LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-825581

ASC FY22 IP Rev 0 Final Document

K. Howard

August 9, 2021

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

Advanced Simulation and Computing

FY22 IMPLEMENTATION PLAN

Version 0

August 9, 2021

NA-ASC-127R-21-Rev.0-IP

LLNL-TR-825581

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	*Chris Clouse, Jason Pruet, Scott Collis*	*Initial*	*Thuc Hoang*	*08/03/21*	*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 DOE National Nuclear Security Administration (NNSA) 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 NNSA to support stockpile requirements. The ultimate goal of the SSP, and thus of 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] Executing 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. ASCI officially evolved into the ASC Program in fiscal year (FY) 2005, but the mission remains essentially the same: provide the simulation and computational capabilities that underpin the ability to maintain a safe, secure, effective nuclear weapon stockpile, without returning to underground nuclear testing.

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 Stockpile Management and Production Modernization programs (formerly Directed Stockpile Work) 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 FY22. The Initial IP, effective <DATE>, 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	FY2022 Target	Endpoint Target
Advanced Simulation and Computing Program	Nuclear Weapons Simulation Capability Progress*	65%	100% (FY2030)

*With OMB approval, the performance measure "Reduced Reliance on Calibration" has been cancelled and replaced with the new performance measure "Nuclear Weapons Simulation Capability Progress" with reporting initiated for FY22.

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.

The ASC Program has been identified as being of the Standard Management category, as defined in the DP Program Execution Instruction.

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 Vision*, 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 FY22 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 Stewardship Capability Delivery Schedule (formerly Predictive Capability Framework), as described in the FY22 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:

The **Integrated Codes** (IC) subprogram develops the mathematical descriptions of the physical processes of nuclear weapon systems and functions. 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 Stockpile Management and Production Modernization 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, e.g., 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 Defense Programs' Experimental Science Program (NA-113), which provides the experimental data that inform development of new models used in simulation codes. PEM also

[2] U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2022 Stockpile Stewardship and Management Plan*, Report to Congress, TBD.

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

The **Verification and Validation** (V&V) subprogram provides evidence that the models in the codes produce mathematically credible answers that reflect physical reality. V&V focuses on establishing soundness in integrated simulation capabilities by collecting evidence that the numerical methods and simulation models are being solved correctly, and whether the simulation results from mathematical and computational models implemented into the codes are in sync with real-world observations. The V&V subprogram funds the critical skills needed to apply systematic measurement, documentation, and demonstration of the ability of the models and codes to predict physical behavior. The V&V subprogram is developing and implementing uncertainty quantification (UQ) methodologies as part of the foundation for the QMU process of weapons assessment and certification. The V&V subprogram also drives software engineering practices to improve the quality, robustness, reliability, and maintainability of the codes that evaluate and address the unique complexities of the stockpile.

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 and development of new codes, V&V will ensure the modifications and new codes are subjected to thorough V&V methodologies.

The **Advanced Technology Development and Mitigation** (ATDM) 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 Stockpile Management and Production Modernization missions 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.

In FY21, ASC began in earnest the transition of production-ready simulation and computing technologies from ATDM to the other ASC subprograms (IC, PEM, V&V, and CSSE) and the subprogram will stand down by end of FY23.

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 continue to pursue advanced R&D in next-generation computing technologies and also 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 FY22 major activities for the ASC Program include:

- Continue to advance 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, as well as secure transportation and production facility processes. Integrate production-ready simulation capabilities from ATDM.
- Further verification suite development for ASC codes focusing on radiation hydrodynamics.
- Support upcoming pegposts in the Stewardship Capability Delivery Schedule (SCDS) by quantifying sensitivities and leveraging the high-energy-density physics suite.
- Engage in helping define and create credible workflow processes and tools to enable V&V and UQ efforts.
- Adapt existing codes to new architectures and migrate current design, engineering, and safety codes to run efficiently on hybrid computer architectures.
- Manage the deployment of the Crossroads system at Los Alamos National Laboratory (LANL) and El Capitan system at Lawrence Livermore National Laboratory (LLNL), both of which will be used to address stockpile stewardship issues and to advance predictive science.
- Steward the Exascale Computing Facility Modernization (ECFM) project at LLNL through the start-of-operation to operation and final design stages. ECFM will support the deployment of the first NNSA exascale system, El Capitan, in FY 2023.
- Transition production-ready, next-generation codes and computing technologies from ATDM to the other ASC subprograms (IC, PEM, V&V, and CSSE). Continue DOE-NCI (National Cancer Institute) collaboration projects.

- Expand the predictive capability assessment suites to include additional UGTs, hydrodynamic tests, and scaled experiments.
- Provide materials models and simulation capabilities that strengthen the ability of the complex to manufacture critical nuclear weapons components.
- Develop and mature capabilities to support activities related to the W93.
- Develop and mature capabilities to support qualification for the W80-4 and the W87-1 Programs.
- Continue ramp-up of the SNL Accelerated Digital Engineering initiative to improve and expedite the weapon system development engineering phase.
- 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) program on the DOE ECI (Exascale Computing Initiative) to advance research and development technologies that will eventually enable procurement and efficient utilization of exascale-class HPC platforms.
- Maintain mentoring program for early career staff.

The drivers of the ASC Program that require the FY22 budget include the following:

- The Nuclear Weapons Council approved the stockpile strategy articulated in the 2018 NPR, which includes completion of 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 simulation and 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, qualification, 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 Stockpile Management and Production Modernization user community. The IDCs currently perform well for general-mission-related activities; however, as the stockpile life is extended and aging takes the current stockpile further away from the data collected from UGTs, maintaining the nuclear weapons stockpile will require IDCs that will be more science-based predictive and use HPC resources more effectively.

A strategic driver for NNSA simulation and computing investment is the global shift in fundamental computing architecture. ASC capabilities that support the Stockpile Management and Production Modernization 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 computing and AI/ML (artificial intelligence/machine learning) sectors 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 continues some R&D partnerships with industry in order to minimize the disruptive technological impacts on the NNSA IDCs.

The ASC Program has developed a new platform 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 of an exascale computing capability for the nation.

IV. Funding Guidance

To support the scope of work contained in this Work Authorization (WA), funding will be distributed through the existing Approved Funding Program (AFP) process described in DOE Order 135.1A and NNSA BOP 001.331, or successor Order/BOP, as applicable. The AFP is adjusted on an as needed basis for the execution of congressionally approved programs, projects, or activities (PPAs). Specific work activities are authorized via this document, with incremental funding plans for each site authorized via the AFP and obligated via formal contract modification. The work contained herein will be funded on an incremental basis. The contractor is authorized to expend up to the dollar amount indicated in the Program/OCL funding table or as otherwise noted in the AFP allotment, whichever is less, and subject to the availability of funds in the M&O contract.

Note: As indicated in Section 4.0, the contractor is required to notify the contracting officer, in writing, when they recognize that they will exceed the estimated cost by 10%, insufficient actual and expected funding is available to continue work, or if their actual funding is insufficient to operate until further expected funding.

Program/Operational Control Level	**President's FY2022 Budget Request**	**FY2022 Operating Target**	**FY2022 Enacted/Full-Year**	**Difference between Request versus Enacted**
ASC Operations	$747.012M	$xxxM	$xxxM	$xxxM

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 tables below list the current ASC Level 1 and Level 2 Milestones for FY22. 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 program targets referenced in the FY22–FY26 Target column may be found in Appendix I.

Table V-1B. ASC Level 2 Milestones for FY22[3]

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY22–FY26 Target
IC	TBD	GPU Port of the FEusion Library	9/30/22	LLNL	IC-1, IC-2
IC	TBD	Create a Persistent Simulation Data Management Capability for Coupled Workflows	9/30/22	LLNL	IC-3, IC-4
IC	TBD	Use Vidya ALE to Guide Mesh Relaxation on Programmatically Relevant Problems	9/30/22	LLNL	IC-4, IC-5
IC	TBD	Demonstration of Modular Turbulence Capability Integrated into LLNL's NextGen Code	9/30/22	LLNL	IC-5
IC	TBD	Optimal Sensor Placement for Reentry Flight Tests for Normal Mechanical Response	9/30/22	SNL	IC-4
IC, V&V	TBD	Automated deployment of credibility evidence for Nuclear Deterrence (ND) designer workflows	9/30/22	SNL	IC-4 V&V-2
IC	TBD	Interface ITS (Integrated Tiger Series) with SGM (Scalable Geometric Modeler) to perform radiation transport on CAD geometries	9/30/22	SNL	IC-2

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

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY22–FY26 Target
IC	TBD	Demonstrate abnormal thermal qualification simulations on ATS-2	9/30/22	SNL	IC-2
IC	TBD	Demonstration of new Monte Carlo transport capabilities in support of the weapons program	9/30/22	LANL	IC-4, IC-5
IC	TBD	Scalable 2D and 3D topology-based meshing for Complex Weapons Applications	9/30/22	LANL	IC-5
PEM	TBD	Aged Plutonium Strength Change	6/30/22	LLNL	PEM-2
PEM	TBD	Age Aware Pu Equations of State	6/30/22	LLNL	PEM-2
PEM	TBD	Pu Kinetics Sensitivities	9/30/22	LLNL	PEM-3
PEM	TBD	Benchmark and assessment of the surrogate rection method for determining unknown (n, n') and (n, 2n) reaction cross sections	9/30/22	LLNL	PEM-3
PEM	TBD	Adapting an electrodynamic model to underwrite safety procedures for assessing the threat of electrostatic discharge events during NW dismantlement at Pantex	9/30/22	LANL	PEM-3
PEM	TBD	Modeling EOS variation due to variations in manufacturing methods	9/30/22	LANL	PEM-2
PEM	TBD	Support Simulations of the effects of surface defects caused by manufacturing on the performance of pits	9/30/22	LANL	PEM-2
PEM, V&V	TBD	Credible Electrochemical Simulations of Full Thermal Batteries	9/30/22	SNL	PEM-3, V&V-1
PEM, V&V	TBD	Detecting Aging of Electronic Devices	9/30/22	SNL	PEM-2
V&V	TBD	Inclusion of modern AGEX and subcritical experiments into modeling suites	9/30/22	LLNL	V&V-1, V&V-2
V&V	TBD	VVUQ Assessment of Capabilities for Modeling the Strength and EOS of Aging Materials in support of the SCDS FY22 Aging Pegpost	9/30/22	LANL	V&V-2, V&V-3
V&V	TBD	Coupled aerothermal ablation for hypersonics (Multiphysics credibility	9/30/22	SNL	V&V-1 V&V-2
V&V	TBD	Deploy a fully featured, usable automation of the multifidelity UQ scheme in DAKOTA	9/30/22	SNL	V&V-1
ATDM, CSSE	TBD	Workflow Portability Across Cloud Services	8/30/22	LLNL	CSSE-2, ATDM-1.d
ATDM	TBD	Embedded Components	9/30/22	SNL	ATDM-1
CSSE	TBD	TOSS4 validation on CTS-2	6/30/22	LLNL	CSSE-2,d

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY22–FY26 Target
CSSE	TBD	Crossroads System Integration Readiness	6/30/22	LANL	CSSE-1a, FOUS-1
CSSE	TBD	Evaluation of Data Flow HW Accelerators for Mission Relevant Mini-Apps and Kernels	9/30/22	SNL	CSSE-1.d
FOUS	TBD	CTS-2 Initial Deployment	6/30/22	LLNL	FOUS-5
FOUS	TBD	CORAL-2 EAS-3 Deployment	6/30/22	LLNL	FOUS-5
FOUS	TBD	Classified Persistent Database Services	9/30/22	LLNL	FOUS-1.a FOUS-5
FOUS	TBD	CTS-2 Installation	9/30/22	LANL	FOUS-5
FOUS, CSSE	TBD	CTS-2 delivery	9/30/22	SNL	FOUS-5
FOUS, IC, ATDM, CSSE, PEM, V&V	TBD	Roadmap for Integrated Sandia ASC DevOps	9/30/22	SNL	IC-1, V&V-3, ATDM-1, CSSE-2, FOUS-1, FOUS-2

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 NNSA's Office of Management and Budget (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. They also conduct other conference calls on a monthly or as-needed basis depending on the urgency and nature of issues.
- **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 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 2021–2031* (August 2021)
- *ASC Computing Strategy* (August 2021)
- *ASC Right Size* (September 2016)
- *ASC Business Plan* (July 2015)
- *ASC Co-Design Strategy* (February 2016)
- *ASC Simulation Strategy* (August 2021)
- Memorandum of Understanding (MOU) between DOE Office of Science (SC) and DOE NNSA Office of Defense Programs (DP) for the coordination of exascale activities (November 2016)
- Memorandum of Understanding between DOE SC, NNSA DP, and the National Cancer Institute (NCI) for the joint research collaboration of advanced computing capability and biomedical research (June 2021)
- *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 Point of Contact

Thuc Hoang
NA-114, Director
Office: 202-586-7050
Fax: 202-586-7754
Thuc.Hoang@nnsa.doe.gov

Headquarters Implementation Plan Focal Point

K. Michael Lang
NA-114
Office: 301-903-0240
Michael.Lang@nnsa.doe.gov

Laboratory ASC Executives

Chris Clouse, LLNL, 925-422-4576, clouse1@lln.gov

Jason Pruet, LANL, 505-667-6577, jpruet@lanl.gov

S. Scott Collis, SNL, 505-284-1123, sscoll@sandia.gov

Laboratory Programmatic Points of Contact

WBS	Title	Contact
1.2.3.1	Integrated Codes	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov Walt Witkowski, SNL, 505-844-3869, wrwitko@sandia.gov
1.2.3.1.1	Engineering and Physics Integrated Codes	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov Walt Witkowski, SNL, 505-844-3869, wrwitko@sandia.gov
1.2.3.1.2	Specialized Codes and Libraries	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov Walt Witkowski, SNL, 505-844-3869, wrwitko@sandia.gov
1.2.3.1.3	Applications and Algorithms Research	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov Walt Witkowski, SNL, 505-844-3869, wrwitko@sandia.gov
1.2.3.1.4	Applications Research for Next-Generation Platforms	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov Walt Witkowski, SNL, 505-844-3869, wrwitko@sandia.gov
1.2.3.2	Physics and Engineering Models	Tom Arsenlis, LLNL, 925-424-2584, arsenlis1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jeffrey Payne, 505-844-4524, jlpayne@sandia.gov
1.2.3.2.1	Materials Response: Equation of State, High Explosives, and Mix and Burn	Tom Arsenlis, LLNL, 925-424-2584, arsenlis1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jeffrey Payne, 505-844-4524, jlpayne@sandia.gov
1.2.3.2.2	Transport, Plasmas, Atomic, Nuclear	Tom Arsenlis, LLNL, 925-424-2584, arsenlis1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jeffrey Payne, 505-844-4524, jlpayne@sandia.gov

WBS	Title	Contact
1.2.3.2.3	Engineering Science, Environments, and Response	Tom Arsenlis, LLNL, 925-424-2584, arsenlis1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jeffrey Payne, 505-844-4524, jlpayne@sandia.gov
1.2.3.2.4	Integrated Modeling and Applications	Tom Arsenlis, LLNL, 925-424-2584, arsenlis1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jeffrey Payne, 505-844-4524, jlpayne@sandia.gov
1.2.3.3	Verification and Validation	Ana Kupresanin, LLNL, 925-422-8348, kupresanin1@llnl.gov Catherine Plesko, LANL, 505-667-2345, plesko@lanl.gov Amanda Dodd, SNL, 925-294-6599, ajbarra@sandia.gov
1.2.3.3.1	Verification and Validation Methods	Ana Kupresanin, LLNL, 925-422-8348, kupresanin1@llnl.gov Catherine Plesko, LANL, 505-667-2345, plesko@lanl.gov Amanda Dodd, SNL, 925-294-6599, ajbarra@sandia.gov
1.2.3.3.2	Verification and Validation Assessments	Ana Kupresanin, LLNL, 925-422-8348, kupresanin1@llnl.gov Catherine Plesko, LANL, 505-667-2345, plesko@lanl.gov Amanda Dodd, SNL, 925-294-6599, ajbarra@sandia.gov
1.2.3.3.3	Data Validation, Archiving, Software Quality Assurance, and Training	Ana Kupresanin, LLNL, 925-422-8348, kupresanin1@llnl.gov Catherine Plesko, LANL, 505-667-2345, plesko@lanl.gov Amanda Dodd, SNL, 925-294-6599, ajbarra@sandia.gov
1.2.3.4	Advanced Technology Development and Mitigation	Teresa Bailey, LLNL, 925-424-6700, bailey42@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	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Chris Werner, LANL, 505-606-1906, cwerner@lanl.gov James R. Stewart, SNL, 505-844-8630, jrstewa@sandia.gov
1.2.3.4.2	Next-Generation Architecture and Software Development	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.4.3	Inter-Agency Co-Design	Teresa Bailey, LLNL, 925-424-6700, bailey42@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	Matt Legendre, LLNL, 925-422-6525 legendre1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.1	Commodity Technology Systems	Matt Leininger, LLNL, 925-422-4110, leininger4@llnl.gov Jim Lujan, LANL, 505-665-0718, jewel@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@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 Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.3	System Software and Tools	Scott Futral, LLNL, 925-422-1658 futral2@llnl.gov Jim Foraker, LLNL, 925-422-0252, foraker1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.4	I/O, Storage Systems, and Networking	Trent D'Hooge, LLNL, 925-423-6100, dhooge1@llnl.gov Debbie Morford, LLNL, 925-423-7086, morford1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov

WBS	Title	Contact
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 Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.6	Beyond Moore's Law	Katie Lewis, LLNL, 925-423-8593, lewis66@llnl.gov Matt Leininger, LLNL, 925-422-4110, leininger4@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.7	Next-Gen Computing Technologies	Matt Legendre, LLNL, 925-422-6525 legendre1@llnl.gov Michael Lang, LANL, 505-500-2993, mlang@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.6	Facility Operations and User Support	Ned Bass, LLNL, 925-422-9389, bass6@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	Scott Futral, LLNL, 925-422-1658, futral2@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	Ned Bass, LLNL, 925-422-9389, bass6@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@lanl.gov Stephen Monk, SNL, 505-284-2811, smonk@sandia.gov
1.2.3.6.3	Common Computing Environment	Ned Bass, LLNL, 925-422-9389, bass6@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@lanl.gov Stephen Monk, SNL, 505-284-2811, smonk@sandia.gov
1.2.3.6.4	Special Purpose Facilities, Systems, Operations, and Support	Ned Bass, LLNL, 925-422-9389, bass6@llnl.gov Jason Hick, LANL, 505-667-4477, jhick@lanl.gov Stephen Monk, SNL, 505-284-2811, smonk@sandia.gov

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:

Thuc Hoang, NNSA ASC Program Director	______________________ Signature	08/03/21 Date

Concurrence:

Chris Clouse, LLNL ASC Executive	*Concurrence received*	07/30/21
Jason Pruet, LANL ASC Executive	*Concurrence received*	08/05/21
S. Scott Collis, SNL ASC Executive	*Concurrence received*	08/03/21

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 or Advanced Architecture Prototype System
AAR	Annual Assessment Review
ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AES	Advanced Encryption Standard
AFP	Approved Funding Program
AI	Artificial Intelligence
ALE	Arbitrary Lagrangian-Eulerian
Alt	Alteration
AMD	Advanced Micro Devices, Inc.
AML	Advanced Machine Learning
AMR	Adaptive Mesh Refinement
AMT	Asynchronous Many Task
ANL	Argonne National Laboratory
API	Application Programming Interface
APS	Advanced Prototype System
ART	Application Regression Testbed
ASC	Advanced Simulation and Computing (formerly ASCI)
ASCI	Accelerated Strategic Computing Initiative
ASCR	Office of Science's Advanced Scientific Computing Research
ASD	Next-Generation Architecture and Software Development
AT	Advanced Technology
ATCC	Advanced Technology Computing Campaign

ATDM	Advanced Technology Development and Mitigation
ATOM	Accelerating Therapeutic Opportunities in Medicine
ATS	Advanced Technology System
ATSE	Advanced Tri-lab Software Environment
B&R	Budget and Reporting
BEE	Build and Execute Environment
BML	Beyond Moore's Law
C2C	Contours to Codes project
CCE	Common Computing Environment
CD	Critical Decision
CDA	Next-Generation Code Development and Applications
CEED	Center for Efficient Exascale Discretizations
CGNS	Computational fluid dynamics General Notation System
CHAI	Copy Hiding Application Interface
CI	Continuous Integration
CMF	Common Model Framework
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
CRADA	Cooperative Research and Development Agreement
CSSE	Computational Systems and Software Environment
CT	Commodity Technology
CTS	Commodity Technology System
DDR	Double Data Rate
DFT	Density Functional Theory
DIMM	Dual Inline Memory Module
DisCom	Distance Computing
DOD	Department of Defense

DOE	Department of Energy
DRAM	Dynamic Random Access Memory
EAP	Eulerian Application Project
EC3E	Exascale Class Computer Cooling Equipment (LANL)
ECFM	Exascale Computing Facility Modernization (LLNL)
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
EMPIRE	ElectroMagnetic Plasma in Radiation Environments
EOS	Equation of State
FAODEL	Flexible, Asynchronous, Object Data-Exchange Libraries
FIR	Fortran Intermediate Representation
FleCSI	Flexible Computation Science Infrastructure for Multiphysics
FOM	Figure of Merit
FOUS	Facility Operations and User Support
FPGA	Field-Programmable Gate Array
FrETT	Friendly Extensible Transfer Tool
GDDR	Graphics Double Data Rate
GPGPU	General-Purpose Graphics Processing Unit
GPU	Graphics Processing Unit
GSK	Glaxo-Smith-Kline
GUFI	Grand Unified File Index
HAMR	Heat-Assisted Magnetic Recording
HBM	High-Bandwidth Memory
HDF5	Hierarchical Data Format 5
HI	Hardware and Integration (ECP)
HPC	High Performance Computing
HPSS	High Performance Storage System
HQ	Headquarters
I/O	Input/Output
IC	Integrated Codes
ICD	Interagency Co-Design

ICR	Independent Cost Review
IDC	Integrated Design Code
IOSS	Input/Output SubSystem
IP	Implementation Plan
ITSM	Information Technology Service Management
KCNSC	Kansas City National Security Campus
KNL	Knights Landing processors
KvN	Koopman-von Neumann
L1	Level 1 (milestone)
L2	Level 2 (milestone)
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LAP	Lagrangian Applications Project
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LCW	Low Conductivity Water
LDMS	Lightweight Distributed Metric Service
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
LLVM	Low-Level Virtual Machine
LSCI	Large Scale Calculations Initiative
M&O	Management and Operations
MFEM	Modular Finite Element Methods
ML	Machine Learning
MLIR	Multilevel Intermediate Representation
MOU	Memorandum of Understanding
MPI	Message Passing Interface
MST	Multisite Target
MMAI	Monitoring, Metrics, Analytics, Integration
MuMMI	Multiscale Machine Learned Modeling Infrastructure
MW	Megawatt

NAS	Network-Attached Storage
NCHPC	Non-conventional High-Performance Computing
NCI	National Cancer Institute
NERSC	National Energy Research Scientific Computing Center
NFS	Network File System
NGCE	Next-Generation Computing Enablement
NGCT	Next-Generation Computing Technologies
NGP	Next-Generation Platform
NIC	Neural-Inspired Computing
NIH	National Institutes of Health
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NPU	Neuromorphic Processing Unit
NRE	Non-Recurring Engineering
NSCC	National Security Computing Center (SNL)
NSCI	National Strategic Computing Initiative
NVME	Nonvolatile Memory Express
NW	Nuclear Weapons
nWBS	National Work Breakdown Structure
OCF	Open Computing Facility
ORNL	Oak Ridge National Laboratory
OS	Operating System
OS/R	Operating System/Runtime
PDU	Power Distribution Unit
PEM	Physics and Engineering Models
PESP	Predictive Engineering Science Panel
PIEP	Pairwise Interaction Extended Point Particle Model
PNNL	Pacific Northwest National Laboratory
PPA	Programs, Projects, or Activities
PSP	Predictive Science Panel
QC	Quantum Computing

QIP	Quantum Information Processing
QMU	Quantification of Margins and Uncertainties
QoS	Quality of Service
R&D	Research and Development
RCE	Remote Computing Enablement
RDMA	Remote Direct Memory Address
RFI	Request for Information
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
ROL	Rapid Optimization Library
RVMA	Remote Virtual Memory Access
SAN	Storage Area Network
SAP	Scalable Applications Preparation
SARAPE	Synchronized Account Request Automated Process
SASI	Standards and Architectures for Storage and I/O
SAW	Sandia Analysis Workbench
SC	Department of Energy's Office of Science
SCC	Strategic Computing Complex (LANL)
SCDS	Stewardship Capability Delivery Schedule
SCF	Secure Computing Facility
SCN	Sandia Classified Network
SCR	Scalable Checkpoint Restart
SDN	Software Defined Network
SFI	Significant Finding Investigation
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SPOT	System for Performance Optimization Tracking
SRN	Sandia Restricted Network
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAT	Stack Trace Analysis Tool

TAU	Tuning and Analysis Utilities
TCE	Tri-lab Computing Environment
TIK	Transient Ignition Kernel
TLCC	Tri-lab Linux Capacity Cluster
TOSS	Tri-lab Operating System Stack
UGT	Underground Test
UQ	Uncertainty Quantification
V&V	Verification and Validation
VTK	Visualization Toolkit
WA	Work Authorization
WAN	Wide Area Network
WBS	Work Breakdown Structure
WSC	Weapon Simulation and Computing
Y-12	Y-12 National Security Complex
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 includes only 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 2020, and through quarter 3, fiscal year 2021, are reflected below for the ATDM subprogram.

- Developed new Spack functionality for developer workflows; significantly refactored MARBL code to better support Spack packaging; worked with MARBL team to convert MARBL build system to Spack. (Appendix I, target ATDM-1.d) (LLNL)
- Rolled out first release of Spack (0.16) including a new Answer Set Programming (ASP)-based concretizer and made numerous improvements and bugfixes as users tested it in production. Implemented binary bootstrapping for the new concretizer and GPG security in preparation for making the new concretizer default in 0.17. (LLNL)
- Applied our floating-point correctness tool suite, FLiT and FPChecker, to more applications and mini-apps; Released FPChecker version 0.2 with better GPU code integration support. (Appendix J, targets CSSE 1.b, IC-1, and ATDM 1.d) (LLNL)
- MFEM developed ATDM-specific GPU tools, algorithms and optimizations in support of the MARBL ATDM L1 Milestone, which successfully demonstrated large-scale GPU performance on ALE multi-physics problems. (Appendix I, target ATDM 1.d) (LLNL)
- Supported integration of Umpire into ASC application codes & libraries. Developed and delivered inter-process shared memory capability in Umpire to allow multiple application processes to share allocations. (Appendix I, target IC-1) (LLNL)
- BEE can now archive, clone, re-run workflows. (LANL)

- Charliecloud fully unprivileged container building on HPC systems. (LANL)
- Fusion release, capability used in section of VnV L2 Milestone: HED Hydrodynamics in the Common Modeling Framework. (LANL)
- Productivity boosts: **2-4x** in performance on Kokkos-aware compilation (***no code change***), **faster** compile times, much **smaller** executables (up to 12x). (LANL)
- Legion Co-designed with Ristra (FleCSI) using control replication to enable SPMD workloads to scale effectively using hierarchical task parallelism. (LANL)
- Provided high quality (production) KOKKOS and KOKKOS KERNELS support and consultation for ASC applications and libraries. (Appendix I, target ATDM-1) (SNL)
- KOKKOS Ecosystem Release 3.4. Provides support for Exascale Computing (complete backends for AMD and Intel GPUs, kernels for AMD and Intel GPUs. (Appendix I, target ATDM-1) (SNL)
- Developed a performance-evaluation suite for VTK-m, along with a driver to run visualization workflows on ATS and ASCR pre-exascale systems. (Appendix I, target ATDM-1) (SNL)
- Completed an in-depth performance analysis of a containerized version of ATDM SPARC application at scale. (Appendix I, target ATDM-1) (SNL)
- Developed and disseminated a Sandia ASC-wide container strategy that details how to integrate several activities supporting and leveraging containers, including the creation of working groups, integration with existing DevOps activities, and further development efforts surrounding unprivileged container build solutions with Podman and Buildah. (Appendix I, target ATDM-1) (SNL)
- Led the addition of a new chapter to the MPI Standard on partitioned communication operations. (Appendix I, target ATDM-1) (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Workflow Portability Across Cloud Services		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 08/30/2022		
ASC nWBS Subprogram: ATDM and CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, targets CSSE-2.a, CSSE-2.c, ATDM-1.d		
Description: This effort will create software, documentation, and a demonstration that illustrates how to run ASC workflows across multiple cloud services. It will identify an application or workflow that is of interest to ASC, which will be containerized and adapted to the cloud. It will demonstrate running this application on LC's PDS (Persistent Data Services) cloud services and on other appropriate cloud services that could include public clouds, other networks, or other labs. Documentation on producing containers, cloud-compatible software, optimization, and any software infrastructure will be made available to ASC code teams.		
Completion Criteria: Delivery of documentation and software that demonstrate containerizing and cloud deployments of ASC software.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: FOUS provided cloud infrastructure, CSSE and ATDM staff.		

Milestone (ID#): Embedded Components		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2022		
ASC nWBS Subprogram: ATDM		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target ATDM-1		
Description: Applications incorporate embedded components to take advantage of data in memory during the simulation to produce results in-situ, and remove barriers to scalability and reduce costs associated with pre-processing, analysis, and post-processing. In this milestone, three embedded components will be investigated: 1. *Mesh refinement within EMPIRE.* Enable exascale simulations, avoid time consuming transfer of large mesh files, and improve/streamline the problem setup. The mesh refinement will access the CAD geometry, so the refined mesh captures the geometries of ND interest. Mesh quality will be evaluated through multiple mesh refinements and improvement strategies will be studied. 2. *Adjoint Sensitivities and Calibration of Hypersonic Flow Problems in SPARC.* Establish a foundation for efficient and scalable calibration and vehicle/trajectory design by investigating adjoint sensitivities and embedded, reduced-space optimization problems in SPARC. 3. *Quantitative Analysis with Catalyst in SPARC.* Identify specific usability barriers and improvements to python interfaces for analyst-defined use cases (e.g., solution visualizations, and compute quantities of interest (QoIs) for V&V and UQ).		
Completion Criteria: 1. Provide and demonstrate in-situ mesh refinement capability in EMPIRE on ND-relevant geometry (e.g., W78) at scale with a target size of ~120B elements. 2. Assess feasibility, accuracy, and efficiency of computing steady-state parametric adjoint sensitivities along with their utility in solving calibration problems on the double-cone problem. 3. Provide and demonstrate in-situ quantitative analysis capabilities with Catalyst in SPARC, e.g., LES flowfield images, LES diagnostics, base pressure/heat transfer correlations, and sectional loads for SPARC, and document implemented interface improvements.		
Customer: EMPIRE, SPARC		
Milestone Certification Method: A program review is conducted, and its results are documented. The "handoff" of the developed capability to EMPIRE and SPARC will be documented.		
Supporting Resources: SGM, Data Propagation Components, Sacado, Tempus, ROL, TRILINOS linear solvers, Zoltan2, ParaView, Catalyst, Kitware		

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)

Next-Generation Computing Enablement (NGCE) is preparing ASC for the next generation of advanced computing technologies. Its project areas include portable user-level scheduling/runtime support for emerging workflows, system-level resource and power management/scheduling, advanced debugging/correctness tools and ROSE compiler technologies. NGCE's activities across these areas complement one another, and the advancements in each area synergistically support the overall project goal: providing readiness for ASC/ATDM applications and simulation workflows on exascale computing systems while mitigating increasing challenges stemming from the introduction of machine learning (ML). NGCE uses a three-pronged approach to exert its impact broadly: 1) significantly advance each individual R&D discipline, 2) drive interdependence among these disciplines and other ASC Program elements through strategic co-design tasks, and 3) provide broad community outreach and communications.

Accomplishments in FY21:

- Provided Flux solutions to a broader set of workflows including LLNL UQ, V&V and coupled workflows via validation suite team support, an ML-based drug design ATOM workflow (a SC20 COVID-19 Gordon Bell finalist), the Pilot2 MuMMI workflow (enabling it to utilize over 600,000 Summit node hours), the MFEM-based ExaConstit portion of an additive manufacturing workflow (demonstrating 4x job throughput improvements). (Appendix I, targets CSSE 1.b, IC-1, and ATDM 1.d)
- Applied our floating-point correctness tool suite including FLiT and FPChecker to more applications/mini-apps; Released FPChecker version 0.2 with better GPU code integration support. (Appendix I, targets CSSE 1.b, IC-1, and ATDM 1.d)
- Completed our analysis of a power monitoring dataset collected from IBM CSM tools on Lassen comprising of over 1.4 million jobs and incorporated the findings into our Flux/Variorum power management solution. (Appendix I, targets CSSE 1.b)
- Improved the robustness of ROSE and its inliner on ASC applications as well as C++20 support within ROSE. (Appendix I, target IC-3)

Planned Activities in FY22:

- Demonstrate Flux's end-to-end support on top of HPE's DataWarp/Rabbit Mock service. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d)
- An end-to-end demonstration of HPE's Common Tools Interface (CTI) within Flux to enable all of the CTI-based debugging tools including STAT, TotalView, as well as HPE's proprietary tools. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d)
- A multi-node demonstration of Flux/Variorum-based power monitoring services and an exploration of novel analysis techniques to characterize a job's power swing. (Appendix I, targets CSSE 1.b/2.b)
- Facilitate broad impact and adoption by disseminating the results of our investigations via various tutorials, presentations, publications, workshops, and conferences. (Appendix I, targets CSSE 1.b/2.b, IC-1, and ATDM 1.d)

RAJA Portability Suite: Abstractions for Portable Execution and Memory Management (LLNL)

RAJA Portability Suite supports the development of performance-portable abstractions that are used by ASC applications to run on all ASC platforms without code rewrites. RAJA provides C++ kernel abstractions for computation on central processing units (CPUs) or GPUs, CHAI provides a managed array abstraction that automatically copies data between CPU and GPU memory based on RAJA execution contexts, and Umpire provides portable resource management focused on facilitating GPU memory allocation sharing in integrated applications. This project supports the development of new features requested by ASC applications in support of ASC milestones and user requests, as well as developing the necessary performance-portable software infrastructure for the ATDM application project.

Accomplishments in FY21:

- Supported ASC applications during "IC usage of PEM libraries" L2 milestone (Appendix I, target ATDM 1.d).
- Continued to support integration of Umpire into ASC application codes & libraries including Cheetah (Appendix I, target IC-1).
- Developed and delivered inter-process shared memory capability in Umpire to allow multiple application processes to share allocations.
- Worked with El Capitan Center of Excellence (COE) partners to report and resolve missing features required for RAJA programming model (Appendix I, target CSSE 1.b).

Planned Activities in FY22:

- Develop and deliver CHAI capabilities to applications in support of ASC milestone and user requirements, including support for overlapping data transfers and enhanced data eviction capability. (Appendix I, target ATDM 1.d)
- Continue to develop and deliver Umpire capabilities to applications in support of ASC milestone and user requirements, including improved monitoring and error identification, and continuing to expand functionality available inside GPU kernels. (Appendix I, target ATDM 1.d)
- Continue interactions with El Capitan Center of Excellence (COE) partners to resolve performance and correctness issues identified on next iteration of early access hardware. (Appendix I, target CSSE 1.b)

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.

DevRAMP is funded by both ATDM and CSSE. Community-targeted effort is reported here, and ASC-targeted effort is reported in the CSSE section.

Accomplishments in FY21:

- Hardened many aspects of Spack binary package management. Enabled secure build-testing of pull requests for Spack for the first time; Improved performance of build caches and binary relocation.
- Rolled out first release of Spack (0.16) including a new Answer Set Programming (ASP)-based concretizer and made numerous improvements and bugfixes as users tested it in production. Implemented binary bootstrapping for the new concretizer and GPG security in preparation for making the new concretizer default in 0.17.

Planned Activities in FY22:

- Roll out release of Spack including compilers and multi-compiler runtime libraries as 1^{st}-class dependencies in the package model in preparation for El Capitan. (Appendix I, target CSSE-2.b)
- Develop new features that simplify complex packages as the ecosystem continues to become more complex. Specifically target bundled packages (Intel / NVIDIA / AMD

/ HPE compilers and scientific libraries and the interfaces they provide; large packages like Trilinos and Axom). Refactor, simplify packages using new features.

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. 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 thousands of downloads per year from over 100 countries. It is freely available at mfem.org and on GitHub.

Accomplishments in FY21:

- Completed major software releases: MFEM-4.2, MFEM-4.3, GLVis-4.0 and GLVis-4.1. (Appendix I, target ATDM 1.d)
- Developed ATDM-specific GPU tools, algorithms and optimizations in support of the MARBL ATDM L1 Milestone, which successfully demonstrated large-scale GPU performance on ALE multi-physics problems. (Appendix I, target ATDM 1.d)
- Implemented ALE discretization improvements for compressible flow in BLAST, including significant improvements in transfer between high-order and low-order-refined simulations, dynamic Adaptive Mesh Refinement (AMR), finite element visualization and mesh optimization methods. (Appendix I, target ATDM 1.d)

Planned Activities in FY22:

- Develop adaptive discretizations for high-order finite element ATDM applications, including algorithms such as dynamic AMR, high-order/low-order-refined transfer and solvers, mesh optimization, finite element visualization, and hpr-adaptivity. (Appendix I, target ATDM 1.d)
- Support MFEM-based ATDM applications in their transition to exascale hardware, including both NVIDIA and AMD GPUs. Leverage the Exascale Computing Project (ECP)/Center for Efficient Exascale Discretizations (CEED) project. (Appendix I, target ATDM 1.d)
- Continue engagement with ATDM application work, develop mini-apps, and provide support for MARBL L2 milestones. (Appendix I, target ATDM 1.d)

Cross-Cutting Extreme-Scale Research (LANL)

This includes research in visualization/analytics and workflows. The work will be coordinated with the related CSSE projects to leverage and extend the expertise, focusing on enabling the ATDM CDA project and use in extreme-scale environments. The project areas include in situ data analytics for large-scale simulations; and 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 Visualization Toolkit (VTK) algorithms for use with ParaView, Cinema, and the VTK-m project.

This NNSA project encompasses the following ECP projects:

2.3.5.xx: Cinema – In Situ Visualization

Accomplishments in FY21:

- Fusion release, capability used in section of VnV L2 Milestone: HED Hydrodynamics in the Common Modeling Framework
- VR: Demo of new Vive Secure visor with three drivers (ParaView, Ensight and SciVista).
- Release of deep learning for image databases from large-scale scientific data (Cinema), using simulation and vis parameters

Planned Activities in FY22:

- Annual release of integrated Cinema toolkit modules and the annual release of Pantheon workflows
- Release of Fusion 1.5. This includes additional capability modules developed in collaboration with LANL scientists and developers, to be integrated into CMF.
- Development and delivery of VR systems and software in support of LANL use cases, with an emphasis on V&V integrated media training.

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:

2.3.5.xx: Legion – Asynchronous data-centric task-based programming model

2.3.5.xx: ATDM Tools – LLVM compiler and runtime tools and QUO thread-level heterogeneity in coupled Message Passing Interface (MPI) applications

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. The co-design area has been funded in CSSE Next-generation Computing Technologies (NGCT).

Accomplishments in FY21:

- Productivity boosts: **2-4x** in performance on Kokkos-aware compilation (***no code change***), **faster** compile times, much **smaller** executables (up to 12x)
- Kitsune+Tapir: explicit awareness of parallelism within the "heart" of the compiler (vs. front-end centric approaches: Kokkos, RAJA, OpenMP)
- Legion Co-designed with Ristra (FleCSI) using control replication to enable SPMD workloads to scale effectively using hierarchical task parallelism

Planned Activities in FY22:

- Parallel reduction support (e.g., Kokkos parallel_reduce)
- Focused support for AMD and NVIDIA architectures -- GPUs and CPUs (e.g., X86 and Arm). Scope includes custom analysis and optimizations, runtime system work, and vendor outreach and engagement.
- Support for Kokkos code generation that matches LANL's use cases/requirements (e.g., multi-dimensional ranges, reductions, etc.). Including use cases wrapped in FleCSI and other frameworks (e.g., Parthenon).

Future-Generation Computing Technologies (LANL)

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

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., Charliecloud, Singularity, Simple Linux Utility for Resource Management (SLURM), LSF, etc.).

This NNSA project encompasses this ECP project:

2.3.5.02: BEE/Charliecloud Containers and Container Orchestration for HPC systems

Accomplishments in FY21:

- BEE can now archive, clone, re-run workflows.
- Charliecloud fully unprivileged container building on HPC systems

Planned Activities in FY22:

- Evaluate and finalize the LANL container registry solution
- Deploy GitLab continuous integration (CI) runners on production HPC systems enabling container building and testing using modern DevOps practices

KOKKOS (SNL)

The goal of the KOKKOS Ecosystem is to provide performance portability for ASC application codes, allowing these codes to obtained good performance on several different next generation computing platforms (CPUs, GPUs, etc.). This projects funds two main components of the KOKKOS Ecosystem, the KOKKOS performance portability library and the KOKKOS KERNELS performance portable math library.

The KOKKOS performance portability library work described here focuses on four categories of work and is aligned with the KOKKOS strategic plan. Category I: Continuous Effort of Application Support & KOKKOS Maintenance; Category II: Support for new hardware architectures, potentially through new Backends; Category III: Optimization of Existing Backends for new Compilers and new iterations of Hardware; Category IV: C++ Standards Work and KOKKOS Community Growth. The KOKKOS KERNELS library work focuses on providing performance-portable sparse/dense linear algebra and graph kernels that utilize the hierarchical memory subsystem expected in current and future HPC architectures. It also focuses on vendor interaction, especially in the context of vendor math libraries.

In terms of ND mission impact, the KOKKOS Ecosystem is now a key component for most of the ASC IC and ATDM application codes. The majority of these codes depend on KOKKOS for obtaining good performance on the variety of next generation architectures. Thus, the KOKKOS Ecosystem indirectly impacts any application of these codes that are run on next generation platforms in the ND mission.

Accomplishments in FY21

- Developed a portable Maximal Distance-2 Independent Set (MIS-2) algorithm (an important graph algorithm for scientific computing) algorithm using KOKKOS, which is 4-15x faster than state-of-the-art for GPUs. The new implementation has been run on NVIDIA/AMD GPUs and ARM/Intel CPUS. It has been integrated with the new algorithm with MueLu multigrid solver and should improve multigrid coarsening on GPUs. (Appendix I, target ATDM-1)

- Added KOKKOS Graph API to help with latency limited execution cases. The feature allows the creation of a dependency graph object between multiple kernels, which can be repeatedly submitted for execution. KOKKOS Graph can leverage the CUDA Graphs API introduced in CUDA 11. (Appendix I, target ATDM-1)

- Improved KOKKOS KERNELS sparse matrix addition for sorted rows on GPU, now ~50% faster overall than NVIDIA cuSPARSE. (Appendix I, target ATDM-1)
- Optimized memory access patterns in KOKKOS KERNELS dense matrix-vector multiplication. In one SIERRA-Thermal Fluids problem on ATS-2, this sped up the overall radiosity solve by 15%. (Appendix I, target ATDM-1)
- Rewrote KOKKOS KERNELS Symmetric Gauss Seidel to support sparse matrices with dense rows such as those appear in SIERRA-Thermal Fluids problems. This supports the ASC level 2 milestone results for SIERRA-TF. The resulting hierarchical parallel algorithm for dense rows results in 6-8x speedup for SIERRA-TF use cases. (Appendix I, target ATDM-1)

Planned Activities in FY22:

- KOKKOS: Provide KOKKOS Training upon request to ASC customers, if the available online recorded tutorials are insufficient. (Appendix I, target ATDM-1)
- KOKKOS / KOKKOS KERNELS: Work with TRILINOS Solvers team to improve performance on AMD GPUs and Intel GPUs. (Appendix I, target ATDM-1)
- KOKKOS: Work with ASC applications to evaluate and improve performance on AMD GPUs. (Appendix I, target ATDM-1)
- KOKKOS KERNELS: Support ATDM and ASC apps for better performance on El Capitan early access platforms. (Appendix I, target ATDM-1)

Scalable Visualization (SNL)

This project provides the necessary R&D to support the data analysis and visualization needs of ATDM applications on ASC advanced technology systems. This work provides development, integration, and deployment of in situ analysis and visualization capabilities as well as the testing, tuning, and deployment of many-core algorithms for visualization to support the direct needs of ATDM applications on the ASC/AT systems. The visualization team will also work closely with ATDM application teams to ensure high-quality analysis and visualization of large-scale datasets generated by our principal applications.

The R&D 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 ECP.

Accomplishments in FY21:

- Developed an implementation plan to enable Catalyst Python development from Python Integrated Development Environments (IDE). The team demonstrated using the Eclipse and Spyder IDE. (Appendix I, target ATDM-1)
- Unified Phactori scripting language for configuration of in situ visualization in support of ASC applications. Deployed with SPARC, SIERRA, and Nalu.
- Updated and deployed Catalyst 5.9.0 in SPARC, Nalu, and SIERRA through IOSS. Catalyst 5.9.0 unified the Python scripting interface between ParaView, Catalyst, and other ParaView tools. (Appendix I, target ATDM-1)
- Began to explore data-management challenges for HPC simulation data sets in cloud services. Completed a prototype to ingest Exodus simulation results into Sandia's DataSEA tool. (Appendix I, target ATDM-1)

Planned Activities in FY22:

- VTK-m: Address performance/scaling and memory overhead issues for in situ analysis and visualization. FY22 will focus on design improvements to the SEACAS IOSS interface to enable references to blocks of simulation memory rather than using temporary memory buffers. (Appendix I, target ATDM-1)
- VTK-m: Continue to evolve the Phactori scripting language for production in situ visualization. (Appendix I, target ATDM-1)
- VTK-m: Identify and ameliorate Sandia analyst barriers to utility, usability, and adoption of ParaView, Catalyst, and VTK-m. Demonstrate a simplified prototype as part of an embedded components L2 milestone. (Appendix I, target ATDM-1)
- Develop prototypes to demonstrate integration of HPC data and cloud services, including assessments of performance and limitations in our current HPC computing environments. (Appendix I, target ATDM-1)

Operating Systems & On-Node Runtimes (SNL)

This project 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 and preparing for next-generation ASC platforms. Priorities in this area include the development of lightweight tasking techniques in the runtime layer, interfaces between the runtime and OS for management of critical resources (including multi-level memory, non-volatile memory, and network interfaces), and resource isolation strategies at the operating system level, e.g., containers, that maintain scalability and performance while providing a more full-featured set of system services. The OS/R technologies developed by this project will be evaluated in the context of ATDM application codes running at large-scale on ASC platforms. Through close collaboration with vendors and the broader community, e.g., through participation in OpenMP and MPI standards participation, the intention is to drive the technologies developed by this project into vendor-supported system software stacks and gain wide adoption throughout the HPC community.

Accomplishments in FY21:

- Completed an analysis of the impact of application perturbation from container infrastructure. Results show that containers running without contention for CPU resources have little impact on application perturbation; however, using Linux control groups to partition resources between concurrent containers can lead to noise-like events that have the potential to significantly degrade application performance. (Appendix I, target ATDM-1)

Planned Activities in FY22:

- MPI communication and GPUs: Characterize how ASC/ATDM workloads use MPI communication in conjunction with GPUs to understand workload performance and opportunities for optimization (Appendix I, target ATDM-1)
- Container and Infrastructure Support: Execute on the container infrastructure and support strategy developed in FY21, continue to partner with vendors, facilities, and application/library developers to leverage container technologies for ASC/ATDM workloads (Appendix I, target ATDM-1)
- OpenMP and MPI Standards advancement: Work through standards committees to shape the direction of the OpenMP and MPI parallel programming models to provide needed capabilities for ASC/ATDM workloads. (Appendix I, target ATDM-1)

Next Generation Contact Modeling (SNL)

Next-Generation Contact Modeling effort applies the Darma-vt library for asynchronous many-task (AMT) scheduling and the KOKKOS library for performance portability to the simulation of contact mechanics on large-scale, heterogeneous computing platforms. The computational cost of simulating contact is a long-standing issue for a wide class of problems relevant to Sandia's missions, and often comprises a significant portion of the overall run time for mechanics simulations. The use of AMT is promising in this context because the optimal load balancing and task management for contact evolve dynamically of the course of a simulation. Darma-vt, used in conjunction with KOKKOS for utilization of on-node accelerators such as GPUs, has the potential to substantially reduce turnaround times for solid mechanics simulations involving contact. This project will evaluate prototype implementations of contact using AMT and provide a roadmap for future implementations in Sandia's production codes.

Accomplishments in FY21

- Comparison of bulk synchronous and Darma-vt asynchronous many-task scheduling (AMT) programming models for contact mechanics in the NimbleSM code. (Appendix I, target ATDM-1)

- Implementation of an abstraction layer that allows for run-time selection of solely MPI or AMT contact strategies in NimbleSM for improved performance and evaluation of parallelization strategies. (Appendix I, target ATDM-1)

Planned Activities in FY22

- This project concludes at the end of FY21.

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 – Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) (LLNL)

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 Collaboration comprises three projects that aim to improve understanding of cancer biology and its application to more effective therapies, of which ASC is funding Pilot 2 which produces a multiscale, multiresolution molecular dynamics simulation framework for CORAL-class machines to study the dynamics of mutated proteins in cancerous cells.

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 FY21:

- Improved UQ models for understanding uncertainties in multiscale models.
- Applied Multiscale Machine-Learned Modeling Infrastructure (MuMMI) to extended RAS interactions and additional protein dynamics to be selected.

Planned Activities in FY22:

- Generalize MuMMI to support multiple resolution models
- Demonstrate accelerated sampling using a Boltzmann generator
- Characterize engagement and dimerization of the RAF kinase domains.

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

DOE/NNSA, the National Cancer Institute, the 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 Cooperative Research and Development Agreement (CRADA) that 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.

Accomplishments in FY21:

- Integrated with automated chemical synthesis and assays to implement an active learning approach that adaptively estimates and reduces model uncertainties as the optimization process progresses.
- Applied ATOM design process to infectious disease and in nonbiological molecular design activities.
- Achieved 501 3(c) status for the Consortium.

Planned Activities in FY22:

- Extend the ATOM molecular design platform for increased chemical design space application and improved translation to human application.
- Apply the molecular design platform to probe molecule design for ATOM member groups, academic collaborators, and biotech collaborators.
- Build the R&D community around ATOM data, models, and platforms.

Advanced Machine Learning (LLNL)

ASC's AML initiative, funded in all three ATDM product groups, aims to integrate modern machine learning techniques that have been used to tackle challenging problems across domains into ASC efforts spanning simulation and computing. To enable success in this area, a portion of this effort includes collaborations with academia, industry, and other government agencies. The AML initiative supports core NNSA applications through application and V&V research. Hardware exploration encompasses near-term hardware advances in neural networks and more speculative hardware, such as neuromorphic technologies.

Accomplishments in FY21:

- Developed an ML solution for automated analysis of radiographs for a specific use case. (IC-5)
- Developed an initial ML surrogate models that can predict simulation behavior per timestep. (IC-2)
- Created a workflow to enable end users to automatically train a neural network for automatic ALE. This workflow has been tested on several test problems and an initial release was completed. An initial 3D version has been implemented. (IC-3)
- Ported ML training for ALE to LBANN for large-scale training. Scaling of data preprocessing and training has been explored and improved. (IC-1)

Planned Activities in FY22:

- Test the generalizability of ML-driven ALE models. Pair this with the existing workflow to allow augmented training when needed for specific problems. (IC-4)
- Demonstrate ML driven ALE in large scale, programmatically relevant calculations. (IC-4)
- Demonstrate surrogate modeling solutions to a programmatically relevant multi-scale calculation. (IC-4)
- Demonstrate the use of the neural network for experimental radiographs. (IC-4)

Advanced Machine Learning (LANL)

The AML project has mission-focused efforts in collaboration with PEM, V&V, and ATDM-CDA. The project conducts studies of Artificial Intelligence (AI), ML, and Deep Learning methodologies applied to multiscale simulation, materials modeling, nuclear data evaluation, turbulence simulations, and radiographs analysis and data-driven approaches that can explain underlying physics. The ASC program is well positioned to leverage investments in experimental facilities, next-generation computer architectures, algorithm development, and simulation data collection initiatives like LSCI to develop ML workflows to utilize multisource, multifidelity data for answering mission-relevant questions from these areas. LANL's CSSE, V&V, PEM, and IC subprograms co-fund these mission-focused research projects.

This project is co-funded in CSSE Next-Generation Computing Technologies (WBS 1.2.3.5.7) and the accomplishments and planned activities are reported there.

Advanced Machine Learning (SNL)

This project covers a broad set of topics that includes, but is not limited to, the use of ML to detect patterns and predict behavior in scientific data; the use of ML for adaptive numerical solvers; the use of ML to understand and improve large-scale system behavior;

the use of ML to automate geometry and mesh design for complex structures; and R&D to understand implications on the use of ML with respect to data correctness, application performance, and various uncertainties that impact decision making. We are focusing initial projects on addressing gaps in physics-constrained ML, ML algorithms in sparse data regimes, verification, validation, and uncertainty quantification for ML, and learning hardware systems in an HPC environment. Some of the activities in this project are also funded by the ATDM ICD and CDA product groups.

Accomplishments in FY21:

- Developed a PINNs (Physics-Informed Neural Networks) implementation able to calibrate machine-learned material models that, in some cases, are not possible to approximate with finite-element methods. (Appendix I, PEM-4)
- Developed time-series classification algorithms for stronglink aging problem that identified active/passive time windows in the training dataset based on device reaction to stress tests. (Appendix I, PEM-4)
- Completed a python-driven workflow to develop finite-element models for porous materials for the mass-production of training data for ML algorithms. (Appendix I, PEM-4)
- Created an extension to the XYCE circuit simulator with the ability to call data-driven device models entirely defined in Python, enabling the ability to utilize recent advancements in machine learning using popular frameworks such as TensorFlow and PyTorch. (Appendix I, PEM-4)

Planned Activities in FY22:

- Develop and expand capabilities in physics-informed machine learning to address shortcomings in material modeling in support of P&EM. (Appendix I, PEM-4)
- Develop methods to rigorously assess machine learning models for data verification and model qualification in support of V&V. (Appendix I, V&V-4)
- Develop interpretable machine learning methods to modernize models used to represent failures in engineering alloys in support of IC. (Appendix I, V&V-4)
- Develop machine learning models to identify causes of component failures in nuclear-weapon surveillance or manufacturing. (Appendix I, PEM-4)

Appendix F: Computational Systems and Software Environment Subprogram (WBS 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. Added to the subprogram starting in FY21 is the Next-Generation Computing Technologies product group, which represents transferred computer science technologies from ATDM ASD.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2020, and through quarter 3, fiscal year 2021, are reflected below for the CSSE subprogram.

- Negotiated modifications to El Capitan build and CORAL-2 NRE contracts. Deployed the second El Capitan early access system to enable tri-lab system software preparations for El Capitan. (Appendix I, target CSSE-1.b) (LLNL)
- Led the Tri-Laboratory process for the next commodity technology systems (CTS-2) procurement, reviewed RFP responses, selected CTS-2 vendor partner, completed negotiations, and awarded CTS-2 contract. (Appendix I, target CSSE-1.c) (LLNL)
- Developed and deployed a spack-based application development environment on rznevada, an El Capitan early access system. (Appendix I, targets CSSE-1.b and CSSE-2.b) (LLNL)
- Improved the portability and interoperability of Flux-enabled workflows by integrating Flux with a wider range of workflow management software systems, which include Maestro, Merlin, and UQP/Themes, Parsl, Swift/T, Radical Pilot as well as domain specific workflow tools such as ATOM's GMD pipeline. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d) (LLNL)
- Assisted the Cheetah team with porting and optimization for Sierra. Demonstrated greater than 10x speedup (Sierra node vs. CTS node) on problems of programmatic interest. (Appendix I, target ATDM-1.d) (LLNL)

- Highly accurate interatomic potentials have been developed for elemental metals including aluminum, tin, gold, and silver. Models for alloys and electron materials are in progress. (LANL)
- Developed AM data visualizer and operator classifier. Allows quick visualization and analysis of splatter as well as labeling by expert needed for ML classification. (LANL)
- Enhanced ARM emulator (ARMie) and used to estimate theoretical vectorizability of several IC applications and delivered report of findings. (LANL)
- Ristra Apps ran at scale on Sierra (P9+Volta), Trinity (Haswell), and Astra (ARM). (LANL)
- FleCSI 2.0 delivered and training class held. (LANL)
- The memory innovation collaborations with Micron have been coordinated with the program's broader advanced architectures efforts including the multi-agency Project 38 innovative architectures activity leading to exploration of new processing-near-memory designs intended to increase the performance of future mission applications and enhance national competitiveness in the high-performance computing area. (Appendix I, target CSSE-1.d) (SNL)
- The first Advanced Architectures Prototype System (AAPS), Vanguard-Astra, was successfully utilized for the ATDM L1 milestone. The 3 NNSA laboratories ran weapons codes on the platform at scale demonstrating the viability of Arm-based architectures for the nuclear deterrence mission. (Appendix I, targets CSSE-1.d and ATDM-1 (SNL)
- Version 11.0 of the Structural Simulation Toolkit (SST) was released with enhanced capability to model future architectural features of future high-performance computing systems. This tool has been utilized to evaluate the performance of new networks, accelerators, and memory sub-systems for the Exascale platforms and beyond. (Appendix I, target CSSE-1.d) (SNL)
- Developed a new capability in Sandia's Next Generation Workflow to enable automated ensemble analysis from computational simulation workflows in support of W80-4 and W87-1 programs. (Appendix I, target CSSE-2.d) (SNL)
- The first of the Intel Loihi neuromorphic architecture testbeds was delivered. Two mini-apps, a random walk (RW) and sparse coding (LCA) mini-app have been implemented and tested and initial performance and scaling studies have been performed on the Loihi system. (Appendix I, target CSSE-3.b) (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): TOSS4 validation on CTS-2		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 06/30/22		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-2.d		
Description: This milestone will port the TOSS4 operating system to the upcoming CTS-2 hardware. While TOSS4 was designed to support likely candidates for CTS-2 hardware, it has only been validated against the currently available CTS-1 hardware. As production CTS-2 hardware becomes available TOSS4 will be ported, tested, and validated as reliable and performant on that hardware. A functional TOSS4 will be deployed along with CTS-2 hardware and validated by ASC users.		
Completion Criteria: The SWL test suite validates TOSS4 on available CTS-2 hardware and an ASC user validates a working CTS-2/TOSS4 environment.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: CSSE and CCE supported staff		

Milestone (ID#TBD): Crossroads System Integration Readiness	
Fiscal Year: FY22	
Completion Date: 06/30/22	
ASC WBS Subprogram: CSSE, FOUS	
Participating Sites: LANL, SNL	
Participating Programs/Campaigns: ASC	
Program Plan Target: Appendix I, target CSSE-1.a, FOUS-1	
Description: Prepare Crossroads Phase 1 for integration into the classified computing environment. Deliver and install system hardware for the initial phase of the Crossroads system. Deliver, test, and demonstrate system software. Complete onsite benchmarks and application performance testing. Prepare Crossroads for onsite integration into the local and remote computing infrastructure, including the user software environment.	
Completion Criteria: Follows the ASC Level 2 milestone criteria for Advanced Technology systems: hardware deliveries from vendor to site are complete, including the basic hardware to integrate "the system" as contractually defined; installation of the system by the vendor onsite to the extent that is contractually required is substantially complete; in general, contractual requirements for formal hardware acceptance have been substantially completed; system software needed for basic operation of the system is delivered, tested, and demonstrated to be operational; vendor has completed onsite benchmark and application performance testing and demonstration; and system is ready to begin onsite integration into local computing environment.	
Customer: NNSA/ASC HQ, tri-lab ASC program managers responsible for ATCCs, SSP, tri-lab weapons applications community	
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.	
Supporting Resources: CSSE, FOUS, ACES	

Milestone (ID#TBD): Evaluation of Data Flow HW Accelerators for Mission Relevant Mini-Apps and Kernels		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2022		
ASC WBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-1.d		
Description: This milestone will provide an initial evaluation of NextSilicon-based data flow accelerators for NNSA-relevant mini-applications and kernels. The analysis is conducted as an exploratory phase for potential technologies to be deployed in the Vanguard Advanced Architectures Prototype System project. The milestone will include an evaluation of the NextSilicon Rev 1.0 Maverick hardware devices, associated hardware simulator as well as the initial software stack and compiler toolchain. This study will provide performance results for at least three microbenchmarks covering compute and memory performance and at least three NNSA-relevant mini-applications or benchmarks. A prototype ATSE-based container will be developed to include the compiler toolchain for NextSilicon devices and standard ATSE libraries and tools. Successful completion of the milestone will include providing evidence of mini-application/benchmark runs and performance on both the NextSilicon Maverick simulator and hardware devices with comparison to at least one other NNSA relevant processor technology. Documentation of best-practices for writing efficient kernels/code and compiling to target the NextSilicon data flow accelerator will be also provided.		
Completion Criteria: 1. Documentation which captures evaluation of at least 3 mini-apps or kernels on NextSilicon simulator and early HW based on availability in FY22. 2. Documentation which captures best-practices for code development and compilation targeting NextSilicon data flow accelerators.		
Customer: ATDM, IC		
Milestone Certification Method: A program review is conducted, and its results are documented. Professional documentation consisting of a report is prepared as a record of milestone completion.		
Supporting Resources: FOUS		

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 FY21:

- Led and managed the CTS-1 contract and FY21 system procurements. (Appendix I, target CSSE-1.c)
- Extensively tested and resolved issues on production CT systems.
- Assisted users of CT systems.
- Led the Tri-Laboratory process for the next commodity technology systems (CTS-2) procurement, reviewed RFP responses, selected CTS-2 vendor partner, completed negotiations, and awarded CTS-2 contract. (Appendix I, target CSSE-1.c)

Planned Activities in FY22:

- Procure and field CTS-2 systems. (Appendix I, target CSSE-1.c)
- Provide production support for CT systems (CTS-1 and CTS-2).

Commodity Systems Planning (LANL)

The scope of the Commodity Systems Planning project is to support the design, acquisition, and delivery 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 FY21:

- Provided programmatic oversight for CT systems.
- Coordinated requirements and participated in the CTS-2 evaluation and award activities. (Appendix I, target CSSE-1.c, FOUS-5)

Planned Activities in FY22:

- Continue to provide programmatic oversight of CT systems.
- Coordinate the delivery and installation activities for CTS-2 systems (Appendix I, target CSSE-1.c, FOUS-5)
- Coordinate the activities for the removal of CTS-1 systems from service.

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 FY21:

- Participated in the CTS-2 proposal contract negotiations, architect decision points, and initial planning for site integration. (Appendix I, target CSSE-1.c)
- Participated in the deployment of the CTS-1+ cluster "Manzano" to the SRN. (Appendix I, target CSSE-1.c)

Planned Activities in FY22:

- Participate in the configuration, delivery and integration of the Commodity Technology System (CTS-2) cluster in support of FOUS milestone ID#TBD. (Appendix I, target CSSE-1.c)

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.

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 is 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 run best on this machine. Running in the classified environment, Sierra is used as a tri-lab resource for tri-lab stockpile stewardship milestones via the Advanced Technology Computing Campaign (ATCC) process, which will be run every six months when the next suite of codes is ushered onto the machine.

Accomplishments in FY21:

- Completed two ATCC processes.
- Updated Sierra system in accordance with security policies.

Planned Activities in FY22:

- Run two additional ATCC processes.
- Upgrade Sierra operating system to Red Hat Enterprise Linux 8.
- Continue to update Sierra in accordance with security policies.

El Capitan Tri-lab Advanced Technology System (LLNL)

In March 2018, the CORAL-2 RFP was released. LLNL co-organized the evaluation of responses in mid-2018. LLNL has signed build and NRE contracts with Cray, Inc., a technology provider that was subsequently acquired by HPE. The NRE contract is enabling key technologies related to the selected system architecture, while LLNL will procure El Capitan, a next-generation supercomputer, that will be delivered and accepted in FY23 through the build contract. Under the El Capitan procurement, LLNL worked with key technology providers to deploy an El Capitan system of significantly over 1 exaflops peak capability to advance science and to ensure national security; the total system capability will be determined through a technical decision point in CY22 that finalizes late binding decisions on the node architecture and technology development related to that decision. LLNL has worked with Cray (now HPE) to complete the late binding decision in Q1 FY20 and the associated contract modification in Q4 FY20. This decision substantially increased the expected delivered performance for mission critical applications over the original plan of record (POR). Subsequently, LLNL has worked with HPE and AMD, the processor partner that was selected in that decision, to refine the node architecture further. An associated contract modification that enhanced the planned usability of the delivered system and again increased the expected delivered performance significantly was completed in Q3 FY21. El Capitan 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. Once this platform has transitioned to the classified environment, it will be 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 FY21:

- Provided technical coordination and contractual management for CORAL-2 NRE and El Capitan contracts. (Appendix I, target CSSE-1.b)
- Continued to evaluate alternatives for node architecture and refined late-binding decision for El Capitan node architecture to reflect technical directions of partners. (Appendix I, target CSSE-1.b)
- Negotiated modifications to El Capitan build and CORAL-2 NRE contracts and submitted them for approval. (Appendix I, target CSSE-1.b)
- Continued application preparations for the El Capitan system through the COE. (Appendix I, target CSSE-1.b)
- Deployed the second El Capitan early access system to enable tri-lab system software preparations for El Capitan. (Appendix I, target CSSE-1.b)

Planned Activities in FY22:

- Continue to provide technical coordination and contractual management for CORAL-2 NRE and El Capitan contracts. (Appendix I, target CSSE-1.b)

- Continue application preparations for the El Capitan system through the COE. (Appendix I, target CSSE-1.b)
- Deploy multiple instances of the third El Capitan early access system to help tri-lab teams prepare for El Capitan. (Appendix I, target CSSE-1.b)

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, R&D efforts and strategic planning for supporting infrastructure. In FY22, this project will continue to focus on the project management, acquisition and deployment oversight of the ASC Crossroads system and accelerated hardware research contracts. The focus in this project also includes the execution of the Crossroads project under the Defense Programs' Program Execution Instruction structure.

Accomplishments in FY21:

- Provided formal project status reporting as required for Crossroads. (Appendix I, target CSSE-1.a)
- Provided ongoing risk assessment and mitigation for the Crossroads project. (Appendix I, target CSSE-1.a)

Planned Activities in FY22:

- Provide formal project status reporting as required for Crossroads. (Appendix I, target CSSE-1.a)
- Provide oversight for the installation and acceptance of the Crossroads system (Appendix I, target CSSE-1.a, FOUS-1)
- Develop programmatic requirements and drives for ATS-5.

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 have continued the ACES partnership to acquire an HPC system in the 2022 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 FY21:

- Completed the Technology Decision Points (i.e., late-binding decisions) for the Crossroads system. (Appendix I, target CSSE-1.a) (LANL)
- Development of Crossroads project Statement of Work (SOW) items with selected vendor (HPE and Intel) (Appendix I, target CSSE-1.a) (SNL)
- Establishment of Crossroads Center of Excellent (COE) launching work to support application and library porting (Appendix I, target CSSE-2.a) (SNL)

Planned Activities in FY22:

- Conduct quarterly reviews report milestone progress towards contractual deliverables. (Appendix I, target CSSE-1.a) (LANL)
- Complete the Crossroads System Integration Readiness milestone. (Appendix I, target CSSE-1.a, FOUS-1) (LANL)
- Development of application/library porting/optimization best-practices for Crossroads system (Appendix I, target CSSE-2.a) (SNL)
- Application and benchmark performance acceptance for Crossroads Phase-I installation (Appendix I, CSSE-2.a) (SNL)
- Support acceptance and benchmarking preparation for full-system installation in early FY23 (Appendix I, CSSE-1.a) (SNL)

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 FY21:

- Analyzed for ECP Hardware Evaluation and end-of-contract ECP PathForward deliverable evaluations. (Appendix I, target CSSE-1.b)
- Coordinated Project 38 activities with the Memory Innovation project and ATSE activities including Arm contracts. (Appendix I, target CSSE-1.b)

Planned Activities in FY22:

- Coordination of Project 38 activities with the new Memory Innovation project and ATSE activities including vendor contracts. (Appendix I, target CSSE-1.b)
- Identification and analysis of existing and emerging ASC application and user requirements for ATS-5 and beyond (Appendix I, target CSSE-1)
- Continue industry R&D engagement including ML/AI and traditional HPC system vendors (Appendix I, targets CSSE-1 and CSSE-3)

Advanced Architecture Test Bed Research and Development (SNL)

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

Accomplishments in FY21:

- Deployed next-generation AMD GPU systems (including MI50 and MI100 GPUs) for KOKKOS programming model and algorithm development/testing (Appendix I, target CSSE-1.b)
- Procured FPGA-based system to evaluate optical-switch based interconnects and specialized accelerators (Appendix I, target CSSE-1.c)
- Procured GraphCore ML/AI-based testbed for application benchmarking and code development in support of the program's advanced machine learning initiative. (Appendix I, target CSSE-3.c)

Planned Activities in FY22:

- Evaluate and prioritize procurement of pre-release hardware from component/system vendors based on NNSA program goals (Appendix I, target CSSE-1.d)
- Participate in early prototyping and evaluation of next-generation systems (ATS3/ATS4 and CTS2) (Appendix I, targets CSSE-1 and CSSE-2)

Collaboration for Innovation in High Performance Memory Systems (SNL)

This project focuses on a critical bottleneck for mission applications: data movement in the computer's memory system or the "memory wall." Much of the computation in our mission application is performance bound by latencies, bandwidth or limitations in concurrency in these memory systems. The NNSA ASC program is currently collaborating with the primary U.S. memory vendor, Micron Technology, to conduct research and development into innovative new memory systems which show promise in breaking through this "memory wall" and significantly increasing the performance of mission applications. The NNSA laboratories will supply mission workload information and determine the impact of potential memory designs. If successful, the U.S. government will gain access to a new differentiating computing technology which creates a competitive advantage for applications critical to national security.

Accomplishments in FY21:

- Collaborated with Micron on development of programming environments for processing-in-memory accelerators and evaluation of performance improvements from processing-in-memory acceleration primitives (Appendix I, target CSSE-1.d)
- Coordinated NNSA lab, Science lab and multi-agency Project 38 activities with the Micron contracts at SNL and PNNL (Appendix I, target CSSE-1.d)

Planned Activities in FY22:

- Evaluation of NNSA-relevant kernel performance utilizing the Micron processing-in-memory simulator (Appendix I, target CSSE-1.d)
- Manage the Micron memory innovation contracts and coordinate with Project 38 and PNNL on related activities (Appendix I, target CSSE-1.d)
- Execute on new memory innovation R&D contract fully engaging the NNSA labs and HQ to optimize the outcomes for the mission (Appendix I, target CSSE-1.d)

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, HPE, 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 FY21:

- Released TOSS 3.7 for TLCC2 and CTS-1/1+ systems. Provided as-needed support for TOSS 3 on Astra. Released TOSS 4.1 and 4.2 for CTS-1/1+ systems and Astra. (Appendix I, targets CSSE-2.c-d)
- Installed TOSS 4 on CORAL-2 EA systems, and integrated the Cray Programming Environment. (Appendix I, target CSSE-2.b)
- Updated Slurm to version 20.11. (Appendix I, targets CSSE-2.b-d)

Planned Activities in FY22:

- Deploy TOSS 4 onto production CTS-2 systems as they arrive and are integrated. (Appendix I, target CSSE-2.d)
- Deploy TOSS 4 onto CORAL-2 EAS-3, and continue extending TOSS to fully support HPE Shasta hardware. (Appendix I, target CSSE-2.b)
- Integrate Flux as an optional TOSS 4 system resource manager. (Appendix I, targets CSSE-2.b/d)
- Integrate tools and technologies such as Secure Boot, SELinux, and munge improvements to increase system surety and enable more flexible use of ASC platforms.

Applications Development Environment and Performance Team (LLNL)

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

Accomplishments in FY21:

- Initiated an applications development environment on the Rznevada El Capitan early access system. (Appendix I, targets CSSE-1.b and CSSE-2.b, L2 Milestone)
- Provided support and improvements for LLNL ASC production platforms including Sierra, CTS-1, and the new Magma compute cluster. (Appendix I, target IC-1)
- The SAP team assisted tri-lab IC teams and users in porting, tuning, and using the Sierra system in ATCC10/11 computing campaigns. (Appendix I, target IC-1)
- Developed tool infrastructures to improve scalability and performance of applications and the code development environment. (Appendix I, target IC-1)
- Supported NRE activities and working groups for the CORAL-2 exascale system preparations. (Appendix I, targets CSSE-1.b and CSSE-2.b)

Planned Activities in FY22:

- Provide and support an initial applications development environment on an El Capitan early access system. (Appendix I, targets CSSE-1.b and CSSE-2.b)
- Improve, test, evaluate, and provide user support for LLNL ASC production platforms including the Sierra system applications development environment. (Appendix I, target IC-1)

- The SAP team will assist tri-lab IC teams and users in porting, tuning, and using the Sierra system in computing campaigns. (Appendix I, target IC-1)
- Continue to improve tools and programming infrastructure to enhance scalability and performance of applications and the code development environment. (Appendix I, target IC-1)
- Support NRE activities and working groups for the CORAL-2 exascale system preparations. (Appendix I, targets CSSE-1.b and 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 applied data analytics and machine learning and an effort in hardware co-design.

The data analytics portion of the project involves both the development and simplification of deployable techniques for "Big Data" data analytics as well as the application of these techniques with targeted collaborations within the program. As such, this includes containerization of other turn-key solutions being developed by the project to deploy/tear down Spark, Hadoop, applied ML libraries, etc., in user space on HPC resources, as well as dedicated analytics infrastructures with streaming data ingest/analysis/visualization capabilities (e.g., the team's Tivan infrastructure). This effort started with application and systems data. But has expanded into supporting production manufacturing. The production work is in collaboration PEM.

Hardware co-design includes collaborations with vendors on hardware design, evaluation, and testing with ASC production simulation applications.

Accomplishments in FY21:

- Using LANL's Tivan big data analytics infrastructure, developed various prototype analyses and "dashboards" for the LAP code team.
- Developed AM data visualizer and operator classifier. Allows quick visualization and analysis of splatter as well as labeling by expert needed for ML classification
- Various analyses for TA-55 hemishell work: scripts, plots, analytics reports, analysis on impacts of variable CMM scan techniques

Planned Activities in FY22:

- ML on CMM hemishell data and radiographs for defect detection
- KCNSC collaboration on ML of melt pool, time over temperature threshold, and other factors.
- Shock & Vibe ML using 10s of TBs of sensor data with Tivan – expected use of Echo State Networks (regression) and CNNs for classification

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 supports collaborations with vendors on emerging advanced architecture hardware and experimental software stacks in a controlled local environment. This includes the analysis of future hardware improvements for ASC applications, including use of architectural simulators and performance tools. This work includes existing efforts with Project 38 and Micron and new collaborations with universities, vendors, and other agencies in hardware advances. Additionally, it also allows for CI from the IC production code teams on these emerging architectures.

Accomplishments in FY21:

- Integration of Sambanova ML hardware with Darwin
- Integration of El Cap early access "like" nodes
- Requirements gathered for ASC gitlab project (with FOUS) supporting all ASC codes with common software engineering, dev-ops, repos, and CI.
- Addition of Vast Filesystem

Installed A64FX testbeds (classified and unclassified) to facilitate hardware evaluation and application performance analysis.

Planned Activities in FY22:

- Addition of Site Crossroads early access hardware for crossroads, and El Capitan.
- Harden gitlab to support all IC workflows and transition all IC teams (Appendix I, target CSSE-1.a)
- Deploy NVidia Arm/A100 software development testbeds.

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 FY21:

- Successfully deployed the PowerAPI on the Fugaku system, supporting a collaboration with RIKEN. Sandia continues to lead the PowerAPI community effort and advance the specification toward improving energy efficiency for HPC systems. (Appendix I, target CSSE-2)
- A paper evaluating two turbo control algorithms on the Astra system was selected as Best Paper at the Cray User Group conference. (Appendix I, target CSSE-2)
- Completed an investigation of container orchestration services and frameworks for broader applicability to ASC applications and workloads. (Appendix I, target CSSE-2)
- Created an approach to effectively exploit persistent memory for an efficient exchange of data in a pipelined workflow. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Advance the state of power management and control on ASC platforms, working with the Power API community and strategic vendor partners. (Appendix I, target CSSE-2)
- Develop OS and runtime support for providing HPC-as-a-service usage models to key Sandia ASC customers. (Appendix I, target CSSE-2)
- Collaborate with FOUS on exploring ways to improve HPC data center efficiency. Prototype bi-directional exchange between the facility and HPC platforms with the initial use case of responding to emergency power events. (Appendix I, target CSSE-2)

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 FY21:

- Released SST 11.0 with new capabilities supporting the program's advanced architectures efforts (Appendix I, target CSSE-1)
- Deployment of new general accelerator interfaces ECP, Vanguard and Project 38 goals (Appendix I, target CSSE-1)
- Deployment of network congestion and network line degradation/failure analysis capabilities in support of ATS procurements and ECP (Appendix I, target CSSE-1)
- Developed new processor models and data flow accelerator models in support of next generation HW options for Vanguard and future ATS systems (Appendix, target CSSE-1)

Planned Activities in FY22:

- Release SST 12.0 with capability improvements to support the program's advanced architectures efforts and collaborations (Appendix I, target CSSE-1)
- Develop exploratory architecture concepts for NNSA-relevant HPC systems, identify capability gaps in SST and complete SST models based on these concepts allowing evaluation (Appendix I, target CSSE-1.c)
- Deploy SST user-space workload modeling capability as a standalone capability (Appendix I, target CSSE-1)
- Develop initial light weight thread-based modeling capability in SST taking greater advantage of new accelerator HW (Appendix I, target CSSE-1.c)

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 FY21:

- Completed a study outlining alternative applications of In-Network Compute Assistance (INCA) for increasing performance via attached network accelerators. (Appendix I, target CSSE-2)
- Advanced and evaluated the Remote Virtual Memory Access (RVMA) network communication paradigm. A paper introducing RVMA was presented at IPDPS and follow-up work with more robust simulation and higher fidelity analysis was also completed. (Appendix I, target CSSE-2)
- Continued to develop the Portals specification, considering major feature enhancements including short- and long-term proposals in cooperation with strategic vendor partners. (Appendix I, target CSSE-2)
- All three DOE Exascale systems are deploying HPE/Cray interconnects that are heavily influenced by Portals technology. (Appendix I, target CSSE-1.b)

Planned Activities in FY22:

- Continue to consider major feature enhancements for Portals, working towards a major version update in collaboration with strategic vendor partners. (Appendix I, target CSSE-2)
- Continue leadership in SmartNICs for HPC, engaging with vendor partners and exploring next-generation SmartNIC designs with a focus on enhancing system software offload possibilities. The long-term goal of this work is enabling production ASC applications to take full advantage of SmartNICs to enhance overall performance without requiring application-level modifications. (Appendix I, target CSSE-1.d)
- Evaluate the ability of RVMA to be a critical support technology for enabling extremely fast communications supporting graph-like GPU task descriptions. The long-term goal of this work is to improve the network performance of accelerator-enabled applications. (Appendix I, target CSSE-1.d)

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 FY21:

- Demonstrated efficient "failure masking" supporting scalable recovery from local hard failures (loss of processes) in a parallel solver, based on extension of the Fenix fault tolerance library. (Appendix I, target CSSE-1.b)

- Demonstrated the automatic and scalable checkpointing capability of resilient extension of KOKKOS with runtime-based and compiler-assisted approaches, applied to representative mini-applications including MiniMD, MiniFE and NinmbleSM (Appendix I, target CSSE-1.b)
- Developed a prototype of the failure emulation library that reflects the failure-data collected from Astra and Cielo supercomputing systems. This new library allows the users to study the mitigations of application program execution under a variety of unreliable computing scenarios. (Appendix I, target CSSE-1.b)

Planned Activities in FY22:

- Demonstrate integrated local error/failure mitigation for both silent errors and hard failures in a prototype parallel numerical solver leveraging the formal modeling and algorithm techniques to derive correctness under asynchronous execution scenarios. (Appendix I, target CSSE-1. b)
- Demonstrate the capability of resilient extension of KOKKOS for mitigating mixture of silent errors and hard failures with hybrid distributed asynchronous applications built with MPI and DARMA. (Appendix I, target CSSE-1.b/2.b)
- Demonstrate the capability of the failure emulation library to evaluate resilience overheads of several key MPI workloads with common failures. (Appendix I, target CSSE-1.b)
- Characterize the reliability of current neuromorphic devices and identify the gaps in current failure mitigation techniques. (Appendix I, target CSSE-3.c)

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 was transitioned as an ATDM project to CSSE in FY19.

Accomplishments in FY21:

- Performance tuned and hardened the DARMA software stack expanding test coverage to 83% with 100s of tests that automatically run with CI that test a wide range of compilers, including GNU/Clang/Intel/NVidia, based on application requirements. (Appendix I, target CSSE-2)
- Implemented support in EMPIRE that use DARMA's tasking and serialization libraries to perform full checkpoint/restarts to support large-scale runs of EMPIRE that require incremental checkpoint-to-disks. (Appendix I, target CSSE-2)
- Implemented new user-level threading (ULT) capabilities in DARMA tasking to support more efficient execution on the GPU so multiple CUDA streams can be used across over-decomposed DARMA tasks to effectively exploit concurrency on the GPU. (Appendix I, target CSSE-2.b)

- Supported development of the distBVH library which uses DARMA tasks to load balance SIERRA solid mechanics contact problems. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Continue to support and performance tune the DARMA software stack for EMPIRE's use case, focusing on improving performance on the GPU and new upcoming architectures. (Appendix I, target CSSE-2.b)
- Prototype using DARMA tasks to improve the load distribution during the matrix assembly stage of the GEMMA electromagnetics application to obtain speedup for target problems that incur imbalances. (Appendix I, target CSSE-2)
- Benchmark the performance benefit and tune distBVH load balancing use of DARMA tasking to improve the load balance for mission relevant contact problems. (Appendix I, target CSSE-2)

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 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 FY21:

- Developed/released HPSS 9.1 and 9.2; developing HPSS 9.3. (Appendix I, target CSSE-2.d)
- Continued driving the LC I/O Strategy: coupling archive with parallel file systems and analyzing HSM interfaces. (Appendix I, target CSSE-2.d)
- Procured unclassified and classified Adaptable Storage Platform for HPSS disk cache and began investigating deployment. (Appendix I, target FOUS-5)
- Deployed HPSS 8.3 to unclassified and classified environments. (Appendix I, target CSSE-2.d)

- Decommissioned unclassified legacy tape infrastructure; 54 Petabytes moved. (Appendix I, target FOUS-5)
- Provided support for deployed archival storage systems. (Appendix I, target CSSE-2.d)

Planned Activities in FY22:

- Develop/release HPSS 9.3 and 10.1. (Appendix I, target CSSE-2.d)
- Continue driving the LC I/O strategy: coupling archive with parallel file systems and analyzing HSM interfaces. (Appendix I, target CSSE-2.d)
- Deploy Adaptable Storage Platforms for HPSS disk cache to the unclassified and classified environments. (Appendix I, target FOUS-5)
- Deploy HPSS 9.2 to unclassified environment. (Appendix I, target CSSE-2.d)
- Decommission classified legacy tape library infrastructure. (Appendix I, target FOUS-5)
- Continue ongoing support for archival storage systems. (Appendix I, 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 FY21:

- Continued to make general performance and management improvements to ZFS-based Lustre. (Appendix I, target CSSE-2.d)
- Finalized ZFS DRAID and merged the feature into the OpenZFS main development branch. (Appendix I, target CSSE-2.d)
- Evaluated the use of Data-on-MDT in Lustre 2.12 to improve small file performance of user codes. (Appendix I, target CSSE-2.d)
- Completed Lustre 2.12 deployment across the center. (Appendix I, target CSSE-2.d)

- Developed client-side software to capture Lustre job statistics in the SONAR data analytics system to manage Lustre and improve user code performance. (Appendix I, target CSSE-2.d)

Planned Activities in FY22:

- Integration and testing of new Adaptable Storage Platform (ASP) Lustre filesystems for the CZ and RZ environments.
- Refine techniques for efficiently migrating Lustre files at scale from existing storage clusters to new ASP based storage clusters.
- Evaluate the use of shared directories in Lustre 2.14 to improve metadata distribution among servers and performance of user codes. (Appendix I, target CSSE-2.d)
- Develop software to use Lustre job statistics in the SONAR data analytics system to manage Lustre and improve user code performance. (Appendix I, target CSSE-2.d)
- Continue to make general performance and management improvements to ZFS-based Lustre. (Appendix I, target CSSE-2.d)

Networking and Testbeds (LLNL)

The Networking and Testbeds 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 FY21:

- Integration of HPE Slingshot Fabric Controller and CrayPE with TOSS 4 on HPE systems. (Appendix I, target CSSE-2.b)
- LLNL ARM systems updated to TOSS 4 in preparation for Astra conversion to TOSS 4. (Appendix I, target ATDM-2)
- TOSS 4 testing on early CTS2 hardware (Appendix I, target CSSE-2.d)

Planned Activities in FY22:

- TOSS 4 support on CTS2 hardware, both node and interconnect (Appendix I, target CSSE-2.d)
- Cassini NIC support in TOSS 4
- Support of CORAL 2 EAS3 hardware in TOSS 4. (Appendix I, target CSSE-2.b)
- Test Nvidia HPC Development kit with BlueField-2 support

Next-Generation HPC Networks (LLNL)

The Next Generation HPC Networks project focuses on an industry partnership on scalable network interconnects. The project seeks a partnership to develop a future generation network technology that can be utilized across multiple system integrators and component providers, support open source software, and impact both future AT and CT systems. The next-generation network will be optimized for traditional HPC workloads as well as emerging AI/ML and data analytics workloads. This project is a multi-year collaboration starting in FY21.

Accomplishments in FY21:

- Began preparations for next-generation HPC networks collaboration and partnership.

Planned Activities in FY22:

- Award an industry contract for research and development and kick off a tri-lab industry partnership in high-performance network for HPC, machine learning, and data analysis.

File Systems, Archival Storage, and Networking (LANL)

The ASC File Systems, I/O and Archival (FSAF) Storage program sub-element exists to meet ASC's requirements for storing, accessing, and analyzing input data, intermediate data, and curated data in order to perform large-scale scientific simulations supporting ASC weapons science and meet weapons performance program goals.

The primary objectives for FSAF are:

- Support ASC HPC users with scientific workflow requirements for performance of file systems, I/O infrastructure, and long-term data storage.
- Design current and next-generation storage systems.
- Conduct applied research to identify and develop promising technologies supporting the access, analysis, and storage of ASC's scientific data sets.
- Maintain LANL's existing approach to archiving by meeting LANL's development obligations to the HPSS.

FSAF impacts the ASC weapons science and performance program goals by providing the storage capability and the ability to extract value from the large data sets, and is responsible for ensuring that ASC data is protected and resilient within the data centers.

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. One project goal is that system users can 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, and to make predictions about the performance of these applications on future HPC platforms.

Accomplishments in FY21:

- Hardware Acceleration for File Systems using ZFS Interface for Accelerators (ZIA) which enables offload of all bandwidth intensive operations (e.g. RAIDZ)
- Improved Fault Tolerance for Storage by developing a new approaches to data protection (SODP and FODP) that provide reliability during failure events
- Enhanced ARM emulator (ARMie) and used to estimate theoretical vectorizability of several IC applications and delivered report of findings
- Continued GPU offload performance work on a EOS package used by IC apps, and related interfaces into the SAP code.

Planned Activities in FY22:

- Developing CONDUIT, the Capacity On Demand User Interface and Tookit, a system for rapidly shuttling data between storage tiers enabling simpler workflows and lowered storage costs.
- C2: Campaign Storage 2.0 (with FOUS) building the successor to MarFS, LANL's disk-based data lake. C2 is designed to tolerate catastrophic failures without data loss while making data more available for efficient analysis
- Continue assisting code teams devops with supporting new compilers and architectures and CI testing development (Gitlab runner based)

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. Further, SNL develops capabilities to facilitate data transfer between the large-scale machines and local sites, as well as within a site. This includes development of the Friendly Extensible Data Transfer Tool (FrETT) for use by the tri-labs and external groups.

Accomplishments in FY21:

- Completed FrETT Transmission Integrity (DTI) capability and HPSS batch/stage callback transfers, leveraging faster Recommended Tape Order (RAO) (Appendix I, target CSSE-1)
- HPSS 9.2 released, including CR302 (Increase crypto strength) to reduce the attack footprint for malicious use of HPSS and also standardizes the cross-realm authentication capabilities. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Provide FrETT/HPSS web-based batch transfer tool. (Appendix I, target CSSE-2)

- Complete HPSS CR302 component: Use industry-standard GSS instead of HPSS home-grown library. (Appendix I, target CSSE-1)

Scalable Input/Output Research (SNL)

The Scalable Input/Output Research project provides the necessary R&D to support anticipated data-management and input/output needs of ASC applications on Advanced Technology Systems. Application-support activities include development, maintenance, and integration of production-level I/O libraries such as IOSS, netCDF, HDF5, and CGNS. This project also provides evaluation of emerging storage technologies to understand the viability of such technologies for ASC mission work, R&D to continue development of user-level data services for in-system data management (e.g., distributed key-value systems), and new work to understand I/O requirements for high-performance data analytics. Success requires close collaboration with IC, ATDM, and the Integrated Workflow Project at SNL, multi-laboratory collaborations (particularly LANL and LLNL) to ensure performant third-party I/O libraries on ATS platforms, and active participation in the broader research community to advance capabilities in data-management and HPDA.

Accomplishments in FY21:

- Completed preliminary analysis of persistent memory hardware for HPC that show a 3x performance advantage of PMEM over NVMe, but at a cost of 10x. We also began investigating persistent-memory simulation through SST and HPE's Quartz simulator. (Appendix I, target CSSE-1.a)
- Used BeeGFS (a user-level file system) to create a first-instance of a machine-learning software stack for ARM, based on ATSE that supports Tensorflow, Torch, and other ML packages. (Appendix I, target CSSE-1.b)
- Worked with production I/O team to integrate the FAODEL data services software into the main branch of SEACAS. (Appendix I, target CSSE-1.b)

Planned Activities in FY22:

- Perform R&D to improve utilization of application-level data services. This activity will extend the capabilities of our data-service software (FAODEL) and explore effective ways to leverage FAODEL for efficient data-management and I/O in application workflows. (Appendix I, target CSSE-1.b)
- Perform R&D to develop enabling capabilities for high-performance data analytics (HPDA). This activity investigates and develops data-management technologies needed to enable effective use of HPDA on HPC platforms and aligns CSSE activities to the new Accelerated Digital Engineering initiative. (Appendix I, target CSSE-1.b)
- Evaluate technology advances in I/O and networking software and hardware to understand their potential impact on ASC applications and I/O software. (Appendix I, target CSSE-1.b)

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 FY21:

- Maintained a consistent and stable data analysis and visualization environment across LC platforms. (Appendix I, target CSSE-2.d)

- Provided operational support for visualization theaters and events, including demonstrations for high-level visits, reviews, and tours.
- 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.

Planned Activities in FY22:

- Support simulation users with data analysis and visualization activities including consulting and the creation of images, movies, and software.
- Maintain the data analysis and visualization environment across LC platforms and networks. (Appendix I, target CSSE-2.d)
- Provide operational support for visualization theaters and provide demo, tour, and special event support.

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 FY21:

- Released new versions of Hopper and Chopper that provide an integrated checksum ability for validating data, plus support for additional data formats, multiple user profiles, and general usability and scalability enhancements. Also developed design for factoring out an open-source version of Hopper that allows external support while also maintaining proprietary elements inside LLNL confines. (Appendix I, target CSSE-2.d)
- Developed and deployed an application for processing the bulk replacement of expiring RSA tokens. For MyLC, expanded HPC data sources and expanded support of IBM schedulers.

Planned Activities in FY22:

- Release new versions of Hopper and Chopper that improve portability by replacing dependencies on external binaries with self-contained code, improve performance when run under VPN, as well as general usability enhancements. Complete separation of Hopper into a version that is Open Source with a plug-in capability for adding in proprietary features. (Appendix I, target CSSE-2.d)

- In Lorenz, incorporate access to Sonar, focusing initially on MyLC's portlets that allow drilling down to job-related data. Build framework components to allow an SCF "task manager interface" that enables Hotline staff to safely perform disk quota changes and other system management tasks. Continue to enhance the MyLC Dashboard with user and staff requested features.

Visualization and Data Analysis (LANL)

Data analysis and visualization are key capabilities in taming and understanding the 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 of ASC end-users for deployed technologies.

The production and facilities component of the project is to provide LANL weapons designers with visualization systems research and support with emphasis on LSCI calculations. 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 Ansys.

The project focuses on integration into the code projects and ensures the production capability of the software components is well-integrated 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 post-processing analysis and visualization.

Accomplishments in FY21:

- Parallel Distributed Memory EAP dump reader for ParaView direct access to rage dumps in ParaView
- EAP LSCI Runs leveraged ParaView Catalyst in situ processing using 70,400 Haswell cores on Trinity (5-8 billion cells)
- Enabled the Ensight visualization application to read Catalyst dump files for IC codes

Planned Activities in FY22:

- Expanded testing of ParaView Catalyst linked against Ristra, EAP, LAP, SAP

- Convert EAP and LAP in situ adapters to Catalyst 2.0 leverages Conduit library adopted by ECP
- Ensure full Catalyst support for FleCSI current implementation specifically supports FleCSALE

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 FY21:

- Slycat ensemble visualization tool nominated for 2021 R&D 100 Award. (Appendix I, target CSSE-2)
- Completed new capability to support 3D mesh analysis in support of analyst needs for Mobile Guardian Transport and B61-12. (Appendix I, target CSSE-1.a)
- Developed a new capability in Sandia's Next Generation Workflow to enable automated ensemble analysis from computational simulation workflows in support of W80-4 and W87 programs. (Appendix I, target CSSE-1.a)
- ParaView 5.9.0 released with a number of Sandia contributions including improved support for the scripting interface for ParaView and Catalyst and improvements to volume rendering. (Appendix I, target CSSE-1.a)

Planned Activities in FY22:

- Enhance and improve ensemble analysis and visualization tools for ASC applications. (Appendix I, target CSSE-1.a)
- Provide integration of analysis and visualization for workflows. (Appendix I, target CSSE-1.a)
- Expand and simplify interface for Catalyst in situ analysis to support quantitative analysis. (Appendix I, target CSSE-1.a)
- Provide production visualization support. (Appendix I, target CSSE-1.a)

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 FY21:

Quantum

- Implement third-generation quantum simulation hardware testbed platform with fast feedback control and integration of classical HPC resources. Evaluated high-speed optical and superconducting control and interconnects for hybrid quantum-classical integration. (Appendix I, target CSSE-3.c)
- Computationally and experimentally demonstrated efficient mapping of algorithms to third-generation quantum simulation platform capabilities, including the NNSA application areas of quantum thermodynamics and quantum many-body dynamics. (Appendix I, target CSSE-3.c)
- Demonstrated connections between quantum computing and machine learning in application areas of interest to NNSA. (Appendix I, target CSSE-3.c)
- Developed key partnerships and quantum computing collaborations with industry and academia. (Appendix I, target CSSE-3.c)

Neuromorphic

- Released an initial version of the proxy application, CRADL, to explore inline high-frequency inference using advanced accelerators. (Appendix I, target CSSE-3.c)

- Develop a model for inline opacity queries of mixed materials using realistic material compositions. (Appendix I, target IC-3)
- Developed a parallelization scheme for the use of neural networks for Material Interface Reconstruction and investigated this solution on two material, 2D simulations. (Appendix I, target CSSE-3.c)

Planned Activities in FY22:

Quantum

- Implement fourth-generation quantum simulation hardware testbed platform with fast feedback control and expand integration of classical HPC resources. Continue development of high-density high-speed optical and superconducting control and interconnects for hybrid quantum-classical integration. (Appendix I, target CSSE-3.c)
- Transition to full-stack co-design for NNSA application areas of quantum thermodynamics and quantum many-body dynamics to drive fourth generation and beyond quantum hardware. (Appendix I, target CSSE-3.c)
- Continue development of machine learning approaches to enhance characterization and control of quantum hardware. (Appendix I, target CSSE-3.c)
- Demonstrate integration of quantum computing and classical/quantum machine learning in application areas of interest to NNSA. (Appendix I, target CSSE-3.c)
- Continue to develop quantum computing collaborations with industry and academia. (Appendix I, target CSSE-3.c)

Neuromorphic

- Use Material Interface Reconstruction application to explore performance on SambaNova. (Appendix I, target CSSE-3.c)
- Port the CRADL miniapp to SambaNova or Cerebras. (Appendix I, target CSSE-3.c)
- Complete the development of opacity models using a representative dataset and deploy an alpha version in a single code. (Appendix I, target PEM-5)

Emerging Technologies (LANL)

This project explores mapping ASC 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 Quantum Annealing, Quantum Gate, and Adiabatic 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.

Accomplishments in FY21:

- Discovery of instance of a trainable quantum machine learning protocol (quantum convolutional neural nets, QCNNs)
- Design and implementation of quantum dynamical simulation algorithms for near-term quantum computers
- New quantum linear system algorithm, state verification, and complexity
- Provided continuous capability of the D-Wave 2000 qubits system. (Appendix I, target CSSE-3)

Planned Activities in FY22:

- Algorithms for nonlinear operations on Quantum Circuits.
- Transition the research access for D-Wave from onsite to remote access to facilitate access to the latest D-Wave technologies and solvers. (Appendix I, target CSSE-3) Study of quantum machine learning algorithms focusing on barren plateau (vanishing gradient) phenomena and robust QML methods.
- Further development of quantum simulation algorithms for near-term QCs.

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.

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.

Sandia's NIC effort investigates viability of a next-generation heterogeneous HPC system that incorporates low-power neuromorphic processors. Specifically, exploring mature neuromorphic and machine learning enhanced HPC capabilities (potential applications, software stack, accelerator hardware, etc.) to justify scoping and procuring an HPC-level heterogeneous test platform coupling state-of-the-art conventional systems (CPUs, GPUs, ARM, etc.) and most suitable available neuromorphic processing units (NPUs). Each of these efforts will be coordinated with other efforts to maximize return on investment.

Accomplishments in FY21:

- Successful completion of L2 Milestone: Neural Mini-Apps for Future Heterogeneous HPC Systems. (Appendix I, target CSSE-3.b)
- In support of the L2 milestone, two mini-aps, a random walk (RW) and sparse coding (LCA) mini-App have been implemented and tested. White papers describing the technical model within each and proposed scaling studies on Loihi are completed. (Appendix I, target CSSE-3.b)
- Completed delivery and testing of new Intel Loihi neuromorphic testbed. (Appendix I, target CSSE-3.b)
- Implemented the Generic Spiking Architecture (GenSA) to enable parameter studies to find design candidates for key spiking algorithms (Appendix I, target CSSE-3.b)
- Designed novel fully static adiabatic logic family (S2LAL). Completed simulations for 2LAL (in CMOS7 and CMOS8) and S2LAL (in CMOS8). Achieved simulated 2aJ energy dissipation/FET/cycle in CMOS7. (Appendix I, target CSSE-3.b)
- Simulated effects of process variation in Tunnel Junctions in Si delta-layer systems. (Appendix I, target CSSE-3.d)

Planned Activities in FY22:

- Provide Tri-Lab access to research Loihi testbed. (Appendix I, target CSSE-3.b)
- Expand and improve neuromorphic software stack FUGU on testbeds. (Appendix I, target CSSE-3.b)
- Develop and test neuromorphic AI algorithms for HPC and edge applications with long-term potential in extremely energy-efficient numerical computing applications. (Appendix I, target CSSE-3.b)
- Develop and test neuromorphic algorithms for numerical analysis and computational physics processes for scientific computing HPC applications. (Appendix I, target CSSE-3.b)
- Develop and test neuromorphic algorithms for large-scale graph analytics for scientific computing HPC applications. (Appendix I, target CSSE-3.b)
- Expand and improve GenSA for co-design studies involving neuromorphic architectures. (Appendix I, target CSSE-3.b)
- Expand scope of adiabatic logic family simulation studies to include a range of voltages, additional processes, more complex circuits. (Appendix I, target CSSE-3.b)

Interagency Quantum Computing (SNL)

This project is focused on creating and testing circuits for noise mitigation and opto-electronic control of information processing on next generation advanced processor design. The project objectives require the development of custom hardware and software control systems as well as advancing an understanding of the physics that governs the noise properties. A goal of the project is to expand the accuracy of material property prediction relevant to NNSA's stockpile stewardship.

Accomplishments in FY21:

- Installed custom electronics control system to control the testbed. This system was an improvement over the previous control system because it is more robust, faster, and provides feedback capabilities to calibrate the hardware during an experiment. (Appendix I, target CSSE-3.c)
- We measured testbed system errors caused by noise and other imperfections in the analog electronic signal train. The collected data will be used to predict future performance and identify areas for hardware improvement. The feedback capabilities of the new control system have been employed to reduce the low-frequency components of these noise sources. (Appendix I, target CSSE-3.c)
- We completed fabrication of new devices to use in the testbed in case we determine that the performance of the current devices is limited by their design. (Appendix I, target CSSE-3.c)

Planned Activities in FY22:

- Increase the operational scope of the testbed by controlling more input and output channels on the device. (Appendix I, target CSSE-3.c)
- Reduce testbed error rate by collecting more data to understand its source and characteristics (magnitude, frequency spectrum). (Appendix I, target CSSE-3.c)
- Identify testbed vulnerabilities to environmental noise. (Appendix I, target CSSE-3.c)
- Package and distribute error and performance data to approved researchers (internal and external) so that they can better understand testbed performance, as well as predict the performance of future systems. (Appendix I, target CSSE-3.c)

Projects for Next-Generation Computing Technologies (WBS 1.2.3.5.7)

The NGCT product is the result of technology transfer from the ATDM subprogram back to the CSSE subprogram. This includes technology developed to bridge the environment and applications from current production efforts to exascale-class simulations. There is a cross-cut of enabling technologies.

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

NGCE is preparing ASC for the next generation of advanced computing technologies. Its project areas include portable user-level scheduling/runtime support for emerging workflows, system-level resource and power management/scheduling, advanced debugging/correctness tools and ROSE compiler technologies. NGCE's activities across these areas complement one another, and the advancements in each area synergistically support the overall project goal: providing readiness for ASC/ATDM applications and simulation workflows on exascale computing systems while mitigating increasing challenges stemming from the introduction of ML. NGCE uses a three-pronged approach to exert its impact broadly: 1) significantly advance each individual R&D discipline, 2) drive interdependence among these disciplines and other ASC Program elements through strategic co-design tasks, and 3) provide broad community outreach and communications.

Accomplishments in FY21:

- Improved the usability, robustness and performance of Flux-based solutions by expanding our collaborations with key software teams including LBNL, LLNL's Validation Suite team, RIKEN's Fugaku team in Japan, ECP ExaAM, ECP ExaWorks, IBM T.J. Watson Research Center and RedHat OpenShift teams, etc. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d)
- Improved the portability and interoperability of Flux-enabled workflows by integrating Flux with a wider range of workflow management software systems, which include Maestro, Merlin, and UQP/Themes, Parsl, Swift/T, Radical Pilot as well as domain specific workflow tools such as ATOM's GMD pipeline. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d)
- Facilitated continuous co-design of El Capitan debugging and performance tools including AMD GPU Address Sanitizer support for Umpire and STAT testing on an El Capitan early access system (EAS)-2 system. (Appendix I, targets CSSE 1.b/2.b, IC-1, and ATDM 1.d)
- Laid a foundation for Flux-based elastic workflows to utilize Kubernetes' technologies better by demonstrating the Fluxion scheduler can efficiently enable HPC scheduling for Kubernetes. (Appendix I, targets CSSE 1.b/2.b)

Planned Activities in FY22:

- Continue to lead the El Capitan tools working group to ensure the bring-up of debugging/performance tools on El Capitan, which includes our Floating Point correctness tool suite and more automated testing of tools on large applications. (Appendix I, targets CSSE 1.b/2.b and ATDM 1.d)
- Bridge further gaps in providing Flux solutions on the workflows on El Capitan early access system (EAS)-2 systems, which includes HPE MPI bootstrapping support within Flux. (Appendix I, targets CSSE 1.b/2.b, IC-1 and ATDM 1.d)
- Engage the Kubernetes scheduling standardization group with a proposal to augment Kubernetes' key interfaces to support HPC workloads more effectively. (Appendix I, targets CSSE 1.b/2.b)

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 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 FY21:

- Led the El Capitan Center of Excellence (COE). Assisted IC and PEM code teams with initial porting efforts to AMD hardware. Identified and reported problems with HPE and AMD compilers, libraries, and development tools. (Appendix I, target CCSE-1.b)
- Assisted the Cheetah team with porting and optimization for Sierra. Demonstrated greater than 10x speedup (Sierra node vs. CTS node) on problems of programmatic interest. (Appendix I, target ATDM-1.d)
- Contributed improvements to Spack to support the use of Spack to simplify developer workflows. Demonstrated this capability with KULL and the AML project. Delivered documentation and tutorial materials. (Appendix I, target CCSE-2.b)
- Participated in co-design and standardization activities including the CORAL-2 working groups and the OpenMP standards committee. (Appendix I, target ATDM-1.d)

Planned Activities in FY22:

- Work with selected IC and PEM code teams including KULL, ARES, ALE3D, and MARBL to execute El Capitan Center of Excellence activities. (Appendix I, target CCSE-1.b)
- Continue to assist in the development of RAJA, CHAI, and Umpire to further the performance portability goals of ASC applications. (Appendix I, target IC-1)
- Continue working with selected IC and PEM code teams to help them meet performance goals on Sierra. (Appendix I, target ATDM-1.d)
- Share knowledge via wikis, talks, publications, etc. (Appendix I, target ATDM-1.d)

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 efficiently deploy developed software 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 FY21:

- Gave Hatchet tutorials to LLNL WSC code teams.
- Expanded hatchet, SPOT, and Caliper into several large WSC codes. (Appendix I, target IC-1)
- Improved deployment and scalability infrastructure of SPOT and Hatchet.
- Added support to Caliper for memory bandwidth measurements.
- Built and release a container-mode for SPOT, allowing other labs the capability to deploy SPOT.

Planned Activities in FY22:

- Develop pre-canned analysis in Hatchet for exploring common performance analysis on sets of runs.
- Build tighter integration between SPOT and Hatchet, allowing users to customize hatchet notebooks with SPOT.
- Integrate Sina database backend (from the LLNL ATDM/ASD Workflow project) into SPOT, which will allow application users to share performance data with developers. (Appendix I, target CSSE-2.b and CSSE-2.d)

User Workflow and Modernization (LLNL)

The User Workflow and Modernization project's mission is to build infrastructure and components that enhance the end-to-end productivity of ASC HPC assets. The project is focused on building solutions to concrete user problems in three areas prioritized by the user community: problem setup, simulation management, and post-processing. The project has three active software projects: Sina, a capability to capture simulation logging and other outputs in a code-agnostic form, create searchable catalogs of simulation campaigns, and leverage this data for orchestrating complex workflows; the contours-to-codes (C2C) and shaping project, focused on building a common geometry specification across our codes that is compatible with the LANL contour format; and ROVER/Devilray, an in-situ hardware-accelerated multigroup ray casting packages for both simulated diagnostics and volume rendering.

Accomplishments in FY21:

- **Code-agnostic Geometry and Shaping:** designed specification language for common shaping syntax, implemented parser and processor and demonstrated in ARES, supported MARBL v1.0 in making all example geometry C2C-based.
- **Code-agnostic outputs and data management:** designed and implemented new curve set feature for Sina common output library, integrated Sina library into Mercury, KULL, and MARBL, enhanced APIs and enhanced Conduit
- **ROVER/Devilray:** assisted with an L1 milestone in creating images on ¼ of Sierra added distributed-memory framework, higher-order isosurfaces and volume rendering

Planned Activities in FY22:

- **Code-agnostic Geometry and Shaping:** work with ARES & ALE3D to integrate common shaping library, upgrade to python 3.8, integrate C2C in ALE3D
- **Code-agnostic outputs and data management:** work with LC Workflow Enablement Group to provide persistent Sina databases, create additional schemas for code outputs, integrate with more codes, pursue further Sina/Kosh integration
- **ROVER/Devilray:** enhanced MARBL support via Ascent package, enhanced higher order mesh support, line outs, develop first scatter approximation for ray effects (with Ardra)

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.

DevRAMP is funded by both ATDM and CSSE. ASC-targeted effort is reported here, and community-targeted effort is reported in the ATDM section.

Accomplishments in FY21:

- Developed new Spack functionality for developer workflows; significantly refactored MARBL code to better support Spack packaging; worked with MARBL team to convert MARBL build system to Spack. (Appendix I, target ATDM-1)
- Completed a study of memory utilization on LLNL machines using Sonar.

Planned Activities in FY22:

- Continue to work with WSC code teams to improve developer workflows and to support internal build caches, pipeline deployments, and containerization. (Appendix I, target IC-1).
- Expand Sonar deployment on the SCF and continue to automate deployment of Sonar services.

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 the Message Passing Interface (MPI), OpenMP, and PMIx. For MPI, we focus on tool support, hybrid architecture support, and fault-tolerance support. In OpenMP, our focus is on the tools interfaces OMPT and OMPD. In PMIx, we maintain leadership as well as participate in working groups, e.g., for developing storage interfaces. However, we participate and monitor developments in all parts of all three 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 SC ASCR program, vendors, and academia, such as planning and technical coordination for vendor and academic contracts and CORAL2 working group efforts.

Accomplishments in FY21:

- Improved SCR for ASC codes and released SCR 3.0; Developed strategy for I/O support on El Capitan system; Developed I/O management support for ML and complex workflows. (Appendix I, target CSSE-2.b)
- Participated in the standardization of MPI 4.0. Contributed use cases to support standardizing MPI features. Developed a specification for the global-restart fault-

tolerance model (Reinit) and led discussion of the model in the MPI fault-tolerance working group. (Appendix I, target IC-1)

- Participated in the OpenMP tools group in the discussion of a DWARF-aware interface to improve the existing OpenMP debugging interface (OMPD). (Appendix I, target IC-1)
- Participated in the MPI Hybrid working group for GPU-triggered communication interface. Participated in CORAL 1 (NVIDIA) and CORAL 2 (Cray/HPE) Communications working groups on the topic. (Appendix I, target IC-1)

Planned Activities in FY22:

- Continue investigating I/O technologies, including work on user-level file systems, and I/O performance measurement. Continue vendor interactions including CORAL2 efforts. (Appendix I, target CSSE-2.b)
- Continue SCR modification for ASC applications and for supporting complex workflow data management. (Appendix I, target IC-1)
- Finalize the global-restart fault-tolerance specification by incorporating feedback from the MPI Forum and checkpoint/restart tools. (Appendix I, target IC-1)
- Continue participation in the MPI Forum meetings to support the standardization of features critical for LLNL applications in the next version of MPI, 4.1. (Appendix I, target IC-1)
- Continue participation in the MPI Hybrid working group in the discussions of MPI interface for GPU-triggered communication, as well as in CORAL 1 (NVIDIA) and CORAL 2 (Cray/HPE) Communications working groups on the topic. (Appendix I, target IC-1, target CSSE-2.b)

Co-Design and Programming Model Research and Development (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.

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, execution, and control models that are consistent with state-of-the-art runtimes, including Legion and asynchronous MPI. The Ristra computer science component of the project builds upon FleCSI to implement multiphysics production applications capable of taking advantage of FleCSI's task and data parallel model. This is in collaboration with IC where the Physics side of the Ristra application is being developed. A well-specified interface allows for a separation of concerns and concurrent efforts between the two projects, providing an environment for rapid development of both new physics and computer science capabilities. This aligns Ristra and FleCSI with ASC's future goals and rapidly evolving mission space.

This project is co-funded by both CSSE and ATDM, specifically the co-design scope is funded in CSSE NGCT. The accomplishments and planned activities are in the ATDM CDA and ASD sections.

Accomplishments in FY21:

- Ristra Apps ran at scale on Sierra (P9+Volta), Trinity (Haswell), and Astra (ARM)
- Substantial offloading to Nvidia GPUs through Kokkos in several Ristra Apps
- Ristra building on gitlab CI/CD and uses cross-cutting Spack build infrastructure to improve developer spin-up time
- FleCSI 2.0 delivered and training class held.

Planned Activities in FY22:

- Preparation for El Cap – support through Kokkos and Legion
- Ristra integration with CMF and associated tools
- FleCSI addition of a new topology to support K-D trees.
- FleCSI addition of Block-Structured AMR for *Eulerian* hydrodynamics.

Advanced Machine Learning (LANL)

The Advanced Machine Learning project has mission-focused efforts in collaboration with PEM, V&V, and ATDM. The project conducts studies of AI, ML, and Deep Learning methodologies applied to multiscale simulation, materials modeling, nuclear data evaluation, turbulence simulations, and radiographs analysis and data-driven approaches that can explain underlying physics. The ASC program is well-positioned to leverage investments in experimental facilities, next-generation computer architectures, algorithm development, and simulation data collection initiatives like LSCI to develop ML workflows to utilize multisource, multifidelity data for answering mission-relevant questions from these areas. LANL's CSSE, V&V, PEM, and IC subprograms co-fund these mission-focused research projects. This project is also co-funded from ATDM interagency co-design.

Accomplishments in FY21:

- Initial investigation of ML for group structures in PARTISN.
- Radiograph analysis team implemented methods for adding physics based constraints into a neural network for radiographic density reconstructions.

- Highly accurate interatomic potentials have been developed for elemental metals including aluminum, tin, gold, and silver. Models for alloys and felectron materials are in progress.

Planned Activities in FY22:

- Demonstrate detailed understanding of the effects of group structures and weight functions on multigroup cross sections for PARTISN.
- Produce database of physics-based in situ signatures for defect detection using in-line pyrometry and thermal imagery for laser powder bed fusion manufacturing.
- LANL will adapt our multi-fidelity machine learning approach to ensembles of aging simulations performed in parameter studies, and will investigate similarities

Engineering Common Model Framework (SNL)

The Engineering Common Model Framework (ECMF) is intended to provide an institutional resource for storing, sharing, and evaluating computational models for stockpile and modernization systems. The ECMF aims to address this challenge by integrating existing technologies at Sandia into a common platform: data archiving using DataSEA and the Simulation Data Management (SDM) database, workflows through the Integrated Workflow (IWF) tool, and workflow regression testing leveraging DevOps capabilities. ECMF is an ASC integrated project that cross-cuts other sub-elements including CSSE and V&V, and is aligned with the Labs Strategic Objective 2.3 to *deliver ND relevant exemplars demonstrating new capabilities including tools, processes, and workflows needed to support greater M&S utilization in qualification activities*.

Accomplishments in FY21:

- Deployed an instance of DataSEA on the SRN CEE Cloud to store the regression test reports. (Appendix I, target CSSE-2)
- Formalized the architecture of ECMF and started prototype development and integration between Jenkins SRN, Docker, SAW, NGW and DataSEA. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Seek opportunities for integration between ADE and ECMF, as well as leveraging and integration across other ASC capabilities and ongoing individual projects (e.g., workflows, model repositories, IWF/SAW/NGW capability). (Appendix I, target CSSE-2)
- Extend set of ND-relevant workflows to support strategic drivers identified by ASC. (Appendix I, target CSSE-2)
- Stand up a pilot of ECMF for testing and evaluation. (Appendix I, target CSSE-2)

Next-Generation Development and Performance Analysis Tools (SNL)

The Next-Generation Development and Performance Analysis Tools project provides a prototyping and support function to Sandia's ASC projects and IC code teams. The project utilizes a cross-section of hardware, system software, tools, and applications expertise to provide initial ports of key performance kernels, libraries and applications to next-generation hardware systems including ATS, CTS and Vanguard platforms. Observations and analysis of performance results are obtained and, in most cases, collaborations with industry vendors and other NNSA laboratories is undertaken to improve hardware or software performance where bugs or deficiencies are found. In addition, this project supports integrated build and testing for TRILINOS and agile components which the next generation application codes are dependent on. This testing is deployed to support a range of ASC platforms including testbeds, Vanguard, CTS and early access and ART platforms for aligned with ATS.

Accomplishments in FY21:

- Initial porting of TRILINOS solver framework to early access AMD/El-Capitan systems (Appendix I, target CSSE-2.b)
- Early evaluation of Intel OneAPI next-generation compiler framework to target Crossroads platform in 2022 (Appendix I, target CSSE-2.a)
- Initial porting of Mantevo mini-applications to NextSilicon data flow simulator to provide analysis to Vanguard-II project (Appendix I, target CSSE-2.c)
- Development of LMDS (lightweight monitoring) connector for KOKKOS-based applications and libraries enabling application progression and kernel time sampling to be recorded at large-scale. This tool development contributes to the CSSE/FOUS LDMS L2 milestone for FY21. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Continued porting and evaluation of TRILINOS and prototype ATDM applications on AMD/El-Capitan systems (Appendix I, target CSSE-2.b)
- Provide additional evaluations of applications and porting results for Crossroads deployment, including support for acceptance runs of the ATDM SPARC application and Sandia benchmark collection used on the system. (Appendix I, target CSSE-2.a)
- Provide porting support for Vanguard-II early access hardware systems (Appendix I, CSSE-2.c)
- Provide prototype results from KOKKOS autotuning tool capabilities (Appendix I, CSSE-2)

Scalable Data Management and I/O (SNL)

This project provides the necessary R&D to support the data management and scalable I/O performance of next-generation applications on ASC advanced technology systems. The data management and I/O work supports the development, maintenance, and integration of production-level I/O libraries such as Input/Output SubSystem (IOSS), netCDF, Hierarchical Data Format 5 (HDF5), and Computational fluid dynamics General Notation System (CGNS). The work also includes development and integration of the emerging storage technologies such as integrated nonvolatile memory (e.g., burst buffers) and unconventional data services (e.g., distributed key-value systems).

The R&D 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 ECP.

Accomplishments in FY21:

- Worked with The HDF Group (THG) and LBL to make improvements to HDF5 Virtual Object Layers (VOL) for caching, asynchronous I/O, and provenance capture. (Appendix I, target CSSE-2)
- Worked with data services team to integrate FAODEL into the master branch of SEACAS. (Appendix I, target CSSE-2)
- Worked with EMPIRE developers to streamline the parallel analysis workflow by using an internal decomposition. This significantly reduced the end-to-end time for EMPIRE runs. (Appendix I, target CSSE-2)
- Finalized design of the Discontinuous Galerkin (DG) API and the low-level storage semantics in Exodus/NetCDF file. (Appendix I, target CSSE-2)
- Added zlib-ng (next-generation) compression for Exodus and CGNS which shows improvements of up to 5x compared to the base zlib library. (Appendix I, target CSSE-2)
- A multi-lab technical team from DOE laboratories and DDN helped establish a statement-of-work and project plan for a contract with DDN to explore innovative new object-storage and file system capabilities for HPC. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Provide production I/O-library support for ASC applications. In collaboration with ASC application teams, this activity provides performance improvements and feature development for Sandia's IOSS library as well as improvements to third-party I/O libraries (e.g., HDF5, netCDF, CGNS) and support of middle-tier storage (e.g., burst buffer) to enable high performance on ATS platforms. (Appendix I, target CSSE-2)
- Perform R&D to develop enabling capabilities for high-performance data analytics (HPDA) and the integration of HPC and Cloud computing. This activity investigates

and develops data-management technologies needed to enable effective use of HPDA on HPC platforms. (Appendix I, target CSSE-2)

- Evaluate technology advances in I/O and networking software and hardware to understand potential impact on ASC applications and I/O software. (Appendix I, target CSSE-2)

Vanguard Advanced Prototype Project (SNL)

The scope of this project is to lead the design, acquisition, and plan for deployment of the Vanguard Advanced Prototype System (APS) to accelerate the maturation of new HPC system architectures for future ASC HPC production platforms. 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 Vanguard is to field large-scale prototypes that are targeted for nuclear weapons (NW) mission workloads including assessing 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 ASC procurements. This project covers all aspects of the technical, programmatic, and procurement planning for APSs. The first Vanguard APS is the Arm-based Astra platform, delivered and deployed in the fall of 2018. Future Vanguard systems will be named as systems are moved from conceptual development into procurement development.

Accomplishments in FY21:

- Supported use of Vanguard-Astra by the NNSA labs ATDM application teams for successful completion of the ASC L1 ATDM milestone application performance and scaling analysis (Appendix I, ATDM-1.c)
- Restarted technical evaluation and technology selection for future Vanguard-II deployment after negotiations of final award for earlier Vanguard-II system became non-viable (due to technology provider roadmap changes). (Appendix I, CSSE-1.c)
- Placed RFQ for restarted revectored Vanguard-II system with procurement and expect to deployment of early access hardware systems during Q4 CY21. (Appendix I, CSSE-1.c)

Planned Activities in FY22:

- Evaluate the performance of applications on coarse grained reconfigurable dataflow accelerated node architecture for Vanguard-II using vendor simulators, engineering samples and production early access hardware on open and restricted networks (Appendix I, target CSSE-1.c)
- Provide Tri-lab users access to Vanguard-II early access hardware systems using SARAPE on OHPC and IHPC networks (Appendix I, target CSSE-1.c)

- Evaluate optical network switching performance using early access Vanguard-II hardware and FPGA-based NIC systems (Appendix I, CSSE-1.c)
- Support ASC L2 milestone for early microbenchmark and kernel performance evaluation on Vanguard-II early access hardware (Appendix I, CSSE-1.c)
- Manage DDN Red technology development contract and coordinate Tri-lab collaboration activities to co-design POSIX support, object storage APIs, and data management capabilities. (Appendix I, CSSE-1.c)

Advanced Tri-lab Software Environment (ATSE) (SNL, LANL, LLNL)

The scope of this project is primarily focused on accelerating the maturity and successfully deploying system software stacks for the Vanguard Advanced Architecture Prototype Systems (AAPS) for ND mission workloads. Beyond this ATSE is a modular, extensible and open HPC stack vehicle to explore innovative new software technologies which are targets for future ASC systems including testbeds ATS and CTS. This software stack effort encompasses aspects of the entire HPC software stack including areas of focus such as OS kernels, runtime systems, libraries supporting inter-processor 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. As a prototype software stack, ATSE for Vanguard-2 will be used to explore new approaches with potential to improve the ASC computing environment, including system software support for managing extreme heterogeneity at the node and system levels, supporting more cloud-like HPC-as-a-service usage models, and deploying user-friendly methods for orchestrating containerized workflows.

The ATSE project will initiate and coordinate technology development engineering contracts with the Vanguard technology providers to improve and optimize their products for ASC workloads. 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 FY21:

- Developed automated workflow for generating reproducible builds of ATSE using Spack; contributed changes needed to support ATSE to the upstream Spack project. (Appendix I, target CSSE-2)
- Developed and deployed ATSE on several advanced architecture testbeds, including "Inouye" Fujitsu A64FX, "Stria" Marvell ThunderX2 + NVIDIA V100 GPU, and "Blake" Intel SkyLake. Demonstrates flexibility of ATSE for supporting diverse hardware. (Appendix I, target CSSE-2)
- Prototyped in-platform unprivileged container build mechanism using Podman. This enables end-users to login to an HPC system and build their own containers directly on the platform, which has not been possible previously. (Appendix I, target CSSE-2)
- Developed new finer-grained ATSE containers and deployed to CEE build and test cluster. (Appendix I, target CSSE-2)

Planned Activities in FY22:

- Develop ATSE system software plan for Vanguard-2 early test systems and FY24 large-scale platform. (Appendix I, target CSSE-2.c)
- Develop and deploy ATSE via containers on Vanguard-2 early test systems on the open (OHPC) and restricted (SRN) networks. (Appendix I, target CSSE-2.c)
- Collaborate with Vanguard-2 vendor to mature their software stack and support the needs of the tri-lab user community. (Appendix I, target CSSE-2.c)
- Explore options for supporting more cloud-like usage models on HPC systems. (Appendix I, target CSSE-2)
- Support the DDN next generation scalable storage collaboration and contract and engage the NNSA labs in the activity (Appendix I, target CSSE-1.c)

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 Astra system 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 inter-processor 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 FY21:

- Evaluated Arm compiler development and tested arm FORTRAN compiler (LANL)

Planned Activities in FY22:

- Continued evaluation of Arm compiler development and testing of arm FORTRAN compiler (LANL)

Advanced Machine Learning (SNL)

The Advanced Machine Learning project has mission-focused efforts in collaboration with IC, PEM, V&V, and ATDM-CDA. The project conducts studies of Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning methodologies applied to multiscale simulation, materials modeling, nuclear data evaluation, turbulence simulations, and radiographs analysis and data-driven approaches that can explain underlying physics. The ASC program is well-positioned to leverage investments in experimental facilities, next-generation computer architectures, algorithm development, and simulation data collection initiatives like Large Scale Calculations Initiative (LSCI) to develop ML workflows to utilize multisource, multi-fidelity data for answering mission-relevant questions from these areas.

Accomplishments in FY21:

- Demonstrated 200x speedup using spatial and temporal parallelism of PyTorch for HPC over serial approach. Results presented at 2021 SIAM CSE meeting. (Appendix I, target CSSE-2)
- Developed and filed a copyright for the CDFG extraction tool for LLVM, a mapping tool to translate ASC AI/ML applications to a control-flow data graph that can be evaluated through a variety of hardware-accelerator simulations. (Appendix I, target CSSE-2)
- Used the Intel Loihi testbed system at Sandia to generate 50k training runs for Markov chain approximation of the diffusion process – relevant for radiation transport mission problems. (Appendix I, target CSSE-3.b)

Planned Activities in FY22:

- Develop foundational capabilities for the deployment of AI/ML software on HPC systems. (Appendix I, target CSSE-2)

- Develop foundational methods to develop estimates for resources, accuracy, and performance of machine learning HW and SW (Appendix I, target CSSE-2)
- Explore the application of AML technologies to data analysis and visualization for mission workflows (Appendix I, target CSSE-2)

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 2020, and through quarter 3, fiscal year 2021, are reflected below for the FOUS subprogram.

- Procured and implemented an HPC next generation classified network and began to converge initial production workloads on the new infrastructure. (Appendix I, targets FOUS-1 and FOUS-5) (LLNL)
- Completed B654 Low Conductivity Water Cooling Loop for future CTS systems. (Appendix I, target FOUS-5) (LLNL)
- Integrated Ruby, an additional CTS-1 compute cluster, in the Open Computing Facility (OCF) collaboration zone to augment existing capacity computing platforms. (Appendix I, target FOUS-5) (LLNL)
- Deployed RZNevada, an El Capitan early access system, to customers on the LLNL unclassified restricted zone (Appendix I, target FOUS-1a) (LLNL)
- Continued operation of Sierra and CTS-1 clusters with sustained high utilization levels despite minimum-safe on-site staff presence during the COVID-19 pandemic. (LLNL)
- Completed Crossroads Installation Design, L2 milestone #7132, for the necessary electrical, mechanical and structural modification and additions required to install the Crossroads system at LANL. (Appendix I, target FOUS-1.b) (LANL)
- Established a new production unclassified restricted enclave for new HPC services in support of tri-lab Remote Computing Enablement (RCE). (Appendix I, target FOUS-1) (LANL)

- Continued production operation of Trinity (ATS-1) and CTS-1 systems Snow, Fire, Ice, and Cyclone in full production use, peaking with over 95% utilization and 99% system availability. (LANL)
- Continued high sustainability operation of ASC systems by improving water quality of reclaimed water used for computer cooling drastically reducing consumption of total water required. (LANL)
- Provided 24x7 HPC operations of ASC systems while transitioning workforce to maximum telework during pandemic. (LANL)
- Deployed CTS-1+ cluster "Manzano" to the Sandia Restricted Network. (Appendix I, target FOUS-5) (SNL).
- Deployed HPC system and application monitoring and analysis pipeline on SNL CTS-1 system as part of L2 milestone: "Integrated System and Application Continuous Performance Monitoring and Analysis Capability". Incorporates LDMS-based system and application data transport, KOKKOS-enabled application performance data, and two-week interactive visualizations of EMPIRE application progress in conjunction with system resource utilization. (L2 to be completed in FY21 Q4) (SNL)
- Installed infrastructure (power/cooling and networking) for CTS-2 systems in conjunction with 3MW power upgrade to the 725-East HPC Facility. (Appendix I, target FOUS-5) (SNL)
- Completed 100Gbps upgrade design, plans, contracts, and procurements for DisCom, including physical path diversity and redundancy for increased resiliency and reliability (in conjunction with RCE) (Appendix I, targets FOUS-1 and CSSE-1) (SNL)
- Restructured HPC utilization reporting in the Workload Characterization Tool to match Defense Programs budget changes and added reporting for HPC Utilization by weapon system. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): CTS-2 Initial Deployment		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 06/30/22		
ASC WBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-5		
Description: Deploy initial CTS-2 scalable units at LLNL and demonstrate utilization by initial users running a code of interest to the ASC program. The first delivery of CTS-2 hardware to the national labs represents a culmination of years of effort by system architects, technical review teams, contract negotiators, and facility engineers. To support this milestone, system integrators will install, configure, and integrate the first scalable units of CTS-2 into the LLNL HPC environment. System software developers will provide a provide a functional and performant development environment, and support consultants will ensure that initial ASC users can successfully run on the system.		
Completion Criteria: An initial user successfully runs an ASC code on the deployed CTS-2 system and demonstrates effective use of the computing resources.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: FOUS, CSSE, IC		

Milestone (ID#TBD): CORAL-2 EAS-3 Deployment		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 06/30/22		
ASC WBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-5		
Description: Deploy a CORAL-2 EAS-3 early access system at LLNL and demonstrate utilization by initial users running a code of interest to the ASC program		
Completion Criteria: An initial user successfully runs an ASC code on the deployed EAS-3 system and demonstrates effective use of the computing resources. The EAS-3 system will be the final early access generation delivery prior to the full El Capitan supercomputer. This system will provide a key platform to prepare codes to run on El Capitan, and it will be a powerful supercomputer on which to run real-world applications of interest to the weapons program. To support this milestone, system integrators will install, configure, and integrate the EAS-3 into the LLNL HPC environment. System software developers will provide a provide a functional and performant development environment, and support consultants will ensure that initial ASC users can successfully run on the system.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: FOUS, CSSE, IC		

Milestone (ID#TBD): Classified Persistent Database Services		
Level: 2	**Fiscal Year:** FY22	**DOE Area/Campaign:** ASC
Completion Date: 09/30/22		
ASC WBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, targets FOUS-1.a and FOUS-5		
Description: Procure and deploy a classified Persistent Database Services offering at Livermore Computing comprised of an Openshift/Kubernetes hardware and software ecosystem, underpinned by container registry and image security scanning infrastructure for curated orchestration of specific database, message broker, and web enabled microservices as required by emerging user-facing workflows in the HPC center.		
Completion Criteria: Demonstrated capability to instantiate a curated containerized backend microservice such as MariaDB, RabbitMQ, Redis, or MongoDB within the PDS ecosystem and access its provided functionality within the HPC center on commonly available login systems for remote client access.		
Customer: ASC		
Milestone Certification Method: (select two then delete one of the three listed below) 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): CTS-2 Installation		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 09/30/22		
ASC WBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-5		
Description: Install hardware fulfilling CTS-2 mission.		
Completion Criteria: Connect compute racks to facility power and cooling.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: FOUS		

Milestone (ID#TBD): CTS-2 delivery and site integration		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 09/30/22		
ASC WBS Subprogram: FOUS, CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-5		
Description: Along with the selected vendor and LLNL, SNL will Configure, deliver, and integrate on-site cluster technology provided by next generation Commodity Technology System (CTS-2) contract (LLNL Subcontract #).		
Completion Criteria: Assuming the current CTS-2 contract cadence continues, a CTS-2 cluster of XSU (X compute nodes) size shall: • have final BOM created to meet SNL site needs • have parts allocated and accounted for at vendor site • be built and tested at vendor site • be delivered by the vendor to SNL Bldg. 725 East • be integrated and deployed to early users on SNL's SRN.		
Customer: ASC		
Milestone Certification Method: A program review is conducted, and its results are documented. Professional documentation consisting of a report is prepared as a record of milestone completion.		
Supporting Resources: FOUS, CSSE		

Milestone (ID#XXXX): Roadmap for Integrated Sandia ASC DevOps		
Level: 2	**Fiscal Year**: FY22	**DOE Area/Campaign**: ASC
Completion Date: 09/30/22		
ASC nWBS Subprogram: FOUS, IC, ATDM, CSSE, PEM, and V&V		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, targets IC-1, V&V-3, ATDM-1, CSSE-2, FOUS-1, and FOUS-2		
This L2 milestone will develop and deliver a strategy and detailed roadmap that when executed across the ASC program, will enable ASC codes to be developed and deployed more quickly, efficiently, and robustly on all key ASC platforms. When the roadmap is executed, codes will be able to be efficiently coupled and integrated, and share common development, test, and deployment environments and corresponding practices. This includes build and test infrastructure and standards as established by the ASC DevOps to minimize developer time spent deploying and releasing credible code. The delivered capability and ongoing support will be used to decrease costs and improve efficiencies of developing and deploying ASC applications. The roadmap will emphasize: • Create a centralized team that manages the build and test infrastructure and standards • Make a standard set of tools available on all production HPC platforms, ART systems, and CEE build and test farms and migrate ASC code teams to this standard • Establish standards and a build and test infrastructure to deploy and release credible code • Build and test hardware strategy, execution, and maintenance teams supporting ASC DevOps needs • Leverage testbed activity associated with heterogenous advanced architecture prototype systems (HAAPS) and the Remote Computing Enablement (RCE) activity • Continuous integration in the ASC DevOps ecosystem with rapid turn-around time To meet this deliverable, the team will work across SNL's ASC program and the TriLab Remote Computing Enablement (RCE) project.		
Completion Criteria: SAND Report describing the roadmap to achieve the ASC DevOps vision within five years.		
Customer: ASC program and ND mission		
Milestone Certification Method: A program review is conducted, and its results are documented. Professional documentation consisting of a report is prepared as a record of milestone completion.		

Supporting Resources: Engagement across the ASC program via the sub-element leads.

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 FY21:

- Continued FY21 procurement support, contract management, and program planning. (Appendix I, targets FOUS-1 and FOUS-5)
- Managed the PSAAP III program for the nine Academic Alliance Centers.
- Managed contracting needs for CORAL2 NRE and Build subcontracts. (Appendix I, targets FOUS-1.a)
- Managed contract negotiation for CTS-2 systems, with expected signing and approval by Q4FY21. (Appendix I, targets FOUS-5)

Planned Activities in FY22:

- Continue FY22 procurement support, contract management, and program planning. (Appendix I, targets FOUS-1 and FOUS-5)
- Support the semiannual Predictive Science Panel (PSP) meeting to be held in March 2022 at LLNL.
- Continue to manage the PSAAP III program for new Academic Alliance Centers.
- Continue contract and procurement support for CORAL2 and CTS-2 systems. (Appendix I, targets FOUS-1.a and FOUS-5)

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 FY21:

- Provided LANL support for HQ.

- Supported PSAAP III collaboration efforts at LANL.
- Supported Presidential Early Career Awards for Scientists and Engineers (PECASE).
- Managed virtual SC support and prepared to host SC21 booth for DOE in FY22.

Planned Activities in FY22:

- Support PSAAP III collaboration efforts at LANL.
- Participate in PI meeting and plan PSP meetings.
- Provide LANL support for HQ.
- Support PECASE
- Host SC21 booth for DOE.

Program Support (SNL)

The Program Support project provides critical ASC SNL program integration, communication, and management within the laboratories and with the external community. Management within the laboratory includes the interface with Sandia's Advanced and Exploratory organization and Stockpile Management, 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 FY21:

- Organized and hosted an FY20 Predictive Engineering Science Panel (PESP) review.
- Performed organizational and host responsibilities for the ASC PI Meeting that was held virtually in May 2021 due to the COVID-19 pandemic, and started planning activities for the 2022 ASC PI Meeting to be held in Monterey, CA.
- Supported the programmatic needs of the PSAAP II program and participated in planning for the PSAAP III program and review of PSAAP III RFP responses.
- Completed the first phase of an integrated road mapping process that cuts across the program and will help ASC more effectively deliver integrated products to the nuclear deterrence mission.

Planned Activities in FY22:

- Organize and host the 2021 PESP review planned for late October 2021.
- Organize and host the ASC PI meeting, to be held in May 2022 in Monterey, California.

- Support the programmatic needs of the PSAAP III program as it continues execution.
- Complete the first draft of the ASC integrated roadmaps and share with Sandia's nuclear deterrence program.
- Support the ECI, including joint planning and execution with DOE/NNSA and DOE/SC laboratories.

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 FY21:

- Participated in a focused technical exchange to improve material data management using Granta MI under JOWOG 28. Installed Granta MI 2020/2021 on a development server to collaborate on Granta schema development and improve Y-12 capabilities in material characterization and diagnostics.
- Collaborated with Lawrence Livermore National Laboratory on a draft Production Simulation Initiative (PSI) proposal to model and simulate numerous manufacturing activities relevant to the CNS mission.
- Participated in the ASC Principal Investigators Meeting and ASC Advanced Machine Learning Initiative Workshop. Discussed machine learning techniques for radiograph interpretation and segmentation of CT volumetric datasets with researchers at LANL, LLNL and SNL.
- Upgraded and maintained modeling software from LLNL and SNL on Linux and Windows platforms at Y-12.

Planned Activities in FY22:

- Refine ALE3D models for sheet hydroforming and thermomechanical manufacturing processes (e.g., welding, forming) at the mesoscale and use PLATO to support the design of additive manufacturing for tooling and fixtures.
- Continue model development using the SNL SIERRA and Dakota software to support predictive simulation on material properties of additively manufacturing materials. Validate simulations through experiments with a tensile and compression stage within a CT system.

- Continue to participate in a focused technical exchange to improve material data management using Granta MI under JOWOG 28. Improve Y-12 capabilities in characterization and diagnostics to quantify material uncertainties and provide new insights into materials.
- Investigate existing alkalai-halide (A-H) molecular dynamics (MD) interaction potentials and calculate basic thermodynamic properties of two example material systems.
- Work with stakeholders to determine the applicability of existing A-H interaction potentials to systems of interest at Y-12.
- Develop testing framework for neural network being used to perform pattern recognition for X-ray scattering experiments ongoing at Y-12.
- Develop A-H interaction potentials from experimental data and density functional theory (DFT) using machine learning.
- Continue simulation work to support development of the total scattering analytical tool and model additional constituents present in the area of interest to determine structures from experimental scattering data.
- Upgrade and maintain modeling software from LLNL and SNL on Linux and Windows platforms at Y-12.

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 FY21:

- Completed the transition of production RSA OTP server instances in all LC environments to RSA AM 8.4.
- Deployed SIEM technologies within classified HPC environments for aggregated log analysis in a broader enterprise context.
- Integrated Ruby, an additional CTS-1 compute cluster, in the Open Computing Facility (OCF) collaboration zone to augment existing capacity computing platforms. (Appendix I, target FOUS-5)
- Increased capacity and enhance the performance of NAS storage for home directories and project workspace. (Appendix I, targets FOUS-1 and FOUS-5)
- Evaluated and investigated new degaussing methods for heat-assisted magnetic recording (HAMR) hard drives that will be used in next-generation file systems.
- Initial integration of ASP Lustre files system (replacement for Porter Lustre file system) a scaled-out preproduction file system serving several clusters in the collaboration zone. (Appendix I, target FOUS-5)
- LC Infrastructure – Integration of NAS flashed based storage to support VM infrastructure environment.
- Procured and deployed hardware and software infrastructure to host Persistent Data Services (PDS), a curated container orchestration environment for on-premises HPC-enabling workflows, in the unclassified restricted zone.

Planned Activities in FY22:

- Fully deploy LC's new identity management system and complete the migration of all workflows. (Appendix I, targets FOUS-1a and FOUS-5)
- Establish Kerberos support for NAS to support system multi-tenancy (Appendix I, targets FOUS-1a and FOUS-5)
- Integrate CTS-2 systems in the Open Computing Facility (OCF) collaboration zone and Secure Computing Facility (SCF) to augment existing capacity computing platforms, subject to hardware availability. (Appendix I, target FOUS-5)
- Integrate new ASP hardware Lustre environments for Open Computing Facility (OCF) collaboration zone and restricted zone environments.
- For LC infrastructure, integrate new Cisco UCS blade hardware for OCF (collaboration zone and restricted zone environments) and SCF environments. (Appendix I, targets FOUS-1.a and FOUS-5)
- Deploy improved Splunk server and storage infrastructure on the OCF and SCF for better data retention and performance for logging, monitoring, system analysis.
- Procure and deploy classified PDS for curated container orchestration of specific database, message broker, and web enabled microservices for HPC enabling workflows.

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 FY21:

- Provided ongoing user support through hotline operations, documentation, and training while in minimum-safe and minimal-normal operations. (Appendix I, targets FOUS-1 and FOUS-5)
- Executed the phase-out of the six PSAAP II centers and the phase-in of the nine new PSAAP III centers.
- Worked with LC Security and Programming teams to create, implement and refine requirements for Forge Rock, the replacement for the IdM Identity Management System. (Appendix I, targets FOUS-1 and FOUS-5)

Planned Activities in FY22:

- Continue providing support through hotline operations, documentation, and training with a focus on assisting users porting to Sierra and CTS-2 systems. (Appendix I, targets FOUS-1 and FOUS-5)
- Redesign, rearchitect and upgrade LC's primary web site hpc.llnl.gov from Drupal 7 to Drupal 9. Migrate frequently viewed tutorials to GitHub.
- Develop documentation and training materials in the use of El Capitan and the CTS-2 class clusters. (Appendix I, targets FOUS-1 and FOUS-5)

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 FY21:

- Procured and implemented an HPC next-generation classified network and began to converge initial production workloads on the new infrastructure. (Appendix I, targets FOUS-1.a and FOUS-5)
- Completed and closed out B654 Low Conductivity Water Project. This project allows for future unclassified CTS-2 and CTS-3 systems to be liquid cooled. Appendix I, target FOUS-5)
- Upgraded B439 for Computing Consolidation Project. This project allows for future classified high side computing through NA-50 investments.
- Developed the design plan for El Capitan site infrastructure. (Appendix I, target FOUS1.a)

- Developed an acquisition strategy plan for El Capitan site infrastructure construction execution. (Appendix I, target FOUS-1.a)
- Completed design and construction site preparation packages for the second El Capitan early access system. (Appendix I, target FOUS-1.a)

Planned Activities in FY22:

- Deploy and transition classified workflows to a 100G DISCOM (Tri-lab data transfer network). (Appendix I, targets FOUS-1 and FOUS-5)
- Rearchitect Open Compute Facility (OCF) critical network infrastructure resources to better align to changing network environments and improve overall redundancy and resiliency.
- Complete design site preparation package for B-453 El Capitan Site Infrastructure project (Appendix I, target FOUS-1.a)
- Take beneficial occupancy of ECFM late 2021 to early 2022 (Appendix I, target FOUS-4)
- Complete the B-654 Expansion design and begin construction (Appendix I, target FOUS-1.a)

Platforms Administration (LANL)

The Platforms Administration project provides production computing administration for ASC computational systems for weapons designers, developers, and engineers. The following activities are included: system hardware and software maintenance and configuration, system troubleshooting and problem resolution, system integration, platform preventative maintenance, platform security, coordination, and outage management activities. Effort in this project begins with the acceptance of delivery and deployment into production of ASC systems, continues with the day-to-day management of ASC systems during their lifecycles, and ends with decommissioning of ASC systems.

Accomplishments in FY21:

- Managed Dedicated Application Times and Dedicated System Times throughout FY21 while maximizing system availability at 99% for CT systems and 99% for Trinity. (Appendix I, target FOUS-5)
- Performed cybersecurity monitoring and managed security patches of all ASC platforms. (Appendix I, target FOUS-5)
- Early preparation for production support of ATS-3 system software. (Appendix I, target FOUS-1.b)
- Participated in selection and award of CTS-2 contract. (Appendix I, target FOUS-5)

Planned Activities in FY22:

- Initial installation and production support of ATS-3 system. (Appendix I, target FOUS-1.b)
- Initial installation and production support of CTS-2 systems. (Appendix I, target FOUS-5)
- Continue to enhance system and user security for AT and CT systems. (Appendix I, target FOUS-5)

High Performance Computing Operations (LANL)

The High Performance Computing Operations project provides 24 × 7 × 365 operations and monitoring of the ASC computing resources, storage systems, network, and visualization resources. The computer operators provide first-tier support for all ASC systems, hardware support and tracking, triage support for problem determination, and management of a leading-edge classroom for user training. Effort in this project is focused on providing timely and productive computational cycles to the ASC user community by maximizing reliability and availability of ASC resources.

Accomplishments in FY21:

- Established and performed hardware support for CTS-1 systems. (Appendix I, target FOUS-5)
- Provided 24 × 7 × 365 Tier-1 support and operations of compute, network, and file system resources. (Appendix I, target FOUS-5)
- Provided continuous hardware maintenance, assessment of system life spans, and associated decommissioning activities for ASC computing equipment. (Appendix I, target FOUS-5)
- Managed HPC operations throughout pandemic with maximum telework posture without disruption of ASC system availability.

Planned Activities in FY22:

- Provide 24 × 7 × 365 Tier-1 support and operations of compute, network, and file system resources. (Appendix I, target FOUS-5)
- Perform continuous hardware maintenance, assessment of system life spans, and associated decommissioning activities for ASC computing equipment. (Appendix I, target FOUS-5)
- Support onsite immersive training classroom for ASC users.
- Manage Classified Removable Electronic Media (CREM) for ASC storage systems.

Facilities for High Performance Computing (LANL)

The scope of the Facilities for High Performance Computing project is to support operations of the computing rooms, mechanical cooling, electrical power distribution, and structural elements necessary to support ASC computing. Activities include long-term facility planning, daily planning and operations, engineering, design, construction support, preventive and corrective maintenance, facility training, computer installation and integration, shipping and receiving, and equipment storage.

Accomplishments in FY21:

- Completed design for installation of Crossroads system. (Appendix I, target FOUS-1.b)
- Completed CTS-2 cooling preparations. (Appendix I, target FOUS-5)
- Completed facility electrical bus duct replacement to remove system electrical installation and integration limitations. (Appendix I, targets FOUS-1.b and FOUS-5)
- Supervised or performed maintenance on HPC facility mechanical and electrical equipment to sustain maximum availability of Trinity and CTS-1 systems. (Appendix I, target FOUS-5)

Planned Activities in FY22:

- Install Crossroads system. (Appendix I, target FOUS-1.b)
- Install CTS-2 system hardware. (Appendix I, target FOUS-5)
- Continue high efficiency computer cooling operations in maximizing use of non-potable water.
- Supervise or perform preventive, predictive, and corrective maintenance on HPC facility mechanical and electrical equipment. (Appendix I, target FOUS-5)

Parallel Infrastructure (LANL)

The Parallel Infrastructure project is responsible for development and production support of the network, NAS, storage area network (SAN), and parallel file systems necessary for ASC users to perform scalable I/O on computational systems and data transfers between computational systems. Activities include WAN (Distance Computing, DisCom) support contract, file system development, network administration, network (home/project) file system management, and parallel (scratch) file system management.

Accomplishments in FY21:

- Supported the RCE working groups to improve user productivity on tri-lab systems. (Appendix I, target FOUS-1)
- Purchased initial campaign storage for supporting the Crossroads and CTS-2 systems in production. (Appendix I, targets FOUS-1.b and FOUS-5)

- Provided top-tier diagnostic and troubleshooting support for standing production file systems and networks, which served especially crucial to keeping remote computing jobs running during maximum telework throughout pandemic.

Planned Activities in FY22:

- Deploy DISCOM 100-Gb upgrade. (Appendix I, target FOUS-1)
- Deploy initial campaign storage for supporting the Crossroads and CTS-2 systems in production. (Appendix I, targets FOUS-1.b and FOUS-5)
- Implement data management features and controls to support user data management. (e.g. CONDUIT and GUFI).

Platforms Tools, Visualization & User Support (LANL)

The Platform Tools, Visualization & User Support project is responsible for software tools on ASC platforms, production visualization services, and direct customer service for local and remote users of ASC/LANL resources. Activities include development and delivery of documentation and training materials for ASC/LANL resources, usage statistics, administrative interface for ASC users, user support services, operational metrics, monitoring of HPC resources, web development, user training and software workshops, large simulation success assurance, resource management (scheduler), parallel runtime environment software and configuration, and production platform performance testing and validation.

Accomplishments in FY21:

- Early development of programming environment for future systems. (Appendix I, target FOUS-1.b and FOUS-5)
- Provided ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Established the new restricted enclave to support network and cybersecurity equivalency to LLNL and SNL in support of RCE. (Appendix I, target FOUS-1)

Planned Activities in FY22:

- Provide ongoing consulting, user support services, documentation, and training for ASC platforms, file systems, and parallel tools infrastructure.
- Promote activities such as improved workflows, and proactive support initiatives that increase efficiency of large-scale simulations on AT and CT systems. (Appendix I, target FOUS-5)
- Identify priority or milestone efforts on AT and CT systems requiring focused or intensive user support, and provided that support through small team direct engagements.

- Continue participation in RCE working groups to improve ASC user productivity on tri-lab systems. (Appendix I, target FOUS-1)

Deep Storage (LANL)

The Deep Storage project provides production archival storage to ASC users. Activities include tri-lab data transfer, long-term mass storage, research and development of archival storage futures, system administration, storage planning, network coordination, and integration with compute systems.

Accomplishments in FY21:

- Continued development of Marchive.
- Decommissioned obsolete tape libraries.
- Provided top-tier diagnostic and troubleshooting support of archival storage for ASC users, especially important during maximum telework posture throughout pandemic to keep file systems from becoming full and unstable.
- Deployed upgraded archival storage hardware.

Planned Activities in FY22:

- Continued development of Marchive.
- Formulation of new user archive storage workflow.
- Provide expert problem diagnosis and troubleshooting for ASC archival storage.

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, analysis, and reporting services. This project supports tri-lab ATS platform resource allocations and coordinates with tri-lab peers in establishing priority scheduling, if required. This project coordinates the integration and deployment of CTS platforms 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.

This project also supports Application Readiness Testbed (ART) platforms (e.g., Mutrino (Trinity), Vortex (Sierra)), at-scale prototype platforms (e.g., Astra), and advanced architecture testbeds.

Accomplishments in FY21:

- Deployed CTS-1+ cluster "Manzano" to the Sandia Restricted Network. (Appendix I, target FOUS-5)
- Participated in the CTS-2 proposal contract negotiations, architecture decision points, and initial planning for site integration. (Appendix I, target FOUS-5)
- Supported successful completion of tri-lab ASC L1 ATDM Milestone work on Astra. (Appendix I, targets ATDM-2.a and FOUS-2)
- Supported COVID-19 Consortium efforts to provide HPC resources to internal/external users.
- Provided program support and coordination for tri-lab ATCCs.
- Standup of next generation advanced architecture testbeds (Cray Shasta, Penguin Fujitsu A64X [Inouye]), Data Analytics Clusters for production HPC Monitoring [Bitzer, Shirley and Shaun], Machine Learning Accelerators ([DodgeCity]) as well as large, advanced storage capabilities (IBM ESS 14PB) [Earle]. (Appendix I, targets CSSE-1, ATDM-2, and FOUS-5)
- Rapid systems reconfiguration to support changing critical needs of code teams including expansion of ART system Vortex to meet increased demand for ART cycles in support of Sierra, redeployment of GPU-enabled resources in support of code team development and testing needs (IBM Power9 – Waterman to Weaver move to OHPC), and rapid implementation and iteration of GPU software capabilities (CUDA and ROCm) (Appendix I, targets CSSE-1, ATDM-2, and IC-1)
- Deployed FrETT 1.2, includes linear Tape Order Recall to improve HPSS tape-stored file transfer performance. (Appendix I, targets CSSE-2 and FOUS-5)
- Deployed new high-performance tape libraries in open, restricted, and classified environments. (Appendix I, target FOUS-5)

Planned Activities in FY22:

- Configure, deliver and install onsite Commodity Technology System (CTS-2) cluster in support of milestone ID#TBD. (Appendix I, target FOUS-5)
- Configure, test and deploy TOSS4 on all production HPC systems. (Appendix I, target FOUS-5)
- Complete deployment of additional systems on the iHPC network. This will provide improved Tri-lab HPC capabilities for the Remote Computing Enablement (RCE) project.
- Deploy new parallel file systems and core Infiniband network equipment purchased in FY21 to support CTS-2 deployments. (Appendix I, target FOUS-5)
- Provide program support and coordination for tri-lab ATCC. (Appendix I, target FOUS-1)

- Standup Vanguard-2 early prototype clusters with an integrated development-to-production workflow. (Appendix I, targets CSSE-1 and FOUS-2)
- Standup of testbed Machine Learning platforms and enhanced storage capabilities and integrations to facilitate cross-platform usage modes (e.g., shared filesystem across testbeds, ASC Build and Test Farm & Production Computing systems) (Appendix I, targets CSSE-1 and ATDM-2)
- Complete large-scale data migrations to new archival storage hardware in restricted and classified environments. (Appendix I, target FOUS-5)
- Deploy multi-host transfer process FrETT/HPSS to increase HPSS transfer scaling. (Appendix I, targets CSSE-2 and FOUS-5)
- Production harden SNL's current HPC monitoring, analysis, and visualization infrastructure to improve query, analysis, and archive (storage and retrieval) performance (Appendix I, CSSE-2, FOUS-5)
- Iteratively improve the production monitoring and analysis infrastructure in collaboration with HPC system administrators and users.
- Provide training and documentation resources about use of the infrastructure by the user and operations community (Appendix I, targets CSSE-2 and FOUS-5)

Advancing HPC Operations (SNL)

The Advancing HPC Operations project develops capabilities to enhance operations of current and future production platforms as well as documentation and training to support their use. Capabilities include HPC monitoring, analysis, visualization, and feedback for improved performance understanding and resource management decisions; improvements in system management methodologies and architectures; and support of advanced application runtime environments (e.g., containers and advanced OS). This project leverages domain knowledge from operations of our advanced platforms and develops and applies new technologies to directly benefit our production operations.

Accomplishments in FY21:

- Deployed enhanced file system metadata tools to improve user data management capabilities for system administration and user data migration. (Appendix I, target FOUS-5)
- Integrated planning for deployment of and systems support for Vanguard-2 exploratory systems. (Appendix I, targets CSSE-1 and FOUS-2)
- Implemented container capabilities on testbeds and ART systems in order to facilitate exploring their use with ATSE / SPACK (SPATSE) and other environments, implementing configuration management using tools such as Ansible. (Appendix I, targets CSSE-1, ATDM-2 and FOUS-2)

- LDMS 4.3.7 integrated with TOSS 3 and TOSS 4 and deployed on SNL 1500 node CTS-1 system. (Appendix I, targets CSSE-1, FOUS-1, and FOUS-5)
- Deployed analyses and dashboards to present HPC SRN users and system administrators live HPC system and application metrics and assist in problem and status identification. (Appendix I, targets CSSE-2 and FOUS-5)
- Developed machine learning models and statistical analyses to identify anomalies within HPC applications and integrated the analyses into our production analysis and visualization framework. (Appendix I, targets CSSE-2 and FOUS-5)

Planned Activities in FY22:

- Deploy user facing file system metadata tools to improve user data management capabilities. (Appendix I, target FOUS-5)
- Expand analyses and dashboards for HPC application and system data to enable better temporal correlation of application performance variation with system factors during run time (Appendix I, targets CSSE-2 and FOUS-5)
- Collaboration within FOUS, CSSE, and SNL Enterprise computing on configuration management strategies (GMI, Ansible Tower), containerization. (Appendix I, targets CSSE-1, ATDM-2 and FOUS-2)
- Prototyping advanced data management capabilities related to metadata, with the goal of improving ease-of-use and automation. (Appendix I, targets CSSE-2 and FOUS-5)
- Continue improvements and associated code releases of LDMS and Scalable Object Store (SOS). (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Apply deployed ML-based models and statistical analyses for anomaly detection to production applications. (Appendix I, targets CSSE-2 and FOUS-5)
- Investigations of capabilities to support Dynamic Resource Management including expansion of Machine Learning based models, statistical analyses, and visualizations as well as infrastructure to enable direct feedback from analytics engines to system and application software. (Appendix I, targets CSSE-2 and FOUS-5)

User Support (SNL)

The User Support project provides user support, development environments (e.g., DevOps infrastructure), 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 FY21:

- Provided service desk (OneStop), tier-3 user support, and service desk improvements for SNL and tri-lab ASC computing system to many thousands of users at high workloads this past year. Continued to monitor metrics and services to ensure improvement initiatives are identified and executed.
- Expanded the centralized ASC DevOps Environment by 1) deploying expanded build and test development platforms to enhance the DevOps pipeline environment for ASC code teams; 2) providing a common set of compilers and MPI across HPC platforms; and 3) providing centralized turnkey Jenkins and CDash services. Facilitated adoption by SIERRA and RAMSES code development projects.
- Continued maintenance of developer tools (profilers, debuggers, and system performance and analysis) for ASC development teams and user workflows.

Planned Activities in FY22:

- Provide Service Desk (OneStop), tier-3 user support, and service desk improvements for SNL and tri-lab ASC computing systems. (Appendix I, targets FOUS-1, FOUS-2, and FOUS-5)
- Expand DevOps capabilities and extend the DevOps Roadmap.
- Maintain developer tools (profilers, debuggers, and system performance and analysis) for ASC development teams and user workflows. (Appendix I, targets CSSE-2, ATDM-1, and IC-1)

Facilities, Networking, and Power (SNL)

The Facilities, Networking, and Power project supports maintenance and improvements to the facilities and infrastructure servicing the HPC systems (CTS and file system servers). 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.

This activity also focuses on reducing overall operating expenses by minimizing cooling and electrical distribution through a comprehensive program of introducing more efficient computer room air conditioning (CRAC) units, using higher voltage electrical source power distribution units (PDUs), exploring alternative energy sources and conservation mechanisms, which include reducing the volume of chilled water required for cooling, reducing overall water utilization, introducing warm-water liquid cooling technologies, and deploying outside air "free cooling" infrastructure and automation.

Accomplishments in FY21:

- Conducted short- and long-term planning to improve infrastructure resilience and power/cooling capacity for 725-East/West HPC data centers via UPS in 725E and generator support for 726 CUB. (Appendix I, targets FOUS-2 and FOUS-5)
- Deployed North route of 100Gbps DisCom WAN links and encryptors between sites in collaboration with LANL and LLNL. (Appendix I, targets FOUS-1 and CSSE-1)
- Deployed unclassified test capability across the production DisCom links and tested new 100Gbps MACSec encryptors (ViaSat) between SNL/CA and SNL/NM in preparation for classified production deployment. (Appendix I, targets FOUS-1 and CSSE-1)
- Improved SNL's ESNet redundancy configurations between sites. (Appendix I, targets FOUS-1 and CSSE-1)
- Improved network capabilities of iHPC in support of the tri-lab RCE. Tri-lab operations agreement in place. (Appendix I, targets FOUS-1 and CSSE-1)

Planned Activities in FY22:

- Start design and construction of additional 4-MW power/cooling upgrade for the 725-E HPC facility. (Appendix I, targets FOUS-2 and FOUS-5)
- Complete installation of equipment to improve infrastructure resilience and power/cooling capacity for 725-East/West HPC data centers via UPS in 725E and generator support for 726 Cooling Utility Building which feeds building 725. (Appendix I, targets FOUS-2 and FOUS-5)
- Complete the South 100Gbps WAN link deployment for DisCom. (Appendix I, targets FOUS-1 and CSSE-1)
- Complete the full deployment of the classified DisCom 100Gbps encryptors for production availability. (Appendix I, targets FOUS-1 and CSSE-1)
- Test and deploy more effective data movement applications (e.g., FrETT), network tuning, and host connectivity to utilize the increased DisCom bandwidth. (Appendix I, targets FOUS-1 and CSSE-1)
- Improve resiliency of LANL's Internet and ESN connection (providing command and control functionality for DisCom) by enabling failover between the North and South DisCom paths (mirroring SNL's resiliency improvements in FY21) (Appendix I, targets FOUS-1 and CSSE-1)
- Design iHPC upgrade to 100Gbps. (Appendix I, targets FOUS-1 and CSSE-1)

Support Contracts for Production Use of the Hierarchical Data Format (HDF) Library (LLNL, SNL, LANL)

Many integrated codes leverage the open source HDF library in order to enable scalable I/O performance while providing comfortable abstraction and organization of scientific and engineering data. This project will provide funding for each laboratory to access qualified, experienced, HDF library developers for resolution of issues and faults. Ongoing support was moved from CSSE to FOUS in FY20 because the work is of production-support nature.

Accomplishments in FY21:

- Enterprise Support (through FOUS):

  Improved documentation, evolved plans for community development, maintained MPI compatibility, improved space (memory) performance, addressed approximately 20 bugs reported by NNSA labs.

- Non-recurring Engineering (through CSSE):

  Migration of HDF5 development to GitHub. Improved CMake build system, reduced compiler warnings for gcc compilers, stood up CDash for HDF5 dev branch on key NNSA systems, parallelized key support tools (e.g., h5diff), multi-threaded prototype for limited scope H5Dread, improved CGNS scalable open behavior and performance, began creation of a performance test suite.

- The THG ES contract provided several benefits in FY21. These include:
  1. Bug reports submitted by multiple individuals have all been resolved and/or the HDF5 developers are currently working to resolve the reported issues.
  2. The timeliness of the support has reduced downtime and increased productivity of the Sandia users of the HDF5 library.
  3. The contract has increased communication between the HDF5 developers and the NetCDF developers as well as increased support of the CGNS developers. Both the NetCDF and CGNS libraries are foundational libraries for Sandia's I/O needs in addition to the HDF5 library.

Planned Activities in FY22:

- Enterprise Support (FOUS):

  Renegotiate our ES license rate down to $50K for FY22. There are components to the ES license that we did not take advantage of (training, ~50 hours of expert consulting)

- Non-recurring engineering (CSSE):

  Build out scalability of key parallel tools (h5diff, h5repack, h5stat), improve compiler warnings for clang, implement parallel page buffering, improve scaling behavior of parallel dataset filters (e.g., compression), produce a performance cookbook with examples, productize any Virtual Object Layer (VOL) connectors of interest to NNSA labs, evolve the performance test suite, provide quality and responsive customer support for NNSA use cases.

- Continue to leverage support contract, improve packaging and distribution of HDF5 and provide priority response to NNSA tri-lab user requests.

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 Tri-lab Linux Capacity Cluster (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. The CCE also coordinates the effort to provide a secure, robust, and performant remote computing environment for tri-lab users. This cross-lab remote computing environment provides users with access to HPC resources at the other two labs that is equivalent to the access at their own lab.

(Appendix I, targets FOUS-1, FOUS-2, and FOUS-5)

System Software Deployment for Commodity Technology Systems

The projects involved in this area include TOSS and Monitoring, Metrics, and Analytics Integration.

TOSS is the software stack that runs on Linux CTS 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 secure, performant, fully functional cluster OS capable of running MPI jobs at scale and provides a common software environment on CT systems across the tri-lab complex. TOSS endeavors to support the growing diversity of HPC architectures, accelerators, and interconnects on future ASC systems; to that end, TOSS also runs on many hardware testbeds. Well-defined processes for release management, packaging, quality assurance testing, configuration management, and bug tracking are used to ensure a production-quality software environment that can be deployed across the tri-lab in a consistent and manageable fashion.

The Monitoring, Metrics, and Analytics Integration project increases the efficiency of NNSA HPC centers and aids future planning with monitoring and analysis. Specifically, the project will: 1) deploy data collection and analysis infrastructure across the HPC center (clusters, applications, facilities, etc.); 2) develop portable analysis techniques (ML and otherwise) that can be applied to data gathered at multiple facilities to 3) derive Figures of Merit (FOMs) from monitored data that can guide and optimize decisions by resource managers, applications, administrators, and management.

Accomplishments in FY21:

- Released TOSS 3.7 for TLCC2 and CTS-1/1+ systems, and TOSS 4.1 and 4.2 for CTS1/1+, Astra, and CORAL-2 EA systems. Continued TOSS 3 support for Astra on an as-necessary basis. (Appendix I, targets FOUS-1, FOUS-2, and FOUS-3)

- Incorporated Lustre server support into TOSS 4.
- Performed initial installation and check-out of TOSS 4 on the CTS-2 vendor test cluster and assisted the evaluation team with the interconnect bake-off. (Appendix I, targets CSSE-2.d and FOUS-5)
- Continued monthly meetings to discuss research topics and collaboration points with tri-lab Monitoring, Metrics, Analytics Integration (MMAI) workgroup members.
- Continued contract with NCSA to deploy performant HPC monitoring and analysis and visualization framework through containers for use across tri-lab infrastructures.
- Collaborated with industry and university partners to develop and evaluate Machine Learning and statistical techniques for identifying anomalous and/or performance-impacting conditions of interest in HPC systems.
- Integrated promising Machine Learning and statistical analyses into production HPC system monitoring frameworks.

Planned Activities in FY22:

- Support the integration of TOSS 4 on CTS-2 clusters as they begin to arrive. (Appendix I, target CSSE-2.d and FOUS-5)
- Continue support of TOSS 3 for TLCC2 and CTS1/1+ systems. Continue as-necessary TOSS 3 support for Astra until it transitions to TOSS 4. Continue development and support of TOSS 4 for CTS1/1+/2, Astra, and CORAL-2 EA systems. (Appendix I, targets CSSE-2.b, CSSE-2d, FOUS-2, and FOUS-5)
- Evaluate gaps and opportunities to increase overall system security posture and ensure the system software supply chain.
- Continue to investigate the effectiveness of multi-cluster Slurm deployments. (Appendix I, target CSSE-2.d, and FOUS-5)
- Explore, develop, and deploy state-of-the-art HPC monitoring, analysis, visualization, and feedback techniques across tri-lab HPC infrastructures.
- Collaborate with tri-lab members on new LDMS features and push of new LDMS releases into TOSS.
- Hold periodic monitoring and metrics meetings to share new results and approaches as nucleation points for collaborating on development and deployment.
- Hold periodic MMAI meetings to
    - share new results and approaches, and
    - form nucleation points for collaborating on development and deployment.

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. Focus areas include improving development and support for common dependencies of performance analysis, testing, and debugging tools. Included in these improvements are continuous performance profiling capabilities and lightweight sampling experiments, both aiming to reduce overhead of profiling applications and improve performance data collection during application runtimes. The debugging efforts include vendor contracts to provide support and training and to enhance the tool capabilities. Tri-lab developed debugging tools such as Stack Trace Analysis Tool (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 MVAPICH2 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.

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

Accomplishments in FY21:

- Managed contracts under the Open Source Contract Maintenance for Performance Analysis Software projects, which enabled workshops and training for, and enhancements of, HPCToolkit, Medium-Weight Profiling Tool, Dyninst, Tuning and Analysis Utilities (TAU), Trenza's Survey Tool, and MUST/Vampir.
- Designed and configured prototype TCE stacks for deployment across various TriLab architectures (Appendix I, targets CSSE 2.d and FOUS-5).
- Developed and upstreamed numerous Spack features in support of tri-lab TCE deployment.
- Created infrastructure that enables sharing TCE deployment capabilities across TriLabs (Appendix I, targets CSSE 2.d and FOUS-5).
- Enabled sharing of common tri-lab container building specifications (Dockerfiles) and container images for CTS-2 acceptance applications. Began work to understand performance and portability optimizations.
- Developed a set of container registry requirements.
- Stack Trace Analysis Tool (STAT) ported to AMD GPUs using ROCgdb and changes adopted into HPE's programming Environment deployment.

Planned Activities in FY22:

- Manage various contracts under the Open Source Contract Maintenance for Performance Analysis Software projects, e.g., Rice University's HPCToolkit, University of Wisconsin-Madison's DyninstAPI, University of Oregon's Tuning Analysis Utilities, Trenza's Survey Tool, and MUST/Vampir.
- Deploy shared TCE development environment to production systems at TriLabs (Appendix I, targets CSSE 2.d and FOUS-5).
- Unify inter-lab differences that remain in TCE stacks (Appendix I, targets CSSE 2.d and FOUS-5).
- Work to improve container building using Spack, including performance and portability optimizations.
- Investigate best practices for HPC containers in production.
- Deploy a containerized production application across the tri-labs.
- Procure a production container registry solution.

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 Synchronized Account Request Automated ProcEss (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.

HPC Metrics and Reporting (Workload Characterization) 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 FY21:

- Added the capability to renew access at LANL and LLNL.
- Added a new site: Nevada National Security Site (NNSS). Trained the processing agent and established resources that site users can request from LLNL and SNL

- Added support for nine PSAAP III universities. Two existing PSAAP II universities were updated to PSAAP III status and added the seven new PSAAP III universities and trained their processing agents. Remaining PSAAP II sites are archived.
- Coordinated Tri-Lab reporting changes based on Defense Programs budget restructure. New reporting structure required changes to the WCTool database for Sandia and LLNL.
- Released multiple versions of WCTool software for Sandia and LLNL for bug fixes and enhancements.

Planned Activities in FY22:

- Continue to operate the Tri-Lab SARAPE process for all remote access account requests. Improve the user interface experience and the remote site processing access and usability as needed.
- Address the reduction in the amount, duration, and accessibility of PII stored or available in the SARAPE Application and Database.
- Investigate opportunities for more automation.
- Continue efforts to implement additional utilization reporting based on weapons systems for the HQ quarterly reporting.
- Address HPC Metrics and Reporting (Workload Characterization) bugs and update software to current versions and security updates.
- Address issues in evolving tri-lab computing environments.

Remote Computing Enablement

The goal of the Remote Computing Enablement (RCE) project is to achieve a remote HPC user experience as close as possible to the local user experience - and use the opportunity to improve the experience of all - to maximize productive utilization of computing resources across the NNSA HPC simulation complex.

The RCE team is comprised of over 50 members from multiple disciplines and management strata across the LANL, LLNL, and SNL HPC centers. Working groups have been formed in several focus areas covering network improvements, software gaps, workflow and continuous integration, login efficiencies, and more.

At its core, RCE is a communication tool for the labs to focus on specific areas of current or potential alignment with a one-to-three-year implementation outlook.

Accomplishments and Planned Activities involve all three laboratories, unless otherwise noted.

Accomplishments in FY21:

- Designed and implemented a tri-lab continuous integration (CI) architecture in the unclassified restricted network. This included deployment of dedicated Gitlab server instances as well as architecture design and coding effort to support CI runners tightly coupled with all three labs' security infrastructure. Limited availability for early adoption users at end of FY21.
- Instantiated a LANL Restricted Enclave (RE) network security space within its HPC center and successfully migrated production platforms.
- Adopted unclassified restricted cross-realm trust model agreement between all three labs. Began cross-realm trust implementation with new LANL RE at LLNL and SNL HPC centers.
- Engaged tri-lab SMEs on coordinated test and debug sessions for WAN networking; improvements made in areas such as tuning and routing.
- Developed via contract code changes and deployed RealVNC server cluster enhancements at LLNL to support load balancing and single sign-on authentication. Compatible client connectivity to RealVNC server deployed at LANL and SNL.
- Incorporated OneID's DUID to act as a foundation for future streamlined user authentication experiences among the tri-lab centers.

Planned Activities in FY22:

- Continue unclassified restricted tri-lab implementation for GA CI capability by a wider base of users. Demonstrate cross-site CI jobs in the unclassified restricted network in preparation for incoming ATS platforms.
- Conduct preliminary aspects of implementing tri-lab classified CI, including seeking security approvals, architecting solution, executing procurements pursuant to appropriate funding, and initial deployment steps.
- Leverage WAN line improvements made on DISCOM by implementing a similar 100Gb capability in the tri-lab iHPC, pursuant to funding; continue to enable iHPC accessibility of tri-lab resources.
- Procure and implement unclassified restricted data transfer infrastructure to model the previous classified environment advancements, pursuant to funding.
- Engage with tri-lab container portability efforts in furthering requirements understanding, build and run methods, and associated user-facing documentation.
- Continue tri-lab effort towards an authentication model that leverages a unified look and feel while allowing authentication call-backs to home institutions building on advancements in cross-realm trust and foundational alignments in tri-lab OneID DUID use.
- Continue to engage tri-lab SMEs for coordinated test and debug sessions needed for WAN network health, inclusive of unclassified ESnet and iHPC as well as classified ESN and DISCOM networks.

- Determine an RCE data storage model of understanding for large datasets at each site while leveraging the advancements currently being deployed in WAN networking.
- Continue to implement the classified DISCOM 100Gb solution to realize full tri-lab mesh production availability of this capability.
- Continue to refine and deploy VNC clients allowing users the convenience of persistent and performant remote desktops into HPC centers.
- Continue abating entry/egress bastion hosts/gateways in access paths of determined “security parity” across the tri-lab WAN.

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 FY21:

- Provided ongoing system maintenance and updates, supported customers in all aspects of the computing environment, perform software and hardware upgrades, and maintained codes on Special Purpose Computing systems
- Installed and refined one way data transfer system to speed data movements from LC/SCF environments to SCIF environments, allowing update frequency on the SCIF systems to move from quarterly to monthly.
- Replaced Active Directory server for special-purpose machine-1 environment, to provide authentication support to isolated environment.
- Made architectural improvements and hardware upgrades to Lustre environment.
- Researched needs for CTS-2 future procurements.

Planned Activities in FY22:

- Continue to provide ongoing system maintenance and updates, support customers in all aspects of the computing environment, perform hardware and software upgrades, and maintain codes on Special Purpose Computing systems.
- Plan for infrastructure improvements and modifications to support CTS-2 purchases in the FY23 to FY25 timeframe.
- Continue to increase redundancy, resiliency, and efficiency of operations in the Special Purpose Computing environment.

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 FY21:

- Deployed a backup file system for CTS-1 Global Security systems for increased availability.
- Supported high availability for CTS-1 Global Security systems.

Planned Activities in FY22:

- Deploy replacement file system for existing end-of-life file system supporting CTS-1 Global Security systems.
- Support any needed facility modifications for SCIFs to host CTS-2 hardware for Global Security.

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 other activities within the ASC program.

Accomplishments in FY21:

- Continued operations and support for tri-lab use of NSCC systems.
- Deployed a new high-performance tape library into the NSCC environment.
- Continue to work with NNSA and DOE/IN on platform refresh requirements and plans.

Planned Activities in FY22:

- Deploy a new centralized Lustre based file system.
- Continue operations and support for tri-lab use of NSCC systems.
- Complete integration of a new high-performance tape library into the NSCC environment.

- Continue to work with NNSA and DOE/IN on platform refresh requirements and plans.

Appendix H: 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**	**FY21**	**FY22**	**FY23**	**FY24**	**FY25**	**Future**	**Contact**
LLNL	Exascale Computing Facility Modernization (ECFM)	Line Item	CD-2/3 12/19/19	111,200	*80,000	30,200	*1,000	0	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	AT System – El Capitan	Capital Equipment, MIE	N/A	600,000	79,000	100,000	125,000	122,000	110,000	64,000	0	Bronis de Supinski 925-422-1062 desupinski1@llnl.gov
LLNL	B453 El Capitan Siting	Minor Construction	N/A	18,000	0	18,000			0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Unclassified CTS-2 System	Capital Equipment, MIE	N/A	10,200	0	0	200	2,000	2,000	2,000	4,000	Ned Bass 925-422-9389 bass6@llnl.gov
LLNL	Classified CTS-2 System	Capital Equipment, MIE	N/A	13,333	0	0	3,333	3,333	3,333	3,333	0	Ned Bass 925-422-9389 bass6@llnl.gov
LLNL	Unclassified El Capitan-like system	Capital Equipment, MIE	N/A	19,700	0	0	0	200	3,900	3,900	11,700	Ned Bass 925-422-9389 bass6@llnl.gov
LANL	Trinity: Acquisition of Trinity (ATS-1) system	Capital Equipment, MIE	CD-4	187,000	187,000	0	0	0	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov

*These numbers contain OPC costs.

LANL	Crossroads: Acquisition of Crossroads (ATS-3) system ($115 M) and water cooling and electrical power installation ($12 M)	Capital Equipment, MIE	N/A	127,000	48,000	45,000	14,000	8,000	6,000	6,000	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	ATS-5 system ($250 M) and water cooling ($18M) and electrical power installation ($12 M)	Capital Equipment, MIE	N/A	250,000 18,000 12,000	0 0 0	0 0 0	0 0 0	5,0000 0 0	25,000 0 0	50,000 4,000 1,000	170,000 14,000 11,000	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	SCC Electrical Upgrade	Minor Construction	N/A	18,400	0	0	18,400	0	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
SNL	Vanguard Phase 2: Acquisition of Vanguard Phase 2 Prototype System	Capital Equipment, MIE	N/A	15,000	2,000	8,000			0	0	0	Robert Hoekstra 505-844-7627 rjhoeks@sandia.gov
SNL	Additional CT Systems	Capital Equipment, MIE	N/A	20,000	0		5,000	5,000	5,000	5,000	0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725 Additional Power	Minor Construction	N/A	20,000	0	0	5,000	5,000	5,000	5,000	0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	ATS-ART System – (El Capitan nodes)	Capital Equipment, MIE	N/A	6,000	0	0			6,000		0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	ATS-ART System – (Crossroads nodes)	Capital Equipment, MIE	N/A	6,000	0	0	0	6,000	0	0	0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	Build & Test System	Capital Equipment, MIE	N/A	5,000	0	0	2,500	2,500		0	0	Tom Klitsner 505-844-1901 tklitsn@sandia.gov

Appendix I: FY22–FY26 Program Targets

This appendix lists the targets for each subprogram for the FY22–FY26 period as found in the *FY18–FY22 ASC Program Plan* and additional or updated targets for FY22–FY26. 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-3 and ATS-4. The target dates are FY23 for Crossroads and FY24 for El Capitan.

IC-2. Improved fidelity physics and numerics capabilities for capturing impact of microscopic variations due to both as-built and as-aged effects. The target date is FY23. Supports the following SCDS pegposts: FY22 Assess Lifetimes and Mitigate Aging, FY23 Enabling Efficient & Flexible Pit Production, FY25 Materials for Future Reentry Environments, FY29 On-Target Assessment & Mitigation.

IC-3. Advanced simulation capability for hostile environments. The target date is FY23. Supports the following SCDS pegposts: FY23 Hostile Mitigation, FY25 Combined Threat Environment Simulation, FY25 Materials for Future Reentry Environments.

IC-4. Improved responsiveness and analysis capabilities. The target date is FY23. Supports the following SCDS pegposts: FY21 Demonstrate Key Responsive Technologies, FY30 Enable 5-year Phase 1–5.

IC-5. Enhanced design of relevant experiments for validation, diagnostic design, and optimized data collection for HEDP. The target date is FY23. Supports the following SCDS pegposts: FY23 Hostile Mitigation, FY23 Enabling Efficient & Flexible Pit Production, FY25 Materials for Future Reentry Environments.

Physics and Engineering Models

PEM-1. Advance the major modeling capabilities identified in the FY20 Hostile Survivability Baseline pegpost to assure survivability in hostile environments in support of the follow-on FY23 Hostile Mitigation pegpost. The target date is FY23.

PEM-2. Advance modeling capabilities for aging and manufacturing assessments (including AM) that capture effects of microstructural features. The target date is FY24.

PEM-3. Extend physics models and simulation methodologies for evaluation of weapon performance and response in combined-physics normal and abnormal environments. The target date is FY25.

PEM-4. In support of AMLI, develop credible and interpretable machine learning (ML) toolkits to enable physics-constrained ML models with quantifiable uncertainties and holistic data assessments. The target date is FY25.

PEM-5. Create an infrastructure for foundational materials modeling and calibration that fully utilizes advanced features of next-generation architectures, machine learning techniques, and experimental full-field data. The target date is FY27.

Verification and Validation

V&V-1. Enable production, manufacturing and responsive weapons development timelines through predictive models, experimental collaborations and integrated V&V/UQ processes. Accelerated computational simulation workflows will be developed to increase turnaround of experiment, prototyping, and design. Predictive models supported by a comprehensive credibility evidence basis will be developed for production applications. Machine Learning informed workflows will be utilized as appropriate. The target is FY26, and intermediate work will support pegposts on the SCDS Modern Materials & Manufacturing and Future Deterrent strands from FY23 onward.

V&V-2. Establish test suites for rigorous Verification, Validation and Uncertainty Quantification of aging and behavior in combined environments to support current and future stockpile assessments. Test suites must integrate into existing workflows for supporting a broad customer base. The target is FY23, and earlier work will support the FY22 Assess Lifetimes and Mitigate Aging SCDS peg post.

V&V-3. Develop Engineering Common Model Frameworks with integrated Verification & Validation, Sensitivity Analysis, and Uncertainty Quantification. The goal is adaptable, efficient, standardized credibility processes and frameworks to support engineering simulation workflows and integration into physics workflows as appropriate. Analyses including thermal, structural, and hostile survivability simulations will be supported for the FY26 Combine Threat Environments Simulation SCDS peg post. The target is FY24.

V&V-4. Develop V&V/UQ framework incorporating ML methods as appropriate that enables agile, credible predictions suitable for stockpile applications. The target is FY24, and ongoing work will support prototype-to-design work in Modern Materials & Manufacturing and Future Deterrent SCDS strands.

V&V-5. Establish a V&V/UQ framework including workflows to support the credible application of Next-Generation Codes on current and emerging platforms. Invest in agile workflows with integrated V&V/UQ processes. Utilize efficient and appropriate UQ techniques that meet new customer needs including execution on novel heterogeneous architectures. The target is FY25 to support subsequent Stockpile Sustainment and Future Deterrent SCDS pegposts.

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. COMPLETED.

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

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

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

e. Close out the ATDM/ECP Software Technology deployment. The target date is FY23.

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

ATDM-3. Interagency Collaborations.

a. Complete Phase 1 of the DOE-National Cancer Institute (NCI) Collaboration Project. The target date is FY19. COMPLETED.

b. Complete Phase 2 (Year 4 and 5) of the DOE-National Cancer Institute (NCI) Collaboration Projects. The target date is FY21. COMPLETED.

c. Continue Year 6 of DOE-National Cancer Institute (NCI) Collaboration Projects. The target date is FY22.

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

Computational Systems and Software Environment

CSSE-1. Platform Acquisition and Deployment.

a. Procure and deploy Crossroads, including Center of Excellence (COE) activities. The target date for system acceptance for Phase 1 is Q3FY22 and classified production service in Q4FY22.

b. Procure and deploy El Capitan, including NRE activities. The target date for system acceptance is Q4FY23 and classified production service in Q2FY24.

c. Procure and deploy CTS-2 systems across tri-labs. The target dates are FY22–FY23.

d. Procure and deploy Vanguard-2 Advanced Prototype System. The target date is FY23.

CSSE-2. Software Environment Deployment.

a. Develop and deploy software and programming environments for Crossroads. The target date is FY21.

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

c. Deploy the Vanguard-2 system with a functioning software stack composed of ATSE, TOSS, and FOUS CCE tools and packages. The target date is FY22.

d. Develop and deploy software and programming environments for CTS-2. The target date is FY22 and FY23.

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, including initial evaluation of neuromorphic Intel Loihi testbed. The target date is FY21.

c. Demonstrate applications of interest on BML-funded hardware. The target date is FY23.

d. Demonstrate applications of interest on BML-funded hardware. The target date is FY25.Facility Operations and User Support

Facility Operations & User Support

FOUS-1. Installation and operation of El Capitan and Crossroads.

a. El Capitan will be deployed to users in FY24. The target date for installation and the beginning of operations is FY23.

b. The Crossroads system will be delivered, installed, and deployed in FY22. The target date for the start of installation preparation is FY21 with installation preparation completing in FY22

FOUS-2. Siting and Operation of Vanguard at-scale advanced technology prototypes. The first prototype system (the Arm prototype system named Astra) is in operation in the SNL classified environment. For Vanguard-2, exploratory technology systems will be deployed in FY22 and FY23, with the at-scale production system scheduled for deployment in FY24.

FOUS-3. Facility preparation for ATS-5 system. Operational improvements to expand the warm-water cooling system and electrical capacity at the Strategic Computing Complex (SCC) to enable up to 50 MW of supercomputing in preparation for ATS-5 system in FY27.

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–FY22.

FOUS-5. Installation and operation of Commodity Technology Systems (CTS) and development of the associated software stack. The CTS-1 platform line will enter its retirement phase in FY21. Planning for the CTS-2 deployments began in FY19. The initial target date for installation, deployment, and operation of CTS-2 is at the beginning of FY22. The target date for continued CCE software development and support is FY21–23.

Appendix J: Codes

Note: The content for the Codes appendix is available upon request from the ASC Program Office.