LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-852676

ASC FY24 Implementation Plan Revision 0

T. Hoang, R. Neely, J. Pruet, J. Gaudioso

August 2, 2023

Disclaimer

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

National Nuclear Security Administration

Advanced Simulation and Computing
FY24 IMPLEMENTATION PLAN

Version 0

August 8, 2023

NA-ASC-127R-23-Rev.0-IP

LLNL-TR-852676

National Nuclear Security Administration

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	Rob Neely, Jason Pruet, Jennifer Gaudioso	Initial	Thuc Hoang	08/08/23	Initial release

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

National Nuclear Security Administration

Contents

I. Overview

The DOE National Nuclear Security Administration (NNSA) Stockpile Stewardship Program (SSP) is an integrated technical program for maintaining the safety, security, 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 previous nuclear tests, 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 nuclear counterterrorism.

The 2022 Nuclear Posture Review (NPR) calls for NNSA to "deliver a modern, adaptive nuclear security enterprise based on an integrated strategy for risk management, production-based resilience, science and technology innovation, and workforce initiatives."[1] Furthermore, "NNSA will establish a Science and Technology Innovation Initiative to accelerate the integration of science and technology (S&T) throughout its activities."[2] Executing this strategy necessitates the continued emphasis on developing and sustaining high-quality scientific and engineering staff, as well as supporting computational and experimental capabilities. These components constitute the foundation of the nuclear weapons program.

[1] *2022 Nuclear Posture Review* October 2022, p.3.

[2] *2022 Nuclear Posture Review*, October 2022, p. 24.

The continued success of the SSP and LEPs is predicated upon the ability to credibly certify the stockpile, without a return to underground nuclear tests (UGTs). Shortly after the nuclear test moratorium entered into force in 1992, the Accelerated Strategic Computing Initiative (ASCI) was established to provide an extensive simulation capability to underpin 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 stewardship mission. The ASCI program was essential to the successful demonstration 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 life 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 FY24. The Initial IP, effective August 9, 2023, should be consistent with the Department's Base Table when operating under a Continuing Resolution (CR). The final IP, effective <DATE>, is consistent with the final, enacted appropriation.

National Nuclear Security Administration

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	FY2024 Target	Endpoint Target
Advanced Simulation and Computing Program	Nuclear Weapons Simulation Capability Progress*	74%	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 FY23 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, including 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 *FY23 Stockpile Stewardship and Management Plan*[3]. 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 five 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 counterterrorism 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

[3] U.S. Department of Energy, National Nuclear Security Administration, *Fiscal Year 2023 Stockpile Stewardship and Management Plan*, Report to Congress, April 2023.

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 **Computational 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 environments. Along with the powerful advanced architecture prototype (AAP), 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 the input/output of data (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 technologies to include quantum, neuromorphic, function-optimized silicon designs and non-complementary metal-oxide-semiconductor (CMOS)-based components and systems.

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 five subprograms (IC, PEM, V&V, CSSE, and FOUS) all contribute to a cohesive set of program deliverables. Highlights of the FY24 major activities for the ASC program include:

- Deploy and support capabilities to help qualify and certify the W80-4 warhead.
- Develop and mature capabilities to support design, qualification, and certification for the W87-1 Program.
- Develop and mature capabilities to support design, qualification, and certification for the W93 Program.
- 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.
- Continue to improve Pu physics models as a function of age and investigate bounding models that would constrain the possible changes to physics as a function of age to support initial deployment of age aware models in FY25.
- 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.
- Implement engineering common model frameworks to enhance common modeling techniques for ASC capabilities.
- Integrate V&V/UQ test suites into existing workflows for supporting a broad customer base.
- Adapt existing codes to new architectures and migrate current design, engineering, and safety codes to run efficiently on hybrid computer architectures.
- Accept and transition the Crossroads platform for classified service. The system will be used to address stockpile stewardship issues and support advanced weapons-relevant prediction science.
- Deliver and accept El Capitan, the NNSA's first Exascale platform. El Capitan will deliver significant performance improvements for ASC codes, enabling new levels of fidelity in complex mission-relevant modeling and simulation activities.
- Connect and optimize remote computing capabilities to support full tri-lab access to the Crossroads and El-Capitan HPC platforms.
- Deploy CTS-2 scalable units (SUs) which will be used to support a broad range of ASC and stockpile stewardship simulation workloads.

- Continue to provide computing infrastructure investments, including minor construction facility modernization efforts, that are required to prepare for the next-generation HPC platforms.
- Transition production-ready, next-generation codes and computing technologies from ATDM to the other ASC subprograms (IC, PEM, V&V, and CSSE). Continue the DOE-NCI (National Cancer Institute) collaboration project in CSSE.
- 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.
- 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.
- Maintain mentoring program for early career staff.

The drivers of the ASC program that require the FY24 budget include the following:

- The Nuclear Weapons Council approved the stockpile strategy, 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, investigate the feasibility of fielding the nuclear explosive package in a Navy flight vehicle, and sustaining the B83-1 bomb past its previously planned retirement date.
- Additionally, NNSA will explore future 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 functions. 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 are 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 missions 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 artificial intelligence/machine learning (AI/ML) 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 FY2024 Budget Request**	**FY2024 Continuing Resolution Base Table**	**FY2024 Enacted/Full Year<Budget or CR>**	**Difference between Request versus Enacted**
ASC Operations	$782.472M	$xxxM	$xxxM	$xxxM
Academic Programs (for PSAAP & CSGF)	$23.830M	$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 2 Milestones for FY24. 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 FY24–FY28 Target column may be found in Appendix I.

Table V-1B. ASC Level 2 Milestones for FY24

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY24–FY28 Target
IC	TBD	Capabilities in support of shallow bubble collapse integrated experiment team	09/30/24	LLNL	IC-2
IC	TBD	Demonstrating relevant simulation capability on Crossroads (ATS-3)	09/30/24	LLNL	IC-1
IC	TBD	LANL Production Readiness	09/30/24	LANL	IC-1
IC	TBD	Demonstrating relevant simulation capability on Crossroads (ATS-3)	09/30/24	LANL	IC-1
IC	TBD	Demonstration of a common infrastructure and workflows for performance measurement for Integrated Codes	09/30/24	LANL	IC-1
IC	TBD	Deploy Sierra applications to ATS-3 and ATS-4	9/30/24	SNL	IC-1
IC/ CSSE	TBD	Conceptual design of data model for radiation-electrical system survivability	09/30/24	SNL	IC-4, CSSE-5
IC/ CSSE/ V&V /PEM	TBD	Develop conceptual design for ModSim digital thread	09/30/24	SNL	IC-4, CSSE-2, V&V-1, V&V-3, V&V-5, PEM-5
PEM	TBD	Multi-scale modeling with alloying effects	09/30/24	LLNL	PEM-2, PEM-5

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY24–FY28 Target
PEM	TBD	Deliver a detonation model based on subscale hot spot physics	09/30/24	LLNL	PEM-1, PEM-2, PEM-3
PEM	TBD	Porting nuclear theory codes on GPU-enabled architectures	09/30/24	LLNL	PEM-5
PEM	TBD	UQ-based Nuclear Data Prioritization Assessment	09/30/24	LANL	PEM-3,4
PEM	TBD	Hierarchic stochastic material radiation model capability within CHEETAH-MC	09/30/24	SNL	PEM-2, PEM-5
PEM	TBD	Constitutive model development for glass ceramic materials	09/30/24	SNL	PEM-3
V&V	TBD	Assess GNDS/GIDI as default nuclear data	09/30/24	LLNL	V&V-5
V&V	TBD	Integration of experimental data into modern workflows	09/30/24	LANL	V&V-3
V&V	TBD	Verification Testing for ASC codes	09/30/24	LANL	V&V-1, V&V-2, V&V-3, V&V-5
V&V/ IC/ CSSE/ FOUS	TBD	ND workflow development and roadmap for secure, scalable data management	09/30/24	SNL	V&V-1, V&V-3, V&V-5, IC-3, IC-4, CSSE-2, FOUS-5
V&V	TBD	Develop and incorporate novel code verification of fundamental equations in Gemma and setup appropriate tests	09/30/24	SNL	TBD
V&V	TBD	Assemble credibility evidence for fire modeling capabilities	09/30/24	SNL	V&V-1, V&V-2
CSSE	TBD	Flux and Rabbit capabilities on El Capitan	08/31/24	LLNL	CSSE-1.b
CSSE	TBD	Machine-learned slope limiters	09/30/24	LANL	CSSE-3.d
CSSE	TBD	Fully Integrate CONDUIT	09/30/24	LANL	CSSE-1.a, FOUS-1.b
CSSE	8539	Crossroads Production Readiness	03/29/24	LANL, SNL	CSSE-1.a, FOUS-1.b
CSSE	TBD	Advancing the utility of Data Flow Accelerators for NNSA computing systems	09/30/24	SNL	CSSE-1.d
CSSE/ FOUS	TBD	Tri-labs investigation of a Unified Software Environment Framework for Future ASC Production Platforms	09/30/24	SNL, LANL, LLNL	CSSE-2

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY24–FY28 Target
CSSE/ FOUS	TBD	Tri-labs investigation of Large Language Models	09/30/24	SNL, LANL, LLNL	CSSE-2
FOUS	TBD	El Capitan compute node integration	09/30/24	LLNL	FOUS-5
FOUS	TBD	SCC Electrical Upgrade Start of Construction	9/30/2024	LANL	FOUS-3
FOUS	TBD	El Capitan ART/Mini System Acquisition and Delivery	09/30/24	SNL	FOUS-1
FOUS	8640	CTS-2 delivery, site integration and deployment	12/31/23	SNL	FOUS-5

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:

- *ASC Strategy 2021–2031* (August 2021)
- *ASC Computing Strategy* (July 2022)
- *ASC Right Size* (September 2016)
- *ASC Business Plan* (July 2015)
- *ASC Co-Design Strategy* (February 2016)
- *ASC Simulation Strategy* (July 2022)
- *ASC Artificial Intelligence for Nuclear Deterrence (AI4ND) Strategy* (August 2023)
- 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 NA-114 Program Points of Contact

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

K. Michael Lang
Headquarters Implementation Plan Focal Point
Office: 301-903-0240
Michael.Lang@nnsa.doe.gov

Laboratory ASC Executives

Robert Neely, LLNL, 925-423-4243, neely4@llnl.gov

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

Jennifer Gaudioso, SNL, 505-284-9489, jmgaudi@sandia.gov

Laboratory Programmatic Points of Contact

WBS	Title	Contact
1.2.3.1	Integrated Codes	James Peltz, NNSA, 202-586-7564, James.Peltz@nnsa.doe.gov 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	Anthony Lewis, NNSA, 202-287-6367, Anthony.Lewis@nnsa.doe.gov Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jill Blecke, 505-844-1997, jblecke@sandia.gov
1.2.3.2.1	Materials Response: Equation of State, High Explosives, and Mix and Burn	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jill Blecke, 505-844-1997, jblecke@sandia.gov

WBS	Title	Contact
1.2.3.2.2	Transport, Plasmas, Atomic, Nuclear	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jill Blecke, 505-844-1997, jblecke@sandia.gov
1.2.3.2.3	Engineering Science, Environments, and Response	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jill Blecke, 505-844-1997, jblecke@sandia.gov
1.2.3.2.4	Integrated Modeling and Applications	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Manolo Sherrill, LANL, 505-665-8559, manolo@lanl.gov Jill Blecke, 505-844-1997, jblecke@sandia.gov
1.2.3.3	Verification and Validation	David Etim, NNSA, 202-586-8081, David.Etim@nnsa.doe.gov Ana Kupresanin, LLNL, 925-422-8348, kupresanin1@llnl.gov Catherine Plesko, LANL, 505-667-2345, plesko@lanl.gov Dan Turner, SNL, 505-845-7446, dzturne@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 Dan Turner, SNL, 505-845-7446, dzturne@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 Dan Turner, SNL, 505-845-7446, dzturne@sandia.gov

WBS	Title	Contact
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 Dan Turner, SNL, 505-845-7446, dzturne@sandia.gov
1.2.3.4.1	Next-Generation Architecture and Software Development	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Timothy Randles, LANL, 505-667-0680, trandles@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5	Computational Systems and Software Environment	Simon Hammond, NNSA, 202-586-5748, simon.hammond@nnsa.doe.gov Matt Legendre, LLNL, 925-422-6525 legendre1@llnl.gov Timothy Randles, LANL, 505-667-0680, trandles@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 Timothy Randles, LANL, 505-667-0680, trandles@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov

WBS	Title	Contact
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 Timothy Randles, LANL, 505-667-0680, trandles@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.5	Post-Processing Environments	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov Timothy Randles, LANL, 505-667-0680, trandles@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 Timothy Randles, LANL, 505-667-0680, trandles@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.7	Next-Generation Computing Technologies	Matt Legendre, LLNL, 925-422-6525 legendre1@llnl.gov Timothy Randles, LANL, 505-667-0680, trandles@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.6	Facility Operations and User Support	K. Michael Lang, NNSA, 301-903-0240 Michael.Lang@nnsa.doe.gov 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 Justin Newcomer, SNL, 505-284-9270, jtnewco@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

WBS	Title	Contact
1.2.3.6.3	Common Computing Environment	Ned Bass, LLNL, 925-422-9389, bass6@llnl.gov Nicholas Jones, LANL, 505-667-9087, njones@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 Jennifer Gonzales, LANL, 505-667-3480, jemx@lanl.gov Stephen Monk, SNL, 505-284-2811, smonk@sandia.gov

PSAAP III Points of Contact	
University at Buffalo, Center for Exascale Simulation of Hybrid Rocket Motors	Prof. Paul Desjardin ped3@buffalo.edu (716) 645-1467 Kelsey Trautwein (Business Analyst) kgtrautw@buffalo.edu (716) 645-1806
University of Colorado, Boulder, Center for Micromorphic Multiphysics Porous and Particulate Materials Simulations with Exascale Computing Workflows	Prof. Richard Regueiro Richard.regueiro@colorado.edu (303) 492-8026 Chris Pappas (Finance Manager) Chris.pappas@colorado.edu (303) 492-9073
University of Illinois, Urbana-Champaign, Center for Exascale-Enabled Scramjet Design	Prof. Jonathan Freund jbfreund@illinois.edu (217) 244-0720 Richard D. Mihm (Manager of Accounting) rmihm@illinois.edu (217) 244-8886
University of Maryland, Solution-Verification, Grid-Adaption and Uncertainty Quantification for Chaotic Turbulent Flow Problems	Prof. Johan Larsson jola@umd.edu (301) 405-5273 Joeshia Brawner (Account) Jbrawne1@umd.edu (301) 405-6960

PSAAP III Points of Contact	
Massachusetts Institute of Technology, Center for the Exascale Simulation of Material Interfaces in Extreme Environments	Prof. Youssef Marzouk ymarz@mit.edu (617) 253-1337 Jean Sofronas (Administrative Assistant) jeans@mit.edu (617) 258-5548
University of New Mexico, Center for Understandable, Performant Exascale Communication Systems	Prof. Patrick Bridges bridges@cs.unm.edu (505) 277-3032 Galen Billings (Administrative Assistant) gbillings@unm.edu (505) 277-3112
Oregon State University, Center for Exascale Monte Carlo Neutron Transport	Prof. Todd Palmer Todd.palmer@oregonstate.edu (541) 737-7064 Matthew Berry (Accountant) Matthew.berry@oregonstate.edu (541) 737-4504
Stanford University, Integrated Simulations using Exascale Multiphysics Ensembles	Prof. Gianluca Iaccarino jops@stanford.edu (650) 906-0302 Corinne Beck (Administrative Associate) corinnebeck@stanford.edu (650) 725-0704
University of Texas, Austin, Exascale Predictive Simulation of Inductively Coupled Plasma Torches	Prof. Robert Moser rmoser@oden.utexas.edu (512) 471-3168 Lisa Gentry (Administrative Support) lgentry@oden.utexas.edu (512) 471-3168

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 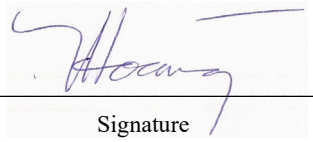
Signature 08/08/23

Concurrence:

Rob Neely, LLNL ASC Executive	*Concurrence received*	08/07/23
Jason Pruet, LANL ASC Executive	*Concurrence received*	08/08/23
Jennifer Gaudioso, SNL ASC Executive	*Concurrence received*	08/08/23

National Nuclear Security Administration

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
AAP	Advanced Architecture Prototype
AAPS	Advanced Architecture and Portability Specialist or Advanced Architecture Prototype System
AAR	Annual Assessment Review
ACES	New Mexico Alliance for Computing at Extreme Scale
ADE	Advanced Digital Engineering
ADEPT	Applications Development Environment and Performance Team
AES	Advanced Encryption Standard
AFP	Approved Funding Program
AI	Artificial Intelligence
ALCF	Argonne Leadership Computing Facility
ALE	Arbitrary Lagrangian-Eulerian
Alt	Alteration
AMD	Advanced Micro Devices, Inc.
AML	Advanced Machine Learning
AMR	Adaptive Mesh Refinement
AMT	Asynchronous Many Task or Advanced Memory Technologies
ANL	Argonne National Laboratory
ANS	American Nuclear Society
API	Application Programming Interface
APS	Advanced Prototype System
ART	Application Regression Testbed or Application Readiness 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
ASIC	Application Specific Integrated Circuit
ASP	Adaptable Storage Platform
AT	Advanced Technology
ATS	Advanced Technology System
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
BOP	Business Operating Procedure
C2C	Contours to Codes project
CAD	Computer Aided Design or Computer Assisted Design
CCE	Common Computing Environment
CD	Critical Decision
CDA	Next-Generation Code Development and Applications
CEED	Center for Efficient Exascale Discretizations
CEESD	Center for Exascale-Enabled Scramjet Design
CFD	Computational Fluid Dynamics
CGNS	Computational fluid dynamics General Notation System
CHAI	Copy Hiding Application Interface
CI	Continuous Integration
CMF	Common Model Framework
CMM	Coordinate Measuring Machine
CMOS	Complementary Metal-Oxide-Semiconductor
COE	Center of Excellence
CORAL	Collaboration of Oak Ridge, Argonne, and Livermore

CONDUIT	Capacity on Demand User Interface and Toolkit
CPEP	Corporate Performance Evaluation Process
CPU	Central Processing Unit
CR	Continuing Resolution
CRAC	Computer Room Air Conditioning
CRADA	Cooperative Research and Development Agreement
CSGF	Computational Science Graduate Fellowship
CSSE	Computational Systems and Software Environment
CSV	Comma Separated Values
CT	Commodity Technology
CTI	Common Tools Interface
CTS	Commodity Technology System
DA	Design Agency
DARMA	Distributed Asynchronous Resilient Models for Applications
DDN	DataDirect Networks
DDR	Double Data Rate
DFT	Density Functional Theory
DIC	Digital Image Correction
DIMM	Dual Inline Memory Module
DISA	Defense Information Systems Agency
DisCom	Distance Computing
DNS	Direct Numerical Simulation
DOD	Department of Defense
DOE	Department of Energy
DP	Defense Programs
DRAM	Dynamic Random Access Memory
DSI	Data Science Infrastructure
DSMC	Direct Simulation Monte Carlo
DUID	DOE Unique Identifier
DYAD	Dynamic and Asynchronous Data
EAP	Eulerian Application Project

EAS	Early Access System
EC3E	Exascale Class Computer Cooling Equipment (LANL)
ECFM	Exascale Computing Facility Modernization (LLNL)
ECI	Exascale Computing Initiative
ECMF	Engineering Common Modeling Framework
ECP	Exascale Computing Project
EMPIRE	Electro-Magnetic Plasma in Radiation Environments
EOS	Equation of State
EQMU	Engineering Quantification of Margins and Uncertainties
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
FY	Fiscal Year
GDDR	Graphics Double Data Rate
GPFS	General Parallel File System
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
HEDP	High-Energy Density Physics
HI	Hardware and Integration (ECP)
HPC	High Performance Computing
HPSS	High Performance Storage System
HQ	Headquarters

I/O	Input/Output
IB	InfiniBand
IC	Integrated Codes
ICD	Interagency Co-Design
ICF	Inertial Confinement Fusion
ICR	Independent Cost Review
IDC	Integrated Design Code
IEP	Integrated Exascale Project
IHPC	Integrated High-Performance Computing (network)
IOSS	Input/Output SubSystem
IP	Implementation Plan
ITS	Integrated Tiger Series
ITSM	Information Technology Service Management
IWF	Integrated Workflow
JDACS4C	DOE-NCI Joint Design of Advanced Computing Solutions for Cancer
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
LBANN	Lawrence Big Artificial Neural Network Toolkit
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LCA	Light Convolutional Autoencoder
LCW	Low Conductivity Water
LDMS	Lightweight Distributed Metric Service
LEP	Life Extension Program
LES	Large Eddy Simulation

LLM	Large-Language Model
LLNL	Lawrence Livermore National Laboratory
LLVM	Low-Level Virtual Machine
LSCI	Large Scale Calculations Initiative
M&O	Management and Operations
MD	Molecular Dynamics
MFA	Multi-Factor Authentication
MFEM	Modular Finite Element Methods
MIR	Material Interface Reconstruction
ML	Machine Learning
MLBLUE	Multilevel Best Linear Unbiased Estimators
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
NASEM	National Academics of Sciences, Engineering and Medicine
NCHPC	Non-conventional High-Performance Computing
NCI	National Cancer Institute
ND	Nuclear Deterrence
NDA	Non-disclosure Agreement
NERSC	National Energy Research Scientific Computing Center
NFS	Network File System
NGCE	Next-Generation Computing Enablement
NGCT	Next-Generation Computing Technologies
NG-HPCN	Next-Generation High-Performance Computer Networking
NGP	Next-Generation Platform
NGW	Next-Generation Workflow

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
OHPC	Open High-Performance Computing (network)
OLCF	Oak Ridge Leadership Computing Facility
OMPD	OpenMP Debugging Interface
OMPT	OpenMP Tools Interface
ORNL	Oak Ridge National Laboratory
OS	Operating System
OS/R	Operating System/Runtime
PA	Production Agency
PDU	Power Distribution Unit
PE	Programming Environment
PEM	Physics and Engineering Models
PESP	Predictive Engineering Science Panel
PIC	Particle-in-Cell
PIEP	Pairwise Interaction Extended Point Particle Model
PINN	Physics-Informed Neural Network
PNNL	Pacific Northwest National Laboratory
PPA	Programs, Projects, or Activities
PSP	Predictive Science Panel
PSAAP	Predictive Science Academic Alliance Program

PX	Pantex Plant
QC	Quantum Computing
QIP	Quantum Information Processing
QMU	Quantification of Margins and Uncertainties
QoS	Quality of Service
R&D	Research and Development
RAM	Random Access Memory
RCE	Remote Computing Enablement
RDMA	Remote Direct Memory Address
RE	Restricted Enclave
RFI	Request for Information
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
ROL	Rapid Optimization Library
ROM	Reduced Order Model
RR	Restricted Region
RVMA	Remote Virtual Memory Access
RZ	Restricted Zone
SAN	Storage Area Network
SAP	Scalable Applications Preparation
SARAPE	Synchronized Account Request Automated Process
SASI	Standards and Architectures for Storage and I/O
SAW	Sandia Analysis Workbench
SC	Department of Energy's Office of Science
SCC	Strategic Computing Complex (LANL)
SCDS	Stewardship Capability Delivery Schedule
SCF	Secure Computing Facility
SCN	Sandia Classified Network
SCR	Scalable Checkpoint Restart
SDM	Simulation Data Management
SDN	Software Defined Network

SFI	Significant Finding Investigation
SLURM	Simple Linux Utility for Resource Management
SNAP	Sn (Discrete Ordinates) Application Proxy
SNL	Sandia National Laboratories
SNSI	Secret National Security Information
SOS	Scalable Object Store
SOW	Statement of Work
SPOT	System for Performance Optimization Tracking
SRD	Secret Restricted Data
SPARC	Sandia Parallel Aerodynamics Re-entry Code
SRN	Sandia Restricted Network
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAT	Stack Trace Analysis Tool
STIG	Security Technical Implementation Guide
SU	Scalable Unit
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
UVM	Unified Virtual Memory
V&V	Verification and Validation
VNC	Virtual Network Computing
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
ZIA	ZFS Interface for Accelerators

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: 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 the 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 computer science technologies for ASC's exascale and post-exascale computing environments.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2022, and through quarter 3, fiscal year 2023, are reflected below for the CSSE subprogram.

- Negotiated modifications to El Capitan build and CORAL-2 NRE contracts including successful "Go" decision to proceed with system procurement from HPE. (Appendix I, target CSSE-1.b) (LLNL)
- Deployed significant portions of the El Capitan system, including infrastructure racks, the capacity storage tier and delivery of the system compute cabinets. Deployed multiple A0 compute nodes to enable further tri-lab system software preparations for El Capitan. (Appendix I, target CSSE-1.b) (LLNL)
- Deployed and supported Flux as the system resource manager on all three CORAL-2 EAS3 machines, and as system resource manager with initial Rabbit support on the Hetchy development system. (Appendix I, target CSSE-1.b) (LLNL)
- Deployed application development environments for El Capitan using early access systems (RZvernal, Tioga, and Tenaya) including compilers, tools, MPIs, and documentation. (Appendix I, targets CSSE-1.b and CSSE-2.b) (LLNL)
- Deployed Tycho (Crossroads Phase-I) system including transition to classified service. The system is now supporting build, test, and development of ASC IC applications ready for campaign use in FY24. (LANL)

- Deployed initial hardware deliveries of the Crossroads Phase-II system, performing initial high-performance interconnect and compute node testing. Acceptance of the system is planned for the first half of FY24, with transition to classified service later that year. (LANL)
- Developed Spack features including better support for HPE/Cray systems, software testing, version handling, and packages with multiple build systems. The Spack CI system and binary cache were scaled out to ~4,600 packages and regularly perform 50-100 thousand builds per week. (Appendix I, Targets IC-1 and CSSE-2) (LLNL)
- Released BEE 0.1.4 to the public on Github and through pip (Python package manager) giving users the ability to manage multiple workflows using BEE on any production systems without administrative privileges or arcane developer environments. (LANL)
- Released BEE 0.1.4 to the public on github and through pip (Python package manager) giving users the ability to manage multiple workflows using BEE on any production systems without administrative privileges or arcane developer environments. (LANL)
- Released Kokkos 4.0 with support for the latest AMD ROCm and NVIDIA CUDA programming tools and environment. (SNL)
- Received first stage of Loihi 2 deliverable from Intel, have successfully set up in SNL-NERL laboratory. (SNL)
- Created software for synthesizing application-driven volumetric benchmark circuits. This has been applied on IBM testbeds and it will be applied to other platforms in FY24 and beyond. (SNL)
- Worked with SambaNova to train a Large Language Model (LLM) for Kokkos code generation. Using OpenAI GPT models, pretraining with selected 20+ Github Kokkos repos. (SNL)
- Deployment of over $30M of contracts to multiple US vendors through a tri-lab procurement activity supporting R&D for innovative new memory technologies. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Flux and Rabbit Capabilities on El Capitan		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 08/31/24		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-1.b		
Description: This effort will build and demonstrate the Rabbit management software and Flux support that will be used on El Capitan. The El Capitan I/O team will develop containerized Rabbit orchestration software that will be deployed to a TOSS-based Kubernetes environment and will manage the system's Rabbit IO nodes. Flux will be extended to interface with the Rabbit control software allowing users to schedule storage for any of the supported I/O modes as part of the compute job. Combined, these capabilities will enable El Capitan users to manage and scale their application's IO workloads on El Capitan.		
Completion Criteria: Flux will demonstrate scheduling and running an application of interest to ASC that utilizes Rabbits for I/O. This demonstration will be scaled relative to the hardware available at the time, show that jobs can be configured and scheduled to utilize Rabbit storage, and will exercise the different Rabbit I/O modes. A report will be generated that shows the state of the Rabbit software stacks and the Flux scheduler on the initial El Capitan deliveries, how I/O scales on the system, and will identify any work that remains.		
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: ASC FOUS and IC collaborators		

Milestone (ID#TBD): Machine-learned Slope Limiters		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2024		
ASC WBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target [CSSE-3.d]		
Description: Assess the potential of using machine-learning to develop slope limiters for use in production codes on mission-relevant calculations. Development will be done in RIOT, an emulator for algorithms employed in the production code xRAGE. The governing equations for compressible flows, the Euler equations, will be targeted in this milestone. Optimization will be driven by the scientific machine learning framework Mystic, using a proven differential evolution algorithm.		
Completion Criteria: 1. Score slope limiters by accuracy, boundedness, positivity, and robustness. 2. Assess above scores over the range of slope limiters, including machine-learned limiters. Evaluate their effectiveness on a suite of canonical test problems. This assessment will be used to decide if implementation in xRAGE is warranted. 3. Prepare a report outlining the results of this study.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: LANL CSSE staff		

Milestone (ID#TBD): Fully Integrate CONDUIT		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2024		
ASC WBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-1.a., FOUS-1.b.		
Description: Full integration of the Capacity On Demand User Interface and Toolkit (CONDUIT), a system for rapidly and atomically shuttling data between storage tiers enabling simpler workflows and lowered storage costs. CONDUIT will further improve data transfers by interacting with job submissions to determine priority of data and adjust resources to ensure timely data migration. Demonstrate the optimized transfer of data between campaign and scratch storage as managed by the workload manager (SLURM) on Crossroads.		
Completion Criteria: A fully integrated, automated, data transfer capability by CONDUIT will be available for all users on Crossroads. This will enable a command line interface, job submission interface, and the optimized data transfer capability. A complete user-guide including a description of expected behavior for successful workflows will be provided.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: LANL CSSE and FOUS staff		

Milestone (ID#8539): Crossroads Production Readiness		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 03/29/2024		
ASC WBS Subprogram: CSSE, FOUS		
Participating Sites: LANL, SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-1a, FOUS-1		
Description: Certify that the Crossroads system meets the requirements for production readiness of an advanced technology computing class system. These requirements include system functionality, system performance, system accessibility and integration, operational support, a demonstrated usage model, user application testing, system reliability, benchmarks, I/O performance testing, and a milestone review.		
Completion Criteria: Follows the ASC Level 2 milestone criteria for Advanced Technology systems: that all of the topics identified in the description above have been successfully demonstrated for advanced technology computing class simulations. These requirements are specifically listed in the usage model for Crossroads, which defines that the system has demonstrated an acceptable production user environment with all the associated support, testing, reliability, and applications use of the system.		
Customer: NNSA/ASC HQ, tri-lab ASC program managers responsible for ATCCs, SSP, tri-lab weapons applications community		
Milestone Certification Method: A milestone review featuring committee members from the tri-lab 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, Crossroads		

Milestone (ID#TBD): Advancing the utility of Data Flow Accelerators for NNSA computing systems		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/24		
ASC WBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-1d		
Description: This milestone will provide a higher readiness level for running NNSA mini-apps and initial mission applications on a data-flow accelerator. The premise of this milestone is to enable and demonstrate a new Kokkos back-end that directly drives an underlying dataflow accelerator. The milestone will go beyond the use of existing Kokkos backends (such as OpenACC or OpenMP) and develop a prototype Kokkos backend implementation in tight coordination with a relevant data-flow vendor(s). Successful completion of this milestone will include both a Kokkos backend prototype and a demonstration of its use in several mini-apps. A detailed performance analysis investigating FOMs will be compared to existing activities as a stretch goal. Documentation of best practices and a plan for follow-on work needed for production-level support of data-flow architectures will also be provided.		
Completion Criteria: Demonstration of Kokkos backend for data-flow accelerator; documentation of best practices and requirements for production-level support of data-flow accelerators.		
Customer: NNSA/ASC HQ, tri-lab ASC program managers		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation consisting of a report will be prepared as a record of milestone completion.		
Supporting Resources: Tri-lab CSSE & FOUS		

Milestone (ID#TBD): Tri-labs investigation of a Unified Software Environment Framework for Future ASC Production Platforms		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2024		
ASC WBS Subprogram: CSSE, FOUS		
Participating Sites: SNL, LANL, LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-2		
Description: The ASC program recognizes that a bottleneck to deploying codes on HPCplatforms is the challenge of disparate software environments on those platforms. This milestone will evaluate and analyze strategies for achieving greater uniformity of user software environments for current and future ASC platforms. The milestone will develop recommendations for ASC on how this may be achieved. Recommendations will take into account how software components from the HPC community (including open-source and vendor products) can be included or tailored/customized for the ASC environment, and what application user and developer requirements can be addressed with the recommendations.		
Completion Criteria: Strategy and implementation document describing recommendations for a unified software environment for future ASC systems.		
Customer: NNSA/ASC HQ, tri-lab ASC program managers		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation consisting of a report will be prepared as a record of milestone completion.		
Supporting Resources:		

Milestone (ID#TBD): Tri-labs investigation of Large Language Models		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2024		
ASC WBS Subprogram: CSSE, FOUS		
Participating Sites: SNL, LANL, LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target CSSE-2		
Description: : Large-language models (LLMs) have an immense potential to revolutionize how data is identified, reasoned about, and analyzed. There are numerous instances where such tools could make drastic differences to the capabilities of ASC both in classified and unclassified environments. In this milestone tri-labs team will evaluate at least one vendor supplied LLM and at least one open-source LLVM for possible uses in processing and analyzing ASC data. The milestone will lead to a set of recommendations on the potential future uses of LLMs within ASC and will document any deficiencies or areas of concern where future research may need to be conducted to improve quality, mission-tailoring, accuracy, scalability, or security.		
Completion Criteria: Evaluation report by tri-lab team of the potential use of LLMs for ASC and broader ND mission capability.		
Customer: NNSA/ASC HQ, tri-lab ASC program managers		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation consisting of a report will be prepared as a record of milestone completion.		
Supporting Resources: IC		

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 systems relative to best practices in Linux cluster deployments today. This objective strives to quickly make these systems robust, useful production clusters for ASC scientific simulation workloads.

Accomplishments in FY23:

- Extensively tested and resolved issues on production CT systems. (Appendix I, Targets CSSE-1.c and FOUS-5)
- Assisted users of CT systems. (Appendix I, Targets CSSE-1.c and FOUS-5)
- Led and managed the tri-lab CTS-2 contract and FY23 system procurements, system delivery schedules, and supply chain issues. (Appendix I, Target CSSE-1.c)
- Fielded deliveries of FY23 CTS-2 production platforms. (Appendix I, Target CSSE-1.c)
- Procured additional CTS-2 platforms. (Appendix I, Target CSSE-1.c)

Planned Activities in FY24:

- Field deliveries of CTS-2 production platforms. (Appendix I, Target CSSE-1.c)
- Procure additional CTS-2 platforms. (Appendix I, Target CSSE-1.c)
- Provide production support for CT systems (CTS-1 and CTS-2). (Appendix I, Targets CSSE-1.c and FOUS-5)

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 FY23:

- Provided programmatic oversight for CT systems
- Continued participation in the CTS-2 procurement activities. (Appendix I, target CSSE-1c, FOUS-5)

Planned Activities in FY24:

- Continue to provide programmatic oversight of CT systems.
- Coordinate the activities for CTS-2 cycles (Appendix I, target CSSE-1c, 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 FY23:

- Deployed initial CTS-2 systems on the restricted network. Working with CTS-2 vendor and tri-labs to stabilize and release systems for production use (Appendix I, targets CSSE-1c and FOUS-5)

Planned Activities in FY24:

- Continue deployment of CTS-2 systems to restricted and classified networks and transition to production use (Appendix I, targets CSSE-1c and FOUS-5)

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 Sierra, a system of over 125 petaflops, with the bulk of the capability delivered from 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 is 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 FY23:

- Completed two ATCC processes.
- Ported TOSS4 to CORAL-1 systems in preparation for security updates.

Planned Activities in FY24:

- Run two additional ATCC processes.
- Deploy TOSS4 to CORAL-1 systems.
- 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 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. LLNL will procure El Capitan, a next-generation supercomputer, for which delivery and deployment began in FY23 and will be completed with system acceptance in FY24. Under the El Capitan procurement, LLNL has worked with key technology providers to deploy an El Capitan system of significantly over 2 exaflops peak capability to advance science and to ensure national security. 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 increases 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. Subsequent additional contract modifications have mitigated risk and increased commonality with the system software to be deployed on El Capitan by formally moving to the use of TOSS (the Tri-Lab Operating System Software), which is also used on ASC Commodity Technology Systems, culminating in the Go contract modification that recognizes the successful completion of the technical checkpoint that converts all target requirements to firm requirements. 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. Early access systems will ensure the effective use of El Capitan upon acceptance and already constitute significant tri-lab resources. Once El Capitan 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 FY23:

- Provided technical coordination and contractual management for the CORAL-2 NRE and El Capitan contracts. (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 significant portions of the El Capitan system including all system infrastructure racks, the capacity storage tier (i.e., the persistent Lustre file system) and compute cabinets with all associated hardware other than the compute blades. (Appendix I, target CSSE-1.b)

- Deployed several A0 compute nodes (i.e., early silicon versions of the El Capitan compute nodes) in El Capitan early access systems to enable further tri-lab system software preparations for El Capitan. (Appendix I, target CSSE-1.b)

Planned Activities in FY24:

- 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)
- Maintain multiple El Capitan early access systems to help tri-lab teams prepare for El Capitan. (Appendix I, target CSSE-1.b)
- Complete siting of El Capitan. (Appendix I, target CSSE-1.b)
- Complete El Capitan acceptance process and enable early user access to El Capitan system. (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 FY24, this project will continue to focus on the project management, acquisition and deployment oversight of the ASC Crossroads system, ATS-5, 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 FY23:

- 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)
- Completed Tri-lab Procurement Strategy document with HQ and apply to ATS-5 procurement activities (Appendix I, target CSSE-1)

Planned Activities in FY24:

- 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)

- Coordinate the activities for the removal of the Trinity system from service.

ATS-5 Tri-lab Advanced Technology System (LANL)

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 2027–2032 timeframe. The primary activity is to lead the design, acquisition, and deployment of the fifth AT system (ATS-5) in the ASC Computing Strategy. 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.

The architecture and design of ATS-5 will provide performance for large-scale applications and high-fidelity simulations 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.

ATS-5 will be designed with the following architectural advancements and major project goals:

- Overcoming the memory wall - continued memory bandwidth performance improvements for tri-lab applications
- Improved efficiency - programmer productivity, energy usage, and increased processor utilization
- Architectural diversity - ensuring that the high-performance computing ecosystem remains vibrant with multiple advanced technology solutions.
- Time-to-solution - advancing strong scaling improvements to tackle the most pressing challenge of major improvements in time-to-solution for NNSA's largest and most complex stockpile simulations.

ATS-5 will replace the Crossroads system sited at LANL but will be used by applications users from the NNSA labs as a tri-lab resource.

Planned Activities in FY24:

- Draft and release the technical specification for the ATS-5 procurement to solicit feedback from the vendor and DOE community on the design and acquisition strategy for ATS-5. (Appendix I, CSSE-1.e)
- Develop and release the benchmarks for the ATS-5 procurement to the vendor community in preparation for release of the RFP. (Appendix I, CSSE-1.e)
- Coordinate with the six major HPC laboratories in the DOE on procurement strategies and language to facilitate a vendor participation in the ATS-5 procurement. (Appendix I, CSSE-1.e)
- Release the ATS-5 RFP and conduct a tri-lab evaluation of responses for eventual contract award. (Appendix I, CSSE-1.e)

- Develop the ATS-5 Acquisition and Management Plan

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 FY22–FY24 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 FY23:

- Installed, accepted, and deployed the ancillary systems for the Crossroads project: Rocinante (unclassified applications development system) and Tycho (interim ATCC system) in the secure. (Appendix I, CSSE-1.a) (LANL)
- Installed the Application Regression Testbed (ART) named Tachi at Sandia. (Appendix I, CSSE-1.a) (SNL)
- Continued application preparations for the Crossroads systems through the COE. (Appendix I, target CSSE-1.a) (LANL)
- Completed the Crossroads Integration Readiness milestone. (Appendix I, target CSSE-1.a, FOUS-1) (LANL/SNL)
- Supported acceptance and benchmarking preparation for full-system installation in FY23 (Appendix I, target CSSE-1.a) (SNL)
- Completed testing and benchmarking for initial acceptance milestone for first Crossroads partition. (Appendix I, target CSSE-1.a) (SNL)
- Completed testing and benchmarking of full Crossroads system in support of the Crossroads Production Readiness milestone. (Appendix I, target CSSE-1.a) (SNL)

Planned Activities in FY24:

- Complete the Crossroads Production Readiness milestone. (Appendix I, target CSSE-1.a, FOUS-1) (LANL/SNL)
- Conduct quarterly reviews report milestone progress towards contractual deliverables. (Appendix I, target CSSE-1.a) (LANL)
- Complete acceptance testing and benchmark performance analysis for Crossroads system. (Appendix I, target CSSE-1.a) (LANL/SNL)
- Support transition of Crossroads platform to full production use on classified network. (Appendix I, target CSSE-1.a) (SNL)
- Support ATCC code deployment on Crossroads platform. (Appendix I, target 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-5 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 R&D projects.

Accomplishments in FY23:

- Coordinated program activities with P38 and international partners including CEA (France) and MEXT/RIKEN (Japan). (Appendix I, target CSSE-3)
- Supported benchmark development for ATS-5 system procurement planning. (Appendix I, target CSSE-1.e)

Planned Activities in FY24:

- Coordinate procurement engagements associated with Vanguard-II, NERSC-10, OLCF-6 and ATS-5. (Appendix I, target CSSE-1)
- Coordinate engagement with Department of Commerce and other agencies related to microelectronics and computing national priorities. (Appendix I, target CSSE-1)

Advanced Architecture Test Bed Research and Development (SNL)

This project will address a critical need for a range of experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy

applications, enable application performance analysis with Mantevo proxy applications, support the Heterogeneous Computing and Programming Model R&D, the Software and Tools for Scalability and Performance projects, and for 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 larger scale. Discussions will continue with Intel, AMD, IBM, NVIDIA, Arm, HPE, 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 FY23:

- Deployment of early Intel SPR hardware including the Tachi early access system in support of Crossroads and CTS-2 activities. (Appendix I, target CSSE-1)
- Deployment of GPU-based test systems in support of El Capitan and future GPU-accelerated systems. (Appendix I, target CSSE-1)
- Deployment of FPGAs and data-flow accelerators to further exploration of reconfigurable computing for the program. (Appendix I, target CSSE-1)

Planned Activities in FY24:

- Participate in early prototyping and evaluation of next-generation systems (ATS5/ATS6 and CTS2+/3). (Appendix I, targets CSSE-1 and CSSE-2)
- Evaluate and prioritize procurement of pre-release hardware from component/system vendors based on NNSA program goals. (Appendix I, target CSSE-1.d)
- Procure system hardware and software supporting Compute-as-a-Service capability for the accelerated digital engineering initiative. (Appendix I, target CSSE-3)

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 and industry partners. This project covers system software components needed to augment Linux and required proprietary operating systems 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) and Flux) to queue and schedule code runs across LLNL systems. LLNL uses TOSS on all of its Linux clusters. This project also funds approximately 60 percent of the manpower required to develop, deploy, and maintain TOSS. The funding LLNL receives for its portion of FOUS TOSS funding accounts for 40 percent of the effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in the FOUS section of this document.

Accomplishments in FY23:

- Continued support of TOSS 3 with quarterly security and bug fix updates. Released TOSS 4.5 and TOSS 4.6 and provided monthly updates. (Appendix I, Target CSSE-2.d)
- Supported the upgrade of LC's CTS-1 and CTS-1+ clusters to TOSS 4 and the integration of new CTS-2 clusters. (Appendix I, Target CSSE-2.d)
- Continued TOSS improvements in support of CORAL2 systems. (Appendix I, Target CSSE-2.b)

Planned Activities in FY24:

- Support the integration of TOSS 4 on El Capitan. (Appendix I, Target CSSE-2.b)
- Continue support of TOSS 4 for CTS-1, CTS-1+, CTS-2, and CORAL2 systems, including minor releases to support new versions of RHEL. (Appendix I, Targets CSSE-2.b and CSSE-2.d)
- Begin development of TOSS 5 based on RHEL 9.x for eventual deployment to CTS and CORAL2 systems. (Appendix I, Targets CSSE-2.b and CSSE-2.d)

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 FY23:

- As part of Milestone 8551, improved the applications development environment for El Capitan using early access systems RZvernal, Tioga, and Tenaya. (Appendix I, targets CSSE-1.b and CSSE-2.b)
- Continued to improve, test, evaluate, and provide user support for LLNL ASC production platforms, including TOSS 4 for CTS and the Sierra system applications development environment. (Appendix I, targets IC-1, CSSE-2.d)
- The Scalable Applications Preparation (SAP) team assisted tri-lab IC teams and users in porting, tuning, and using the Advanced Technology Sierra system in two tri-Lab Computing Campaigns (ATCC-14, ATCC-15). (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 FY24:

- Continue development of the applications development environment for El Capitan using early access systems and the delivered El Capitan system. (Appendix I, targets CSSE-1.b and CSSE-2.b)
- Maintain and provide user support for LLNL ASC production platforms including the Sierra system applications development environment. (Appendix I, target IC-1, CSSE-2.d).

- The SAP team will assist tri-lab IC teams and users in porting to RedHat 8 based TOSS 4 and using the Sierra system in computing campaigns. (Appendix I, target IC-1)
- Continue to improve tools and programming infrastructure through support of NRE activities and working groups for the CORAL-2 Exascale system. (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 large-scale applications. It includes a strong effort in applied data analytics and machine learning.

This effort is closely tied to the Enabling Manufacturing (EM) effort which works to apply various ASC-developed tools and technologies to weapons program manufacturing efforts around LANL. Data sources, including images, text (e.g., CSV or reports), and raw data from instruments, are growing in size and complexity. ASC possesses a unique expertise and ability to assist in studying these datasets. This exercise ranges from applying or developing visualization tools, performing statistical analysis, or applying machine learning techniques in close discussion with data owner stakeholders to understand what the knowledge discovery goals are. The production work is in collaboration with PEM.

Accomplishments in FY23:

- Completed analysis of proton radiography (prad) image data using various machine learning and image processing techniques. Report written and delivered.
- Numerous CMM Analytics Toolkit (CACTI) improvements including:
- Integrated with Data Science Infrastructure (DSI)
- Developed cross-dataset analytics for trend analysis with classic (e.g., interpretable) machine learning techniques.
- Designed new scan methodologies and analyzed mathematically for optimizing on various techniques (e.g., path length, number of stylus lifts, spacing between points, number of points scanned, etc.). Working with stakeholders to evaluate choices and prototype implement in hardware.

Planned Activities in FY24:

- Utilize multiple manufacturing data sources to connect back to the design for the purpose of evaluating as built parts. For example, CMM + radiography + CT + simulation codes.
- Evaluation of various vendor-tools for manufacturing workflow including Digital Twins (DTs) via OmniVerse, and TAO for scalable image anomaly detection.

- Using a combination of pRad (experimental) and synthetic image (ASC simulation tools), learn accurate inference of damage from radiographs. The approach will use a neural network technique which the team has prototyped in FY23 and will give increased efficacy of radiographs in predictive simulation capabilities.
- Conduct analysis of 3D Pu-based microstructures in the presence of impurities using the Tusas framework. Preliminary testing has been done on Frontier and these simulations will be on Sierra. Technique involves use of GPUs and will fold in machine learning advances in FY24.
- Model performance improvements of workloads using advanced memory technologies using existing hardware and simulation capabilities. Codesign prototype compiler and software infrastructure to support these technologies.
- Measure impact of branching in a variety of workloads on commodity technologies (CPUs, GPUs) with a particular focus on Monte Carlo, irregular data structure manipulation, and AI (graph neural network algorithms). Develop synthetic benchmarks that represent these workloads in a compact form and develop a plan for designing technologies to address bottlenecks in these workloads.

Advanced System Test Beds (LANL)

The Advanced System Test Beds project provides test bed hardware and software for research investigations in support of the ASC IC/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. The testbed is a key resource for ASC IC DevOps and Performance Engineering cross-cutting projects providing an environment to develop and test both CI workflows and performance experiments on leading-edge hardware. Additionally, the testbed is a resource for ASC HPC projects exploring areas such as provisioning, workflow orchestration, and system monitoring.

Accomplishments in FY23:

- Continued to support advanced Machine Learning (ML) hardware from SambaNova and Grok
- Began planning to move ASC testbeds from their existing location in the Advanced Computing Laboratory (ACL) to the Strategic Computing Complex (SCC)
- Supported the IC DevOps project by deploying test and development instances of GitLab, the Jacamar CI runner, and the Quay container registry.
- Deployed early access Intel systems and Intel compiler suites in support of the Crossroads ATS-3 system (Appendix I, CSSE-1.a)
- Deployed early access NVIDIA Grace/Hopper nodes and the NVIDIA programming environment.

Planned Activities in FY24:

- Deliver classified ASC testbed in the Strategic Computing Complex (SCC) environment with initial focus on Large Language Model (LLM) support utilizing emerging Machine Learning (ML) hardware technology (e.g., SambaNova DataScale).
- Develop plan for migrating unclassified ASC testbed out of the Darwin cluster into the SCC environment.
- Deliver ongoing unclassified ASC testbed support hosted in the Darwin cluster as driven by program requirements until efforts can be migrated to the SCC.
- Evaluate new hardware technologies (including AI hardware), algorithms, and codes using an agile / rapid response approach.

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 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 operating systems (Kitten), lightweight runtime system (QThreads), and high-performance network stack (Portals communication protocol) development with a production HPC computing stack.

Accomplishments in FY23:

- Produced a detailed study examining the ability to achieve reproducible results from pseudo-random number generation algorithms used to seed machine learning algorithms. (Appendix I, target CSSE-2)
- Completed the development and analysis of a tool for maintaining a target minimum power usage using on a high-performance computing system while minimizing interference with running applications. (Appendix I, target CSSE-2)

- Participated in OpenMP architecture review board continuing in incorporate capabilities which sure current and future DOE systems. (Appendix I, Target CSSE-2)

Planned Activities in FY24:

- 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)
- Enhance container and container orchestration technologies to improve performance, scalability, and usability for ASC platforms and workloads. Continue to partner with vendors, facilities, and application/library developers to improve container-based infrastructure. (Appendix I, target CSSE-2)
- Update and enhance the QThreads runtime library for more robust and efficient programming model support of new and emerging hardware capabilities. (Appendix I, target CSSE-2)
- Actively participate in the OpenMP Architecture Review Board to shape the direction of the OpenMP parallel programming model to provide needed capabilities for ASC applications and workloads. (Appendix I, target CSSE-2)

High Performance Computing Hardware Architecture Simulation (SNL)

Structural Simulation Toolkit (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 FY23:

- Released SST 13.0 with capability improvements to support the program's advanced architectures efforts and collaborations. (Appendix I, target CSSE-1)
- Initial rollout of user support plan and user support group to efficiently support a larger set of internal and external users. (Appendix I, target CSSE-1)

- Deployed new infrastructure for advanced RISC-V and accelerated system models in support of Project 38, Vanguard, AGILE, and other NNSA-supported projects. (Appendix I, target CSSE-1)
- Infrastructure for fully automatic generation of computation models for Kokkos skeleton applications, in support of full-system modeling. (Appendix I, target CSSE-1)

Planned Activities in FY24:

- Release SST 14.0 with capability to support the program's advanced architectures, testbeds and prototype activities. (Appendix I, target CSSE-1)
- Significantly improve scalability and performance of the SST core to allow increased scale of high-fidelity simulations. (Appendix I, target CSSE-1)
- Host user support group meetings and continue to improve infrastructure for users of SST. (Appendix I, target CSSE-1)

Interprocess Communication System Software Stack (SNL)

The Interprocess Communication System Software Stack project will develop capabilities to enable performance and scalability of ASC applications on current and future high-performance interconnection networks on extreme-scale platforms. This project will concentrate on characterizing application requirements with respect to functionality and performance for intra-application data movement as well as application network transfers to external I/O services. It will also provide a low-level network programming interface appropriate for current-generation network hardware as well as more advanced next-generation hardware with more sophisticated network interface capabilities and functionality. As applications explore alternative programming models beyond the current distributed memory MPI model, the low-level network programming interface must evolve to include the ability to provide very lightweight one-sided data transfer operations, while continuing to enable efficient two-sided message-based transfers.

Accomplishments in FY23:

- Developed the Configurable Messaging Benchmark (CMB) that measures the potential impact of different fine-grained communication mechanisms on application performance. (Appendix I, target CSSE-2)
- Studied the use of machine learning techniques running on a SmartNIC to predict network traffic to explore the use of offloading system services with minimal application perturbation. (Appendix I, target CSSE-2)
- Completed a detailed analysis of network hardware-level design optimization for an implementation of the MPI Partitioned standard using the InfiniBand Verbs API. (Appendix I, target CSSE-2)
- Published a study that presented a novel analytic model of fine-grained communication, explored the implications of that model for communication

performance, and compared and contrasted model behavior with results from a series of fine-grained communication benchmarks. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- 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)
- Actively participate in the MPI Forum to shape the direction of the MPI standard and MPI implementations to provide capabilities needed for current and future ASC applications and workloads. (Appendix I, target CSSE-2)
- Engage with vendor partners to explore next-generation SmartNIC designs with a focus on enhancing system software offload possibilities, including machine learning and data movement services. The long-term goal of this work is to enable ASC applications to take full advantage of SmartNICs to enhance overall performance without requiring application-level modifications. (Appendix I, target CSSE-2)
- Build on the foundations provided by the MiniMod framework and the Configurable Messaging Benchmark to gain insights into the benefits and drawbacks of various communication methodologies (remote memory access, two-sided, fine-grained), especially when used with accelerator (e.g., GPU) technologies. (Appendix I, target CSSE-2)

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 FY23:

- Demonstrated the effectiveness of checksum-based resilience techniques adapted from HPC applications for correcting bit flips in embedded algorithms in ND weapon systems. (Appendix I, target CSSE-2)
- Integrated Kokkos-Resilience with NimbleSM. (Appendix I, target CSSE-2)
- Studied I/O characteristics of AMT-based checkpointing w.r.t. task over-decomposition, load balancing, and checkpoint efficiency. (Appendix I, target CSSE-2)
- Completed modifications to Kokkos-Resilience to work with Kokkos 4.0. improving our code stability by adding tests and automated CI/CD. (Appendix I, target CSSE-2)

- Development, evaluation, documentation of common neuromorphic computing failures and incorporating into failure data repository. (Appendix I, target CSSE-2)
- Completed analysis of GPU HBM (High Bandwidth Memory). (Appendix I, target CSSE-2)
- Incorporated neuromorphic failures into the current failure data repository at Sandia. (Appendix I, target CSSE-2)
- Completed several modifications to the NimbleSM distBVH contact library to enable resilience in conjunction with the LFLR project thus enabling recovery during the contact search phase. (Appendix I, target CSSE-2)
- Performed NimbleSM experiments that analyzed performance of the contact search algorithm including scaling experiments, Kokkos with OpenMP acceleration and considering over-decomposition. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Demonstrate scalable local recovery for handling additional error scenarios and more realistic HPC and embedded application types. (Appendix I, target CSSE-2)
- Conduct an initial exploration of evaluating resilience techniques using SST and evaluating co-design tradeoffs. (Appendix I, target CSSE-2)
- Prototype resilient process recovery using DARMA/VT on top of FENIX/ULFM. (Appendix I, target CSSE-2)
- Explore process recovery and VT runtime restoration on potentially migrated data on a representative test application. (Appendix I, target CSSE-2)
- Characterize reliability for available neuromorphic computing hardware. (Appendix I, target CSSE-2)
- Evaluate the impacts of realistic failure injection on a neuromorphic computing platform using several key workloads and evaluating mitigation overheads. (Appendix I, target CSSE-2)
- Evaluate neuromorphic workload performance from a reliability perspective and identify resilience areas of interest. (Appendix I, target CSSE-2)
- Initial GPU support for the NimbleSM contact capability to align with SIERRA integration. (Appendix I, target CSSE-2)
- Develop more sophisticated contact enforcement methods that better reflect the current state-of-the-art in SIERRA and the workload balance between search and enforcement. (Appendix I, target CSSE-2)

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 FY23:

- Improved VT efficiency (reducing overheads) for small tasks to improve the EMPIRE's performance in the case of few particles in a mesh block. (Appendix I, target CSSE-2)
- Adopted C++17 making DARMA easier to use for end-users (including sending data, reductions, etc.). (Appendix I, target CSSE-2)
- Implemented a centralized, memory-aware load balancer able to redistribute GEMMA workloads under a tight memory constraint. (Appendix I, target CSSE-2)
- Created an initial implementation of tasking in GEMMA with the necessary code to move matrix sub-blocks under a tight memory constraint. (Appendix I, target CSSE-2)
- Developed a new modeling approach for GEMMA workloads suitable for predicting the amount of work across the range of problems. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Full asynchronous execution of EMPIRE on GPUs, overlapped with CPU work. (Appendix I, target CSSE-2)
- Implementation of tasking in GEMMA with online load balancing using the new memory-aware load balancer. (Appendix I, target CSSE-2)
- Implementation of new task visualizer in C++ that can be run online in VT to visualize workloads for applications. (Appendix I, target CSSE-2)

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 project 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: ISO C++ language standard 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 FY23:

- Released Kokkos 4.0 with support for the latest ROCm and CUDA environment. (Appendix I, target CSSE-2)
- Collaborated with vendors (e.g., Intel, AMD, and NVIDIA) to support application needs. (Appendix I, target CSSE-2)
- Successfully completed Kokkos and Kokkos Kernels deliverables for ATDM KPP-3. (Appendix I, target CSSE-2)
- Supported successful SPARC KPP-2 milestone with better performance of block Jacobi and block Tridiagonal solvers on the AMD architectures compared to V100. (Appendix I, target CSSE-2)
- Developed a new SpGEMM interface in order to support both AMD and NVIDIA interfaces and provide greater flexibility to Trilinos. (AppendixI, target CSSE-2)
- Delivered several presentations at SIAM CSE23: “The Future of Linear Algebra in the C++ Standard”; “Batched Sparse Linear Solvers in Kokkos Kernels”; “Trilinos: Portable and Scalable Linear Algebra using MPI+Kokkos”.
- Developed a new C++ class “mdspan” that was accepted into the C++23 draft standard and is considered one of the key new features in C++23.

Planned Activities in FY24:

- Provide Support via Email, Slack and Github to ASC customers. (Appendix I, target CSSE-2)
- Provide Kokkos Training upon request to ASC customers, if the available online recorded tutorials are insufficient. (Appendix I, target CSSE-2)
- Collaborate with maintainers of C++ standard libraries to integrate the C++23 mdspan feature into the production libraries. Work on linear algebra support for the ISO C++26 standard. (Appendix I, target CSSE-2)

- Develop backend to support data-flow accelerators in support of CSSE L2 milestone. (Appendix I, target CSSE-2)
- Update Kokkos as needed to support new vendor software stacks on ASC relevant platforms including El-Capitan, Crossroads, and NVIDIA Grace Hopper systems (Appendix I, target CSSE-2)
- Add explicit support for unified memory GPU platforms such as NVIDIA Grace-Hopper and AMD MI300 based systems. Evaluate usability and performance of approach in applications. (Appendix I, target CSSE-2)
- Develop new kernels/solver algorithms and optimize existing ones in support of ASC applications on El Capitan and Crossroads. Specific focus on techniques to improve strong scaling issues on GPU platforms such as El Capitan. (Appendix I, target CSSE-2)
- Develop GPU stream-based kernels (including linear algebra building blocks) in the Kokkos Kernels library (Appendix I, target CSSE-2)
- Develop batched sparse methods to support the needs of ASC IC customers on latest NVIDIA and AMD platforms. Evaluate using new batched preconditioners as part of this effort. (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 FY23:

- Developed and released HPSS 10.2; began development of HPSS 10.3. (Appendix I, target CSSE-2.d)
- Explored using multithreaded POSIX-capable transfers to a distributed archive filesystem. (Appendix I, target CSSE-2.d)
- Moved all disk-resident data to newly deployed classified Adaptable Storage Platform HPSS disk cache. (Appendix I, target FOUS-5)

- Designed and procured Adaptable Storage Platform configuration for HPSS Core Server and Metadata Disk deployment to the unclassified and classified environments. (Appendix I, target CSSE-2.d and FOUS-5)
- Deployed HPSS 9.3. (Appendix I, target CSSE-2.d)
- Began investigation of HPSS RAIT technology. (Appendix I, target CSSE-2.d)
- Continued support for archival storage systems. (Appendix I, target FOUS-5)

Planned Activities in FY24:

- Develop/release HPSS 10.3 and 11.1. (Appendix I, target CSSE-2.d)
- Prototype new tool to automate archiving user data. Explore integration of new tool to automate retrieving user data with Flux. (Appendix I, target CSSE-2.d)
- Explore/test HPSS features which make tape reads more performant (e.g. buffered tape marks, fast positioning, etc.). (Appendix I, target CSSE-2.d)
- Deploy Adaptable Storage Platform configuration for HPSS Core Server and Metadata Disk deployment to the unclassified and classified environments. (Appendix I, target CSSE-2.d and FOUS-5)
- Investigate deployment of HPSS 10.3 to unclassified environment. (Appendix I, target CSSE-2.d)
- Further investigations of HPSS RAIT technology. (Appendix I, target CSSE-2.d)
- Continue 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 FY23:

- Deployed and migrated data to new RZ file system. (Appendix I, Target FOUS-5)
- Tested and hardened Lustre 2.15 for center-wide deployment readiness. (Appendix I, Target CSSE-2.d)

- Prepared for Merced, El Capitan's Lustre filesystem, readiness. Integrated HPE's slingshot network into Lustre stack, tested Lustre 2.15, TOSS4 integrated onto HPE storage hardware, disk management and monitoring, and performance testing and tuning. (Appendix I, Target CSSE 2.b)
- Integrated Merced, El Capitan's Lustre filesystem. (Appendix I, Target CSSE 2.b)
- Continued to make general performance and management improvements to ZFS-based Lustre. (Appendix I, target CSSE-2.d)

Planned Activities in FY24:

- Integrate, migrate data to, and deploy, new SRD file system. (Appendix I, Target FOUS-5)
- Deploy Lustre 2.15 in the OCF environment to provide higher streaming I/O and metadata performance. (Append I, Target CSSE-2.d)
- Evaluate use of sharded directories in Lustre 2.15 to improve metadata distribution among servers. (Appendix I, target CSSE-2.d)
- Develop software to use Lustre job statistics in the SONAR and Elasticsearch data analytics systems to manage Lustre. (Appendix I, target CSSE-2.d)
- Integrate Merced, the Lustre filesystem for the El Capitan exascale system (Append I, Target CSSE-2.b)
- Work with HPE to complete El Capitan exascale system's near-node rabbit solution. (Append I, Target CSSE-2.b)
- Configuration and test of self-encrypting hard drives for Lustre filesystem. (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 FY23:

- Integrated Slingshot support into TOSS. (Appendix I, Target CSSE-2.b)

- Tested long haul InfiniBand between buildings for Lustre SAN. (Appendix I, Target CSSE-2.c)

Planned Activities in FY24:

- Test and evaluate Mellanox NDR network technologies.
- Test and evaluate NVIDIA H100 GPUs with ARM and X86_64 host processors.
- Test and evaluate Intel Ponte Vecchio GPU technologies.

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 FY23:

- Managed NG-HPCN contract and FY23 milestone deliverables.
- Organized several tri-lab technical deep dive meetings on various co-design topics.

Planned Activities in FY24:

- Manage NG-HPCN contract and FY24 milestone deliverables.
- Continue strong tri-lab co-design engagement with the selected NG-HPCN industry partner.

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.

The Data Science Infrastructure (DSI) effort, a meta data collection, storage and search project, is part of the FSAF sub-element. The DSI project focuses on gathering meta data from existing scientific activities, storing this meta data and (pointers to) associated data in a database, as well as supporting the querying and access of this meta data and its associated data. The DSI project supports meta data reader and writer abstractions. These abstractions support the reuse of scientific data without requiring DSI users to have to learn specific application meta data formats or storage system details. Project goals include integrating the DSI framework into ASC simulations, machine learning activities, and testing workflows.

Accomplishments in FY23:

- Released first open-source version of DSI software.
- Completed the deployment of DSI software and 3PB of DSI-enabled storage in the secure.
- Completed integration of singularity-eos into the EAP codes, with similar work underway for the SAP application. An initial port of SAP and its Spack environment to the RZvernal early access system, including GPU offload support, was completed in time for the Spring 2023 El Capitan hackathon.
- Provided SME assistance with VPIC – part of the acceptance test suite, helped with vendor MPI issues, and COE and early user access support.
- Completed SAP performance report for the performance engineering project. The Bueno framework utilized by the project continues to be maintained/enhanced as needed.
- Continued development of software projects supported by AR team members including Quo-Vadis, Open MPI and OpenPMIx.

Planned Activities in FY24:

- Development of an SqlAlchemy backend to support the representation of structured and unstructured metadata in DSI.

- Continued DSI development for machine learning, ensembles, and performance workflows.
- ASC code team support will focus on El Capitan and Crossroads, dependent on the needs of the ASC code teams.
- A release of Quo-Vadis and Spack support for building/deploying Quo-Vadis.
- At least one bug-fix release of Open MPI 5.0.x and related OpenPMIx releases.

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 FY23:

- Integrated FrETT into Sandia Analysis Workbench (SAW) providing cold storage for job workflow results. (Appendix I, target CSSE-2)
- Achieved promising results in testing container registry on HPSS/FUSE. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Contribute to the HPSS version 11 release series, to include CR302 related enhancements: update Unix Authentication protocol. (Appendix I, target CSSE-2)
- Add Kerberos authentication to HPSS S3 component to provide secure cloud service from HPSS cold storage. (Appendix I, target CSSE-2)

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 FY23:

- Partnered with several ADE (Accelerated Digital Engineering) projects to develop object storage capabilities. (Appendix I, target CSSE-2)
- Demonstrated conversion of IOSS Exodus data from a PLATO topology optimization tool into data appropriate for the DataSEA data-management tool. Also demonstrated movement of IOSS data directly to S3 object storage. (Appendix I, target CSSE-2)
- Deployed RAM-based BeeOND (BGFS on-demand) user-level file system on Sandia's CTS systems, enabling the configuration of large-scale shared-cache arrangements like those encountered in ML training. (Appendix I, target CSSE-3)
- Completed successful functionality tests for the DDN RED cluster. The first beta release is imminent and will include S3 functionality that should prove useful for some cloud-oriented workloads. (Appendix I, target CSSE-2)
- Continued improvements to IOSS and ParaView support for high-order Discontinuous Galerkin data formats. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Perform R&D to improve utilization of application-level data services focusing on evaluating effectiveness of on-premises cloud data storage for the purposes of simulation data management. (Appendix I, target CSSE-2)
- Perform R&D to develop enabling capabilities for data-management in support of Digital Engineering and Advanced Machine Learning workflows. (Appendix I, target CSSE-2)
- 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-2)

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 peta- and exascale 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 develops and supports hardware, software tools, services and facilities for managing, visualizing, analyzing, and presenting scientific data. The visualization hardware architecture team engages in planning, test bed prototyping, testing of systems and components, and procurement and integration of new systems. The project exploits the latest capabilities of clustering hardware, GPUs, and parallel storage systems. Hardware capabilities include three production visualization servers (Pascal, Tron, and Vertex) and VNC (Virtual Network Computing) servers in three security zones. These VNC servers support visualization and remote use of interactive tools. The project installs, maintains, and consults on software visualization tools, including resource management tools, movie players, animation, and visualization packages. Users of Livermore Computing systems have the opportunity to partner with visualization experts from the project, who can augment, design, or create visualizations (including animations, movies, and images) of the users' data. The project provides operational support for high-resolution display devices and facilitates demonstrations and presentations on these displays.

Accomplishments in FY23:

- Maintained data analysis and visualization environment across LC platforms while transitioning hardware, vis software, and users to TOSS4. (Appendix I, Target CSSE-2.d)

- Maintained user access to VNC servers & desktops, allowing users to easily connect to and launch graphics on our compute systems, while migrating VNC servers to new hardware and to TOSS4. (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 and other LC users with visualization and data analysis activities, including consulting and creation of visuals and movies for presenting and analyzing scientific data.

Planned Activities in FY24:

- Maintain the data analysis and visualization environment across LC platforms and networks. (Appendix I, Target CSSE-2.d)
- Continue support for VNC servers enabling remote connections & an environment pre-configured for graphics. (Appendix I, Target CSSE-2.d)
- Provide operational support for high-resolution displays and provide demo, tour, and special event support. (Appendix I, Target CSSE-2.d)
- Support LC users with data analysis and visualization activities via software consulting and the creation of images and movies. (Appendix I, Target CSSE-2.d)

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 FY23:

- In Hopper, made significant progress towards splitting the application into components, factoring out proprietary and core logic into separate, standalone repos. The end goal being an open-source version of Hopper that can be used freely throughout the computing world, which we can also extend with LLNL-proprietary elements for use on site. Also investigated possible support for object stores in Hopper, incorporated support for HPSS 9.3, developed multiple CI capabilities for streamlining the building and scanning processes, and made extensive usability and

performance improvements as dictated by hardware and OS changes within the Lab. (Appendix I, Target CSSE-2.d)

- In Lorenz, overhauled the "change quota" feature for Hotline consultants to use a more streamlined and robust architecture, introduced an improved user experience for the Task Manager Interface, incorporated new tasks for simplifying Hotline activities, and made significant improvements to the Staff Away and System Support Matrix utilities. Continued to enhance the MyLC Dashboard with user and staff requested features. (Appendix I, Target CSSE-2.d)

Planned Activities in FY24:

- For Hopper, install the open-source portions of Hopper into GitHub and develop the necessary collaboration, tracking, and communication ecosystem. Develop CI pipelines both in GitHub and within LC for mirroring open-source components and building the various Hopper instances. Maintain the application across hardware and OS enhancements and continue to improve usability and performance. (Appendix I, Target CSSE-2.d)
- For Lorenz, expand job management capabilities in Lorenz to support the Flux resource manager. Continue to enhance the MyLC Dashboard with user and staff requested features. (Appendix I, Target CSSE-2.d)

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 for application simulation and development efforts, and ongoing expert support of ASC end-users for deployed technologies.

The application visualization 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 application visualization project focuses on integration into the code projects and ensures the quality of the software components so they can be well-integrated into LANL's HPC infrastructure. This is also the path for newly developed visualization and analysis research to make its way into the production HPC 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 FY23:

- Supported simulation developers and users with data visualization and analysis software, facilities, support, training, and feature extensions by maintaining and improving visualization software, responding to support and training requests with qualified personnel, and studying and improving gaps in capability and functionality.
- Developed and improved Ristra's FleCSI in situ support and worked with tri-lab community to improve common interfaces for the next generation in situ infrastructure while maintaining and improving capability and support for existing data analysis capability in LANL's EAP, LAP, and SAP codes.

Planned Activities in FY24:

- Support simulation developers and users with data visualization and analysis software, facilities, support, development, tutorials, and maintenance with an emphasis on designing new in situ and post processing workflow processes to effectively leverage the Crossroads design.
- Collaborate with tri-lab to develop and productionize catalyst 2.0, ParaView's next generation standardized in situ interface for ASC simulation codes.

Cross-Cutting Extreme-Scale Research (LANL)

LANL Cross-Cutting Extreme-Scale Research project focuses on research in data extraction, visualization, analytics and end-to-end workflows, with an emphasis on support for ensembles and support for LANL's Data Science Infrastructure (DSI) for use in extreme-scale environments. The research is coordinated with other CSSE projects to leverage and extend cross-cutting expertise. Project areas include in-situ data analytics for large-scale simulations, developing reusable workflows including in situ data reduction approaches and post-processing visual analysis of reduced-size data extracts via Cinema, and developing high-performance, portable, Visualization Toolkit (VTK) algorithms for use with ParaView, Cinema, and CMF workflows.

Accomplishments in FY23:

- Data Capture and Extraction: Release of generalized Synthetic Radiograph ParaView in-situ workflow in CMF in collaboration with the ICF team. This includes multi-material lookup and can run on any simulation results that are in VTK format. This is an expansion of capability from FY22 in-situ workflows.
- ECP Delivery: Delivered KPP-3's in support of ECP project finalization, including Cinema workflows, in-situ data extraction support, and UIs for Data Science Infrastructure (DSI) project.

- Data Analysis and Ensemble Support: Release of initial Cinema viewer toolkit, built on the Cinema Engine. The viewer toolkit supports development of applications for comparative visualization of simulations, investigation of parameter sets for AI/ML, and UIs in support of DSI project.

Planned Activities in FY24:

- Deliver Cinema viewer toolkit, supporting updated Cinema:Explorer and Cinema:Viewer applications, in support of DSI customers, and for general Cinema database analysis and visualization.
- Develop AI/ML support capability for data extraction, compression, and reconstruction within the Cinema toolkit.
- Deliver Capture and Extraction workflows, including ongoing support for projects such Synthetic Radiograph workflow and Equation of State (EOS) workflows.

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 FY23:

- Released a new production version of the Slycat ensemble analysis tool, which includes new algorithms for mesh metrics. (Appendix I, target CSSE-2)
- Provided ensemble analysis to the ASC ADE DetNet project. (Appendix I, target CSSE-2)
- Extended Slycat workflow plugin to support Dakota HDF5 results file processing to remote, unsupervised, HPC workflows that are created and launched by Sandia's SAW and Next Generation Workflow (NGW). (Appendix I, target CSSE-2)
- Developed communications protocol for distributed rendering, including handling for arrays of structures of multiple dynamic arrays. (Appendix I, target CSSE-2)
- Developed and presented Catalyst training at SPARC Stakeholder's meeting. (Appendix I, target CSSE-2)

- Addressed usability and adoption barriers in python user interface to Catalyst. (Appendix I, target CSSE-2)
- Established regular Catalyst community meetings with Kitware and LANL; DoD and CEA also engaged. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Enhance and improve existing, and explore new, ensemble analysis and visualization tools for ASC engineering applications. (Appendix I, target CSSE-2)
- Integrate analysis and visualization tools to support digital engineering workflows. (Appendix I, target CSSE-2)
- Improve Catalyst in situ analysis interface support for quantitative assessments used for stockpile assessment. (Appendix I, target CSSE-2)
- Enhance Catalyst adoption across user communities such as SPARC, Sierra Solid Mechanics, EMPIRE, and others. (Appendix I, target CSSE-2)
- Improve VTK-m performance on AMD GPUs and Intel CPUs through improvements to Kokkos. (Appendix I, target CSSE-2.e)
- Improve the rendering performance of VTK-m, focused on scalability. (Appendix I, target CSSE-2.e)
- Provide production visualization support for stockpile stewardship calculations and early design studies for future systems. (Appendix I, target CSSE-2)

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 computing 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 FY23:

- Implemented fifth-generation quantum simulation hardware testbed platform with fast feedback control and integrated control of characterization equipment for just-in-time measurements of classical hardware. (Appendix I, Target CSSE-3.d)
- Expanded characterization and control of quantum hardware and improved average performance with continued development of machine learning approaches and slow feedback. (Appendix I, Target CSSE-3.d)
- Began developing methods for co-designing emulation circuits for ASC applications. (Appendix I, Target CSSE-3.d)
- Collaborated with SNL and LANL to develop an ASC roadmap for quantum computing hardware and algorithm development. (Appendix I, Target CSSE-3.d)
- Performed a comprehensive survey of available superconducting quantum computing processors available for purchase. (Appendix I, Target CSSE-3.d)
- Worked on adapting custom gate design to generate more robust gates on the trapped ion platform using SNL's Quantum Scientific Computing Open User Testbed (QSCOUT). (Appendix I, Target CSSE-3.d)

- Conducted an extensive evaluation of credibility assessment methods for neural networks, including ensembles, autoencoders, and DeltaUQ (LLNL tool), on atomic physics data when interpolating and extrapolating. DeltaUQ shows high correlation between expected variance and model error for predictions outside of the training domain. (Appendix I, Target CSSE-3.d)
- Worked with the SambaNova team to support a portable version of the ECP ExaLearn CosmoFlow model. (Appendix I, Target CSSE-3.d)
- Worked to support automatic mixed-precision solutions for LBANN benchmarks. (Appendix I, Target CSSE-3.d)

Planned Activities in FY24:

- Implement next-generation quantum simulation hardware testbed platform by installing superconducting quantum processor. (Appendix I, Target CSSE-3.d)
- Implement a parallel-in-time integration schemes for optimal control methods. Demonstrate deployment of classically optimized multi-qubit, multi-level control pulses on the quantum testbed. (Appendix I, Target CSSE-3.d)
- Develop an interface to integrate ASC QuDIT (Quantum Device Integration Testbd) device data into the classical optimal control loop to enable open-loop gate optimization. (Appendix I, Target CSSE-3.d)
- Co-design proof-of-concept emulation circuitry to explore optimized analogue quantum simulation methods to solve problems in nuclear physics. (Appendix I, Target CSSE-3.d)
- DeltaUQ will be incorporated into a training/retraining/inference workflow to determine when the atomic physics model can be used and datapoints for retraining. It is expected that modification will be necessary to better assess the credibility of the neural network where the model performs poorly, although the query point is within the training domain. (Appendix I, Target CSSE-3.d)
- The CosmoFlow model and the Hermit (atomic physics) model will be ported to the SN-30 (SambaNova) systems for evaluation. These new systems will be available on higher security networks than our current systems, allowing for evaluation of more realistic training/retraining/inference workflows. (Appendix I, Target CSSE-3.d)
- Work with the Cerebras open-source Large Language Model to evaluate training and inference pipelines on the CS-2 and El Capitan EAS systems. Develop a retraining pipeline for these models to evaluate performance. (Appendix I, Target CSSE-3.d)

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. Our neuromorphic efforts are directed at the understanding of neuromorphic computing methodology and the implementation of machine learning (ML) methods both in

simulation and natively, on-chip. Our ML efforts are targeted at improving computational fluid dynamics simulations via the optimization of flux limiters for shock capturing codes. Efforts in quantum computing investigate a wide range of questions and problems in quantum computing from quantum algorithm development for fault-tolerant quantum computation to quantum algorithms for near-term, so-called Noisy Intermediate-Scale Quantum (NISQ) hardware. We additionally research issues in quantum machine learning (QML) and its application to quantum sensing protocols.

Accomplishments in FY23:

- 2022 (FY23) R&D 100 Award neuromorphic backpropagation algorithm on Intel Loihi neuromorphic chip.
- Publication of "Machine Learning Changes the Rules for Flux Limiters," a study demonstrating high-accuracy, machine learning flux limiters for fluid simulation.
- High-profile publications on quantum simulation, QML for quantum sensing, understanding generalization in QML, and group-invariant QML.
- Benchmarking study of NISQ devices giving clear guidelines for prioritization of algorithmic application efforts.
- Completed LANL-QCI CRADA Phase 1 Qatalyst report comparing results for QUBO-based Graph Partitioning and Community Detection algorithms using QCI and LANL approaches.
- New quantum linear system solver algorithm with order-of-magnitude reduction in quantum resources.
- Development of the ASC Quantum Computing Program Plan in collaboration with SNL, LLNL.

Planned Activities in FY24:

- Neuromorphic algorithms for implementing machine learning algorithms and testing on new Intel Loihi 2 neuromorphic.
- Assessment of machine learned flux limiters for 1D Euler's equation on xRAGE. Potential LANL ASC Level 2 Milestone.
- Subcontract with University of Colorado, Boulder for machine learned flux limiters for supersonic flow.
- Assess resource requirements for simulation of nonlinear classical system using quantum linear systems algorithm.
- Symmetry-enhanced schemes for quantum simulation of condensed matter and material sciences systems.
- QC for atomistic scales, such as opacity calculations

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 computation.

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 FY23:

- Received first stage of Loihi 2 deliverable from Intel, have successfully set up in NERL laboratory. (Appendix I, target CSSE-3)
- Featured article on neuromorphic computing [Kramer, David. "A computing hardware approach aspires to emulate the brain." Physics Today 76, no. 1 (2023): 23-26.]. (Appendix I, target CSSE-3)
- SNL neural computing featured as part of DOE Consumer Electronics Showcase (CES) booth. (Appendix I, target CSSE-3)
- Integrated GenSA (Generic Spiking Architecture) component with SST, Fugu and N2A. (Appendix I, target CSSE-3)
- Published "Logical and Physical Reversibility of Conservative Skyrmion Logic" in IEEE Magnetics Letters. (Appendix I, target CSSE-3)
- Published "Strong quantization of current-carrying electron states in delta-layer systems", D. Mamaluy, J.P. Mendez, Solid State Electronics (Appendix I, target CSSE-3)

Planned Activities in FY24:

- Develop Fugu and Neural Mini App package, with particular focus on integrating into SpiNNaker2 hardware and software stack and Loihi 2 Lava software stack. This will also include extending Fugu to incorporate a canonical learning rule. (Appendix I, target CSSE-3)

- Advance neuromorphic scientific computing algorithms, particularly Monte Carlo sampling, with a particular emphasis on taking advantage of extra hybrid conventional / non-conventional capabilities offered by Loihi 2 and SpiNNaker 2. (Appendix I, target CSSE-3)
- Develop spiking neuromorphic algorithms to solve Markov decision processes and reinforcement learning; building on neuromorphic graph algorithms and Markov Chains. Implement in Fugu as a future neuromorphic mini-app. (Appendix I, target CSSE-3)
- Explore emerging neuromorphic discrete optimization algorithms on SpiNNaker2 platform. (Appendix I, target CSSE-3)
- Design a new adiabatic CMOS test chip based on results of present testing work and benchmarking adiabatic vs. conventional CMOS analysis using GAA device data. (Appendix I, target CSSE-3)
- Develop backends capable of fully utilizing next-generation neuromorphic machines (Loihi-2 and SpiNNaker-2). These support general mathematical models, going beyond the limited Leaky-Integrate-and-Fire (LIF) dynamic of earlier machines. (Appendix I, target CSSE-3)
- Develop an efficient APAM FET concept for digital cryogenic computing with a primary focus on investigating the intrinsic gate delay, the switching speed and other device characteristics of these novel APAM devices. (Appendix I, target CSSE-3)

Interagency Quantum Computing (SNL)

This project is focused on understanding the potential impacts that quantum computers will have on computational science as they mature from noisy intermediate-scale quantum (NISQ) architectures into fault-tolerant application-scale quantum (FASQ) architectures. This involves creating and assessing (A) hybrid quantum-classical algorithms on NISQ platforms, including the supporting software infrastructure, (B) benchmarks for testbed platforms, and (C) novel quantum architectures and algorithms. This effort is synergistic with ASC/PEM activities on quantum-enhanced materials modeling, but it also aims to explore applications outside of this scope.

Accomplishments in FY23:

- Initiated capability/software development in support of realizing a hybrid quantum-classical algorithm for unbiasing continuous-time quantum Monte Carlo calculations. (Appendix I, target CSSE-3)
- Created software for synthesizing application-driven volumetric benchmark circuits and applied on IBM testbeds. (Appendix I, target CSSE-3)
- Capability advancement for the detailed resource analysis of quantum algorithms implemented on FASQ architectures based on lattice surgery. (Appendix I, target CSSE-3)

- Explored extensions of the adiabatic state preparation algorithm for accelerating/enhancing general-purpose observable estimation. (Appendix I, target CSSE-3)

Planned Activities in FY24:

- Implement numerical simulations of the hybrid quantum-classical algorithm for unbiasing continuous-time quantum Monte Carlo calculations executed on NISQ hardware. (Appendix I, target CSSE-3)
- Understand the impact of limited access to non-Clifford operations on early fault-tolerant pre-FASQ architectures and incorporate into benchmark design for testbeds with 100s to 1,000s of physical qubits. (Appendix I, target CSSE-3)
- Quantify runtime speedups associated with algorithmic/architectural advances from FY23, focusing particularly on the types of parallelism that seem achievable on early FASQ architectures. (Appendix I, target CSSE-3)
- Create and analyze a new algorithm for general-purpose observable estimation using advances from FY23. (Appendix I, target CSSE-3)

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 crosscut of enabling technologies.

DOE-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. ASC is funding the LLNL-led ADMIRRAL project which produces a multiscale, multiresolution molecular dynamics simulation framework for CORAL-class machines to study the dynamics of mutated proteins in cancerous cells. 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 FY23:

- Generalized Multiscale Machine-Learned Modeling Infrastructure (MuMMI) code, API, and data structures to make it easy to add new scales and toggle between them, including the expansion of self-healing framework to allow for online analysis, communication, and processing of more general data sets.
- Demonstrated use of ML latent space to represent biological phenomena and generate realistic hypothetical configurations.
- Multilayer perceptrons were trained to transform ML configurations into spatial coordinates (tertiary structures).
- A new sample mechanism was developed to explore the ML latent space to generate new, previously-unseen but realistic protein configurations.

- A molecular model of autoinhibited BRAF-14-3-3 complex bound to KRAS on the membrane was created is varies states for the full campaign run.
- Created ultra-coarse grained model using self-healing capability.
- A full campaign running MuMMI in 4 scales (macromodel, ultra-coarse grained, coarse grained, and all atom was run on Summit.
- Parts of MuMMI are ported to Frontier and Tioga in preparation to run on El Capitan in FY24. (Appendix I, Target CSSE-2.b)

Planned Activities in FY24:

- Run a large sized MuMMI campaign to capture RAS-RAF complex rearrangement on the membrane.
- Demonstrate the viability of using many short simulations guided by ML to develop long timescale predictions, i.e., encode the dynamics of the simulation at short scales and combine such information from large ensembles to predict the long-term behavior using long short-term memory (LSTM) neural networks.
- Port MuMMI and application code to El Capitan/Frontier type architecture. (Appendix I, Target CSSE-2.b)

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 programming model and runtime development/standardization, user-level scheduling/runtime support for emerging workflows, system-level resource and power management/scheduling, and advanced debugging/correctness tools. 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 FY23:

- Demonstrated Flux's component-by-component capabilities on top of the DataWarp/Rabbit early-access software as well as early integration with Rabbit hardware. (Appendix I, Target CSSE 2.b)
- End-to-end demonstration of CTI under Flux using CTI-based tools including STAT, TotalView, HPE valgrind4HPC etc. (Appendix I, Target CSSE 2.b)

- Demonstrated Scalable multi-node Flux/variorum-enabled power monitoring services, showing our solution fills the critical technical gaps that exist in pre-exascale power monitoring tools including IBM's Cluster System Management (CSM). (Appendix I, Targets CSSE 1.b and 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 and 2.b)

Planned Activities in FY24:

- Collaborate with vendors and standards bodies on next-gen support for programming models including OpenMP, C++, SYCL, and portability layers like RAJA (Appendix I, Target CSSE 1.b)
- Demonstrate Flux Rabbit scheduling capabilities on final DataWarp/Rabbit hardware and software. (Appendix I, Target CSSE 2.b)
- Expand Flux support for complex affinity and co-scheduling on AMD GPU platforms, especially El Capitan. (Appendix I, Target CSSE 2.b)
- Disseminate the results of our investigations via various tutorials, presentations, publications, workshops, and conferences.

Next Generation Partnerships and Outreach (LLNL)

LLNL partners with a wide range of vendors, researchers, developers, and other collaborators. This new project will support those partnerships, education, outreach, and sustainability efforts for next generation software that was developed during the Exascale Computing Project. The project leader will work with the CSSE Leader to continue the management efforts of the now retired Architectures and Software Development program element. This includes coordinating communication across the LLNL projects in the Next Generation Computing Technologies program element, as well as to the partnership and outreach work.

Planned Activities in FY24:

- Coordinate LLNL Next Generation discussions, reviews, and activities across the NGCT program element.
- Participate in FY24 software sustainability planning and management.
- Provide technical leadership for vendor partnerships outside the ATS and CTS activities. (Appendix I, Target CSSE-2.b and CSSE-2.d)
- Support Livermore Computing and Integrated Codes collaborations, planning, reviews, and publications. (Appendix I, Targets CSSE-2 and IC-1)
- Support the El Capitan Center of Excellence activities including outreach to code teams. (Appendix I, Target IC-1)
- Run ORNL contract for development and support of Jacamar.

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. ProTools works closely with the HPC tools research community to develop and harden tooling, while working with ASC application developers to identify needs and integrate tools with applications. ProTools supports tools such as Caliper, which provides application-integrated performance data collection, Hatchet and Thicket, which provide programmatic interfaces for analyzing performance data, and SPOT, which provide performance visualization. The ProTools team also works with hardware vendors to design and deploy tools for ATS platforms.

Accomplishments in FY23:

- Released Hatchet 2023.1.0 with improved query language for call trees, and enhancements to the native Caliper reader interface. (Appendix I, Targets CSSE-2.b and CSSE-2.d)
- Released initial Thicket version 2023.2.0, and published HPDC paper. Thicket is a python-based tool for analysis and visualization of ensembles of runs. (Appendix I, Targets CSSE-2.b and CSSE-2.d)
- Worked with CORAL2 collaborators on designing and deploying vendor tools for El Capitan, including TotalView, gdb4hpc, OmniTrace, OmniPerf, HPE Perf Tools, and address sanitizer. Evaluated new AMD performance profiling interfaces. (Appendix I, Target CSSE-2.b)
- Improved Caliper support for AMD GPU performance profiling. Added integration with HPE Perf Tools to Caliper. (Appendix I, Target CSSE-2.b and CSSE-2.d)

Planned Activities in FY24:

- Continue adding statistical analysis, performance modeling, and visualization techniques for multi-experiment performance data in thicket. (Appendix I, Target CSSE-2.b and CSSE-2.d)
- Add support in Hatchet and Thicket for tracking performance across nightly tests.
- Work with CORAL-2 Tools Working Group to move tools to new AMD performance profiling interfaces. Deploy updated AMD performance tools onto El Capitan. (Appendix I, Target CSSE-2.b)
- Extend Caliper to support GPU annotations. (Appendix I, Target CSSE-2.b)

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 provides long term capability development for tools in these areas, complementing efforts focused on integration of these tools into user workflows. The project enhances the Sina and Kosh data management packages and supports enhancement of workflow orchestration tools such as Merlin and Maestro, to improve their scalability and usability. The project also contributes to the development of code agnostic setup tools such as C2C (code-agnostic contours), Klee (code agnostic shaping) and supports the enhancement and productization of code-to-code geometry transfer tools.

Accomplishments in FY23:

- Enhanced the Sina visualization package with additional utilities for easily visualizing Sina data catalogues, created a Sina Fortran interface, and improved Sina documentation and tutorial materials.
- Enhanced Maestro documentation and tutorial materials.
- Maintained and enhanced C2C C++ library to conform to C++ standard and fix bugs, TOSS4 migration, and improved documentation for the C2C tools.

Planned Activities in FY24:

- Integrate the Sina C++ library (for code outputs) into the Axom project build environment, removing third party dependency and improving interoperation with other infrastructure libraries.
- Explore scalability of workflow orchestration tools and identify areas of improvement to support full scale AI and ensemble workflows on El Capitan.
- Work towards Maestro 2.0 and greater interoperability of Maestro and Merlin.
- Support and enhance code agnostic setup and geometry transfer tools.

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. DevRAMP works closely with other ATDM teams.

Accomplishments in FY23:

- Made major breakthrough in Spack's concretizer performance to enable multiple configurations of the same package to exist in the same solver environment. This work included optimizations to cycle detection and to the way conditional rules are

modeled in our solver. This work enables better and more maintainable support for Python, and better modeling of compilers as dependencies.

- Scaled Spack CI infrastructure to 100,000 jobs per week, and significantly increased reliability and robustness of Spack CI infrastructure. Added monitoring infrastructure to detect pipeline errors and reduced overall system job error rate from 14% to 2%, where the 2% are mostly real build failures. This drastically improved pipeline turnaround time for developers and package contributors.
- Developed new Spack features to support MARBL and other code teams, including environment and package requirements, further enhanced git versioning, more predictable concretization, and UI/UX improvements. (Appendix I, Target IC-1)
- Supported EAS and El Capitan environments with continued participation in the CORAL2 packaging working group and rollout of TCE environment on EAS systems. Deployed first HPE/Cray pipeline in CI. (Appendix I, Target CSSE-2.b)

Planned Activities in FY24:

- Harden pipelines and work on curating a core binary stack for Spack, so that users can deploy most of their stack from vetted binaries without building.
- Develop support for deploying MPI binaries spliced with the system MPI on bare metal. Leverage this support for other packages (GPUs, etc.).
- Continue to improve performance and UX of Spack, working closely with WSC code teams and code developers. (Appendix I, Target IC-1)
- Ensure that Spack supports El Capitan when it is deployed, as well as TCE on ATS and CTS, and code team developer environments. Develop needed features for TCE, WSC/CP, and workflow teams. (Appendix I, Targets CSSE-2.b and CSSE-2.d)

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

This project covers three key items critical for ASC's next generation computing: programming models standards, storage and I/O, and system benchmarking. The standards work is focused 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 tool's interfaces (OMPT and OMPD). In PMIx, we participate in working groups, e.g., for improving the community standard document. However, we participate and monitor developments in all parts of all three standards 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. System benchmarking is performed as part of the ATS system acceptance and involves close vendor interactions to ensure the procured ATS system meets LLNL's performance requirements. System benchmarking is also used for system procurement activities,

requesting performance Figures of Merit (FOMs) on the proposed systems from the vendors for evaluation and system selection by LLNL.

Accomplishments in FY23:

- Improved SCR for asynchronous data movement and support on El Capitan system; I/O management for complex workflows. (Appendix I, target CSSE-2)
- Performed readings of the Reinit specification with support for tools and MPI sessions for MPI in the Fault Tolerance Working Group. (Appendix I, target IC-1)
- Participated in the MPI Hybrid working group for GPU-triggered communication interface. Participated in CORAL and CORAL2 Communications working groups. (Appendix I, target IC-1)
- Performed benchmarking of the EAS3 systems as part of the EAS3 system acceptance. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Continue investigating I/O technologies, including work on user-level file systems, and I/O performance measurement. Continue vendor interactions including CORAL2 efforts. Continue SCR modification for El Capitan, ASC applications, and for supporting complex workflow data management. (Appendix I, target CSSE-2)
- Continue participation in the MPI Forum meetings to support the standardization of features critical for LLNL applications. Perform a formal reading of the global-restart fault tolerance specification in a plenary at the MPI Forum and propose its standardization for MPI 5.0. (Appendix I, target IC-1)
- Continue participation in the MPI Hybrid working group for partitioned and GPU-triggered communication and interaction with the Sessions Model. Participate in CORAL2 Communications working group. (Appendix I, target IC-1, target CSSE-2)
- Continue participation in the OpenMP Tools Group with a focus on improving debugging support for OpenMP via OMPD and a DWARF-aware interface. (Appendix I, target IC-1)
- Perform benchmarking of the El Capitan system as part of the system acceptance. Prepare the benchmarking suite for the ATS-6 RFP release. (Appendix I, target CSSE-2)

Advanced Machine Learning (LLNL)

ASC's AML initiative aims to integrate modern machine learning techniques that have been used to tackle challenging problems across domains, including Large Language Models (LLMs), into ASC efforts spanning simulation and computing. The AML initiative evaluates work from LDRDs and other research to support core NNSA applications at production scales. Hardware exploration encompasses near-term hardware

advances in neural networks and more speculative hardware, such as neuromorphic technologies.

Accomplishments in FY23:

- Developed a sampling method that uses clustering to maximize data diversity. This method has shown data reduction that leads to increased training accuracy with multiple datasets.
- Successfully applied Enzyme (from MIT) to automatically differentiate complex non-linear plasticity models. This is a key step towards facilitating the implementation of an implicit ALE method. More work is needed in Enzyme to handle GPU memory management.
- Developed a retraining workflow for high-fidelity opacity calculations within a multi-physics code.
- Further developed tools for AI-informed radiography analysis of fractures.
- Completed a Tri-lab L2 Milestone on data infrastructure requirements.

Planned Activities in FY24:

- High-fidelity ML surrogate for opacities:
    - Integrate neural network with an AI-driven design optimization workflow.
    - Integrate solutions for Equation of State predictions that are thermodynamically consistent with the opacities.
    - Demonstrate the full training/retraining/inference workflow integrated into the design optimization workflow on EAS-3 hardware
- Evaluate neural network performance for multi-scale hydrodynamics coupling across parameter variations
- Evaluate OpenLLaMA, or another open-source Large Language Model, for LLNL's Knowledge Management program.
- Investigate the usefulness of Large Language Models for HPC code development, such as porting to accelerators and programming language translation.

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 developing a workflow orchestration system that allows standard containerized HPC applications to run on any HPC infrastructure as well as any cloud infrastructure. In BEE the lower level of the HPC system is abstracted such that the science teams can manage and archive complex workflows for reusability and provenance. 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, Flux, PBSPro).

This project will be partially supported by post-ECP software-sustainability funding.

Accomplishments in FY22 Q4 and FY23:

- BEE was used to successfully orchestrate multi-step workflows on Frontier using containerized applications.
- Built VASP container and successfully ran tests on LANL testbed Darwin and Chicoma and are actively engaged with science teams using VASP for developing workflows.
- Released BEE 0.1.4 to the public on GitHub that is now installable through pip giving users the ability to manage multiple workflows using BEE on any production systems without administrative privileges or arcane developer environments. This effort included developing and implementing Continuous Integration testing.
- Investigated using Exaworks PSI/J for alternative support for backends and are currently integrating PSI/J in parallel with development and continuing improvements of our current Slurm, LSF and Flux backends. Using PSI/J may give a quicker path to incorporation of any new job scheduling systems in the future.
- Successfully accomplished the required ECP KPP-3 activities.

Planned Activities in FY24:

- Ensure Resilience of all BEE components to include recovery of monitoring and continuation of workflows for any system or front-end disruptions.
- Integrate the use of BEE with the V&V Common Model Framework (CMF) for production workflows.
- Work with LANL DevOps to manage workflows such as scaling studies of production codes (FLAG) since in-house GitLab runners are currently overtaxed.
- Investigating RO_Crate and other means of storing and retrieving provenance data as part of the workflow orchestration system.

Compiler and Programming Model Research (LANL)

Compilers and supporting runtime systems provide the glue between programming languages and the capabilities of the targeted hardware resources including processors, memory subsystems, and accelerators. The software architecture of modern compilers has

yet to evolve to match the rate of innovation appearing at the hardware level. This leaves the low-level details of parallelism, data movement, and concurrency opaque to the compiler as it performs analysis and optimizations. Our efforts aim to modernize compilers, supporting runtime layers, and overarching toolset to recognize, analyze, transform, and directly optimize parallelism, and data movement to improve performance, platform portability, and developer productivity. Another goal of this effort is to provide mechanisms and models for mapping data and processor for executing code that are beneficial to achieving performance portability and flexibility without having to hard-code the platform details into application codes.

This project also serves as the lead institution for NNSA's (ASC's) effort to provide a modern, open-source compiler infrastructure for Fortran. This serves to reduce overall programmatic risk from declining support for Fortran from the vendor community at a time where Fortran codes continue to provide a core set of capabilities across ASC as well as DOE and the international scientific community.

This project will be partially supported by post-ECP software-sustainability funding.

Accomplishments in FY23:

- Kitsune+Tapir outperforming hand-coded CUDA by 2-3x
- Completed Legion co-design with FleCSI of multi-color accessors
- Completed FleCSI co-design for Kokkos support

Planned Activities in FY24:

- Optimization of Legion gather copy performance leading to better scalability and Legion tracing for copy operations leading to reduced runtime overheads
- Production hardening of FleCSI/Legion/Kokkos on pre-El Capitan hardware
- Improve Kitsune+Tapir support for AMD GPUs and Grace-Hopper architecture
- Continue leading the ASC Flang LLVM-based Fortran compiler effort

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.

FleCSI is a middleware software layer that addresses multiple challenges faced in the development of large-scale scientific software. FleCSI separates the computational-science aspects of an application (the physics being simulated) from the computer-science aspects (the interfacing to the underlying hardware). It facilitates the addition of new physics routines to an extant application structure. It also enables applications to exploit hybrid CPU-GPU systems and manycore processors without excessive code surgery.

The FleCSI CSSE project has two components. *FleCSI Core* incorporates state-of-the-art libraries for managing data motion, communication, and parallelism. These include Legion, HPX, MPI, Kokkos, and OpenMP. Based on these, FleCSI Core provides highly customizable, distributed data structures. *FleCSI Specializations* provide a bridge between a physics application and the FleCSI Core, customizing the distributed data structures for the application's specific physics models in a co-design fashion. The FleCSI project is collaborating closely with LANL's Ristra application project (ASC IC) to develop new physics and new computer-science capabilities in support of ASC's future goals and rapidly evolving mission space.

Accomplishments in FY23:

- After two years of active development, FleCSI v2.2.0 was finally released. This version not only is the first to support 100% of the original FleCSI's functionality (before the near-complete rewrite), but it extends the original FleCSI with new features, richer data structures, a more object-oriented API, support for more run-time libraries, and numerous performance enhancements and bug fixes.
- FleCSI now fully supports distributed, multidimensional arrays that store different amounts of data per element. This is needed, for example, by physics applications in which a varying number of materials are present in each spatial region.
- FleCSI is now compatible with recent versions of Legion with better HIP support (needed for El Capitan) following bug fixes on both sides.
- Performance of CPU-GPU communication with the MPI backend has been improved substantially. For example, execution time of a Poisson solver running on an El Capitan development node with eight AMD GPUs has dropped from 23.33 seconds to 2.96 seconds.
- In support of Ristra's Moya application, FleCSI Specializations has implemented a fully functional unstructured mesh specialization for staggered grid hydrodynamics. This supports state-of-the-art 1-D, 2-D, and 3-D unstructured meshes that represent quantities not only at vertices, edges, faces, and cells but also at corners and sides, as needed by the staggered-grid hydrodynamics algorithm.
- Implemented serial X3D support for the Moya application (Burton specialization).

Planned Activities in FY24:

- Improve performance when using the Legion back end by eliminating redundant data transfers and by more intelligently distributing tasks among computational resources.
- Improve portability by collaborating with GPU vendors to ensure that their compilers provide correct implementations of the C++ features required by FleCSI.
- Improve programmability by integrating more Kokkos loop-parallelism features directly into FleCSI.
- Improve applicability by providing a set topology that represents unordered sets of non-interacting physics particles.

- Improve robustness by supporting checkpoint/restart of internal FleCSI state and providing hooks to the application for it to checkpoint/restart its own, non-FleCSI state.
- Improve scope by devising a means of efficiently expressing deterministic transport in terms of FleCSI mechanisms.
- Improve accessibility by completing a FleCSI users guide and other documentation and by offering FleCSI training to users.
- Develop a multi-material model that supports execution on accelerators (GPUs) with material-centric or cell-centric access patterns (possibly both) depending on the requirements of the Ristra project.
- Provide a specialization of the set topology to support material-point method in the Moya application code.
- Extend the Moya input and I/O formats to support X3D (pre-partitioned).
- Add solver support for turbulence models in Moya using Whisk.
- Extend the FleCSolve package with support for algebraic multigrid (AMG).

Advanced Machine Learning (LANL)

The Advanced Machine Learning project has mission-focused efforts in collaboration with PEM, V&V, and ATDM-CDA (to end of FY22). The project conducts studies of Artificial Intelligence (AI), ML, and Deep Learning methodologies applied to, materials modeling, radiographic analysis, advanced uncertainty quantification, and improving and accelerating simulation workflows. 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, multi-fidelity 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.

Accomplishments in FY23:

- Completed the review of eight 3-year AI/ML pathfinding projects. Two will continue in FY24.
- Participated in the writing of the DOE AI4SES strategy document
- Participated in the writing of the ASC AI4ND strategy document
- Worked with AI/ML hardware vendors to explore the suitability of their products to accelerate AI/ML and traditional multi-physics applications

Planned Activities in FY24:

- We will coordinate with the DDMD framework from headquarters and the soon-to-be-released AI4ND strategy document. DDMD refers to a notional breakdown of the weapons lifecycle: discovery, design, manufacturing, and deployment. The framework lists the major areas and the connections between them where we believe that AI can produce major advances.
- We are seeking proposals on the use of AI/ML to address the design or design optimization of a complex engineered dynamic system. This should be a system of programmatic interest.
- We are issuing a call for proposals intended to explore and advance intelligent aids for programming and software engineering in support of nuclear security. All work should focus on an ASC problem of interest or be convincingly transferrable.

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 crosscuts other sub-elements including CSSE and V&V.

Accomplishments in FY23:

- Released next version of the ECMF prototype on the SCN, where team implemented automatic rafting for SRN-to-SCN. (Appendix I, targets, V&V-3 and CSSE-2) [ECMF-2]
- Incorporated ND-relevant exemplars into ECMF with embedded credibility and workflows on SCN for system qualification. (Appendix I, targets, V&V-3 and CSSE-2) [ECMF-1]

Planned Activities in FY24:

- Demonstrate ability to plug-and-play other ND-relevant exemplars into ECMF. (Appendix I, targets, V&V-3 and CSSE-2) [ECMF-1].
- Demonstrate use of ECMF for rigorous, multi-team peer review of high consequence modeling activities for system qualification. (Appendix I, targets, V&V-3 and CSSE-2) [ECMF-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 FY23:

- Hosted 2 'hackathons' for El Capitan Center of Excellence activities (Appendix I, target CSSE-1b)
- Developed benchmarks and support for Crossroads acceptance testing (Appendix I, target CSSE-1a)
- Developed benchmarks and proxy applications for ATS-5 (Appendix I, target CSSE-1e)
- Support DevOps and optimized, containerized builds of Trilinos for code teams (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Support development of benchmarks and proxy applications for ATS-5 and ATS-6 (Appendix I, target CSSE-1e)
- Develop plan for the 30+ ASC code teams to adopt the baseline ASC Unified Environment including a common software stack and improve gathering of usage statistics. (Appendix I, target CSSE-2)
- Support the deployment of codes on Crossroads for production usage (Appendix I, target CSSE-1a)
- Support the porting and optimization of codes for the El Capitan platform working with the COE (Appendix I, target CSSE-1b)
- Support DevOps and optimized, containerized builds of Trilinos for code teams (Appendix I, target 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 FY23:

- Developed a zero-copy API in SEACAS IOSS to support passing memory references to stored LLNL Conduit data on read by ParaView/Catalyst API 2 to be deployed with the Catalyst API 2 database in SEACAS IOSS. (Appendix I, target CSSE-2).
- Implemented prototype zero-copy API in IOSS library for use by Paraview. (Appendix I, target CSSE-2).
- Improved parallel communication paths in IOSS for complex models containing many element blocks compared to processor count. (Appendix I, target CSSE-2).
- Worked with NetCDF, CGNS, and HDF5 developers to improve scalability and I/O performance; submitted as PRs and accepted into their code base. (Appendix I, target CSSE-2).
- Improved support for Discontinuous Galerkin fields in IOSS in EMPIRE ATDM application and ParaView visualization. (Appendix I, target CSSE-2).

Planned Activities in FY24:

- Provide performance improvements and feature development for Sandia's IOSS and Exodus libraries in collaboration with ASC application teams, as well as working with vendors to address performance and bug fixes for third-party I/O libraries (e.g., HDF5, NetCDF, CGNS, ADIOS2). (Appendix I, target CSSE-2)
- Extend Catalyst API2 database implementation in SEACAS IOSS to be 100% compatible with all IOSS features enabling consistency between ParaView Exodus and CGNS readers and data sent to Catalyst. (Appendix I, target CSSE-2)
- Evaluate DDN next generation technologies to develop a viable storage-system solution for future ASC production platforms. (Appendix I, target CSSE-2)

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-II will be used to explore new approaches with the 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.

Accomplishments in FY23:

- Released an initial “standard” ATSE distribution as open source. (Appendix I, target CSSE-2)
- Demonstrated a next-generation containerized workflow on Vanguard-II. (Appendix I, target CSSE-2)
- Contributed to the development of a tri-lab programming environment. (Appendix I, target CSSE-2)
- Developed and deployed ATSE via containers on a Vanguard-II early test system. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Develop, refine, and support ATSE software stack and container support for Vanguard-I, Vanguard-II, and other prototype platforms of interest on SRN and SON networks. (Appendix I, target CSSE-2)
- Work with Vanguard hardware and software vendors and tri-lab user community to mature software offerings. (Appendix I, target CSSE-2)
- Collaborate with tri-lab community on a common HPC software stack and build/continuous integration infrastructure. (Appendix I, target CSSE-2)

- Support the DDN next generation scalable storage collaboration and contract and engage NNSA and Office of Science labs in the activity. (Appendix I, target CSSE-2)
- Leverage ongoing vendor collaborations to explore novel means of workflow construction for heterogeneous platforms. (Appendix I, target CSSE-2)

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.

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 FY23:

- Completed a Gaussian Process (GP)-based multi-fidelity surrogate constitutive model that combines experimental uniaxial stress data demonstrating that the MF approach was roughly twice as accurate as a single fidelity (SF) method. (Appendix I, target CSSE-3)
- Developed environment to capture training data to identify electromagnetic slots in mechanical systems and identified a specific geometry idealization use case for electromagnetic slots. (Appendix I, target CSSE-3)
- Organized a hackathon with SambaNova for the min-trisolve that resulted in the co-design of a new triangular solve algorithm. (Appendix I, target CSSE-3)
- Worked with SambaNova to train a Large Language Model (LLM) for Kokkos code generation. Using OpenAI GPT models, pretraining with selected 20+ Github Kokkos repos. (Appendix I, target CSSE-3)

- Implemented functional PCA (fPCA) on the stronglink time-series classification exemplar to automate feature extraction in support of a full comparative analysis on the variability of ML classification trust scores. (Appendix I, target CSSE-3)
- Performed model experimentation on several open source pre-trained LLMs (EleutherAI GPT-NeoX and Pythia, Meta LLaMa, Databricks Dolly, StabilityAI StableLM). Deployed (closed access) 2 on-prem chatGPT-ish web interfaces at Sandia. (Appendix I, target CSSE-3)

Planned Activities in FY24:

- Develop and expand capabilities in physics-informed machine learning to address shortcomings in material modeling in support of PEM. (Appendix I, target CSSE-3)
- Develop methods to rigorously assess machine learning models for data verification and model qualification in support of V&V. (Appendix I, target CSSE-3)
- Develop machine learning methods to modernize and create efficiencies in meshing and numerical solvers for electromagnetics in support of IC. (Appendix I, target CSSE-3)
- Explore co-design opportunities for AI-hardware relevant to the ASC program. (Appendix I, target CSSE-3)
- Explore opportunities to integrate, deploy, and train on-premises large-language models in support of the broader ND and ASC mission. (Appendix I, target CSSE-3)
- Further develop integration of data-infrastructure advances for digital engineering and advanced machine learning workflows. (Appendix I, target CSSE-3)

Accelerated Digital Engineering (SNL)

The CSSE Accelerated Digital Engineering (ADE) project supports R&D activities to streamline the delivery of ASC application capabilities to the analyst and designer communities by leveraging and enhancing computing-as-a-service capabilities. Modern cloud and web-based technologies have great potential to deliver advanced simulation and computing capabilities as a service, without the need for end-users to install software, navigate complex computing environments, or have expertise in interacting directly with the computing resources underlying sophisticated modeling and simulation applications and workflows. Near instantaneous results can be delivered to facilitate rapid design space exploration and provide seamless extension to large-scale HPC runs for investigating promising solutions. This project intends to move beyond the traditional HPC batch scheduling model by applying modern computing technologies to deliver faster, simpler, and more productive digital engineering capabilities to Sandia's ND community.

Accomplishments in FY23:

- Deployed OpenShift Kubernetes testbed system on Sandia Restricted Network (Appendix I, target CSSE-2)

- Deployed GPU resources to production OpenShift Kubernetes systems to support Accelerated Digital Engineering computing workloads. (Appendix I, target CSSE-2)
- Developed Helm support for DetNet Software-as-a-Service Application to simplify deployment to multiple Kubernetes clusters. (Appendix I, target CSSE-2)
- Benchmarked production S3 object storage performance from HPC and Cloud computing resources to inform future networking connectivity needs. (Appendix I, target CSSE-2)

Planned Activities in FY24:

- Develop reusable templates for exposing containerized ASC Modeling and Simulation Codes via web-based Application Programming Interfaces to facilitate deployment on OpenShift Kubernetes. (Appendix I, target CSSE-2)
- Investigate approaches for supporting Directed Acyclic Graph job workflow dependencies on OpenShift Kubernetes and apply to ADE computing workloads. (Appendix I, target CSSE-2)
- Deploy ADE computing workloads in the classified environment and support early users from Sandia's ND analyst and designer communities. (Appendix I, target CSSE-2)
- Develop S3 object storage backend for Sandia IOSS library and apply to ADE computing workloads. (Appendix I, target CSSE-2)

Projects for the Advanced Memory Technologies and Vendor R&D Product (WBS 1.2.3.5.8)

High-performance systems deployed by the ASC demand the very highest levels of aggregate calculation performance, reliability, and component endurance, combined with optimized cost and energy consumption. The Advanced Memory Technology and Vendor R&D product supports projects which work directly with leading US vendors to codesign future components, systems, and software to meet the needs of the ASC user community. Joint vendor/industry projects are a significant component of the program and often run over multiple years to ensure strong on-going relationships and focus on specific technology outcomes.

Vendor & Multi-Agency Collaborations and Benchmarking (LLNL)

This project covers R&D, vendor contracts, hardware, and software evaluations for next-generation systems including support for the Advanced Memory Technologies and Project 38 collaborations, through which ASC can help drive the research and development in memory, networking, and other technologies needed in the post-exascale era. To help direct these vendor collaborations, this effort also develops benchmarks representing problems-of-interest to ASC, which can be shared with vendors and used to evaluate and design hardware and software on future HPC systems.

Accomplishments in FY23:

- Close interaction with HPE to benchmark the EAS-3 system for system acceptance, facilitate interaction between application teams and HPE to evaluate the performance optimizations applied to LLNL's benchmarks.
- Technical oversight team participation in AMT contracts. Evaluated and provided constructive feedback on reports submitted by AMT awardees and on follow-on/new proposals submitted by vendors for FY24 and beyond funding. Focus on monitoring progress relative to proposed tasks and schedule, applicability, usability, and overall impact on LLNL applications.
- Provided feedback on P38 Micron milestone and other progress reports.
- Installed experimental CXL testbed and began evaluation.

Planned Activities in FY24:

- Close interaction with HPE to benchmark El Capitan for system acceptance, using TOSS and Flux. Facilitating interaction between application teams and HPE to evaluate El Capitan-specific performance optimizations applied to LLNL's benchmarks.

- Continue to participate in technical oversight of existing and new AMT contracts. Provide feedback on relative benefits of new memory technologies for LLNL applications
- For P38, continue evaluation of Micron's experimental testbed. In addition to performance evaluation of benchmarks, provide feedback and direction on their language and tools to facilitate use of near memory computing in larger applications.
- In collaboration with the LBNL P38 Mosaic project, integrate the zhw floating point encoder and decoder into the P38 heterogeneous System on Chip framework.
- CXL evaluation. Use high-level memory access characteristics of a P38 chiplet (e.g., bandwidth and latency profiles/distributions) to assess potential benefits of CXL memory for accelerators with those characteristics.

Vendor & Multi-Agency Collaborations and Benchmarking (LANL)

This project covers R&D, vendor contracts, hardware and software evaluations for next-generation systems including the Advanced Memory Technologies and Project 38 collaborations, through which ASC can help drive the research and develop in memory, networking, and other technologies needed in the post-exascale era. To help direct these vendor collaborations, this effort also develops benchmarks representing problems-of-interest to ASC, which can be shared with vendors and used to evaluate and design hardware and software on future system.

Accomplishments in FY23:

- Continued interactions with Micron and SNL on advanced memory technologies
- Micron shipped an advanced PCIE-CXL hardware testbed platform in June 2023 for their proposed compute near memory accelerator architecture. This is currently being installed at the USRC facility at LANL and plans are under way to make it more widely accessible to various researchers across the lab for evaluations.
- To support the Advanced Memory Technology effort LANL worked with vendors such as SK Hynix to develop advanced memory controller technology to address sparse memory access patterns in ASC codes, making these codes more efficient on CPU and GPU architectures.
- Participated in SNL, LANL joint workshop on microelectronics (Jan 2023). Discussed materials science, AI and simulation research avenues to support the Chips Act initiatives.
- OpenRoads based synthesis, placement and routing ASIC flow was applied to LANL's parametrized HDL architecture for atomistic MD. A 2 FPU unit with hundreds of millions of transistors finishes synthesis, placement and much of the routing in around 70 hours on a typical high-end laptop with the Sky130 node. The flow, however, does not seem to complete with Skywater's open-source 130nm standard cells and fails towards the end of detailed routing. This is most likely a bug

in the OpenRoads tools, we are continuing to look at the causes but isolating a simple instance of the bug is proving to be very difficult. We are also porting Skywater 90nm standard cells for use with OpenRoads flow.

- NDA is being entered into between LANL and Intel Corporation for access to their 12nm/16nm foundry shuttle nodes under a Chips Act related initiative. Expecting to have the NDA in place sometime around start of FY24.

Planned Activities in FY24:

- Continue interactions with Micron and SNL on advanced memory technologies.
- Make Micron PCIE+CXL accelerator testbed and NDGSim simulator accessible to researchers at LANL for benchmarking.
- Finalize NDA with Intel Corp for access to their latest foundry shuttle nodes (beginning with 12nm/16nm for example). Port their standard cells over to OpenRoads flow, and target the atomistic MD architecture for synthesis, placement and routing for potential prototype fabrication.
- Continue targeting Skywater's 90nm and 130nm nodes for OpenRoads flow.
- Begin development of Pipedream software for novel dataflow based graphical high-level synthesis (HLS) of custom hardware accelerators for science algorithms.

Vendor & Multi-Agency Collaborations and Benchmarking (SNL)

This project supports engagement on a range of tri-lab R&D activities including benchmarking for new system procurements, collaboration and evaluation of vendor R&D contracts and support for multi-agency collaboration activities. Strategic tri-lab engagement on these activities is critical for the future of the ASC program. As we move beyond Exascale, the program is increasing investment in innovative new memory and computing technologies including R&D contracts with domestic computing vendors. This team supplies the necessary benchmarking for these technologies, engages in collaboration with vendor partners as well as other government agencies with a vested interest in exploiting cutting edge computing technologies.

Accomplishments in FY23:

- Supported benchmarking development for ATS-5 procurement activities (Appendix I, target CSSE-1e)
- Coordinated with P38 and multi-agency ML/AI activities. (Appendix I, target CSSE-3)

Planned Activities in FY24:

- Development of benchmarks and proxy applications for future Vanguard, ATS and CTS procurements (Appendix I, target CSSE-1)

- Support multi-agency collaboration for future technologies. (Appendix I, target CSSE-3)

Advanced Memory Technologies (AMT) R&D (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 developing R&D contracts with computing vendors 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 FY23:

- Deployment of over $30M of contracts to multiple US vendors through a tri-lab procurement activity supporting R&D for innovative new memory technologies. (Appendix I, target CSSE-3)
- Management of initial contract deliverables for AMT contracts. (Appendix I, target CSSE-3)
- Development of tri-lab and HQ plan for ongoing AMT activities. (Appendix I, target CSSE-3)

Planned Activities in FY24:

- Release RFQs for additional AMT contracts and activities (Appendix I, target CSSE-3)
- Continued management of AMT contracts and collaboration with vendors to develop innovative memory technologies for future ASC systems. (Appendix I, target CSSE-3)

Projects for the Advanced Architecture Prototype Systems Product (WBS 1.2.3.5.9)

The growing diversity and complexity of high-performance computing hardware (including novel accelerators for machine learning), represents a significant and growing risk for the ASC program if thorough analysis of options cannot be adequately performed prior to large-scale system deployments. The Advanced Architecture Prototype Systems (AAPS) product within ASC, allows the program to take a calculated higher risk, higher reward path to deploying moderately sized systems to support evaluation of promising future high-performance computing technologies. Through early prototyping and moderate deployments, ASC code teams are able to evaluate hardware and software technologies at a much earlier point in their lifecycle providing feedback to developers, and vendors on any issues of concern long before such technologies are adopted in production environments. These activities help to lower the risk of performance or functionality mismatches before novel technologies are deployed in AT or CT systems. In addition, the AAPS product helps to demonstrate the viability of off-roadmap technologies for mission applications, allowing them to transition into production use at an accelerated pace. The result is that the ASC program is able to adopt high impact technologies much earlier than originally planned.

Vanguard-II Advanced Architecture Prototype System (SNL)

The scope of this project is to lead the design, acquisition, and plan for deployment of the Vanguard Advanced Architecture Prototype System (AAPS) 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 AAP systems. The first Vanguard AAPS is the Arm-based Astra platform, delivered and deployed in the fall of 2018. Future Vanguard systems will be named as systems as they are moved from conceptual development into procurement development.

Accomplishments in FY23:

- Evaluated the performance of applications on coarse grained reconfigurable dataflow accelerated node architecture for Vanguard-II using vendor simulators, engineering samples. (Appendix I, target CSSE-1.d)
- Delivery of DDN storage technology development contract and coordination of tri-lab collaboration activities. (Appendix I, CSSE-3)

- Evaluated Ctrl-IQ workflow orchestration software on testbed hardware. (Appendix I, CSSE-1d)

Planned Activities in FY24:

- Evaluate additional high risk/high reward technologies for Vanguard-II and -III systems procurements. (Appendix I, CSSE-1.d)
- Oversee DDN storage technology development focused on scalable performance improvements. (Appendix I, CSSE-3)
- Support the development of a Kokkos-based backend for data-flow accelerators (Appendix I, CSSE-1.d)
- Evaluate the performance of next generation of data-flow accelerators with available simulators and early HW in collaboration with vendor (Appendix I, CSSE-1.d)
- Collaborate with vendor on algorithms and implementations of sparse/dense linear algebra and graph kernels on data flow hardware. Evaluate vendor provided implementation (when available), advise vendor on state-of-the-art on algorithmic work, evaluate Kokkos Kernels implementations using vendor provided stack. (Appendix I, target CSSE-2)

Appendix F: 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 2022, and through quarter 3, fiscal year 2023, are reflected below for the FOUS subprogram.

- Completed the El Capitan site infrastructure project to provide electrical, mechanical, and structural modifications required to install the system (LLNL)
- Deployed significant portions of the El Capitan system including all system infrastructure racks and the capacity tier (i.e., the persistent Lustre file system) (LLNL)
- Deployed Racer, a new 20 PB Lustre file system in the LC Restricted Zone, replacing an older system. All user data was successfully migrated from the old system with minimal user disruption. (LLNL).
- Completed TOSS 3 to TOSS 4 conversions across the entire unclassified environment and began TOSS 4 conversions of major systems in the classified computing environment. (LLNL)
- Continued to provide 24x7x365 operational support for all production HPC resources and maintain consistently high system utilization levels and customer productivity. (LLNL)
- Completed CTS-2 production readiness for open and secure production capability with hardware accessible to users and complete integration with existing storage and networking environment. (LANL)
- Deployed Gitlab Continuous Improvement runner functionality at LLNL for LANL users demonstrating value of tri-lab Remote Computing Enablement (RCE) for ASC users. (LANL)

- Continued production operation of Trinity (ATS-1) and CTS-1 systems Snow, Fire, Ice, and Cyclone in full production use, peaking with over 94% utilization and 99% system availability. (LANL)
- Completion of minor construction SCC Electrical Upgrade design and procurement of long-lead infrastructure in preparation for ATS-5. (LANL)
- Provided 24x7 HPC operations of ASC systems and infrastructure to meet ATCC goals. (LANL)
- Increased use of AppSysFusion by multiple application teams' (e.g., EMPIRE, DARMA) developers and analysts for diagnostics and resource allocation decisions, System Administrators using AppSysFusion for Lustre performance and utilization information on SRN. (SNL)
- Released LDMS 4.3.9 - 4.3.11 including internally and externally contributed features for upcoming DOE exascale systems (e.g., slingshot monitoring, Avro-Kafka store, and monitoring Intel's Distributed Asynchronous Object Storage (DAOS)). (SNL)
- Led technical efforts on the resolution of two systemic problems that benefit CTS-2 installations at all sites and overcame the expected new system integration efforts to deploy the CTS-2 cluster Amber into a production ready state. (SNL)
- Collaborated with NA-82 entities and LLNL to specify, configure and start the procurement process for a 6SU CTS-2 HPC cluster to be delivered in FY24. (SNL)
- Established ASC DevOps build, test, and deploy framework and Spack tooling configurations under the ASC Unified Environment (AUE). Demonstrated unified environments viability by delivering a shared complier, MPI, and package shared across three individual codes teams. (SNL)

Milestone (ID#TBD): El Capitan Compute Node Integration		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/24		
ASC WBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-5		
Description: Take delivery of El Capitan compute nodes and prepare the system for integration into the computing environment. Test and demonstrate that the system software is functional. Complete benchmarks and application performance testing.		
Completion Criteria: Hardware deliveries are complete. Installation of the system by the vendor and LLNL personnel is substantially complete. System software needed for operation of the system is installed and functional, including. log collection, system health monitoring, remote hardware and power management, and node provisioning and booting capabilities. Testing in support of final system acceptance has been completed.		
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, CSSE		

Milestone (ID#TBD): SCC Electrical Upgrade Start of Construction		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2024		
ASC WBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target FOUS-3		
Description: Construction activities initiated on minor construction SCC Electrical Upgrade (SEU) project to prepare for increased facility electrical power distribution for ATS-5.		
Completion Criteria: Mobilization complete for construction subcontractor, and site preparation complete for start of construction.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: FOUS		

Milestone (ID#TBD): El Capitan ART/Mini System Acquisition and Delivery		
Level: 2	**Fiscal Year**: FY24	**DOE Area/Campaign**: ASC
Completion Date: 09/30/23		
ASC WBS Subprogram: FOUS, CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix I, target [FOUS-1		
Description: Mini El Capitan System (El Dorado) will be acquired and delivered. Site preparation activities (e.g., location, power, cooling) will be completed.		
Completion Criteria: System acquisition terms will be completed. System will be delivered and sited.		
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# 8640): CTS-2 delivery, site integration and deployment		
Level: 2	**Fiscal Year**: FY23	**DOE Area/Campaign**: ASC
Completion Date: 09/30/23		
ASC WBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-5		
Description: Work is in support of the Commodity Technology System (CTS-2) contract (LLNL Subcontract # B640169). Severe supply chain constraints continue, but the current schedule should hold. Along with the selected vendor and LLNL, SNL has completed the following activities in FY22: final site-specific systems configuration and early parts allocations. Activities in FY23 include final parts allocations, system build at vendor site, delivery of the systems to SNL, installation and testing in the 725 East data center, final integration, and deployment as an ASC HPC resource.		
Completion Criteria: An 8SU (1496 compute nodes) CTS-2 cluster shall: • have all cluster material allocations complete and at vendor integration facility • be built and tested at vendor facility and made ready for delivery to SNL • be delivered by the vendor to SNL Building 725 East • be installed and made operational by the vende~~o~~r and turned over to SNL • be integrated and deployed to early ASC and ND users on SNL's SRN by SNL		
Customer: ASC 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: FOUS, CSSE		

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 FY23:

- Continued FY23 procurement support, contract management, and program planning. (Appendix I, targets FOUS-1 and FOUS-5)
- Managed the PSAAP III program for new Academic Alliance Centers.
- Managed contract and procurement support for CORAL2 and CTS-2 systems. (Appendix I, targets FOUS-1.a and FOUS-5)
- Organized and hosted the ASC Principal Investigator (PI) meeting, held in May 2023 at Y-12

Planned Activities in FY24:

- Continue FY24 procurement support, contract management, and program planning. (Appendix I, targets FOUS-1 and FOUS-5)
- Continue to manage the PSAAP III program for new Academic Alliance Centers, and prepare for PSAAP IV request for proposals.
- 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 FY23:

- Provided LANL detailee support to HQ.
- Supported the programmatic needs of the PSAAP III program as it continues execution.

- Supported the ECI, including joint planning and execution with DOE/NNSA and DOE/SC laboratories.
- Provide program direction to subprogram elements at LANL.

Planned Activities in FY24:

- Organize and host the ASC PI meeting.
- Provide LANL detailee support to HQ.
- Support the programmatic needs of the PSAAP III program as it continues execution.

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 Weapon Engineering and Production Program (which includes advanced and exploratory systems, stockpile systems, modernization programs, and production) 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 FY23:

- Delivered scenario-based combined environments significance matrix. Organized and hosted tri-lab and NNSA HQ AI4ND (previously DDMD) Workshop at Sandia NM.
- Participated in organization of the ASC PI Meeting that was held in May 2023 at Y-12, partnering closely with HQ, LANL, and LLNL.
- Organized and hosted PSAAP III Forum, Albuquerque, NM and PSAAP IV Pre-Proposal Meeting, Houston, TX. Supported programmatic needs of the PSAAP III program as it continues execution.
- Supported the National Academies of Sciences, Engineering, and Medicine (NASEM) post-exascale review in close partnership with HQ, LANL, and LLNL.

Planned Activities in FY24:

- Continue supporting ASC integration and partnership with stockpile and modernization teams.
- Plan and organize the next Predictive Engineering Science Panel (PESP) review.
- Continue supporting the programmatic needs of the PSAAP III program. Organize and host PSAAP IV Proposal Review and Selection meetings.

- Pursue partnership opportunities with DOE/NNSA and DOE/SC laboratories to support the National AI/ML strategy.

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 FY23:

- Converted all unclassified compute and infrastructure TOSS 3 systems to TOSS 4, and began classified system conversions (Appendix I, target FOUS-5)
- Installed CTS-2 systems in all computing environments to augment existing capacity computing platforms. (Appendix I, target FOUS-5)
- Partially deployed Elastic Stack as an eventual replacement for Splunk to achieve improved query performance and advanced machine learning capabilities. (Appendix I, targets FOUS-1a and FOUS-5)
- Refreshed LC NAS disaster recovery environment with new systems on OCF and SCF. (Appendix I, targets FOUS-1.a and FOUS-5)
- Enhanced converged HPC/Cloud capabilities. Implemented LaunchIT, a tool for user self-provisioning of curated containers in the HPC center. Implemented an air-gapped repository design for operating a multifunction cloud-native build capability in a classified environment.
- Implemented IB SAN improvements, including investigation of long-run IB networks to connect B453, B451, and B654 SANs to eliminate need for Lustre routers. Refreshed old IB SAN hardware infrastructure (Appendix I, targets FOUS-1.a and FOUS-5)
- Investigated immersion cooling technologies. Forged partnerships with representative vendors, developed safety plans, and ordered initial hardware for evaluation.

Planned Activities in FY24:

- Deploy replacements for old classified Lustre file systems with new hardware, including migration of existing data. (Appendix I, targets FOUS-1.a and FOUS-5)
- Fully deploy center-wide monitoring solution based on Elastic Stack as a replacement for Splunk. (Appendix I, targets FOUS-1.a and FOUS-5)
- Migrate LC web and collaboration tool to containers to allow for rapid updates and reduced downtime.
- Enhance and deploy LC infrastructure container storage to provide fault tolerance, and disaster resiliency for building outages due to power or cooling failures. (Appendix I, targets FOUS-1.a and FOUS-5)
- Upgrade LC's converged infrastructure hardware allowing for increased network speed and bandwidth, larger memory, and higher CPU based applications supporting LC infrastructure. (Appendix I, targets FOUS-1.a and FOUS-5)
- Deploy VAST NFS environment on SCF to allow for NFS high-speed scratch storage. (Appendix I, targets FOUS-1.a and FOUS-5)
- Deploy VNC replacement servers to allow for higher capacity workloads enhancing user productivity and overall use of the service.
- Further leverage DevOps tools such as Ansible and Gitlab in the deployment of LC systems to achieve a consistent and more reliable deployment of LC compute, file systems and infrastructure environments. (Appendix I, targets FOUS-1.a and FOUS-5)

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 FY23:

- Provided ongoing user support through hotline operations, documentation, and training while in minimum-safe and minimal-normal operations. (Appendix I, targets FOUS-1.a and FOUS-5)
- Began development of documentation and training materials in the use of El Capitan and the CTS-2 class clusters. (Appendix I, targets FOUS-1.a and FOUS-5)
- Provided support to the 9 ASC PSAAP 3 Alliance centers, with a focus on porting codes to TOSS 4 on the Quartz CTS-1 cluster.

Planned Activities in FY24:

- Continue providing support through hotline operations, documentation, and training with a focus on assisting users porting to EAS-3 and CTS-2 systems. (Appendix I, targets FOUS-1.a and FOUS-5)
- Continue development of documentation and training materials in the use of EAS-3, El Capitan and the CTS-2 clusters. (Appendix I, targets FOUS-1.a and FOUS-5)
- Continue to provide support to the 9 ASC PSAAP 3 Alliance centers, with a focus on porting codes to EAS-3 (Tioga), El Capitan and CTS-2 architectures.

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 FY23:

- DisCom 100G Upgrade – Completed Phase-3 - transitioned production workloads over to the full 100G DisCom path end-to-end.
- Open Compute Facility (OCF) Network Re-architecture – Transitioned majority production workloads over to the new IP address ranges.
- Open Compute Facility (OCF) firewall consolidation – Improved network security architecture while migrating and consolidating multiple Juniper Networks firewalls to single Palo Alto Network next-generation firewall hardware.

- Completed El Capitan Site Infrastructure Project Construction in B-453 in conformance to the design completed in FY22. (Appendix I, target FOUS-1.a)
- Completed Site Infrastructure Project Design for Tuolumne, the unclassified analog to El Capitan in B-453. (Appendix I, target FOUS-1.a)

Planned Activities in FY24:

- Complete Site Infrastructure Project for Tuolumne, the unclassified analog to El Capitan in B-453, in conformance to the design completed in FY23 (Appendix I, target FOUS-1.a)
- Complete Open Compute Facility (OCF) Network Re-architecture project – Migrate the remaining production workloads over to new IP address ranges.
- DISCOM iHPC 100G Network Upgrade – Complete Phase-1 procurement of required hardware.

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 FY23:

- Establish new infrastructure to support DevOps oriented approach to managing secure systems. (Appendix I, target FOUS-1.b and FOUS-5)
- Production use of new infrastructure to support DevOps oriented approach to managing open systems. (Appendix I, target FOUS-1.b and FOUS-5)
- Completed CTS-2 production readiness for open and secure production capability with hardware accessible to users and complete integration with existing storage and networking environment. (Appendix I, target FOUS-5)
- Enhancement of utilization and availability of AT and CT systems. (Appendix I, target FOUS-1.b and FOUS-5)
- Continued production operation of Trinity (ATS-1) and CTS-1 systems Snow, Fire, Ice, and Cyclone in full production use, peaking with over 94% utilization and 99% system availability. (Appendix I, target FOUS-1.b)

Planned Activities in FY24:

- Decommissioning of Trinity (ATS-1) and CTS-1 systems Snow, Fire, and Ice. (Appendix I, target FOUS-1.b)
- Continued operation of production CTS-1 Cyclone system. (Appendix I, target FOUS-5)
- Enhancement of utilization and availability of AT and CT systems. (Appendix I, target FOUS-1.b and 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 FY23:

- Provided 24/7/365 Tier-1 support and operations of network and file system resources. (Appendix I, target FOUS-5)
- Continued upgrade of facilities monitoring sensors, network, and services provided for the Strategic Computing Complex.
- Managed Classified Removable Electronic Media (CREM) for ASC storage systems.
- Supported the onsite immersive training classroom for ASC users.

Planned Activities in FY24:

- 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 CREM from creation to destruction and disposition for storage systems supporting ASC simulations.

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 FY23:

- Installed Crossroads and CTS-2 system hardware. (Appendix I, targets FOUS-1.b & FOUS-5)
- Identified and began implementation of new efficiencies in computer cooling (installation of cooling tower agitators) and electrical operations (ordered equipment for chiller plant economizer) to minimize impact on water and power resources.
- Supervised and performed preventive, predictive, and corrective maintenance on HPC facility mechanical and electrical equipment. (Appendix I, target FOUS-1.b and FOUS-5)

Planned Activities in FY24:

- Begin facility preparation for hosting the ATS-5 system at the SCC. (Appendix I, target FOUS-3)
- Initiate implementation of new efficiencies in computer cooling (cooling tower agitators) and electrical operations (chiller plant economizer) to minimize impact on water and power resources.
- Operation of Crossroads and CTS systems for production use supporting ASC simulations. (Appendix I, targets FOUS-1.b & 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 FY23:

- Implemented prototype for GUFI and initial CONDUIT data management features and controls to support user data management.

- Deployed initial campaign storage for supporting the Crossroads and CTS-2 systems in production. (Appendix I, targets FOUS-1.b and FOUS-5)
- Supported the RCE working groups with parallel infrastructure required for tri-lab systems. (Appendix I, target FOUS-1)

Planned Activities in FY24:

- Production support for GUFI and CONDUIT data management services to support ASC simulation workflows.
- Design for supporting machine learning workflows and infrastructure.
- Support the RCE working groups with parallel infrastructure required for tri-lab systems. (Appendix I, target FOUS-1)

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 FY23:

- Identified priority or milestone efforts on AT and CT systems requiring focused or intensive user support, and provided that support through small team direct engagements.
- Participated and lead a programming and runtime environment RCE working group for tools and techniques that improve ASC user productivity on tri-lab systems. (Appendix I, target FOUS-1)
- Deployed Gitlab Continuous Improvement runner functionality at LLNL for LANL users demonstrating value of tri-lab RCE for ASC users. (Appendix I, target FOUS-1)
- Supported production visualization for ASC simulations. (Appendix I, targets FOUS 1.b & 5)

Planned Activities in FY24:

- Participate in RCE working groups for tools and techniques that improve ASC user productivity on tri-lab systems. (Appendix I, target FOUS-1)

- Continue service for production visualization of ASC simulation results supporting stockpile stewardship. (Appendix I, targets FOUS 1.b & 5)
- Management and resolution of ASC user challenges on infrastructure and platforms to accomplish Annual Assessment Report and ASC milestones in other program elements. (Appendix I, targets FOUS 1 & 5)

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 FY23:

- Continued development of Marchive to enable next generation simulation archiving for ASC users.
- Implemented new archive user storage workflow on production campaign and archival storage for Crossroads and CTS systems. (Appendix I, targets FOUS 1.b & 5)
- Provided top-tier diagnostic and troubleshooting support of archival storage for ASC users.

Planned Activities in FY24:

- Improve archival storage capabilities for data supporting ASC verification and validation.
- Continue development of production functionality for Marchive to enable next generation digital archive capabilities for ASC users.
- Provide top-tier diagnostic and troubleshooting support of archival storage for ASC users.

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., Tachi (Crossroads), Vortex (Sierra)), at-scale prototype platforms (e.g., Astra), and advanced architecture testbeds.

Accomplishments in FY23:

- Increased use of AppSysFusion by multiple application teams' developers and analysts (e.g., EMPIRE, DARMA) for diagnostics and resource allocation decisions, System Administrators using AppSysFusion for Lustre performance and utilization information on SRN. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Multiple successful LDMS and AppSysFusion education and outreach activities: LDMS User Group Conference (LDMSCON) held in 2022 and 2023, and tutorials and/or presentations given to HPC production system administrators and EMPIRE, RAMSES, DARMA, ECMF teams. (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Expanded deployments including LDMS on all SRN production systems. Modules for collecting Darshan I/O data into AppSysFusion deployed on Eclipse and Stria, further interoperability of DOE labs tools. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Stood up new hardware and developed a process to upgrade to HPSS v9.3, in consultation with IBM, and executed that in the test environment in preparation for FY24 production upgrades (Appendix I, target FOUS-5)
- Led technical efforts on the resolution of two systemic problems that benefit CTS-2 installations at all sites and overcame the expected new system integration efforts to deploy the CTS-2 cluster “Amber” into a production ready state. (Appendix I, target FOUS-5)
- Brought TOSS4 into our image management system and deployed on Amber. The roadmap is defined for its deployment to Sandia’s legacy HPC systems. (Appendix I, target FOUS-5)
- Delivered and installed Crossroads ART platform, “Tachi”. (Appendix I, target FOUS-1)
- Facilitated frequent hardware and software changes to Vanguard II early systems and other CSSE testbed systems. (Appendix I, target FOUS-2)
- Provided support of local LANL and LLNL ATS ART platforms. (Appendix I, target FOUS-1)
- Provided program support and coordination for tri-lab ATCC.
- Collaborated with NA-82 entities and LLNL to specify, configure and start the procurement process for a 6SU CTS-2 HPC cluster to be delivered in FY24.
- Provided operational support of existing HPC systems and associated capabilities.

Planned Activities in FY24:

- Continue expansion of the application user base using AppSysFusion to address performance issues and to inform resource allocation decisions on CTS and ATS systems. (Appendix I, targets FOUS-1, FOUS-2 and FOUS-5)
- Expand deployments of AppSysFusion and LDMS on Sandia and external systems. (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Develop enhancements to LDMS and to the configuration and installation processes of LDMS, in collaboration with DOE labs and vendors, for operation on upcoming CTS and ATS systems. (Appendix I, targets FOUS-1, FOUS-5, and CSSE-2)
- Provide training, education, user group meetings to educate users and system administrators on capabilities, new features, and directions of AppSysFusion and LDMS. (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Deploy an integrated Starfish/FUSE solution for directed and/or automated data transfers to HPSS. (Appendix I, targets FOUS-5, and CSSE-2)
- Complete tasks for HPSS hardware/OS/software upgrades across all environments to update to RHEL8 and HPSS v9.3. (Appendix I, target FOUS-5)
- In support of ASC L2 milestone # TBD, collaborate with LLNL to plan for, integrate and complete early installation of the Eldorado system. (3 rack El Capitan) (Appendix I, target FOUS-1)
- Complete TOSS4 installation on legacy HPC systems in all four security environments. (Appendix I, target FOUS-5)
- Specify requirements and start procurement process for the 2nd ASC funded CTS-2 system. (Appendix I, target FOUS-5)
- Apply knowledge from ASC CTS-2 Amber cluster deployment to the deployment of non-ASC CTS-2 clusters including Condo (institutional), Hops (institutional) and Dark Mead (NA-82). (Appendix I, target FOUS-5)
- FrETT test deployment in SNL RR/LLNL RZ/LANL RE. (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Continuing support of Vanguard II test systems. (Appendix I, target FOUS-2)
- Production support of the Crossroads ART system, Tachi. (Appendix I, target FOUS-1)
- Deploy the NA-82 funded 6SU CTS-2 based cluster (DarkMead) and associated supporting infrastructure. (Special Projects)
- Retire legacy TLCC2 HPC system (including EoL filesystems) and migrate user operations to next-generation platform. (Special Projects)
- Continue operations and support of NSCC computing environment. Refine operations strategy around continued HPC capabilities in this environment. (Special Projects)

- Engage benefitting program stakeholders to develop a long-term capital investment strategy for future platforms. (Special Projects)

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 FY23:

- Released LDMS 4.3.9 - 4.3.11 including internally and externally contributed features for upcoming DOE exascale systems (e.g., slingshot monitoring, Avro-Kafka store, and monitoring Intel's Distributed Asynchronous Object Storage (DAOS)). (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Released the Kokkos-LDMS Connector as part of the open-source Kokkos code base.. The Kokkos-LDMS Connector enables users to inject Kokkos kernel timing information into the LDMS transport. This collaborative work motivated a refactoring of the functionality in Kokkos to more generally support such interoperability. (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Integration of collaboratively developed ML-based analyses into the production AppSysFusion infrastructure for exploring application profiling and anomaly detection. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Organized and participated in international meeting to advance HPC Operations through Monitoring and Operational Data Analytics (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Installed and configured a next-generation FUSE/Starfish testbed system and tested new features and compatibilities for future production rollout. (Appendix I, targets FOUS-5, and CSSE-2)
- Extensive development of Ansible and Git-based cluster management tooling on testbeds resulting in very rapid stand-up and change capability with deterministic results. (Appendix I, target FOUS-2)
- Implementation of advanced user workflow software (vendor CIQ) for containerized and cloud technology workloads. (Appendix I, target FOUS-2)
- Increased the usage of user facing file system metadata tools (e.g., Starfish) to improve user data management capabilities and improve workflow.

Planned Activities in FY24:

- Continue improvements and associated code releases of LDMS and Scalable Object Store (SOS). (Appendix I, targets FOUS-1, FOUS-2, FOUS-5, and CSSE-2)
- Continue development and expansion of AppSysFusion infrastructure, interoperability, analyses, and visualizations. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Expand and assess feasibility of and requirements for time-series and ML-based analyses in the production AppSysFusion infrastructure for application profiling and anomaly detection. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Prototype direct feedback mechanisms from LDMS and AppSysFusion analyses to system and application software to explore data-driven operations opportunities. (Appendix I, targets FOUS-2, FOUS-5, and CSSE-2)
- Set up a test environment for a Unified Namespace (UNS) using Starfish and HPSS Rumbler to determine viability, using HPSS v10.2. (Appendix I, targets FOUS-5 and CSSE-2)
- Build web-based container registry tool over HPSS/FUSE to support SNL testbed container-based platform deployment. (Appendix I, target FOUS-2)
- Continue vendor codesign explorations in advanced system management and functionalities for Vanguard II (e.g., Ctrl-IQ workflows, Warewulf system provisioning, DDN storage). (Appendix I, target FOUS-2)
- Continue migration, started in FY2023, of testbed systems from ad-hoc system management to an Ansible-based, Git-backed system management stack. (Appendix I, target FOUS-2)
- Support deployment of containerized workflows on testbed systems, including scientific applications and services like LDMS. (Appendix I, target FOUS-2)
- Investigate and deploy a test environment for running HPC/AI/ML mixed workloads within a traditional HPC production system.

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 FY23:

- Provided service desk, Tier-3 user support, and service desk improvements for SNL and tri-lab ASC computing systems.
- Maintained and ensured developer tools (profilers, debuggers, and system performance and analysis) were available for ASC development teams and user workflows.
- Established ASC DevOps build, test, and deploy framework and Spack tooling configurations under the ASC Unified Environment (AUE). Demonstrated unified environments viability by delivering a shared complier, MPI, and package shared across three individual codes teams.

Planned Activities in FY24:

- Provide service desk, 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)
- 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)

Software Environment Deployments (SNL)

The ASC DevOps integration effort is executing a vision to move SNL's 30+ code teams to similar software stacks to improve the overall efficiency of the delivery of ASC Codes to the Nuclear Deterrence Mission. This includes working with the tri-lab community to expand a consistent software stack across the Labs. In addition, an effort to have an annual development and operations (DevOps) conference as part of the annual National Labs Information Technology (NLIT) summit to significantly increase lab-wide communications at technical levels. This includes continued engagement with the Tri-lab programming environments initiative.

Accomplishments in FY23:

- New project in FY24

Planned Activities in FY24:

- Develop plan for the 30+ code teams to adopt the baseline ASC Unified Environment which includes the same compiler and MPI software stack. (Appendix I, targets CSSE-2, ATDM-1, and IC-1)
    - Develop an adoption plan for each team.
    - Determine which subset of teams will adopt the AUE by based on the plan.

- Add data collection for each execution of an ASC code on HPC clusters to get consistent usage statistics across the 30+ code teams. (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 FY23:

- Started additional 4MW power project to 725East data center.
- Completed power and cooling site prep for CTS-2 system.
- Started design work on cooling system for 725East data center.
- Completed the South 100Gbps WAN link deployment for DisCom.
- Continued design of iHPC upgrade to 100Gbps.
- SRN HPC network began the migration into an isolated network security zone.

Planned Activities in FY24:

- Complete additional 4MW power project to 725East data center.
- Finalize design and complete cooling project for 725East data center.
- Update DisCom network service monitoring and fault isolation and create a standard operating procedure for carrier and site interactions and communications.
- Establish network telemetry and analytics capabilities for DisCom to enable 100G network tuning.
- Upgrade, diversify and consolidate DisCom black router hardware.
- Replace RedEagle COMSEC equipment ahead of December 2023 mandatory upgrade.

- Improve isolation and reduce implicit access between SRN HPC and local (non-HPC) site resources.
- Update iHPC 100G design (with feedback and lessons learned from FY23 SRD upgrade) and complete upgrade.

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

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 FY23:

- Maintained ES license at $50K and leveraged support contract, improved packaging, and distribution of HDF5 and provided priority response to NNSA tri-lab user requests.
- SNL successfully leveraged this contract to resolve a number of issues related to the CGNS effort.

Planned Activities in FY24:

- Continue ES license support contract for FY24 at $50K.
- Continue to leverage support contract, improve packaging and distribution of HDF5 and provide priority response to NNSA tri-lab user requests.

System and Environment Administration and Operations (PX)

This project provides necessary operational support for a reliable Linux/GPU environment to enable production process modeling in support of the Production Science Initiative. 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.

Accomplishments in FY23

- Pursued support for the High-Performance Compute cluster was pursued through the Enterprise Estimate Review Board (EIIRB) process which ranked this work as second on the Pantex Priority List. This resulted in the formation of a Demand Review Team (DRT). This review prepared Business Partner Review (BRP) concluded that the plan was technically sound and that IS&S could support the work technically, but would need to pursue a staff augmentation contract in order to have the man power. Initial

estimates were about one year to get the contract in place. Documentation has been prepared and reviewed in support for this staff augmentation contract.

- Identified additional software most likely to be useful in nonproliferation and sought initial approval for its use. This software includes MNCP, Monte Carlo N-Particle code used for general-purpose transport of many particles, as well as GADRAS, Gamma Detector Response and Analysis Software used for analyzing radiation detector responses.

Planned Activities in FY24:

- Finalize the installation and set-up of a five-compute node, one GPU node high performance compute cluster on the unclassified network. This system will initially support the objectives laid out over the next five years, but will also be made available to others as needed.
- Explore additional software developed by other NNSA, DOE, and other governmental agencies for their applicability to Pantex.
- Evaluate existing techniques to interface multiple code bases and recommend a set of best practices, selecting software used to model and interpret signals in seismic and infrasound regions as an initial test case.

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.

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 FY23:

- Continued support of TOSS3 for TLCC2 and CTS1/1+ systems with quarterly updates. (Appendix I, targets CSSE-2.b-d, FOUS-2, and FOUS-5)
- Continued support of TOSS4 for CTS1/1+/2 systems with monthly updates. Released TOSS 4.5 (based on RHEL 8.7) and 4.6 (based on RHEL 8.8). Supported the

integration of TOSS4 on CTS-2 clusters. (Appendix I, targets CSSE-2.b-d, FOUS-2, and FOUS-5)

- Collected requirements for the next version of TOSS and began evaluation of RHEL 9 in preparation for TOSS5. (Appendix I, targets CSSE-2.b-d, FOUS-1.a, FOUS-2, and FOUS-5)
- Collected requirements for the next version of TOSS and began evaluation of RHEL 9 in preparation for TOSS 5. (Appendix I, targets CSSE-2.b-d, FOUS-1.a, FOUS-2, and FOUS-5)
- Integrated three new LDMS versions (4.3.[9-11]) into TOSS which provided a variety of new features and upgrades including Kafka bus and Variorum stores, new samplers for the slingshot network and process tracking, improved logging capabilities, and more flexible dataset management. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Organized and led a two-day CCE-wide meeting in February at SNL where the community came together to share tri-lab progress, discuss collaboration opportunities, and create new connections within and between subgroups. NNSA HQ members commented this was a worthwhile meeting and should be repeated yearly. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Continued university research collaborations and explored machine learning anomaly detection with less data, anomaly detection with time-series clustering, and dynamic resource management with a total of six accepted publications over the course of the year. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Added community analyses to the MMAI common analysis repository and moved the repo to a more accessible area within the LLNL gitlab. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- LLNL and SNL planned a database comparison experiment for continuous storage and retrieval of HPC monitoring data. This will result in a test harness that the community will be able to use and add onto to compare performance of databases on their own hardware for their use cases. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Created and configured a Slurm REST API server for community collaboration on common Slurm polling mechanisms. (Appendix I, targets CSSE-2, FOUS-1 and FOUS-5)

Planned Activities in FY24:

- Continue limited support of TOSS3 on legacy tri-lab systems until the RHEL 7 End-Of-Life in 3Q24. (Appendix I, targets CSSE-2.b-d, FOUS-2, and FOUS-5)
- Continue support of TOSS4, including integration on future CTS-2 systems and El Capitan. (Appendix I, targets CSSE-2.b-d, FOUS-1.a, FOUS-2, and FOUS-5)

- Continue development of TOSS5, based on RHEL 9. (Appendix I, targets CSSE-2.b-d, FOUS-1.a, FOUS-2, and FOUS-5)
- Explore, develop, and deploy state-of-the-art HPC monitoring, analysis, visualization, and feedback techniques across tri-lab HPC infrastructures. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)
- Collaborate with tri-lab members on new LDMS features and push of new LDMS releases into TOSS. (Appendix I, targets CSSE-2, FOUS-1, and FOUS-5)

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. Developer Tools Confidence Suite (DTCS), testing tools, and containerized environments similarly promotes operational efficiency across the tri-lab systems.

Accomplishments in FY23:

- Provided workshops for HPCToolkit, TotalView, and TAU, and more are planned.
- Teamed with developers for Rice University to integrate HPCToolkit with ATS-3 Crossroads Center of Excellence, worked heavily on ATS-4 El Capitan architecture, and collaborated with SNL ASC DevOps Core team to support correct HPCToolkit deployments across a wide range of HPC and Linux platforms.
- Provided ongoing support and enhancements for DyninstAPI.

- Deployed Trenza Survey tool at SNL to evaluate and demonstrate impact with application code teams and perform comparisons on CTS-2/ATS-3 systems. Potential license at LANL being discussed.
- Continued to develop Pavilion2 features (such as plotting plug-in), test integrations, and dashboards. Pavilion2 used heavily as primary test driver for ATS-3 Crossroads and related systems and acceptance testing efforts. LANL dedicated system time testing utilizes Pavilion2.
- Incorporated cross-platform testing enhancements (focused on preparing for CTS-2) into Developer Tools Confidence Suite (DTCS) and CDash results updated nightly.
- Provided tri-lab Computing Environment (TCE) is now part of the default Programming Environment on ATS-3 Crossroads. TCE releases are developed and deployed targeting Cray Programming Environment releases on LANL's Cray EX systems.
- Identified network fabric related MPI layer constraints that impact unprivileged containers.
- Built and successfully ran an unprivileged MPI container that: a) is modeled around network fabric related MPI layer constraints, and b) worked on both TOSS3 systems with OPA/IB and Shasta Slingshot 11 fabrics at LANL, SNL, and LLNL.

Planned Activities in FY24:

- Continue to manage contracts for HPCToolkit and DyninstAPI support out of Rice University and the University of Wisconsin-Madison, providing cross-platform HPC performance analysis tools on tri-lab systems.
- Continue Pavilion2 development in support of continuous testing, acceptance testing, performance testing and broader support for tri-lab resources.
- Tri-lab Container Environment study. Create a base image with a representative application and shared environment stack and get it running across the tri-labs. Initially target CTS-2 and change as appropriate.
- Pursue TCE convergence and deployment of a shared ASC Programming Environment infrastructure for ubiquity and shared capability in the tri-labs.
- Continue tri-lab collaboration efforts to track and visualize the various packages each lab is providing to its respective user community to facilitate the determination of a common set of packages all three labs can provide to their users.

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 implementations 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 FY23:

- Continued SARAPE set up for migration of code and data in preparation for modernization/upgrade to Ruby 3.1.2 and Rails 7.0.4.
- Completed prioritized process improvements for SARAPE and onboarded additional resources.
- Successfully planned, coordinated, and delivered significant software improvements to the Workload Characterization Tool (WCTool) resulting in improved quarterly tri-lab utilization data reports to HQ.

Planned Activities in FY24:

- Continue to operate the tri-lab SARAPE process for all remote access account requests and continue to implement SARAPE process improvements.
- Continue to deliver quarterly utilization reports to DOE HQ. Plan, coordinate and implement additional utilization reporting as requested by DOE HQ. 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 FY23:

- Deployed to production a fully designed (i.e., including redundant capability) and senary-way operational 100Gbps DisCom network to the four tri-lab sites (LLNL; LANL; SNL, CA; SNL, NM) representing a 10-fold increase in bandwidth on this data transfer and visualization backbone.
- Deployed to production a senary-way ESN classified frontend network bandwidth of 10Gbps representing a 10-fold increase in bandwidth.
- Deployed a senary-way production GA cross-site production CI in the restricted unclassified component of the iHPC network space (RE/RR/RZ), allowing users to conduct CI pipeline actions at remote-to-them HPC centers.
- Updated, signed and inplemented the Tri-lab IHPC Service Level Agreement for RCE Continuous Integration.
- SNL deployed a new GitLab system in the SNL RR to replace the initial prototype GitLab system for integration with larger restricted unclassified iHPC tri-lab projects (e.g., cross-site CI, code repo mirroring for cross-site co-development, etc.) and integrated with SARAPE account management system.
- SNL tested RealVNC server in the SNL RR login nodes to support LLNL RZ and LANL RE users' VNC clients.
- LLNL established an iHPC trust model with LANL, with all required security approvals and configuration changes, such that incoming SSH IHPC connections from LANL to the LLNL RZ do not require LLNL RZ tokens.
- LANL established an iHPC trust model with LLNL and SNL on ihpc-gate to accept incoming SSH iHPC connections unprompted with a previously existing local MFA Kerberos credential from those institutions on that gateway host.
- SNL established an iHPC trust model with LANL and LLNL on ihpccluster to accept incoming SSH iHPC connections unprompted with a previously existing local MFA Kerberos credential from those institutions on that gateway cluster.
- LANL planned and procured a restricted unclassified data transfer cluster (DTC) to align with DTCs in place at SNL and LLNL for purposes of tri-lab data transfer and file system access endpoints.
- Tri-labs have fully integrated the DOE OneID DUIDs into each of their local LDAP infrastructures to be used for authentication purposes in consideration of an NNSA

tri-lab HPC user base (e.g., used most recently for cross-site CI and potential future use cases such as a ubiquitously available distributed single file system).

- LLNL prototyped RealVNC server configuration allowing LANL remote client Kerberos credential authentication.
- LANL deployed RE-FTA DTC into iHPC RE enclave for upcoming senary-way WAN data transfer capability with LLNL and SNL.

Planned Activities in FY24:

- Deploy a classified cross-site CI capability modeling the FY23 unclassified restricted deployment of cross-site CI, to include security approvals, network access requirements, and application capability deployments.
- LLNL and SNL to works towards establishing a restricted unclassified iHPC trust model for unprompted MFA-based SSH command line (CLI) connectivity to include identification of blockers and potential solutions towards modeling the functionality recently made available in the LLNL/LANL and LANL/SNL bi-directional SSH CLI relationships.
- Investigate allowing remote desktop IP ranges across the iHPC network directly to RE/RR/RZ web server fronts (i.e., https protocol traffic).
- Conduct initial study of a tri-lab file system (POSIX) and object store (S3 endpoint) singular resource capability, with distribution across all three HPC centers as a singular namespace to include initial use case study, vendor engagement, and proof of concept analysis; may include leveraging DOE OneID DUID integration work previously completed at all sites.
- Tri-labs to deploy data transfer tools in the restricted unclassified iHPC ubiquitously across sites, modeling same environment deployed in FY20 in the classified environment.
- LANL to remove RE gateway known as ihpc-gate for both inbound and outbound iHPC traffic gatekeeping purposes, favoring access to frontend nodes of clusters.
- Pursuant to funding, upgrade iHPC data transfer and visualization network to 100Gbps, leveraging line upgrades associated with and the following the deployment model of the classified DisCom WAN 100Gbps upgrade completion of FY23.
- Deploy to production an LLNL RealVNC server configuration allowing LANL remote client Kerberos credential authentication coinciding with TOSS4 upgrades on same architecture.
- Tri-lab HPC centers continue user-facing web documentation refresh to demonstrate congruity with shared services via a common set of FAQs, common-application use examples, etc.
- Continue to engage with CCE tri-lab container portability efforts in furthering tri-lab build and run methods, cross-site container registry access, and user-facing documentation.

- Continue to seek opportunities to ease remote user burdens with respect to lifting gateways, allowing routing to avoid complicated tunnelling machinations, enabling use of local MFA for remote resources, and related streamlining.

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 ~~P~~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 FY23:

- Replaced Lustre storage (for air-gapped network)
- Procured spare parts for out-of-warranty CTS-1 systems
- Upgraded infrastructure upgrade including power, seismic isolation bases, cooling, and overhead cable trays

Planned Activities in FY24:

- 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.
- 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 FY23:

- Completed special purpose program coordination and handoff for facility modifications to SCIF to host Crossroads/CTS-2 hardware for Global Security.
- Transitioned staff supporting special purpose programs to continue providing mission support
- Installed and replaced file system supporting a special purpose program
- Deployed updated CTS-1 hardware

Planned Activities in FY24:

- Integrate new infrastructure and compute hardware into existing environment and prepare for production use
- Initiate construction of facility modifications in SCIF to host new Crossroads/CTS-2 hardware
- Transition user codes to new system architecture

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.

Activities for this project are included under Projects for the Systems and Environment Administration and Operations Product (WBS 1.2.3.6.2) Production Computing Services (SNL) project for FY23 and beyond.

APPENDIX G: Academic Programs

Computational Science Graduate Fellowship (CSGF)

The goal of the DOE Computational Science Graduate Fellowship (CSGF) program is to cultivate the next generation of scientists and engineers in computational sciences. For NNSA, CSGF supports the ASC and Stockpile Modernization missions by establishing academic programs for multidisciplinary simulation science and providing students an opportunity to develop weapons codes through open science applications. The NNSA CSGF program is managed by the Krell Institute and jointly funded with the DOE Office of Science's Advanced Scientific Computing Research program.

The DOE CSGF fosters a community of enthusiastic and committed doctoral students, alumni, DOE laboratory staff and various scientists who desire to have an impact on national security and energy missions while advancing their research. It increases collaboration between NNSA national security laboratories, the fellows, and their universities by enhancing the fellows' research experience at the National Laboratories via access to unclassified high-performance computing systems and exposing them to the broader, multi-disciplinary research activities at the laboratories. The program also provides a yearly stipend, tuition fee coverage, and academic allowance.

Accomplishments in FY23:

- Selected another record level of fellows and provided benefits in STEM fields that use high performance computing to solve complex science and engineering problems.
- Held a successful annual CSGF Program Review that highlighted incoming and ongoing CSGF fellows' research work performed during the year.

Planned Activities in FY24:

- Collaborate with DOE Office of Science in funding new cohort of fellows to be trained as next-generation leaders in computational science.
- Support CSGF community of energetic and committed doctoral students, alumni, and DOE/NNSA laboratory staff who together serve as a support system for the new and current fellows.

PSAAP III Centers

The following FY23 Accomplishments and FY24 Planned Activities are for the PSAAP III Centers, which are activities that started in September 2020.

National Nuclear Security Administration

University of Colorado at Boulder

Center for Micromorphic Multiphysics Porous and Particulate Materials Simulations with Exascale Computing Workflows

Accomplishments in FY23:

- Developed a methodology for pressing mock High Explosive (HE) (F50 sand in Kel-F binder, and IDOX crystals in Estane binder) to mimic the stress condition experienced in larger samples.
- Conducted single grain quasi-static and mini-Kolsky bar experiments on individual F50 sand grains, IDOX crystals, composite F50-sand/Kel-F, IDOX/Estane mock HE, Kel-F and Estane binder cylinders, and also stress-relaxation experiments on the latter two.
- Used the Center's newly developed open-source GPU-accelerated Ratel code for production quasi-static simulations, and developed plans to incorporate an implicit Material Point Method (MPM) capability. In the interim, LLNL's GEOS-MPM will be used for "fast" dynamic simulations and SNL's LAMMPS-granular for the quasi-static press and mechanical behavior simulations.
- Refined the workflow to generate segmented images of grain-scale geometry using Segmentflow for initial Ratel finite element (FE) and GEOS-MPM discretization, and applied it to segment IDOX grains from Estane binder and voids.
- Fitted probability density functions from experiments and simulations of single grain compression to failure to slope data (related to Young's modulus) and peak data (related to failure strength).

Planned Activities in FY24:

- Apply methodology for Uncertainty Quantification (UQ)-based calibration of pure binder and single grain behavior to IDOX grains and Estane binder, extending the FY23 work on F50 sand grains and Kel-F binder.
- Apply the annual predictive simulation workflow to 0.5in diameter cylinders of IDOX/Estane mock HE.
- Implement an in-situ micromorphic filter in Ratel-FEM, using Machine Learning (ML) to determine various functional forms for micromorphic elastic free energy. The in-situ micromorphic filter will streamline data collection (minimize output), and parallelize the filter to run along with Ratel.
- Implement implicit MPM within Ratel and a gradient-enhanced damage formulation on top of the large deformation micromorphic elasto-plastic constitutive model in Tardigrade-MOOSE. This will provide a smeared failure modeling capability within the upscaled, continuum-scale micromorphic Tardigrade-MOOSE code.

University of Illinois at Urbana-Champaign

Center for Exascale-Enabled Scramjet Design

Accomplishments in FY23:

- Assessed predictive capability for the experimental supersonic combustor and identified weak points to direct future efforts.
- Used the Center's MIRGE approach to demonstrate a platform-portable multi-physics simulation using Python based drivers, loop-abstraction enabled lazy evaluation, and hardware-specific execution code generation.
- Integrated additional physics of carbon-fiber composite combustor wall material oxidation and degradation with supersonic combustion code.
- Designed variable geometry scramjet target to further assess predictive confidence.
- Co-designed physics-targeted and physics-integrated low-dimensional experiments and simulations for model integration, validation, and uncertainty quantification.

Planned Activities in FY24:

- Assess integrated predictive capability for a variable geometry scramjet with novel high-temperature composite combustor wall materials.
- Instrumentation and analysis to study and improve code performance, both its runtime efficiency and compile time efficiency including its impact on overall workflow.
- End-to-end UQ of the full prediction, building on a hierarchy of single- and coupled-physics validation simulations and their laboratory counterparts.
- Continue to train computer scientists, computational scientists and experimentalists broadly in predictive science, leveraging the unique environment of an integrated PSAAP center.

Stanford University

Integrated Simulations using Exascale Multiphysics Ensembles

Accomplishments in FY23:

- Developed a detailed verification plan aimed at both solution and code verification with both single- and multi-physics targets. Progress on unit test cases is combined with investigation of full system simulations by using a sequence of computational grids ranging from few thousand nodes up to 200M nodes.
- Successfully completed sensitivity analysis of full system simulations involving grid refinement, inlet conditions, sub grid models, and chemical mechanisms. In addition, the effect of geometrical fidelity has been studied by comparing non-reactive flow simulations with a fully unstructured grid solver.

- Enhanced the Continuous Integration (CI) framework by incorporating explicit dependencies between HTR solver and Legion repositories; increased the software coverage with unit modules and full regression tests.
- Introduced a new validation metric to assess the accuracy of the integrated simulations; the focus is the shock wave position and speed induced by the initial energy deposition. Comparisons between simulations and experiments are extremely promising.
- Demonstrated the first data-driven model of laser-based ignition consisting of a gaussian mixture variational auto-encoder (GMVAE). The model uses an embedded classifier to distinguish the successful ignition cases and uses both experimental and high-fidelity simulations as training data.

Planned Activities in FY24:

- Introduce performance monitoring and analysis within the Continuous Integration framework, and expand platform converge to include multiple CPU and GPU clusters at Stanford and LLNL extending current exploratory activities on AMD platforms.
- Carry out the first series of experimental tests with liquid propellent and characterize the conditions in the chambers and quantify the energy deposited and interactions with the liquid droplets. Construct the first ignition map.
- Complete the V&V studies corresponding to the multiphase flow modeling framework for multicomponent, compressible fluid. Carry out the first full system simulations mirroring the experimental tests with liquid propellant.
- Complete the verification and deployment of three low-fidelity models of ignition based respectively on grid coarsening, ignition kernel transport model, GMVAE. Carry out linear and non-linear multifidelity assessments and compare the prediction accuracy and cost.
- Complete the transition between HTR and HTR++ (an updated version that uses the Legion API C++ interface) and demonstrate both performance and portability on existing DOE platforms. HTR++ targets for the year are demonstrations of improved geometrical fidelity (multi block) and multiphase flow modeling capability (liquid propellents).

University of Texas at Austin

Exascale Predictive Simulation of Inductively Coupled Plasma Torches

Accomplishments in FY23:

- Resolved discrepancies between glow discharge predictions by developing an enhanced, six-species argon chemistry model and validating it against measurements in a new, advanced glow discharge facility which was developed to enable the use of a Langmuir probe and reduce uncertainty in voltage and current measurements.

- Resolved discrepancies between magnetic field predictions and measurements in the torch by deploying a custom-built B-dot probe, which was necessary because the original, commercially purchased probe was not properly designed or calibrated for the torch operating frequency.
- Equipped the torch simulator with new modeling capabilities, particularly additional lower-fidelity models aimed at reducing overall simulation time and used it to perform torch simulations at a range of power conditions.
- Completed deterministic and DSMC Boltzmann solvers with Coulombic collisions. The solvers extend the current state-of-the-art for low-temperature plasmas (Bolsig+) in that they support transient solves, arbitrary accuracy poloidal and azimuthal discretizations with and without Coulombic interactions. Integrated the Boltzmann solvers with PyKokkos and Parla, to support multiple GPUs using both Parla and MPI.
- Redesigned and reimplemented Parla scheduler in C++ to allow scalability with finer-grain tasks. Parla now supports scheduling of multi-GPU tasks. Developed CrossPy a multi-GPU array data structure to make programming easier on heterogeneous architectures.
- Developed, implemented, tested, and open sourced a new cycle detection algorithm on unsymmetric graphs, with application to sweep-based methods for radiation transport.

Planned Activities in FY24:

- Complete blow out experiments and simulations for the torch operating with argon feed gas, including resolving any validation issues remaining for the argon case after FY23.
- Extend the modeling, simulations, and measurements to the case of nitrogen feed gas, which will involve development, implementation, and validation of chemistry, transport, and radiation models for nitrogen as well as new experiments in both the torch and glow discharge facilities.
- Deploy a new experimental facility for probing shock-induced thermal non-equilibrium in nitrogen and air plasmas, using the ICP torch as a plasma source.
- Complete software and algorithmic integration of electron Boltzmann and torch codes. Investigate possible discrepancies between transient and quasi steady-state electron transport, the effect of spatial gradients, and plasma sheath boundary conditions.
- Develop and test a novel fast adaptive solver for radiation transport using sweep algorithms, and deploy it on heterogeneous architectures.
- Add new capabilities in the Boltzmann solvers to support spatial coupling, dynamic adaptivity in position and velocity, and algorithmic load balancing across heterogeneous nodes.

- Improve scheduler policies of Parla by building partial task graph information and using memorizing and performance monitoring.

University at Buffalo

Center for Exascale Simulation of Hybrid Rocket Motors

Accomplishments in FY23:

- Generated a detailed reaction mechanism for C16H34 paraffin (n-hexadecane) combustion using an automated method and validated its performance for ignition delay time predictions vs. literature data and prior detailed models from LLNL.
- Expanded a one-dimensional, micro-mechanical model for shear-driven droplet formation that accurately predicts both the final droplet volume and pinch-off time for a large range of inlet velocities, density and viscosity ratios, and surface tensions
- Demonstrated new near wall-modeling of turbulent flows with high levels of surface blowing.
- Developed scalable level set interface tracking / volume of fluid solver to explore Kelvin-Helmholtz instability wave growth.
- Created a pipeline for running an ensemble of ABLATE slab burner simulations and training different surrogate models (GPR/multiscale), which allows for forward propagation of the uncertainty in chemical kinetics and inlet conditions to the Quantities of Interest (QoIs: temperature, regression rate, radiative heat flux).

Planned Activities in FY24:

- Extend detailed reaction mechanism to provide a skeletal mechanism for paraffins up to C32H66 and incorporate formation pathways for soot precursors.
- Plan to use surrogate droplet model, combined VOF fluid model, to generate Lagrangian droplets that are essential for describing highly ablating fuels
- Incorporate new FGMs for solid fuel combustion into the ANN-based CHEMTAB framework to significantly reduce computational cost
- Perform intermediate scale sounding rocket simulations and collect validation data for the hybrid paraffin wax - O2 system
- Develop integrated UQ in final QoI (e.g., regression rates) for the full modeling/simulation pipeline including flows, atomization processes and chemistry.

Massachusetts Institute of Technology

Center for the Exascale Simulation of Material Interfaces in Extreme Environments

Accomplishments in FY23:

- Completed a first full-system simulation of the oxidation of a Hf surface and propagation of the Hf/HfO2 boundary, by integrating automated DFT data generation, fitting of machine learning (ML) interatomic potentials, molecular dynamics simulations, and a new off-lattice kinetic Monte Carlo scheme.
- Developed a new "proper orthogonal descriptor" (POD) interatomic potential, which produces compressed invariant descriptors of atomic environments via reduced basis techniques and demonstrated substantial performance gains over competing ML potentials.
- Created new composable workflows in Julia for interatomic potential construction, fitting, and UQ.
- Developed new backends for the OpenCilk task-parallel compiler and Kitsune to target NVIDIA and AMD GPUs.

Planned Activities in FY24:

- Refine and deploy kernel-based active learning strategies to create improved ab initio datasets for the surface-to-bulk oxidation of hafnium.
- Develop a new "+U"-inspired Hubbard model functional that will allow DFT calculations to better capture gas-phase quantum chemistry relevant to surface oxidation.
- Develop efficient, accurate, and robust reduced order models for hypersonic flow models with non-equilibrium chemistry.
- Benchmark and improve the scalability of POD potential implementations in LAMMPS.

University of Maryland

Solution-Verification, Grid-Adaption and Uncertainty Quantification for Chaotic Turbulent Flow Problems

Accomplishments in FY23:

- Computed approximate sensitivities for turbulent flows using three completely different approaches, all developed within the Center.
- Extended the error estimation and grid-adaptation method for turbulence simulations to high-order finite-element-type codes.
- Organized a collaborative workshop focused on solution-verification and convergence in turbulence simulations.

Planned Activities in FY24:

- Implement the most rigorous sensitivity estimation method in the Kokkos-based Center solver and apply method at large Reynolds numbers using Lassen or similar.

- Perform a formal comparison of all the sensitivity methods in terms of cost, accuracy, strengths, and weaknesses.
- Develop algorithm for the automatic creation of block-adapted curvilinear meshes in complex geometries based on error estimation inputs only.
- Continue development of sensitivity algorithms for higher accuracy and lower computational cost.

University of New Mexico

Center for Understandable, Performant Exascale Communication

Accomplishments in FY23:

- Designed and implemented a benchmark suite for approaches to GPU-triggered regular communication on modern accelerator-based hardware, and assessed the performance characteristics of different communication triggering approaches on current vendor hardware and software.
- Assessed and quantified the highly variable nature of communication demands in the LANL XRage application and created a communication benchmark capable of recreating the and other irregular communication patterns to support research on improved communication primitives for production DOE applications.
- Created a distributed high-order fluid interface model implementation that uses brute-force calculation of far-field forces and designed a cut-off based approach the optimizing far-field force calculations that can be used to benchmark and guide optimization of global-communication strategies in HPC applications.
- Designed and prototyped new MPI abstractions to support global optimization of neighbor data exchanges and integrated and tested these abstractions in the HYPER solver library.

Planned Activities in FY24:

- Create and begin curation of an MPI communication benchmark suite whose characteristics are informed by assessment of communication characteristics in DOE mission applications.
- Model and measure the impact of improved GPU communication latency on strong scaling of key DOE benchmarks and applications
- Design initial low-level abstractions and APIs that decompose current MPI communication primitives and support the creation of specialized high-level communication primitives to secure performance-portability at application-level abstractions.
- Release of initial high communication system design and implementation course materials, including example hackathon materials, and lecture contents.

Oregon State University

Center for Exascale Monte Carlo Neutron Transport

Accomplishments in FY23:

- By the end of FY23, all graduate students directly and indirectly affiliated with CEMeNT will have participated in a Tri-Lab internship.
- The Quasi-Monte Carlo approach achieved the capability to solve 3D multigroup neutron transport problems within MC/DC.
- Python bindings have been developed for Harmonize, enabling pure-Python implementation of on-GPU asyncronous programs, for integration with MC/DC.
- Integrated iterative Quasi-Monte Carlo, time-dependent weight windows, and variance-deconvolution UQ into MC/DC, which has ~10 users/developers on the CEMeNT team. Additionally, MC/DC has been successfully used to explore a new sensitivity calculation technique - the Derivative Source Method - demonstrating its merit as an exploratory tool for rapid prototyping of methods.
- Drafted and delivered to the Tri-Lab Sponsor Team (TST) a document that describes both measures of success and approaches to compute them, for staged variations of the small modular reactor (SMR) challenge problem. Generated baseline data for the formative versions of the challenge problem with both MC/DC and Shift.
- Eight peer-reviewed full conference papers have been accepted for presentation at the upcoming American Nuclear Society Topical Meeting on Mathematics and Computations in Niagara Falls, Canada. Six journal articles are also in preparation, to be submitted by the end of FY23.

Planned Activities in FY24:

- Focus strongly on consistent measurement of success metrics as various algorithmic improvements mature, assessing which combinations of approaches work well for the stages of the challenge problem.
- Proceed with development and implementation of advanced weight-windows algorithms. Development and implementation of domain decomposition algorithms. Proceed with development and analysis of multilevel hybrid MC methods.
- Continue to work toward continuous energy collision physics, incorporation of geometric visualization, and enable machine learning tuning of hardware parameters for MC/DC execution.
- Flesh out performance of embedded UQ and sensitivity approaches (variance deconvolution, Derivative Source Method) in MC/DC to propagate cross-section uncertainties through calculation of local and global quantities of interest.
- Begin work on compressive sensing algorithms, exploring the effect of memory reduction on calculation effectiveness.

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**	**FY23**	**FY24**	**FY25**	**FY26**	**FY27**	**Future**	**Contact**
LLNL	Commodity Technology System (CTS)-2 System*	Capital Equipment, MIE	N/A	70,000	10,000	20,000	20,000	20,000	0	0	0	Matthew L. Leininger 925-4224110 leininger4@llnl.gov
LLNL	Unclassified ATS-4-like System, LLNL*	Capital Equipment, MIE	N/A	19,700	0	200	3,900	3,900	3,900	3,900	3,900	Ned Bass 925-422-9389 bass6@llnl.gov
LLNL	AT System – El Capitan*	Capital Equipment, MIE	N/A	600,000	287,667	158,654	95,793	57 ,886	0	0	0	Bronis de Supinski 925-422-1062 desupinski1@llnl.gov
LLNL	Bldg 453 Sierra Retirement	Minor Construction	N/A	10,000	0	0	0	10,000	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 453 Tuolumne Site Prep	Minor Construction	N/A	6,000	0	3,000	3,000	0	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 451 Power and Cooling Improvements/Modifications	Minor Construction	N/A	7,500	0	0	0	750	6,750	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 453 Power and Cooling Improvements/Modifications	Minor Construction	N/A	7,500	0	0	0	750	6,750	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	AT System – ATS-6	Capital Equipment, MIE	N/A	250,000	0	0	0	0	2,000	8,000	240,000	Bronis de Supinski 925-422-1062 desupinski1@llnl.gov
LANL	Crossroads: Acquisition of Crossroads (ATS-3) system	Capital Equipment, MIE	N/A	115,000	95,000	8,000	6,000	6,000	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY23**	**FY24**	**FY25**	**FY26**	**FY27**	**Future**	**Contact**
LANL	ATS-5: Acquisition of ATS-5 system	Capital Equipment, MIE	N/A	250,000	0	0	2,000	15,000	100,000	95,500	37,500	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	Future Supercomputing Infrastructure	Line Item	N/A	350,000	57	1,500	3,000	30,000	60,000	60,000	196,000	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	Crossroads Installation Project	Minor Construction	N/A	12,000	12,000	0	0	0	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	SCC Electrical Upgrade	Minor Construction	N/A	18,400	18,400	0	0	0	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	ATS-5 Cooling Installation	Minor Construction	N/A	12,000	0	0	0	500	11,500	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	ATS-5 Electrical Installation	Minor Construction	N/A	12,000	0	0	0	500	11,500	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	CTS-3 Installation	Minor Construction	N/A	5,000	0	0	0	0	0	0	5,000	Jason Hick 505-667-4477 jhick@lanl.gov
SNL	Additional CT Systems	Capital Equipment, MIE	N/A	20,000			5,000	5,000	5,000	5,000		Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725E 4MW Power Expansion	Minor Construction	N/A	5,600	500	3,500	1,600					Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725E Cooling Capacity Expansion	Minor Construction	N/A	9,600		5,400	4,200					Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725E Long Term Power and Cooling	Minor Construction	N/A	20,000							20,000	Tom Klitsner 505-844-1901 tklitsn@sandia.gov

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY23**	**FY24**	**FY25**	**FY26**	**FY27**	**Future**	**Contact**
SNL	725E Expansion	Minor Construction	N/A	5,000					5,000			Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	ATS-ART System (El Capitan Nodes)	Capital Equipment, MIE	N/A	16,400			16,400					Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	ATS-ART System (OHPC Crossroads Nodes)	Capital Equipment	N/A	2,500			2,500					Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	Build & Test System	Capital Equipment	N/A	7,000			3,000			4,000		Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	Digital Engineering Integration Laboratory (DEIL)	Minor Construction	N/A	25,000				5,000	20,000			Tom Klitsner 505-844-1901 tklitsn@sandia.gov

Appendix I: FY24–FY28 Program Targets

This appendix lists the targets for each subprogram for the FY24–FY28 period. 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. Builds upon the FY22 Assess Lifetimes and Mitigate Aging SCDS pegpost and supports the following pegposts: 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 Materials for Future Reentry Environments, FY26 Combined Threat Environment Simulation.

IC-4. Improved responsiveness and analysis capabilities. The target date is FY23. Builds upon the FY21 Demonstrate Key Responsive Technologies SCDS pegpost and supports the FY32 First Production Unit in 5_Years pegpost.

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 FY21 Hostile Survivability Baseline SCDS 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. Support responsive deterrent capabilities 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 evidence basis will be developed for production applications. Machine Learning informed workflows will be utilized as appropriate. The target is FY27.

V&V-2. Establish and maintain test suites for rigorous Verification, Validation and Uncertainty Quantification of relevant engineering and physics phenomena 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 FY24.

V&V-3. Develop Engineering and Physics Common Model Frameworks with integrated V&V/UQ and Sensitivity Analysis. 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 Combined Threat Environments Simulation SCDS pegpost. Develop integration between common model frameworks and experimental data repositories to allow modern workflows to seamlessly integrate experimental data into deterrent simulation capabilities and to ensure distinction between data used for calibration/training and that used for validation. The target is FY26.

V&V-4. Incorporate ML methods that enable agile predictions suitable for stockpile applications. The target is FY28, and ongoing work will support prototype-to-design work in Modern Materials & Manufacturing and Future Deterrent SCDS strands.

V&V-5. Extend the V&V/UQ frameworks to support the 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.

Computational Systems and Software Environment

CSSE-1. Platform Acquisition and Deployment.

a. Deploy and accept Crossroads, including Center of Excellence (COE) activities. The target date for full-system acceptance is Q1FY24 and classified production service in Q2FY24.

b. Procure and deploy El Capitan, including NRE activities. The target date for system acceptance is Q2FY24 and classified production service in Q4FY24.

c. Deploy CTS-2 platform across tri-labs and transition to production service. The target dates are FY22-FY24.

d. Procure and deploy Vanguard-II Advanced Architecture Prototype System. The target date is FY25.

e. Initiate planning and initial acquisition steps for ATS-5 in FY24 for a full-system deployment at LANL. The target system acceptance date is FY27.

CSSE-2. Software Environment Deployment.

a. Develop and deploy software and programming environments for Crossroads. The target date is FY24.

b. Develop and deploy software and programming environments for El Capitan. The target date is FY24.

c. Deploy the Vanguard-II system with a functioning software stack composed of ATSE, TOSS, and FOUS CCE tools and packages. The target date is FY25.

d. Develop and deploy software and programming environments for CTS-2. The target date is FY22 - FY25.

e. Develop and deploy software and programming environments for ATS-5. The target date is FY27.

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

a. Demonstrate applications of interest on BML-funded hardware. The target date is FY22 – FY25.

b. Demonstrate applications of interest on BML-funded hardware, including initial evaluation of neuromorphic Intel Loihi testbed. The target date is FY24.

c. Demonstrate hybrid quantum-classical algorithm relevant to mission science using latest quantum device designs. The target date is FY27-FY29.

d. Evaluate and demonstrate the application of machine learning models for ASC mission relevant problems. The target date is FY23 – FY26.

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 FY24.

b. The Crossroads system, Phase 1, will be delivered, installed, and deployed in FY23, with Crossroads full system deployed in FY23. 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-II, exploratory technology systems will be deployed in FY22 and FY23, with the at-scale production system scheduled for deployment in FY25.

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. FOUS-4. Installation and operation of ATS-5 and ATS-6.

a) The planning phase of ATS-5 will begin in FY23. The target date for deployment of ATS-5 is FY27 with transition to classified service in FY28.

b) The planning phase of ATS-6 will begin in FY25. The target date for deployment of ATS-6 is FY29 with transition to classified service in FY30.

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 FY22. Planning for the CTS-2 deployments began in FY19. The initial installation, deployment, and operation of CTS-2 started in FY22. The target date for continued CCE software development and support is FY21–24.

FOUS-6. Design and build the Digital Engineering & Integration Laboratory (DEIL). This facility will enable ASC and NA-12 staff integration for high priority stewardship programs (e.g., W93 and W87R). The current target is to have the facility design phase completed in FY25 and all construction completed in FY26.

Appendix J: Codes

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

National Nuclear Security Administration

Appendix K: ASC Milestone Reporting Requirements

Introduction:

This document specifies reporting requirements for milestone to be considered complete by the Advanced Simulation and Computing (ASC) program. These requirements will ensure consistency between the various elements of the program and allow the program to defend the high quality of the NNSA laboratories' work, while allowing considerable flexibility for how projects and research are managed and executed by local lab managers and principal investigators.

General Requirements:

1. Milestone review dates (including any mid-year and final reviews) and review committee membership shall be shared no later than one quarter in advance with the cognizant federal program manager (FPM) who is responsible for the portfolio(s) involved with the milestone.
2. Milestone review committees shall contain at least one member from a different NNSA ASC program unless an alternative agreement is made with the cognizant FPM ahead of time.
3. All documentation must be received 5 business days before the milestone due date. Should some final documents not be available for the FPM to use for Milestone Reporting Tool (MRT) statusing, the appropriate lab subprogram executive shall work with the cognizant FPM to determine and comply with an interim reporting process, at least five (5) business days before the milestone due date.

Required Documentation:

Evidence of milestone completion sent to the federal ASC program shall consist of the following documentation:

1. Certification Memorandum;
2. Material provided to the review committee as presentations or papers; and
3. Milestone Report or journal article.

All evidence shall be emailed by the appropriate lab program element executive, or delegate, to the federal ASC program team, which includes the ASC program director, the FPM(s) responsible for the portfolio involved in the milestones, supporting laboratory detailees, and the support service contractor(s) responsible for archiving the milestone documents at NNSA HQ. Classified documents shall be emailed to the addressees on the appropriate system and an unclassified email shall be sent as an alert that the milestone documents have been sent on the high side.

NOTE: If a milestone involves the completion of a construction project (minor construction and/or line-item construction), the FOUS FPM and lab subprogram element executive may agree to waive, in writing, the above-referenced evidence of milestone completion. If waived, the laboratory shall provide a completed construction memo and any paperwork otherwise required by the local site office, laboratory management and, if necessary, any evidence required by the Office of Infrastructure (NA-90).

Documentation Descriptions:

1. **Certification Memorandum**: Signed by the review panel chair or other credible sources addressed to the cognizant lab program manager. This memorandum shall include the following information:
   - Review panel committee members' names and affiliations, date(s) of review, milestone title with MRT ID number - or alternately the name of the certifier, date of certification, milestone title with MRT ID number;
   - Descriptions of the closure criteria and evaluation of how the work passed/failed/exceeded milestone scope; and
   - Technical recommendations for future work at the discretion of the review committee.

2. **Milestone Report**: A description of the work and results prepared by the milestone team.
   - *Executive Summary*—A self-contained synopsis of the report: a summary of what was completed, the closure criteria presented to the committee, and the results.
   - *Introduction*: A summary of the milestone, context of what problem was addressing (background), and context for why the work plan was developed.
     - What was the purpose/objective of the work?
     - Why was the work important in a broader context to the program/mission?
   - *Methods*: A description of the materials, procedures, computational platform, computation time, etc.
   - *Results/Discussion*: An interpretation of the results/findings and what they mean for the future. If the milestone is used for hardware demonstrations, or, system or application readiness demonstrations, then performance results supporting the conclusion must be included.
   - *Conclusion*: A summary of the results as they relate to the context set in the introduction and any plan for future work or follow-up milestones.
   - *Bibliography/References*: A citation of the material provided to the review committee and any other useful references for the record.