



Advanced Simulation and Computing

FY18 IMPLEMENTATION PLAN

Version 0

September 8, 2017

NA-ASC-127R-17-Rev.0-IP

LLNL-TR-738237

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	<i>Michel McCoy, Bill Archer, Scott Collis</i>	<i>Initial</i>	<i>Mark Anderson, Thuc Hoang</i>	<i>9/12/17</i>	<i>Initial release</i>

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

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

Contents

I. OVERVIEW	5
II. CORPORATE PROGRAM GOALS.....	7
III. MAJOR ACTIVITIES.....	8
IV. FUNDING GUIDANCE.....	14
V. DESCRIPTION OF PLANNED ACTIVITIES	15
VI. REPORTING REQUIREMENTS.....	18
VII. KEY EXECUTION YEAR REFERENCE DOCUMENTS.....	19
VIII. MAJOR RISKS AND HANDLING STRATEGIES	20
IX. POINTS OF CONTACT.....	21
X. APPROVALS	27
APPENDIX A: KEY TERMS.....	28
APPENDIX B: INTEGRATED CODES SUBPROGRAM (WBS 1.2.3.1).....	36
APPENDIX C: PHYSICS AND ENGINEERING MODELS SUBPROGRAM (WBS 1.2.3.2).....	37
APPENDIX D: VERIFICATION AND VALIDATION SUBPROGRAM (WBS 1.2.3.3).....	38
APPENDIX E: ADVANCED TECHNOLOGY DEVELOPMENT AND MITIGATION SUBPROGRAM (WBS 1.2.3.4)	39
Accomplishments.....	39
Level 2 Milestone Descriptions	41
Projects for the Next-Generation Architecture and Software Development Product (WBS 1.2.3.4.2)	45
Projects for the Inter-Agency Co-Design Product (WBS 1.2.3.4.3).....	66
APPENDIX F: COMPUTATIONAL SYSTEMS AND SOFTWARE ENVIRONMENT SUBPROGRAM (1.2.3.5)	70
Accomplishments.....	70
Level 2 Milestone Descriptions	72

Projects for the Commodity Technology Systems Product (WBS 1.2.3.5.1).....	74
Projects for the Advanced Technology Systems Product (WBS 1.2.3.5.2).....	76
Projects for the System Software and Tools Product (WBS 1.2.3.5.3)	82
Projects for the Input/Output, Storage Systems, and Networking Product (WBS 1.2.3.5.4)	90
Projects for the Post-Processing Environments Product (WBS 1.2.3.5.5)	96
Projects for Beyond Moore's Law (WBS 1.2.3.5.6)	101

**APPENDIX G: FACILITY OPERATIONS AND USER SUPPORT
SUBPROGRAM (WBS 1.2.3.6)** **105**

Accomplishments.....	105
Level 2 Milestone Descriptions	107
Projects for the Collaborations Product (WBS 1.2.3.6.1).....	110
Projects for the System and Environment Administration and Operations Product (WBS 1.2.3.6.2)	115
Projects for the Common Computing Environment Product (WBS 1.2.3.6.3)	125
Projects for the Special Purpose Facilities, Systems, Operations, and Support Product (WBS 1.2.3.6.4).....	131

APPENDIX H: ACADEMIC ALLIANCE CENTERS..... **133**

APPENDIX I: CONSTRUCTION AND CAPITAL EQUIPMENT **139**

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 incorporates nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. The suite of data analyzed comes from activities including stockpile surveillance, experimental research, and development and engineering programs. This integrated national program requires the continued use of experimental facilities and the computational capabilities to support the SSP missions. These component parts, in addition to an appropriately scaled production capability, enable the National Nuclear Security Administration (NNSA) to support stockpile requirements. The ultimate goal of the Stewardship Program, and thus the Advanced Simulation and Computing (ASC) Program, is to ensure that the U.S. maintains a safe, secure, and effective strategic deterrent.

The ASC Program is a cornerstone of the SSP, providing simulation capabilities and computational resources to support the annual stockpile assessment and certification process, 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 work requires a balance of resources, including technical staff, hardware, simulation software, and computer science solutions.

The ASC Program focuses on increasing the predictive capabilities in a three-dimensional (3D) simulation environment while maintaining support to the SSP. The Program continues to improve its unique tools for understanding and 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 from the original test base. While the focus remains on the U.S. nuclear weapons program, 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.

A Nuclear Posture Review (NPR) has been initiated by the President, and when finished this report will reaffirm the role of the NNSA in maintaining the U.S. nuclear deterrent. In areas essential for stockpile life extensions and stewardship, key investments will continue to be 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.

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. 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 FY18. The Initial IP effective October 1, 2017, should be consistent with: 1) the Department's Base Table when operating under a Continuing Resolution (CR); and 2) the final signed appropriation or full-year CR once enacted.

¹ 2010 Nuclear Posture Review Report, April 2010, p. 42.

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	FY2018 Target	Endpoint Target
Advanced Simulation and Computing Program	Reduced Reliance on Calibration	63%	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 FY18 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 FY18 Stockpile Stewardship and Management Plan². 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 composed of 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 NNSA's Science Program, which provides the experimental data that inform development of new

² U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2018 Stockpile Stewardship and Management Plan—Biennial Plan Summary*, Report to Congress, 2017.

models used in simulation codes. PEM also includes activities that are directly aligned with initiatives to support non-stockpile nuclear security missions, some examples of which are provided in the IC section above.

The **Verification and Validation** (V&V) subprogram provides evidence that the models in the codes produce mathematically credible answers that reflect physical reality. V&V focuses on establishing soundness in integrated simulation capabilities by collecting evidence that the numerical methods and simulation models are being solved correctly, and whether the simulation results from mathematical and computational models implemented into the codes are in sync with real-world observations. 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.

This area delivers science-based assessments of the predictive capability and uncertainties in ASC integrated performance, engineering, and specialized methods to support the needs of the Stockpile Stewardship Program. 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 methodologies. This will be a major focus area for the V&V subprogram in FY18.

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 ASC integrated design codes (IDCs) is falling significantly when they are used on the latest HPC platforms, and this trend is expected to accelerate unless mitigated. Three major challenges to address through investments in this subprogram include: 1) the radical shift in computer architecture, 2) maintenance of the current millions of lines of IDCs 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 FY18 major activities for the ASC Program include:

- Further develop nuclear performance assessment codes for boost and secondary performance; safety codes to address multi-point safety issues; and engineering assessment codes for hostile, normal, and abnormal environments.
- Adapt existing codes to new architectures, and migrate current design and safety codes to run efficiently on hybrid computer architectures.
- Bring to Production Readiness the CT systems and complete the Trinity-Knight's Landing (Phase 2) system for the tri-labs' production computing environment in early FY18 to address stockpile stewardship issues and to advance predictive science.
- Continue to adapt/develop models for components built by advanced/adaptive manufacturing techniques.
- Further develop reactive flow models for HE detonation and burn that capture grain scale material heterogeneity and are computationally efficient.
- Deploy ASC Sierra platform at Lawrence Livermore National Laboratory (LLNL), with anticipated system acceptance in Q4 FY 2018, to address stockpile stewardship issues and to advance predictive science.
- Steward two major facility modernization efforts, one at Los Alamos National Laboratory (LANL) and another at LLNL, through the initial stage of the acquisition process. These two facilities will support deployment of advanced architecture systems on the path to exascale and at exascale.
- Expand the portfolio of the ATDM subprogram to include new simulation capabilities to evaluate hostile environment response, accelerated development of next-generation IDCs, and increased research and development (R&D) investment with the computer industry to address exascale technical challenges.
- 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.
- Broaden development of V&V protocols for algorithms running on hybrid HPC architectures.
- Implement quality assurance controls to ensure material and nuclear databases are correctly updated and maintained.
- Coordinate with the Department of Energy (DOE)/Advanced Scientific Computing Research (ASCR) office on the DOE Exascale Computing Initiative (ECI) to advance research and development technologies that will eventually enable procurement and utility of an exascale-class HPC platform.
- Maintain mentoring program for early career staff.

The ASC budgetary increase in FY18 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](#),” which directs NNSA to transition the stockpile to fewer warhead and bomb types. 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 ATDM subprogram consolidates the investments Congress directed, starting in FY14, for exascale activities into a unified portfolio 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 Work Authorization (WA), funding will be distributed through the existing Approved Funding Program (AFP) process described in DOE Order 135.1A and NNSA BOP 001.331, or successor Order/BOP, as applicable. The AFP is adjusted on an as needed basis for the execution of congressionally approved programs, projects, or activities (PPAs). Specific work activities are authorized via this document, with incremental funding plans for each site authorized via the AFP and obligated via formal contract modification. The work contained herein will be funded on an incremental basis. The contractor is authorized to expend up to the dollar amount indicated in the Program/OCL funding table or as otherwise noted in the AFP allotment, whichever is less, and subject to the availability of funds in the M&O contract. Note: As indicated in Section 4.0, the contractor is required to notify the contracting officer, in writing, when they recognize that they will exceed the estimated cost by 10%, insufficient actual and expected funding is available to continue work, or if their actual funding is insufficient to operate until further expected funding.

Program/Operational Control Level	President's FY2018 Budget Request	FY2018 CR Operating Target	FY2018 Enacted/Full-Year CR	Difference between Request versus Enacted
ASC Operations	\$709.244M	\$663.284M	\$663.284M	
ASC Construction	\$25.000M			

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 the subprograms' Level 2 Milestones. Successful progress toward completing these milestones is tracked on a quarterly basis. Progress toward completion of subprogram deliverables contributes toward an aggregate assessment of the program'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 FY18. 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 Milestones for FY18³

Sub-Program	ID#	Milestone Title	Complete Date	Site
IC	TBD	Develop and Demonstrate an Improved Workflow in the Integrated Code System	9/30/18	LLNL
IC	TBD	Demonstrate Progress on GPU Performance with ASC Production Codes	9/30/18	LLNL
IC, V&V	TBD	Verification and Validation of Fracture Modeling Capabilities with SPH	9/30/18	LLNL
IC, ATDM	TBD	Capability Assessment for Simulating Weapons Performance in Limited Hostile Environments	9/30/18	LANL
IC	TBD	Thermal/Mechanical Modeling for Crash and Burn Use Cases	9/30/18	SNL
IC	TBD	Electrical Analysis Calibration Workflow Capability Demonstration	12/31/17	SNL

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
PEM	TBD	HE Models for Non-Ambient Temperatures and Corner Turning	9/15/18	LLNL
PEM	TBD	Initial Delivery of GND/GIDI to WCI	9/30/18	LLNL
PEM, IC	TBD	Delivery of Improved Code Capabilities and Physics Models Supporting FY19 Predictive Capability Framework Peg Post	9/30/18	LANL
PEM, IC	TBD	Improve Replication of In-service Mechanical Environments	9/30/18	SNL
V&V	TBD	Modeling of X-ray Driven Ablative Response Experiments	9/30/18	LLNL
V&V	TBD	Statistics Software for Detonator Testing	6/30/18	LANL
V&V	TBD	Validation of the Models for Acoustic Vibrations During Reentry Using Data Provided by Sandia's Delivery Environments Program	9/30/18	SNL
ATDM	TBD	Demonstration of Modular Thermonuclear (TN) Burn in Multiphysics Applications	9/30/18	LLNL
ATDM	TBD	Integrate and Enhance Support for HPC Development Cycle	9/30/18	LLNL
ATDM, CSSE	TBD	Model of Faults of Large-scale Advanced Architectures	12/31/17	LANL
ATDM	TBD	Demonstrate ATDM Radiation-Hydrodynamics Capabilities	9/30/18	LANL
ATDM	TBD	Demonstrate ATDM Hostile IEPM Physics Capability	9/30/18	SNL
ATDM	TBD	Demonstrate ATDM Reentry Physics Capability	9/30/18	SNL
ATDM, CSSE, FOUS	TBD	Initial Capability of an Arm-based Advanced Architecture Prototype System and Software Environment	9/30/18	SNL
ATDM, IC, CSSE	TBD	Tri-lab Co-design Milestone: Improved Strategies for Development and Performance of Codes on ATS Platforms	9/30/18	SNL, LANL, LLNL

Sub-Program	ID#	Milestone Title	Complete Date	Site
CSSE	TBD	Sierra Applications Development Environment	8/31/18	LLNL
CSSE	TBD	Local Failure Local Recovery (LFLR) Resiliency for Asynchronous Many Task (AMT) Programming and Execution Models	9/30/18	SNL
FOUS	TBD	ECFM Design	9/30/18	LLNL
CSSE, FOUS	TBD	Sierra Initial Delivery (ID) System Acceptance	6/15/18	LLNL
FOUS	TBD	Submit Exascale-Class Computer Cooling Equipment Critical Decision 2/3 Package	12/31/17	LANL

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 provide 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 ASC documents are incorporated by reference:

- *ASC Strategy 2016–2026* (October 2016)
- *ASC Computing Strategy* (May 2013)
- *ASC Right Size* (September 2016)
- *ASC Business Plan* (July 2015)
- *ASC Co-Design Strategy* (February 2016)
- *ASC Simulation Strategy* (June 2017)
- MOU between DOE Office of Science (SC) and DOE NNSA Office of Defense Programs for the coordination of exascale activities (November 2016)
- *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 DOE Exascale Computing Project, which is available upon request from the ASC Program Office.



IX. Points of Contact

Headquarters Programmatic Points of Contact

Mark Anderson

NA-114, Director

Tele.: 202-586-5746

FAX: 202-586-7754

Mark.Anderson@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

Scott Collis, SNL, 505-284-1123, sscoll@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 James R. Stewart, SNL, 505-844-8630, jrstewa@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 James R. Stewart, SNL, 505-844-8630, jrstewa@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 James R. Stewart, SNL, 505-844-8630, jrstewa@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 James R. Stewart, SNL, 505-844-8630, jrstewa@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	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@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

WBS	Title	Contact
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
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 kmatous@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:

Mark Anderson,
NNSA ASC Program Director



Mark Anderson 9/5/17

Signature

Date

Concurrence:

Michel McCoy,
LLNL ASC Executive

Concurrence received

9/5/17

Bill Archer,
LANL ASC Executive

Concurrence received

8/31/17

Scott Collis,
SNL ASC Executive

Concurrence received

9/6/17

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
AID	Advanced Infrastructure for Debugging
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
ART	Application Regression Test Bed
ASC	Advanced Simulation and Computing (formerly ASCI)
ASIC	Application Specific Integrated Circuit
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

AWS	Amazon Web Services
B&R	Budget and Reporting
BE	Behavioral Emulation
BEE	Build and Execute Environment
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
CEED	Center for Efficient Exascale Discretizations
CFD	Computational Fluid Dynamics
CGNS	Computational fluid dynamics General Notation System
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
DAG	Directed Acyclic Graph
DisCom	Distance Computing
DOE	Department of Energy
DRAM	Dynamic Random Access Memory
DSW	Directed Stockpile Work
EAP	Early Access Program; Eulerian Application Project
ECCCE	Exascale-Class Computer Cooling Equipment

ECFM	Exascale Computing Facility Modernization
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
ECP-ST	Exascale Computing Project-Software Technologies
EMP	Electromagnetic Pulse
EMPIRE	ElectroMagnetic Plasma in Radiation Environments
ESSIO	Extreme-Scale Storage and Input/Output
F-SEFI	Fine-Grained Soft Error Fault Injection Tool
FEM	Finite Element Method
FET	Field Effect Transistor
FleCSI	Flexible Computation Science Infrastructure for Multiphysics
FOUS	Facility Operations and User Support
FPGA	Field-Programmable Gate Array
GPFS	Global Parallel File System
GPGPU	General-Purpose Graphics Processing Unit
GPU	Graphics Processing Unit
HAANA	Hardware Acceleration of Adaptive Neural Algorithms
HBM	High Bandwidth Memory
HIO	High Performance Input/Output
HPC	High Performance Computing
HPCDE	HPC Developer Ecosystem
HPSS	High Performance Storage System
HQ	Headquarters
I/O	Input/Output
IOSS	Input/Output SubSystem
IB	Infiniband
IC	Integrated Codes
ICD	Interagency Co-Design
IMEX	Implicit/Explicit
ICE	Illinois Coding Environment
IDC	Integrated Design Code

IEMP	Internal Electromagnetic Pulse
IP	Implementation Plan
IRB	Institutional Review Board
ITSM	Information Technology Service Management
JBOD	Just a Bunch of Disk
KC-NSC	Kansas City National Security Campus
KNL	Knights Landing processors
KVS	Key Value Store
L2	Level 2 (milestone)
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LANSCE	Los Alamos Neutron Science Center
LAP	Lagrangian Applications Project
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LDMS	Lightweight Distributed Metric Service
LDRD	Laboratory Directed Research and Development
LEP	Life Extension Program
LFLR	Local Failure Local Recovery
LIB	Laser-Induced Breakdown
LLNL	Lawrence Livermore National Laboratory
LLVM	Low-Level Virtual Machine
LOCA	Library of Continuation Algorithms
M&O	Management and Operating
MD	Molecular Dynamics
MDHIM	Multidimensional Hashing Indexing Middleware
ML	Machine Learning
MOU	Memorandum of Understanding
MPI	Message Passing Interface
MPICH	MPI over CHameleon
MST	Multisite Target



NAS	Network-Attached Storage
NCHCP	Non-conventional High-Performance Computing
NCI	National Cancer Institute
NERL	Neural Exploration & Research Laboratory
NFS	Network File System
NGBB	Next-Generation Backbone
NGCE	Next-Generation Computing Enablement
NGP	Next-Generation Platform
NIC	Neural-Inspired Computing; Network Interface Controller
NNSA	National Nuclear Security Administration
NNTI	Nessie Network Transport Interface
NPO	NNSA Production Office
NPR	Nuclear Posture Review
NRE	Non-Recurring Engineering
NSCC	National Security Computing Center
NSCI	National Strategic Computing Initiative
NVMe	Non-Volatile Memory Express
NVRAM	Non-Volatile Random Access Memory
nWBS	National Work Breakdown Structure
NWEATIM	Nuclear Weapons Engineering Analysis Technical Interchange Meeting
O SS	Open SpeedShop
OCR	Open Community Runtime
ORNL	Oak Ridge National Laboratory
OS	Operating System
OS/R	Operating System/Runtime
PCF	Predictive Capability Framework
PDE	Partial Differential Equation
PDX	Patient-Derived Xenograft
PECASE	Presidential Early Career Award for Scientists and Engineers
PEM	Physics and Engineering Models
PESP	Predictive Engineering Science Panel

PI	Principal Investigator
PIC	Particle in Cell
PIDX	Portable Information Data eXchange
PIEP	Pairwise Interaction Extended Point Particle Model
PIM	Processor in Memory
PIRT	Phenomena Identification and Ranking Table
PPT	Performance Prediction Toolkit
PSAAP	Predictive Science Academic Alliance Program
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
RMA	Remote Memory Access
RTS	Runtime System
RZ	Restricted Zone
SAN	Storage Area Network
SAP	Scalable Applications Preparation
SARAPE	Synchronized Account Request Automated Process
SASI	Standards and Architectures for Storage and I/O
SAW	Sandia Analysis Workbench
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
SEER	Surveillance Epidemiology and End Results
SFI	Significant Finding Investigation
SGEMP	System Generated Electromagnetic Pulse
SHMEM	Symmetrical Hierarchical MEMory
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
SQL	Structured Query Language
SQS	Special Quasi-random Structure
SRN	Sandia Restricted Network
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAPL	Standard Template Adaptive Parallel Library
STAT	Stack Trace Analysis Tool
SU	Scalable Unit
SUPER	SUSTained Performance, Energy & Resilience (SCiDAC Institute)
TAU	Tuning and Analysis Utilities
TCE	Tri-lab Computing Environment
TIK	Transient Ignition Kernel
TLCC	Tri-Lab Linux Capacity Cluster
TLS	Transport Layer Security
TOSS	Tri-Lab Operating System Stack
TRL	Technology Readiness Level
UB	Uncertainty Budget
UGT	Underground Test
UQ	Uncertainty Quantification
V&V	Verification and Validation



VERTS	Verification Test Suite
VM	Virtual Machine
VPIC	Vector Particle-in-Cell
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)
YAML	Yet Another Multicolumn Layout
ZFS	Zettabyte File System



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 Initiative (ECI) and aligned with the National Strategic Computing Initiative (NSCI) Executive Order. Performing the ATDM work within the scope of the DOE ECI allows for broader engagement in co-design activities and provides a conduit to HPC vendors to enable next-generation, advanced computing technologies to be of service to the stockpile stewardship mission.

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 2016, and through quarter 3, fiscal year 2017, are reflected below for the ATDM subprogram.

- Conducted contract negotiations for Path Forward, issued awards, and provided technical coordination and contractual management. (LLNL)
- Made extensive progress toward a next-generation software environment, through coordinated efforts with the Centers of Excellence (COE), standards committees, participation in Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) Non-Recurring Engineering (NRE) activities, and work in the areas of debugging, scheduling, I/O, and power. (LLNL)
- Improved software testing and deployment with a center-wide continuous integration service for users, package build/testing infrastructure, and many new Spack features to address code team needs. (LLNL)
- Partnered with code teams preparing for next-generation platforms on development and implementation of strategies for parallelism, data movement, and unified memory, as well as the extension of RAJA and Copy Hiding Application Interface (CHAI). (LLNL)

- Deployed a suite of development tools for performance visualization and analysis, correctness verification, and resilience, including Scalable Checkpoint Restart (SCR), mpiFileUtils, Spack, Archer, Caliper, and GOTCHA. (LLNL)
- Deployed and tested MarFS in concert with the FOUS subprogram (campaign storage—an intermediate storage tier between the parallel file system and archive storage). (LANL)
- Completed initial investigation and mitigation of software fault injection into the FleCSALE RISTRA/CDA proxy application. (LANL)
- Implemented “Control Replication” in Legion to support apparently sequential semantics in a fully distributed environment. (LANL)
- Completed on-site co-design development of Flexible Computation Science Infrastructure for Multiphysics (FleCSI) and Legion, which enabled a demonstration of multinode FleCSI on Legion via a Lax–Wendroff-based proxy application in collaboration with CDA. (LANL)
- Tracked and evaluated Flang (an open-source FORTRAN compiler) with ASC codes of interest. (Tri-lab)
- Developed a performance-portable sparse matrix–matrix multiplication algorithm that is significantly faster than vendor-provided libraries on both Knight’s Landing (KNL) and NVIDIA architectures. Vendors are adopting the new algorithms into their libraries. (SNL)
- Achieved an order-of-magnitude performance improvement for auto-decomposition of large-scale meshes. (SNL)
- Delivered a powerful embedded transient sensitivity analysis capability to the ATDM applications, which resulted in dramatically improved performance and accuracy. (SNL)
- Demonstrated large-scale visualization on the Intel KNL architecture by developers at SNL, Kitware, and Intel. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Integrate and Enhance Support for HPC Development Cycle		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 9/30/18		
ASC nWBS Subprogram: ATDM ASD		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
<p>Description: This effort will integrate and enhance a unified tool set focused on improving the efficiency of the entire software lifecycle (develop, test, deploy, monitor, analyze, tune). This builds on the FY17 efforts of the HPC Developer Ecosystem team for creating a common infrastructure for DevOps tasks such as monitoring, databases, continuous integration, and software deployment.</p>		
<p>Completion Criteria: An ASC code team member will verify that the team has demonstrated a unified tool set focused on improving the efficiency of the entire development cycle: develop, test, deploy, monitor, analyze, tune.</p>		
Customer: ASC		
<p>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.</p>		
Supporting Resources: ASD, CSSE, and FOUS Staff		

Milestone (ID#TBD): Model of Faults of Large-Scale Advanced Architectures		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 12/31/17		
ASC nWBS Subprogram: ATDM, CSSE		
Participating Sites: LANL		
<p>Description: This milestone will use the system log, error data, and machine learning to produce a fault model for Trinity. The goal is to produce the methodology that can then be applied to future large-scale machines to predict the level of resilience applications, system software, and hardware that is needed to run successfully. Specifically, this milestone looks toward ATS machines and exascale. This will advance ASC co-design efforts by impacting one or more of the following:</p> <ul style="list-style-type: none"> • Determining fault modes of existing and future machines • Determining the amount and type of data needed to produce these models • Providing a resilience target for Next Generation Code and IC codes • Evaluating resilience of current applications with fault injection 		
<p>Completion Criteria: Build a fault model for Trinity. Evaluate the resilience of at least one IC or ATDM code using a fault injector. Develop the specifications of the data required for ATS systems to enable building fault models.</p>		
Customer: IC/ATDM		
<p>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.</p>		
Supporting Resources: ATDM, CSSE staff		

Milestone (ID#TBD): Initial Capability of an Arm-based Advanced Architecture Prototype System and Software Environment		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 09/30/2018		
ASC nWBS Subprogram: ATDM/ASD, CSSE, FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
<p>Description: This milestone will deliver an initial capability for an Arm-based Advanced Architecture Prototype High Performance Computing (HPC) System and associated Software (SW) Environment. The initial capability will include vendor source selection; an executed contract and associated statement of work (SOW); an early delivery small-scale system suitable for developing the software environment; a full description of the software environment, including vendor and NNSA laboratory contributions, and its plan for its integration; and a beta delivery of the software environment running on the early delivery system. Sandia will organize and lead a tri-lab team to develop the requirements, architecture, and components of the software environment for the final system to be delivered in CY2019. This effort includes evaluation of the potential of an OpenHPC-based software stack and possibly other open standards approaches to meet future mission needs.</p>		
<p>Specific Deliverables:</p> <ol style="list-style-type: none"> 1) Release of Request for Proposals, including the specification of the technical requirements and targets for the final system to be delivered in CY2019. 2) Technical review of the proposals by a multilab committee (Sandia, LANL, LLNL, invited SC lab participants), and source selection. 3) Negotiated and signed contract for the final Arm prototype system. 4) Initial released report for version 1.0 of the Arm SW environment, documenting the requirements, architecture, and components of the software environment for the final Arm prototype system. 5) Early (beta) delivery and testing of the core software stack running on a small-scale, Arm-based testbed system, demonstrated by running a representative sample of tri-lab mini-apps. 		
<p>Completion Criteria: This milestone will be completed when the above-mentioned delivery, testing, and report are complete.</p>		
<p>Customer: ASC CSSE and FOUS programs at SNL, LANL and LLNL</p>		
<p>Milestone Certification Method:</p> <p>A program review is conducted and its results are documented. A report is prepared as a record of milestone completion.</p>		
<p>Supporting Resources: Advanced Architecture Testbeds project, Sandia Institutional Computing Program (facilities and infrastructure)</p>		

Milestone (ID#TBD): Tri-lab Co-design Milestone: Improved Strategies for Development and Performance of Codes on ATS platforms		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 9/30/18		
ASC nWBS Subprogram: ATDM, IC, CSSE		
Participating Sites: SNL, LANL, LLNL		
Participating Programs/Campaigns: ASC		
<p>Description: Obtaining the highest levels of performance from complex NNSA/ASC codes is becoming increasingly difficult because of the use of complex programming language features, the use of deep portability abstractions, the complexities of multiple interacting runtimes, and substantial increases in hardware complexity that frustrate efficient code generation. This FY18 milestone will perform an analysis of issues relating to optimizing the use of development toolchains including compilers, runtimes, and source-to-source translation to obtain higher performance on NNSA platforms in areas such as code generation, optimization, binary size, and compile/link times. The milestone will focus on a set of agreed upon metrics in this area and analyze them in the context of ASC-relevant application kernels, libraries, or full applications. Additionally, the milestone will develop potential strategies or recommendations for developers to improve NNSA/ASC application performance and/or reduce compile/link times. We will also engage with vendors to identify possible improvements for their toolchains and runtimes.</p>		
<p>Completion Criteria:</p> <ol style="list-style-type: none"> 1) Tri-lab report or annotated presentation detailing strategies for optimized performance utilizing the ATS and/or CTS development environments with a focus on the development toolchain. 2) Build and/or runtime performance improvements will be demonstrated for an ATDM or an IC code at each laboratory. 3) Deliver a suite of examples and tests that demonstrate significant development toolchain issues experienced by our codes allowing vendors and tool developers to better optimize their tools. 		
<p>Customer: ASC</p>		
<p>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.</p>		
<p>Supporting Resources: ASD, CSSE, and FOUS Staff</p>		

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

The Next-Generation Architecture and Software Development (ASD) 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 (CDA). The ASD 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 (NGCE) 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. NGCE project areas include system-level software, power-aware computing, debugging tools, and resource management. 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.

In FY17, this project also included work on performance analysis and tools, data analysis tools, as well as storage and I/O projects targeting new software architectures for defensive I/O and associated storage. These two elements will migrate to other projects in FY18.

Accomplishments in FY17:

- Made extensive progress toward a next-generation software environment through coordinated efforts with COE, standards committees, participation in CORAL NRE activities, and work in the areas of debugging, scheduling, I/O, and power.
- Participated in vendor and COE activities as well as *Forward efforts; coordinated on tools with the institutional Center of Excellence (iCOE) activities.
- Conducted a large-scale power variation study on new CTS-1 systems at scale.
- Released PRUNERS Toolset 1.0, which includes four individual tool components: Archer (an Open Multi-Processing (OpenMP) data race detector), ReMPI (a record and replay tool for Message Passing Interface (MPI)), NINJA (a network injector to

expose unintended MPI message races), and FLiT (a tester for compiler-induced floating-point variability).

- Used PRUNERS on LLNL codes such as ARES, HYDRA, MERCURY, and ParaDiS, which helped to identify the root causes of some highly elusive bugs in these codes.
- Advanced work on the next-generation resource manager, FLUX, and delivered a version capable of managing ensembles of tasks within a single job, e.g., for UQ and V&V workloads. Achieved a 133× increase in job throughput.
- Refined FLUX's rich resource model and demonstrated its flexibility to support various resource selection algorithms pertaining to exascale environments.

Planned Activities in FY18:

- Expand power and power variation studies to Sierra and Sierra Early Access (EA) systems.
- Understand the impact of on-node noise on ASC workloads.
- Build an HPC bug taxonomy based on the difficult-to-diagnose errors reported to the Advanced Infrastructure for Debugging (AID) team and identify further debugging gaps in exascale environments.
- Extend Archer to cover programming model-level concurrency and ReMPI to cover nondeterministic sources beyond MPI (e.g., OpenMP, use of clocks, etc.).
- Demonstrate the effectiveness of our fully hierarchical scheduling directly on UQPipeline-based UQ simulations of one or more Weapons and Complex Integration (WCI) applications.
- Harden the rich resource model as a new service to FLUX (a next-generation resource and job management framework) and improve scalability using a concept of pruned graph searching.
- Continue participation in vendor, COE, and co-design activities, as well as technical management of academic contracts and other external interactions.

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 Advanced Architecture and Portability 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 General-Purpose Graphics Processing Unit (GPGPU) programming, many-core programming, I/O, and parallel application development.

Accomplishments in FY17:

- Worked with deterministic transport codes ARDRA and TETON to help develop and implement a porting strategy. For ARDRA, we continued to develop and improve the nested loop abstraction within RAJA, and helped refactor the code to use it with CHAI. For TETON, we worked in coordination with NVIDIA and IBM to resolve compiler and programming model (CUDA FORTRAN and OpenMP 4.5) interoperability issues between KULL (C++) and TETON (FORTRAN) on Minsky nodes, then helped devise and begin implementing a strategy for parallelism and data movement.
- Worked with the ARES code team to help implement their unified memory porting strategy. Found a critical build error that was preventing compilation on early access Minsky nodes. Developed Clang compiler plugins to detect and warn on otherwise opaque RAJA and ARES implementation errors to speed development and limit debugging. Conducted performance analysis to identify bottlenecks.

Planned Activities in FY18:

- Work closely with multiphysics/hydrodynamics code teams (KULL, ARES, ALE3D, MARBL) to carry out cross-cutting activities identified by team leads as their highest priorities in achieving portable performance on Sierra.
- Continue to work with Monte Carlo transport code MERCURY and deterministic transport codes ARDRA and TETON to help with porting to Sierra.
- Continue to assist in the development of RAJA, CHAI, and Umpire (a resource management library for memory management in heterogeneous architectures) as needed to further the performance portability goals 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 the established set of proxy applications.

Production-Quality Tools Development Project (LLNL)

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 and AAPS teams on the other.

Accomplishments in FY17:

- Delivered the System for Performance Optimization Tracking (SPOT) tool, including a new visualization interface, to the ALE3D code team. The team is using it to track performance history of their nightly tests.
- Supported Spack adoption by the ARES, MERCURY, and KULL code teams with feature development, training, and support.
- Ported the SCR tool to Cray systems for use on Trinity by the MERCURY team, who have adopted SCR for their general production runs.
- Developed the prototype implementation for OMPD, a shared-library plugin that enables a debugger to inspect and control execution of an OpenMP program, which is needed to present a debugging standard to the OpenMP standards committee.
- Worked with researchers on deployment and releases of several tools, including SCR, mpiFileUtils, Spack, Archer, GOTCHA, and Caliper.

Planned Activities in FY18:

- Develop and support a centralized repository of Spack packages for WCI code teams, which will allow teams to share build recipes and packaging knowledge.
- Integrate SCR into WCI applications in preparation for using Sierra's burst buffers. Initial target applications include ARES, PF3D, and HYDRA.
- Integrate and support SPOT, a tool for performance tracking and visualization, into more WCI applications. Initial targets are ARES and ARDRA. Integrate the GOTCHA tool into Caliper, which will allow advanced profiling capabilities in tools using always-on performance analysis.
- Expand SONAR data collection capabilities with job queue and I/O history, and work with researchers on data analytics and predictive capabilities for the SONAR data.
- Continue the porting of tools to Sierra. Remaining high-value tools include Caliper, Stack Trace Analysis Tool (STAT), and memory profiling tools.
- Continue the ongoing effort of identifying the next generation of research tools, and working with code teams, AAPS, and management to prioritize tools for productization.

User Workflow and Modernization (LLNL)

The project's mission is to build infrastructure and applications that enhance the end-to-end productivity of HPC assets. The project focuses on building solutions to problems in three areas prioritized by users: problem setup, simulation management, and post-processing. The team mitigates risk by providing a set of components that encompasses best practices and common operations, allowing for custom workflows as needed. There are three active software projects: Siboka, components and infrastructure supporting user workflows, being demonstrated in the VV4ALE3D application; C2C, the contours-to-

codes project, focused on building a common geometry specification across codes that is compatible with the Los Alamos National Laboratory (LANL) contour format; and ROVER, a hardware-accelerated, multigroup ray casting package for simulated diagnostics and volume rendering.

Accomplishments in FY17:

- Siboka: Moved the VV4ALE3D application from prototype to alpha-release candidate, incorporating enhanced Structured Query Language (SQL) simulation metadata schema, VisIt post-processing, and a configurable dashboard for results. The SQL schema and tools were enhanced and evaluated on 50,000 Hydra simulations and made available to early adopters. Prototyped the SimManager tool, which parses and creates workflows as specified by a Yet Another Multicolumn Layout (YAML) specification. With SNL, ORNL, and LANL, released a Simulation Data management requirements document.
- C2C: Augmented the LANL format to meet LLNL requirements and developed tools to extract contour information from ARES and convert it to the LANL format, and vice versa. Implemented capabilities in PMESH for reading the LANL format.
- Nearly completed integration of a lightweight parameter study tool into UQ Pipeline.
- ROVER: Completed implementation of the multigroup raytracing functionality; VisIt integration will occur in summer 2017.

Planned Activities in FY18:

- Generalize the VV4ALE3D application and Siboka to target other use cases and evaluate usability. Work with the VisIt team to add more visualization to the simulation dashboard. Finish parameter study integration in UQ Pipeline. Participate in the tri-lab simulation data management effort.
- C2C: Collect and prioritize user requirements to guide development in one or more of the following areas: integration of conversion and parsing tools into more applications, tools for viewing and editing contour information, a central contour repository, and exploration of different back-end options.
- Take initial steps toward a problem setup data portal. Explore either the capabilities of Granta to store nonmaterial information, or data virtualization to bring data sources together behind a common Application Programming Interface (API) and access methodology.
- Continue ROVER multigroup ray-trace capability enhancements, and explore in-situ use cases and integration into the HADES simulated radiography code.

DevRAMP: Reproducibility, Analysis, Monitoring, & Productivity (LLNL)

The HPC Developer Ecosystem (HPCDE) team has expanded its scope to include productivity multipliers for all stages of the HPC software development pipeline. Its prior mission included common infrastructure for developer operations tasks such as

monitoring, databases, continuous integration, and software deployment. In addition to these, the HPCDE team is adding data analysis and tuning to its list of focus areas. This leverages data from the prior monitoring efforts, and it allows the team to focus on improving the efficiency of the entire development cycle: Develop, Test, Deploy, Monitor, Analyze, and Tune. The new DevRAMP team, like HPCDE before it, focuses on areas where the mission needs of simulation developers overlap with compute center needs, reducing duplication and creating a world-class development platform. It works closely with other ATDM teams.

Accomplishments in FY17:

- Developed a package-testing infrastructure for Spack, enabling build tests in Amazon Web Services (AWS).
- Shared an increased fraction of Spack maintenance effort with other laboratories such as Argonne National Laboratory (ANL).
- Migrated WCI's packages to the latest version of Spack; started a central repository for internal packages.
- Deployed the Bamboo CI in the Secure Computing Facility (SCF); automated agent deployment in the Collaboration Zone (CZ)/Restricted Zone (RZ) with portlet.
- Developed a custom authentication/authorization module for the Cassandra performance DB.
- Completed the first version of the ScrubJay data analytics tool for querying Livermore Computing (LC) facility data.
- Used Apollo, an autotuning extension of the RAJA portability framework, with Caliper and machine learning (ML) to speed up CleverLeaf more than 4 times and ARES by 15%.
- Employed network simulation tools to allow analysis of entire cluster workloads.
- Performed memory analysis to classify access patterns of application kernels using similar benchmarks.

Planned Activities in FY18:

- Implement new Spack features to support creation of “environments”; develop a new concretizer to enable faster and more robust dependency resolution; secure binary packaging.
- Develop a secure model to deploy JupyterHub and Jupyter notebooks with end-to-end transport layer security (TLS); develop a library for notebooks to talk to running simulations; contribute this work back to the Jupyter project so that it can be used broadly by other labs.
- Collaborate with ECP on multi-laboratory efforts to deploy Continuous Integration.
- Expand access to the Sonar cluster/Cassandra performance data base with new authentication.

- Develop and deploy the Jupyter interface for ScrubJay.
- Interface with ECP teams to integrate Caliper more broadly.
- Build out memory analysis tools to handle full applications.
- Develop new Apollo autotuning capabilities with AAPS; engage more code teams.
- Continue to work with LC on network and workload evaluation.

Math Libraries – MFEM (LLNL)

The LLNL ATDM Mathematical Libraries efforts include work on the MFEM library that is focused on providing high-performance finite element discretizations to high-order ATDM applications. One component of this effort is the preparation of the MFEM finite element library for exascale platforms by using mathematical algorithms/implementations that exploit increasing on-node concurrency. Enhancements in MFEM and the MFEM-based BLAST code, intended to provide efficient discretization components for ATDM-specific physics targeting LLNL’s MARBL application, are also a main objective for this project. MFEM is an open-source finite element library with nearly 3,000 downloads per year from over 70 countries. It is freely available at mfem.org and on GitHub.

Accomplishments in FY17:

- Released the mfem-3.3 software (at <http://mfem.org>), including many features relevant to ATDM applications and updated documentation.
- Performed general parallel conforming and nonconforming mesh refinement on unstructured meshes targeting parallel unstructured Adaptive Mesh Refinement (AMR). Optimized memory consumption of the nonconforming mesh data structures to about 30%. Work on a unique parallel anisotropic refinement feature is nearly complete and has been demonstrated in MFEM’s Shaper miniapp.
- Implemented many Arbitrary Lagrangian-Eulerian (ALE) discretization improvements for compressible flow in BLAST: localized bounds for monotonicity preservation during remap; improved element-based Flux-Corrected Transport capability; improved closure model computations; and also developed a new type of smoothing for the initial length scale h_0 . ALE discretization work was documented in a paper, *High-order multi-material ALE hydrodynamics* (submitted to SISC), and in presentations at international conferences.
- Actively engaged with ATDM application work: participating in weekly MARBL meetings and in discussions on BLAST reorganization for graphics processing units (GPUs) and advanced partial assembly algorithms; implemented low-order-refined (LOR) mesh constructors in MFEM for MARBL’s Level 2 (L2) milestone; worked on integration of Axom (a software toolkit for physics simulation code)/Sidre (datastore) in MFEM and BLAST.

Planned Activities in FY18:

- Improve parallel nonconforming AMR implementation for better scaling, efficiency, and handling of anisotropic adaptivity.
- Demonstrate the benefits of nonconforming AMR algorithms in MFEM examples and in the BLAST code. Document results in a paper or technical report.
- Continue to be actively engaged with GPUs and advanced partial assembly algorithms work in BLAST, coordinating with and leveraging work in MFEM and the ECP/Center for Efficient Exascale Discretizations (CEED) project.
- Continue the development of ALE discretization improvements for compressible flow in BLAST. This may include completing the capability to perform multigroup radiation diffusion simulations and documenting rad-hydro algorithms in a journal publication.
- Continue to be actively engaged with ATDM application work and provide support for MARBL L2 milestones.

Tri-lab Flang Subcontract (LLNL)

Flang is an open-source LLVM-native FORTRAN compiler being developed under subcontract by NVIDIA through tri-lab ASC support. FY18 represents year 3 of the planned 4-year project, and will build upon the open-source release in FY17 by expanding the set of language features provided, continuing to pursue performance optimizations. The subcontractor will work closely with DOE applications in the NNSA and across the DOE ECP teams to test advanced functionality, performance on emerging advanced architectures, and to prioritize work toward meeting both the NNSA mission needs and gaining broad community adoption. A primary goal of this year is to see Flang receive significant uptake in the HPC and LLVM communities so that the project may continue beyond FY19 with significantly reduced reliance on ASC funding.

This subcontract is managed by LLNL, with technical input and oversight provided by LANL and SNL.

Accomplishments in FY17:

- Achieved all subcontract milestones on time, and contract deliverables and payments were managed by LLNS in a timely manner.
- Engaged in bi-weekly planning calls between the tri-lab team and the subcontractor, regular installation and testing of releases, and feedback on prioritization.
- Subcontractor performed first release of Flang to the open source community via github under the Apache 2.0 license.
- Subcontractor implemented significant performance improvements to bring performance on-par with the GNU FORTRAN compiler.

Planned activities in FY18:

- LLNL/LLNS: Continue successful oversight of the subcontract procurement, including an update to the Phase 3 (FY18) SOW to better reflect current priorities for NNSA and ECP.
- Tri-lab team: Continue working closely with the subcontractor on testing, prioritization, and providing feedback that will help Flang take root in the open source community

Programming Model Standards and Architectures for Storage and I/O (SASI) (LLNL) – New in FY18

This project covers two key items critical for ASC's next generation computing: programming models standards, and storage and I/O. This is a new project in FY18, having been split out from the NGCE project.

The primary programming models standards work for this project is on MPI and OpenMP. For MPI, we focus on interfaces for supporting tools, including MPI_T and MPIR, and for fault tolerance, including Reinit. In OpenMP, our focus is on the tools interfaces OMPT and OMPD. However, we will participate and monitor developments in all parts of both standards in addition to our focus areas. Additionally, we will participate in community programming model efforts to ensure that they support ASC needs.

The I/O portion of this project targets the design and implementation of a next-generation software stack for storage and I/O and includes work on checkpointing, user-level file systems, and burst buffer management. This project coordinates next-generation activities both within the ASC program and externally. It includes interactions with the DOE's ASCR, vendors, and academia, such as planning and technical coordination for vendor and academic contracts and CORAL working group efforts.

Accomplishments in FY17 (under NGCE):

- Updated the SCR API and implementation to better suit several LLNL application efforts, including ARES and MERCURY.
- Continued work on file systems futures, including: development on Unify Framework for user-level file systems, including release of BurstFS burst buffer file system; leadership of a Dagstuhl seminar on User-Level Filesystems; and continued development on the I/O tracing framework DissectIO.
- Participated in standardization activities for both MPI and OpenMP in leading roles, as well as initiatives within these standardization efforts.
- Continued work on advanced and scalable debugging with a focus on reproducibility and race detection/avoidance for both MPI and OpenMP applications.
- Continued playing a leadership role in the MPI Forum and the OpenMP language committee.

- Developed and released a prototype of the OMPD debug API for OpenMP, which helped drive the standardization of this API that is critical to ASC.

Planned Activities in FY18:

- Actively participate in standardization efforts for MPI and OpenMP.
- Continue LLNL's file system futures effort investigating I/O technologies, software, and methods for the exascale timeframe, including work on Unify framework, user-level file systems, and DissectIO.
- Extend the SCR API to support more application needs, including general file output.
- Continue vendor interactions including participation in CORAL NRE Burst Buffer discussions and PathForward efforts.

Cognitive Simulation Project (LLNL)

The LLNL Cognitive Simulation Initiative aims to integrate modern Machine Learning techniques that have been used to tackle challenging problems across domains into ASC simulation efforts. Several areas have been identified and explored in this area, from in-situ analytics to improved physics models to the use of surrogate models to quickly explore design space. By incorporating these activities into a coordinated effort, we can ensure that we are leveraging the experiences of others and we can invest in frameworks, when and where they are applicable, to increase the effectiveness of longer-term investigations. Additionally, this effort allows us to build a workforce that combines knowledge of physics, parallel computing, and data science, which will strategically enable us to address future challenges and innovate in a data-rich, exascale environment. This effort will be closely coordinated with the Workflow Modernization Project.

Accomplishments in FY17:

- This is a new project.

Planned Activities in FY18:

- Provide initial accuracy results for the use of existing machine learning techniques to predict meeting high/low thresholds for mechanical strength of HE lots. (Q2)
- Provide initial accuracy results for the use of deep reinforcement learning for mesh relaxation. (Q3)
- Work with the WSC Workflow team and the design community to determine current and future data availability and steerability in a UQ environment. Use this information to create a 5-year plan for integration of Cognitive Simulation into UQ. (Q4)

Exascale Computing Facilities Modernization Enhanced Infrastructure Planning (LLNL)

This effort will expand the conceptual design report (CDR) completed in CD-1 for ECFM by creating enhanced planning alternatives and options to provide the most cost-effective solution to add 40 MW of power and 8,000 tons of cooling to B453. The project will develop a detailed cost/benefit analysis of efficient and sustainable distribution and transmission options.

Accomplishments in FY17:

- This is a new project.

Planned Activities in FY18:

- Provide detailed cost/benefit analysis of distribution and transmission options.
- Complete ECFM Enhanced Infrastructure Planning.

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 FY17:

- Conducted contract negotiations and issued awards.
- Provided technical coordination and contractual management for PathForward projects.

Planned Activities in FY18:

- Provide technical coordination and contractual management for PathForward projects.

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 visualization and analytics, advanced storage, architectural analysis, workflows, and resilience, and the systems interfaces for these. This work will be the conduit to other areas of the Exascale Computing Project, PathForward, 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. The project areas include most CSSE areas: system data and application monitoring, scalable storage, containers and virtualization for applications, in situ data analytics for large-scale simulations and resilience for application fault injection, and system fault domain characterization. This NNSA project encompasses the following ECP projects:

- 1.3.4.02: Data Analysis Initiative and HXHIM – application focused key-value store
- 1.3.5.01: Cinema – In Situ Visualization
- 1.3.7.01: RESL – Resilience characterization

Accomplishments in FY17:

- Deployed and tested MarFS, a scalable near-POSIX file system over Cloud Objects, in concert with the FOUS subprogram (campaign storage—an intermediate storage tier between the parallel file system and archive storage).
- Integrated Darshan (I/O logging tools) within LANL's PFTool suite and as a component within multiple LANL platforms.
- Completed design and initial implementation of LetGo: A Lightweight Continuous Framework for HPC Applications Under Failures.
- Completed initial investigation and mitigation of software fault injection into the FleCSALE RISTRA/CDA proxy application.

Planned Activities in FY18:

- Add arbitrary data types and operators to the Cinema database, such as support for 1D and 2D plots.
- Demonstrate and deliver accelerated ParaView, Cinema, and vtk-m workflow with an ECP exascale application.
- Arbitrary extract visualization, analysis, and dynamic query support for the Cinema database, including reduced-size geometric representations.

- Integrate new and existing functionality and demonstrate its application with the LANL CDA application.
- Extend HXHIM, a next-generation key-value storage middleware for storing simulation data, demonstrating storage performance across a variety of storage back-ends and across a variety of interconnection networks.
- Implement a scale-out monitoring infrastructure capable of meeting HPC requirements for data capacity, longevity, and ingest rate.
- Document existing workflows to identify opportunities to characterize the software packages in use by integrated codes and identify passive monitoring techniques for monitoring IC executing efficiency.
- Develop a methodology (and tool) for anomaly detection across the data center, considering multiple data sources.
- Investigate the relationship between facilities data and machine health.
- Extend the data analytics and resilience modeling for DOE systems (Q1 FY18 L2).

Co-Design and Programming Model Research (LANL)

This project contains the forward-looking research for advanced computing technologies at extreme scale. Co-design and programming model research are the basis of these investigations in support of ASC code needs on future hardware. This ASC project encompasses the following ECP projects:

- 1.3.1.02: Legion – Asynchronous data-centric task-based programming model
- 1.3.2.01: ATDM Tools – Low-Level Virtual Machine (LLVM) compiler and runtime tools and QUO thread-level heterogeneity in coupled MPI applications
- 1.3.3.12: FleCSI – Topology and execution framework and analysis for next-generation codes

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 is the Flexible Computational Science Infrastructure (FleCSI) project, whose objective is to provide a high-level programming system interface for the development of multiphysics production code applications. FleCSI provides a task-based distributed-memory and a fine-grained data-parallel execution model that is designed to allow portability between various system architectures. Additionally, FleCSI provides data and control models that are consistent with state-of-the-art runtimes, including Legion and asynchronous MPI.

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 FY17:

- Implemented “Control Replication” in Legion to support apparently sequential semantics in a fully distributed environment.
- Demonstrated multinode FleCSI on Legion with a Lax–Wendroff-based proxy application in collaboration with CDA.
- Completed prototyping and unit tests of MPI backend (Risk Mitigation) closely aligned with Legion requirements.
- Tracked and evaluated Flang (an open-source FORTRAN compiler) with ASC codes of interest.

Planned Activities in FY18:

- Implement FleCSI C++ support for optimized compilation (L2).
- Implement Legion task support for other runtimes (OpenMP, etc.) and Legion tools based on LLVM.
- Perform Flang vetting based on ASC applications.
- Implement control replication for FleCSI under Legion.
- Add dependent partitioning for FleCSI and implement FleCSI extensions, including support for Dynamic physics and I/O in support of CDA RISTRA.

Future-Generation Computing Technologies (LANL)

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

Develop advanced virtualization and container technology in support of future workflows with large multiphysics codes and complex input and output stages. This includes instrumented workflows, workflows integrated with in situ data processing, and abstraction of the underlying HPC systems and resources. These technologies will also be developed in support of continuous integration to allow simplified development and regression testing environments.

The Build and Execute Environment (BEE) project is currently working toward a runtime environment that allows standard containerized HPC applications to run on any HPC infrastructure as well as any cloud infrastructure. In addition, we will develop the required standards and procedures needed to interface containers to their facility-specific runtime execution systems (e.g., BEE, Singularity, etc.).



The Machine Learning (ML) component has two main thrusts: investigating in-line ML support for multiscale simulation and cognitive simulation. Specifically, the multiscale work will focus on an in-line ML for mesoscale and macroscale model inference for the dynamic response materials; the cognitive simulation will take an exploratory look at inference of simulation, adaptive simulation parameter steering, and related methods.

The performance prediction work is focusing on discrete event simulation. The work will suggest application parameters for optimizing HPC resource allocation and in vetting algorithmic choices for applications. Additionally, we will conduct general performance analysis of existing and emerging architectures.

This NNSA project encompasses these ECP projects:

1.3.6.09 BEE – Container and Virtualized Application Workflow Environments

1.3.8.XX (TBD) – LANL Machine Learning and Application Performance

Accomplishments in FY17:

- Released the BEE container launcher.
- Made BEE containers portable between HPC and Amazon systems with checkpoint restart.
- Demonstrated the BEE launcher using Vector Particle-in-Cell (VPIC) and Paraview.

Planned Activities in FY18:

- Achieve integration with LANL's Charlie Cloud secure container environment.
- Develop BEE containers for the CDA RISTRA project for development and continuous integration.
- Deliver containerized applications to integrated codes and/or ECP, with input (meshing/common model) and output (in situ/post-process).
- Perform initial investigation of ML for multiscale and cognitive simulation.
- Complete the LANL IC application model for the performance prediction toolkit (PPT) and finalize the hardware model.

Programming Models and Abstractions (SNL)

This project explores programming 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 Asynchronous Many Task (AMT) programming models (DARMA) with an associated data-warehouse capability. The techniques are tested and demonstrated

through mini-apps that help inform both hardware co-design efforts and integrated codes projects.

Accomplishments in FY17:

- Used Kokkos to achieve performance portability in many applications within the DOE Exascale Computing Project (ECP).
- Introduced a “lambda capture” feature into the ISO/C++17 language standard and NVIDIA CUDA language (this is important for readability and portability).
- Matured Kokkos’ on-node parallel task-directed acyclic graph (DAG) capability, which is performance-portable to GPU architectures – the first such capability.
- Highlighted interesting challenges for uncertainty quantification at scale using the Multi-Level Monte Carlo mini application and identified this as an important use case for next-generation programming model co-design with DARMA.
- Implemented a data layout in Kokkos Kernels that improves the memory use and vectorization in the Intel KNL architecture.
- Demonstrated the performance and productivity potential for particle-in-cell algorithms in preliminary studies using the DARMA abstraction layer.
- Demonstrated good strong scaling of the DARMA C++ abstraction layer for asynchronous many-task runtimes in a preliminary study.

Planned Activities in FY18:

- Research, develop, enhance, and optimize Kokkos back-ends for ATS-1, ATS-2, Arm, and other NGP platforms, with emphasis on runtime performance, compile-time performance, and maintainability.
- Provide high-quality (production) Kokkos support and consultation for ASC applications and libraries, including additional documentation.
- Integrate EMPIRE, DARMA/CHARM++, Kokkos, and Resource Manager.
- Integrate into Kokkos collaboratively developed SIMD data types and math operators for robust vectorization.
- Develop a DARMA application kernel for PIC and SPARC (within the full application).
- Implement MPI interoperability (supporting hand-offs between MPI and DARMA phases of work).
- Perform DARMA API and runtime work to satisfy application developer feature requests (subsetting, termination detection algorithms, and hierarchical collections).

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 FY17:

- Added support for reading structured and unstructured mesh data in the CGNS (Computational fluid dynamics General Notation System) file format, the de-facto standard for computational fluid dynamics (CFD) analysis data, to the IOSS (Input/Output SubSystem) library package.
- Completed the transition of the Trilinos finite element tools for basic element topologies to a new library implementation based on Kokkos data structures, resulting in a path forward for performance portability on next-generation architectures. Completed the integration into discretization tools and EMPIRE.

Planned Activities in FY18:

- Time Integration: Develop, integrate, and demonstrate the IMEX (Implicit/Explicit) method in SPARC and EMPIRE.
- Discretization: Transition the SPARC reacting Gas Finite Volume kernels to use Kokkos, the Panzer library package discontinuous Galerkin method, Finite Basis tools for PIC, and componentize EMPIRE.
- Gather requirements and refactor the abstractions for the Nonlinear-Integrate and demonstrate the LOCA (Library of Continuation Algorithms) library package continuation driver for SPARC.

Next-Generation Software Technologies (SNL)

The challenges posed by next-generation hardware are being addressed and mitigated largely by software technologies. These technologies include 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. Because these technologies are foundational, they need to be well-designed and tested, and robust to failure. Improving these aspects of the technologies will be a focus in FY18. Machine learning techniques provide an opportunity to make significant improvements in the interface between users and algorithms, creating the possibility for a next-generation workflow (a project focused on machine learning for next-generation workflows will be new for FY18).

Accomplishments in FY17:

- Improved thread affinity in Qthreads multithreading library, contributed to standardizing memory management in OpenMP, and ported the Kitten lightweight operating system to Cray XC systems with support for system call forwarding to enable better Linux compatibility and high-speed network access (by the OSR team).
- Devised a new communication engine called OpBox that allows users to express a series of communication operations as a state machine that reacts to specific network events (by the Data Warehouse team).
- Performed application profiling and performance analysis activities to optimize performance for the latest NGP platforms, by the Application Performance Team (APT).
- Demonstrated that burst buffer resources made of Non-Volatile Memory express (NVMe) devices can dramatically accelerate I/O performance.
- Developed a technique for automatically leveraging hierarchical parallelism on NVIDIA GPU architectures by the ATDM Analysis Components team, resulting in improved partial differential equation (PDE) assembly performance and thread scalability.

Planned Activities in FY18:

- Develop production ready line implicit solvers, multilevel algorithms, and preconditioning for SPARC.
- Provide resource management for Kokkos-DARMA integration, ATS-1/2 run time system optimization, container and lightweight operating system evaluation, and communication characterization of ATDM apps and libraries.
- Continue work with the OpenMP Architecture Review Board standards committee, MPI forum, and the new Universal Test Interface specification.
- Perform production hardening for Data Warehouse capabilities for ATS-1 (SPARC: NGP Performance).
- Demonstrate data check-pointing capability for applications.
- Adapt Viz technologies and ATDM I/O technologies to ATS-2 node architecture.
- Prototype SPARC visualization.
- Use machine learning for improved adaptivity of solvers for large-scale PDE codes.
- Use machine learning to support and improve mesh design, generation, and pre-processing.
- Use machine learning for optimization of resource management in system software.

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 many 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 FY17:

- Established automated testing of ATDM aero (SPARC) and Drekar codes against development versions of Trilinos on the Advanced Technology Test Bed machines.
- Adopted and greatly refined the usage of the JIRA issue tracking tool and the JIRA Portfolio tool to manage the planning and execution of all the SNL ATDM projects.

Planned Activities in FY18:

- Coordinate and prioritize higher-priority DevOps tasks in the Coordinated DevOps for ATDM process, involving other teams supporting DevOps.
- Continue to develop and demonstrate a scalable approach for the common build, test, and integration framework to bring together different ATDM projects and codes.
- Deploy the common framework for build, integration, and test system to more ATDM codes based on priority.
- Continue extensions of CMake, CTest, and CDash for ATDM priorities with Kitware contract.
- Deploy Intel Phi architecture available during this timeframe if architectural improvements or differences are significant.
- Deploy the IBM Power/NVIDIA GPGPU architecture available during this timeframe if architectural improvements or differences are significant.

- Deploy or build on the initial Arm development and test environment.
- Implement targeted improvements in hardware infrastructure as necessary.
- Purchase software licenses in support of the ATDM test and development environment.

Arm Vanguard Advanced Prototype System Deployment (SNL)

The scope of this project is to lead the design, acquisition, and plan for deployment of the Arm Vanguard Advanced Prototype System (APS). The project takes into consideration mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design, integration, and operation of the platform.

The goal of the Vanguard project is to field a large-scale prototype of an Arm-based system that is optimized for NW mission workloads including providing performance for large-scale applications in support of the NNSA program's most challenging problems. Successful execution of this project will result in a new architectural option for ATS procurements. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

Accomplishments in FY17:

- This is a new project in FY18.

Planned Activities in FY18:

- Conduct face-to-face vendor design presentations.
- Release the Vanguard project RFP to vendors.
- Conduct tri-lab review of vendor responses.
- Select a vendor and negotiate a contract.
- Deliver a prototype software development platform.
- Coordinate site preparation activities.

Arm System Software Development and Deployment (SNL, LANL, LLNL)

The scope of this project is to accelerate the maturity and successfully demonstrate a system software stack for the Arm Vanguard project that is optimized for HPC and NW mission workloads. This software stack effort encompasses aspects of the entire HPC software stack including areas of focus such as operating system (OS) kernels, runtime systems, libraries supporting interprocessor communications that implement one or more massively parallel programming models, libraries supporting I/O, system and resource management, and resilience. Another critical area of focus is the application software

development environment required for a usable system, including compilers, optimized libraries, debugging, and performance profiling tools.

The software environment requirements will be defined with the expectation that an integrated software environment will be a collaboration between the system vendor and the NNSA laboratories. The vendor will have responsibilities under contract for delivering core elements of the software stack necessary for a viable integrated system. The laboratories are also expected to contribute tools and capabilities to integrate into an overall tri-lab software environment for the prototype system. There is also the potential for a multi-way collaboration in the software development environment. Producing quality high-performance compilers and tools for Arm will require contributions from multiple interested parties and has already been identified as a key gap. As lead laboratory for the Arm prototype system effort, Sandia will organize a tri-lab team that will define the detailed software plan for the Arm prototype system.

Accomplishments in FY17:

- This is a new project in FY18.

Planned Activities in FY18:

- Complete a gap analysis for existing Arm software stack options.
- Complete a tri-lab Arm software stack design and implementation plan.
- Complete negotiations with vendor(s) for software stack deliverables and collaboration areas.
- Deploy an OpenHPC-based software stack on the Arm prototype software development platform.

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 with NIH (LLNL and LANL)

DOE is partnering with the National Cancer Institute (NCI) of the National Institutes of Health (NIH) 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 comprises three pilot projects that aim to improve understanding of cancer biology and its application to more effective therapies. Pilot 1 develops advanced machine learning methods for large-scale DOE clusters to predict tumor drug response. Pilot 2 produces a multiscale, multiresolution molecular dynamics simulation framework for CORAL-class machines to study the dynamics of mutated proteins in cancerous cells. Pilot 3 applies natural language processing and deep learning to a suite of large cancer databases to optimize cancer therapies.

LLNL

In Pilot 2, LLNL will develop an unsupervised machine learning ecosystem that analyzes simulation data, recognizes biologically relevant models, creates new hypotheses and abstracts, creates new systems for finer-scale simulations, and returns newly learned results to higher-scale simulations.

Accomplishments in FY17:

- Pilot 1 – Developed baseline machine learning methods for predicting drug response in NCI-60 cell lines. Implemented a new autoencoder approach to deep feature learning and began initial tests.
- Pilot 2 – Developed and implemented baseline cell membrane molecular dynamics models. Simulated membrane dynamics and validated versus experimental membrane data. Initiated design activities for new multiscale molecular dynamics methods.

- Pilot 3 – Prepared for Surveillance Epidemiology and End Results (SEER) data access and Institutional Review Board (IRB) reviews. Developed preliminary trajectory estimation tools using locally available data sets.

Planned Activities in FY18:

- Pilot 1 – Complete and scale-up deep feature learning methods. Apply to drug response prediction in large-scale preclinical data sets.
- Pilot 2 – Develop full-scale simulation of multiple RAS proteins interacting with a cell membrane. Obtain initial results from new multiscale molecular dynamics methods.
- Pilot 3 – Apply computational trajectory modeling tools to NCI SEER data sets.

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 FY17:

- Worked to build a repository of curated data sets from available databases (NCI-60, CCLE, GDC, MPACT, and others) to be used for all Pilot 1 tasks.
- Employed clustering algorithms to validate the consistency of the data sets originating from these diverse databases.
- Explored machine learning models for drug sensitivity predictions.
- Engaged in analyzing the data using traditional machine learning based on the AJCC cancer screening manual, and assessing accuracy. A manuscript is in progress.
- Worked to devise an architecture for extremely high-quality information extraction from a subset of pathology reports by identifying and tagging the reports for which the extraction is uncertain.
- Designed UQ as a layer on top of the implementation of ORNL-published multi-objective deep-learning classification code; this is still being finalized. (This is joint work funded by the ECP CANDLE project and Pilot 3).
- Produced a first implementation of Rademacher error to better quantify the generalization errors (this is joint work between Pilot 1, Pilot 3, and the ECP CANDLE project).

Planned Activities in FY18:

- Integrate pathway and other mechanistic models with machine learning models.
- Develop a Bayesian framework for uncertainty quantification and experiment design.

Accelerating Therapeutic Opportunities in Medicine (ATOM) partnership (LLNL)

DOE, National Cancer Institute, University of California, and Glaxo-Smith-Kline (GSK) have established a co-design partnership focused on new high-performance computing architectures and approaches for molecular design and predictive biology that will accelerate drug design processes. LLNL will lead overall research activities in a multi-party CRADA and develop the computational components of the partnership. The desired result is the design and demonstrated application of new integrated computational-experimental approaches to molecular design for drug discovery and development that will accelerate access to new cancer therapeutics. The CRADA will create a shared computing environment for predictive modeling of drug efficacy, safety, and pharmacokinetics and a complementary experimental approach that is based on rapid, miniaturized, human-relevant biological testing. New active learning approaches to integrated development workflows will be developed and quantitatively evaluated. The project will lead to improved methods for molecular scale simulation and new techniques for integrating machine learning, simulation, and high-performance computing that will have broad impact across a variety of scientific applications.

The concept of operations of the project is to bring a team of computational scientists, biologists and biomedical engineers, and drug development experts together in a common location (planned to be at the University of California, San Francisco campus) with access to state-of-the-art computational and laboratory resources (including reach-back to home organizations) to develop a common ecosystem of tools and methods to support development of integrated molecular design workflows. LLNL research efforts will focus in three technical areas, which are listed in the FY18 Planned Activities section below.

Accomplishments in FY17:

- Completed the ATOM R&D plan.
- Established the project team and set up the UCSF location.
- Performed initial data transfers and computational tool development.

Planned activities in FY18:

- Mechanistic prediction algorithms and models – Develop a collection of multiscale mechanistic models to predict the pharmacokinetics, efficacy, and safety liabilities (on/off target) of small molecule therapeutic candidates in healthy and diseased human systems. Define the data and computing hardware requirements to run such models.



- Data-driven prediction algorithms and models – Develop a collection of data-driven machine learning models using all available data sources to help predict the pharmacokinetics, efficacy, and safety liabilities (on/off target) of small molecule therapeutic candidates.
- Active learning methods and workflows – Develop active learning methods and workflows for the intended discovery and development processes.

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

The mission of this national subprogram is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of the NNSA. This subprogram 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. In addition, this subprogram will evaluate potential weapon applications of computing technologies that go beyond Moore's Law scaling and Von Neumann architectures.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2016, and through quarter 3, fiscal year 2017, are reflected below for the CSSE subprogram.

- Procured 83 Scalable Units of CTS-1 systems for the NNSA labs (funded by ASC and institutional computing programs), supported and evaluated technologies for CTS-1, and tracked technologies for CTS-2. (LLNL)
- Received, integrated, and accepted new ASC-funded CTS-1 systems: Opal, Jade, Quartz, Mica, and Topaz. (LLNL)
- Delivered Applications Development Environment for Sierra Early Access systems and conducted applications preparation for Sierra through COE and user training. (LLNL)
- Delivered storage, file systems, and data management advances including High Performance Storage System (HPSS) 7.5.1, Lustre 2.8, and new versions of Hopper. (LLNL)
- Developed a Beyond Moore's Law roadmap with tri-lab partners, supported the TrueNorth system, and evaluated algorithms and applications on the neuromorphic system. (LLNL)
- Ran two Advanced Technology Computing Campaigns (ATCC) on Sequoia and supported applications teams. (LLNL)

- Designed, implemented, and deployed the initial data systems monitoring architecture and software stack. (LANL)
- Received an SC16 Best Video award for the collaborative work in visualizing a tsunami generated by an asteroid impact. This work was in collaboration with IC under the visualization of a planetary defense simulation ensemble. (LANL)
- Received, integrated, and accepted new ASC-funded CTS-1 systems Snow, Fire, and Ice. (LANL)
- Trinity KNL nodes were transitioned and merged with the classified Trinity system (Haswell nodes). Trinity returned to ATCC use with both types of processors. (LANL)
- Completed the Critical Decision (CD)-1/3a documentation and obtained concurrence and released the Crossroads Request for Proposal (RFP) for vendor responses. (LANL)
- Completed and released the ATS-3/Crossroads RFP and completed source selection, with vendor negotiations commenced. (SNL)
- Deployed Cavium ThunderX2 Arm Testbed. (SNL)
- Received, integrated, and accepted new ASC-funded CTS-1 systems: Cayenne and Serrano. (SNL)
- Applied Trinity Advanced Power Management (APM) NRE technologies to characterize workloads running on Trinity Haswell and Knights Landing partitions in support of the L2 milestone. (SNL)
- Provided SNL's Portals communication technology, which is the basis for new interconnection networks being designed and built for next-generation systems by three hardware vendors. (SNL)
- Developed a new detailed model in Structural Simulation Toolkit (SST) for advanced nonvolatile memory technologies that shows great promise for future HPC platforms. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Sierra Applications Development Environment		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 8/31/18		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
<p>Description: To coincide with the siting and installation of the Sierra Initial Delivery system (one-quarter size), the Scalable Applications Preparation (SAP) effort within the Applications Development Environment and Performance Team (ADEPT) will deploy an application development environment that enables initial porting, tuning, and characterization of ASC application codes. We will extend the CORAL EA system development environment to support new features provided with the Power9 Central Processing Unit (CPU), NVIDIA Volta GPU hardware, and new system software. The team will evaluate MPI-direct-from-GPU messaging, unified memory malloc, and burst-buffer related features. Third-party tools and software for debugging and performance analysis will also be deployed and evaluated for expected capabilities.</p>		
<p>Completion Criteria: Deployment and documentation of the Sierra development environment for the system users.</p>		
Customer: ASC Sierra early users		
<p>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.</p>		
Supporting Resources: CSSE, Sierra system		

Milestone (ID#TBD): Local Failure Local Recovery (LFLR) Resiliency for Asynchronous Many Task (AMT) Programming and Execution Models		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 9/30/2018		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
<p>Description: Resilience is an imminent issue of next-generation platforms (NGP) because of reliability concerns based on the architectural complexity and the constraint in system power budget. Under such unreliable computing systems, introducing failure mitigations at the runtime and application layers is essential for executing ASC applications. Our study involves the development of a prototype resilient AMT system to enable multiple task replay, replication strategies, and other AMT-based failure mitigations. The prototype code will facilitate 1) evaluating the scalability, performance, and costs for different AMT resilience options, and 2) assessing the accuracy/cost tradeoffs of the application-specific failure detection and mitigation schemes in the AMT context. Our empirical study will be performed with emulated failures either on a large-scale HPC system or simulated architecture to represent NGP. The outcome is expected to provide definitive guidance to the design roadmap of the AMT programming framework (e.g., ATDM-DARMA) development as well as the resilience-driven co-design of NGP. Additionally, the milestone activity will provide initial capability for AMT resilience on ASC apps.</p>		
<p>Specific Deliverables:</p> <ol style="list-style-type: none"> 1. Prototype implementation of resilience schemes for the asynchronous many-task model, including task replication and replay. 2. An analysis of the accuracy-cost tradeoffs, scalability, performance and costs for multiple AMT resilience options. 3. A report to inform the code development roadmap guiding the Sandia/ASC strategy for AMT resilience for NGP. 		
<p>Completion Criteria: This milestone will be completed when the above-mentioned analyses are complete and a report is provided summarizing the key results and recommendations concerning AMT programming framework development for NGP.</p>		
<p>Customer: ASC application developers and the ATDM DARMA and Kokkos projects.</p>		
<p>Milestone Certification Method: A program review is conducted and its results are documented. A report is prepared as a record of milestone completion.</p>		
<p>Supporting Resources: Trinity, DARMA team, Structural Simulation Toolkit (SST)</p>		

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 FY17:

- Procured, stabilized, and moved into production 10 additional CTS scalable units.
- Extensively tested and resolved issues on production CTS systems.
- Assisted users in transitioning to CTS platforms.

Planned Activities in FY18:

- Lead and manage the CT systems contract and any FY18 CT system procurements.
- Perform technical evaluation of CT systems including selected cooling and networking configurations.

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 FY17:

- Received, integrated, and accepted new ASC-funded CTS-1 systems: Snow, Fire, and Ice.
- Transitioned ASC workloads from older systems to the new CTS-1 systems.

Planned Activities in FY18:

- Stabilize the CTS-1 systems to improve system availability.

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 FY17:

- Continued lease payments for two SNL systems, Cayenne and Serrano.
- Negotiated corrective hardware additions for a power problem identified in CTS-1 cabinets.
- Identified a slow liquid coolant leak issue in CTS-1 nodes. Worked with vendors to identify the root cause and schedule corrective actions.

Planned Activities in FY18:

- Complete final payments on lease agreement for Cayenne and Serrano.

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 multicore and multithreading 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 is ushered onto the machine.

Accomplishments in FY17:

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

Planned Activities in FY18:

- 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 FY17:

- Provided technical coordination and contractual management for Sierra NRE and contracts.
- Completed deployment of the Sierra Early Access Systems, which are helping applications to prepare for Sierra.
- Continued application preparations for the Sierra system through the Sierra COE.
- Completed Sierra's clustered file system technical checkpoint and finalized file system selection.
- Conducted the Sierra prototype system technical checkpoint and finalized Sierra system configuration.
- Participated in activities related to ANL's Aurora system, within the CORAL procurement.

Planned Activities in FY18:

- Provide technical coordination and contractual management for Sierra NRE and contracts.
- Continue application preparations for the Sierra system through the Sierra COE.
- Deploy and accept the initial Sierra Initial Delivery system.
- Begin science runs on the Sierra Initial Delivery system.
- Deploy the full Sierra system and begin its acceptance.
- Establish the ATCC process for FY19 Sierra full-system access.
- 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 FY18, this project will focus on the project management of



the ASC Trinity and Crossroads systems. The focus in this project also includes the execution of DOE Order 413.3b.

Accomplishments in FY17:

- Continued to provide program and project management for computing platforms, including requirements gathering and analysis.
- Developed infrastructure design in anticipation of CTS-2 and Crossroads systems.

Planned Activities in FY18:

- 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 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 FY17:

- Brought Trinity Haswells into full production and ran the initial ATCC-1.
- Delivered and installed the Xeon Phi, or KNL, phase of Trinity.
- Successfully completed the Trinity KNL Acceptance.
- Worked with tri-lab code teams to support improved code performance, analyst workflows, and proposal process and workload prioritization for simulations on future ATCCs.
- Deployed the Trinity KNL partition (TR2) on unclassified network.
- Conducted an Open Science period on TR2.

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



Planned Activities in FY18:

- Continue ATCC support.
- Complete the Trinity Production Readiness Level 2 milestone.
- Integrate Trinity Haswell (TR1) and TR2 partitions on the LANL classified network.
- Release the integrated Trinity system for general availability and ATCC-4.

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 when defining 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 HPC 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 will provide performance for large-scale applications in support of the NNSA's most challenging stockpile stewardship 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 the NNSA labs as a tri-lab resource.

Accomplishments in FY17:

- Completed the CD-1/3a documentation and obtained concurrence.
- Released the Crossroads RFP for vendor responses.
- Completed source selection.
- Initiated negotiations with selected vendor.

Planned Activities in FY18

- Complete negotiations and draft a SOW for the selected vendor.
- Conduct an Independent Cost Review of the negotiated system.

- Develop the CD-2/3b documentation and gain concurrence in early CY18.
- Complete the formal contracts for the Crossroads systems.
- Define and initiate NRE work scope with the Crossroads vendor.

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 the 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 FY17:

- Participated in the evaluation and selection of Pathforward contracts.
- Completed analysis of ECP RFI responses to understand potential 2021/2023/2025 architectures.

Planned Activities in FY18:

- Support technical project management of Pathforward contracts.
- Analyze potential system configurations for the ATS-3 Crossroads platform.

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. Currently, it is more important to explore a diverse set of architectural alternatives than to push large scale. Discussions will 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 FY17:

- Deployed latest Intel Xeon and Xeon Phi test beds supporting ACES procurements.
- Deployed ATS-2 prototype of IBM Power9 and NVIDIA Pascal.
- Deployed Cavium ThunderX2 Arm platform.

Planned Activities in FY18:

- Complete the RFP for the multi-Petaflop Arm system “Vanguard” procurement.
- Acquire the latest ATS-2 testbed system.
- Upgrade the Cavium system with the latest Arm.

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, and 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 Tri-Lab Operating System Stack (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 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 FY17:

- Released updates to TOSS 2 (versions 2.5-4 through 2.5-9, and 2.6-1 through 2.6-2) and TOSS 3 (version 3.0-1, 3.0-2, 3.1-1, and 3.1-2), which included security updates and bug fixes.
- Utilized the SONAR software architecture to continuously collect and analyze system data produced by the Lightweight Distributed Metric Service (LDMS) and performance measurements from WCI proxy applications.

- Increased ConMan scalability to provide console management and logging for CTS-1 systems.
- Migrated from the Moab job scheduler to a SLURM-only configuration on all LLNL Linux clusters.

Planned Activities in FY18:

- Provide ongoing TOSS software development and support.
- Develop/deploy TOSS 2.X for legacy systems (based on Red Hat Enterprise Linux (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.
- Update SLURM on TOSS 3 systems from v2.3.3 to v16.05 or newer.
- Add OpenSSL 1.1.0 support to the MUNGE authentication service.

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 runtime library work.

Accomplishments in FY17:

- Delivered the applications development environment for the Sierra EA compute systems, which completes the associated L2 milestone.
- Provided additional CTS-1 tools, compilers, and software library improvements and updates through the Tri-Lab Computing Environment (TCE).
- Supported users and tri-lab applications teams on the ATCC efforts on Sequoia.
- Developed and conducted workshops on the use of LC HPC systems and tools, including a “Getting Started on CORAL EA Systems” session.
- Completed and published an empirical study evaluating the sensitivity of HPC proxy applications to multirail networking to help inform the procurement of network technology at LLNL.

Planned Activities in FY18:

- Continue code development environment support on all LLNL ASC platforms.

- Deploy, test, evaluate, and provide user support for the initial CORAL Sierra system applications development environment.
- Conduct Sierra Applications Preparation (SAP) activities to assist tri-lab IC teams and their users in porting, tuning, and using the Sierra ATC system.
- Develop tool infrastructures to improve scalability and performance of applications and the code development environment.
- Conduct network investigations with micro-benchmarks and archive performance results to characterize Sierra dual-rail Infiniband and track historic CTS-1 trends.

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. It 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 aspects 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. Current activities include investigation of distributed systems, power, monitoring, and techniques for power-capping and scheduling of HPC systems.

Characterizing hardware reliability in HPC systems remains a challenge. 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 nonvolatile storage will be studied in support of future hierarchical storage systems (for example, burst buffer architectures).

Accomplishments in FY17:

- Designed, implemented, and deployed the initial monitoring architecture and software stack.

Planned Activities in FY18:

- Deploy the full production HPC system monitoring system.
- Conduct Los Alamos Neutron Science Center (LANSCE) Beam testing for resilience of modern HPC node hardware.

- Develop machine learning methodology (and tool) for extracting actionable knowledge from ticketing system.

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 with experimental software stacks.

Accomplishments in FY17:

- Integrated Early Access Sierra-type hardware.
- Expanded support for KNL compute nodes and investigated emerging Arm-based processors for computational nodes.
- Added support for Docker to the Darwin test bed using LANL's Charlie Cloud software.
- Added support for D-Wave and TrueNorth development to the Darwin cluster.

Planned Activities in FY18:

- Integrate prerelease Sierra systems (P9 with Volta).
- Integrate the Arm test bed for I/O and system software stack testing.

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 FY17:

- Applied Trinity APM NRE technologies to characterize workloads running on Trinity Haswell and Knights Landing partitions in support of L2 milestone.
- Oversaw Trinity APM NRE contracts with Cray and Adaptive Computing. Evaluated APM NRE Power API deliverables on Trinity test beds and worked with vendors to resolve issues. Prepared for deployment on Trinity.
- Set up an archiving method for Trinity TR2 Open Science system monitoring data, collecting three months of information available for post-analysis.
- Released several new versions of the Power API specification, incorporating feedback from Trinity APM NRE vendors and the HPC community.

Planned Activities in FY18:

- Demonstrate Trinity APM NRE Power Band and Ramp Rate Control capability on Trinity.
- Analyze system monitoring data collected during a three-month Trinity TR2 Open Science period to analyze power usage behavior of individual applications.
- Evaluate performance and stability of the OpenHPC software stack on an advanced architecture system and identify gaps.
- Evaluate container and virtualization solutions for supporting HPC + data analytics workflows.

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/algorithms 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 FY17:

- Released three versions of SST with version 7.0 being the latest.

- Developed new memory models for non-volatile random access memory (NVRAM), high bandwidth memory (HBM), and the latest DDR, which are validated with vendor hardware including ATS-1.
- Modeled Si photonics-based reconfigurable dragonfly interconnect.
- Collaborated with ECP Hardware Technology Design Space Evaluation effort.

Planned Activities in FY18:

- Deploy a validated GPU model.
- Deploy an initial Arm ThunderX2/3 model.
- Analyze performance on potential ATS-3 architectures.
- Analyze performance of potential ATS-5 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 FY17:

- Provided SNL's Portals communication technology, which is the basis for new interconnection networks being designed and built for next generation systems by three hardware vendors, and includes updates to Portals, MPI and OpenSHMEM software support, with Sandia OpenSHMEM forming the basis for vendor (Intel) delivered Symmetrical Hierarchical MEMORY (SHMEM) implementations
- Conceptualized and developed new techniques to support MPI internal message matching support for new and existing processor architectures (Xeon and Phi)
- Collaborated with LANL to improve Open MPI remote memory access (RMA) support for multi-threaded codes, demonstrating scalability on both Trinity TR1 and TR2.

Planned Activities in FY18:

- Advance Portals technology to meet future generation networking requirements, including guiding the next major generation of vendor interconnects. Release a new version of Portals encompassing these changes.
- Explore more advanced Portals related approaches to networking through exploration of network interface controller (NIC)-local active messages and the complexity that is needed versus what can be delivered, which is expected to inform our vendor partners of the five- to seven-year window to fully revised networking hardware.
- Continue work with MPI optimizations for emerging compute architectures, which will be performance and/or memory consumption centric. Push these optimizations into future OpenMPI releases and, where possible, provide MPICH with guidance on similar implementations.
- Continue development of OpenMPI and Sandia OpenSHMEM in cooperation with our vendor partners. Release new versions of OpenMPI and Sandia OpenSHMEM throughout FY18.

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 FY17:

- Performed an analysis of the entire lifetime of the Cielo leadership-class system and determined the most common failures have unknown origins. This characterization data is being used to inform vendors to perform proper root-cause analysis.
- Developed (in partnership with Rice University) a resilient extension to the Open Community Runtime (OCR) stack that can recover lost data and tasks at process/node failures locally, as part of our exploration of AMT resilience solutions.
- Made an extensive survey of resilience techniques for parallel asynchronous task runtimes. A paper documenting the work will be published in August 2017.
- Explored (in collaboration with Rutgers University) resilience extension of the tuple-space framework for HPC systems. Tested the framework with random read-write, random Directed Acyclic Graph (DAG), and in-situ workflows.
- Developed a transparent error-mitigation methodology based on a “resilient view” of solver mesh data for improved code maintainability.

- Adapted linear algebra checksums in a more efficient approach based on physical quantities for precise local error detection and recovery by rollback.

Planned Activities in FY18:

- Determine the influence of hardware aging effects on DRAM and SRAM failure rates using lifetime-of-Cielo failure data. (Q1)
- Evaluate the efficacy of failure prediction mechanisms to aid in mitigating failures using Cielo data. (Q4)
- Collaborate with LANL to ensure proper failure tracking is incorporated on the Trinity HPC platform.
- Analyze and characterize the first year of production of hardware failure data from the Trinity platform.
- Deliver the L2 milestone on the performance analysis of LFLR on AMT, to include: (a) evaluation of multiple resilience options for AMT for HPC systems, or simulated computing systems using SST; and (b) a prototype of a resilient AMT runtime. (Q4)
- Conduct performance tuning of Fenix in collaboration with ECP OpenMPI project.
- Collaborate with the DARMA project team to plan DARMA's resilience capability.
- Implement and perform a moderate-scale evaluation of conservation-based error detection in one or more prototype parallel unstructured solvers, to include an example modeled on MiniAero, MiniFE, or NALU.
- Support the LFLR L2 Milestone on AMT resilience options by contributing scenarios for invalid-state (e.g., silent-error) detection, and by guiding experiments to help inform its role in future LFLR.

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 FY17:

- 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 while simultaneously planning and developing a minor release.

⁵ See <http://www.hpss-collaboration.org/index.shtml>.

- Deployed HPSS 7.5.1 to first production HPSS environments in the world, leveraging DB2 replication to significantly reduce archive outage impact from a weekend to only 10 hours.
- Deployed HPSS tape mover clusters to modernize remaining HPSS archival subsystems and retired existing five-year-old tape movers without user interruptions.
- Evaluated, selected and procured HPSS Core Server platforms to refresh existing five-year-old systems that are catastrophically failing.
- 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 FY18:

- 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 and simultaneously develop, test, and release HPSS 7.5.2.
- Deploy HPSS Core Server platforms to refresh existing five-year-old systems.
- Evaluate storage ecosystem hardware and architect future archives using technology that will continue to cost-effectively satisfy archive capacity and bandwidth needs.
- 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 Lustre open source file system 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 FY17:

- Participated in and led aspects of the continuing community software development efforts of Lustre 2.8 and beyond.

- Augmented drive monitoring and management in support of Just a Bunch of Disk (JBOD) arrays under Lustre.
- Enhanced Zettabyte File System (ZFS)-based Lustre performance.
- Supported LLNL production Lustre deployments.
- Actively participated in CORAL NRE File System and Burst Buffer working groups in preparation for the Sierra machine.

Planned Activities in FY18:

- Support the development, testing, deployment, and production use of new Lustre versions in classified and unclassified environments.
- Collaborate with Intel Corporation on auto-splitting directory feature which will enable manageable and efficient use of DNE2 multiple metadata server functionality.
- Evaluate external-to-Lustre metadata indexing solutions including to facilitate purge policy enforcement and data provenance management.
- 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 FY17:

- Tested Nvidia Pascal with RedHat 7 on x86_64, and Power8. Nvidia collaboration and testing with early access CUDA toolkit 8.0.
- RedHat collaboration for Omni-Path integration into RHEL 7.3.
- Intel KNL hardware updated and tested with RedHat 7.
- Multi-path iscsi support for stateless nodes on commodity clusters.

Planned Activities in FY18:

- Cavium ThunderX2 testing with RHEL7/TOSS 3.
- AMD Firepro testing with TOSS 3 and ROCm (a platform for GPU-enabled HPC and ultrascale computing) support on x86_64 and ARM64.
- Conduct Intel Skylake and AMD Naples processor testing on TOSS 3 (if available).

- Testing Security Enhanced Linux (SELinux) on TOSS 3 in a production environment.
- Test CPU isolation on TOSS 3.

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 campaign storage (MarFS), 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 closely with application teams and HPC 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 ASC HPC users can make productive use of the systems. 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 FY17:

- Continued the hardening and testing of MarFS.
- ARTeam supported integrated code work with porting to the KNL partition of Trinity.
- Demonstrated running the EAP Trinity Demonstration Problem at its full scale using High-Performance Input/Output (HIO) to DataWarp with an N-M I/O pattern (number of processes–number of files, where $N > M$).
- Released a HDF5-HIO bridge and demonstrate running Lagrangian Applications Project (LAP) with this code on Trinity.
- Completed integration of HIO into VPIC to enable burst buffer utilization for VPIC science runs on Trinity KNL.

- LANL's Application Readiness Team has resolved multiple issues related to the Slurm migration of Trinity at LANL. This includes advanced task affinity settings and a mapping of features in the prior resource manager (ALPS) to the Slurm toolset.

Planned Activities in FY18:

- Transition Trinity full system to production use for ASC codes.
- Performance improvements in ASC codes through improved usage of MarFS, DataWarp, and/or KNL processors.
- HPSS Development activities including release testing and feature development.
- Develop Techniques for enhanced data management. This task includes developing and using tools, such as UnionFS, Starfish, and/or Robinhood, to provide an enclave wide view of user data.
- Embedded team support for IC projects EAP and LAP for performance and porting to Trinity and Sierra.
- Continue DataWarp readiness and performance testing.
- Continued development of HIO library to support Trinity and Sierra.
- Communications performance support activity.
- Provide support and assistance for IC codebases in achieving correct MPI (and other communications software) configurations and diagnosing communications related performance issues.

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 FY17:

- Installed test systems for evaluation of HPSS Version 7.5.2.
- Supported development and debugging activities of the HPSS Collaboration and feature determination for versions currently in development.

Planned Activities in FY18:

- Complete local integration testing and performance evaluation of HPSS Version 7.5.1
- Deploy Version 7.5.1 in production across all Sandia networks.

Scalable Data Services (SNL)

The Scalable Data Services project provides low technology readiness level (TRL) R&D to investigate and develop data management and analysis services with potential benefit to large-scale ASC applications on current and future ATS and CTS platforms. The scope of the data services project includes the exploration of multi-laboratory collaborations in software-defined storage, the integration and use data-analysis services by HPC applications, and the development of enabling capabilities for management of data for application workflows.

Accomplishments in FY17:

- Completed a technical report evaluating the viability of DOE-developed software-defined storage capabilities for ASC.
- Started working with an ATDM system software team and Cray to develop a general-purpose query service that uses the new Cray Minerva software. This software enables cloud-like services to be deployed concurrently with HPC/MPI jobs on Cray architectures through a lightweight virtualization layer.

Planned Activities in FY18:

- Design an analytics integration service. This work leverages an ATDM-funded system software effort to deploy container-based analytics on Cray HPC systems. The role of this project is to define and develop interfaces to expose these capabilities as a service for HPC applications and systems software.
- Explore enabling capabilities for scientific workflows. This activity will explore the use of the ATDM-funded Data Warehouse as a data-management layer for application workflows.

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 FY17:

- Supported 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.
- Maintained a stable data analysis and visualization environment across LC platforms.
- Provided operational support for all visualization theaters, including demonstrations for high-level visits.
- Completed initial phase of testing and development of our machine learning framework for predicting mesh tangling failures in ALE simulations that are prevalent in ASC simulation codes using ALE.
- Developed initial visual analytics using parallel coordinates plot for high-dimensional machine learning training data and an interactive tree visualizer for random forest model output.

Planned Activities in FY18:

- Maintain the data analysis and visualization environment across LC platforms and networks.
- Support simulation users with data analysis and visualization activities including consulting and the creation of images, movies, and software.
- Provide operational support for three visualization theaters.
- Pursue research activities in topological methods and data compression algorithms and tools.
- 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 FY17:

- Released new versions of Hopper which include a new feature for performing complex, data-driven selections using an intuitive UI; among other things, this allows the performance of smart deletes within large datasets. Additional new features include automatically generating debug and diagnostic reports, expanding directory viewing flexibility, and general usability enhancements.
- Extended the MyLC dashboard to support a secure, managed means for help desk personnel to perform routine system administration tasks directly. Deployed an application to assist with revalidation of HPC resources, including project areas in file systems.

Planned Activities in FY18:

- Release new versions of Hopper and Chopper which can restore previous sessions, prevent duplicate copy jobs, provide additional support for working with sites which have restrictive gateway proxy systems, and contain general usability enhancements.
- Deploy a web application that helps the LC operations team keep track of the immediate issues and their status, and which will facilitate communication between operational shifts. Investigate a Hotline diagnostic tool that will automatically perform a set of standard checks on behalf of a user when they call in with a problem. Deploy a thorough testing framework to verify data and software integrity. Continue enhancing the MyLC Dashboard with user and staff requested features.

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 focuses on integration into the code projects and ensures the production capability of the software components are well integrating into LANL's HPC

infrastructure. This is also the path for newly developed visualization and analysis research to make its way into a production environment. 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 software framework) post-processing analysis and visualization.

Accomplishments in FY17:

- The collaborative work in visualizing tsunami generated by asteroid impact resulted in a SC16 best video award. This work was in collaboration with IC under the visualization of a planetary defense simulation ensemble.
- Visathon, held at LANL resulting in performance enhancements in software rendering for Paraview, Ensight, and VisIt on Trinity for both the KNL and Haswell architectures.

Planned Activities in FY18:

- Support the integrations, support, and maintenance of in situ capabilities in the EAP and LAP codes.
- ASC end-users Production visualization and support.
- Applied and targeted Visualization and Data Analysis (VDA) research in support of ASC codes including sampling and analytics development and integration for EAP/LAP.

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 VTK, the in-situ data analysis library (Catalyst) for coupling directly with running simulation codes, and a tool for analysis of large ensembles of simulation runs

(Slycat). 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 ASC advanced technology and commodity technology platforms.

Accomplishments in FY17:

- Completed and integrated single sign-on capabilities for authentication into the Slycat ensemble analysis tool. This greatly simplifies the ability for a user to interactively analyze and visualize large, distributed ensembles of simulation results.
- Completed initial integration of Slycat, ParaView, and Catalyst into the Sandia Analysis Workbench (SAW). This enables the use of some of our analysis and visualization tools for application workflows using SAW.
- Released ParaView 5.4.0 with nearly a dozen new features for Sandia codes, including improved performance for navigating directories with more than a million files, support for CGNS meshes (needed by ATDM and ASC/IC), and improved capabilities for generating visualization animations.

Planned Activities in FY18:

- Expand the use/capabilities of in situ analysis, focusing on a redesign of our in situ interface to simplify the integration of analytics into our HPC codes. We will also integrate compression algorithms, developed through recent Laboratory Directed Research and Development (LDRD) investments, into our in situ tools.
- Continue integration and maintenance of analysis and visualization tools with the Sandia Analysis Workbench.
- Enhance and improve ensemble analysis and visualization tools on ASC platforms. The primary focus for FY18 is on developing and integrating a capability for analyzing ensembles of videos generated from simulation results.
- Continue ParaView, Catalyst, and Slycat releases with production support.

Projects for Beyond Moore's Law (WBS 1.2.3.5.6)

The Beyond Moore's Law (BML) 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. 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."

Beyond Moore's Law Computing (LLNL)

This 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 FY17:

- *Roadmap Development* – The three NNSA Labs developed a roadmap that was delivered to the NNSA ASC CSSE program.
- *Testbed Support* – The three NNSA Labs have demonstrated that they have access to the proper tools for developing and running on the TrueNorth system. Additionally, we have coordinated training and documentation through Confluence and are sharing Corelets (small programs written in an object-oriented language for neuromorphic computing) through BitBucket.
- *Algorithm and Application Evaluation* – The algorithm and application evaluation has focused on scalability, accuracy, and effectiveness in simulations. Neuromorphic systems are being evaluated for applications from image analysis to ALE controls to numerical solutions. Additionally, a testing framework for Quadratic Unconstrained Binary Optimization (QUBO) was developed to evaluate the distribution of solutions and the effects of scaling on accuracy.
- Developed master equation level simulations of open quantum system dynamics in quantum annealing and other quantum computing architectures.

- Established quantum computing hardware laboratory by developing and implementing prototype testbed platform for studying quantum emulation.
- Developed a design for system with all-to-all connectivity.
- Developed industry contacts in quantum computing and established initial informal collaborations.

Planned Activities in FY18:

Testbed Support

- Continue to support access and documentation for TrueNorth. (Q4)
- Provide access to IBM's next-generation simulator and documentation when it is available. (Q3)

Algorithm and Application Evaluation

- Develop and share implementations of QUBO to evaluate scaling and accuracy and provide a way to compare against D-Wave. (Q2)
- Provide an investigation of image analysis for ASC problems with a neuro-inspired framework. (Q3)
- Evaluate applications on next generation simulators from IBM or others. (Q4)

Nonconventional HPC

- Investigate nonconventional hardware technologies through testbed purchases and evaluations.

Quantum Computing

- Implement second generation quantum simulation hardware testbed platform.
- Computationally investigate efficient mapping of algorithms to second generation quantum simulation platform capabilities, including the NNSA application areas of quantum thermodynamics and quantum many-body dynamics.
- Explore connections between quantum computing and machine learning in application areas of interest to NNSA.
- Continue to develop quantum computing collaborations with industry and academia.

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 twofold. First, the project looks to create tools and programming models to simplify 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 FY17:

- Successful demonstration of a compressive sensing/sparse coding algorithm on the D-Wave 2X for inferring the radial density dependence of 3D objects from synthetic radiographs. A classical QUBO solver (GUROBI) was unable to converge to a solution after more than 9 hours of execution whereas the D-Wave obtained the correct solution in only a few seconds.
- D-Wave was used to generate sparse representations of thumbnail images for automatic labeling. As with the above radiographic example, the D-Wave can infer sparse representations in only a fraction of the time required by GUROBI.
- Implemented a Sparse Prediction Machine based on a Deep Sparse Generative Autoencoder (DSGA) applied to High-Definition ImageNet Video on TrueNorth and Trinity.
- Near complete design of the MD-Application Specific Integrated Circuit (ASIC)/MD-Field-Programmable Gate Array (FPGA) chip.

Planned Activities in FY18:

- Continued development of coupled quantum and traditional HPC computing.
- Use of TrueNorth and deep learning to inference of hydrodynamics simulation results from existing simulation data.
- Fabricate and test the MD-ASIC.
- Characterize an ASC proxy application over a wide range of precision, including full precision. And investigation in reduced precision computing.

Non-Conventional Computing Technologies (SNL)

This project explores the potential value of emerging device technologies, computer architecture concepts, and models of computing. It spans three technical thrusts: Non-conventional High-Performance Computing (NCHPC), Neural-Inspired Computing (NIC), and Quantum Information Processing (QIP). The latter two thrusts (NIC and QIP) entail highly nonconventional models of computing that have the potential to provide exceptional computational capability on particular classes of computations. These approaches to computing are unlikely to support efficient, general-purpose computing.

NCHPC 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. Unlike the Advanced Architecture projects (above), NCHPC entertains advantageous changes at level throughout the system

stack. Three driving objectives of NCHPC area: 1) enabling on-going improvement of performance of general-purpose computing – it remains to be determined how to best to define “performance” for future computing systems; 2) greatly increasing energy efficiency; and 3) maintaining manageable cost of ownership, which includes capital cost, operating (power) cost, and user cost; i.e., minimizing disruption to productive work. NCHPC is defined by additions to CMOS-based CPUs, not exclusion of them.

Accomplishments in FY17:

- Identified an approach to extend SST to be applicable to emerging architectural concepts.
- Established detailed design for 10-nm node, end-of-roadmap CMOS field effect transistor (FET) and characterized its properties.
- Demonstrated through simulation the potential of a special-purpose processor in memory (PIM) microarchitecture, “Superstrider,” with custom algorithms to relieve the “von Neumann memory bottleneck” for memory-bound applications.
- Conducted neural algorithms research that focused on transitioning into the ASC mission the low-TRL technologies in four key areas that are coming out of Sandia’s Hardware Acceleration of Adaptive Neural Algorithms (HAANA) Grand Challenge and other LDRD projects. This work resulted in several conference publications and identified specific efforts to pursue in FY18 to best impact HPC.
- Considered how to represent special quasi-random structure (SQS) generation on the D-Wave quantum annealer ISING platform at LANL, which has led to algorithmic improvements that also make the classical calculation more efficient.

Planned Activities in FY18:

- Begin implementing the plans for extending SST for simulating emerging architectures.
- Develop plan for a “universal charge device simulator” and begin implementation.
- Further develop new architectural concepts: SuperStrider to improve efficiency of memory-bound applications and adiabatic circuits for recycling signal energy.
- Pursue application of neural-inspired computing techniques to HPC for facilitating operation of supercomputers, including their multi-level memories, and for speeding up simulations of physical systems.
- Continue neural computing testbed investigations.
- Continue investigating the effectiveness of ISING for selected discrete optimization applications.

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

This subprogram provides two critical enablers for the effective use of ASC tri-lab computing resources: 1) physical facility and operational support for reliable, cross-lab production computing and storage environments, and 2) a suite of user services. 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 subprogram.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2016, and through quarter 3, fiscal year 2017, are reflected below for the FOUS subprogram.

- Integrated, deployed and supported three Sierra early delivery systems to users (RZ, CZ, SRD).
- Integrated, deployed and supported multiple CTS-1 Systems (RZ, CZ, and SRD) including access for the Alliances on the CZ. (LLNL)
- Completed Sierra facility preparation projects: Building 453 Power Modernization and Advanced Development Required for Cooling Loop. (LLNL)
- Replaced the Restricted Zone's CryptoCard authentication service with one based on RSA. (LLNL)
- Implemented and deployed building redundant network design within HPC facilities. (LLNL)
- Integrated Snow, Fire, Ice into production environment enabling ASC codes to transition to the new CTS-1 systems. (LANL)
- Replaced Global Security systems with new Hail and Lysander CTS-1 systems. (LANL)
- Achieved Critical Decision-0 and Critical Decision-1 for the Exascale Class Computer Cooling Equipment Project. (LANL)



- Deployed new monitoring infrastructure for compute systems into production enabling new analysis of compute operations. (LANL)
- Replaced scheduler with Slurm across CTS-1 and ATS-1 systems. (LANL)
- Transitioned Serrano and Cayenne CTS-1 clusters to production status on the Restricted and Classified networks respectively. (SNL)
- Supported joint SNL/LANL Level 2 Milestone #6017 Trinity System Production Readiness. (SNL)
- Completed L2 Milestone #6019 Application of Performance Analysis Tools on SNL ASC Codes. (SNL)
- Completed transition of DisCom network long distance provider from CenturyLink to L3. (SNL)
- Added physical isolation system in the National Security Computing Center (NSCC) (i.e., wire caging with locks) to provide additional security access controls. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): ECFM Design		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 9/30/18		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Exascale Computing Facility Modernization (ECFM) Design – This effort will expand the advanced infrastructure planning report created in the ECFM Advanced Infrastructure Planning project by providing detailed construction project drawings and specifications in order to execute the project and achieve CD-2 and proceed to CD-3.		
Completion Criteria: Detailed architectural and engineering construction package including Title II drawings and specifications with estimate that sets the project baseline for construction.		
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: FOUS staff		

Milestone (ID#TBD): Sierra Initial Delivery (ID) System Acceptance		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date: 6/15/18		
ASC WBS Subprogram: CSSE, FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The initial first quarter of the Sierra machine will be sited, installed, and partial acceptance will be performed.		
Completion Criteria: The LC will demonstrate compute, network, and file system functionality of the first quarter of the Sierra system, allowing partial system acceptance.		
Customer: ASC Sierra early 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		

Milestone (ID#TBD): Submit Exascale-Class Computer Cooling Equipment Critical Decision 2/3 Package		
Level: 2	Fiscal Year: FY18	DOE Area/Campaign: ASC
Completion Date 12/31/2017		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: Having completed design, this milestone will submit required documentation for readiness for construction approval of exascale-class computer cooling equipment project.		
Completion Criteria: Required CD-2/3 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, Exascale-Class Computer Cooling Equipment (ECCCE) Project staff		



Projects for the Collaborations Product (WBS 1.2.3.6.1)

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

Program Support (LLNL)

The Program Support project provides service to the ASC program. Program Support services include procurement and contracting, project management, and meeting support. These services are in support of both tri-lab and LLNL-only activities, including collaborations with academic, industrial, and other government agencies.

Accomplishments in FY17:

- Continued FY17 procurement support, contract management, and program planning.
- Supported the annual HPC Operations Review meeting with Office of Science.
- Supported biannual Predictive Science Panel (PSP) meetings.
- Managed the Predictive Science Academic Alliance Program (PSAAP) II program for the six Academic Alliance Centers.
- Supported SC17, FASTMath/ParaDiS, and SUstained Performance, Energy & Resilience (SUPER) Institutes.
- Supported Scientific Discovery through Advanced Computing (SCIDAC) projects HAXTON and NUCLEI (Nuclear Computational Low-Energy Initiative).

Planned Activities in FY18:

- Continue FY18 procurement support, contract management, and program planning.
- Support annual HPC Operations Review meeting with SC.
- Support biannual PSP meetings.
- Assist in project management of the six PSAAP II Centers.

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 FY17:

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

Planned Activities in FY18:

- Support PSAAP II collaboration efforts at LANL.
- Participate in PI and PSP meetings.
- Provide LANL support for HQ.
- Manage the SC18 booth for DOE.

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 FY17:

- Organized and hosted a Predictive Engineering Science Panel (PESP) meeting in October 2016.
- Managed Leidos contract to provide organizational support and program coordination for HQ and tri-lab activities.
- Organized and led the ASC Principal Investigators (PI) meeting at the Naval Postgraduate School in May 2017.
- Organized attendance, booth, and meeting logistics for the SC16 conference.
- Supported programmatic needs of the PSAAP II program and the DOE ECP.
- Published the ASC strategy, ASC 2016 Right-Size study, and the ASC Simulation Strategy.

Planned Activities in FY18:

- Organize and host an FY18 PESP review.
- Support the ASC PI meeting to be held in 2018.
- 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.

Biological Planetary Protection (SNL)

This computational co-design challenge is a collaboration with NASA-Jet Propulsion Laboratory (JPL) and 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 FY17:

- Allocated funding and finalized plan February 15, 2017.
- Held kickoff meeting at Sandia on March 15, 2017, and weekly telecons have ensued.
- Formulated the Phenomena Identification and Ranking Table (PIRT).
- Developed sample problem Verification Test Suite (VERTS).
- Designed Computation Multi-scale architecture.

- Evaluated existing SIERRA random field capabilities for simple (beam bending) problem.

Planned Activities in FY18:

- Design element enrichment infrastructure.
- Develop element enrichment via subscale finite element method (FEM).
- Develop subscale FEM problem dispatch to GPU.
- Implement FEM multiscale informed enrichment for fracture.

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 FY17:

- NNSA Production Office approval was received for the FARO Freestyle3D handheld scanner, which will be used to scan detailed 3D objects to include in virtual reality training modules.
- The Wet Vacuum System in the building 9212 Headhouse Fan Room was selected for a virtual reality training pilot. Twenty-seven 3D laser scans were made of the Wet Vacuum System. The laser scans are being converted into 3D poly meshed models.
- Free energies of solvation were calculated for a range of metal ions. These values correlated well with experimental observation. These types of calculations will be used to assess the solubility of impurities and choice of lithium salt to be used in material processing.
- Determined the efficacy of different levels of theory for predicting vibrational modes of molecules of interest to material processing. These calculations can assess the molecular environment of metal ions such as lithium.

Planned Activities in FY18:

- Implement a virtual reality (VR) training module on the Wet Vacuum System for operator certification.

- Specify, procure and install a virtual reality training station for the Jack Case Center (JCC) training room.
- Upgrade the VR laboratory equipment for mobile scanning and high-performance processing of data cloud images from laser scanners. Investigate alternative systems and software for delivering VR content.
- Utilize the SNL Sierra toolkit to develop a prototypical model of the escape of a gas under pressure as it is vented from one or more drilled holes in a small pipe, in collaboration with the Kansas City National Security Campus (KC-NSC).
- Create a prototypical model of the diffusion of hydrogen in graphite using SNL's molecular dynamics simulator, LAMMPS.
- Continue to support Y-12 representative participation in Nuclear Security Enterprise (NSE) activities such as the Nuclear Weapons Engineering Analysis Technical Interchange Meeting (NWEATIM), the ASC PI meeting, and other NSE-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 FY17:

- Integrated and deployed three Sierra early delivery systems (RZ, CZ, SRD) to users.
- Integrated and deployed CTS-1 Systems (RZ, CZ, and SRD) including access for the PSAAP II Alliances on the CZ.
- Provided ongoing 24 × 7 support for CTS-1, Tri-lab Linux Capacity Cluster (TLCC)-2 systems, Sequoia, Vulcan, Lustre, and NAS file systems.
- Retired the Sierra Dell Capacity Cluster, Juno, Muir, Ansel, rzZeus, and rzMerl.
- Deployed two new unclassified Lustre file systems in the unclassified in support of CTS-1 systems.
- Selected a replacement identity management tool.
- Replaced the Restricted Zone's CryptoCard authentication service with one based on RSA.

Planned Activities in FY18:

- Perform initial unclassified deployment of the Sierra system.
- Install and deploy a new classified Lustre system.
- Provide ongoing 24 × 7 support for Sierra early delivery, CTS-1, TLCC-2, BGQ, Lustre, and NAS file systems.

- Deploy new identity management system and infrastructure, and retire the existing IdM account management.
- Configure and deploy a host-based intrusion detection system on select infrastructure and production systems.

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 FY17:

- Provided ongoing support services for hotline operations, documentation, and training.
- Deployed new LC website that includes general information and user portal.
- Continued to develop in-house knowledge and expertise in accelerator technologies and their software libraries for providing support to the user community.
- Developed documentation and training materials on the use of the new CTS-1 clusters and Sierra Early Access (EA) machines.
- Assisted users in migrating from TLCC-2 systems to new CTS-1 machines.
- Migrated from the Front Range to the Service Now Incident Management system.

Planned Activities in FY18:

- Continue to provide ongoing support services for hotline operations, documentation, and training.
- Continue to develop in-house knowledge and expertise in accelerator technologies and their software libraries for providing support to the user community.

- Continue to assist users in migrating from TLCC-2 systems to CTS-1 systems. Assist users with porting their codes to the Sierra EA systems.
- Develop documentation and training materials on the use of the Sierra and Sierra EA systems.
- Perform process improvements in accounts management, including migrating to paperless processing of remaining forms and consolidation of accounts databases.
- Develop and deploy a utility for logging user environment and linked library data for each job run on LC clusters improving the ability to debug user jobs and monitor library usage.

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 FY17:

- Completed Sierra facility preparation projects: Building 453 Power Modernization and Advanced Development Required for Cooling Loop.
- Completed design of Sierra mechanical and electrical infrastructures.
- Began build of Sierra mechanical and electrical infrastructures.
- Implemented and deployed building redundant network design within HPC facilities.
- Obtained Critical Decision-0 approval for the Exascale Computing Facility Modernization (ECFM) Project.
- Deployed support networks for CTS-1 and Sierra early delivery HPC machines.
- Upgraded external network bandwidth to 40G within the HPC and evaluated and procured HPC internet facing firewalls.
- Continued to evaluate 100G and greater networking technologies and encryption for use in the HPC SANs and WAN.

Planned Activities in FY18:

- Work with NNSA to complete an Independent Project Review (IPR) and obtain CD-1 approval for ECFM.
- Complete ECFM Advanced Infrastructure Planning.

- Begin ECFM Design.
- Complete B-654 Air Handling Unit (AHU) Installation.
- Complete supporting Center network architecture, design, and hardware procurements for Sierra HPC Machines.
- Deploy 40G HPC Internet facing firewall cluster and evaluate and procure New 100G HPC edge router.
- Evaluate and enhance network monitoring and security compliance capabilities.
- Upgrade WAN encryption from 1G to 10G.

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 FY17:

- Supported HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers.
- Deployed new hardware infrastructure for monitoring compute systems to enable new HPC operations and analysis.
- Upgraded systems to latest versions of TOSS-3.
- Replaced scheduler with Slurm across CTS-1 and ATS-1 systems.
- Deployed MarFS for ATS-1 production usage.

Planned Activities in FY18:

- Support HPC systems by conducting ongoing daily system and storage administration, with continuous monitoring of production systems and infrastructure servers.
- Monitor MarFS usage and decide on sufficiency or need for upgrade.
- Refresh archival storage hardware for improved resiliency.
- Explore new methods of system and user security to include Continuous Security Monitoring.

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, Fire, Ice, Snow, and Moonlight, and data storage and retrieval systems such as HPSS. In addition to monitoring all components $24 \times 7 \times 365$, 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 FY17:

- Provided $24 \times 7 \times 365$ 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.
- Folded the Trinity Phase 2 system into the production operations, monitoring, and maintenance program.
- Supported installation of Phase 2 of the Trinity platform.

Planned Activities in FY18:

- Provide $24 \times 7 \times 365$ Tier-1 support, operations, and monitoring of compute and file system resources.
- Provide continuing hardware maintenance, assessment of system life spans, and guidance on hardware maintenance burdens for computing equipment.

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 (SCC) at LANL for

deploying these CT systems. The integration and deployment activities will focus on the following areas: system/OS, file systems, interconnect, network, regression testing, and monitoring.

Accomplishments in FY17:

- Supported integration, testing, and diagnostic activities for final acceptance of the Trinity system.
- Provided continuing integration and top-tier diagnostic and troubleshooting support for standing production compute platforms and file systems.

Planned Activities in FY18:

- 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 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 FY17:

- Provided ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Produced new user help website to document new systems and services (DataWarp, KNLS, and MarFS).
- Trinity roadshow for assisting ASC Tri-Lab users with productive use of ATS-1.

Planned Activities in FY18:

- Provide ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Produce new data storage policy to provide framework for how to use different tiers of storage.

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 FY17:

- Approved Alternate Selection and Cost Range, Critical Decision-1 (CD-1), for the Exascale-Class Computer Cooling Equipment Project.
- Provided networking and facilities support, and notification services for the ASC user community.

Planned Activities in FY18:

- Provide ongoing operations, maintenance, and configuration of electrical and mechanical systems for ASC computing platforms and facilities.
- Submit to HQ the Critical Decision-2 (CD-2) packet for the Exascale-Class Computer Cooling Equipment Project.
- Create and submit to HQ the Mission Need Statement for the Beyond Exascale Infrastructure Project.
- Begin planning for installation of the Crossroads system.

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, Application Regression Test Bed (ART) platforms including Mutrino (Trinity) and a small early delivery Arm-based platform in FY18.

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 FY17:

- Provided ongoing production operations and support for ASC systems at SNL (CTS-1 and TLCC-2).
- Transitioned Serrano and Cayenne CTS-1 clusters to production status on the Restricted and Classified networks respectively. Transitioned TLCC-2 systems (Pecos and Chama) off ASC operations support.
- Upgraded all Sandia cluster systems to TOSS 3 and converted to SLURM-only batch management.
- Supported joint SNL/LANL Level 2 Milestone #6017, Trinity System Production Readiness.

Planned Activities in FY18:

- Support new early delivery Arm system in preparation for acquisition and support of Arm-based prototype system.
- Provide ongoing production operations of Cayenne, Serrano, HPSS systems, high performance parallel file systems, WC Tool, Synchronized Account Request Automated Process (SARAPE), and the advanced architecture test beds.
- Complete transition of all file systems to fully redundant power and cooling configuration supported by centralized Uninterruptable Power Supply and auxiliary computing center diesel generator.

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 FY17:

- Supported evaluations of performance on Trinity system, both Phase I and Phase II, for Sandia codes. Provided user support for SNL and tri-lab ASC computing systems, primarily Trinity and Sequoia, but including Sierra early access systems starting in June 2017.
- Ramped up support for Trinity Phase 2 deployment and transition to SLURM-only scheduling.
- Continued development of new HPC Web portal and redeployed the ACES Web portal.
- Completed L2 Milestone #6019 Application of Performance Analysis Tools on SNL ASC Codes.

Planned Activities in FY18:

- Continue to 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 LLNL's Sierra platform.
- Complete transition to new service management tool.
- Continue improvement 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 FY17:

- Concluded transition of DisCom long distance provider from CenturyLink to L3, which was delayed approximately 5 months from the original schedule.
- Upgraded High Performance Computing core network devices, replacing 6-year-old Cisco switches with new Arista switches.
- Relocated the Systems Operations Center to provide additional data center floor space for HPC and Enterprise systems.

Planned Activities in FY18:

- Establish initial operational environment for early delivery Arm Testbed platform.
- Prepare newly constructed 725-East HPC facility for anticipated siting of the Arm-based prototype system in FY19.
- Relocate existing advanced architecture test beds from Room 230 Computing Annex to Room X-50 and renovate vacated space in Room 230 to current computing facility standards.

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, starting 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 systems. The system must meet CCE requirements for providing a common software environment on CT systems 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 more effective response to system conditions.

Accomplishments in FY17:

- Deployed the latest version of TOSS-3 on CTS-1 machines.
- Installed new Continuous Security Monitoring solution on secure CTS-1 machines.
- Released TOSS 2.6, TOSS 3.0, and TOSS 3.1. Provided concurrent support for both TOSS 2 and TOSS 3 with monthly updates that included new kernels, security updates, and bug fixes. Performed additional releases as necessary to address critical

bugs and security issues. Expanded the use of Pavilion in the TOSS 3 release QA process.

- Built TOSS 3 for the 64-bit PowerPC Little-Endian architecture in support of potential future non-x86 CTS platforms and to demonstrate the feasibility of maintaining TOSS as a multi-architecture OS.
- Continued tri-lab discussion of migrating CT systems to run SLURM-only, resulting in a decision to do so. Negotiated a tri-lab support contract with SchedMD for this effort.
- Released LDMS v3.4.0, including new samplers and enhancements to enable creation of new metrics and metric sets that use functional combinations of metrics. Deployed LDMS on current Advanced Architecture Testbed(s) in collaboration with application analysts to identify and collect metrics of interest for performance and resource utilization characteristics for applications across platforms as they evolve. Also released Baler 2.0.
- Collaborated on HPC Application Monitoring efforts, sharing current status and identifying further collaborative opportunities, paths forward, and gaps. Discussed topics including facilities, platform, and application monitoring, large data analytics, and monitoring infrastructure related needs and efforts.
- Developed a tri-lab approach to gaining insight on application performance variation, identified candidate applications at each lab, and agreed on performance figures of merit for each. Began planning experiments and data collection.

Planned Activities in FY18:

- Continued support of current TOSS software stack on CTS-1 machines.
- Provide ongoing TOSS software development and support; develop/deploy TOSS 2.X (based on RHEL 6) for legacy systems and TOSS 3.X (based on RHEL 7) for CTS-1 systems.
- Add support to TOSS 3 for new architectures and hardware that will be relevant to future CTS platforms.
- Continue SLURM support efforts through tri-lab collaboration. Support the migration to running SLURM-only on TOSS clusters, and resolve any issues that may arise.
- Utilize our monitoring, analysis, and visualization tools and agreed upon performance figures of merit on production tri-lab clusters, to correlate tri-lab application performance variation with monitored system characteristics such as network congestion, file system contention, and node level phenomena.
- 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, assessment of containerized environments, 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. These frameworks and others will add to a new thrust of continuous performance collection that will add to the monitoring capability. Debuggers include scalability work with debugger vendors to further scale and enhance their products, 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, Tuning and Analysis Utilities (TAU), MUST/Vampir correctness checking with MPI usage/performance optimization tool, and others.

Accomplishments in FY17:

- Progressed in keeping up with architecture evolution on the O|SS development contract; ported and tested O|SS to Power8 (non-GPU version), with new features covering OpenMP integration and memory analysis and installed on Trinity test systems; began porting to Power8+GPU testbeds. Accepted all the milestones in the contract including; I/O experiments working with Lustre (Sonexion) and Burst Buffer NVRAM on the Cray XC-40 with Haswell and KNL processors on the Trinity Platform; developed Milestones for the new development contract.
- Released debugger - STAT 3, which has improved thread-debugging support and adds some new data-centric debugging features.
- RogueWave delivered the final milestone in BIGCAT II, which included a TotalView version with Python support and a most robust replay engine that was hardened against Tri-Lab codes, and formulated plans for post-BIGCAT II that includes better support for evolving programming models such as Kokkos and RAJA.
- Implemented Arm contractor SOW with Arm for distributed debugging tool (DDT) support to debug MPI code – forge better support for KNL Threaded applications.
- Provided MVAPICH and Open MPI enhancements and bug fixes, released support for Open MPI, ensuring that both MPI implementations worked well for TOSS3

implementation on CTS-1 machines. Packaged MVAPICH2-2.2, Open MPI 2.0.0, and Intel MPI 5.3 and Intel MPI 2017 for the cross lab TCE software distribution. Collected data in the MPI Tuning Study for test suite on CTS-1 for default TOSS 3 Open MPI, began gathering data for TOSS 3 MVAPICH.

- Completed OpenSource SW Support / training (Vampir, MUST, Score-P, Scalasca, O|SS) and contracts (O|SS, MUST, OpenMP). Incorporated MUST into the tri-lab release of the TCE software stack.
- Began effort to establish a common use programming environment for CTS systems through Tri-Lab Computing Environment (TCE), including a cross-lab build farm for a standardized software deployment process.
- Developed compiler benchmarking base suite consisting of an OpenMP benchmark and a modest collection of vectorization tests, currently on a LANL Git-environment with plans to open up for tri-lab submissions.

Planned Activities in FY18:

- Continue development and support efforts for debuggers, performance analysis tools, and MPI as programming models and architectures evolve, including Open source development and support contracts, and new additions including the Pruner tools set.
- Integrate with the Monitoring project to utilize programming environment frameworks and tools to enable continuous metric collection from applications.
- Continue focus on updating functionality of the tool ecosystem to support upcoming ATS architecture needs, including continued testing and integration for accelerator support, porting to OpenPower, Nvidia GPUs, and Intel's KNL, and broader collaboration on efforts such as the compiler benchmarking suite.
- Continue integration of the tri-lab tools build environment and broader integration through cross-lab deployment of the TCE environment, support of Spack HPC Software package manager, and assessment of containerized environment deployment.

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 SNL 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 (WC Tool) 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 FY17:

- Began testing the SARAPE application in a virtual server environment, which will provide greater application stability, increased performance, and better administrator management.
- Developed additional automation of database updates, which will increase request processing efficiency and accuracy for Sandia processing agents, as the number of SARAPE requests is increasing due to more numerous collaborations.
- Delivered the Metrics capability to processing agents at each SARAPE site, allowing custom metrics reports to be generated for each site.
- Began initial deployment and testing of the MySARAPE interface, which will allow users to view their SARAPE requests for Sandia resources for the previous two years.
- Added several new SARAPE resources, including early-access Sierra hardware platforms at LLNL.
- Integrated Homeland Security Presidential Directive 12 (HSPD-12) authentication as a login option, providing users at all NW sites simplified access to SARAPE.
- Deployed multiple versions of WC Tool for SNL and completed work with LLNL to host WC Tool for them at SNL.
- Delivered quarterly tri-lab workload characterization usage reports to NNSA (WC Tool).

Planned Activities in FY18:

- Continue to operate and improve the tri-lab SARAPE process for all remote access account requests and implement evolutionary improvements.
- Implement domain-restricted access controls within SARAPE rather than at a network level, to increase flexibility in managing an expanded customer base.
- Improve the MySARAPE (Phase 2) user experience through redesign, simplification, and expanded information; for example, provide discretionary actions such as user ability to clone a previous request to a renewal.
- Implement Trusted Identity, a new Sandia cyber-visitor role with streamlined approval using the DOE OneID database; gather feedback from users and processing agents and develop improvements.
- 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; address WC Tool bugs and update software to current versions and security updates.

Projects for the Special Purpose Facilities, Systems, Operations, and Support Product (WBS 1.2.3.6.4)

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 FY17:

- Provided HPC procurement, system administration, and operational support.
- Finished Facility modifications to support CTS-1 systems and infrastructure.
- Deployed CTS-1 machine into new computer room.
- Deployed an older existing system into the new computer room.

Planned Activities in FY18:

- Provide HPC procurement, system administration, and operational support.
- Decommission legacy HPC system and infrastructure.

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 FY17:

- Integrated and accepted new CTS-1 Hail and Lysander systems.
- Transitioned applications from older systems to new CTS-1 Global Security systems.

Planned Activities in FY18:

- Continued stabilization and maximizing availability for new CTS-1 Global Security systems.
- Retirement and decommissioning of former Global Security production systems.

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 FY17:

- Provided operations and support for Tri-Lab use of NSCC systems.
- Upgraded COI infrastructure to provide enhanced single sign on security. Added physical isolation system (i.e., wire caging with locks) to provide additional security access controls per customer requirements.

Planned Activities in FY18:

- Continue operations and support for Tri-Lab use of NSCC systems.
- Plan for acquisitions and retirements of platforms based on mission need.

Appendix H: Academic Alliance Centers

Accomplishments listed below represent those of FY17 for the PSAAP II Alliance Centers.

University of Utah

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

Accomplishments in FY17:

- Reworked the I/O subsystem in the Uintah runtime to include Portable Information Data eXchange (PIDX).
- Further advanced the Uintah runtime to be more deeply Kokkos and C++11 compliant, and demonstrated scalability with a 350 million core-hour simulation of the 1000-MWe boiler on MIRA and TITAN through an INCITE award.
- Performed top-down V/UQ to define error budget for each brick in the center hierarchy.
- Developed and validated new char oxidation, mineral matter transformation, and radiation emissivity models to be computationally efficient and within our target uncertainty budget.
- Made predictions for a) optimizing an existing 1000-MWe coal-fired boiler that was newly commissioned and made operational by an industrial partner, and b) reducing the footprint of a new-generation pressurized, oxygen-coal fired gasifier for pre-combustion carbon capture.

Planned Activities in FY18:

- Advance the PIDX I/O subsystem integration in the Uintah runtime including checkpoint/restart.
- Perform next iteration (year 4) top-down V/UQ with nonlinear response surface to a) quantify error budget for each physics component in our application, and b) identify model-form in need of improvement.
- Perform V/UQ analysis of each integrated systems experiment to a) identify adequacy of experimental data for quantifying uncertainty in overarching problem, and b) extrapolate uncertainty to overarching prediction.
- Together with our industrial partner, use Uintah/Arches with V/UQ to minimize gas-side thermal imbalance for a 1000-MWe coal-fired boiler.

- Hold a short course for advanced graduate students and professionals (from universities, national laboratories, and industry) interested in using the center's hybrid Bound-to-Bounds Bayesian approach to hierarchical top-down and bottom-up V/UQ for engineering applications.

University of Illinois, Urbana-Champaign

Center for Exascale Simulation of Plasma-Coupled Combustion

Accomplishments in FY17:

- Used the noninvasive annotation framework ICE (Illinois Coding Environment) to run a full “Golden Copy” PlasComCM prediction application at scale on multiple NNSA and other systems (including all the CTS-1 system Quartz), optionally invoking Moya, our Just-in-Time recompilation tool.
- Developed PlasCom2, a repackaging of the computational kernels to facilitate development as a standard-language “Golden Copy” application and expose opportunities for optional computer science tool application.
- Predicted with quantified uncertainty the plasma-coupled combustion ignition threshold of a fluid jet. This entailed over 300 runs of the full-scale ignition configuration.
- Developed a coupled experiment-prediction plan for applying physics models to high-speed engineering combustion applications.
- Made multiple visits to NNSA labs, including additional PSAAP internships, several of which have generated new and ongoing collaborations.

Planned Activities in FY18:

- Use ICE to launch (optionally) additional performance-enhancing tools while allowing further code development via the Golden Copy approach with only standard languages and constructs.
- Predict with quantified uncertainty Transient Ignition Kernel (TIK) characteristics in a non-premixed Mach 2.5 combustion flow.
- Make prediction runs with PlasCom2, invoking CS tools for performance and HW flexibility.
- Advance experimental diagnostics to quantify early times of our laser-induced breakdown (LIB) ignition seed.
- Continue to foster close interactions with NNSA labs through additional student internships, service by lab personnel on Ph.D. dissertation committees, collaborative publications, and visits by Illinois staff and faculty.

Stanford University

Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment

Accomplishments in FY17:

- Demonstrated Soleil-X scalability on O(10,000) CPUs and O(500) GPUs using Stanford exascale stack: i.e., the Legion runtime, the Listz domain-specific language, and all the required middleware.
- Implemented particle-particle collision models and developed enhanced point-particle treatments for two-coupling with fluid momentum and heat transfer.
- Performed 1000-run ensemble UQ of fully-integrated simulations, which involve high-fidelity radiation models (DOM), particle collision, and direct representation of fluid turbulence.
- Completed first heat-on experiment and carried out detailed comparisons with simulations, which include uncertainties due to 14 sources.
- Compared low- and high-fidelity radiation models in Soleil-MPI, and validated the DOM transport reference experimental data involving turbulent multiphase flows.

Planned Activities in FY18:

- Implementation and verification of high-fidelity radiation models (DOM and MCRT) in the Soleil-X software framework.
- Implementation and verification of particle-particle collision models in the Soleil-X software framework.
- Development, implementation, and verification of point particle correction models for momentum/heat transfer and particle-wall collisions using results and analysis of particle-resolved simulations.
- Validation of novel dynamic (i.e., parameter-free) subgrid scale models for wall-bounded particle-laden turbulent flows.
- Complete experiments for high-mass loading and carry out detailed comparisons with corresponding simulations including UQ.

University of Florida

Center for Compressible Multiphase Turbulence

Accomplishments in FY17:

- Improved our fundamental understanding of shock-induced particulate turbulent flows through microscale and mesoscale simulations. Developed novel gas-particle coupling models (such as PIEP – Pairwise Interaction Extended Point Particle Model) for two-way coupled compressible multiphase flow simulations.

- Developed CMT-nek as a higher-order shock-capturing discontinuous Galerkin spectral element methodology with fully implemented Lagrangian point-particle capability. Also developed novel load balancing algorithms for efficient computation of rapidly evolving Euler-Lagrange simulations.
- Enhanced the Behavioral Emulation (BE) methodology to scale beyond the device level, performing experiments on, and predictive BE simulations for, millions of cores.
- Fully developed Uncertainty Budget (UB) framework for model improved following rigorous uncertainty reduction, and its demonstration at the microscale and mesoscale.
- Close interaction with NNSA researchers through student internship, validation experiments, numerical/physical models, held joint workshop on multiphase flows, and extensive use of computational tools developed at the labs.

Planned Activities in FY18:

- Incorporate the PIEP model into CMT-nek, as well as complete ongoing work on CMT-nek development into a compressible Navier-Stokes solver using spectral elements, including validation and verification efforts.
- Extend the ongoing work on load balancing to pareto-optimization of performance and energy, while satisfying power and thermal constraints. Translate these co-designed developments from the proxy app to CMT-nek.
- Conduct a large suite of microscale simulations of shock-contact interaction with a bed of particles for varying parametric and geometric conditions, especially at conditions relevant to explosive dispersal of particles.
- Fully demonstrate the UB framework in the Sandia shock tube and Arizona State Expansion Fan problems. Propagate these mesoscale uncertainties to perform UQ of the demonstration problem.
- Fully extend the UB framework to guide SST-based exascale simulations and FPGA-based exascale emulations for predicting CMT-nek performance on emerging and future exascale architectures.

Texas A&M University

Center for Exascale Radiation Transport

Accomplishments in FY17:

- After many difficulties related to cross section data, neutron room return, and conflicting AmBe neutron source spectra in the literature, we calibrated and validated our fourth impurity model (IM1-D).
- We completed experiments through Level 1 (of the hierarchical experimental/simulation problem roadmap).

- We significantly improved our load-balancing algorithm for extruded triangular meshes.
- We have developed a new spatial multigrid method for the standard form of the Sn equations that is showing great promise.
- We made a partial public release of the STAPL (Standard Template Adaptive Parallel Library) framework for developing parallel programs in C++.

Planned Activities in FY18:

- Complete experiments, simulations, and UQ analysis for Levels 2–4 of the problem roadmap.
- Reduce size of Year 5 experiment to accommodate a box enclosure.
- Implement parallel uncollided and last-collided flux options.
- Identify parameters for each Year 5 Quantity of Interest (QoI) via PDT adjoint calculations.
- Make a full public release of STAPL software.

University of Notre Dame

Center for Shock Wave Processing of Advanced Reactive Materials

Accomplishments in FY17:

- Advanced spectral wavelet code and demonstrated error control, advanced Eulerian-Lagrangian solver, and developed PASTA-DDM-UQ solver.
- Performed impact simulations using Google-based volumetric billboards in Ni/Al systems and developed chemo-thermo-mechanical mini-app, iSIM.
- Released HPX-5 v4.0; further developed TTL (Tensor Template Library), and developed photon based MTL-HPX (Matrix Template Library).
- Demonstrated cubic boron nitride (c-BN) synthesis of TiN+B material system.
- Hardened the code base and improved the production performance by more than 2x, and advanced verification coverage and improved software engineering practices.

Planned Activities in FY18:

- Develop 3D wavelet code and HPX+ runtime system software specification for parallelization.
- Develop multiscale, multigrid solver using Google-based volumetric billboards, and advance chemo-thermo-mechanical coupled solver.
- Improve HPX+ scalability and performance, improve MTL/TTL-HPX libraries, replace MPI calls in physics codes and multiscale communications to HPX-photon, and develop HPX-light CPU/GPU scheduling for PGFem3D.

- Validate advanced QOIs (shock speed, final microstructure, temperature, etc.) for confined impact test of Ni/Al powders, and calibrate chemo-thermo-mechanical poro-visco-plastic material properties and determine parametric uncertainties.
- Perform full system demonstration simulations, validate them against thermo-mechanical experiments for Ni/Al powders, and perform advanced sensitivity and UQ studies.

Appendix I: Construction and Capital Equipment

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

Site	Title/Description	Classification	Last CD Completed	\$ in Thousands								Contact
				TEC	Prior Years	FY18	FY19	FY20	FY21	FY22	Future	
LLNL	B-453 Power Modernization	Construction-GPP	N/A	4,875		2,000						Anna Maria Bailey 925-423-1288 bailey31@lbl.gov
LLNL	B-453 Advanced Development Cooling Loop	Construction-GPP	N/A	1,500		100						Anna Maria Bailey 925-423-1288 bailey31@lbl.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@lbl.gov
LLNL	B-453 Sierra Site Prep	Construction-GPP	N/A	9,500		7,500	2,000					Anna Maria Bailey 925-423-1288 bailey31@lbl.gov
LLNL	ECFM TEC – CD-0 Project Data Sheet Information	Line Item	CD-0	95,000		3,000	40,000	44,000	8,000			Anna Maria Bailey 925-423-1288 bailey31@lbl.gov
LLNL	Commodity Technology System (CTS-1)	Capital Equipment, MIE		40,000	25,550	10,000	4,450					Matt Leininger 925-422-4110 leininger4@lbl.gov
LLNL	Advanced Technology System (ATS-2)	Capital Equipment, MIE		170,000	96,000	53,200	11,000	6,800	3,000			Bronis de Supinski 925-422-1062 desupinski1@lbl.gov

Site	Title/Description	Classification	Last CD Completed	\$ in Thousands								Contact
				TEC	Prior Years	FY18	FY19	FY20	FY21	FY22	Future	
LANL	Trinity: Acquisition of Trinity (ATS-1) system	Capital Equipment, MIE	CD-2/3b	182,000	171,600	5,400	5,000	0	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Crossroads: Acquisition of Crossroads (ATS-3) system (\$170M) and water cooling and electrical power installation (\$12M)	Capital Equipment, MIE	CD-1/3a 6/29/16	182,000	9.0	134,000	61,250	60,200	21,600	8,000	8,550	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, Line	CD-1 5/31/17	66.989		2,338	22,000	42,651	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	ATS-5 system (\$500M) and water cooling and electrical power installation (\$50M)	Capital Equipment, MIE		550,000	0	0	0	1,000	1,000	20,000	488,00 0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Future Supercomputing Infrastructure: 40MW of additional power and cooling	Capital Equipment, Line		140,000	0	0	1,000	3,000	15,000	40,000	81,000	Jason Hick 505-667-4477 jhick@lanl.gov