms on TrueNorth architecture, with emphasis on approaches with potential positive impact on simulation for stewardship.

Planned Activities in FY17:

- Develop a plan with input from SNL and LANL and establish collaborations for evaluating machine learning and numerical algorithms of NNSA interest on TrueNorth and other emerging neuromorphic architectures.
- Demonstrate access to the TrueNorth test bed from LANL and SNL.
- Demonstrate algorithms and applications on TrueNorth and generate evaluation reports for each application. These applications are intended to show use of machine learning techniques along with algorithms that are currently in use in ASC code development.
- Evaluate quantum annealing algorithms and methods addressing NNSA/ASC problems of interest, including on the D-Wave class of architectures.
- Design and test quantum emulation approaches to solving specific physical problems of NNSA interest, focusing on Qubit arrays at LLNL.
- Evaluate emerging QC commercial efforts and planned products, assess timelines for potential ASC application impact, and develop industry collaborations.

Emerging Technologies (LANL)

This project explores algorithm mapping of areas of interest to emerging "Beyond Moore's Law" technologies, focusing mainly on Quantum and Neuromorphic computing. Efforts in quantum computing are two fold. First, the project looks to create tools and programming models to simply the programming of these emerging computers. Second, the project looks to explore computational methods that can aid in basic algorithms of interest to ASC. LANL will work to create collaborations with university and other external partners. Management of the research, collaborations, and operation of the D-Wave system is also included in this scope. The Neuromorphic efforts look to provide insight via machine learning into large simulation datasets of interest to ASC and will explore other applications of Neuromorphic computing that may aid in the ASC mission.

Accomplishments in FY16:

- Procured and sited D-Wave Quantum computing system.
- Hosted multiple D-Wave instructional classes for LANL and DOE.

Planned Activities in FY17:

- Begin investigating mapping of algorithms to D-Wave programming model.
- Manage system and allocations to tri-lab and external collaborators.
- Collaborate on Neuromorphic research on the TrueNorth system at LLNL.

Non-Conventional Computing Technologies (SNL)

This project explores the potential value of emerging device technologies, computer architecture concepts, and models of computing for performing large-scale scientific computing applications a decade or more in the future. The investigations under this project span three technical thrusts: Beyond-Moore General-Purpose Computing (BMGPC), Neural-Inspired Computing (NIC), and Quantum Information Processing (QIP).

BMGPC is digital or analog computer systems designed to use CMOS transistors and/or alternate "bit-level" devices and new architectural concepts to support highly energy-efficient, general-purpose "classical" computing. BMGPC also entails architectural changes, such as moving processing close to the data; new techniques, such as using error correction for logic; and new models of computing, such as probabilistic computing. Three driving objectives of BMGPC are 1) greatly increasing energy-efficiency, 2) identifying new paths for exponential scaling of system performance; and 3) maintaining the ability to execute existing ASC applications codes.

The latter two thrusts (NIC and QIP) entail highly non-conventional models of computing that have the potential to provide exceptional computational performance on particular classes of computations. These approaches to computing are unlikely to support efficient, general-purpose computing.

Accomplishments in FY16:

- Completed training (10–15 SNL staff) on using the D-Wave system.
- Submitted several papers for presentation at the IEEE International Conference on Rebooting Computing, to take place in Oct. 2016.
- Wrote an article on applications of artificial neural networks to general purpose computing, which was published in the Jun. 2016 issue of *IEEE Computer Magazine*.

Planned Activities in FY17:

- Contribute to LLNL-led effort to finalize the roadmap for neural-inspired computing R&D.
- Investigate using the D-Wave Systems computer for materials simulation and data analysis.
- Implement algorithms on the TrueNorth platform at LLNL and evaluate performance.
- Develop and test a set of SST components for simulation of TrueNorth and other NIC architectures employing post-CMOS devices.
- Extend and apply computational simulation capabilities for characterizing emerging switching devices.
- Extend PIMS computer system design to support error-corrected logic.
- Draft a technical plan for the design and development of a "beyond exaFLOPS" computer.

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 2015, and through quarter 3, fiscal year 2016, are reflected below for the FOUS subprogram.

- Achieved Building 654 “Beneficial Occupancy” May 13, 2016. (LLNL)
- Made progress on integrating CTS-1 machine on classified network. (LLNL)
- Continued to provide 24x7 support for TLCC-2 systems, Sequoia, and Vulcan. (LLNL)
- Developed with Splunk the ability to correlate various authentication logs to identify issues and potential threats. (LLNL)
- Finished migration, reorganization, and consolidation of existing hotline tools and documentation to the Confluence server. (LLNL)
- Completed the Open Science effort on Trinity (Phase 1 of ATS-1). (LANL)
- Deployed production version of Fire CT system clusters. (LANL)
- Completed conceptual design for Exascale-Class Computer Cooling Equipment (ECCCE) project. (LANL)
- Launched new dashboard for ASC user support. (LANL)
- Renegotiated Wide Area Network contract and completed switchover to new supplier. (SNL)

- Took delivery of and installed two 6-scalable unit (SU) CTS-1 machines: Cayenne on the SNL classified network (SCN) and Serrano on the SNL restricted network (SRN). (SNL)
- Supported tri-lab Level 2 milestone, *Deploy a Common Computing Environment for CTS-1 Platforms*, by successfully integrating TOSS onto an early deployment CTS-1 machine. Integration was done remotely from SNL while the CTS-1 machine was still at Penguin's factory location. (SNL)
- Deployed improved user support model in anticipation of Trinity General Availability, including expanded phone coverage throughout the HPC service desk shift. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Sierra "Early Delivery System" Deployed on Secret Restricted Network		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 6/30/17		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The Sierra early delivery system will be installed and an early user will successfully run a code on the system.		
Completion Criteria: An early user successfully runs a code on the Sierra early delivery system installed on the SCF network.		
Customer: ASC Sierra early delivery 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: CSSE, IC, ATDM		

Milestone (ID#TBD): Complete Sierra Facilities Electrical and Mechanical Design		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Sierra electrical and mechanical design will be complete. Necessary contracts will be written and let to initiate construction.		
Completion Criteria: Sierra electrical and mechanical design drawings complete and approved. Contract for outside construction services written.		
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.		
Supporting Resources: IBM must provide Sierra technical specifications no later than December 2016 to enable timely completion of this milestone. FOUS staff, Facilities and Infrastructure staff, vendor services contract, outside purchase order subcontracts will also support the milestone.		

Milestone (ID#TBD): Next-Generation Infrastructure Plan		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 6/30/17		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: Produce a plan incorporating necessary elements (user support, networking, storage, facilities) for satisfying HPC infrastructure demands for FY18 and beyond computational systems. The plan will address a changing storage paradigm, incorporate new networking capabilities, and ensure LANL user support scales with the demands of pre-exascale workloads.		
Completion Criteria: A report or viewgraphs proposing a multi-year strategy to address future needs.		
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: LANL FOUS staff, CCS Division staff, and X Division staff for review		

Milestone (ID#TBD): Submit Exascale-Class Computer Cooling Equipment Critical Decision 2/3 Package		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: Having completed conceptual design, this milestone will submit required documentation for readiness for construction approval of exascale-class computer cooling equipment project.		
Completion Criteria: CD-2/3 required materials submitted to Headquarters.		
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: LANL FOUS staff, CCS Division staff, and X Division staff for review.		

Milestone (ID#TBD): Application of Performance Analysis Tools on SNL ASC Codes		
Level: 2	**Fiscal Year**: FY17	**DOE Area/Campaign**: ASC
Completion Date: 9/30/17		
ASC nWBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: Using currently available tools, specifically including those developed and/or provided within the CCE, characterize the performance of at least two codes (to be determined) on CTS-1 and Trinity platforms. Identify strengths and gaps in the performance analysis tool set and, in concert with application readiness support for new platforms, provide insights into possible paths for increasing platform productivity via code performance improvements.		
Completion Criteria: Deliver analysis results to product owners of the targeted codes and publish usage guidelines for the application of the various tools.		
Customer: SNL Trinity ATCC and CTS-1 projects and users, IC, ATDM.		
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 user and application support staff, IC code-team experts, CSSE and ATDM subject-matter experts.		

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 FY16:

- Continued FY15 procurement support, contract management, and next-generation computing planning.
- 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 for the six Academic Alliance Centers.
- Supported SC16, FASTMath/ParaDiS, and SUstained Performance, Energy & Resilience (SUPER) Institutes.
- Supported Scientific Discovery through Advanced Computing (SCIDAC) projects HAXTON and NUCLEI.

Planned Activities in FY17:

- Continue FY16 procurement support, contract management, and program planning.
- Support annual HPC Operations Review meeting with SC.
- Support bi-annual PSP meetings.
- Manage PSAAP II program for the six Academic Alliance Centers.
- Support SC17, FASTMath/ParaDiS, and SUPER Institute, Fernbach Fellowship, Rensselaer Polytechnic Institute Fellowship, and Summer Student Program.
- 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 FY16:

- Supported PSAAP II collaboration efforts at LANL.
- Participated in PI and PSP meetings.
- Provided LANL support for HQ.

Planned Activities in FY17:

- Support PSAAP II collaboration efforts at LANL.
- 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 laboratories and with the external community. Management within the laboratory includes the interface with DSW, including the LEPs and Alts, as well as day-to-day management of ASC program activities. 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 FY16:

- Organized and hosted a mid-year Predictive Engineering Science Panel (PESP) meeting and moved the primary review to the first quarter of the fiscal year to provide a better interface to other laboratory reviews.
- Managed Leidos contract to provide organizational support and program coordination for HQ and tri-lab activities.
- Supported the ASC Principal Investigators (PI) meeting in Las Vegas in Feb. 2016.
- Organized attendance, booth, and meeting logistics for the SC15 conference.

- Supported programmatic needs of the PSAAP II program and the DOE ECP.

Planned Activities in FY17:

- Organize and host the FY17 PESP meeting.
- Support ASC program activities to develop and publish a new ASC strategic document and produce a right-size program baseline.
- Organize and lead the ASC PI meeting to be held in late summer or fall of 2017.
- Support the programmatic needs of the PSAAP II program.
- Support the ECI, including joint planning and execution with DOE/NNSA and DOE/SC laboratories.
- 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 organizational 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 Nuclear Weapons Complex ASC-related activities is covered.

Accomplishments in FY16:

- Adapted Finite Element Discrete Ordinates Radiation Transport module to COMSOL Multiphysics.
- Completed proof of principle Godiva bare metal sphere with eigenvalue and Kobayashi benchmark for fixed source.
- Applied virtual reality, 3D modeling, and 3D laser scanning to provide a virtual immersive environment to support operations training.
- Set up the zSpace virtual reality system and approved it for the Y-12 network. The system is based on the Virtalis Visionary Render virtual reality software and provides an engaging environment to view and interact with 3D stereoscopic scenes.

- Developed a process to translate laser-scanned images of Y-12 facilities into 3D poly meshed models. The models can be scripted and programmed to simulate hazardous area access and training scenarios in a safe and highly realistic environment.
- Procured a FARO Freestyle hand scanner and set up to make detailed scans of process equipment and control panels.
- Built 3D computational model of plutonium electrorefining equipment (with designs similar to those used at LANL and AWE) and determined the primary current distribution during operation. The studies showed areas of localized polarization on electrodes consistent with experimental results observed on operating Pu electrorefiners. Presented the results at the JOWOG 22-2 focused exchange at AWE-Aldermaston in Sep. 2016.
- Presented "Y-12 ASC Activities Overview" at the 2016 Principal Investigator's Meeting and discussed collaborations for production agencies with KC-NSC.

Planned Activities in FY17:

- Acquire COMSOL license for HPC cluster and refine fissile solution COMSOL model to cluster.
- Investigate several benchmark excursion experiments; perform an optimization study.
- Investigate other applications of CAD-based finite element coupled radiation transport and multiphysics.
- Continue to work with LANL and LLNL pyroprocessing experts to develop models for transuranic electro-refining cells. Discuss alternative designs during JOWOG 22-2 that can be implemented to minimize localized polarization of the electrodes in Pu electrorefiners.
- Complete pyroprocessing studies that will concentrate on evaluating possible design changes to support resumption of electrolytic transuranic processing at LANL.
- Continue to support Y-12 representative participation in Nuclear Weapons Complex activities such as NWEATIM, the ASC PI meeting, and other Nuclear Weapons Complex-wide technical interchanges.

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 FY16:

- Began integrating CTS-1 machine on classified network (projected completion FY16Q4).
- Continued to provide 24x7 support for TLCC-2 systems, Sequoia, Vulcan, Lustre, and NAS file systems.
- Completed deployment of new NAS home directory hardware in open computing facility (OCF) and SCF centers.
- Deployed new non-"purgeable" project space in OCF and SCF environments.
- Documented requirements for identity management and began evaluating tool options.
- Developed with Splunk the ability to correlate various authentication logs to identify issues and potential threats.
- Implemented real-time blacklisting of hosts attempting to break into LC systems.
- Implemented the RSA middleware software layer as a web service.

Planned Activities in FY17:

- Integrate three Sierra early delivery systems to CZ, RZ, and SRD networks (FY16Q4).
- Integrate additional SRD CTS-1 machine, RZ CTS-1 machine, and CZ CTS-1 machine with five scalable units of CZ machine dedicated to the Alliances.

- Provide ongoing 24x7 support for CTS-1, TLCC-2 systems, Sequoia, Vulcan, Lustre and NAS file systems.
- Retire Sierra, Juno, and Muir.
- Deploy file system resources in support of unclassified CTS-1 (FY16Q4).
- Select replacement identity management tool and convert existing workflows.
- Deploy a new instance of the RSA one-time password service for LC's RZ.
- Automate security service verification and validation.

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 FY16:

- Continued to provide ongoing support services for hotline operations, documentation, and training.
- Finished migration, reorganization, and consolidation of existing hotline tools and documentation to the Confluence server.
- Made significant progress in the reorganization and updating the internal and external LC websites.
- Developed in-house knowledge and expertise in accelerator technologies and their software libraries in preparation for providing support to the user community in their use.

- Provided requirements to Service Now administrators in preparation of moving from the Front Range Incident Management system to Service Now.

Planned Activities in FY17:

- Continue to provide ongoing support services for hotline operations, documentation, and training.
- Deploy new internal and external LC websites.
- Continue to develop in-house knowledge and expertise in accelerator technologies and their software libraries in preparation for providing support to the user community.
- Develop documentation and training materials on the use of the new CTS-1 commodity clusters and Sierra early delivery machines.
- Assist users in migrating from TLCC-2 systems to new CTS-1 machines.
- Migrate from Front Range to the Service Now Incident Management system (dependent on Chief Information Officer's organization's (LivIt) readiness).

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 FY16:

- Achieved Building 654 Beneficial Occupancy and completed construction.
- Completed demolition of under floor for SCF CTS-1 machine siting.
- Completed power installation for SCF CTS-1 machine.
- Began Sierra facility preparation projects: Building 453 Power Modernization and Advanced Development Required for Cooling Loop.
- Completed network design and hardware procurement for extending 100G network to Building 654.
- Completed architecture, design, and hardware procurements for implementing building redundant networking within the HPC facilities.
- Continued to enhance network monitoring tools and Splunk statistics gathering.

- Continued to evaluate 100G and greater networking technologies and encryption for use in the HPC SANs and WAN.

Planned Activities in FY17:

- Complete Sierra facility preparation projects: Building 453 Power Modernization and Advanced Development Required for Cooling Loop.
- Complete design of Sierra mechanical and electrical infrastructures.
- Begin build of Sierra mechanical and electrical infrastructures.
- Implement and deploy building redundant network design within HPC facilities.
- Deploy support networks for CTS-1 and Sierra early delivery HPC machines.
- Upgrade external network bandwidth to 40G within the HPC and internet facing firewalls.
- Continue to evaluate 100G and greater networking technologies and encryption for use in the HPC SANs and WAN.

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 FY16:

- Supported HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers.
- Completed implementation of 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.

Planned Activities in FY17:

- Support HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers.

- Transition the Trinity Phase 2 system into production operation.
- Provide continuing security enhancements and requisite software stack upgrades on deployed compute platforms, file systems, 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. 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 FY16:

- Provided 24x7x365 Tier-1 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 Phase 2 platform and Trinity KNL test system.
- Supported installation of the Fire CTS-1 system.
- Supported decommissioning of the Cielo and Typhoon platforms, and associated file systems.
- Supported the acquisition of several mission-essential equipment components.

Planned Activities in FY17:

- Provide 24x7x365 Tier-1 support, operations, and monitoring of compute and file system resources.
- Provide support for installation of Snow and Ice CTS-1 systems.
- Provide continuing hardware maintenance, assessment of system life spans, and guidance on hardware maintenance burdens for computing equipment.
- Fold the Trinity Phase 2 system into the production operations, monitoring, and maintenance program.
- Support installation of Phase 2 of the Trinity 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 Strategic Computing Complex at LANL (SCC) for deploying these CT systems. The integration and deployment activities will focus on the following areas: system/OS, file systems, interconnect, external network including next-generation backbone (NGBB), regression testing, and monitoring.

Accomplishments in FY16:

- Integrated the Trinity Phase 1 system into the production computing environment, and entered this system into production service in support of ATCC-1.
- Accepted delivery of the Trinity Phase 2 system, and integrated this system into the production computing environment.
- Accepted the Fire CTS-1 system, and integrated the system in the production computing environment.
- Provided continuing integration and top-tier diagnostic and troubleshooting support for standing production compute platforms and file systems.

Planned Activities in FY17:

- Support integration, testing, and diagnostic activities for final acceptance of the Trinity system.
- Complete the integration of all CTS-1 platforms (Ice and Snow) into production service.
- Support design and requirements gathering activities for the future Crossroads system.
- 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 FY16:

- Developed a user dashboard website for improved user awareness of status and availability of HPC resources.
- Completed and published CTS-1 machine documentation.
- Provided user consulting, support, and notification services for the ASC user community.

Planned Activities in FY17:

- Facilitate user transition from Cielo to the Trinity AT system during the first and second ATCCs.
- 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, distance computing (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 FY16:

- Completed conceptual design for the exascale-class computer cooling equipment project.
- Demonstrated the reliable operation of the newly installed warm-water cooling system for the Trinity platform.
- Provided ongoing operations, maintenance, and configuration of electrical and mechanical systems for ASC information technology systems and facilities.
- Decommissioned the power and networking infrastructure for the Typhoon and Cielo computing platforms.

Planned Activities in FY17:

- Complete system and project design, and initiate construction, for the exascale-class computer cooling equipment project.
- 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 advanced technology platform resource allocations and coordinates with tri-lab peers in establishing priority scheduling, if required. This project coordinates the integration and deployment of CT systems into SNL's production computing environment, in collaboration with WBS 1.2.3.6.3 CCE. Support of CCE common service and environment decisions and configuration management activities are also provided. FOUS supports administration and operations activities associated with the ACES project, ART platforms including Mutrino (Trinity) and Muzia (Cielo).

This project has expertise in operating CT 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 FY16:

- Took delivery and installed two 6-SU CTS-1 machines: Cayenne on the SCN and Serrano on the SRN.
- Supported tri-lab Level 2 milestone, *Deploy a Common Computing Environment for CTS-1 Platforms*, by successfully integrating TOSS onto an early deployment CTS-1 machine. Integration was done remotely from SNL while the CTS-1 machine was still at Penguin's factory location.
- Provided ongoing operations support for all production platforms and test bed systems, including: Pecos, Chama, HPSS systems, file systems, WC Tool, the Synchronized Account Request Automated Process (SARAPE), and the advanced architecture test beds.
- Integrated LDMS into Trinity system to enhance data gathering and support efforts aligned with power management research.

Planned Activities in FY17:

- Provide ongoing production operations of Cayenne, Serrano, HPSS systems, file systems, WC Tool, SARAPE, and the advanced architecture test beds.
- Remove Pecos and Chama from ASC support.
- Develop plans for infrastructure technology refresh of networks, file systems, and storage devices to be accomplished in FY18.
- Continue integration of LDMS onto Trinity to include instrumenting burst buffer, KNL, and interconnect.

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 Service Management (ITSM) principles and practices; 2) ITSM incident, problem, change, 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 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 FY16:

- Provided user support for SNL and tri-lab ASC computing systems, primarily Cielo and Sequoia, but including Trinity starting in Jan. 2016.
- Deployed improved user support model in anticipation of Trinity General Availability, including expanded phone coverage throughout the HPC service desk shift.
- Continued to develop expertise in support of next-generation architectures, software environments, and new debugging tools.
- Continued development of HPC Web portal and ACES Web portal.

Planned Activities in FY17:

- Provide user support for SNL and tri-lab ASC computing systems.
- Continue to develop expertise in support of next-generation architectures and software environments.
- Ramp up support for Trinity Phase 2 and Sierra.
- Complete review of service management tool set and any transition to new tool(s).
- Continue development of HPC and ACES Web portals.

Facilities, Networking, and Power (SNL)

The Facilities, Networking, and Power project supports maintenance and improvements to the facilities and infrastructure servicing the HPC systems (CT 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 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 EStar Awards from the DOE Office of Sustainability Support.

Accomplishments in FY16:

- Removed two TLCC-1 platforms from the computing facility.
- Modified the computing facility to support two 6-SU CTS-1 systems, bringing in overhead 480-V power and running under-floor feed pipe for liquid cooling.
- Renegotiated WAN contract and completed switchover to new supplier.

Planned Activities in FY17:

- Initiate design activity for expansion of Building 725 Data Center for future HPC platforms.
- Continue facilities support for all production systems and advanced architecture test bed systems.
- Relocate the Systems Operations Center to provide additional data center floor space for HPC systems.

- Add additional power and cooling to isolate HPC infrastructure from Enterprise infrastructure.

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 CT 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 activity enables increasingly 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 that can impact the performance of both applications and throughput on those platforms, both current and future; 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 FY16:

- Released updates to TOSS (version TOSS 2.4-3, 2.4-4, and 2.4-7) that included security updates and bug fixes.
- Developed TOSS 3 (based on RHEL 7) with a beta release used for hardware check-out at the CTS-1 vendor site; verified functionality of TOSS 3 by successfully recompiling and running the full complement of Synthetic Workload applications on the system at the vendor site.

- Completed deployment of 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 utilizing the Pavilion test framework and employed them during CTS-1 hardware check-out at the vendor site.
- Completed test and debug phases of the Lightweight Distributed Metric Service (LDMS) v3.

Planned Activities in FY17:

- 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 machines (based on RHEL 7.X).
- Continue SLURM support efforts through tri-lab collaboration; continue investigation of running SLURM-only on TOSS systems.
- Develop identified collaborative system software tasks, including integration of virtualization, 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 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 O|SS, TAU, MUST/Vampir, and Valgrind.

Accomplishments in FY16:

- Continued efforts in collaboration with AT system projects to further integrate the tools ecosystem to support both CT and AT architectures. With Trinity, it included successful integration and KNL support. Broader tools assessment for Sierra in concert with CORAL tools working group. Other efforts in this area included TAU support and enhancements in the MUST MPI correctness tool.
- The O|SS development contract has made progress for providing threading support using the OMPT tools API. They have included specialized support for successful and optimal installation of LANL's AT system, Trinity. Work includes: CUDA accelerator analysis support development, improving the O|SS frontend capabilities, continuing to port the build system to Cmake, and pursing a Spack build of O|SS to ease installation complexity. There have also been continued efforts at the labs to further improvements for the memory profile experiment.
- Delivered enhanced debugging capabilities through the TotalView BIGCAT collaboration to enhance ReplayEngine support for ASC applications and scope requirements for supporting debugging of performance portability layers in use by the tri-labs.
- The MPI tuning study and infrastructure for CTS-1 with SNL applications was completed by SNL and included varying MPI implementations and tuning parameters with resulting data collated to draw tactical and strategic conclusions impacting CCE MPI, CCE TOSS, and the application teams. Praetor was the Python framework developed to shepherd the study. It required minimal user interaction with built in resilience to support test failures. It is currently at version 1 and will be released as open source.
- Provided MVAPICH and Open MPI enhancements and bug fixes, support for MPI_THREAD_MULTIPLE and native PMIx support; added PSM2 support in MVAPICH2 to support CTS-1; and ported PMI2 to SLURM for faster MPI startup in CTS-1.

Planned Activities in FY17:

- Continue development and support efforts for debuggers, performance analysis tools, and MPI as programming models and architectures evolve.
- Continue focus on updating functionality of the tool ecosystem to support upcoming ATS architecture needs, including continued testing and integration for accelerator support, port to OpenPower, Nvidia GPUs, and Intel's KNL.
- Continue integration of the tri-lab tools build environment and broader integration of the Spack HPC software package manager to enhance efficiency for support and developer environments.
- Broaden collaboration on compiler benchmark suite with COEs and engage vendors.

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, and to provide common tools among the tri-labs for usage reporting and resource management.

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 cross-site authentication and resource approval through enhancements to the SARAPE system. SARAPE 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 tri-lab tools required to manage accounts within and among the tri-labs. As part of the inter-site HPC deployment, it offers a service catalog through which collaborators can view and request accounts and services available in the shared environment.

The Workload Characterization Tool (WCTool) provides the ability for common resource usage reporting to ASC HQ and within the labs, and additionally assists in optimizing management of ASC computing resources.

Accomplishments in FY16:

- Delivered quarterly tri-lab workload characterization reports to NNSA (WC Tool).
- Released WCTool version 2 at LLNL.
- Set up SARAPE application in a virtual server environment and applied reverse proxy, to provide greater application stability and increased performance.
- Implemented domain-restricted access controls within SARAPE rather than at a network level, to increase flexibility in managing an expanded customer base.
- Gathered initial requirements for producing metrics from the SARAPE database.
- Enhanced automated data gathering using SAML/CryptoCard login as an initial step to providing improved MySARAPE information to registered users.

Planned Activities in FY17:

- 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.
- Update WCTool to accommodate new operating system and software package requirements.
- Continue to operate and improve the tri-lab SARAPE process for all remote access account requests and implement evolutionary improvements.
- Complete HSPD-12 authentication integration, to meet DOE requirements..

- Enhance metrics available to SARAPE host processing agents, based on feedback to initial deployment of the capability.
- Improve MySARAPE user experience through redesign, simplification, and expanded information.

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 FY16:

- Provided HPC procurement, system administration, and operational support.
- Cleared new HPC computer room space to prepare for modifications and accreditation.

Planned Activities in FY17:

- Provide HPC procurement, system administration, and operational support.
- Install electrical and networking for single scalable unit CTS-1 machine.
- Integrate and operate CTS-1 machine in new computer room.

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 FY16:

- Provided ongoing operations, maintenance, and administration of compute facilities, platforms, and file systems.
- Provided ongoing support and expanded deployment of software resources.
- Planned and procured Hail and Lysander.
- Completed facilities work to prepare for installation of Hail and Lysander.

Planned Activities in FY17:

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

Special Purpose Computing (SNL)

The National Security Computing Center (NSCC) at SNL provides CT-class computing, 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 FY16:

- Obtained security plan approval for special category computing.
- Provided ongoing support for tri-lab use of NSCC capabilities and resources.
- Decommissioned Cielo del Sur.

Planned Activities in FY17:

- Provide support for tri-labs use of NSCC resources, including implementation of new capabilities.
- Expand NSCC utilization to include other DOE/NNSA program offices (NA-80).
- Upgrade file system infrastructure to accommodate future platforms.

Appendix H: Academic Alliance Centers

Accomplishments listed below represent those of FY16.

University of Utah

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

Accomplishments in FY16:

- Completed a massive rewrite of the Uintah runtime to incorporate KOKKOS and multiple C++11 constructs.
- Developed a hybrid UQ methodology for achieving predictivity with extrapolation that integrates Bayesian and Bound-to-Bound methodologies.
- Developed physics components for coal char oxidation with mass transfer that integrates into the LES framework.
- Made predictions of a 1000MWe coal-fired boiler that was newly commissioned and made operational by an industrial partner.
- Developed and taught two new graduate courses on center themes (V/UQ and coal combustion) to students at the three participating universities.

Planned Activities in FY17:

- Rework the I/O subsystem in the Uintah runtime to include Portable Information Data eXchange (PIDX).
- Perform V/UQ analysis on coal laboratory, burner and full-system scales to identify consistent char, devolatilization, and mineral matter transformation models.
- Advance the development of mineral matter transformation and radiation models to be both computationally efficient and to reduce the model form uncertainty.
- Perform first round of engineering optimization of operating parameters on the 1000MWe coal-fired boiler while working with industrial partner.
- Teach the two center graduate courses (V/UQ and coal combustion) and expand its reach to include not only students at the three participating universities but to also include participants from the three DP labs.

University of Illinois, Urbana-Champaign

Center for Exascale Simulation of Plasma-Coupled Combustion

Accomplishments in FY16:

- Demonstrated computer science tools on the principal application PlasComCM. These include tools for optimizing just-in-time re-compilation, over-decomposition for balancing and robustness, detailed vectorization analysis, multihardware portability, multirate time integration, and a general annotation pre-processing orchestration framework.
- Increased physics fidelity of predictive models based on UQ and detailed bench-top diagnostics.
- Advanced experimental diagnostics to inform model development.
- Hosted the Workshop on Exascale Software Technologies (WEST) in Albuquerque, which provided a platform to broadly discuss forthcoming tools, challenges, and best practices to realize performance at exascale.
- Made multiple visits to NNSA labs by XPACC staff, including placement of eight PhD students in internships, several of which have fostered ongoing interactions.

Planned Activities in FY17:

- Apply the center's computer science performance tools on a full-scale predictive simulation.
- Finish repackaging of PlasComCM (into PlasCom2) as a standard software (C++, Fortran, MPI, OpenMP) application that facilitates flexible performance at exascale though the use of interoperable CS tools with embedded UQ.
- Implemented UQ for full-scale prediction, informed by low-dimensional physics-targeted experiments.
- Evaluate options for engineering impact of predicting plasma-coupled combustion for high-speed propulsion systems.
- Continue to foster close interactions with NNSA labs through additional student internships, staff and faculty visits, and the planning for a possible follow-on to WEST.

Stanford University

Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment

Accomplishments in FY16:

- Implemented scalable high-fidelity radiation transport in Soleil-MPI.
- Demonstrate Center application (Soleil-X) using complete software stack (Liszt/Legion) on heterogeneous (CPU/GPU) and multinode settings.

- Assessed effect of particle-size uncertainty on flow/particle exit temperature.
- Performed first validation of flow and particle statistics against Stanford experiment.
- Developed and verified a converging point-particle model for momentum transfer between fluid and particles.

Planned Activities in FY17:

- Demonstrated Soleil-X scalability on O(100) nodes of Titan using combinations of multicore CPU GPU.
- Performed UQ of full-integrated simulations (with low-fidelity radiation transport).
- Completed first heated experiment and comparisons with Soleil-MPI.
- Compared low- and high-fidelity radiation models in Soleil-MPI.
- Implemented collision model and development point-particle treatment for heat transfer.

University of Florida

Center for Compressible Multiphase Turbulence

Accomplishments in FY16:

- Conducted direct numerical simulations of shocks interacting with multiple particles, including structured arrays and random packs of particles. The foundation for a new drag model called the Pairwise Interaction Extended Point Particle (PIEP) model was developed and tested.
- Developed auto-tuning methods for simultaneously optimizing for software and hardware parameters and demonstrated them using the gem5 simulator; also developed performance and energy models for key kernels on multicore and GPU architecture, and used the models to develop performance and energy-based load balancing algorithms for hybrid architectures.
- Developed CMT-bone as a proxy app for CMT-nek, validated it using the Veritas tool from LLNL, and developed an initial CPU-GPU implementation.
- Enhanced the Behavioral Emulation (BE) methodology to scale beyond the device level, performing experiments on, and predictive BE simulations for, thousands of cores on the Cab and Vulcan systems at the LLNL.
- Fully integrated the Uncertainty Budget team with the other teams, finding multiple opportunities to implement surrogate multifidelity techniques. This enabled applying the full range of UQ and reduction to the mesoscale shock tube simulations, and initial UQ for CS and exascale activities.

Planned Activities in FY17:

- Focus on implementing a new Lagrangian-based volume fraction model, and a collision/compaction model to capture more accurately the early moments of the demonstration problem.
- Conduct additional simulations of shock interaction with particles at lower Mach numbers to emulate SNL experiments as well as conduct simulations at blast conditions.
- Complete ongoing work on the discontinuous Galerkin treatment of high-order operators and make CMT-nek into a compressible Navier-Stokes solver using spectral elements.
- Develop optimized implementations of CMT-bone and evaluate performance energy tradeoffs. Develop load-balancing algorithms for particle-in-cell (PIC) problems for hybrid multicore processors and implement them.
- Perform UQ for demonstration problem and fingers phenomenon, to guide new experiments in the micro-, meso-, and macroscales, and to validate by blind prediction CS and exascale models.

Texas A&M University

Center for Exascale Radiation Transport

Accomplishments in FY16:

- Completed calibration and validation of Impurity Model 1.
- Implemented extruded unstructured triangular mesh capability in PDT transport code.
- Achieved efficiency of roughly 70% for weak scaling of PDT/STAPL on O(1M) processors.
- Implemented initial load-balancing algorithm for unstructured meshes.
- Released STAPL (Standard Template Adaptive Parallel Library) to NNSA labs.

Planned Activities in FY17:

- Recalibration and Validation of Impurity Model 1 with new experimental configuration isolating the experiment from neutron room return.
- Take delivery of the neutron generator and make it operational.
- Complete experiments and UQ analysis for next two levels of the UQ hierarchy.
- Extend load-balancing algorithm to efficiently treat highly non-uniform meshes.
- Improve efficiency of diffusion-based convergence acceleration methods in PDT.

University of Notre Dame

Center for Shock Wave Processing of Advanced Reactive Materials

Accomplishments in FY16:

- Developed poro-visco-plastic constitutive model for metal powders, and advanced image-based modeling pipeline for micro-to-macro simulations.
- Developed MTL-HPX library, and released HPX-5 v3.0.
- Developed spectral Wavelet code that eliminates non-essential collocation points, an Eulerian-Lagrangian solver, and *PGFem3D* MPI+HPX with residual vector and stiffness matrix multithreading.
- Performed full-system demonstration simulations and validated them against experiments.
- Performed predictive damage multiscale simulations using up to 128,000 cores.

Planned Activities in FY17:

- Develop chemo-thermo-mechanical coupled solver.
- Complete image-based modeling framework and demonstrate both bottom-up and top-down approaches.
- Improve HPX scalability and performance, and develop HPX++.
- Validate Quantities of Interest (for example, shock speed and final microstructure) for confined impact test of Ni/Al powders.
- Finalized DAKOTA immersion with the C-SWARM framework and implement Bayesian inference for parameter estimation.

National Nuclear Security Administration

Appendix I: Construction and Capital Equipment

The following table shows current ASC construction projects and capital equipment purchases.

				$ in Thousands							
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**FY17**	**FY18**	**FY19**	**FY20**	**FY21**	**Future**	**Contact**
LLNL	B-654 Livermore Computing Facility	Construction-GPP	N/A	9,875							Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B-453 Power Modernization	Construction-GPP	N/A	4,875	2,000						Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B-453 Advanced Development Cooling Loop	Construction-GPP	N/A	1,500	100						Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B-654 AHU and Adiabatic Cooling Solution	Construction-GPP	N/A	1,500	1,100	400					Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B-453 Sierra Site Prep	Construction-GPP	N/A	4,875	2,875	2,000					Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Commodity Technology System (CTS-1)	MIE		44,000	11,000	10,763	1,000				Matt Leininger 925-422-4110 leininger4@llnl.gov
LLNL	Advanced Technology System (ATS-2)	MIE		170,000	21,900	56,000	53,200	11,000	6,800	3,000	Bronis du Supinski 925-422-1062 desupinski1@llnl.gov
LLNL	Commodity Technology System (CTS-1)	MIE			10,500	10,500	8,498				Matt Leininger 925-422-4110 leininger4@llnl.gov

				$ in Thousands							
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**FY17**	**FY18**	**FY19**	**FY20**	**FY21**	**Future**	**Contact**
LLNL	Advanced Technology System (ATS-2)	MIE			21,900	56,000	53,200	11,000	6,800	3,000	Bronis du Supinski 925-422-1062 desupinski1@llnl.gov
LANL	Trinity: Acquisition of Trinity (ATS-1) system	Capital Equipment, MIE	CD-2/3b	170,000	17,106	7,426	5,132	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Crossroads: Acquisition of Crossroads (ATS-3) system	Capital Equipment, MIE	CD-0 7/2015	170,000	9,000	13,400	55,250	54,200	21,600	16,550	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	PIIC: Physical Infrastructure Integration for Crossroads (PIIC)	Capital Equipment	CD-0	9,000	0	0	0	9,000	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Exascale Class Computer Cooling Equipment Project: Expansion of the SCC warm-water cooling capacity to 33MW prior to arrival of the Crossroads system	Capital Equipment, MIE	CD-0 9/2015	38,000	4,100	17,000	14,000	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	Title TBD: Expansion of SCC power capacity to facilitate use of all 40 MW available to the SCC wall, prior to the arrival of ATS-5	Capital Equipment, Line	TBD	TBD	0	0	0	0	TBD	TBD	Jason Hick 505-667-4477 jhick@lanl.gov
SNL	Construct Addition to Bldg. 725. Additional computing space (~18,000 sq. ft.) for HPC systems consolidation.	Construction - GPP	TBD	9,700	2,000	7,700	0	0	0	0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov