Advanced Simulation and Computing

FY19 IMPLEMENTATION PLAN

Version 0

August 8, 2018

NA-ASC-127R-18-Rev.0-IP

LLNL-TR-756207

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	*Michel McCoy, Bill Archer, Scott Collis*	*Initial*	*Mark Anderson, Thuc Hoang*	*8/6/2018*	*Initial release*

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

Contents

I. Overview

The Stockpile Stewardship Program (SSP) is an integrated technical program for maintaining the safety, surety, and reliability of the U.S. nuclear stockpile. The SSP 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.

The 2018 Nuclear Posture Review (NPR) report reaffirmed 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: “Maintain and enhance the computational, experimental, and testing capabilities needed to annually assess nuclear weapons.”[1] To execute this strategy necessitates the continued emphasis on 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.

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

[1] *2018 Nuclear Posture Review Report*, February 2018, p. 11.

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 FY19. The Initial IP, effective August 8, 2018, should be consistent with the Department's Base Table when operating under a Continuing Resolution (CR). The final IP, effective date TBD, is consistent with the final, enacted appropriation.

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	FY2019 Target	Endpoint Target
Advanced Simulation and Computing Program	Reduced Reliance on Calibration	71%	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 FY19 Stockpile Stewardship and Management Plan[2]. ASC also provides critical capabilities informing efforts to extend the life of the nuclear stockpile.

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

The ASC Program is 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 models used in simulation codes. PEM also includes activities that are directly aligned

[2] U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2019 Stockpile Stewardship and Management Plan—Biennial Plan Summary*, Report to Congress, <Date>.

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 FY19.

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 Department of Energy (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 FY19 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.

- Deliver improved code capabilities and physics models supporting FY19 Predictive Capability Framework Peg Post.
- Incorporate additional weapons systems into the V&V UQ methodology for the Annual Assessment Review (AAR).
- Combine different types of NTS diagnostic data to provide higher-quality data interpretation for code validation.
- Develop validation methodologies for circuit-level modeling of mixed technologies in combined radiation environments.
- Adapt existing codes to new architectures, and migrate current design and safety codes to run efficiently on hybrid computer architectures.
- With system acceptance in Q4 FY18, complete transition of the ASC Sierra platform at Lawrence Livermore National Laboratory (LLNL) to the classified computing environment 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 construction and final design stages, respectively. These two facilities will support deployment of advanced architecture systems on the path to exascale and at exascale.
- Sustain the ATDM investment in 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 DOE/Advanced Scientific Computing Research (ASCR) office on the DOE 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 decrease in FY19 forces the program to decrease its investment in its collaboration with the National Cancer Institute (NCI), and vendor R&D partnerships and next-generation software technologies in the DOE Exascale Computing Project.

The drivers of the ASC Program that require the FY19 budget include the following:

- The Nuclear Weapons Council approved the stockpile strategy articulated in the 2018 Nuclear Posture Review, which includes completion of the W76-1 LEP by FY2019, life-extending the B61-12 and completing the W88 Alteration (Alt) by FY2024, life-extending the W80-4 to support the Long Range Standoff program by FY2031, pulling the W78 warhead replacement forward by one full year to support fielding the Ground Based Strategic Deterrent by FY2030, investigating the feasibility of fielding a nuclear-capable Navy flight vehicle, and sustaining the B83-1 bomb past its previously planned retirement date. Additionally, NNSA will explore future ballistic missile warhead requirements based on the threats posed by potential adversaries. Successful execution of all these responsibilities 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 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 FY2019 Budget Request**	**FY2019 Operating Target**	**FY2019 Enacted/Full-Year Budget**	**Difference between Request versus Enacted**
ASC Operations	$656.401M	$xxxM	$xxxM	
ASC Construction	$47.000M			

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 FY19. A more comprehensive list (including milestone description and evaluation/exit criteria) is included in the individual subprogram detail in the appendices. The description of the targets referenced in the FY18–FY22 Target column may be found in Appendix J.

Table V-1. ASC Level 2 Milestones for FY19[3]

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY18–FY22 Target
IC	TBD	Mix Model Advancements for Production Use	9/30/19	LLNL	None
IC	TBD	Toward a Shared Capability for Problem Setup for ASC Applications	9/30/19	LLNL	IC-3
IC	TBD	Demonstration of Programmatic Applications on Sierra	9/30/19	LLNL	IC-1
IC	TBD	Parallelization of 2D AMR	9/30/19	LLNL	IC-3
IC	TBD	ATS-1 Goal for Integrated Design Codes	9/30/19	LANL	IC-1
IC	TBD	Delivery of Improved Code Capabilities Supporting FY21 Predictive Capability Framework Peg Post	9/30/19	LANL	IC-2
IC	TBD	Sensitivity and Uncertainty Workflow of Full System SIERRA Models Supporting High Consequence Applications	9/30/19	SNL	IC-3
IC, V&V	TBD	Demonstration of Model-Based Design for NW Digital Controller Using Formal Methods	9/30/19	SNL	IC-3
IC, V&V	TBD	Automated Meshing for Analysis	9/30/19	SNL	IC-3

[3] Factors such as FY18 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	FY18–FY22 Target
PEM	TBD	Polymer Modeling	8/1/19	LLNL	PEM-2, PEM-5
PEM	TBD	Computation and Code Comparison of EOS for Light Binary Materials	3/31/19	LLNL	PEM-2, PEM-5
PEM	TBD	Generation of Higher Fidelity Mid-Z and High-Z Opacity Tables	8/31/19	LANL	PEM-1, PEM-5
PEM, IC	TBD	Demonstrate Ability to Assess Crash and Burn Scenarios	9/30/19	SNL	PEM-5
PEM, V&V	TBD	Predict Thermal Protection System (TPS) Material Response to Hypersonics Environments with Quantified Uncertainties Using Legacy Material Models and Methods	9/30/19	SNL	PEM-5
V&V	TBD	Development and Testing of Advanced Adaptive Sampling and a Fully Bayesian Approach to Uncertainty Quantification of Weapons Predictions	9/30/19	LLNL	V&V-4
V&V	TBD	Modeling of Warhead Thermo-Structural Response to a Hostile Neutron Assault	9/30/19	LLNL	V&V-5
V&V	TBD	Analytical Process Development and Execution of System Response Prediction to Combined Environments	6/30/19	LANL	IC-4
V&V	TBD	Multi-Scale Electrical UQ and Model Abstraction	9/30/19	SNL	V&V-4
ATDM	TBD	Demonstration of Dynamic, Two-way Coupling Between Modular Physics Packages	9/30/19	LLNL	None
ATDM	TBD	Demonstration Physics Capabilities for Selected DSW-relevant Applications	9/30/19	LANL	ATDM-1.b
ATDM	TBD	SPARC credibility, performance, and scalability	9/30/19	SNL	ATDM-1, ATDM-2
ATDM	TBD	Validation of EMPIRE Cavity-SGEMP Simulation Results with ICF Experimental Data	9/30/19	SNL	ATDM-4
ATDM	TBD	Portable Environment for Evaluating New Hardware	8/31/19	LLNL	ATDM-2.c
ATDM, CSSE, FOUS	TBD	Acceptance of an Arm-based Advanced Architecture Prototype System (named “Astra”) and Software Environment Development	9/30/19	SNL, LANL, LLNL	ATDM-2
ATDM, IC, CSSE	TBD	Tri-lab Co-design Milestone: Impact of Supercomputer Interconnects on Application Performance	9/30/19	SNL, LANL, LLNL	ATDM-1

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY18–FY22 Target
CSSE	TBD	Development Environment for Test Bed Systems	8/31/19	LLNL	CSSE-1.b, ATDM-2.a
CSSE, ATDM	TBD	Container Environments for Production Applications with Complex Workflows	9/30/19	LANL	CSSE-2.a
CSSE	TBD	Deploy SST capability to simulate node architectures with CPUs and GPUs	9/30/19	SNL	CSSE-3
FOUS	TBD	Large-scale Sierra Run	3/31/19	LLNL	FOUS-1
FOUS	TBD	High Performance Tape Archive Infrastructure Deployed to Unclassified Computing Environment	8/31/19	LLNL	FOUS-1, FOUS-5
FOUS	TBD	Exascale Computing Facility Modernization (ECFM) Design	6/30/19	LLNL	FOUS-4
FOUS	TBD	Improvements for ASC Data Management	8/31/19	LANL	FOUS-1, FOUS-5
FOUS, ATDM	TBD	Place "Astra" (Vanguard 1) in Operation and Configured for OHPC and SRN Advanced Prototype Work	9/30/19	SNL	FOUS-2
FOUS	TBD	Site Preparation, Integration, and Operation of Mini-Sierra ART System	9/30/19	SNL	FOUS-1

VI. Reporting Requirements

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

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

VII. Key Execution Year Reference Documents

The following ASC documents are incorporated by reference:

- *FY18 to FY22 ASC Program Plan* (July 2018)
- *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)
- Memorandum of Understanding (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 Robert Lowrie, LANL, 505-667-2121, lowrie@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 Robert Lowrie, LANL, 505-667-2121, lowrie@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 Robert Lowrie, LANL, 505-667-2121, lowrie@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 Robert Lowrie, LANL, 505-667-2121, lowrie@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 Robert Lowrie, LANL, 505-667-2121, lowrie@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 Marianne Francois, LANL, 505-667-4237, mmfran@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 Marianne Francois, LANL, 505-667-4237, mmfran@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 Marianne Francois, LANL, 505-667-4237, mmfran@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 Marianne Francois, LANL, 505-667-4237, mmfran@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 Michael Steinkamp, LANL, 505-667-4420, steinmj@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 Michael Steinkamp, LANL, 505-667-4420, steinmj@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 Michael Steinkamp, LANL, 505-667-4420, steinmj@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 Michael Steinkamp, LANL, 505-667-4420, steinmj@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

WBS	Title	Contact
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
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 Ken Alvin, SNL, 505-844-9329, kfalvin@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	Bert Still, LLNL, 925-423-7875, still1@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@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 ________________________ 8/7/2018
Signature Date

Concurrence:

Michel McCoy, LLNL ASC Executive	*Concurrence received*	8/7/2018
Bill Archer, LANL ASC Executive	*Concurrence received*	8/7/2018
(Erik Strack for) ***Scott Collis***, SNL ASC Executive	*Concurrence received*	8/7/2018

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.

1D	One Dimensional
2D	Two Dimensional
3D	Three Dimensional
AAPS	Advanced Architecture and Portability Specialist
A&E	Architectural and Engineering
AAR	Annual Assessment Review
ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AFP	Approved Funding Program
ALE	Arbitrary Lagrangian-Eulerian
Alt	Alteration
AMD	Advanced Micro Devices, Inc.
AMR	Adaptive Mesh Refinement
AMT	Asynchronous Many Task
ANL	Argonne National Laboratory
API	Application Programming Interface
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
ATOM	Accelerating Therapeutic Opportunities in Medicine

ATSE	Arm Tri-lab Software Environment
B&R	Budget and Reporting
BE	Behavioral Emulation
BEE	Build and Execute Environment
BLAS	Basic Linear Algebra Subprograms
BML	Beyond Moore's Law
BVSS	Best Value Source Selection
CBTF	Component-Based Tool Framework
C2C	Contours to Codes project
CCE	Common Computing Environment
CD	Critical Decision
CDA	Next-Generation Code Development and Applications
CDR	Conceptual Design Report
CEED	Center for Efficient Exascale Discretizations
CGMD	Coarse-Grain Molecular Dynamics
CGNS	Computational fluid dynamics General Notation System
CHAI	Copy Hiding Application Interface
CI	Continuous Integration
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
DAG	Directed Acyclic Graph
DisCom	Distance Computing
DOE	Department of Energy
DRAM	Dynamic Random Access Memory
DSW	Directed Stockpile Work

EA	Early Access
EAP	Eulerian Application Project
EC3E	Exascale Class Computer Cooling Equipment Project (LANL)
ECFM	Exascale Computing Facility Modernization
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
EMPIRE	ElectroMagnetic Plasma in Radiation Environments
EOS	Equation of State
ESAAB-E	Energy Systems Acquisition Advisory Board Equivalent
F-SEFI	Fine-Grained Soft Error Fault Injection Tool
FAODEL	Flexible, Asynchronous, Object Data-Exchange Libraries
FleCSI	Flexible Computation Science Infrastructure for Multiphysics
FOUS	Facility Operations and User Support
FPGA	Field-Programmable Gate Array
FSI	Future Supercomputing Infrastructure
GPFS	Global Parallel File System
GPU	Graphics Processing Unit
HDF5	Hierarchical Data Format 5
HEDP	High Energy Density Physics
HIO	High Performance Input/Output
HPC	High Performance Computing
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

IP	Implementation Plan
ITAR	International Traffic in Arms Regulations
ITSM	Information Technology Service Management
JBOD	Just a Bunch of Disk
KCNSC	Kansas City National Security Campus
KNL	Knights Landing processors
L2	Level 2 (milestone)
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LAP	Lagrangian Applications Project
LBANN	Livermore Big Artificial Neural Network
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LCM	Laboratory for Computational Mechanics
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
M&O	Management and Operating
MD	Molecular Dynamics
MFEM	Modular Finite Element Methods
MHD	Magnetohydrodynamics
ML	Machine Learning
MOU	Memorandum of Understanding
MPAS	Model for Prediction Across Scales
MPI	Message Passing Interface
MPICH	MPI over CHameleon
MST	Multisite Target

MW	Megawatt
NAS	Network-Attached Storage
NCHPC	Non-conventional High-Performance Computing
NCI	National Cancer Institute
NERSC	National Energy Research Scientific Computing Center
NERL	Neural Exploration & Research Laboratory
NetCDF	A set of software libraries and machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data.
NFS	Network File System
NGCE	Next-Generation Computing Enablement
NGP	Next-Generation Platform
NIC	Neural-Inspired Computing
NIH	National Institutes of Health
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NRE	Non-Recurring Engineering
NSCC	National Security Computing Center
NSF	National Science Foundation
NSCI	National Strategic Computing Initiative
NW	Nuclear Weapons
nWBS	National Work Breakdown Structure
O|SS	Open|SpeedShop
OHPC	Sandia Open High Performance Computing network
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
PII	Personally Identifiable Information
PSAAP	Predictive Science Academic Alliance Program
PSP	Predictive Science Panel
QC	Quantum Computing
QIP	Quantum Information Processing
QMU	Quantification of Margins and Uncertainties
QUBO	Quadratic Unconstrained Binary Optimization
R&D	Research and Development
RAIT	Redundant Array of Inexpensive Tape
RANS	Reynolds-Averaged Navier-Stokes
RFI	Request for Information
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
ROL	Rapid Optimization Library
RV	Reentry Vehicle
SAN	Storage Area Network
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
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
SOS	Scalable Object Store
SOW	Statement of Work
SPOT	System for Performance Optimization Tracking
SQL	Structured Query Language
SQS	Special Quasi-random Structure
SRN	Sandia Restricted Network
SRP	Stockpile Responsiveness Program
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAT	Stack Trace Analysis Tool
STPU	Space-Time data Processing Unit (SNL)
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
TMOP	Target-Matrix Optimization Paradigm
TOSS	Tri-Lab Operating System Stack
TRL	Technology Readiness Level
UB	Uncertainty Budget
UGT	Underground Test
UQ	Uncertainty Quantification
V&V	Verification and Validation
VM	Virtual Machine

VTK	Visualization Toolkit
WA	Work Authorization
WAN	Wide Area Network
WBS	Work Breakdown Structure
WCI	Weapons & Complex Integration (at LLNL)
WSC	Weapon Simulation and Computing
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 ECI and aligned with the 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 2017, and through quarter 3, fiscal year 2018, are reflected below for the ATDM subprogram.

- Released Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) 2 Request for Proposal (RFP) and evaluated responses, selected the El Capitan proposal, and began negotiations on the El Capitan Non-Recurring Engineering (NRE) and build contracts. (LLNL)
- Shepherded programming model standards needed by ASC codes and tools, including OMPD for OpenMP debugging and MPI_T for Message Passing Interface (MPI) performance analysis. (LLNL)
- Developed, maintained, and released a large suite of software supporting ASC applications including mfem, RAJA, Umpire, Spack, System for Performance Optimization Tracking (SPOT), Caliper, Scalable Checkpoint Restart (SCR), and others. (LLNL)
- Partnered with ASC code teams to integrate performance portability infrastructure into codes, achieving both impressive performance on Sierra and preparing for next-generation systems. (LLNL)
- Received Critical Decision-1 (CD-1) approval for the Exascale Computing Facility Modernization (ECFM) project. (LLNL)
- Added tools support for static analysis of Flexible Computation Science Infrastructure for Multiphysics (FleCSI) programs/applications, and C++ language extensions to

support architecture-agnostic, fine-grained data-parallel operations (forall). (LANL) (Appendix J, target ATDM-1)

- Developed two supercomputer reliability models validated on Trinity: Complex (usable for next-generation systems) and Simpler (hierarchical, composable, FIT-rate-based). (LANL) (Appendix J, target CSSE-1)
- Initial demonstration of CharlieCloud Containers with LANL applications in the open environment. (LANL) (Appendix J, target CSSE-2)
- Released the CharlieCloud Build and Execute Environment (BEE) launcher and demonstrated the BEE-CharlieCloud launcher on a LANL production machine. (LANL) (Appendix J, target CSSE-2)
- Defined four thrust areas in Machine Learning (ML) for ASC and scoped projects in these four areas: materials modeling, nuclear data evaluation, turbulence simulations, and image analysis using radiographs. (LANL) (Appendix J, target PEM-4)
- Demonstrated the feasibility of using the Singularity container framework to support running an unmodified container on a laptop, Amazon EC2 instance, and a large-scale Cray machine. (SNL) (Appendix J, target ATDM-1)
- Added support for ATS-2 (IBM Power + NVIDIA Volta) systems and Arm platforms, preparing for the 2018 DOE HPC systems Sierra (LLNL) and Astra (Sandia National Laboratories, SNL). (SNL) (Appendix J, targets ATDM-1, ATDM-2)
- Delivered Kokkoskernels-based line-implicit solvers to SPARC, resulting in an eight-fold solver speedup on the Graphics Processing Units (GPUs). (SNL) (Appendix J, target ATDM-1)
- Completed significant improvements to the Input/Output SubSystem (IOSS) library, including specialized data-decomposition methods and improved application-programmer interfaces (APIs) for SPARC and hybrid structured/unstructured meshes. (SNL) (Appendix J, target ATDM-1)
- Released the Advanced Tri-lab Software Environment (ATSE) Version 1.0. (Tri-labs) (Appendix J, target ATDM-2)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Portable Environment for Evaluating New Hardware		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 8/31/19		
ASC nWBS Subprogram: ATDM		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target ATDM-2.c		
Description: This effort will build a ubiquitous always-on performance analysis tool suite that will be integrated with an ASC code. Previous work in the Spot project built infrastructure that tracked the performance of an application's nightly tests; this work will expand that infrastructure to track the performance of arbitrary application runs. Performance data collection will be tightly integrated into the target application's workflow, which will feed web-based visualizations developed by the Spot team. The tool suite will include components for measuring and visualizing wall time, MPI, I/O, and performance counters on CTS-1 and Sierra systems.		
Completion Criteria: An ASC code team member will verify that the team has delivered a ubiquitous performance analysis tool suite that is integrated with an ASC code.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: ATDM and IC		

Milestone (ID#): Acceptance of an Arm-based Advanced Architecture Prototype System (named “Astra”) and Software Environment Development		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 9/30/19		
ASC nWBS Subprogram: ATDM		
Participating Sites: SNL (lead), LANL, LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target ATDM-2		
Description: This milestone will complete delivery and acceptance of the Arm-based Vanguard project “Astra” HPC platform, as well as a system software stack environment for future large-scale prototype systems. The scope of this stack includes the full range of software needed to operate Astra including system management, base operating system, network stack, parallel I/O subsystems, file system integration, job scheduling and resource management, runtime systems, and the user-facing programming environment. A tri-lab team will develop and test the software stack in collaboration with Arm technology vendors, including HPE, Arm Inc., Cavium, and Linaro, using a set of NNSA codes running on Astra while it is initially operating in unclassified networks. Successful completion of this milestone will include providing the evidence needed to show that Astra is ready to be moved to a classified network. The milestone will include the completion of Version 2 of the ATSE release that is distributed to the tri-lab community and installed on the Astra platform and related advanced architecture systems. Preliminary porting and testing of ASC IC applications will also be included as part of this milestone.		
Completion Criteria: This milestone will be completed when the above-mentioned demonstrations, benchmarks, and analyses are complete, and a report or presentation is provided summarizing the key results and recommendations. The results of this milestone will inform the decision on when Astra moves to the classified network, but this milestone’s success does not depend on Astra transitioning to the classified network in FY19.		
Customer: ASC programs at SNL, LANL, and LLNL		
Milestone Certification Method: A program review is conducted and its results are documented. A report is prepared as a record of milestone completion.		
Supporting Resources: FOUS, Advanced Architecture Testbeds project (SNL), Sandia Institutional Computing Program (facilities and infrastructure), CSSE Scalable Systems Software.		

Milestone (ID#TBD): Tri-lab Co-design Milestone: Impact of Supercomputer Interconnects on Application Performance

Level: 2 | **Fiscal Year**: FY19 | **DOE Area/Campaign**: ASC

Completion Date: 9/30/19

ASC nWBS Subprogram: ATDM, IC, CSSE

Participating Sites: SNL, LANL, LLNL

Participating Programs/Campaigns: ASC

Description: Next-generation architectures are increasingly relying on novel network topologies and sophisticated network features to meet DOE performance requirements. However, the performance trade-offs between purchasing more capable (and more expensive) networking features or additional compute nodes are not fully understood. The FY19 co-design milestone will examine the impact of network interconnects on the performance of individual applications as well as overall lab workloads. The milestone team will work with vendors to analyze DOE workloads and applications to quantify the performance impacts of network options. Work will include evaluation of various network tapering strategies as well as new topologies enabled by optical interconnects. The milestone team will also study proposed solutions for on-node bottlenecks to internode communication such as data motion between accelerators and network endpoints.

Completion Criteria:

A Tri-lab report containing:

1) A study of how network tapering on the ATS and/or CTS machine impacts performance of lab workflows with recommendations for future procurements.
2) Tests of on-node hardware and software strategies to optimize data movement to/from interconnect endpoints including performance impacts on at least one ASC code per lab.
3) An evaluation of novel interconnect topologies enabled by optical interconnects with performance estimates for lab workflows.

Customer: ASC

Milestone Certification Method:

A milestone review is conducted, and its results are documented.

A report is prepared as a record of milestone completion.

Supporting Resources: ASD, CSSE, and FOUS Staff

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

The 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 ATDM CDA product. 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 ASCR program, 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.

Accomplishments in FY18:

- Released CORAL2 RFP, evaluated responses, selected a vendor, and began negotiations on NRE and build contracts. (Appendix J, targets CSSE-1.b, CSSE-2.b)
- Expanded power and power variation studies to Sierra and Sierra Early Access (EA) systems. (Appendix J, target ATDM-1.d)
- Studied the impact of on-node noise on ASC codes. (Appendix J, target ATDM-1.d)
- Flux demonstrated the effectiveness of fully hierarchical scheduling on Sierra as a representative complex exascale workload, and for UQ-Pipeline-based simulations for ASC codes. (Appendix J, target ATDM-1.c)
- Expanded debugging and code correctness capabilities for ASC codes via tool efforts by the Pruners and ROSE teams. (Appendix J, target ATDM-1.d)
- Continued participation in vendor meetings, Center of Excellence (COE), and co-design activities for CORAL, as well as technical management of academic contracts and other external interactions. (Appendix J, target ATDM-1.c)

Planned Activities in FY19:

- Complete negotiations on and award the El Capitan NRE and build contracts. (Appendix J, targets CSSE-1.b, CSSE-2.b)
- Provide technical coordination and contractual management for the El Capitan NRE, COE, and build contracts. (Appendix J, targets CSSE-1.b, CSSE-2.b)
- Participate in NRE and build subcontract activities for Argonne National Laboratory's (ANL) A21 system and ORNL's Frontier system. (Appendix J, targets CSSE-1.b, CSSE-2.b)
- Harden the ROSE Tool infrastructure through code-specific automated testing to insure robust support for ASC applications. (Appendix J, target ATDM-1.c)
- Develop debugging and code correctness tools through the Pruners tools and with the LANL Fine-Grained Soft Error Fault Injection Tool (F-SEFI) team's DisCVar framework. (Appendix J, target ATDM-1.d)
- Expand power measurement and control capabilities through porting msr-safe/libmsr, GEOPM, and Exascale Computing Project (ECP) Power Steering to additional architectures of interest to the ASC platform effort. (Appendix J, targets CSSE-2.a, CSSE-2.b)
- Enhance capabilities of the Flux hierarchical resource scheduler to address relevant ASC workloads for exascale, including evaluation of machine-learning-driven workflows. (Appendix J, target ATDM-1.d)

Advanced Architectures and Portability Specialists Project (LLNL)

LLNL ATDM supports the development of next-generation simulation capabilities with the Advanced Architecture and Portability Specialists (AAPS) team, which disseminates advanced architecture expertise to code teams and supports portability abstractions such as RAJA and Umpire. Rather than relying upon each code team to independently stay abreast of the developments in architecture, programming models, and kernel optimization needed to make efficient use of new hardware, ASC code developers have access to an agile knowledge and labor pool of computational/computer scientists skilled in scaling applications on new, cutting-edge hardware. The team shares knowledge through a code repository, wiki, seminars, hackathons, and publications. The team includes specialists in key areas such as GPU programming, many-core programming, I/O, and parallel application development.

Accomplishments in FY18:

- Integrated RAJA improvements and UMPIRE-based memory management into ARES. Delivered clang-based compiler plug-ins that identify common coding errors. Worked with the Modular Finite Element Methods (MFEM) and Blast teams to prototype RAJA optimizations in the Laghos proxy application, which improved performance by over 10 times. Started incorporating lessons learned from Laghos back into MARBL. (Appendix J, target ATDM-1.d)

- Ported TETON to GPUs, achieving up to 8× speedup in the nonlinear solver. Enabled ARDRA on GPUs by refactoring data structures to support arbitrary partitioning of problem dimensions. Lessons learned from the Quicksilver proxy app guided the MERCURY team's GPU refactoring. (Appendix J, target ATDM-1.d)
- Developed OpenMP target offload and ROCm backends for RAJA. Refactored CHAI to improve design, performance, and usability. Released Umpire as open source. Contributed to RAJA enhancements including shared memory views and improved performance for reductions and nested loops. (Appendix J, target IC-1)
- Supported CORAL-2 benchmarking. Identified and reported 40+ compiler bugs to IBM and NVIDIA. Contributions to the OpenMP standard included C++17 support, deep copy, and a reworking of defaultmap. (Appendix J, target ATDM-1.d)

Planned Activities in FY19:

- Continue working with multiphysics code teams including Kull, ARES, ALE3D, and MARBL toward portable performance on Sierra. (Appendix J, target ATDM-1.d)
- Work with PEM codes such as LEOS, Cheetah, MCGIDI, Cretin, and MSLIB to help them meet performance goals on Sierra. (Appendix J, target ATDM-1.b)
- Continue to assist in the development of RAJA, CHAI, and Umpire to further the performance portability goals of ASC applications. (Appendix J, target IC-1)
- Maintain leadership in architecture co-design by working with vendors; continue investigating programming models via proxy apps. (Appendix J, target ATDM-1.d)
- Share knowledge via wikis, talks, publications, etc. (Appendix J, target ATDM-1.b)

Production-Quality Tools Development Project (LLNL)

LLNL has built up a strong R&D portfolio in tools for performance analysis, debugging, correctness verification, power-aware computing, and resilience support. Enabling ASC code teams, especially in the context of ATDM efforts, to leverage and efficiently deploy software developed that leverages these tools is essential but also requires significant efforts to ensure the tool software is hardened to production quality, maintained on ASC systems, and documented for end users. The ProTools development team ensures that these goals are met, working closely with the R&D efforts on one side and user support and AAPS teams on the other to deliver hardened reliable tools to code teams.

Accomplishments in FY18:

- Alongside devRAMP, worked with ARES to modernize their Spack recipes and create a centralized repository for LLNL ASC codes. (Appendix J, target ATDM-1.c)
- Ported SCR to Sierra and integrated it with ARES. (Appendix J, target ATDM-1.c)

- Integrated Spot into ARDRA and began ARES integration. Spot 1.0 tracks performance changes through a code's nightly tests. (Appendix J, target ATDM-1.c)
- Continued development of Caliper, focusing on production-required features needed by ASC codes. Released the initial version of Gotcha, which is infrastructure for always-on performance tools. (Appendix J, target ATDM-1.c)
- Alongside the Standards and Architectures for Storage and I/O (SASI) project, shepherded the OpenMP debugging standard, OMPD, into likely acceptance in the OpenMP 5.0 standard. (Appendix J, target IC-1)
- Working with AAPS, integrated Umpire into ARES. (Appendix J, target IC-1)

Planned Activities in FY19:

- Integrate Caliper into ARES and MARBL. Expand Caliper with new measurement capabilities needed for Spot 2.0. (Appendix J, target ATDM-1.c)
- Develop and deliver Spot 2.0 to the ARES and MARBL code teams, who will use the tool to track performance of arbitrary application runs. Create visualizations corresponding to the new capabilities in Caliper. (Appendix J, target ATDM-1.c)
- Support the Sonar project by deploying data analytics tools, which will automatically analyze site-wide monitoring data for anomalies. (Appendix J, target ATDM-1.c)
- Continue supporting the ATDM AXOM project by integrating UMPIRE into ASC codes. (Appendix J, target ATDM-1.c)
- Continue the ongoing effort of identifying the next generation of research tools, and working with ASC code teams, AAPS, and management to prioritize tools for productization. (Appendix J, target ATDM-1.d)

User Workflow and Modernization (LLNL)

The Workflow project builds infrastructure to enhance the end-to-end productivity of ASC HPC assets. The project focuses on building solutions to concrete user problems in three areas: problem setup, simulation management, and post-processing. The team builds best-practice components that encompass common operations, allowing users to build custom workflows while reducing duplicated work. The Workflow project has three active software projects: Siboka, components and infrastructure supporting user workflows; the Contours to Codes (C2C) project, which focuses on building a common geometry specification compatible with LANL contours; and ROVER, an in-situ multigroup ray casting package for both simulated diagnostics and volume rendering.

Accomplishments in FY18:

- Siboka released Sina (simulation insight and analysis) v1.0, which imports, exports, and queries metadata in Structured Query Language (SQL) and NoSQL databases. The team also released Mnoda, a schema for ingesting simulation information into Sina. Sina and Mnoda imported tens of thousands of records, with millions of scalars and paths, into SQL and NoSQL databases at LLNL. The team demonstrated analytics on the data with JupyterHub. (Appendix J, target IC-3)

- The C2C team released v0.5.2 with C++ support, Draco (PMESH) integration, and a repository structure for assembly files that handles use cases not supported by the LANL format. The tool is now in use by early adopters and being considered as part of a broad geometry interchange effort in LLNL's Weapons & Complex Integration (WCI) organization. (Appendix J, target IC-3)
- ROVER was released as open source, and underlying geometry algorithms were incorporated into Visualization Toolkit (VTK)-m. ROVER now provides native ray-tracing for cylindrical coordinate problems. Work is underway to demonstrate the capability in-situ in with the ECP ALPINE and ATDM MARBL projects. (Appendix J, target IC-3)
- The project contributed to a prototype web portal called syflux, incorporating lessons learned with the VV4ALE3D, and the ability to manage data and media sets and link those with problem setups, documentation, and source data. (Appendix J, target IC-3)

Planned Activities in FY19:

- Siboka: develop C/C++ Mnoda library and integrate it with ASC codes. Explore integration with Caliper and Sonar. Enhance Sina for use-cases in Engineering and Design Physics. Continue tri-lab collaborations. (Appendix J, target IC-3)
- C2C: integrate C++ parser library into ASC application codes, work with community to integrate tools as part of geometry interchange effort, explore modeling and repository use cases. (Appendix J, target IC-3)
- ROVER: integrate Alpine and demonstrate in-situ capability in MARBL. Develop methods for ray tracing higher-order finite element meshes. (Appendix J, target IC-3)
- Web applications: explore resource integration with Granta and document management systems, explore collaborations with performance analysis applications, and use-cases for Jupyter notebooks in the data center. (Appendix J, target IC-3)

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

DevRAMP includes productivity multipliers for all stages of the HPC software development pipeline. This includes infrastructure for developer operations tasks such as monitoring, databases, continuous integration, and software deployment, as well as data analysis and tuning. DevRAMP focuses on improving the efficiency of the entire development cycle: Develop, Test, Deploy, Monitor, Analyze, and Tune, particularly in 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 FY18:

- New Spack efforts include “environments” that enable new user workflows and new concretizer for faster and more robust dependency resolution. Worked with Kitware to add binary packaging and cloud-based testing. Worked with LLNL Weapon Simulation and Computing (WSC) teams to modernize Spack usage. (Appendix J, target ATDM-1.c)
- Deployed JupyterHub with end-to-end transport layer security and developed a library for notebooks to talk to running simulations. (Appendix J, target ATDM-1.c)
- Spearheaded the ECP Continuous Integration (CI) technical specification and RFP and worked with ECP to select a vendor and manage the contract.
- Wrote a new authentication mechanism for the Sonar database. Continued to work with Livermore Computing (LC) on network and workload evaluation with Sonar. (Appendix J, target ATDM-1.c)
- Developed new Apollo autotuning capabilities with AAPS. (Appendix J, target IC-1)
- Worked with LC on containerized service deployment using Kubernetes.

Planned Activities in FY19:

- Complete Spack environments work by adding “view” support. Add features for multiuser deployments, parallel builds, and fine-grained architecture optimization. Create a shared package repository and internal binary mirror for WSC. (Appendix J, target ATDM-1.c)
- Continue to co-lead the ECP CI effort. Work with Onyx Point on secure Gitlab deployment at LC and other facilities.
- Contribute Jupyter hardening work upstream so that it can be maintained externally and used broadly by other labs. Develop a notebook interface for ScrubJay.
- Work with LC and customers to standardize web technologies across applications deployed in LC. Develop an authentication mechanism for staff- and user-provisioned containerized services in LC.
- Secure Kafka implementation on Sonar for multitenant usage. (Appendix J, target ATDM-1.c)
- Deploy internal CDash instances in LC using Kitware’s Docker-based CDash distribution.

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 FY18:

- Software releases: mfem-3.3.2 and mfem-3.4, which included features for ATDM applications. MFEM joined the xSDK project, the FASTMath SciDAC institute, and the OpenHPC distribution. (Appendix J, target ATDM-1.a)
- Improvements in parallel nonconforming adaptive mesh refinement (AMR) implementation: construction of parallel interpolation, parallel refinement and rebalancing, demonstrated excellent weak and strong scalability during runs on the full Vulcan BG/Q machine at LLNL (about 400k MPI tasks); updated AMR benchmarks; demonstrated the benefits in the Laghos miniapp; incorporating into the BLAST code; documented results and algorithms in a submitted journal publication. (Appendix J, target ATDM-1.a)
- Arbitrary Lagrangian-Eulerian (ALE) discretization improvements for compressible flow in BLAST: initial support for multigroup rad-diffusion; high-order mesh optimization based on the target-matrix optimization paradigm (TMOP) from the ETHOS project; improved closure model. (Appendix J, target ATDM-1.a)
- Actively engaged with ATDM application work: participating in weekly MARBL and BLAST meetings; supporting MARBL's Level 2 (L2) milestones; help with GPU code refactoring, leveraging advanced partial assembly algorithms work in MFEM and ECP/Center for Efficient Exascale Discretizations (CEED); Laghos, a new MFEM-based proxy for BLAST developed under the CEED co-design center in the ECP, was added as a CORAL-2 procurement benchmark, listed also in ECP's proxy app list. (Appendix J, target ATDM-1.a)

Planned Activities in FY19:

- Improve parallel nonconforming AMR with features such as parallel anisotropic AMR; triangle/tet bisection in nonconforming meshes and space-filling curve coarse element reordering for parallel load balancing. Demonstrate large-scale parallel dynamic AMR in an ATDM-relevant application. (Appendix J, target ATDM-1.b)
- Continue engagement with GPUs and advanced partial assembly algorithms work in BLAST, coordinating with and leveraging work in MFEM and the ECP/CEED project. (Appendix J, target ATDM-1.b)
- Continue engagement with ATDM application work and provide support for MARBL L2 milestones. (Appendix J, target ATDM-1.b)

Tri-lab Flang Subcontract (LLNL)

Flang is an open-source, Low-Level Virtual Machine (LLVM)-native FORTRAN compiler being developed under subcontract by NVIDIA through tri-lab ASC support. FY19 represents the final year of the four-year contract and will build upon the prior three years of effort by continuing to expand the set of language features provided, pursuing performance optimizations, rewriting large sections of code to be LLVM-compliant, establishing a regular release cycle for support at DOE facilities, engaging additional external contributors, and finishing work on an OpenMP 4.5 GPU-offload capability. A public roadmap for development is published at https://github.com/flang-compiler/flang/wiki/Roadmap and is regularly updated as priorities emerge. 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 stated goal of this contract was to deliver a version of Flang that would be self-sustaining in the community, and plans are in place to establish ongoing funding support through ECP/ASCR funding post-FY19 as part of that.

This subcontract is managed by LLNL, with technical input and oversight provided by LANL, SNL, and other DOE stakeholders in the ECP. This NNSA project encompasses the following ECP project:

2.3.5.06 ECP FLANG project

Accomplishments in FY18:

- Established Contributor License Agreements to allow for external contributions, with Arm Holdings being an initial contributor as part of their strategy to use Flang in their production HPC stack.
- Implemented a binary release schedule that closely tracks LLVM through v6, optimized math intrinsics, improved DWARF debugging, and added OpenMP 4.5 support on multicore central processing units (CPUs). Began work on OpenMP target offload to GPUs.
- Had regular interactions with multiple NNSA application teams. Held community-building activities (ECP Annual Meeting, EuroLLVM, others).
- Rewrote front-end parser in LLVM-compliant C++, with full support for the Fortran2018 standard. Migrating source to baseline C++ for better integration with LLVM.
- Implemented priority F2008 features including SUBMODULE and BLOCK.

Planned activities in FY19:

- Complete support for OpenMP 4.5 target offload to GPUs. Add Spack support and additional platforms, e.g., Mac OS X or Windows. Implement new F18 parser with benchmarking and performance optimizations.
- Continue converting to C++ to support additional LLVM features.

- Continue to integrate/test Flang with each LLVM release on a timely basis and engage the LLVM community to push toward adoption of Flang as a first-class LLVM subproject.

Programming Model Standards and Architectures for Storage and I/O (SASI) (LLNL)

This project covers two key items critical for ASC's next generation computing: programming models standards, and storage and I/O. The standards work is on MPI and OpenMP. For MPI, we focus on tool support, including the MPI_T and MPIR interfaces, and for fault tolerance, including Reinit and fault-tolerance support in MPI. In OpenMP, our focus is on the tools interfaces OMPT and OMPD. However, we participate and monitor developments in all parts of both standards. 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 ASCR program, vendors, and academia, such as planning and technical coordination for vendor and academic contracts and CORAL working group efforts.

Accomplishments in FY18:

- Maintained leadership role in the MPI and OpenMP standards bodies. (Appendix J, target IC-1)
- Co-developed the OMPD specification (debugging OpenMP), which was accepted in a preview of OpenMP 5.0. (Appendix J, target IC-1)
- Progressed on the MPI_T Events specification for notifying tools about MPI implementation events. (Appendix J, target IC-1)
- Extended the SCR API to support application needs, including general file output. Componentized SCR front-end. (Appendix J, target ATDM-1.a)
- Continued work on file system futures: prototyping I/O support for machine learning, evaluated commercial burst buffers, participated in National Science Foundation (NSF) and ASCR I/O workshops, and continued development of DissectIO I/O tracer. (Appendix J, target CSSE-2.c)

Planned Activities in FY19:

- Continue leadership in MPI and OpenMP standardization. Improve OMPD and OMPT to support advanced tool functionality. Participate in the MPI Forum's fault-tolerance working group. Present the MPI_T Events specification to the MPI Forum for a vote. (Appendix J, target IC-1)

- 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 I/O performance measurement. (Appendix J, target CSSE-2.2)
- Continue SCR modification for ASC applications for support of Sierra platform. (Appendix J, target IC-1)
- Continue vendor interactions, including participation in CORAL, CORAL2, and PathForward efforts. (Appendix J, targets CSSE-1, ATDM-1.c, ATDM-2.a)

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 FY18:

- Presented work to the Predictive Science Panel (PSP).
- Developed a strategic plan, incorporating feedback from the PSP and lessons learned from early investigations.
- Wrote a white paper, presented a talk, and followed up with potential Predictive Science Academic Alliance Program (PSAAP) III participants.
- Implemented a one-dimensional (1D) k-L Reynolds-Averaged Navier-Stokes (RANS) model and demonstrated the ability to solve the discrete adjoint for an unsteady Rayleigh-Taylor surrogate problem.
- Provided feature extraction for image-based neural networks for a scaling experiment using Livermore Big Artificial Neural Network (LBANN).

Planned Activities in FY19:

- Integrate ALE tangling prediction and mitigation into two ASC codes, with regression testing. (Appendix J, target IC-3)
- Analyze and reduce errors introduced by machine learning techniques for numerical solutions, such as equation of state (EOS) interpolation. (Appendix J, target IC-2)

- Begin the evaluation of large-scale machine learning frameworks for ASC applications. (Appendix J, target IC-3)
- Early demonstration of applying machine learning techniques to construct a regression model for the RANS correction term. (Appendix J, target IC-2)

Exascale PathForward Project (LLNL)

The PathForward project is an element of the DOE ECP Hardware and Integration (HI) focus area and is intended to accelerate the R&D of critical hardware technologies needed for exascale computing. This is a three-year project ending early in 2020. In 2017, LLNL awarded six commercial R&D contracts to AMD, Cray, IBM, Intel, HPE, and NVIDIA. The contracts are being administered by LLNL on behalf of ECP and under technical project management of the ECP HI team and personnel of six DOE national laboratories (ANL, Lawrence Berkeley National Laboratory (LBNL), LLNL, LANL, ORNL, and SNL).

Accomplishments in FY18:

- Provided management and oversight of the PathForward project for ECP.
- Provided technical coordination and contractual management for the R&D contracts along with personnel from other DOE national laboratories. (Appendix J, target CSSE-1.b)
- Planned and held two reviews of PathForward progress. (Appendix J, target CSSE-1.b)
- Completed an assessment document of PathForward progress, titled *Q4 CY2017 PathForward Assessment: WBS 2.4.1, Milestone PM-HI-1050*, dated January 2, 2018.

Planned Activities in FY19:

- Continue providing management and oversight of the PathForward project for ECP.
- Continue providing technical coordination and contractual management for the R&D contracts along with personnel from other DOE national laboratories. (Appendix J, target CSSE-1.b)
- Plan and hold two reviews of PathForward progress. (Appendix J, target CSSE-1.b)

Cross-Cutting Extreme-Scale Research (LANL)

This is a cross-cutting applied research project that spans the CSSE 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 projects to leverage and extend the expertise, focusing on enabling the ATDM CDA project and extreme-scale environment. The project areas include in situ data analytics for large-scale simulations and resilience for application fault injection, and system fault domain characterization; developing an exascale in situ workflow including in situ data reduction approaches and post-processing visual analysis of reduced-size data extracts via Cinema; and developing high-performance, portable, multithreaded VTK algorithms for use with ParaView, Cinema, and the VTK-m project.

This NNSA project encompasses the following ECP projects:

2.3.4.02: Cinema – In Situ Visualization

2.3.5.02: RESL – Resilience characterization

Accomplishments in FY18:

- In concert with FOUS, implemented a scale-out monitoring and analytics infrastructure capable of meeting HPC requirements for data capacity, longevity, and ingest rate.
- Developed two supercomputer reliability models validated on Trinity: Complex (usable for next-generation systems) and Simpler (hierarchical, composable, FIT-rate-based). (Appendix J, target CSSE-1)
- Evaluated vulnerability of applications and developed an error detection/correction mechanism; for a cost of 2.4% in performance we corrected 100% of faults. (Appendix J, target ATDM-1)
- Investigated fault injection into the FleCSALE ATDM-CDA proxy application. (Appendix J, target ATDM-1)
- Specified a RESTful API for creation of dependent database metadata and artifacts that are created on demand, utilizing analysis algorithms that operate on an existing Cinema database.

Planned Activities in FY19:

- Use the Resilience model to evaluate and compare prospective ECP architectures. (Appendix J, target CSSE-1)
- Add arbitrary data types and operators to the Cinema image database including direct support for 1D and two-dimensional (2D) plotting. (Appendix J, target IC-3)
- Demonstrate and deliver accelerated ParaView, Cinema, and vtk-m workflow with the NGC exascale application, including plots, acceleration, and querying. (Appendix J, target IC-3)

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:

2.3.1.02: Legion – Asynchronous data-centric task-based programming model

2.3.2.02: ATDM Tools –LLVM compiler and runtime tools and QUO thread-level heterogeneity in coupled MPI applications

2.3.3.02: 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 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 FY18:

- Completed integration of Fortran task support in Legion to support FleCSI use cases.
- Added tools support for static analysis of FleCSI programs/applications, and C++ language extensions to support architecture-agnostic, fine-grained data-parallel operations (forall). (Appendix J, target ATDM-1)
- Implemented a common shared-memory cache for QUO. (Appendix J, target PEM-4)
- Flang source code became available to the public on May 17, 2018. The Model for Prediction Across Scales (MPAS) performance spreadsheet was updated to reflect the most recent performance data. Concluded the quarterly progress review with bugs and performance evaluation. (Appendix J, targets IC-1, ATDM-2)
- Finished OpenMP/libomp backend target to support tasking, parallel-for, and reduce constructs. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Complete initial implementation of a dynamic application (multimaterial, adapting mesh) using control replication. (Appendix J, target ATDM-1)

- Develop FleCSI static analysis tool that identifies interface calls and emits errors for incorrect syntax. (Appendix J, target ATDM-1)
- Use FleCSI framework to satisfy Ristra multiphysics application deliverables. Deliver radiation-hydrodynamics simulation code that can solve a selected problem set. (Appendix J, target ATDM-1)
- Complete implementation of the Kitsune back-end for GPU target. (Appendix J, target ATDM-1)

Future-Generation Computing Technologies (LANL)

This project includes high-risk, high-reward research for future systems, including research on virtual and containerized environments, advanced in situ analytics, and machine learning.

A key component of this work is to 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.

The 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 ML component has efforts in collaboration with PEM, V&V, and ATDM-CDA. Investigating in-line ML support for multiscale efforts in EOS and material models. 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. Initial thrust areas include studies of ML methodologies applied to materials modeling, nuclear data evaluation, turbulence simulations, and image analysis using radiographs.

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.

This NNSA project encompasses these ECP projects:

2.3.5.02: BEE/CharlieCloud Containers and Container Orchestration for HPC systems; LANL Machine Learning and Application Performance

Accomplishments in FY18:

- Initial demonstration of CharlieCloud Containers with LANL applications in the open environment. (Appendix J, targets IC-3 and CSSE-2)

- Released the CharlieCloud BEE launcher and demonstrated the BEE-CharlieCloud launcher on a LANL production machine. (Appendix J, targets IC-3 and CSSE-2)
- Defined four thrust areas in machine learning for ASC and scoped projects in these areas. (Appendix J, target PEM-4)

Planned Activities in FY19:

- Demonstrate ASC Lagrangian Applications Project (LAP) application in a container with complex setup and in situ visualization running on a production HPC system in the secure environment. (Appendix J, targets IC-3 and CSSE-2)
- Conduct initial studies of ML methodologies applied to materials modeling, nuclear data evaluation, turbulence simulations, and image analysis using radiographs. (Appendix J, target PEM-4)

Programming Models and Runtimes (SNL)

This project focuses on four categories of work associated with the Kokkos performance portability library; the work is aligned with the Kokkos strategic plan. The categories are:

- I: Continuous Effort of Application Support and Kokkos Maintenance
- II: Support for new hardware architectures, potentially through new Backends
- III: Optimization of Existing Backends for new Compilers and new iterations of Hardware
- IV: C++ Standards Work and Kokkos Community Growth

Accomplishments in FY18:

- Added support for ATS-2 (IBM Power plus NVIDIA Volta) system, preparing for Sierra (LLNL) and Astra (SNL). (Appendix J, targets ATDM-1, ATDM-2)
- Co-designed with AMD a new backend for AMD GPUs using the ROCm software stack to support AMD's Path Forward project within ECP. (Appendix J, target ATDM-1)
- Implemented prototype remote memory space capability, which will allow a seamless integration of PGAS programming paradigms into Kokkos-based applications. The current prototype has two back ends: SHMEM and QUO. (Appendix J, target ATDM-1)
- Supported the adoption of Kokkos by numerous application customers, including Sandia ATDM and across multiple DOE Laboratories. (Appendix J, target ATDM-1)
- Advanced C++ standard proposals for multidimensional arrays with layouts and atomic reference objects through the ISO committee. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Provide high-quality (production) Kokkos support and consultation for ASC applications and libraries. (Appendix J, target ATDM-1)

- Harden and optimize the ROCm-based AMD GPU back end, develop a prototype back end for the Intel ECP Path Forward architecture, and improve the existing prototype Remote Memory Space capabilities. (Appendix J, target ATDM-1)
- Optimize existing back ends for improved hardware and compiler capabilities, focusing on the OpenMP runtime on Intel platforms. (Appendix J, target ATDM-1)
- Engage the C++ standards committee to further the adoption of successful Kokkos concepts into the standard. (Appendix J, target ATDM-1)

Development Tools (SNL)

The development environment and infrastructure is often referred to as DevOps. DevOps in this context includes all the software infrastructure development, testing support, integration, and deployment work in support of the ATDM software application and component development teams. Primary responsibilities for the team are to: (1) coordinate and prioritize tasks for the various teams that provide "DevOps" support for ATDM codes, applications, and customers; (2) develop and help deploy shared build, test, and install infrastructure across the ATDM codes and projects; (3) define and support development, testing, integration, and other related workflows for ATDM projects. The performance analysis component of this project is focused on delivering a deeper understanding of the performance bottlenecks for our ATDM codes and strategies to improve that performance on pre-exascale and exascale HPC architectures. This project is essential to the success of the SNL ATDM application activities and seeks to create ties with similar activities for the broader set of ECP applications.

Accomplishments in FY18:

- Assessed vectorization levels achieved in SPARC (with TRILINOS solver stack), which is at nearly 60% of floating point operations. (Appendix J, target ATDM-1)
- Deployed automated check-in testing (pull request) for TRILINOS. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Coordinate and support General DevOps for ATDM. (Appendix J, target ATDM-1)
- Deploy scalable build, integration, and test framework. (Appendix J, target ATDM-1)
- Conduct TRILINOS integration and multiplatform CI testing for ATDM apps. (Appendix J, ATDM-1)
- Conduct performance analysis support studies by the Application Performance Team for SPARC, including assessing SPARC and TRILINOS kernel execution times. (Appendix J, target ATDM-1)
- Provide application/programming environment support and application bring-up on Arm systems in preparation for Astra. (Appendix J, target ATDM-2)

- Develop and improve Kokkos performance tools based on requirements from the EMPIRE and TRILINOS development teams. (Appendix J, target ATDM-1)
- Support scaling runs and performance breakdown/analysis of EMPIRE and SPARC on thousands of KNL nodes on Trinity. (Appendix J, target ATDM-1)
- Support initial scaling runs on ATS-2 LLNL Sierra systems or equivalent Sandia prototype systems. (Appendix J, targets ATDM-1, ATDM-2)

Math Libraries (SNL)

This project develops and integrates scalable, modular, and cross-cutting software infrastructure components for ATDM and other future exascale applications that utilize, where appropriate, ATDM Core CS components: Kokkos performance portability, Sacado/Rapid Optimization Library (ROL) embedded sensitivity analysis and optimization technology, DARMA asynchronous many-tasking, and DataWarehouse.

The project combines algorithmic R&D with delivery of interoperable software components that are expected to be crucial capabilities that will enable Sandia's ATDM application codes to be performance portable across next-generation computing architectures such as GPUs and Xeon Phis. This work will include integration of these components into the application codes and improving their design and interfaces for mission relevant use cases.

Accomplishments in FY18:

- Delivered KokkosKernels-based line-implicit solvers to SPARC, resulting in an 8× solver speed-up on GPUs. (Appendix J, target ATDM-1)
- Coupled the Tempus time integrator to the EMPIRE/Particle in Cell (PIC) application, enabling higher-order time integration schemes. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Develop, integrate, baseline, and support agile components in ASC/ATDM applications. (Appendix J, target ATDM-2)
- Show strong and weak scaling of MueLu for steady-state hypersonic RANS physics for a stockpile reentry system. (Appendix J, target ATDM-1)
- Demonstrate scalability of the electromagnetics solver for an EMPIRE problem of interest on GPU architecture. (Appendix J, targets ATDM-1, ATDM-2)
- Continue solver algorithm development for SPARC with Wyoming and document algorithmic approaches and gains. (Appendix J, target ATDM-1)

Data and Visualization (SNL)

The scalable data-management tasks provide the necessary R&D to support the I/O performance and portability requirements of the ATDM applications and system software. Without this R&D, ATDM applications will not be able to effectively manage and store data in the more complex AT systems. The work includes support and maintenance of production-level I/O libraries such as IOSS, netCDF, and HDF5, as well as the research-grade Data Warehouse—I/O middleware that provides efficient data movement and management of transient data within an HPC system. This data-management layer leverages available system memory, nonvolatile memory (e.g., burst buffers), and parallel file systems and is available as both a user-level library and a persistent service to HPC applications, application workflows, and various runtime systems (e.g., AMT systems). The R&D performed will be evaluated and deployed on existing ASC platforms (primarily AT systems), Sandia's advanced architecture testbed, and ASC-developed simulation capabilities (where necessary) to demonstrate proof-of-concept software and develop production-capable packages for use by the ATDM program with eventual deployment to the IC program as well as the broader needs of the Exascale Computing Project.

Accomplishments in FY18:

- Completed significant improvements to the IOSS library including specialized data-decomposition methods and improved APIs for SPARC and hybrid structured/unstructured meshes. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Improve production IOSS supporting SPARC, EMPIRE, and some of the newly added ATDM applications. (Appendix J, target ATDM-1)
- Develop capabilities for integrated mesh-refinement that adds spatial awareness to the particle interfaces for our particle I/O methods. (Appendix J, target ATDM-1)
- Develop a technology demonstration for the I/O research products (e.g., FAODEL) to evaluate utility and performance on AT platforms. Appendix J, target ATDM-1)
- Production integration of ParaView/Catalyst in-situ visualization for SPARC.
- Complete and demonstrate Catalyst-integrated TuckerMPI support for compression of nonrectilinear meshes.
- Develop Kokkos-enabled tensor-based compression algorithms using TuckerMPI.
- Continue to develop machine-learning capabilities to support automated meshing, adaptive solvers, and "smart" system software.

Software Ecosystem and Delivery (SNL)

This project consists of two efforts or subprojects. The purpose of the first subproject is to broaden the impact of programming model technologies developed in ATDM by integrating these technologies into an existing ASC code suite, evaluating the ability of these technologies to be incrementally added into a legacy code base, and to deliver performance improvements on current- and next-generation computing platforms. This subproject focuses on using and evaluating Kokkos and asynchronous many task (AMT) programming models in the NGP Contact code. The second subproject focuses on the design, implementation, and evaluation of operating system and runtime system (OS/R) interfaces, mechanisms, and policies supporting the efficient execution of the ATDM application codes on next-generation ASC platforms. Priorities in this area include the development of lightweight tasking techniques that integrate network communication, interfaces between the runtime and OS for management of critical resources (including multilevel memory, nonvolatile memory, and network interfaces), portable interfaces for managing power and energy, and resource isolation strategies at the operating system level that maintain scalability and performance while providing a more full-featured set of system services.

Accomplishments in FY18:

- Demonstrated the feasibility of using the Singularity container framework to support running an unmodified container on a laptop, Amazon EC2 instance, and a large-scale Cray machine. (Appendix J, target ATDM-1)

Planned Activities in FY19:

- Coordinate with ATDM DevOps and Sandia HPC support teams to develop a methodology for utilizing containers for testing and deploying mission applications and demonstrate at scale. (Appendix J, target ATDM-1)
- Improve exploitation of node-level task-Directed Acyclic Graph (DAG) parallelism in Kokkos using the Qthreads backend on ASC platforms and advanced architecture testbeds. (Appendix J, target ATDM-1)
- Examine the relationship between MPI resource usage and application scalability. (Appendix J, target ATDM-1)
- Optimize GPU execution performance of contact algorithms in the NimbleSM application. (Appendix J, target ATDM-1)
- Integrate a face-face contact algorithm into the NimbleSM application, representative of the primary contact algorithm used in SIERRA/SM. (Appendix J, target ATDM-1)
- Add frictional contact capability to NimbleSM architected for CPU or GPU use, and compare performance for ATS-2 hardware. (Appendix J, target ATDM-1)

Arm Vanguard Advanced Prototype System Deployment (SNL)

The scope of this project is to lead the design, acquisition, and plan for deployment of the Vanguard project's Arm-based Advanced Prototype System (APS), named Astra. 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 Arm-based APS is to field a large-scale prototype of an Arm-based system that is optimized for nuclear weapons (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 FY18:

- Completed vendor surveys and released the formal RFP for the system. (Appendix J, target ATDM-2)
- Held a tri-lab technical proposal review and source selection. (Appendix J, target ATDM-2)
- Completed vendor negotiations and formal Statement of Work (SOW). (Appendix J, target ATDM-2)
- Executed a contract for the system with delivery dates and milestone payment schedule. (Appendix J, target ATDM-2)

Planned Activities in FY19:

- Complete delivery, installation, and acceptance of Astra system hardware. (Appendix J, target ATDM-2)
- Complete initial setup and make the system available on the open network for science application porting and testing. (Appendix J, target ATDM-2)
- Complete needed security approvals and transition Astra to the restricted network. (Appendix J, target ATDM-2)
- Initiate porting and testing of ASC application codes up to International Traffic in Arms Regulations (ITAR) restricted levels. (Appendix J, target ATDM-2)

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-based Vanguard Phase 1 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 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 FY18:

- (LANL) Initially tested CharlieCloud containers on an Arm system. Initial benchmarking of PARTISN on Arm. Initial evaluation of LLVM on Arm. (Appendix J, target ATDM-2)
- (SNL) Completed initial release of ATSE definition document. (Appendix J, target ATDM-2)
- (SNL) Completed porting and testing of OpenHPC to Commanche testbed. (Appendix J, target ATDM-2)
- (SNL) Released ATSE Version 1.0 (Appendix J, target ATDM-2)

Planned Activities in FY19:

- (LANL) Test CharlieCloud containers on a full application on an Arm system. LLVM development targeting Arm. (Appendix J, target ATDM-2)
- (SNL) Continue development of ATSE, leading to Version 2.0 release. (Appendix J, target ATDM-2)
- (SNL) Test current ATSE on Astra and optimize performance for open science apps on open network. (Appendix J, target ATDM-2)
- (SNL) Initiate testing ATSE on Astra and optimizing performance for integrated mission applications on restricted network. (Appendix J, target ATDM-2)

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. 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. Algorithms and methods developed in the partnership will advance DOE capabilities for predictive modeling in both biomedical and DOE Program applications.

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 FY18:

- Demonstrated the first implementation of our true multiscale code connecting continuum to coarse-grain molecular dynamics (CGMD) models.
- First use of all 3,540 nodes of the new Sierra machine (April 2018), spawning 14,000 individual molecular dynamics (MD) simulations.
- Manuscript accepted for talk, "Distinguishing between Normal and Cancer Cells Using Autoencoder Node Saliency," to the ISC HPC 2018 Workshop. Describes a novel method called autoencoder node saliency, which presents a semi-supervised

learning method for identifying nodes in an autoencoder, which highlight important genes associated with distinct cancer types.

- Developed a novel semi-supervised feature learning method called center-loss/center-distance. Results were presented at "The 22nd Annual Signal & Image Sciences Workshop" in Livermore, showing how the method reduced a 17,000-gene feature set down to 1000 features, and was used to improve prediction performance.

Planned Activities in FY19:

- Use the newly optimized capability to perform millisecond simulations of RAS on the lipid bilayer and dissect the relationship between the lipid bilayer and the RAS protein and having the system validated by new experiments.
- Extend the model to allow introduction of new protein(s).
    - Develop "bootstrap" algorithm for model development.
    - Understand the interaction between the two different proteins and the lipid bilayer.
- Develop novel machine learning algorithms to explore multidimensional space and more efficiently guide the direction of the overall multiscale simulation on the fly.
- Develop ML techniques to optimize entire simulation/experimental campaign around RAS on membranes (e.g., "ML on the outside").
- Begin extension of code to incorporate a third resolution level (fully atomistic scale) in the future, with attendant ML machinery to initiate multilevel simulations.
- Refine existing deep learning models for transfer learning to apply the existing cell line-based models to make predictions on patient-derived xenograft (PDX) data.
- Make prediction models available for broader use in the cancer and materials research communities.

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 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 FY18:

- Trained drug response models using two types of models (random forest and Bayesian multitask multikernel) for each of the five cell lines data sets (NCI-60, CTRP, GDSC, CCLE, gCSI). The models were used to: 1) estimate intraset performance and interset (cross-comparisons or transfer learning) performance, and 2) predict drug responses for PDX models. (Appendix J, target ATDM-3)

- Devised two novel approaches to estimate effective drug combination therapies. One approach is based on models that predict the likelihood of a pair of drugs to be synergistic. Another approach uses a mechanism-based mathematical modeling strategy for the discovery of effective drug combinations to achieve durable inhibition of ERK signaling in RAF and/or RAS-driven cancers while avoiding (or significantly delaying) signaling recovery and resistance. (Appendix J, target ATDM-3)

Planned Activities in FY19:

- Complete efforts to quantify the uncertainty of the predictions made by the drug response models. (Appendix J, target ATDM-3)
- Generalize the mechanism-based mathematical model to include cell line specific information. (Appendix J, target ATDM-3)

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 (UCSF) 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 FY19 Planned Activities section below.

Accomplishments in FY18:

- Completion of four-party CRADA between LLNL, Frederick National Lab for Cancer Research, UCSF, and GSK to govern partnership R&D.
- Approval of R&D plan for joint research on predictive models for molecular interactions in drug design applications.
- Completion of baseline computational models for drug safety and pharmaceutical performance using LLNL HPC systems.
- First demonstration of hybrid predictive model capability that integrates data-driven machine learning models with mechanistic simulation models. Limited training set sizes limit the performance of purely data-driven modeling approaches. The addition of simulated features to the data sets can enhance the prediction performance of these models.

Planned activities in FY19:

- Optimize and characterize performance of hybrid prediction models integrating machine learning and simulation models for molecular interactions.
- Develop new molecular representation methods and generative modeling algorithm to support drug development and material design applications.
- Develop and demonstrate an integrated active learning approach that automatically specifies priority experimental measurements needed to improve model prediction performance and minimize uncertainty.

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 2017, and through quarter 3, fiscal year 2018, are reflected below for the CSSE subprogram.

- Accepted the Sierra Initial Delivery system and started Gordon Bell and early science runs. (LLNL)
- Developed the application development environment and documentation for Sierra users. (LLNL)
- Released Request for Information (RFI) for LC Exascale Storage Series, which will position LC's storage capability for exascale. (LLNL)
- Procured and deployed Ulna, a Cavium ThunderX2 Arm testbed, with Arm support from Tri-Lab Operating System Stack (TOSS). (LLNL)
- Ran two Advanced Technology Computing Campaign (ATCC) campaigns on Sequoia and supported application teams. (LLNL)
- Implemented anomaly detection and other analytics for the Production DAC, the systems logging data analytics cluster for the open environment. (LANL) (Appendix J, target CSSE-2)
- Advanced fast neutron detector work: captured data, analyzed it, developed a machine learning algorithm to differentiate neutrons from gammas. (LANL) (Appendix J, target CSSE-1)
- Installed HPE Arm testbed for application porting, toolchain, and performance investigations. (LANL) (Appendix J, target ATDM-2)

- ARTeam tuned and performed troubleshooting of the High Performance Input/Output (HIO) burst-buffer library in support of the large-scale 3D simulation runs on Trinity KNL. (LANL) (Appendix J, target CSSE-2)
- Initial containerized in situ workflow shown with LAP and CharlieCloud containers. (LANL) (Appendix J, targets CSSE-2 and IC-3)
- Deployed early Power9/Volta GPU nodes to support advanced application work for ATS-2/Sierra. (SNL) (Appendix J, target CSSE-1)
- Completed memory subsystem and interconnect analysis supporting technical decisions for exascale systems procurements. (SNL) (Appendix J, target CSSE-1)
- Resilient on-node AMT programming model design, implementation, and analysis. (SNL)
- Developed a new video-analysis capability (VideoSwarm) for the Slycat ensemble analysis software, driven by the needs of NW analysts. (SNL)
- Developed Fugu, a C++ platform for linking conventional programming paradigms with neuromorphic algorithm implementations and compilers for neuromorphic hardware. (SNL) (Appendix J, target CSSE-3)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Development Environment for Test Bed Systems		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 8/31/19		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, targets CSSE-1.b, ATDM-2.a		
Description: This effort will build a portable development environment that can be quickly deployed to test bed systems, enabling users to evaluate new architectures with familiar tools and compilers. After deploying this new portable development environment across several test beds, it (and the test beds) will be evaluated by building and running a code of interest to ASC. This effort is necessary because LLNL's existing development environment, Tri-Lab Computing Environment (TCE), is well-liked and works well on x86-based systems but takes significant targeted effort to port to new systems, such as Sierra. It has not been worth the effort to repeat the large TCE porting exercise for each test bed. However, Spack is now enabling us to build and manage large software ecosystems with ease. Instead of relying on a team of lab personnel to port and build software for each test bed system type, we can leverage the same work already being done in the much larger Spack open-source community. Specifically, this effort will work with the Spack community to add additional features and packages needed to recreate TCE. We will use Spack to deploy this new development environment onto test beds that will at least include CaviumX2 processors and AMD Firepro GPUs. Finally, we will work with a member of the ASC community to build and run a code of interest to ASC on the test beds and evaluate both the test bed and development environment.		
Completion Criteria: A portable development environment will be deployed on multiple test bed systems, and an evaluation will have been performed.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: LLNL CSSE and ATDM/ASD staff.		

Milestone (ID#TBD): Container Environments for Production Applications with Complex Workflows		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 9/30/19		
ASC WBS Subprogram: CSSE, ATDM		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target CSSE-2.a		
Description: Demonstrate the capability to run production codes with complex workflows including setup and in situ processing on production LANL systems from a containerized environment. This effort involves moving CSSE research software to production systems at LANL. In particular, key elements of the CharlieCloud container runtime and BEE container orchestration software will be made available on production HPC systems at LANL. This will be demonstrated with a minimum of one production code from IC or ATDM including problem setup and in situ processing of the results all from a common container environment.		
Completion Criteria: Demonstration of production quality run on an AT or CT production LANL HPC system: • Compile and build an IC simulation code in container with all dependencies • Demonstrate a simulation run with preprocessing • Demonstrate a simulation run with in situ processing of results • Demonstration checkpoint and restart from container Stretch goals: • Show a containerized workflow portable between ATS and CTS • Show containerized workflow for multiple IC and ATDM applications		
Customer: IC/ATDM		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: IC/FOUS		

Milestone (ID#TBD): Deploy SST Capability to Simulate Node Architectures with CPUs and GPUs		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 09/30/19		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target CSSE-3		
Description: This work is to complete and deploy capability to analyze heterogeneous node architectures consisting of x86 CPUs and relevant GPUs to support co-design and performance analysis of ASC computing platforms running ASC workloads. This capability will be deployed through the Structural Simulation Toolkit (SST).		
Completion Criteria: Publish a report documenting use of SST with GPU modeling capability on an ASC problem of interest. Also, update SST documentation for use of this new capability.		
Customer: ASC, NNSA		
Milestone Certification Method: A program review is conducted and its results are documented. A report is prepared as a record of milestone completion.		
Supporting Resources: FOUS (production cycles to run SST analyses).		

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 FY18:

- Extensively tested and resolved issues on production systems.
- Assisted users in utilizing and transitioning to CT platforms.
- Led and managed the CT systems contract and FY18 system procurements.
- Completed CTS-1+ architecture update study and incorporated two new architectures into the contract.
- Started CTS-2 market survey and procurement planning activities. (Appendix J, CSSE-1.c)

Planned Activities in FY19:

- Lead and manage CT systems contract (CTS-1) and any FY19 CT system procurements.
- Complete CTS-2 market survey and prepare draft SOW for public release. (Appendix J, CSSE-1.c)

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 FY18:

- Stabilized the CTS-1 systems to improve system availability.
- Continued operational support for CTS system.

Planned Activities in FY19:

- Receive, integrate, and accept a new ASC-funded CTS-1 system: Cyclone.
- Provide operational support for CTS systems.

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 FY18:

- Accepted a 6-Scalable-Unit (SU) CTS-1 Eclipse cluster.
- Procured and received a 2-SU expansion for the Eclipse cluster.
- Completed final payments on lease agreement for Cayenne and Serrano.

Planned Activities in FY19:

- Integrate and accept a 2-SU Eclipse expansion.

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-petaFLOPS 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-petaFLOPS system has a staggering 1.6 million processor cores with a total possible 102 million hardware threads all operating simultaneously. This type of parallelism dictated new directions in supercomputing. 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 ATCC process, which is run every six months when the next suite of codes is ushered onto the machine.

Accomplishments in FY18:

- Ran two ATCC processes.
- Continued to investigate optimal performance tuning for specific codes. (Appendix J, target IC-1)

Planned Activities in FY19:

- Run an ATCC process.
- Retire the Sequoia system.

Sierra Tri-Lab Advanced Technology System (LLNL)

In November 2014, LLNL signed a contract with IBM to begin to deliver Sierra, a next-generation supercomputer, in 2017 with acceptance in 2018. Under the CORAL procurement, LLNL has worked with IBM, NVIDIA, and Mellanox to deploy a Sierra system of over 125 petaFLOPS, with the bulk of the capability in NVIDIA V100 (Volta) GPUs 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. Codes that offload the bulk of their computation to the GPUs will run best on this machine. Once this platform has transitioned to the classified environment, it will be used as a tri-lab resource for tri-lab stockpile stewardship milestones via the ATCC process, which will be run every six months when the next suite of codes is ushered onto the machine.

Accomplishments in FY18:

- Provided technical coordination and contractual management for Sierra NRE and build contracts.
- Continued application preparations for the Sierra system through the Sierra COE. (Appendix J, target IC-1)
- Deployed and accepted the initial Sierra Initial Delivery system. (Appendix J, target FOUS-1)
- Began science runs on the Sierra Initial Delivery system that included calculations of interest to the ASC Program and two submissions to the Gordon Bell Award. (Appendix J, target FOUS-1)
- Deployed the full Sierra system and began its acceptance. (Appendix J, target FOUS-1)
- Established the ATCC process for FY19 Sierra full-system access. (Appendix J, target IC-1)
- Participated in NRE and build subcontract activities for ANL's A21 system as part of the broader CORAL collaboration activities. (Appendix J, target CSSE-1)

Planned Activities in FY19:

- Provide technical coordination and contractual management for Sierra NRE and build contracts.
- Continue application preparations for the Sierra system through the Sierra COE. (Appendix J, target IC-1)
- Complete science runs on the full Sierra system. (Appendix J, target FOUS-1)
- Transition the full Sierra system to the classified environment. (Appendix J, target FOUS-1)
- Begin the ATCC process for FY19 Sierra full-system access. (Appendix J, target IC-1)

National Nuclear Security Administration

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 FY19, this project will focus on the project management of the ASC Trinity and acquisition of the Crossroads systems. The focus in this project also includes the execution of Crossroads under the Defense Programs' Program Execution Guideline structure.

Accomplishments in FY18:

- Developed 10-year projections for power and cooling of projected LANL HPC systems for site-wide planning of utility power and water.

Planned Activities in FY19:

- Continue to provide program and project management for computing platforms, including requirements gathering and analysis. (Appendix J, target CSSE-1.a)

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 LANL and SNL.

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 FY18:

- Completed the Trinity Production Readiness Level 2 milestone.
- Integrated Trinity Haswell (TR1) and TR2 partitions on the LANL classified network.

---

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

- Released the integrated Trinity system for general availability and ATCC-4.
- Closed out the Trinity Project and completed CD-4.
- Supported the Trinity IPT.

Planned Activities in FY19:

- Continue ATCC support.

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 2022–2026 timeframe. The primary activity is to lead the design, acquisition, and deployment of the third AT system (ATS-3) in the ASC 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.

Los Alamos and Sandia are continuing the ACES partnership to acquire an HPC system in the 2021 timeframe (to be sited at LANL).

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 FY18:

- Reevaluated the mission needs for Crossroads to pursue a new procurement. (Appendix J, target CSSE-1.a)
- Updated the technical specifications for an updated RFP given changes in the HPC market. (Appendix J, target CSSE-1.a)
- Transitioned the project planning for Crossroads from a DOE 413.3b-based plan to a Defense Program Acquisition and Management plan. (Appendix J, target CSSE-1.a)
- Released a revised RFP for Crossroads for delivery of a system in FY21. (Appendix J, target CSSE-1.a)

Planned Activities in FY19:

- Conduct an Independent Cost Review of the negotiated system. (Appendix J, target CSSE-1.a)
- Develop the Acquisition and Management Plan and pursue concurrence in early FY19. (Appendix J, target CSSE-1.a)
- Complete the formal contract for the Crossroads system. (Appendix J, target CSSE-1.a)

Arm Processor Feature Acceleration (LANL)

The objective of this effort is to facilitate the R&D necessary to accelerate key functionalities in an HPC Arm-based processor for availability in potential future AT and CT systems that will enable increased efficiencies of processor cycles for NW Mission workloads. The scope of this work includes the contracting of R&D efforts along with reviews of progress and milestones associated with the acceleration of these features for final evaluation in the 2020 timeframe.

Accomplishments in FY18:

- Negotiated deliverables for the various R&D efforts necessary to accelerate key features in an HPC Arm processor.
- Completed contract negotiations and award.

Planned Activities in FY19:

- Continue to monitor progress of the various engineering deliverables and refine as appropriate.
- Evaluate quarterly milestone progress.

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 FY18:

- Led ECP/Hardware Evaluation and support of the ECP/Pathforward projects. Many of these activities are significantly influencing the ATS-3 and ATS-4 procurement activities. (Appendix J, target CSSE-1)
- Collaborated and led a multiagency exploration activity for innovative/semicustom architectures optimized for our mission workloads.

Planned Activities in FY19:

- Support the completion of a strategic plan for the multiagency collaboration on innovative/semi-custom architectures optimized for our mission workloads.
- Deliver technical analysis supporting the ATS-3 and ATS-4 procurement activities including NRE. (Appendix J, target CSSE-1)

Advanced Architecture Test Bed Research and Development (SNL)

This project will address a critical need for a range of experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy applications, enable application performance analysis with Mantevo proxy applications, support the Heterogeneous Computing and Programming Model R&D, the Software and Tools for Scalability and Performance projects, and for SST validation efforts. These test bed systems are made available for "test pilot" users who understand the experimental nature of these test beds. 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 FY18:

- Deployed early Power9/Volta GPU nodes to support advanced application work for ATS-2/Sierra.
- Deployed early Cavium ThunderX2 systems from HPE and Cray in support of the Vanguard project.

Planned Activities in FY19:

- Deploy early testbeds of systems that are expected to be options for Crossroads and Vanguard Phase 2. (Appendix J, target CSSE-1)
- Update existing testbeds as available to the latest prerelease hardware from computer vendors. (Appendix J, target CSSE-1)

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 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 (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 FY18:

- Continued development and support for TOSS 2 on legacy systems, and for Red Hat Enterprise Linux 7 (RHEL7)-based TOSS 3 with 3.2 and 3.3 releases. (Appendix J, CSSE-2.c)
- Added aarch64 ARM support to TOSS 3 (Appendix J, ATDM-2.a)
- Updated SLURM to 17.02 in TOSS 3. (Appendix J, CSSE-2.c)
- Worked with RedHat to deploy and mitigate the performance impacts of Spectre/Meltdown security patches.

Planned Activities in FY19:

- Extend TOSS support on non-CTS-1 hardware such as ROCm and AMD CPUs. (Appendix J, CSSE-2.b)
- Integrate Spack into the TOSS build farm for easier deployments on next-generation systems. (Appendix J, ATDM-2.a)
- Continue support for TOSS 2 and TOSS 3 and continue development for TOSS 3. (Appendix J, CSSE-2.c)

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 FY18:

- Developed initial application development environment and documentation for early users and testers of the Sierra system as part of the L2 Milestone 6459 effort. (Appendix J, CSSE-2.b)
- Maintained the Sequoia production application development environment in support of ATCC efforts, and expanded availability of C++ language capabilities through the CLANG compiler.
- Supported CTS-1 users and code developers' efforts with a productive and capable user environment. (Appendix J, target CSSE-2.c)
- Provided technical evaluations for development environment aspects of the CORAL 2 system. (Appendix J, target CSSE-2.b)
- Expanded availability and capability of code development correctness tools through wider release of the LLNL-developed Pruners tools and provided highly scalable tools to users and system administrators for file management through the Distributed I/O Tools set. (Appendix J, target CSSE-2.b)

Planned Activities in FY19:

- Continue code development environment support on all LLNL ASC platforms. (Appendix J, target IC-1)
- Improve, test, evaluate, and provide user support for the production Sierra system applications development environment. (Appendix J, target IC-1)

- Conduct Sierra Applications Preparation activities to assist tri-lab IC teams and their users in porting, tuning, and using the Sierra AT system in production, including introductory training at tri-lab sites. (Appendix J, target IC-1)
- Develop tool infrastructures to improve scalability and performance of applications and the code development environment. (Appendix J, targets CSSE-2.b–c, IC-1)
- Support NRE activities and technical evaluations preparing for the CORAL 2 exascale system procurement. (Appendix J, targets CSSE-1.b, CSSE-2.b)

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 FY18:

- Implemented anomaly detection and other analytics for the Production DAC, the systems logging data analytics cluster for the open environment. (Appendix J, target CSSE-1)
- Applied machine learning to systems logs for unsupervised anomaly detection. (Appendix J, target CSSE-1)
- Advanced fast neutron detector work: captured data, analyzed it, developed a machine learning algorithm to differentiate neutrons from gammas. (Appendix J, target CSSE-1)

Planned Activities in FY19:

- Work on initial development of SpARQ jobs to perform analytics on the Production DAC, the systems logging data analytics cluster for the open and secure environments. (Appendix J, target CSSE-1)
- Develop early warning prediction of job outcome based on syslog. Develop a baseline for job outcome prediction using keyword counts. (Appendix J, target CSSE-1)

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 FY18:

- Obtained early access to Sierra-type nodes for local investigations on porting and performance. (Appendix J, target IC-1)
- Installed an HPE Arm testbed for application porting, toolchain, and performance investigations. (Appendix J, target ATDM-2)

Planned Activities in FY19:

- Acquire and support cutting-edge node technology for ASC application development. (Appendix J, target IC-1)
- Baseline existing testbed nodes with LANL proxy and full applications. (Appendix J, target CSSE-1)

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, 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 runtime system (Qthreads), and high-performance network stack (Portals communication protocol) development with a production HPC computing stack.

Accomplishments in FY18:

- Demonstrated dynamic power management in the SPARC application running on Trinity using the Power API.
- Conducted an evaluation of the Singularity containerization solution running on Trinity and Amazon EC2.

Planned Activities in FY19:

- Lead the Power API community and work with vendors and other labs to deploy implementations.
- Port and deploy a lightweight kernel operating system on a 64-bit Arm processor testbed system in preparation for testing on the Astra platform.
- Demonstrate the use of containers for creating composed applications built up of discrete simulation, analysis, and/or visualization components.

High Performance Computing Hardware Architecture Simulation (SNL)

The SST is a suite of tools enabling multiscale computer architecture simulation to meet the needs of HPC software/hardware co-design. The SST consists of a core set of components that enable parallel discrete-event simulation; high-fidelity networking, memory, and processor components; and coarse-grained simulation components that capture essential elements of machine performance with low computational cost. Future HPC systems and the applications designed to utilize them are impacted by a variety of considerations, including scalability of applications, ease-of-programming, memory and network latencies becoming more imbalanced relative to computation rates, data corruption and its propagation, frequency of interrupts, power consumption, and overall machine cost. SST is designed to allow each of these parameters to be explored, permitting the consideration of a broad space of potential architectural and application/algorithmic designs. The goal is for the SST components to be extended and enhanced by a community of simulator developers, including academic, industrial, and government partners. An even larger community is expected to be the users of SST, including algorithm developers, architecture designers, and procurement team members.

Accomplishments in FY18:

- Completed memory subsystem analysis supporting technical decisions for exascale systems procurements.
- Completed interconnect analysis supporting technical decision for exascale systems procurements.

- Initiated a contract to develop GPU modeling capability for SST.
- Developed the initial model for reconfigurable computing elements, which are potential exascale architectural features.

Planned Activities in FY19:

- Support the L2 milestone deploying initial GPU modeling capability in SST and explore ATS-2 and ATS-4 architectural performance estimates. (Appendix J, target CSSE-1)
- Support exercising SST for novel and beyond-Moore's-Law computing applications. (Appendix J, target CSSE-3)
- Address technical questions for ATS-3 and ATS-4 architectural features and NRE. (Appendix J, target CSSE-1)

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.

Accomplishments in FY18:

- Co-authored a paper describing a portable programming environment for offloading packet processing, which was nominated for Best Paper at the SC17 conference.
- Developed several MPI optimizations for message matching specific to many-core architectures and integrated into the open source Open MPI implementation.

Planned Activities in FY19:

- Implement and demonstrate a new prototype software notification mechanism for one-sided communication that improves performance and provides resilience.
- Continue to advance Portals technology to meet future-generation networking requirements, including guiding the next major generation of vendor interconnects.

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 FY18:

- Completed studies of fault characterization on recent ASC Cielo production capability system.
- Implemented physics-based checksums in a conjugate gradient solver using 3D unstructured mesh, showing the ability to detect injected errors with lower overhead.
- Demonstrated task-level error detection for stencils, and guided techniques for resilient AMT implementation, simulation, and analysis.
- Performed resilient on-node AMT programming model design, implementation, and analysis.

Planned Activities in FY19:

- Carry out a systematic study of bit-patterns using a recently purchased proprietary DRAM failure injection card to determine the likelihood of silent data corruption for SEC-DED and ChipKill ECC schemes.
- Utilize initial Trinity failure data to reevaluate the overheads of current failure mitigation techniques.
- Perform initial integration of error-detection capabilities as part of scalable resilience mechanisms applicable to one or more production ASC codes.
- Extend the resilient on-node AMT prototype for enabling task-based message replay in the MPI+AMT model.
- Explore the latest MPI standard and its future direction for designing resilient MPI+AMT integration.

AMT Programming Models and Runtimes (SNL)

The DARMA project is exploring AMT programming and execution model abstractions designed to isolate the applications programming layer from the AMT runtime layer. This project is transitioning from an ATDM project to a CSSE project starting in FY19. The accomplishments shown for FY18 are from the ATDM activity.

Accomplishments in FY18:

- Presented an ECP tutorial at the February all-hands meeting, maintained interactive on-line tutorial.
- Gave overview presentations at Supercomputing, ECP all-hands, PMCD, and HiHat monthly meetings.
- Provided abstractions for MPI/OpenMP/Kokkos interoperability.
- Developed an active message layer based on MPI+OpenMP (rather than Charm++) for improved interoperability and ability to leverage vendor implementations.
- Developed experimental front-end abstractions for fine-grained concurrency (subsetting) and termination detection.
- Integrated the existing PIC mini-app with HPCCG solver for debugging and performance tuning of MPI interoperability.

Planned Activities in FY19:

- Execute performance tuning and debugging of MPI interoperability on AT-class machines (e.g., Trinity), leveraging performance analysis tools.
- Integrate MPI-interoperable DARMA abstractions with EMPIRE and demonstrate on testbed-scale platforms (e.g., Mutrino).
- Develop and evaluate load balancing heuristics and algorithms (theoretical performance bounds and experimental results) for scales up to AT-class platforms.

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 (LLNL, SNL, and LANL), ORNL, LBNL, and IBM. LLNL provides development, deployment and support of archival storage software for tri-lab ASC customers on unclassified and classified networks. It includes the selection, procurement, maintenance of archival storage software/hardware/media, ongoing technology refresh, and data stewardship. Locally developed HPSS software provides scalable, parallel, archival storage interfaces and services to the tri-labs. A diverse array of hardware is integrated beneath HPSS supplying the performance necessary to offload data from ASC platforms, facilitating computation. This includes disk arrays, robotic tape subsystems, servers, networks, and hundreds of petabytes of tape media, all of which contribute to enable high-speed parallel transfers into an efficiently scaled out data store.

Accomplishments in FY18:

- Continued ongoing HPSS software development: HPSS 7.5.1p1 and HPSS 7.5.1p2 (patch releases); HPSS 7.5.2 (major feature release). HPSS Rumbler Change Request Proposal was defined as part of future LC I/O Strategy for coupling archive with parallel file systems. (Appendix J, target CSSE-2.c)
- Deployed TOSS3 HPSS Core Server platforms and metadata disk to refresh old systems in both unclassified and classified networks. (Appendix J, target CSSE-1.c)

- Began evaluating storage ecosystems to architect a future archive using technology that will cost-effectively satisfy capacity and bandwidth. (Appendix J, target CSSE-1.b)
- Kicked off the LC Exascale Storage Series, which surveyed vendors on the state of large-scale storage; began evaluating RFI responses. (Appendix J, target CSSE-1.b)
- Upgraded OCF Disk/Tape Mover Nodes to TOSS3. (Appendix J, target CSSE-2.c)
- Provided support for deployed archival storage systems. (Appendix J, target CSSE-2.c)

Planned Activities in FY19:

- Develop and release HPSS 7.5.3. (Appendix J, target CSSE-2.c)
- Continue engaging with LC I/O Strategy: coupling archive with parallel file systems, exploring HPSS Disk Cache becoming LC's "warm" data level, and analyzing HSM interfaces. (Appendix J, target CSSE-2.c)
- Research exascale-capable archiving software. Review DOE-SC Distributed Archival Storage System RFI results and market surveys. (Appendix J, target CSSE-2.c)
- Upgrade Secure Computing Facility (SCF) Disk and Tape Mover Nodes to TOSS3. (Appendix J, target CSSE-2.c)
- Investigate deploying HPSS 7.5.2/3 to the unclassified network. (Appendix J, target CSSE-2.c)
- Deploy OCF next-gen tape infrastructure (library, drives, media). Exercise options on FY18 RFP option for SCF next-gen tape procurement. (Appendix J, target FOUS-5)
- Continue ongoing support for archival storage systems. (Appendix J, target FOUS-5)

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 FY18:

- Enhanced Zettabyte File System (ZFS)-based Lustre performance. (Appendix J, target CSSE-2.c)
- Supported LLNL production Lustre deployments. (Appendix J, target CSSE-2.c)
- Tested, improved, and packaged Lustre 2.10, a long-term support release, for TOSS 3 in preparation for first production deployment. (Appendix J, target CSSE-2.c)
- Evaluated Starfish metadata indexing solutions, including to facilitate purge policy enforcement and data provenance management. (Appendix J, target CSSE-2.c)
- Actively participated in CORAL 2 working groups and technical evaluation teams in preparation for El Capitan. (Appendix J, target CSSE-2.b)
- Completed the Project Quota feature for ZFS-based Lustre to improve management of space accounting. (Appendix J, target CSSE-2.c)

Planned Activities in FY19:

- Continue to make general performance and management improvements to ZFS-based Lustre. (Appendix J, target CSSE-2.c)
- Evaluate the use of Progressive File Layouts to improve data distribution among servers and performance of user codes. (Appendix J, target CSSE-2.c)
- Improve Lustre utilities, removing obsolete functionality and improving user interfaces. (Appendix J, target CSSE-2.c)
- Collaborate to finalize the ZFS Distributed parity RAID (DRAID) feature for next-generation Lustre Just a Bunch of Disk (JBOD) systems. (Appendix J, target CSSE-2.c)

Networking and Test Beds (LLNL)

The Networking and Test Beds project provides research, performance testing, capability testing, and analysis for new processors, file systems, networks, 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 FY18:

- Deployed Ulna and added TOSS 3 support for its ThunderX2 Arm processors. (Appendix J, target CSSE-2.c)

- Added AMD ROCm and AMD hardware support in TOSS. (Appendix J, target CSSE-2.c)
- Deployed AMD Epyc and Intel Skylake for limited user testing. (Appendix J, target CSSE-2.c)
- Tested Lustre 2.10 in preparation for production. (Appendix J, target CSSE-2.c)

Planned Activities in FY19:

- Add ROCm TOSS support for Arm64. (Appendix J, target CSSE-2.c)
- Test future software and new interconnect hardware from Cray for integration into TOSS. (Appendix J, target CSSE-2.c)
- Deploy the AMD Epyc system with AMD GPUs to continue to improve the AMD ecosystem. (Appendix J, target CSSE-2.c)
- Deploy production ThunderX2 hardware in support of Vanguard. (Appendix J, target ATDM-2.a)

File Systems, Archival Storage, and Networking (LANL)

Capabilities of the Archival and File Systems components of the project include online file systems such as the NFS complex and enterprise-wide supercomputer file systems, global parallel file system (GPFS) development, deployment and management, scalable I/O (SIO) middleware development and support, storage area network (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 (ARTeam) capabilities are consolidated in this project, addressing issues with HPC customers' applications production-run readiness on current and incoming computing systems at the tri-labs. Working with subsystem teams such as systems management, file systems and I/O, archive, and tools, the Application Readiness team identifies causes of unexpected behavior and deploys fixes in production. One project goal is that system users are able to make productive use of the systems with their applications to solve their problems. Another goal of the project is to analyze the performance of customers' applications both to improve performance on current production platforms, as well as to make predictions about the performance of these applications on future HPC platforms.

Accomplishments in FY18:

- ARTeam successfully supported the large-scale 3D simulation runs on Trinity KNL. (Appendix J, target IC-1)
- ARTeam tuned and performed troubleshooting of the HIO burst-buffer library in support of the large-scale 3D simulation runs on Trinity KNL. (Appendix J, targets IC-1 and CSSE-2)

Planned Activities in FY19:

- Continue development of MarFS with extensions for global indexing of data, ensuring LANL's users can efficiently move data into and out of MarFS for storage terms longer than one campaign. (Appendix J, targets IC-1 and CSSE-2)
- Deploy DAC, the systems logging data analytics cluster for the secure environment. (Appendix J, target CSSE-1)
- ARTeam will provide systems support for application teams porting to Sierra. (Appendix J, target IC-1)

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 FY18:

- Supported development and completion of HPSS Version 7.5.1.
- Supported numerous bug fixes to HPSS production releases.
- Supported development for HPSS Version 7.5.2.
- Placed Redundant Array of Inexpensive Tape (RAIT) into production.

Planned Activities in FY19:

- Complete development and test for HPSS Version 7.5.2.
- Continue coding and maintenance support of current versions.

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 multilaboratory 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 FY18:

- Completed a technical report assessing the production I/O needs of the ASC Integrated Codes program at Sandia.
- Published and presented work describing requirements for metadata management to support application-workflow productivity.

Planned Activities in FY19:

- 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 FY18:

- 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 consistent and stable data analysis and visualization environment across LC platforms. (Appendix J, target CSSE-2.c)
- Provided operational support for visualization theaters, including demonstrations for high-level visits, reviews, and tours.
- Implemented and delivered visual analytics capabilities for simulation data and machine learning features using web technologies, including D3.js, through the Lorenz MyLC framework. (Appendix J, targets IC-2, IC-3)

Planned Activities in FY19:

- Maintain the data analysis and visualization environment across LC platforms and networks. (Appendix J, target CSSE-2.c)
- Support simulation users with data analysis and visualization activities including consulting and the creation of images, movies, and software.
- Provide operational support for visualization theaters and provide demonstration, tour, and special event support.
- Exploit research results in data analysis and visualization for ASC simulations, including topological methods and machine learning algorithms. (Appendix J, target IC-3)

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 FY18:

- Released new versions of Hopper, which can restore previous sessions, prevent duplicate copy jobs, provide additional support for working with sites that have restrictive gateway proxy systems, detect which features are most used, and contains general usability enhancements. (Appendix J, target CSSE-2.c)
- Developed a Task Management Interface within Lorenz for enabling Hotline staff to perform system tasks securely and in a controlled manner. Implemented a Lorenz regression testing portlet and data integrity checker. Completed the Hotpad 2.0

release for managing RSA tokens, including conversion to a REST-based middleware service. Completed numerous Lorenz infrastructure and security improvements. (Appendix J, target CSSE-2.c)

Planned Activities in FY19:

- Release new versions of Hopper and Chopper, which include expanded file transfer wizard capabilities, smart favorites for helping to navigate connections, site file verification, and general usability enhancements. (Appendix J, target CSSE-2.c)
- Extend the Lorenz job management capabilities to include Sierra-class systems. Release version 3 of the Task Manager Interface, featuring optimizations and additional tasks. Deploy a web application that facilitates the LC operations team's knowledge exchange between and during operational shifts. Continue enhancing the MyLC Dashboard with user- and staff-requested features. (Appendix J, targets CSSE-1.b, CSSE-2.c)

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 FY18:

- Supported visualization of IC's 3D simulations with Ensight and Paraview. (Appendix J, targets IC-3 and CSSE-2)

- Showed initial containerized in situ workflow with LAP and CharlieCloud containers. (Appendix J, targets IC-3 and CSSE-2)

Planned Activities in FY19:

- Show containerized in situ workflow with LAP, Eulerian Application Project (EAP) and CharlieCloud containers. (Appendix J, targets IC-3 and CSSE-2)
- Continue support for in situ for large-scale 3D simulations on LAP and EAP. (Appendix J, targets IC-3 and CSSE-2)

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.

Accomplishments in FY18:

- Released ParaView 5.5.1 with several new features for ASC codes, including updates to OSPRay ray-traced rendering (benefiting Trinity users) and large file support for Computational fluid dynamics General Notation System (CGNS) (needed by ATDM and IC).
- Developed a new video-analysis capability (VideoSwarm) for the Slycat ensemble analysis software, driven by the needs of NW analysts.
- Completed production releases of Slycat, ParaView, and Catalyst for the Sandia Analysis Workbench (SAW).
- Developed and improved in situ visualization capabilities for SPARC, SPARTA, Nalu, and other ASC codes.

Planned Activities in FY19:

- 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.
- Integrate compression algorithms into our in situ tools.

- Continue integration and maintenance of analysis and visualization tools with the SAW, focusing on support for the FY19 IWF Level 2 milestone.
- Enhance and improve ensemble analysis and visualization tools on ASC platforms. The primary thrusts will be expanded use cases for Slycat and VideoSwarm, and development of ML-inspired clustering capabilities.
- 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 FY18:

- Held a BML Neuromorphic tri-lab meeting, where all labs presented progress and areas for continued collaboration. (Appendix J, target CSSE-3.a)
- Created a high-level Python interface to create matlab corelets. This interface and corelets are available to LANL and SNL collaborators. (Appendix J, target CSSE-3.a)
- Created 3200 Quadratic Unconstrained Binary Optimization (QUBO) solutions of 320 graphs on TrueNorth to evaluate the effects of low precision and quantizing requirements. Evaluated accuracy and an approximation of energy usage, including the effects of the stochastic leak. Presented initial results at SIAM PP. (Appendix J, target CSSE-3.a)
- Continued discussions with IBM for their next generation of neuromorphic simulators and hardware, which is not yet available. Started discussions with Intel about gaining access to their Loihi neuromorphic chips for the evaluation of QUBO solutions. (Appendix J, target CSSE-3.b)

Planned Activities in FY19:

- Increase the scalability of QUBO solutions on TrueNorth beyond 64-node graphs. (Appendix J, target CSSE-3.a)

- Develop a tool to import neural networks trained in common machine learning tools, such as TensorFlow or PyTorch, into TrueNorth. (Appendix J, target CSSE-3.a)
- Evaluate Sandia's Fugu tool chain for abstraction between applications and hardware, allowing for greater exploration of neuromorphic hardware. (Appendix J, target CSSE-3.a)

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 simply the programming of these emerging computers. Second, the project looks to explore computational methods that can aid in basic algorithms of interest to ASC. LANL is targeting both Quantum Annealing and Quantum Gate approaches. LANL is creating collaborations with university and commercial 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. Other BML efforts include custom hardware application specific integrated circuits (ASICs) and Approximate Computing.

Accomplishments in FY18:

- An automated quantum algorithm discovery applied to a quantum Fourier transform showed a 43% reduction in gate count. (Appendix J, target CSSE-3)
- Initially implemented a directed algorithmic approach for solving QUBO on the IBM neuromorphic TrueNorth system at LLNL. QUBO problems can currently be solved on the D-Wave, allowing comparisons of the two architectures. (Appendix J, target CSSE-3)

Planned Activities in FY19:

- Conduct a quantum computing summer school, to be held at LANL over 10 weeks, and including speakers from D-Wave, IBM, Rigetti, MIT, Microsoft, and IonQ. (Appendix J, target CSSE-3)
- Further maturation of initial design of matrix multiplication on neuromorphic hardware including IBM TrueNorth and Intel Liohi. (Appendix J, target CSSE-3)

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 a digital or analog computer system 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 are: 1) enabling ongoing 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 FY18:

- Formulated and implemented a hybrid, continuous-discrete optimization algorithm for sparse basis selection as a QUBO problem, which selects basis functions in polynomial surrogate models. (Appendix J, target CSSE-3)
- Designed and characterized a unit cell of a multiplier circuit that executes on the Ising platform at LANL. (Appendix J, target CSSE-3)
- Confirmed that 64-site Special Quasi-random Structures (SQS) exhibit the variety of local chemical environments sufficient to give a realistic description in an infinite random material, (Appendix J, target CSSE-3)
- Implemented on the IBM TrueNorth platform a spiking cross-correlation algorithm. (Appendix J, target CSSE-3)
- Developed Fugu, a C++ platform for linking conventional programming paradigms with neuromorphic algorithm implementations and to compilers for neuromorphic hardware. (Appendix J, target CSSE-3)
- Acquired for SNL's Neural Exploration & Research Laboratory (NERL) a SpiNNaker neuromorphic platform, a prominent neuromorphic platform.
- Completed a beta version code of our Whetstone pipeline for converting Keras/tensorflow-defined deep learning networks into a spiking version compatible with neuromorphic platforms. (Appendix J, target CSSE-3)
- Determined through quantum charge-transport simulations the major electrical characteristics of the "ultimate CMOS," a 10-nm physical gate length device to be used as a point of reference for assessing any beyond-CMOS technology.
- Designed a non-von Neumann implementation of SuperStrider based on a serial sorting network, shown in simulation to be 2.5 times faster than a traditional parallel implementation.

Planned Activities in FY19:

- Continue implementing extension to SST for simulating SNL's space-time data processing unit (STPU) and other neural-inspired, spike-based architectures. (Appendix J, target CSSE-3)
- Continue implementing a universal charge device simulator. (Appendix J, target CSSE-3)
- Further develop new architectural concepts: SuperStrider demonstration on a field-programmable gate array (FPGA), and adiabatic circuits for recycling signal energy.
- Deploy the Fugu neural computing programming interface as a Tri-lab resource and to foster trilab collaboration on neuromorphic platforms. (Appendix J, target CSSE-3)
- Continue development and implementation of Whetstone and combine it with our context-dependent deep learning (DL) results described previously.
- Continue neural computing testbed investigations. (Appendix J, target CSSE-3)
- Continue investigating the effectiveness of Ising for selected discrete optimization applications. (Appendix J, target CSSE-3)

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 2017, and through quarter 3, fiscal year 2018, are reflected below for the FOUS subprogram.

- Completed site preparation for Sierra, accepted the first 25% of Sierra compute nodes, and installed and integrated all 4320 nodes. (LLNL)
- Continued to provide 24 × 7 support for Sierra early delivery, CTS-1, Tri-Lab Linux Capacity Cluster (TLCC)-2, BGQ, Lustre, and NAS file systems. (LLNL)
- Integrated and deployed a new 15-petabyte Lustre filesystem on the SRD network. (LLNL)
- Received CD-1 approval for ECFM from NA-1. (LLNL)
- Deployed and supported SLURM 17.02 on all CTS-1 clusters. (LLNL)
- Obtained approval of Critical Decision 2/3 for performance baseline and start of construction of the Exascale Class Computer Cooling Equipment (EC3E) Project. (LANL) (Appendix J, target FOUS-3)
- Provided production support for Trinity critical to the success of a large 3D boundary pushing problem during ATCC-4 and -5. (LANL) (Appendix J, target FOUS-1)
- Operated CTS-1 Fire and Ice systems in full production use peaking with over 95% utilization and 99% system availability. (LANL)
- Demonstrated success of leveraging monitoring infrastructure and analytics to improve system operations. (LANL)
- Produced new data storage policy to provide framework for how to use different tiers of storage. (LANL) (Appendix J, target FOUS-1)

- Deployed the 6-SU CTS-1 Eclipse system and ordered a 2-SU expansion. (SNL)
- Diagnosed, replaced, and repaired on-chip water cooling components for CTS-1 platforms Cayenne and Serrano. (SNL)
- Conducted evaluation of performance impacts to Sandia codes due to mitigation for Spectre/Meltdown system vulnerabilities. (SNL)
- Upgraded Mutrino Application Readiness Testbed System to support code migration and optimization work for Trinity. (SNL)
- Completed early evaluation of 100 Gb/s encryptor functionality and performance. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Large-scale Sierra Run		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 3/31/2019		
ASC WBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-1		
Description: The Sierra system will be deployed on the SCF network, and an early user will successfully run an ASC code on at least 75% of the system.		
Completion Criteria: An early user successfully runs an ASC code on Sierra on the SCF at large scale.		
Customer: ASC Sierra users		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: CSSE, IC, ATDM		

Milestone (ID#TBD): High Performance Tape Archive Infrastructure Deployed to Unclassified Computing Environment

Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC

Completion Date: 08/31/2019

ASC WBS Subprogram: FOUS

Participating Sites: LLNL

Participating Programs/Campaigns: ASC

Program Plan Target: Appendix J, targets FOUS-1, FOUS-5

Description: This milestone represents the first of a two-phase deployment. Phase one deployment (this milestone) occurs in the unclassified HPC environment. This solution and deployment thereof will represent a template to be used again for the classified environment in a near future follow-on procurement (phase two). The current outgoing long-lived tape infrastructure consisting of technology ranging in age from 5 to 15 years old will be completely modernized such that the HPC data archives are well positioned for the exascale era. Phase one is the culmination of the planning, procurement, development, integration, testing, and deployment of a new generation of high-performance archive tape infrastructure in the Livermore Computing unclassified HPC center. This will be brought about by updating tape libraries to units that are expected to be supported for more than a decade and installing tape drives and associated media that represent approximately a doubling of individual native cartridge capacities. Aggregate bandwidth to tape is expected to improve, reaching 30 GB/s and beyond. Efficiencies are anticipated by more tightly coupling overall library footprint with recent and continuing advancements of magnetic tape systems that yield ever-increasing gigabits per square inch. The environment recalibration will yield maintenance and support cost reductions associated with the streamlining of operations with more contemporary hardware.

Completion Criteria: Late FY18 procured unclassified tape infrastructure for LC's HPC high-performance data archives operating in production in the unclassified environment.

Customer: ASC/LLNL

Milestone Certification Method:

Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.

The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.

Supporting Resources: FOUS

Milestone (ID#TBD): Exascale Computing Facility Modernization (ECFM) Design		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 6/30/19		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-4		
Description: The general scope of the project is to modify and upgrade the civil, structural, architectural, mechanical, and electrical systems of Livermore Computing's most capable High Performance Computing Center, Building 453 (B453) and an adjacent approximately 3.5 acres of previously developed infill disturbed area. These modifications and upgrades would allow for continued use of B453 to support the increasing demands of the next generation high-performance computers. These modifications and upgrades would also provide enhanced power and cooling for the current and future ATS computers. The project will provide an additional 40 MW of power and the required cooling to support the load increase.		
Completion Criteria: Receipt of final construction documents, including drawings, specifications, cost estimates, design narrative, reports, and calculations where architects and engineers-of record are required to have a current license in the State of California, and are required to duly seal, sign, and date final documents in accordance with the California Business and Professions Code, sections 5535 through 5538, and sections 6730 through 6749. In addition to the architect or engineer-of-record, an independent engineer or architect of the same discipline (may be of the same firm) is required to check and sign each document to be released for construction.		
Customer: ASC Program		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: FOUS		

Milestone (ID#TBD): Improvements for ASC Data Management		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 8/31/19		
ASC WBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, targets FOUS-1, FOUS-5		
Description: Identify user requirements for managing data across multiple tiers of storage.		
Completion Criteria: Identify and document data usage and recommend management of data based on system capabilities.		
Customer: LANL		
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 ASC Program Office, NNSA FOUS Program Manager, ASC computational system users		

Milestone (ID#TBD): Place "Astra" (Vanguard 1) in Operation and Configured for OHPC and SRN Advanced Prototype Work		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 9/30/19		
ASC nWBS Subprogram: FOUS, ATDM		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-2		
Description: "Astra" is an Arm-based Advanced Architecture Prototype High Performance Computing (HPC) System. Including site preparation, the HPC system team will prepare and deploy Astra and its associated system and application software onto the OHPC network and then onto the SRN. The system will be fully functional within the given security environments.		
Completion Criteria: Astra will be a functional HPC cluster capable of successfully executing the tests, demonstrations, benchmarks, and application codes specified by the ATDM program element.		
Customer: ASC programs at SNL, LANL, and LLNL		
Milestone Certification Method A report 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: Advanced Architecture Testbeds project, Sandia Institutional Computing Program (facilities and infrastructure), CSSE Scalable Systems Software		

Milestone (ID#TBD): Site Preparation, Integration, and Operation of Mini-Sierra ART System		
Level: 2	**Fiscal Year**: FY19	**DOE Area/Campaign**: ASC
Completion Date: 9/30/19		
ASC nWBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-1		
Description: Sandia will complete deployment and production availability of an Application Readiness Testbed (ART) system consisting of hardware and a software stack conforming to the Sierra ATS sited at LLNL. The system will support the porting and testing of Sandia's codes and their applications against this new architecture. Site preparation includes water cooling and power infrastructure. The system will be integrated into the Sandia computing environment, including network, security, and DevOps. The milestone is complete when the system is in production and available to users.		
Completion Criteria: The mini-Sierra ART system is generally available to Sandia code development teams and computational analysts. The system and its environment shall be capable of supporting development and testing of Sandia codes, including nightly regression tests (according to application readiness).		
Customer: ASC program and Sierra ATS platform users		
Milestone Certification Method A report 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: IBM (platform vendor), Sandia Facilities, CSSE, ATDM, IC, Sandia DevOps		

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 FY18:

- Continued FY18 procurement support, contract management, and program planning.
- Supported the Presidential Early Career Award for Scientists (PECASE) awardees.
- Supported SC18, 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).
- Supported biannual PSP meetings.
- Supported detailee assignment.
- Managed the PSAAP II program for the six Academic Alliance Centers.

Planned Activities in FY19:

- Continue FY19 procurement support, contract management, and program planning.
- Support biannual PSP meetings.
- Assist in project management of the six PSAAP II Centers and planning for initiating the next phase of the program (PSAAP III).

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 FY18:

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

- Provided LANL support for HQ.
- Managed the SC18 booth for DOE.

Planned Activities in FY19:

- Support PSAAP II collaboration efforts at LANL and planning for PSAAP III.
- Participate in PI and PSP meetings.
- Provide LANL support for HQ.
- Provide support for the SC19 booth for DOE.

Program Support (SNL)

The Program Support project provides critical coordination and integration activities essential to the success of ASC, with a focus on providing 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. 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 and the exascale initiative (with DOE/SC), is also included in this project.

Accomplishments in FY18:

- Supported the 2018 ASC PI meeting.
- Supported the programmatic needs of the PSAAP II program and participated in planning for PSAAP III program and the PSAAP III Pre-Proposal meeting.
- Supported the ECI, including joint planning and execution with DOE/NNSA and DOE/SC laboratories.
- Supported 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.

Planned Activities in FY19:

- Organize and host an FY19 Predictive Engineering Science Panel (PESP) review.
- Support the ASC PI meeting to be held in 2019.
- Support the programmatic needs of the PSAAP II program and participate in planning for PSAAP III program and review of PSAAP III RFI responses.

- Support the ECI, including joint planning and execution with DOE/NNSA and DOE/SC laboratories.

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 FY18:

- Developed a novel implementation of the multiscale GFEM global-local (GFEMgl) algorithm with MPI-parallel execution, to provide a basis for the large system crash analyses required for assessment of planetary protection requirements.
- Applied the "nonlocal failure" capability currently in SIERRA to a representative failure problem (Sandia Fracture Challenge 2) to evaluate this capability for use in Europa lander crash and fragmentation analyses.

Planned Activities in FY19:

- Extend advanced mesh convergent ductile fracture and failure methods (e.g., phase field), evaluate them on representative problems, and work with JPL colleagues to apply them to full system models.
- Develop GFEMgl for dynamic explosion/crash use cases, including material damage, failure, and fracture. Apply GFEMgl to representative problems and evaluate its effectiveness for full system lander crash and fragmentation analyses.

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 FY18:

- Utilized the SNL SIERRA toolkit and CUBIT software 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 (KCNSC)
- Utilized the SNL Dakota software to support UQ for simulation and modeling at Y-12.
- Utilized the LLNL ALE3D and VisIt software to support multiphysics modeling efforts on rolling and joining operations at Y-12.
- Worked with KCNSC to access Granta:MI remotely from Y-12 to gather material property data from the Materials Properties Development and Standardization (MMPDS) Handbook and studied KCNSC's use of the software to organize material data and its pedigree with an emphasis on exotic materials of interest to Y-12.
- Retired the Penguin Altus 1600SA cluster after 11 years of operation.
- Completed laser scanning of the Y-12 Wet Vacuum System. The 26 scans were converted into a 3D virtual reality model. The model is being scripted to simulate training and procedural scenarios.
- Held a demonstration of FL Simulators Inc. ForkLift Training Simulator at Y-12. Y-12 Training is proceeding with procurement of the ForkLift virtual reality training system.
- Completed testing of importers into our virtual reality application for a number of CAD file formats. The Uranium Processing Facility Process Support Facility has successfully been imported into a virtual reality model that we can walk through and interact with the 3D hardware.
- Approved a Telecommunication Proposal for the Oculus Rift VR headset. The hardware has been purchased and is being tested.

Planned Activities in FY19:

- Continue model development using the SNL SIERRA and Dakota software to support predictive simulation on material properties of additively manufactured materials.
- Continue model development using the LLNL ALE3D, VisIt, and MIDAS software and explore the use of LLNL codes to convert CT voxel data with associated material property data into "as-built" models.
- Investigate the use of UQ methodologies to support image-based metrology of CT voxel data and mesh translations by providing prediction intervals for boundary and defect detection within computed imagery.
- Extend the Sandia molecular dynamics code LAMMPS to perform direct calculation of entropy and free energy to reduce the computational effort to model the thermodynamic and physical properties of materials of interest to Y-12.
- Conduct research on the potential for using augmented reality systems to support Y-12 operations.
- Work with the Uranium Processing Facility (UPF) and Y-12 Operations to develop virtual reality models of the UPF that can be used for readiness reviews.
- Submit Telecommunication Proposals to use virtual and augmented reality systems in the Y-12 training facilities.
- Investigate techniques and software solutions to link design metadata into our virtual reality models. Demonstrate linking and accessing metadata of virtual components in the VR models.

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 FY18:

- Completed initial unclassified deployment of Sierra. (Appendix J, target FOUS-1)
- Completed deployment of the SCF Lustre file system. (Appendix J, target FOUS-5)
- Provided ongoing 24 × 7 support for Sierra early delivery, CTS-1, TLCC-2, BGQ, Lustre, and NAS file systems. (Appendix J, target FOUS-5)
- Deployed new identity management system internally and prepared for retirement of the existing IDM account management tool.
- Configured and deployed a host-based intrusion detection system on select infrastructure and production systems. (Appendix J, target FOUS-5)
- Upgraded OCF core security services to TOSS-3. (Appendix J, target FOUS-5)
- Completed disaster recovery implementation for core security and infrastructure.
- Retired Lustre systems (marzen, stout) and decommissioned Herd, Bigfoot, and GDO.
- Deployed new classified hard disk shredder to significantly reduce disk disposal overhead and mitigate security exposures.

Planned Activities in FY19:

- Complete classified deployment of Sierra and its file system. (Appendix J, target FOUS-1)
- Deploy new OCF and OCF Lustre file systems. (Appendix J, target FOUS-5)

- Provide ongoing 24 × 7 support for Sierra (unclassified and classified), TLCC-2, CTS-1, BGQ, Lustre, and NAS file systems. (Appendix J, target FOUS-1)
- Complete the remaining workflows and the production deployment of the new identity management system.
- Complete TOSS 3 platform upgrade of core security, collaboration tools, and web services in the classified and unclassified environments.
- Retire remaining NetApp Lustre file systems (Lambic, Bock, Pilsner, Porter, Cider) and decommission Cab, Zin, and Vulcan production systems.
- Refresh LC's Web Application Firewall technology.

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 FY18:

- Provided ongoing support through hotline operations, documentation, and training.
- Continued to develop in-house knowledge and expertise in accelerator technologies and their software libraries for providing support to the user community.
- Continued to assist users in migrating from TLCC-2 systems to CTS-1 systems. Assisted users with porting their codes to the Sierra EA systems.
- Developed documentation and training for Sierra. (Appendix J, target FOUS-1)
- Implemented process improvements in accounts management, including migrating to paperless processing of some forms and consolidation of accounts databases.

- Developed a utility for logging user environment and linked library data for jobs run on LC clusters, improving LC's ability to debug user jobs and monitor library usage.

Planned Activities in FY19:

- Continue providing support through hotline operations, documentation, and training with a focus on assisting users porting to Sierra. (Appendix J, target FOUS-1)
- Perform process improvements in accounts management, including migrating to a new Identity Management system for processing user accounts, groups, and banks.
- Develop and deploy a utility to monitor CPU load on production cluster login nodes to identify users who are adversely impacting login node performance.

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 FY18:

- Completed supporting Center network architecture, design, and hardware procurements for Sierra HPC machines. (Appendix J, target FOUS-1)
- Deployed 40G HPC Internet facing firewall cluster and evaluated, procured, and installed new 100G HPC edge router.
- Evaluated and enhanced network monitoring and security compliance capabilities.
- Completed site preparation for Sierra.
- Completed B-654 Air Handling Unit Installation. (Appendix J, target FOUS-5)

Planned Activities in FY19:

- Design and test new HPC network architecture.
- Research and improve network security and compliance capabilities.
- Increase HPC science DMZ bandwidth to 100G.

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 FY18:

- Evaluated suitability of home and project file systems for user workloads and recommended necessary upgrade for FY19.
- Supported performance impact analysis of Spectre and Meltdown security patches.
- Provided production support for Trinity campaign storage, NFS, Trinity scratch, and archival storage critical to success of large 3D boundary pushing problem during ATCC-4 and -5. (Appendix J, target FOUS-1)
- Planned refresh of secure archival storage hardware.
- Explored new methods of system and user security to include Continuous Security Monitoring. (Appendix J, target FOUS-1)

Planned Activities in FY19:

- Expand home and project file systems to adequately support user workloads through FY20. (Appendix J, target FOUS-1)
- Continue to enhance system and user security including Continuous Security Monitoring for secure AT and CT systems. (Appendix J, target FOUS-1)
- Deploy secure HPSS replacement hardware. (Appendix J, target FOUS-1)

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 × 7 × 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 FY18:

- Operated CTS-1 Fire and Ice systems in full production use, peaking with over 95% utilization and 99% system availability.

- Conducted discussions and planning for CTS-1+ options and CTS-2. (Appendix J, target FOUS-5)
- Continued support of Crossroads (ATS-3) procurement. (Appendix J, target FOUS-1)
- Demonstrated success of leveraging monitoring infrastructure and analytics to improve system operations.

Planned Activities in FY19:

- Complete Crossroads (ATS-3) procurement. (Appendix J, target FOUS-1)
- Procure the Cyclone (CTS-1) system and transition to operations.
- Monitor systems and coordinate with the EC3E Project to maintain 24 × 7 operations and high availability of AT and CT systems. (Appendix J, target FOUS-3)
- Provide 24 × 7 × 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 FY18:

- Managed Dedicated Application Times and Dedicated System Times throughout FY18 while maximizing system availability at 99% for CT systems and 98% for Trinity.
- Completed a major upgrade to Trinity system software (UP05) and security and system patches to maintain ATS-1 and CTS-1.
- Identified a replacement firewall for the network supporting CT systems.

- In concert with CSSE, completed a production monitoring solution and demonstrated its ability to scale in providing analytical information to improve system operations.

Planned Activities in FY19:

- Replace the firewall for the network supporting CT systems.
- Support design and requirements-gathering activities for the future Crossroads system. (Appendix J, target FOUS-1)
- 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 FY18:

- Provided ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Produced a new data storage policy to provide the framework for how to use different tiers of storage. Identified and executed a new support model, enabling success for a boundary pushing 3D problem on the Trinity system. (Appendix J, target FOUS-1)

Planned Activities in FY19:

- Identify priority or milestone efforts on AT or CT systems requiring focused or intensive user support.
- Develop new training opportunities for expanding knowledge of best practices in using AT or CT systems, especially for code performance improvements.
- Provide ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Identify user requirements for improved data management of the different tiers of storage. (Appendix J, target FOUS-1)

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 be operational at all times, the project provides on-call support after hours and on weekends for facility related issues.

Accomplishments in FY18:

- Provided ongoing operations, maintenance, and configuration of electrical and mechanical systems for ASC computing platforms and facilities. (Appendix J, target FOUS-1)
- Achieved Critical Decision-2/3 (CD-2/3) for the EC3E Project. (Appendix J, target FOUS-3)
- Created and requested to submit the Mission Need Statement for the Future Supercomputing Infrastructure Project.

Planned Activities in FY19:

- Begin planning for installation of the Crossroads system. (Appendix J, target FOUS-1)
- Oversee construction of the EC3E Project to enable startup and commissioning. (Appendix J, target FOUS-3)
- Submit the Mission Need Statement and Program Requirements Document for the Future Supercomputing Infrastructure (FSI) Project.

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 Vanguard at-scale prototype platforms.

Accomplishments in FY18:

- Completed planning for the initial installation of Vanguard/Astra cabinets into Sandia External Collaboration Network. (Appendix J, target FOUS-2)
- Deployed the 6-SU CTS-1 Eclipse system and ordered 2-SU expansion.
- Replaced/repaired on-chip water cooling devices for CTS-1 platforms Cayenne and Serrano.
- Upgraded the Mutrino Application Readiness Testbed System to support code migration and optimization work for Trinity.
- Implemented SLURM scheduling features to improve resource allocation and system throughput.

Planned Activities in FY19:

- Integrate Astra, the Arm-based Vanguard platform, into the Sandia Restricted Environment. (Appendix J, target FOUS-2)
- Deploy the Vanguard/Astra file system and associated infrastructure. (Appendix J, target FOUS-2)
- Expand use of monitoring data available from SPLUNK, SLURM, and Lightweight Distributed Metric Service (LDMS) to expose production workloads, system status, file system free capacity, and job priorities via a web-based dashboard.
- Deploy mini-Sierra ART system. (Appendix J, target FOUS-1)
- Acquire infrastructure and implement initial core capabilities for the centralized SNL/ASC Dev/Ops environment.
- Collaborate on the Tri-lab CTS-2 market surveys and procurement planning. (Appendix J, target FOUS-5)
- Integrate the new Sandia Classified Network (SCN) file system into the ASC CTS environment.
- Acquire and deploy the new SRN file system.
- Support advanced testbed systems (HAAPS) and environment.

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 leverages Information Technology Service Management (ITSM) best practices to deliver support capabilities and services including: 1) direct user support through a service desk and as-needed tiered support; 2) ITSM incident, problem, change, and knowledge management processes and tools; 3) training facilities, equipment, and training services; and 4) a web portal for HPC-related information, real-time data, and documentation. The project also funds tri-lab user support activities and collaborative efforts such as ACES and PSAAP II.

Accomplishments in FY18:

- Provided user support for SNL and tri-lab ASC computing systems, including commodity technology systems, advanced technology systems Trinity and Sequoia, Trinity and Sierra advanced readiness systems, and advanced architecture testbed systems. (Appendix J, target FOUS-1)
- Completed development of new HPC and ACES tri-lab web portals using the JOOMLA content management system.
- Conducted evaluation of performance impacts to Sandia codes due to mitigation for Spectre/Meltdown system vulnerabilities.
- Completed transition to the SmartIT Service Desk service management tool.
- Supported ongoing development of Sandia's Integrated Workflow environment.

Planned Activities in FY19:

- 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. (Appendix J, target FOUS-1)
- Implement a support model for the new Astra prototype system. (Appendix J, target FOUS-2)
- Continue evaluation and customization of SmartIT tool.
- Continue improvement of HPC and ACES web portals. (Appendix J, target FOUS-1)
- Support ongoing development of Sandia's Integrated Workflow environment.

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 FY18:

- Completed early evaluation of 100 Gb/s encryptor functionality and performance.
- Demonstrated Version 1 of the FrETT data transfer tool with collaboration of ORNL.
- Placed HPSS Version 7.5.1 with RAIT capability into production.
- Completed testing and evaluation of cold plate cooling technology with NREL to inform future system acquisitions.
- Relocated and re-engineered the advanced testbed (HAAPS) environment.

Planned Activities in FY19:

- Prepare the new 725E HPC facility for Vanguard/Astra installation. (Appendix J, target FOUS-2)
- Test thermosyphon performance for 725E liquid cooling systems. Expand thermosyphon infrastructure if appropriate. (Appendix J, targets FOUS-2, FOUS-5)
- Deploy a high-speed network to support 725E HPC facility systems on the SRN and SCN. (Appendix J, targets FOUS-2, FOUS-5)
- Test cross-site performance of 100 Gb/s encryptor for possible deployment in FY19.
- Upgrade HPSS to Version 7.5.2.

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 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 FY18:

- Developed, deployed, and released TOSS 3.2 and TOSS 3.3. Provided concurrent support for both TOSS 2 (based on RHEL 6) and TOSS 3 (based on RHEL 7) with monthly updates that included new kernels, security updates, and bug fixes. Performed additional releases as necessary to address critical bugs and security issues.
- Added support to TOSS 3 for the Arm aarch64 architecture (in addition to x86_64 and ppc64le). Worked with RedHat to ensure that technologies of interest for future CT systems (such as AMD processors and GPUs (including ROCm), GPUDirect, and

Mellanox CX6) will be well-supported in RHEL and TOSS. (Appendix J, target FOUS-2)

- Added support to TOSS 3 for SLURM 17.02 (in addition to SLURM 2.3) as part of the tri-lab migration to running SLURM-only on all clusters.
- LDMS v3.4.6 included in TOSS distribution. LDMS v4, Baler (log analysis tool) v4, and Scalable Object Store (SOS) v4 were split into three projects for maintainability.
- Worked with Cray and other collaborative partners to run a set of network congestion experiments, while collecting network performance counter data for post run analysis, on a four-cabinet Cray XC40 system in support of detection and mitigation of congestion problems in the Cray Aries Dragonfly network.
- Worked with IBM in Austin to validate LDMS data collection on Power 9 using the IBM specialized performance counter kernel module.
- Used LDMS for characterizing a particular application's CPU, network, and memory usage on TLCC2 and CTS-1 platforms to develop recommended scheduling changes and enable increased throughput and resource utilization efficiency. Demonstrated periodic (0.1–10 Hz) collection of MPI performance metrics on Xeon and KNL machines with a particular simulation code and quantified overhead.
- Continued to enhance LDMS: Extended the PAPI hardware performance counter sampler that enables users to request defined counter menus at application launch. Transport and visualization of Kokkos performance data. Developed Grafana-based dashboards to enable visualization of raw and processed LDMS metric data to support run-time and post processing analysis and troubleshooting. Ported the NFS v3 metrics sampler to TOSS v3. Developed a new OmniPath LDMS host interface sampler. Extended LDMS source code and install scripts to provide all necessary information to third parties developing plugins using only the installed LDMS libraries and headers.

Planned Activities in FY19:

- Provide ongoing TOSS software development and support; develop/deploy TOSS 3.x (based on RHEL 7) to CTS-1 machines and continue maintenance of TOSS 2.x (based on RHEL 6) until legacy systems are retired.
- Continue to evaluate and integrate hardware technologies of interest for future CT systems. (Appendix J, target FOUS-5)
- Maintain and support SLURM through tri-lab collaboration.
- Investigate challenges, benefits, and opportunities for running TOSS on non-CT ASC systems.
- Evaluate new base operating system versions, additional package collections (EPEL, OpenHPC, etc.), and upcoming technologies (container runtimes) for inclusion into the next major TOSS release (TOSS 4).

- Release LDMS v4 including development of enhancements to enable more efficient and automated configuration and operation and higher-fidelity security posture (e.g., per-user authentication).
- Develop more application-focused data collection, transport, storage, analysis, visualization, and feedback to application and system software.
- Deliver continued performance and functionality enhancements to storage and analysis components of monitoring and log analysis software.
- Continue development of machine learning based detection of congestion in shared resources and anomalous behaviors of both dedicated and shared compute resources.
- Port all functionality to new architectures as appropriate (e.g., Arm, Power). (Appendix J, target FOUS-2)
- Work with vendors to improve instrumentation, APIs, and documentation to make monitoring and analysis more productive.
- Publish in conferences and workshops as appropriate.

Programming Environment Development/Support for Tri-Lab Systems

The goal of the Programming Environment Project is 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 emerging architectures (processors and accelerators), new systems technology and evolving programming models. This project entails software integration, feature enhancements, installation, training, support for vendor provided tools, open source software tools development, and assessment of containerized environments

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. The debugging efforts include vendor contracts to provide support and training and to enhance the tool capabilities. Tri-lab developed debugging tools such as STAT and the PRUNERS toolset are continually evolved to handle the state-of-the-art architectures and programming models.

The MPI integration and scaling efforts provide development support to the communities to add and fix features in both MVAPICH and Open MPI. 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 are in close collaboration to provide the technical guidance for the contracts. This currently includes tools such as O|SS, Tuning and Analysis Utilities (TAU), the MUST/Vampir correctness checking with MPI usage/performance optimization tool, and others.

Another focus area is adoption of a common compute environment across the three labs. This will be accomplished through a tri-lab tools build environment, broader adoption of the Tri-lab Compute Environment (TCE), and increased usage of the Spack HPC Software package manager.

Accomplishments in FY18:

- Continued development of performance analysis and debugging tools, including Open|SpeedShop, MUST, STAT, and the PRUNERS toolset. Promoted broader adoption of these tools at all three labs.
- Developed LDPXI as a parallel application performance survey tool and began using as a profiler and as a correctness/verification benchmark for other profilers.
- Developed new optimizations and bug fixes for MVAPICH2 and OpenMPI. Released new versions and deployed at the labs. Continued MPI tuning study through the MPIPERF microbenchmarks.
- Deployed the XALT2 prototype and evaluated to monitor library usage at LANL. Worked through shared issues through discussion amongst the three labs.
- Brought up Arm test clusters and began evaluation of the architecture. (Appendix J, target FOUS-2)
- Began evaluation of several containerized environments solutions as they are becoming increasingly of interest to users.
- Renewed contracts for MUST, MVAPICH2, TAU, Paratools, HPC Toolkit, Open|SpeedShop.

Planned Activities in FY19:

- Continue development and support efforts for debuggers, performance analysis tools, and MPI through open-source development and support contracts. Increase tri-lab collaboration and deployment of programming environment software, such as the PRUNERS toolset (which has gained user adoption at LLNL) and LDPXI (which has gained adoption at Sandia); both are of interest to the tri-labs.
- 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. (Appendix J, target FOUS-2)
- Broaden collaboration on efforts such as the compiler benchmarking suite (including Developer Tools Confidence suite), tri-lab tools build environment, the TCE, and the Spack HPC Software package manager.
- Assess 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 (WCTool) provides the ability for common resource usage reporting to ASC HQ and within the labs, and additionally assists in optimizing management of ASC computing resources.

Accomplishments in FY18:

- Increased control and protection of personal information. Removed SARAPE "Guest Login" and requiring user login via SAML (LLNL/LANL), HSPD-12, or Sandia CryptoCard, due to enhanced security requirements. Developed a new SARAPE external home page that will facilitate Trusted Identity login based on the DOE OneID database. Improved management of Personally Identifiable Information (PII) to reduce time personal information is stored in SARAPE database.
- Began evaluation of new SARAPE access control methods. Under consideration are: implementing domain restrictions within SARAPE, and allowing users listed in DOE's OneID database who meet the criteria for access.
- Defined requirements for MySARAPE Phase 2 improvements and began developing the updated interface.
- Increased automated processing of SARAPE requests to improve accuracy and efficiency. Deployed new SARAPE dev and quality servers with upgraded RedHat OS. Began detailed documentation of Sandia SARAPE processes.
- Deployed multiple versions of WCTool for SNL and completed work with LLNL to host WCTool for them at SNL.
- Collaborated among the tri-labs to discuss SLURM deployment and address issues, documented metrics definitions and methods, discussed Tri-Lab development collaboration for R/LaTeX HQ report generation solution.
- Compiled and delivered quarterly Tri-Lab workload characterization usage reports to NNSA.

Planned Activities in FY19:

- Continue to operate the tri-lab SARAPE process for all tri-lab remote access account requests and implement evolutionary improvements.
- Improve management of Trusted Identity, a new Sandia cyber-visitor role with streamlined approval using the DOE OneID database; develop and deploy improvements based on feedback from users and processing agents.
- Continue to improve the MySARAPE user interface experience, based on user and processing agent feedback.
- Investigate adding other DOE labs as SARAPE sites, e.g., the National Energy Research Scientific Computing Center (NERSC), ANL, and ORNL.
- Continue efforts to meet new and/or expanded tri-lab ASC HQ reporting requirements with WCTool enhancements; investigate and develop streamlined ASC HQ reporting tools; address WCTool bugs and update software to current versions and security updates. Address issues in evolving tri-lab computing environments.

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 FY18:

- Provided HPC procurement, system administration, and operational support.
- Decommissioned legacy HPC system and infrastructure.

Planned Activities in FY19:

- Procure and install additional CTS-1 based compute clusters.
- Procure and install Lustre file systems to replace old systems.
- Decommission old compute clusters and Lustre clusters.
- Provide HPC procurement, system administration, and operational support.

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 FY18:

- Continued stabilization and maximized availability for new CTS-1 Global Security systems.
- Retired and decommissioned former Global Security production systems.

Planned Activities in FY19:

- Deploy a replacement storage system supporting CTS-1 Global Security systems.
- Provide 24 × 7 support for ongoing operations of CTS-1 Global Security systems and associated infrastructure.
- Assist with requirement determination for increasing resources available to the Global Security community.

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 FY18:

- Provided operations and support for tri-lab use of NSCC systems.
- Deployed upgraded file systems and security environment for NSCC systems.

Planned Activities in FY19:

- Continue operations and support for tri-lab use of NSCC systems.

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 FY18:

- Improved performance and applications portability and performance in center software (Arches, Hypre, and Uintah).
- Incorporated performance analysis into in situ visualization with lightweight wrappers in Uintah-X.
- Implemented improved ash deposition rate, properties, and heat transfer models to reduce most significant model form uncertainty.
- Demonstrated center V&V/UQ strategy to produce model form uncertainty in overarching predictive quantities of interest within 5%.

Planned Activities in FY19:

- Move to a performant and portable full problem on a post-Titan architecture, using for example all GPUs, while extending the infrastructure to broader method types.
- Extend and enhance performance analysis and in situ visualization.
- Perform targeted validation experiment to reduce model form uncertainty in physics components of the overarching problem.
- Harden the V&V/UQ strategy to produce extrapolated uncertainty for predictive quantities of interest.

University of Illinois, Urbana-Champaign

Center for Exascale Simulation of Plasma-Coupled Combustion

Accomplishments in FY18:

- Made quantitative prediction of laser-induced Transient Ignition Kernel (TIK) characteristics in an inventive high-speed flow configuration, with multiple quantities of interest designed to assess predictive capacity and uncertainty procedures for future combustion applications.
- Used the Illinois Code Environment (ICE) to selectively and noninvasively invoke community tools (PIPs, Chill, etc.).

- Finalized PlasCom2 Golden Copy design; completed inert simulation on our annual prediction meshes.
- Designed and conducted advanced diagnostic experiments in stand-alone laser-induced breakdown (LIB) experiments for model development and validation.
- Placed graduates at NNSA labs, set up PSAAP II student internships, involved lab personnel in Ph.D. dissertation evaluation, and conduced and published collaborative research with lab personnel.

Planned Activities in FY19:

- Use PlasCom2 Golden Copy to predict, with end-to-end quantified uncertainty, transient flame evolution in a cavity recess under a high-speed flow.
- Selectively and comparatively invoke CS tools using ICE and evaluate them on our full-scale coupled-physics PlasCom2 predictions.
- Use CS tools in node-to-node comparisons on advanced architectures.
- Use ultra-fast (femtosecond) diagnostics to close models for early LIB ignition seeds, a key uncertainty in FY18 efforts.

Stanford University

Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment

Accomplishments in FY18:

- Concurrently measured and validated fluctuating quantities (temperature, transmission. and particle concentration)
- Implemented full physics models in Soleil-X and demonstrated portability on Summit and Titan.
- Developed, implemented, and verified point particle correction models for momentum/heat transfer and particle-wall collisions using results and analysis of particle-resolved simulations.
- Used combined UQ/Validation and modeling to capture the effect of inhomogeneous wall heating.
- Developed and demonstrated fast compression techniques for I/O and boundary condition treatment.

Planned Activities in FY19:

- Assess the homogenization error for the radiation transmission through the particle phase.
- Demonstrate scalability of the Soleil-X solver with full physics capabilities.
- Evaluate parallelism options in ensemble analysis for the multifidelity UQ process.

- Validate Soleil-X in high-power, high-loading conditions.
- Develop low-fidelity models for particle phase.

University of Florida

Center for Compressible Multiphase Turbulence

Accomplishments in FY18:

- Released CMT-nek as a powerful compressible multiphase flow code co-designed for exascale performance.
- Performed two iterations of the demonstration problem (Frost case) with CMT-nek and completed the mesoscale simulations of Sandia's particle curtain experiment for the iterative forensic UQ analysis.
- Worked closely with collaborators at Eglin Air Force base and carried out the proposed macroscale experiments at the Eglin blast pad.
- Completed a shock velocity campaign to determine the variability caused by using the shockwave as a triggering mechanism and to determine whether the shock is accelerating or decelerating in our region of measurement.
- Fully developed enhanced Behavioral Emulation (BE) methodology and tools to perform predictive simulations for over a million cores in preparation for algorithmic and parametric design-space exploration (DSE) of CMT-nek on notional architectures using SST- and FPGA-based tools.

Planned Activities in FY19:

- Carry out full-scale, full-physics blast pad simulations of the demonstration problem with CMT-nek. Conduct a large suite of mesoscale simulations of Eglin's mesoscale experiments with CMT-nek for UQ.
- The co-design team will implement nonideal equations of state, the PIEP (pairwise-interaction extended point-particle force model), include shock detectors, positivity-preserving solution limiters, sponge layers or other robust boundary conditions capable of transmitting shock waves out of the domain, and surface fluxes for discontinuous artificial viscosity into CMT-nek.
- At ASU, we will supplement the current data sets with more particle sizes and initial bed heights. We will take data in very specific instances in time to further study the shock tube experiment and reduce variability.
- Fully demonstrate the use of BE methods and tools to support DSE for ongoing development of CMT-nek and the exploration of notional/near-future HPC architectures.

- Fully demonstrate the Uncertainty Budget (UB) framework to microscale and ASU vertical shock tube simulations and experiments and propagate these uncertainties to finalize the UQ of the demonstration problem.

Texas A&M University

Center for Exascale Radiation Transport

Accomplishments in FY18:

- Completed the Level 1 and Level 2 experiments in the V&V/UQ hierarchy and found that the Level 1 results did not agree with the hypothesis of negligible model error.
- Determined that the discrepancy was caused by material compositions differing substantially from manufacturer's specifications. Designed and launched an effort to infer the correct compositions with quantified uncertainties.
- Revised the experimental plan in response to the declaration of vendor that they could not deliver the promised pulsed neutron generator. Despite this and the material-composition setback, developed detailed plans to complete the first pass through all levels of our V&V/UQ hierarchy (including the Year-5 experiment) during summer 2018.
- Developed a new residual Monte Carlo method for estimation of numerical error in the PDT calculations and developed an implementation plan that is expected to be completed by fall 2018.
- Verified the parallel uncollided flux algorithm outside of PDT and plan implementation and testing in PDT by fall 2018.

Planned Activities in FY19:

- Refine Year-5 experiment suite based upon results from the initial cycle (simulations-experiments-UQ) and execute measurements, simulations, and V&V/UQ on the refined experiments.
- Test and improve the residual Monte Carlo error estimation method.
- Improve the parallel performance model to include the load-leveled partitioning and use the resulting model for run-time optimization of the tradeoff between sweep efficiency and load balancing on unstructured meshes.
- Implement new Sweep-Plane Data Structure in PDT, which will reduce overhead for all meshes and enable efficiency gains that are crucial for unstructured meshes.
- Develop, analyze, implement, and test new sweep algorithms to utilize “accelerator” hardware such as GPUs.

University of Notre Dame

Center for Shock Wave Processing of Advanced Reactive Materials

Accomplishments in FY18:

- Developed Design and Analysis of Computer Experiments (DACE) for poro-visco-plastic materials under impact conditions and advanced chemical reaction kinetics for Ni-Al system.
- Performed full system DACE studies and identified favorable pressure conditions for material synthesis.
- Released C-SWARM codes to GitHub and released Docker container for PGFem3D.
- Advanced verification coverage, improved software engineering practice, and optimized PGFem3D-HPX-lite with extended scaling to 64,000 cores.
- Developed 3D Multi-Resolution Wavelet Toolkit (MRWT).

Planned Activities in FY19:

- Deploy multiscale, multifidelity machine learning emulator for chemo-thermo-mechanical Ni-Al composites.
- Finalize development of MRWT and parallelize it. Further optimize PGFem3D with MTL/TTL-HPX libraries and extend the scaling.
- Deploy Python-based hybrid multiscale engine for macro- and microcoupling and perform detailed simulations of chemically reactive Ni-Al powders.
- Validate advanced QoIs (shock speed, final microstructure, temperature, etc.) for UQ certified DACE and calibrate chemo-thermal-mechanical, poro-visco-plastic material properties using Bayesian inference and determine parametric uncertainties.

Appendix I: Construction and Capital Equipment

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

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY19**	**FY20**	**FY21**	**FY22**	**FY23**	**Future**	**Contact**
LLNL	ECFM	Line Item	CD-1	CD-1 Baseline Cost Range is 58,000–125,000	6,500	21,500	49,000	30,500	17,500	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LANL	Trinity: Acquisition of Trinity (ATS-1) system	Capital Equipment, MIE	CD-4	187,000	177,000	5,000	5,000	0	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Crossroads: Acquisition of Crossroads (ATS-3) system ($115 M) and water cooling and electrical power installation ($12 M)	Capital Equipment, MIE	CD-1/3a 6/29/2016	127,000	4,000	20,000	24,000	45,000	14,000	8,000	12,000	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	ATS-5 system ($500 M) and water cooling and electrical power installation ($50 M)	Capital Equipment, MIE		550,000	0	0	0	1,000	1,000	20,000	488,000	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-2/3 4/3/2018	49,887	23,764	24,000	1,696	427	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	Future Supercomputing Infrastructure: 80 MW power and cooling and 500 lb/sq ft structural support	Capital Equipment, Line		183,000	0	0	2,000	4,000	28,000	41,000	108,000	Jason Hick 505-667-4477 jhick@lanl.gov

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY19**	**FY20**	**FY21**	**FY22**	**FY23**	**Future**	**Contact**
SNL	Vanguard Phase 1: Acquisition of Astra system	Capital Equipment, MIE		25,000	19,000	6,000						Ken Alvin 505-844-9329 kfalvin@sandia.gov
SNL	Vanguard Phase 2: Acquisition of Vanguard Ph. 2 Prototype system	Capital Equipment, MIE		25,000			12,000	13,000				Ken Alvin 505-844-9329 kfalvin@sandia.gov
SNL	CTS Additional Systems	Capital Equipment, MIE				0	5,000	0	5,000	4,000		Tom Klitsner 505-844-1901 tklitsn@sandia.gov

Appendix J: FY18–FY22 Program Targets

This appendix lists the targets for each subprogram for the FY18–FY22 period as found in the *FY18–FY22 ASC Program Plan*. Most of the targets apply to all laboratories, but several are specific to either the physics laboratories (LLNL and LANL) or the engineering laboratory (SNL). The targets represent areas of focused work intended to achieve a specific deliverable by the target date. In many cases, additional work will be required in the focus area to continue to improve the capability, or the target may be the first in a series of targets in the focus area. Because the targets represent specific deliverables, they do not cover the entire program. Multiple areas not covered by the targets require significant research and development but are not expected to achieve major deliverables in the period. Other work not covered by the targets is driven by near-term issues encountered during LEPs, Alts, and the AAR.

The objectives for the ASC Program, along with the targets for the subprograms, represent a prioritized set of activities over this period necessary to maintain the simulation capability for stockpile stewardship and broader nuclear security issues, and to begin to address the remaining gaps in predictive capability.

Integrated Codes

IC-1. Demonstrate performance portability for all IDCs on ATS-1, ATS-2, and ATS-3. This is necessary to maximally support the DSW community by productively using all ASC platforms. The target dates are FY18 for Trinity, FY20 for Sierra, and FY23 for Crossroads, appropriately coinciding with machine delivery.

IC-2. Improved fidelity physics and numerics capabilities for 3D simulations. Near-term plans include 3D applications demonstrations together with the necessary code upgrades, e.g., 3D production capability of IDCs on ATS platforms. The target date is FY19.

IC-3. Enhance simulation workflow to support and improve simulation efficiency for current and future stockpile activities. The target date is FY19. Simulation workflow enhancements will support efficient modeling of stockpile systems, enable validation with respect to relevant experiments, and facilitate practical uncertainty quantification. Workflow tools are necessary to support the Stockpile Responsiveness Program (SRP).

IC-4. Initial simulation capability for hostile environments. The target date is FY20, coinciding with a relevant PCF pegpost. The target includes production code and workflows enabling a complete capability for hostile environment simulation.

IC-5. Enhanced modeling of HEDP relevant validation experiments. The relevant modeling capabilities are 3D resistive magnetohydrodynamics (MHD) and the physics necessary to model 3D direct and indirect-drive ICF implosions. These capabilities are required to validate the simulation tools for high energy density physics (HEDP), and to permit training designers and engineers for environments that were traditionally associated with nuclear tests. The target date is FY22.

Physics and Engineering Models

PEM-1. Assess current capabilities and deliver improved capability to assure survivability in hostile environments. The target date is FY20.

PEM-2. Provide initial modeling capabilities for aging and manufacturing assessments that capture structure features in new and aged materials. The target date is FY20.

PEM-3. Revamp foundational materials modeling infrastructure to fully support and utilize next-generation architectures. The target date is FY21.

PEM-4. Develop UQ/machine learning toolkits to enable physics models and holistic data assessments. This encompasses projects by several different subcomponents, such as strength and damage and nuclear data. The target date is FY22.

PEM-5. Improve physics models relevant to full range of applications. This includes improved modeling of multiphysics response to combined abnormal environments, expanding current inline opacity capabilities to support modeling certification efforts and hostile environments, and implementing phase aware material models for strength and ejecta. Target date is FY22.

Verification and Validation

V&V-1. Enhance the verification and validation suite for nuclear and non-nuclear components, including components produced by advanced manufacturing techniques. The target date is FY19.

V&V-2. Extend the verification and validation infrastructure to include next-generation IDCs. The target date is FY20.

V&V-3. Develop an uncertainty bounding approach for physical processes. The target date is FY20.

V&V-4. Develop a comprehensive simulation uncertainty estimation methodology. The target date is FY21.

V&V-5. Develop a validation suite for the hostile survivability simulation component. The target date is FY21.

Advanced Technology Development and Mitigation

ATDM-1. Develop and deploy next-generation application codes and required software infrastructure.

a. Complete abstraction infrastructure and modular component technologies design and initial implementation. The target date is FY18.

b. Demonstrate viability of new ATDM computer science infrastructure and initial integration of sufficient physics/engineering features to simulate a real DSW problem. The target date is FY19.

c. Demonstrate and evaluate ATDM code performance on available ASC large-scale systems. The target date is FY20.

d. Optimize ATDM code performance on available ASC large-scale systems. The target date is FY22.

ATDM-2. Deploy advanced architecture prototype systems and develop accompanying tri-lab software environment.

a. Deploy the Vanguard-Phase 1 system, named Astra, with a functioning software stack, leveraging FOUS Common Computing Environment (CCE) effort. The target date is FY19.

b. Deploy the Vanguard-Phase 2 system with a functioning software stack, leveraging FOUS CCE effort. The target date is FY21.

ATDM-3. Interagency Collaborations. Complete Phase 1 of the DOE-National Cancer Institute (NCI) Collaboration Project. The target date is FY19.

ATDM-4. Develop and deploy new capabilities for hostile environment simulations. Target date is FY21.

Computer Systems and Software Environment

CSSE-1. Platform Acquisition and Deployment

a. Procure and deploy Crossroads, including Non-Recurring Engineering (NRE) activities. The target date is FY21.

b. Procure and deploy El Capitan, including NRE activities. The target date is Q1FY23.

c. Procure and deploy CTS-2 systems across tri-labs. The target dates are FY21 and FY22.

CSSE-2. Software Environment Deployment

a. Develop and deploy software and programming environments for Crossroads. The target date is FY20.

b. Develop and deploy software and programming environments for El Capitan. The target date is FY22.

c. Develop and deploy TOSS for CT-2. The target date is FY20.

CSSE-3. Beyond Moore's Law (BML) Computing Research

a. Demonstrate applications of interest on BML-funded hardware. The target date is FY19.

b. Demonstrate applications of interest on BML-funded hardware. The target date is FY21.

Facility Operations and User Support

FOUS-1. Installation and operation of Sierra and Crossroads. Sierra will be deployed to users in FY18. The target date for installation and the beginning of operations is FY18. The Crossroads system will be delivered and installed in FY21 and deployed in FY22. The target date for the installation preparation is FY21.

FOUS-2. Preparation for advanced technology prototype (at scale). The first prototype system (the Vanguard Arm prototype system named Astra) is scheduled for installation in FY18 and full deployment in FY19. Completion of Building 725-East is scheduled for FY18 in time for the system delivery. Astra is expected to be transitioned to the classified environment in FY20.

FOUS-3. LANL Exascale Class Computer Cooling Equipment Upgrade. Expand the warm-water cooling system at the Strategic Computing Complex (SCC) to enable up to 33 MW of supercomputing. The upgrade is essential to closing the current capability gap for warm-water cooling prior to installation of the Crossroads system in FY21.

FOUS-4. Exascale Facility Modernization Project (ECFM), facility prep for ATS-4 system. The target date for facility modernization design and construction activities is FY19–FY21. Deployment of ATS-4 is targeted for FY23.

FOUS-5. Installation and operation of CTS-2 and development of the associated software stack. Planning for the CTS-2 deployments will begin in FY19. The initial target date for installation, deployment and operation of CTS-2 is FY21. The FOUS subprogram is also responsible for developing the tri-laboratory system software stack for the CT systems, including the operating system and utility software. The target date for CTS-2 software is early FY21.