Advanced Simulation and Computing

FY25 IMPLEMENTATION PLAN

Version 0

August 5, 2024

NA-ASC-127R-24-Rev.0-IP

LLNL-TR-867762

Version History

Version #	Implemented By	Revision Date	Approved By	Approval Date	Reason
0	Rob Neely, Scott Doebling, Jennifer Gaudioso	Initial	Thuc Hoang	8/05/24	Initial release

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

Contents

I. Overview

The DOE National Nuclear Security Administration (NNSA) Stockpile Stewardship Program (SSP) is an integrated technical program for maintaining the safety, 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. The ASCI program was essential to the successful demonstration of the SSP, providing critical nuclear weapons simulation and modeling capabilities. ASCI became the ASC program in fiscal year (FY) 2005 and the mission remains essentially the same to today: to provide the simulation and computational capabilities that underpin the NNSA's ability to maintain a safe, secure, effective nuclear weapon stockpile, without returning to underground nuclear testing.

The capabilities that the ASC program provides play a vital role in the nuclear security enterprise and are necessary for fulfilling the stockpile stewardship and life extension programs assigned to the 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 FY25. The Initial IP, effective August 12, 2024, 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.

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	FY2025 Target	Endpoint Target
Advanced Simulation and Computing Program	Nuclear Weapons Simulation Capability Progress*	79%	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 FY25 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.

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 FY2025 Budget Request**	**FY2025 Enacted/Full Year<Budget or CR>**	**Difference between Request versus Enacted**
ASC Operations	$879.500M	$xxxM	$xxxM
Academic Programs (for PSAAP & CSGF)	$23.261M	$xxxM	$xxxM

National Nuclear Security Administration

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 FY25. 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 J.

Table V-1B. ASC Level 2 Milestones for FY25

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY25–FY29 Target
IC	TBD	Towards Exascale Capable Adaptive Mesh Refinement (AMR) Methodologies and Calculations	9/30/25	LLNL	TBD
IC	TBD	Further Applications of Configuration Transfer Enabling New Capability and Enhanced Verification	9/30/25	LLNL	IC-2, IC-4
IC	TBD	Open-source Software Advancements Supporting Integrated Code Capabilities on El Capitan	6/01/25	LLNL	IC-1
IC	TBD	Container-based workflows	9/30/25	LANL	IC-1, CSSE-2.a
IC	TBD	Demonstration of a Fast Sweeping Program Burn Capability	9/30/25	LANL	IC-2
IC	TBD	Code preparations towards El Capitan	9/30/25	LANL	IC-1
IC	TBD	Enhancing Developer Productivity	TBD	SNL	TBD
IC	TBD	MC Rad Transport Consolidation	TBD	SNL	TBD
PEM	TBD	Quantification of Pu EOS Uncertainty Using NIF and Z Experiments	9/30/25	LLNL	PEM-5, V&V-1

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY25–FY29 Target
PEM	TBD	Use of ab initio Molecular Dynamics with Relevance to Manufacturing and Aging of Plutonium	9/30/25	LLNL	PEM-2, PEM-5
PEM	TBD	Analysis of Alternative Insensitive High Explosive Options	6/30/25	LLNL	PEM-2
PEM	TBD	Uncertainty Quantification of High Explosive Models	9/30/25	LANL	PEM-2,3,4,5
PEM	TBD	Characterizing, propagating, and archiving uncertainty estimates associated with strength model parameters	9/30/25	LANL	PEM-2, 3, 4, 5
PEM	TBD	Electromagnetic EM Consolidation	9/15/25	SNL	PEM-2, PEM-4
PEM	TBD	Production Simulation Workflow for Distortion Compensation in Metal Additive Manufacturing	9/30/25	SNL	PEM-2
V&V/ PEM/ IC	TBD	Assess Direct-Processed GNDS/GIDI as Default Nuclear Data	6/01/25	LLNL	V&V-5
V&V	TBD	Verification and Validation of 3D Capabilities	9/30/25	LANL	V&V-1, V&V-2
V&V	TBD	ROM-enabled Multifidelity UQ	9/30/25	SNL	V&V-1, V&V-4
V&V	TBD	Verify Thermal-Structural Coupling of System Temperatures and Ablated Shapes for Reentry Simulations	9/30/25	SNL	V&V-1 V&V-5
CSSE	TBD	Tri-lab AI/ML	9/30/25	LLNL LANL SNL	TBD
CSSE/ FOUS	TBD	Converged HPC Center-Wide Data Analytics	8/31/25	LLNL	FOUS-1, FOUS-5, CSSE-1.a
CSSE	TBD	El Capitan System Readiness	3/30/25	LLNL	TBD
CSSE	TBD	Integrate DSI and ASC capabilities on HPC resources	9/30/25	LANL	CSSE-2.c, CSSE-2.d
CSSE/ IC/ FOUS	TBD	Computing-as-a-Service for Digital Engineering	TBD	SNL	TBD
CSSE	TBD	Computing Platform Engineering for Digital Integration	9/30/25	SNL	TBD
FOUS	TBD	SCC Electrical Upgrade commissioning complete	09/30/25	LANL	FOUS-3

Sub-Program	ID#	Milestone Title	Complete Date	Site	FY25–FY29 Target
FOUS	TBD	El Dorado Standup and Deployment	9/30/2025	SNL	FOUS-1

VI. Reporting Requirements

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

- **Quarterly Program Reviews**. M&O contractors report Level 2 milestone status to Headquarters (HQ) using the Milestone Reporting Tool. In preparation for each quarterly review, each site and Federal Program Manager will assess the status of each milestone by providing a score (Green, Yellow, Red, Blue, or Black). In addition, supporting details for the assessment of each milestone will be reported and any programmatic risk will be identified.
- **Monthly Financial Reporting**. Monthly cost/financial expenditure data will be reported by the 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 Tri-Lab Procurement Engagement Model* (March 2023)
- *ASC Artificial Intelligence for Nuclear Deterrence (AI4ND) Strategy* (January 2024)
- 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.

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

Scott Doebling, LANL, 505-667-6950, doebling@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 Jimmy Fung, LANL, 505-667-1407, fung@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@sandia.gov
1.2.3.1.1	Engineering and Physics Integrated Codes	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Jimmy Fung, LANL, 505-667-1407, fung@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@sandia.gov
1.2.3.1.2	Specialized Codes and Libraries	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Jimmy Fung, LANL, 505-667-1407, fung@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@sandia.gov
1.2.3.1.3	Applications and Algorithms Research	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Jimmy Fung, LANL, 505-667-1407, fung@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@sandia.gov
1.2.3.1.4	Applications Research for Next-Generation Platforms	Teresa Bailey, LLNL, 925-424-6700, bailey42@llnl.gov Jimmy Fung, LANL, 505-667-1407, fung@lanl.gov Erik Strack, SNL, 505-284-9270, oestrac@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 Darby Jon Luscher, LANL, 505-665-2848, djl@lanl.gov Cole Yarrington, SNL, 505-845-0974, cdyarri@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 Darby Jon Luscher, LANL, 505-665-2848, djl@lanl.gov

WBS	Title	Contact
		Cole Yarrington, SNL, 505-845-0974, cdyarri@sandia.gov
1.2.3.2.2	Transport, Plasmas, Atomic, Nuclear	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Darby Jon Luscher, LANL, 505-665-2848, djl@lanl.gov Cole Yarrington, SNL, 505-845-0974, cdyarri@sandia.gov
1.2.3.2.3	Engineering Science, Environments, and Response	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Darby Jon Luscher, LANL, 505-665-2848, djl@lanl.gov Cole Yarrington, SNL, 505-845-0974, cdyarri@sandia.gov
1.2.3.2.4	Integrated Modeling and Applications	Marina Bastea, LLNL, 925-424-2803, bastea1@llnl.gov Darby Jon Luscher, LANL, 505-665-2848, djl@lanl.gov Cole Yarrington, SNL, 505-845-0974, cdyarri@sandia.gov
1.2.3.3	Verification and Validation	David Etim, NNSA, 202-586-8081, David.Etim@nnsa.doe.gov Karlene Maskaly, LLNL, 925-422-5526, maskaly2@llnl.gov Gowri Srinivasan, LANL, 505-665-3157, gowri@lanl.gov Tricia Gharagozloo, SNL, 925-337-7866, peghara@sandia.gov
1.2.3.3.1	Verification and Validation Methods	Karlene Maskaly, LLNL, 925-422-5526, maskaly2@llnl.gov Gowri Srinivasan, LANL, 505-665-3157, gowri@lanl.gov Tricia Gharagozloo, SNL, 925-337-7866, peghara@sandia.gov

WBS	Title	Contact
1.2.3.3.3	Data Validation, Archiving, Software Quality Assurance, and Training	Karlene Maskaly, LLNL, 925-422-5526, maskaly2@llnl.gov Gowri Srinivasan, LANL, 505-665-3157, gowri@lanl.gov Tricia Gharagozloo, SNL, 925-337-7866, peghara@sandia.gov
1.2.3.4	Capabilities for Nuclear Intelligence	Paul Miller, LLNL, 925- 423-6455, miller3@llnl.gov Kris Eriksen, LANL, 505-667-2295, keriksen@lanl.gov Tom Hafenrichter, SNL, 505-845-7038, thafenr@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 John Patchett, LANL, 505-665-1110, patchett@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 John Patchett, LANL, 505-665-1110, patchett@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 John Patchett, LANL, 505-665-1110, patchett@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 John Patchett, LANL, 505-665-1110, patchett@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 John Patchett, LANL, 505-665-1110, patchett@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.7	Next-Generation Computing Technologies	Todd Gamblin, LLNL, 925-422-9319 gamblin2@llnl.gov John Patchett, LANL, 505-665-1110, patchett@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.8	Advanced Memory Technologies and Vendor R&D	Matt Leininger, LLNL, 925-422-4110 leininger4@llnl.gov John Patchett, LANL, 505-665-1110, patchett@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov

WBS	Title	Contact
1.2.3.5.9	Advanced Architecture Prototype Systems Product	Matt Legendre, LLNL, 925-422-6525 legendre1@llnl.gov John Patchett, LANL, 505-665-1110, patchett@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.10	Data Science and Artificial Intelligence	Katie Lewis, LLNL, 925-423-8593 lewis66@llnl.gov John Patchett, LANL, 505-665-1110, patchett@lanl.gov Rob Hoekstra, SNL, 505-844-7627, rjhoeks@sandia.gov
1.2.3.5.11	Compilers, Programming Models and Runtimes	John Gyllenhaal, LLNL, 925-424-5485 gyllenhaal1@llnl.gov John Patchett, LANL, 505-665-1110, patchett@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
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

WBS	Title	Contact
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
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

PSAAP III Points of Contact	
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

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 *FY 2024 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 six subprograms:

The **Integrated Codes** (IC) subprogram develops the mathematical descriptions of the physical processes of nuclear weapon systems and functions. Combined with weapon-specific input data created by the nuclear weapons designers and engineers, this allows detailed simulations of nuclear weapons performance assessment without the need for underground nuclear testing. The IC subprogram funds the critical skills needed to develop, maintain, and advance the capabilities of the large-scale integrated simulation codes needed for the following SSP and Stockpile Management and Production Modernization activities: annual assessment; LEP design, qualification, and certification; SFI resolution; and safety assessments to support transportation and dismantlement. In addition, these capabilities are necessary for a host of related requirements such as nuclear 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' Office of Experimental Science (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 2024 Stockpile Stewardship and Management Plan*, Report to Congress, November 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 **Capabilities for Nuclear Intelligence** (CNI) subprogram is responsible for advancing and adapting SRT&E capabilities developed for the weapons program to serve the needs of the intelligence community, thereby enabling sound technically based performance, safety, and vulnerability assessments of foreign nuclear weapon activities. CNI focuses on activities related to non-stockpile weapons training, high explosive knowledge development, weapon modeling advancements and computational platforms, weaponization studies, and experimental capabilities and assessments.

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. It 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.

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 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 six subprograms (IC, PEM, CNI, V&V, CSSE, and FOUS) all contribute to a cohesive set of program deliverables. Highlights of the FY25 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 enhance verification suite 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.
- Demonstrate improved and validated nuclear data from machine learning techniques integrated into material response and plasma models.
- 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.
- Transition into classified production service 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 additional 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.
- Complete transition of next-generation codes and computing technologies from the previously funded Advanced Technology and Development Mitigation (ATDM) program element 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.
- Continue investment in AI model development and infrastructure, including testbeds and site-wide digital infrastructure.
- Maintain mentoring program for early career staff.

The drivers of the ASC program that serve as programmatic anchors for the six subprograms requiring the FY25 budget include the following:

- NNSA's FY 2025 Presidential Budget Request outlines the Office of Defense Programs' Weapons Activities scope and budget requirements. Highlighted stockpile management goals include the completion of life-extending the B61-12 and completing the W88 Alteration (Alt) 370 by FY2025, execution of the B61-13 Phase 6.3 and 6.4 (Development/Production Engineering), continuation of Phase 6.4 (Production Engineering) activities for the W80-4 LEP, continuation of Phase 6.3 (Development Engineering) activities for the W87-1 Modification Program, and completion of Phase 2 (Feasibility Study and Design Options) for the W93 Program. The W93 will also commence Phase 2A (Feasibility and Cost Study).
- 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 ASC integrated design codes (IDCs) will open design options for subsystem components for future LEPs.
- Investments in AI infrastructure, model development, and testbeds will support the Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.

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 to minimize the disruptive technological impacts on the ASC IDCs.

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 8/02/24

Concurrence:

Rob Neely,
LLNL ASC Executive *Concurrence received* 7/30/24

Scott Doebling,
LANL ASC Executive *Concurrence received* 8/02/24

Jennifer Gaudioso,
SNL ASC Executive *Concurrence received* 7/30/24

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
ASF	Application-System Fusion (aka AppSysFusion)
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	Data Direct 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
TLPM	Tri-lab Procurement Model
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

National Nuclear Security Administration

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: Capabilities for Nuclear Intelligence (WBS 1.2.3.4)

Note: The content for the CNI subprogram is available upon request from the ASC Program Office.

National Nuclear Security Administration

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

The mission of this national subprogram is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of the NNSA. This subprogram strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms. Along with 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, including evaluating mission use cases for Quantum computing, artificial intelligence technologies and future customized compute packages or chiplets.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2023, and through quarter 3, fiscal year 2024, are reflected below for the CSSE subprogram.

- Sited all El Capitan hardware. Accepted the El Capitan Early Delivery and Intermediate Delivery Systems. Developed and deployed system software including TOSS-4, Flux, Rabbit software, and TCE-2. (LLNL)
- Implemented next-generation quantum simulation hardware testbed platform by installing superconducting quantum processor and supporting 5 connected qubits. Calibration routines and gates were made available to users. (LLNL)
- Continued collaboration with AWS on Spack, Flux, and containerized versions of MARBL and ALE3D. Began discussions on Google Distributed Cloud Air-gapped system. Ran successful procurement of Oxide on-premises cloud system for Tri Labs. (LLNL)
- Successful procurement of Oxide on-premises cloud system for tri lab users. (LLNL)
- Worked with AMD Flang team to successfully compile LLNL's Goulash GPU stress tests, which yielded unexpectedly strong Fortran OpenMP GPU performance at 82% of C++ equivalent. (LLNL)
- Released TOSS 5.0, based on RHEL 9.4, to support forward-looking systems such as testbeds, as well as El Capitan. (LLNL)

- Delivered unclassified ASC testbed (called Jarvis) in the Strategic Computing Complex (SCC) environment hosting both SambaNova's standalone DataScale system and SambaNova Suite system. (LANL)
- Developed an unsupervised-learning approach to analysis of laser-induced breakdown spectroscopy (LIBS) data used to study metal alloy material composition. This approach augments the existing supervised-learning approach to provide a robust suite of processes. (LANL)
- Completed initial simulation and modeling of performance improvements using advanced memory technologies using a combination of existing hardware and simulation capabilities. Codesigned prototype compiler changes and software infrastructure to support these technologies. (LANL)
- Released Data Science Infrastructure (DSI) v1.0 with read, write, query/find, and data movement capabilities to improve the I/O component of simulation workflow, data creation, and curation. (LANL)
- Released SST version 14.0. Improved the scalability and performance of SST core to allow increased scale of high-fidelity simulations. Built out initial infrastructure in SST for embedding dataflow graphs onto grid-based spatial architectures. Organized and coordinated the first SST user group with over 45 participants. (SNL)
- Received, announced, and initial deployment of world's largest neuromorphic system: Intel HalaPoint 1 billion neuron system as a part of a multi-year collaboration with Intel. (SNL)
- SNL placed a contract with Cerebras to train a 70B parameter model of open, but ASC-relevant data sources.
- Purchased and deployment of four wafers of a Cerebras CS-3 system for installation at Sandia. (SNL)
- Prototyped with NextSilicon an initial set of capabilities in a customized Kokkos backend to allow better scalability and performance on Gen-2 hardware. (SNL)
- Two significant Kokkos Ecosystem releases (4.2, 4.3) provided increased AMD GPU support in preparation for El Capitan (4.4 scheduled for summer). (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Tri-lab AI/ML		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 9/30/25		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL, SNL, LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target [ID-#]		
Description: Tri-lab Programming Environments Development		
Completion Criteria: xxx		
Customer: xxx		
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. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: xxx		

Milestone (ID#TBD): El Capitan System Readiness		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 09/31/25		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target CSSE-1.a, CSSE-2.a		
Description: : This milestone will certify that El Capitan is ready for integration into the classified computing environment. System hardware and software will be benchmarked and tested for correctness and reliability. Plans for system support and accessibility will be documented and communicated to the tri-lab community. At least three unclassified ASC applications will be demonstrated on the platform.		
Completion Criteria: Hardware deliveries from vendor to site are complete and on-site installation of the system by the vendor is substantially complete to the extent that is contractually required; contractual requirements for formal hardware acceptance have been substantially completed; system software needed for basic operation of the system is delivered, tested, and demonstrated to be operational; vendor has completed onsite benchmark and application performance testing and demonstration; and the system is ready to begin onsite integration into local computing environment. A tri-lab review will be conducted and the milestone results documented according to ASC milestone requirements.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: CSSE, FOUS, IC, PEM, V&V, SNL and LANL		

Milestone (ID#TBD): Integrate Data Science Infrastructure and ASC capabilities on HPC resources		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2025		
ASC WBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target CSSE-2.c, CSSE-2.d		
Description: LANL ASC simulation codes generate large datasets including ensembles of data. The Data Science Infrastructure (DSI) project focuses on data-driven approaches to make data more readily available to programmatic scientific users. DSI workflows leverage metadata stored in data-agnostic databases, supported by an API to simplify searching and accessing data across simulation runs, filesystems and environments. In this milestone, we will develop and integrate DSI capabilities with available ASC/HPC capabilities such as CONDUIT and GUFI, demonstrating across HPC resources targeted by ASC clients.		
Completion Criteria: 1. Demonstrate integration(s) on relevant HPC resources. 2. Prepare a report and presentation on new integrations. 3. Provide comprehensive API documentation and user guide.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: CSSE		

Milestone (ID#TBD): Computing-as-a-Service for Digital Engineering		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: XX/XX/XX		
ASC WBS Subprogram: CSSE, IC, FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target [ID-#]		
Description: Deploy CaaS capability and infrastructure in support of Accelerated Digital Engineering requirements.		
Completion Criteria: xxx		
Customer: xxx		
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. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: CSSE, FOUS, IC		

Milestone (ID#TBD): Computing Platform Engineering for Digital Integration		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 09/30/2025		
ASC WBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC IC, FOUS		
Program Plan Target: Appendix J, target CSSE		
Description: Sandia is launching a new initiative focused on platform engineering. This is a cross program effort that ties into the Digital Engineering Initiative and is intended to deliver a unified computing environment which supports model digital engineering workflows and integration with all sources of data whether it be experimental or modsim. This milestone team will deliver strategic and implementation plans for this activity which demonstrate the cross program integration as well as the requirements, gaps and dependencies on the broader computing and networking infrastructure across the weapons complex which will allow coordination and data sharing with the other labs and the production sites.		
Completion Criteria: 1. Deliver strategic and implementation plans for the CPE initiative. 2. Elucidate requirements and gaps that are beyond the span of control for the ASC program requiring cross-program coordination.		
Customer: ND		
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: SNL Enterprise Computing, CEE LAN		

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 FY24:

- Extensively tested and resolved issues on production CT systems.
- Led and managed the tri-lab CTS-2 contract and FY24 system procurements, system delivery schedules, and supply chain issues. (Appendix J, target CSSE-1.b)
- Fielded deliveries of FY24 CTS-2 production platforms. (Appendix J, target CSSE-1.b)
- Procured additional CTS-2 platforms. (Appendix J, target CSSE-1.b)

Planned Activities in FY25:

- Field deliveries of CTS-2 production platforms. (Appendix J, target CSSE-1.b)
- Procure additional CTS-2 platforms. (Appendix J, target CSSE-1.b)
- Provide production support for CT systems. (Appendix J, target FOUS-5)
- Develop market survey and procurement strategy for future CT systems. (Appendix J, target CSSE-1.e)

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 FY24:

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

Planned Activities in FY25:

- Continue to provide programmatic oversight of CT systems.
- Coordinate the activities for CTS-2 cycles (Appendix J, 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 FY24:

- 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 J, targets CSSE-1c and FOUS-5)
- Successfully deployed the CTS-2 cluster Amber into production completing FY24 ASC L2 milestone # 8640 (SNL)
- Specified requirements and started procurement process for the 2nd ASC funded CTS-2 system. (Appendix J, target FOUS-5)

Planned Activities in FY25:

- Continue deployment of CTS-2 systems to restricted and classified networks and transition to production use (Appendix J, targets CSSE-1c and FOUS-5)
- Deploy 2nd ASC funded CTS-2 systems with GPU and Standard processors. (Appendix J, target FOUS-5)

- Deployment of ASC CTS-2 SCN System (6SU) – "Nitro" (HPC)
- Deployment of ASC CTS-2 SCN GPU System (16 nodes) – "Cascade" (AI/ML and HPC)

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 FY24:

- Ran two additional ATCC campaigns.
- Updated Sierra in compliance with security policies.

Planned Activities in FY25:

- Complete a final ATCC campaign.
- Retire Sierra system.

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 has enabled key technologies related to the selected system architecture and continues to provide benefits to applications through its Center of Excellence (CoE) activities. LLNL is procuring El Capitan, a next-generation supercomputer, for which delivery and deployment began in FY23 and will be completed with system acceptance in FY25. 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 increased the expected delivered performance for mission critical applications over the original plan of record (POR). Subsequently, LLNL has worked with HPE and AMD, the processor partner that was selected in that decision, to refine the node architecture further. An associated contract modification that enhanced the planned usability of the delivered system and again increased the expected delivered performance significantly was completed in Q3 FY21. 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 converted 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 have ensured that effective use of El Capitan will begin soon after acceptance when it has transitioned to the classified environment. At that point, 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 FY24:

- Provided technical coordination and contractual management for the CORAL2 NRE and El Capitan contracts. (Appendix J, targets CSSE-1.a)
- Negotiated modifications to El Capitan build and CORAL-2 NRE contracts and submitted them for approval. (Appendix J, targets CSSE-1.a)
- Continued application preparations for the El Capitan system through the CoE. (Appendix J, targets IC-1 and CSSE-2.a)
- Maintained multiple El Capitan early access systems to help tri-lab teams prepare for El Capitan. (Appendix J, targets CSSE-1.a, IC-1)
- Completed siting of all El Capitan hardware. (Appendix J, target CSSE-1.a)
- Completed acceptance of the El Capitan Early Delivery and Intermediate Delivery Systems and enabled early user access to the El Capitan system. (Appendix J, target CSSE-1.a)

Planned Activities in FY25:

- Continue to provide technical coordination and contractual management for CORAL-2 NRE and El Capitan contracts. (Appendix J, target CSSE-1.a)
- Continue application preparations for the El Capitan system through the CoE. (Appendix J, targets IC-1, CSSE-1.a, and CSSE-2.a)
- Maintain multiple El Capitan early access systems to facilitate final preparations for El Capitan. (Appendix J, target CSSE-1.a)
- Complete the El Capitan Acceptance Process. (Appendix J, target CSSE-1.a)
- Deliver El Capitan system to production use. (Appendix J, targets CSSE-1.a and FOUS-1)

ATS-6 Tri-Lab Advanced Technology System Procurement (LLNL)

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 2030 - 2035 timeframe. The primary activity is to lead the design, acquisition, and deployment of the sixth AT system (ATS-6) 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.

ATS-6 will be sited at Lawrence Livermore National Laboratory but will serve the tri-lab and broader ASC community.

Planned Activities in FY25:

- Gather tri-lab mission and application requirements for future Advanced Technology System deployments.
- Develop tri-lab vendor engagement plan to evaluate potential technology options in the 2030 timeframe.
- Develop draft RFP documents according to the Tri-lab Procurement Model (TLPM) process.

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 FY24:

- Provided formal project status reporting as required for Crossroads.
- Provided ongoing risk assessment and mitigation for the Crossroads project.
- Facilitated the market surveys, the development of benchmarks, and drafted the initial technical specifications for ATS-5 (Appendix J, target CSSE-1d)

Planned Activities in FY25:

- Coordinate the activities for the removal of the Trinity system from service.
- Coordinate the development of the Acquisition, and Project Management documents along with preparing the FFRDC review for ATS-5 (Appendix J, target CSSE-1d)

ATS-5 Tri-lab Advanced Technology System Procurement (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.

Accomplishments in FY24:

- Drafted and released 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 J, target CSSE-1d)
- Developed and released the benchmarks for the ATS-5 procurement to the vendor community in preparation for release of the RFP. (Appendix J, target CSSE-1d)
- Coordinated 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 J, target CSSE-1d)

Planned Activities in FY25:

- Develop the ATS-5 Acquisition and Management Plan
- Release the ATS-5 RFP and conduct a tri-lab evaluation of responses for eventual contract award. (Appendix J, target CSSE-1d)
- Help plan and participate in an FFRDC review for the ATS-5 procurement/project in order to submit a formal report as part of the FY22 NDAA Due Out (Appendix J, target CSSE-1d)

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 FY24:

- Supported transition of Crossroads platform to full production use on classified network. (Appendix J, target CSSE-1)
- Supported ATCC code deployment on Crossroads platform. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Prepare for ATS-5 procurement and assist in Technical Advisory Team formation. (Appendix J, target CSSE-1)
- Ensure baselining for ATS-5 benchmarking is satisfactory to represent NNSA mission code needs. (Appendix J, target CSSE-1)

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 FY24:

- Coordinated procurement engagements associated with Vanguard II & III, N10, and ATS-5. (Appendix J, target CSSE-1)
- Engaged with Cornelis Networks on next-gen Omnipath network investigation and assist Cornelis on advanced DES capabilities using SST. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Ensure the initial Vanguard II system is fielded and tested at Sandia. (Appendix J, target CSSE-1)
- Investigate viable pathways for future AI-focused hardware to intersect with traditional ASC mission computing needs. (Appendix J, target CSSE-1)

Advanced Architecture Test Bed Research and Development (SNL)

This project will address a critical need for a range of experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy applications, enable application performance analysis with Mantevo proxy applications, support ~~the~~ Heterogeneous Computing and Programming Model R&D, the Software and Tools for Scalability and Performance projects, and for 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 FY24:

- Procured and investigated several hardware and software solutions supporting Accelerated Digital Engineering initiatives focused on Services capabilities. (Appendix J, target CSSE-1)
- Evaluated pre-release hardware and components from system vendors with specific focus future on NNSA platforms. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Procure and evaluate experimental hardware, including CPU, GPU, and/or AI testbed systems. (Appendix J, target CSSE-1)
- Deploy additional dataflow architectures and accelerators to provide additional continuity to reconfigurable computing investigations. (Appendix J, target CSSE-1)

Projects for the System Software and Tools Product (WBS 1.2.3.5.3)

This 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 FY24:

- Supported the integration of TOSS 4 on El Capitan, including enablement of the MI300A APU and Slingshot interconnect, as well as improvements to firmware management and the scalability of power- and console-management tools. (Appendix J, target CSSE-2.a)
- Continued support of TOSS 4 for CTS-1, CTS-1+, CTS-2, and El Capitan systems, including the release on TOSS 4.7 (based on RHEL 8.9) and 4.8 (based on RHEL 8.10). (Appendix J, targets CSSE-2.a and CSSE-2.c)

- Released TOSS 5.0, based on RHEL 9.4, to support forward-looking systems such as testbeds, as well as El Capitan. (Appendix J, targets CSSE-2.a-c, CSSE-2.e)

Planned Activities in FY25:

- Continue support of TOSS 4, based on RHEL 8, for existing CTS-1/1+/2 systems as well as El Capitan. (Appendix J, targets CSSE-2.a-c)
- Continue development and support of TOSS 5, based on RHEL 9, including minor releases to support new versions of RHEL, in support of its deployment on El Capitan and as an option for CTS-2 systems. (Appendix J, targets CSSE-2.a-c)
- Begin development of TOSS 6, expected to be based on RHEL 10, with a focus on supporting future platforms and management models. (Appendix J, targets CSSE-2.a-c, CSSE-2.e)

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 HPC platforms, supports user productivity, participates in 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 integrated by this project include compilers, linkers, debuggers, power and resilience, performance and memory tools, interfaces to the parallel environment, and associated runtime library work.

Accomplishments in FY24:

- Advanced the applications development environment for El Capitan using early access systems and the delivered RZadams system. (Appendix J, targets CSSE-1.a and CSSE-2.a)
- Supported users of LLNL ASC production platforms, including the Sierra system applications development environment. (Appendix J, target IC-1).
- The ADEPT Scalable Applications Preparations (SAP) team assisted tri-lab IC teams and users in porting to RedHat 8 based TOSS 4 for CTS-2 and using the Sierra system in computing campaigns. (Appendix J, target IC-1)
- Improved tools and programming infrastructure through support of NRE activities and working groups for the CORAL-2 Exascale system. (Appendix J, targets CSSE-1.a and CSSE-2.a)

Planned Activities in FY25:

- Continue to deploy and improve the applications development environment for the delivered El Capitan system. (Appendix J, target CSSE-2.a)

- Support users of LLNL ASC production platforms, including El Capitan, CTS-2 and the Sierra systems; on applications development, porting and HPC workflow issues. (Appendix J, target IC-1, CSSE-2.a).
- Continue to improve tools and programming infrastructure through support of NRE and LLNL research and development activities. (Appendix J, targets CSSE-1.a and CSSE-2.a)

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 provides performance visualization. The ProTools team also works with hardware vendors to design and deploy tools for ATS platforms.

Accomplishments in FY24:

- Worked with CORAL-2 Tools Working Group on El Capitan tool preparation, with focuses on improving the AMD infrastructure for performance tools and enabling Sanitizer tools. (Appendix J, target CSSE-2.a)
- Continued supporting performance analysis tools Caliper, Adiak, Gotcha, Hatchet, and Thicket. (Appendix J, targets CSSE-2.a, CSSE-2.c, and IC-1)
- Began Jupyter-based SPOT replacement, which will enable performance visualization in Jupyter notebooks. Expected to be easier to maintain than web-based SPOT and portable to other sites.

Planned Activities in FY25:

- Complete CORAL-2 Tools Working Group. Final thrusts will be around PC-sampling in GPU kernels and hardening AMD-provided performance tools. (Appendix J, target CSSE-2.a)
- Deploy Jupyter-based SPOT replacement. Will allow code teams to build customizable visualizations for tracking performance across many applications runs, such as tracking nightly performance test results or conducting performance studies.
- Continue to support in-house developed performance tools. (Appendix J, targets CSSE-2.a, CSSE-2.c, and IC-1)

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 FY24:

- Developed an unsupervised-learning approach to analysis of laser-induced breakdown spectroscopy (LIBS) data used to study metal alloy material composition. This approach augments the existing supervised-learning approach to provide a robust suite of processes.
- Delivered an AI-empowered image analysis workflow for use in explosive experiments. The tool aids the experimentalist in identifying first and last light in high resolution streak camera results from precision explosive experiments.
- Prototyped a transformer (ML) and echo state network (non-spiking neural network) approach to identifying good vs. bad welds using input power. This method seeks to assist the production agency with in-situ detection of anomalies to aid in production.
- Developed performance improvements to a neuromorphic anomaly detection parallel tool (PetaVision) that greatly enhances its performance on non-GPUs, allowing it to run on more conventionally available architectures at reasonable performance.

Planned Activities in FY25:

- Study new computed tomography (CT) datasets from manufactured parts with machine learning analysis tools (e.g. PetaVision) to evaluate efficacy of the ML tools to aid radiographers in detection and quantifying defects.
- Develop various enhancements to CMM Learning Analysis Workflow (CLAW) used on coordinate-measuring machine (CMM) datasets including ParaView integration, "clocking" and "stacking" assistants, and catalog new datasets from the national security enterprise.
- Deliver a prototype for anomaly detection in microscopy data needed by staff who have to manufacture and certify parts meet stringent tolerances for manufacturing defects.

- Perform initial analysis of polymer and metal additive manufacturing datasets and develop a machine learning approach to learning material properties such as stress, strain, and predict other behaviors under load.

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 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 FY24:

- Delivered unclassified ASC testbed (called Jarvis) in the Strategic Computing Complex (SCC) environment hosting both SambaNova's standalone DataScale system and SambaNova Suite system.
- Delivered ongoing unclassified ASC testbed support hosted in the Darwin cluster including additional Grace-Grace and Grace-Hopper nodes to support ASC IC teams preparing for Venado as well as specific AMD Genoa SKUs for ATS-5 benchmarking.
- Evaluated new hardware technologies (including AI hardware), algorithms, and codes through collaboration with other ASC and LANL LLM and ML projects utilizing both testbed and institutional computing resources.

Planned Activities in FY25:

- Deliver classified ASC testbed in HPC SCC environment with initial focus on LLM support utilizing emerging ML hardware technology (e.g., SambaNova DataScale).
- Deliver ongoing unclassified ASC testbed support hosted in the Darwin or Jarvis cluster driven by program requirements. Leverage collaborations with CCS and HPC.
- Develop active collaboration with LANL NSAI and LDRD SIC AI/ML efforts to help ASC stay aligned with progress in these areas. Deliver testbed solutions where appropriate.
- Continue active collaboration with ASC IC and ASC CSSE Systematic Improvement by Design projects with the goal of leveraging appropriate development testbed resources for ongoing path-forward research.

Systematic Improvement by Design (LANL)

System Improvement by Design (SID) is focused on new, aggressive, and comprehensive HPC design. The combination of increasing demands for computing with the technology and market challenges in HPC necessitates an agile and responsive approach to design. SID will spearhead deeper design and evaluation of applications, software, and hardware technologies for next-generation computational platforms and workloads (Mod/Sim, AI/ML, etc.) in a rapidly evolving technological landscape. High level goals for the project include:

- Establishing a well-defined framework for evaluating cost-benefit tradeoffs that can be leveraged for ATS-5 technologies and beyond.
- Successfully demonstrating a tailoring approach to hardware design that allows critical ASC applications and codes to anticipate and successfully adapt to ATS-5 and beyond.
- Informing LANL's ASC strategy on future technologies.

Accomplishments in FY24:

- Completed initial simulation and modeling of performance improvements using advanced memory technologies using a combination of existing hardware and simulation capabilities. Codesigned prototype compiler changes and software infrastructure to support these technologies.
- Completed initial measurement of the impact of code branching in Monte Carlo on both CPU and GPU technologies. Completed initial measurement of the impact of code branching in Large Language Model inference-based workloads on CPU architectures.
- Evaluated a wide variety of new hardware technologies including latest generation CPU and GPU architectures as well as AI-specific technologies using a rapid response approach with the Darwin testbed.
- Communicated computing roadmap in multiple venues / forms including a presentation at the JOWOG-34 ACS meeting.

Planned Activities in FY25:

- Complete initial development of sparse memory accelerator in simulator.
- Complete prototype code changes in UME and EAP patterns for memory acceleration.
- Complete initial development of optimized prefetch, non-temporal loads, and/or branch prediction methods and architectures.
- Complete report on the impact of hardware tailoring on workloads of interest to the ASC program.

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. 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.

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. (Appendix J, target CSSE-2)
- 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. (Appendix J, target CSSE-2)

Accomplishments in FY24:

- Contributed to the OpenMP Architecture Review Board's Language Committee toward better support in the OpenMP API for ASC platforms and applications. (Appendix J, target CSSE-2)
- Streamlined and refactored the QThreads runtime library for more robust and efficient execution and programming model support. (Appendix J, target CSSE-2)
- Worked with Vendor partner CIQ to demonstrate multi-container workflows in the Fuzzball orchestration framework, supporting client-server use cases. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Continue participation in the OpenMP Architecture Review Board's Language Committee and engagement with vendors and the open-source community to improve support for ASC platforms and applications in the OpenMP API. (Appendix J, target CSSE-2)
- Develop an implementation of the Senders/Receivers interface for C++ structured parallelism using the QThreads runtime library. (Appendix J, target CSSE-2)
- Work with Vendor partner CIQ to demonstrate scalability of the Fuzzball orchestration framework to large numbers of jobs. (Appendix J, target CSSE-2)
- Identify gaps in system software support for secure multitenancy (nonexclusive node allocation) in ASC platforms. This work supports the ASC AI R&D roadmap's data center operations goals. (Appendix J, 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 FY24:

- Released SST version 14.0. (Appendix J, target CSSE-1)Improved the scalability and performance of SST core to allow increased scale of high-fidelity simulations. (Appendix J, target CSSE-1)
- Built out initial infrastructure in SST for embedding dataflow graphs onto grid-based spatial architectures. (Appendix J, target CSSE-1)Organized and coordinated the first SST user group with over 45 participants. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Develop a ready-to-simulate mini-app suite that is aligned with NNSA vendor R&D collaborations and procurement efforts. The effort will also involve augmenting workload modeling infrastructure to increase the efficiency of model development and encompass a larger range of applications. (Appendix J, target CSSE-1)
- Enable full-scale mini-app evaluation on detailed (e.g., node) simulations. Potential mechanisms to evaluate may include multi-fidelity methodologies, application-level checkpoints, sampling, and/or leveraging application analysis tools. (Appendix J, target CSSE-1)
- Develop and implement an open-source governance document for the SST software. (Appendix J, target CSSE-1)

Interprocess Communication System Software Stack (SNL)

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

Accomplishments in FY24:

- Joined and participated in the Ultra Ethernet Consortium (UEC), a group comprising major network vendors and cloud providers (Microsoft, HPE, Intel, Meta, etc.) and charged with developing a new high-performance communication standard to serve machine learning, artificial intelligence, and traditional scientific workloads. (Appendix J, target CSSE-2)
- Presented and published paper on using statistical analysis techniques to determine communication middleware settings to improve the performance of scientific codes.
- Presented and published paper on optimizing a partitioned communication (a new communication method introduced into the Message Passing Interface standard) on NVIDIA InfiniBand networks; this paper was selected as the best paper for the conference. (Appendix J, target CSSE-2)
- Presented and published paper on using lightweight machine learning techniques to model network traffic for purposes of enabling offloaded traffic scheduling on intelligent networks. (Appendix J, target CSSE-2)
- Presented and published paper on utilizing our Configurable Messaging Benchmark (CMB) to assess the effectiveness of different implementations and configurations of partitioned communication. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Continue to represent our interests within the Ultra Ethernet Consortium (UEC) and coordinate this representation with other DOE laboratory members as the consortium moves towards a version 1.0 of the standard. (Appendix J, target CSSE-2)
- Introduce proposal to the MPI Forum for extending the Message Passing Interface (MPI) standard to enhance remote memory access (RMA) procedures, which previous work has shown capable of achieving better performance than traditional point-to-point messaging. (Appendix J, target CSSE-2)
- Apply ML natural language techniques to the problem of intelligently scheduling network traffic and allocating network resources. (Appendix J, target CSSE-2)
- Develop comprehensive evaluation of communication characteristics of CTS-2 and ATS-3 machines under different memory configurations and networks to assess opportunities for enhancing performance of scientific codes. (Appendix J, 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 FY24:

- Extended SNL-internal failure data repository/simulator to account for additional failure modes found on recent emerging platforms (neuromorphic and HBM memory). (Appendix J, target CSSE-1)
- Implemented scalable pseudo-local recovery of a prototype stencil solver from hard failures, with an implementation combining Fenix and MPI-ULFM, for scales at which conventional global recovery becomes infeasible. (Appendix J, target CSSE-1)
- Developed checkpoint-restart serialization capabilities compatible with Kokkos Resilience for the DARMA/vt serialization library Magistrate, enabling resilience for future ATS or CTS apps using DARMA/vt. (Appendix J, target CSSE-1)
- Extended the LDMS framework to utilize the extended Berkeley Packet Filter (eBPF) to allow for collection of Linux kernel-internal metric data. (Appendix J, target CSSE-1)
- Developed checksum-based implementations of Kalman filters for navigation-based applications, to investigate their resilience to processor errors and comparisons with error-free execution. (Appendix J, target CSSE-1)
- Developed fault-tolerant Communication-Avoiding Conjugate Gradient (CA-CG) using physics-based checksums to investigate failure-masking and pseudo-local recovery of iterative solvers experiencing silent errors. (Appendix J, target CSSE-1)
- Performed initial analysis of the reliability challenges facing neuromorphic computing platforms. (Appendix J, target CSSE-1)
- Demonstrated new scheduling algorithms with error injection for Kokkos Resilience OpenMP resilient execution spaces in mini-apps. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Perform periodic analysis of hardware failures on available ATS and CTS platforms to: (1) identify common failure modes, (2) put these failures in historical context against previous generation HPC systems, and (3) identify gaps in current hardware-based failure metrics. (Appendix J, target CSSE-1)
- With the instrumented hardware failures, use simulation to identify failure performance impacts a set of key applications running on ATS and CTS systems.

- Develop implementations of transparent hard failure recovery for exemplar apps (e.g. SPARC, LAMMPS) on CTS/ATS platforms. (Appendix J, target CSSE-1)
- Integrate FENIX process-resilience capability into a Kokkos Resilience release to provide online recovery for future ATS and CTS Kokkos apps. (Appendix J, target CSSE-1)
- Develop prototype implementations of fault-tolerant algorithms targeting ND embedded applications, e.g. Kalman filters, FFTs. (Appendix J, target CSSE-1)
- Design low-overhead fault-tolerant linear algebra kernels to explore advanced architecture codesign, and study fault-tolerance of AI/ML workflows. (Appendix J, target CSSE-1)
- Collaborate with Kokkos Tools in order to provide concrete fault metrics and reporting for Kokkos applications running on ATS or CTS systems, for the purposes of system administration, maintenance, and research into future systems. (Appendix J, target CSSE-1)
- Implement resilient execution spaces in Kokkos Resilience for accelerators in ATS and CTS systems. (Appendix J, target CSSE-1)
- Collaborate across platform to identify concrete application-level fault metrics and reporting mechanisms for workloads running on ATS or CTS systems to better understand application failures that do not having corresponding hardware failures. (Appendix J, target CSSE-1)

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), 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 FY24:

- Developed and released HPSS 10.3 and 11.1. (Appendix J, target CSSE-2.c)
- Explored HPSS features to make tape reads more performant. Resulted in enabling Recommended Access Order drive feature and initial deployment of batch staging tools. (Appendix J, target CSSE-2.c)
- Deployed Adaptable Storage Platform configuration for HPSS Core Server and Metadata Disk deployment to the unclassified and classified environments. (Appendix J, target CSSE-2.c and FOUS-5)

- Performed initial investigation to deploy HPSS 11.1 to unclassified environment. (Appendix J, target CSSE-2.c)
- Performed a Logical Block Protection Verify background operation across majority of data-full tapes in OCF environment. (Appendix J, target CSSE-2.c)
- Continue support for archival storage systems. (Appendix J, target FOUS-5)
- Tested archival object storage access (S3 protocol). (Appendix J, target CSSE-2.c)

Planned Activities in FY25:

- Develop/release HPSS 11.2 and 11.3. (Appendix J, target CSSE-2.c)
- Deploy HPSS 11.1 to unclassified environment. (Appendix J, target CSSE-2.c)
- Deploy tape library management software upgrade (LumOS) across unclassified and classified. (Appendix J, target FOUS-5)
- Investigate next generation tape drive technology refresh procurement. (Appendix J, target FOUS-5)
- Prototype new tool to automate archiving user data. Explore integration of new tool to automate retrieving user data with Flux. (Appendix J, target CSSE-2.c)
- Continue exploring HPSS features which make tape reads more performant (e.g. buffered tape marks, fast positioning, etc.). (Appendix J, target CSSE-2.c)
- Continue support for archival storage systems. (Appendix J, target FOUS-5)
- Evaluate deployment of archival object storage (S3 protocol). (Appendix J, target CSSE-2.c)

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 FY24:

- Integrated, migrated data to, and deployed, new S/RD Lustre file system (lustre2). (Appendix J, target FOUS-5)

- Deployed Lustre 2.15 in the OCF environment to provide higher streaming I/O and metadata performance. (Appendix J, target CSSE-2.c)
- Explored use of sharded directories in Lustre 2.15 to improve metadata distribution among servers. (Appendix J, target CSSE-2.c)
- Developed software to use Lustre job statistics in the Elasticsearch data analytics system to manage Lustre. (Appendix J, target CSSE-2.c)
- Integrated and achieved acceptance of Merced, the Lustre filesystem for the El Capitan exascale system (Appendix J, target CSSE-1.a)
- Developed and deployed new "pairwise" resource management solution for Merced to greatly improve performance and stability. (Appendix J, target CSSE-1.a)
- Worked with HPE to have El Capitan exascale system's near-node rabbit solution ready for early user access. (Appendix J, target CSSE-2.a)
- Configuration and test of self-encrypting hard drives for Lustre filesystem. (Appendix J, target CSSE-2.c)
- Deployed ZFS 2.2 which made general performance and management improvements to ZFS-based Lustre. (Appendix J, target CSSE-2.c)

Planned Activities in FY25:

- Integrate, migrate data to, and deploy new CZ Lustre file system (lustre1) and new SRD Lustre file system (lustre1). (Appendix J, target FOUS-5)
- Deploy Lustre 2.15 in the classified environment to provide higher streaming I/O and metadata performance. (Appendix J, target CSSE-2.c)
- Evaluate and decide if want to deploy sharded directories in Lustre. (Appendix J, target CSSE-2.c)
- Develop software to use Lustre job statistics in the Elasticsearch data analytics systems to manage Lustre. (Appendix J, target CSSE-2.c)
- Work with HPE to complete El Capitan exascale system's near-node rabbit solution. (Appendix J, target CSSE-2.a)
- Tool development and testing of self-encrypting hard drives for Lustre filesystem. (Appendix J, target CSSE-2.c)
- Continue to make general performance and management improvements to ZFS-based Lustre. (Appendix J, target CSSE-2.c)
- Evaluate and possibly deploy project quotas in LC Lustre file systems (Appendix J, target FOUS-5)

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 FY24:

- Evaluated Mellanox NDR for production. Future GPU systems from CTS2 will be using NDR. (Appendix J, targets CSSE-2.c and CSSE-1.e)
- Tested NVIDIA H100 GPUs in a 4-way configuration on CTS2 hardware. This included support of gdrcopy, peermem and UCX. TOSS required support of MOFED on H100 nodes to see full functionality. MOFED support is disabled by default on most systems but can be enabled where applicable. Not defaulting to MOFED was made with input from multiple labs that run TOSS in production. (Appendix J, targets CSSE-2.c)

Planned Activities in FY25:

- Evaluate Intel Granite Rapids system against CTS2 nodes. (Appendix J, target CSSE-1.e)
- Purchase latest NVIDIA CPU / GPU systems for testbed and deploy with TOSS 5. Use for user evaluation against other GPU platforms for simulations and AI/ML workloads. (Appendix J, target CSSE-2.c)
- Integrate Cornelis OPX support into TOSS 5. (Appendix J, target CSSE-2.c)

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 FY24:

- Managed NG-HPCN contract and FY24 milestone deliverables. (Appendix J, targets CSSE-1.e)
- Organized several tri-lab technical deep dive meetings on various co-design topics.

- Executed additional NG-HPCN contract milestones in the Architecture and Simulation work package.

Planned Activities in FY25:

- Manage NG-HPCN contract and FY25 milestone deliverables. (Appendix J, targets CSSE-1.e)
- Continue strong tri-lab co-design engagement with the selected NG-HPCN industry partner.

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

This project covers two key items critical for ASC's next generation computing: programming models standards, and storage and I/O. The standards work is 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.

Accomplishments in FY24:

- Evaluated and tested El Capitan Rabbit architecture; Improved SCR to support on El Capitan system and complex workflows; UnifyFS improvement for scalability. (Appendix J, target CSSE-2.a)
- Participated and represented LLNL in the MPI Forum meetings of Oct/23, Dec/23, and Mar/24. Attended regularly the MPI Forum Fault Tolerance Working Group meetings and drafted modifications to the Reinit proposal, which include support for the Sessions model to recover from failures. Attended MPI Forum Hybrid Communication Working Group meetings and discussed proposals for improving communication on GPU systems and improving memory allocations. Also attended the MPI Tools Working Group meetings and discussed improvements to the MPI-T and QMPI interfaces. (Appendix J, target CSSE-2)
- Attended PMIx monthly and quarterly meetings. Attended the Implementation Agnostic Working Group meetings. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Continue investigating I/O technologies, including work on user-level file systems, and I/O performance measurement. Continue SCR modification for El Capitan, ASC applications, and for supporting complex workflow data management. (Appendix J, target CSSE-2.a)

- Modify the specification of the Reinit fault tolerance proposal to include support for both the Sessions model and the World model in a single API. Read the specification in an MPI Forum meeting in FY25. Continue our participation in the MPI Forum meetings in FY25. Continue representing LLNL in discussions on MPI tool support and GPU support such as device-triggered partitioned communication. (Appendix J, target CSSE-2)

Continue attending PMIx monthly, quarterly, and Implementation Agnostic Working Group Meetings and contributing to improvements to the PMIx Standard. (Appendix J, target CSSE-2)

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 FY24:

- Continued the development of the Capacity On Demand User Interface and Toolkit (CONDUIT) for rapidly and atomically shuttling data between HPC storage tiers, in support of FY24 L2 milestone #9030.
- Added a Quo-Vadis package to Spack.
- Finalized a new bugfix release (version 5.0.3) of OpenMPI.
- Enabled PAGOSA to reduce memory used by equation of state tables by adding support for EOSPAC's shared memory feature.
- Developed a new Spack environment for the Lagrangian Applications Project using OpenMPI on Crossroads.
- Integrated XCAP codes into xRAGE.
- Improved FleCSolve (FleCSI solvers) and integrated it into host codes based on the FleCSI framework.

Planned Activities in FY25:

- Continue to provide code support (bug fixes and feature requests) for data management utilities (CONDUIT, pftool, GUFI) which includes implementing advanced features requested and designed throughout early use of the solutions.
- Collaborate with industry and academia to evaluate emerging storage system technologies using realistic ASC workloads.
- Study the use of state-of-the-art columnar data analytics technology paired with computational storage to perform fast, efficient analysis of ASC's scientific data sets.
- Continue interactions with the HPSS consortium to provide development, maintenance, and testing to support ASC archival storage.
- Release a production version of Quo-Vadis, including multi-node mixed-mode support and improved mapping policies for mixed-model couplings.
- Release Open MPI 6.0.0 featuring improved GPU support plus support for MPI 4 big count and MPI 4.1 GPU memory kind.
- Assist the ASC IC Safety Applications Project with integrating XCAP libraries into PAGOSA.
- Support Open MPI as an option on ATS-4 systems, through the Spack environment of one or more ASC IC codes.

Production Input/Output Services (SNL)

- HPSS provides a Kerberos enhanced S3 interface to meet lab authentication requirements. This extends HPSS scalable cold storage to lab software that utilizes S3 in a secure environment. (Appendix J, target CSSE-1)
- FrETT includes a containerized web server to provide a user-friendly data transfer service. This also provides transfer and SMSS storage integration for lab products which can utilize web-based services (e.g., Sandia Analysis Workbench). (Appendix J, target CSSE-1)

Planned Activities in FY25:

- HPSS will include data path encryption for parallel data transfers. Increased lab network security requirements will necessitate scalable, high performance, secure data transfer capability. (Appendix J, target CSSE-1)

Scalable Data Management and Input/Output Support (SNL)

The Scalable Data Management and Input/Output Support 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 data-management capabilities for digital engineering. Success requires close collaboration with IC, digital engineering, 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.

Accomplishments in FY24:

- Developed a container for continued development of EMPIRE and IOSS/S3 cloud storage and developed abstractions of two RAMSES workflow exemplars to enable provenance and single source of truth concepts (when data may be replicated across storage systems). Both accomplishments are key advances towards a performant data model for RAMSES. (Appendix J, target CSSE-2)
- Successfully executed EMPIRE tests writing data directly to S3 object-storage using the IOSS to S3 backend. This is an important step toward a cloud-storage integration with high-performance computing modeling and simulation codes. (Appendix J, target CSSE-2)
- Collaborated with Empire Fluid and ParaView/Kitware teams to provide explicit support for Discontinuous Galerkin Fields in IOSS and Exodus. The DG and enhanced field capability is now ready for client use. (Appendix J, target CSSE-2)

- Collaborated with the LDMS team to integrate Darshan and other I/O performance tools into LDMS to assess performance and scalability issues in production codes. (Appendix J, target CSSE-2)
- Worked with DDN cross-API access team to identify an interface to allow data in the same dataset to be read and written by both S3 and POSIX operations. Both parties contributed to the API specification which is now in the statement of work for DDN. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Perform R&D to develop enabling capabilities for data-management in support of Digital Engineering and Computing as a Service (CaaS). This activity aligns with the Accelerated Digital Engineering initiative. (Appendix J, target CSSE-3.d)
- Begin a joint IC, CSSE implementation of the RAMSES data model designed in the FY24 level 2 milestone. The implementation will include concepts of data pedigree/provenance, intermediate cloud-like data storage during workflow execution of highest priority RAMSES use cases. (Appendix J, target CSSE-2.d,e)
- Evaluate technology advances in I/O and networking software and hardware to understand their potential impact on ASC applications and I/O software. (Appendix J, target CSSE-1.b)
- Provide production I/O-library support for ASC applications. (Appendix J, CSSE-2)
- Work with DDN to develop a viable storage-system solution for ATS platforms (Appendix J, CSSE-2 b,d)

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. Hardware capabilities include three production visualization servers (Pascal, Tron, and Vertex) and VNC (Virtual Network Computing) servers in three security zones (CZ, RZ, and SCF). These VNC servers support visualization and remote use of interactive tools, allowing users to easily connect to and launch graphics on our compute systems. 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 vis experts from the project, who can augment, design, or create visualizations (animations, movies, and images) of the users' data. The project provides support for high-resolution display devices and facilitates demonstrations and presentations on these displays.

Accomplishments in FY24:

- Maintained data analysis and visualization environment across LC platforms. (Appendix J, target CSSE-2.c)
- Maintained/improved user access to VNC servers. (Appendix J, target CSSE-2.c)
- Migrated CZ, RZ, and SCF VNC servers to TOSS4 and facilitated gradual user transitions from TOSS3 to TOSS4 (via simultaneous support).

- Improved user experience & stability on SCF VNC server with VNC Session Manager software and increased node count on server.
- Assisted LANL with their own VNC deployment and supported October 2023 LANL COOP with NICE DCV & Ensight access, setup, & instructions.
- 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 visualizations of scientific data.

Planned Activities in FY25:

- Maintain visualization environment across LC platforms and networks. (Appendix J, target CSSE-2.c)
- Continue VNC server support, enabling remote connections & an environment pre-configured for graphics. (Appendix J, target CSSE-2.c)
- Provide support for high-resolution displays and for demos, tours, and special events (Appendix J, target CSSE-2.c)
- Maintain high standards for visualization infrastructure, upgrading as needed (Appendix J, target CSSE-2.c):
    - Migrate SCF VNC server to new hardware during FY25.
    - Replace `Pascal` with new CTS-2 system supporting visualization.
    - Assess replacing our high-resolution display used for demos and tours.
- Support LC user vis needs with software consulting and visualization creation. (Appendix J, target CSSE-2.c)

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 FY24:

- In Hopper, the project completed the split of the application into components, factoring out proprietary and core logic into separate, standalone repos. This allowed the creation of an open-source version of Hopper that can be used throughout the computing world, which we can extend with LLNL-proprietary elements for use on our site. (Appendix J, targets CSSE-2.a and CSSE-2.c)
- Completed development of Hopper CI capabilities for evaluating vulnerabilities of 3rd party libraries. Created site files used for configuring the app based on the computing resources at the site. (Appendix J, targets CSSE-2.a and CSSE-2.c)
- In Lorenz, completed an initial implementation of Flux support into MyLC's job capabilities. This allowed monitoring of Flux jobs in addition to SLURM jobs. (Appendix J, targets CSSE-2.a)
- Developed Lorenz features that extend the change-quota portlet to provide Hotline staff with more and easier-to-use capabilities when managing quotas on behalf of users. (Appendix J, targets CSSE-2.c and FOUS-5)

Planned Activities in FY25:

- In Hopper, add analysis of file interchange service connection problems to give users very specific error messages. (Appendix J, targets CSSE-2.a and CSSE-2.c)
- Fine-tune Hopper strategies and CI for managing the open-source version of Hopper core. Maintain the application across hardware and OS enhancements and continue to improve usability and performance. (Appendix J, target CSSE-2.a and CSSE-2.c)
- In Lorenz, enhance the Job Management functionality to be able to more fully utilize newer concepts made possible by Flux. (Appendix J, target CSSE-2.a)
- Refresh the MyLC dashboard UI. Develop v2 of the Lorenz REST API and perform major library updates. Continue to enhance the MyLC Dashboard with user and staff requested features. (Appendix J, targets CSSE-2.a and CSSE-2.c)

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 FY24:

- Maintained a consistent and stable interactive, batch, and in situ visualization environment on LANL supercomputers.
- Maintained and supported the LANL Visualization corridor: power wall theater, cave virtual environment (La Cueva Grande), and desktop stereo access for select users.
- Consistently developed new content for signature facilities while supporting ASC core mission users.

Planned Activities in FY25:

- Continue maintaining and supporting interactive, batch, and in situ visualization environments, both software and hardware (visualization corridor), for LANL ASC users.
- Explore 3D display technologies for interactive visual exploration and explanation by ASC and their users.

Cross-Cutting Extreme-Scale Research (LANL)

LANL's 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). LANL science workflows create complex sets of data that are increasingly difficult for scientists to understand without *assisted* analysis and visualization capabilities that take advantage of the fast-moving frontier of AI and ML. This project develops and supports forward-looking, cross-cutting technologies for in-situ data analysis and exploration workflows, and collaborates with production visualization and DSI projects to deliver production solutions that directly impact the analysis being done by LANL scientists. 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 algorithms for use with ParaView, Cinema, and CMF workflows.

Accomplishments in FY24:

- Delivered additional capability in Cinema toolkit, supporting updated Cinema:Explorer and Cinema:Viewer applications, in support of DSI requirements, and for general Cinema database analysis and visualization.
- Developed AI/ML support capability for data analysis, characterization, and reconstruction within the Cinema toolkit. This included the release of Synema project repository, and pycinema capabilities for executing trained models and visualization of results.
- Supported Capture and Extraction workflows, such Synthetic Radiograph workflow, in response to user requests.

Planned Activities in FY25:

- Deliver Cinema capability, including pycinema toolkit, viewers and applications, in support of DSI customers, and for general Cinema database analysis and visualization.
- Prototype AI/ML components into data workflows to assist in understanding AI foundation models, and curation, evaluation, and understanding of complex ensembles of data.
- Deliver Cinema-assisted workflows supported by pycinema releases, addressing LANL use cases, and made available (as appropriate) in CMF for production analysis and visualization.

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 FY24:

- Released open source SlyPI (Slycat Python Integration) Library, including documentation and data analysis examples. (Appendix J, target CSSE-2)

- Developed containerized “as a service” proxy application to demonstrate embedded components of Slycat analyses and visualizations within web-based mod-sim tools, such as DetNet or TABS. (Appendix J, target CSSE-2)
- Deployed ParaView/Catalyst 5.12 for Sierra codes to CEE, CTS1, and CTS2. SPARC deployment available on CTS1 and CTS2. (Appendix J, target CSSE-2)
- Used Omniverse to visualize E3SM SCREAM data in VR. (Appendix J, target CSSE-2)
- Improved the pipeline from ParaView to Omniverse; presented the approach and results at NVIDIA’s GTC 2024. (Appendix J, target CSSE-2)
- Released Ensight 2024 R1 on supported platforms. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Extend Slycat visualization capabilities emphasizing incorporation of machine learning methods into analysis and visualization of scientific machine learning models. (Appendix J, target CSSE-2)
- Expand Slycat user community and adoption by demonstrating usage in Digital Engineering and Cloud environments. Includes ongoing modernization efforts to support target infrastructure and environments. (Appendix J, target CSSE-2)
- Continue Paraview, Catalyst, and Ensight releases with production support. Three concurrent releases of Paraview and Catalyst, and two for EnSight. (Appendix J, target CSSE-2)
- Deploy, support, and expand adoption of Catalyst to SPARC and SIERRA codes. (Appendix J, target CSSE-2)
- Provide three training sessions for Paraview GUI, Paraview Python, and Catalyst with additional training for Ensight and Slycat as time and resources permit. (Appendix J, target CSSE-2)
- Develop and demonstrate AR/VR technology on datasets relevant to ASC. (Appendix J, target CSSE-2)
- Develop and demonstrate distributed renderer improvements on ASC-relevant visualizations. (Appendix J, 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 FY24:

Quantum (Appendix J, target CSSE-3.c)

- Implemented next-generation quantum simulation hardware testbed platform by installing superconducting quantum processor and supporting 5 connected qubits. Calibration routines and gates were made available to users.
- Collaborated with industry on better implementation of their quantum system.
- Identified proof of principle nuclear physics emulation system – neutrino interaction in supernovae collapse. Begun circuit design for emulation.
- Completed galvanic tunable coupler design for micromachined cavities -- first design of its kind for 2.5D cQED; comparable coupling to existing 3D couplers (~10s MHz); amenable to 2D fabrication techniques in silicon / sapphire
- Developed a full integration of the C++ Quandary optimal control code into a python environment, including its deployment, testing and maintenance on the QuDIT control computer. Developed a time-parallelization strategy to speed up gradient computations for optimal control utilizing many-core HPC platforms.

Neuromorphic (Appendix J, target CSSE-3.b)

- Participation and leadership in the LLNL AI3 AICoE with both Cerebras and SambaNova, focusing on the development of a Cognitive Simulation workflow with the upgraded CS-2 and SN30 systems.
- Developed a scalable AI benchmark for leadership-class HPC+AI systems and future procurements. By the end of FY24, it is expected that we will have a scale-invariant AI benchmark that is optimized for profiling the end-to-end performance of next-generation leadership-class advanced technology HPC systems and their software stack.
- Developing ExaLearn and ExaGAN models for the Cerebras and SambaNova systems.
- Mapping Lulesh to DaCe and then to the Cerebras CS-2.
- Investigated credibility assessment algorithms for neural networks related to atomic physics calculations. This is still an active area of investigation.

Planned Activities in FY25:

Quantum (Appendix J, target CSSE-3.c)

- Implement next-generation quantum simulation hardware testbed platform by bringing up 10 connected qubits on superconducting quantum processor.
- Implement multi-qubit quantum optimal control.
- Demonstrate warm dense matter simulation on 5 qubits.
- Simulate circuit for nuclear physics quantum emulation and determine fabrication feasibility.
- Enable GPU utilization for Quandary to speed up classical compute time for optimal control. Develop generalized carrier frequency analysis for optimal control on flux-tunable qubits. Implement a pulse-level variational quantum eigensolver in Quandary.

Neuromorphic (Appendix J, target CSSE-3.b)

- Continue exploration of the CosmoFlow proxy application on the SambaNova and Cerebras AI accelerators. Extend work to CosmoGAN generative neural network.
- Explore usage of AI accelerators for LLMs and DarkStar surrogate models.
- Complete a scalable AI benchmark work and publish a paper.
- Support the integration of AI accelerators with LC HPC, and support their accessibility for researchers, including updating user-facing software stacks.
- Establish a roadmap and execution plan for tighter integration of AI accelerators into LC to support mission space.

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. Recent hardware advances indicate that the time horizon for fully fault-tolerant quantum computation is within 3-5 years. Our quantum computing research focuses on both theoretical and applied studies designed to understand and test the capabilities of quantum computing concepts and hardware. The work has impact on our understanding of best use of these computational architectures for ASC-related applications. Our neuromorphic computing research focuses on both theoretical and applied studies designed to understand and test the capabilities of neuromorphic computing low power AI/ML. The work has impact on our understanding of best use of these computational architectures for ASC related applications. Here, we additionally study how machine learned flux limiters, integral to the stability of computational fluid dynamics codes, can improve a flux limiter's accuracy. This work focuses on developing and expanding the applicability of these methods to be of more value to the ASC applications.

Accomplishments in FY24:

- Reduced resource requirements for quantum linear systems algorithm by an order of magnitude.
- Improved quantum algorithms for solving partial differential equations using quantum linear systems algorithm.
- Developed optimal, symmetry-enhanced schemes for quantum simulation of condensed matter and material sciences systems.
- Developed neuromorphic algorithms for implementing machine learning algorithms and testing on new Intel Loihi 2 neuromorphic processor.
- Assessed machine learned probabilistic flux limiters for Burger's equations and machine learned slope limiters for 1D Euler's equation in xRAGE.
- Established a subcontract with University of Colorado, Boulder to study machine learned flux limiters for hypersonic and supersonic flows.
- Established a new fault-tolerant quantum simulation collaboration for nuclear dynamics and materials EOS in co-design with QuEra, a vendor of neutral atom quantum computers with error correction capabilities.

Planned Activities in FY25:

- Assess resource requirements for fault-tolerant nuclear scattering simulation on QuEra neutral atom quantum computer.
- Assess resource requirements for simulation of nonlinear classical systems using quantum linear systems algorithm through a subcontract with PsiQuantum.
- Monitor NISQ hardware offerings and conduct benchmark testing.

- Investigate neuromorphic algorithms for implementing transformer attention heads and testing on new Intel Loihi 2 neuromorphic.
- Study machine learned slope limiters for hypersonic and supersonic flow and Euler equations extended to reactions through a subcontract with CU Boulder.
- Integrate Bayesian inference within flux limiter methods for solving multidimensional fluid dynamics problems.

Non-Conventional Computing Technologies (SNL)

- Created and validated the Content Block Reduction method based "CBR2D/3D" Predictive Quantum Simulator for investigating the state-of-the-art and beyond CMOS device technologies including inelastic scattering model for conducting properties at room temperature. (Appendix J, target CSSE-3)Received, announced, and initial efforts with world's largest neuromorphic system.
- Received first, single chip, SpiNNaker2 evaluation hardware. (Appendix J, target CSSE-3)
- Developed and released open-source neural random walker code. Developed the infrastructure to automatically generate fixed-point implementations of neural algorithms, a necessary step for supporting low-level assembly code (DASM) on Loihi-2. (Appendix J, target CSSE-3)

Planned Activities in FY25:

- Investigate suitability of *CMOS-compatible* Atomic Precision Advanced Manufacturing (APAM) technologies for developing 2 Transistor-1 Capacitor (2T-1C) device-based Neural Processing Units. (Appendix J, target CSSE-3)
- Receive and stand up large-scale SpiNNaker2 testbed. (Appendix J, target CSSE-3)
- Benchmark billion neuron Loihi performance. (Appendix J, target CSSE-3)
- Explore mixing HPC and neuromorphic simulations using a single tool to distribute work based on appropriateness of match between structure of model and machine capability. (Appendix J, target CSSE-3)
- Develop neural-inspired AI algorithms for mission applications. (Appendix J, 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 FY24:

- Continued developing supporting software in preparation of our FY26 materials aging hybrid quantum-classical testbed demonstration. (Appendix J, target CSSE-3)
- Published initial assessment of the limitations of algorithmic primitives related to physics simulations on early fault-tolerant architectures. (Appendix J, target CSSE-3)
- Advanced capabilities with in-house and external resource estimation in pursuit of better understanding the prospects for achieving useful parallelism in simulation applications. (Appendix J, target CSSE-3)

Planned Activities in FY25:

- Investigate whether hardware accelerators can be used to scale up our planned FY26 materials aging hybrid quantum-classical testbed demonstration. (Appendix J, target CSSE-3)
- Develop logical application benchmarking protocols on small color code architectures. Propose their implementation on available quantum testbeds. (Appendix J, target CSSE-3)
- Begin campaign to identify quantum algorithms with super-quadratic quantum advantages in more general scientific computing applications. (Appendix J, 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 FY24:

- Concluded Campaign 1 exploring the initial stages of the release of auto-inhibition of BRAF. The campaign involved 95,000 coupled simulations utilizing 14.5M GPU-hours on 2048 nodes of the Lassen supercomputer and generating 157 TB of simulation data.
- Analysis of Campaign 1 data has revealed unexpected conformation changes in BRAF leading to the release of auto-inhibition.
- Substantial re-parameterization of ML model to capture greater conformational changes associated with the fully released BRAF protein.
- Initiated Campaign 2 exploring the creation of fully open BRAF protein on the membrane – creating the "unfurled" structure that would allow dimerization.

Planned Activities in FY25:

- In our final year, we will expand the ML component to incorporate a conditional generative capability, in preparation for Campaign 3.
- Perform Campaign 3 exploring the dimerization of BRAF, leading to the creation of an active (signaling) complex. This is a large-scale (Frontier/El Capitan class) computing campaign anticipated to use close to 1M node-hours on Frontier.
- Wrap up and submit the remaining (7-10) manuscripts.
- Disseminate data arising from simulation campaigns.
- Open-source the new code components that were developed for Campaigns 2 and 3.

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 FY24:

- Demonstrated Flux Rabbit scheduling on production rabbit hardware and software on Hetchy. (Appendix J, target CSSE-1.a)
- Expanded support for affinity management on AMD platforms, as well as integration with mpibind, and prepared for MI300a support with experiments for co-scheduling in progress. (Appendix J, target CSSE-2.a)
- Significant advancements in OpenMP support for required use-cases through working with the standards body as well as a move to improve the overall offload ecosystem through transitioning libomptarget to llvm-offload. (Appendix J, target CSSE-2.a)

Planned Activities in FY25:

- Work with standards bodies and vendors on continued improvement of programming models and compilers for next-generation systems, mainly through OpenMP, C++, and portability libraries like RAJA and desul (Appendix J, target CSSE-2.a)
- Demonstrate flux Rabbit scheduling at scale on El Capitan (Appendix J, targets CSSE-1.a and CSSE-2.a)

- Continue efforts to improve the feature set and ergonomics of flux system instance for users and site administrators (Appendix J, target CSSE-1.a, CSSE-2.a, CSSE-2.c)
- Continue work with vendors on tool and affinity support for El Capitan and future deployments (Appendix J, target CSSE-1.a)

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.

Accomplishments in FY24:

- Continued collaboration with AWS on Spack, Flux, and containerized versions of MARBL and ALE3D. Discussed future on-prem hardware and S/RD cloud plans and collaborated with IT and NA-IM on an S/RD cloud pilot expected to start in FY25. (Appendix J, targets CSSE-1 and IC-1)
- Began discussions with Google on their Google Distributed Cloud Air-gapped (GDCH) system, which provides many cloud capabilities on-prem, without a requirement to connect to a public cloud region. (Appendix J, targets CSSE-1)
- Ran a successful on-premises cloud system procurement resulting in the purchase of a half-rack (16 sleds) Oxide computer system for LC + tri-labs. (Appendix J, targets CSSE-1)

Planned Activities in FY25:

- Partner with industry to evaluate system architectures for future procurements and to advance computing technologies of interest.
- Work with SD code teams to evaluate usefulness of GovCloud and S/RD cloud for cross-NNSA-Enterprise workloads, focusing on problems of interest to Y-12. (Appendix J, target CSSE-1 and IC-1)
- Continue to explore AWS, GDCH, Azure, and other cloud vendor options for S/RD on-prem and cloud capabilities, focusing on both open source and programmatic collaborations. (Appendix J, target CSSE-1)
- Collaborate with Oxide Computer Company to enable multi-rack support as well as tight integration with Flux and the LC HPC Center. Enable code teams to more easily deploy services with their HPC jobs. (Appendix J, target CSSE-1)

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. To enable multiple codes to use the same input, the project supports code agnostic setup tools such as C2C (code-agnostic contours) and Klee (code agnostic shaping). It also supports the enhancement and productization of code-to-code curve and surface transfer tools to enable pipelines of codes in different physics regimes.

Accomplishments in FY24:

- Integrated Sina C++ library into the Diablo engineering analysis code using Fortran to C++ interface. (Appendix J, target IC-1)
- Productized the Echnida library with improved DevOps and build infrastructure. (Appendix J, target CSSE-2.a, CSSE-2.c, and IC-1)
- Created new depth-first execution order for Maestro DAGs to support high impact user need. (Appendix J, target CSSE-2.a, and CSSE-2.c)
- Added ability to easily integrate Sina databases with standard Python ORM tools, enabling integration with model management capabilities being developed in Strategic Deterrence.

Planned Activities in FY25:

- Enhance Maestro with a SQL database for tracking workflow status and improve Maestro restart capabilities using this database. (Appendix J, target CSSE-2.a, and CSSE-2.c)
- Support further productization of curve and surface transfer tools. (Appendix J, target IC-1)
- Move the Sina C++ library into the Axom build and source code environment, integrate with additional Engineering analysis codes, work towards normalization of Sina outputs across ASC codes.

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 FY24:

- Added compiler runtime and libc handling to Spack, enabling the smooth installation of binaries across different Linux systems and compiler versions. (Appendix J, target CSSE-2)
- Developed a performant implementation of error tracing for Spack's concretizer, enabling much more detailed error messages that identify root causes of issues. (Appendix J, target CSSE-2)
- Improved the UX of Spack for developers, including features needed by MARBL and other code teams. Improved environment UI, added support for conflicts and strong preferences in config, and simplified querying not-yet-installed packages. (Appendix J, targets CSSE-2 and IC-1)
- Continued to test Spack on LLNL EA and El Capitan systems, as well as on TCE and in code team developer environments. Developed and reviewed/merged needed features for TCE, WSC/CP, and workflow teams. (Appendix J, target CSSE-2.a)

Planned Activities in FY25:

- Complete implementation of compilers-as-dependencies, enabling users to start on a system with no compiler and install a compiler from scratch. (Appendix J, target CSSE-2)
- Enable bare-metal binary installation for MPI- and GPU-dependent packages with concretizer and build farm improvements. (Appendix J, target CSSE-2)
- Curate Spack's binary cache and enable binary installation by default. (Appendix J, target CSSE-2)
- Continue to enable Spack on El Capitan as it comes online, as well as on TCE and in ASC code team environments. (Appendix J, target CSSE-2.a)

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 FY24:

- Improved resilience of major BEE components.
- Began integration of BEE with the Common Modeling Framework (CMF) for production workflows, including a working demonstration of running a BEE workflow from CMF. Implemented a Python API to build the Common Workflow Language (CWL) specification.
- Worked with ASC IC DevOps to manage workflows such as scaling studies of production codes (FLAG) since in-house GitLab runners are currently overtaxed. Enabled remote submission of BEE workflows from Continuous Integration pipelines to support distributed testing. Demonstrated a distributed FLAG workflow on multiple nodes.
- Studied RO-Crate for storing and retrieving provenance data and found it to be a useful addition to the workflow orchestration system.

Planned Activities in FY25:

- Expand support for Continuous Integration (CI) workflows, including distributed parallel workflow on production systems that starts BEE from CI, and running BEE from CI on a remote system.
- Automate deployment of expanded workflows of a production application from LANL’s Common Model Framework (CMF).
- Continue to improve resilience of BEE and the workflows it orchestrates, including the ability to pause running workflows when stopping BEE and the ability to consolidate workflow graphs.
- Start developing a plan to integrate BEE with Data Science Infrastructure (DSI).
- Expand outreach to LANL ASC users who will use BEE as part of their production workflow.

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 FY24:

- Released FleCSI 2.3.0 adding an N-tree topology category (fully GPU-enabled); a more powerful interface to multidimensional structured grids; several utility routines for working with topology data, including support for rendezvous operations between meshes; a complete, documented set of public header files; a more flexible system for initialization; and additional support for ranges. FleCSI 2.3.0 also supplies a number of bug fixes, particularly with structured-grid auxiliaries (and halo exchanges thereof) and Legion performance.
- Together, the FleCSI and Ristra projects successfully ran Ristra's Moya application atop FleCSI on El Capitan-style hardware (MI300 GPUs, both A0 and A1 steppings) during Center of Excellence hackathons in December 2023 and March 2024. In both hackathons little effort was needed to get up and running on the new hardware. Between December and March, the collective efforts of Legion, Kokkos, FleCSI, and Ristra resulted in a 2x speedup in our benchmark execution both on the older MI250 and newer MI300 processors. Due to combined hardware and software improvements, we're observing a 3x speedup relative to our earlier benchmarking effort.
- Integrated dynamic state tracking to reduce unnecessary data movement between the CPU and GPU when using FleCSI's MPI backend. Data determined not to have been modified since last copied from CPU to GPU does not need to be re-copied to the GPU before the GPU can use it, which results in faster and more efficient data storage.

- Improved the performance and correctness of FleCSI's logging facility.
- Improved the memory efficiency of FleCSI's Legion mapper, which maps tasks to processor cores, thereby enabling larger problems to be run.
- Worked closely with NVIDIA's compiler team to address bugs in and correct limitations of NVIDIA's nvcc compiler. Based on this collaboration, NVIDIA is making changes to their compiler for FleCSI's benefit, and we are exploring workarounds for their remaining incompatibilities.
- Extended the FleCSolve package with support for algebraic multigrid.
- Supported the integration of the PUMA libraries into Ristra's Moya code.
- Produced a range of example specializations (e.g., finite-difference, finite-volume) that are already being used as starting points for other projects.
- Integrated Catalyst visualization and NetCDF output into example specializations.

Planned Activities in FY25:

- Complete support for checkpoint/restart and work with clients (in particular, Ristra) to integrate this into their applications.
- Exploit Legion's support for dynamic task variants, which lets developers provide multiple versions of a task (e.g., a CPU-optimized version and a GPU-optimized version). FleCSI's task mapper will decide at task-launch time which variant to run based on performance characteristics and resource availability.
- Implement asynchronous loop termination to enable checking for convergence in iterative solvers while concurrently performing the next iteration of the loop.
- Provide efficient storage for cells containing multi-material data and fast and flexible access to those data.
- With Louisiana State University, construct a FleCSI backend based on LSU's HPX run-time library. The intention is to provide an alternative to FleCSI's Legion and MPI backends, striking a different balance between performance and flexibility.
- Speed up ghost copies of dense fields in the context of GPU computation on AMD systems when using the Legion backend.
- Introduce a set topology, intended to be used by ASC applications that simulate non-interacting particles. (FY24's N-tree topology supports particles with direct interactions.)
- Make more sophisticated use of the Kokkos library, better exploiting nodes containing multiple GPUs, supporting multiple Kokkos streams on a single device, launching multiple OpenMP instances from Kokkos (to benefit execution on Crossroads), and letting FleCSI clients control where data-parallel code should run.

- Enhance the FleCSI task mapper to support running MPI-based libraries concurrently with FleCSI code and reduce copies by more carefully managing memory spaces shared between subsequent tasks.
- Avoid CPU/GPU data copies on systems such as El Capitan in which CPUs and GPUs physically share memory.
- Take advantage of contiguous ghost cells on unstructured meshes to enable increased vectorization and therefore performance.
- Let tasks access per-color non-field data—essentially like global variables in an MPI program.
- Begin adding an interface and developing a specialization to support Discrete Ordinates (SN) methods.
- Support ragged fields (e.g., varying material per cell) on GPUs.
- Aggregate task-dependency information so Legion can more efficiently schedule and place tasks.
- Implement pre-partitioned X3D mesh support.
- Extend fvm example application with AMR and multi-material support.
- Implement a particle-in-cell example application with an associated specialization based on the FleCSI narray and set topologies.

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 integrated ASC project that crosscuts other sub-elements including CSSE and V&V.

Accomplishments in FY24:

- Developed core infrastructure capabilities including a UI/UX frontend for users to interact with ECMF by registering models, defining automated scheduling to run models, and to capture and store model status. (Appendix J, target CSSE-2)
- Demonstrated the initial capability to run a full system classified model on our classified network and documented lessons learned. (Appendix J, target CSSE-2)
- Implemented a Scrum methodology on the ECMF team including sprint planning, biweekly Scrum meetings, sprint reviews, retrospectives, and backlog management. (Appendix J, target CSSE-2)

- Aligned with laboratory digital engineering strategic efforts by partnering with Accelerated Digital Engineering (ADE). (Appendix J, target CSSE-2)
- Participate in a L2 Milestone (#9018) titled 'ND Workflow Development & Roadmap for Secure Scalable Data Management'. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Integrate with ADE project planning and participate with FY25 pathfinders to ensure tighter integration with ASC and WE&P mission objectives. (Appendix J, target CSSE-2)
- Demonstrate the ability to plug-and-play other ND-relevant exemplars into ECMF and to increase the number of SAW SDM workflows that are sustained. (Appendix J, target CSSE-2)
- Produce provenance standards for workflows and their data stored with ECMF.
- Release the next version of the ECMF prototype. (Appendix J, target CSSE-2)

Next-Generation Development and Performance Analysis Tools (SNL)

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

Accomplishments in FY24:

- Provided performance analysis and system guidance for advanced systems. (Appendix J, target CSSE-2)
- Performed and participated in formal benchmarking efforts (DOE MEXT, InPEx, contracts). (Appendix J, target CSSE-2)
- Provided extensive focus on ATS-3 Acceptance and Standup, ATS-4 early access (COE efforts, app teams), and initial ATS-5 planning. (Appendix J, target CSSE-2)
- Investigation of bio apps for Exascale systems continues to uncover new/non-standard HPC performance issues (RDMA, UPC++, PGAS). (Appendix J, target CSSE-2)

- Working on base container image support on HPC platforms. (Appendix J, target CSSE-2)
- Improving APT's integration with basic research and built conduits with universities, hosted several interns as a result. (Appendix J, target CSSE-2)
- Prototyped and had and initial deployment effort of a strategic improvement to the Trilinos AutoTester, moving towards the usage of GitHub Actions and containers to better support the open-source development of Trilinos, as well as poise it better for usage of external cloud-like resources in the future. (Appendix J, target CSSE-2)
- Supported software sustainability of the partitioned communication interface in Open MPI by creating and contributing unit tests, identifying and fixing functionality issues, and designing and developing a higher performance partitioned communication module. (Appendix J, target CSSE-2)
- Defined the new NNSA Integration effort with input from Sandia, LANL, and LLNL. Participating in the broader software stewardship organization effort around ecosystem integration. (Appendix J, target CSSE-2)
- Proposed a three-pronged effort between late FY24 – FY27 focused on (1) metrics collection for DevOps tools and processes, (2) approaches for identifying and quantifying technical debt, and (3) research strategies to reduce identified technical debt. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Continue offering dedicated performance analysis / system guidance for advanced systems. (Appendix J, target CSSE-2)
- Deeper engagement with formal benchmarking efforts via DOE MEXT, InPEx, and ATS procurements. (Appendix J, target CSSE-2)
- Focus on ATS-4 hardening and ensuring Sandia apps build, run, & scale well on El Capitan and/or Eldorado. (Appendix J, target CSSE-2)
- Revamp proxy app R&D efforts to build integration w/ Sandia codes and communicate needs to wider community, including with more meaningful Biology proxies for Exascale. (Appendix J, target CSSE-2)
- Ensure base container images with Sandia software can be supported & performant on El Capitan and/or Eldorado systems. (Appendix J, target CSSE-2)
- Continue towards better "foundation" integration for Trilinos into HPSF, including the completion transition of Trilinos to AutoTester2 (a container-based CI system from SEMS that integrates well with GitHub Actions), as well as enhanced security scanning for better software provenance. (Appendix J, target CSSE-2)
- Will provide continued support of Open MPI through partitioned communication by optimization, fixing any functionality issues that arise, and exploring extensions (including a more general interface, and partitioned collectives) and extending the

remote memory access interface to support more HPC application use cases. (Appendix J, target CSSE-2)

- Determine NNSA software project needs more granularly and work towards increased participation in the scientific software ecosystem, as well as capturing greater benefits from the ecosystem for the benefit of NNSA codes. (Appendix J, target CSSE-2)
- Develop novel AI-assisted approaches to reduce existing and future ASC software technical debt for decades to come. (Appendix J, target CSSE-2)

Advanced Tri-lab Software Environment (ATSE) (SNL)

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 FY24:

- Delivered updates to the ATSE stack build process that enables better separation of platform-specific details, enabling better compartmentalization of vendor proprietary information. (Appendix J, target CSSE-2)
- Delivered to users several iterations of containers for NextSilicon software stacks, Spack versions, and general package updates. (Appendix J, target CSSE-2)

- Collaborated with Tri-lab on an L2 milestone regarding shared development infrastructure for ASC programming environments, which has included standing up cloud infrastructure for stack development activities and running experiments in deploying portable ATSE-based containers between HPC systems among the Tri-Labs. (Appendix J, target CSSE-2)
- Collaborated with RISC-V LDRD project to port ATSE to a RISC-V processor. (Appendix J, target CSSE-2)
- Prototyped with NextSilicon an initial set of capabilities in a customized Kokkos backend to allow better scalability and performance on Gen-2 hardware. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Support newly delivered Vanguard II and Vanguard Next systems as they arrive by providing software stacks (via ATSE containers and base operating system support) and evaluation (via benchmarking, debugging, and troubleshooting). (Appendix J, target CSSE-2)
- Demonstrate an HPC operating environment (leveraging CIQ Fuzzball as appropriate) that enables separation of software concerns at the container boundary, enabling simplified system management and portable high-performance containers. (Appendix J, target CSSE-2)
- Develop ATSE into a modern, extensible, container-first software base to enable modification and selective upgrading by existing apps teams while providing an always up-to-date starting point for new uses. (Appendix J, target CSSE-2)
- Work with NextSilicon to continue advancing and hardening Kokkos capabilities on Gen-2 hardware. (Appendix J, target CSSE-2)

Systems Benchmarking (SNL)

The CSSE Systems Benchmarking project supports current systems benchmarking activities across the Trilabs, including dedicated performance analysis of CTS and ATS systems. These benchmarking activities can identify gaps in machine software ecosystems as they are maintained, and validate sustained performance over time. This project also feeds baselines for investigation in future ASC procurement activities, including upcoming ATS procurements, testbeds, prototypes, and AMT program investments. Overall, the Systems Benchmarking items are designed to specifically facilitate technical engagements across the Trilabs and DOE in shared benchmarking, analysis, and new workload studies.

Accomplishments in FY24:

- Investigated performance differences across Intel Sapphire Rapids (SPR) deployments across ATS-3 and CTS-2 resources and identified areas and sources of noise are impacting job performance. (Appendix J, target CSSE-2)
- Baseline shared Trilab workloads for an AMT app analysis and tracked contractual updates for the AMT vendor engagements. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Investigate novel ways to create representative mini apps and code validation suites that can help communicate ASC needs to the wider vendor industry. (Appendix J, target CSSE-2)
- Assist in the ATS-5 technical evaluation as needed and help upgrade existing benchmarking apparatus for use in ATS-6. (Appendix J, target CSSE-2)

Software Sustainability (SNL)

The CSSE Software Sustainability project supports the sustainment of software capabilities developed under the Exascale Computing Project.

Accomplishments in FY24:

- Supported software sustainability of the partitioned communication interface in Open MPI by creating and contributing unit tests, identifying and fixing functionality issues, and designing and developing a higher performance partitioned communication module. (Appendix J, target CSSE-2)
- Defined the new NNSA Integration effort with input from Sandia, LANL, and LLNL. Participating in the broader software stewardship organization effort around ecosystem integration . (Appendix J, target CSSE-2)
- Added new LAMMPS capabilities to expose more parallelism for small systems and Kokkos support for "tiled" halo exchange communication and hybrid bond interactions. Currently adding support for an alternative method to pack/unpack MPI buffers to reduce launch latency (in progress). Collaborated with Cerebras for molecular dynamics work. (Appendix J, target CSSE-2)
- Proposed a three-pronged effort between late FY24 – FY27 focused on (1) metrics collection for DevOps tools and processes, (2) approaches for identifying and quantifying technical debt, and (3) research strategies to reduce identified technical debt.

Planned Activities in FY25:

- Will provide continued support of Open MPI through partitioned communication by optimization, fixing any functionality issues that arise, and exploring extensions (including a more general interface, and partitioned collectives) and extending the remote memory access interface to support more HPC application use cases.
- Determine NNSA software project needs more granularly and work towards increased participation in the scientific software ecosystem, as well as capturing greater benefits from the ecosystem for the benefit of NNSA codes.
- Improve the performance of the LAMMPS molecular dynamics code on exascale hardware and sustain development efforts to meet the needs of internal, tri-lab, and external customers.

- Develop novel AI-assisted approaches to reduce existing and future ASC software technical debt for decades to come.

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 FY24:

- Designed and developed Continuous benchmarking infrastructure to effectively and reproducibly communicate ASC workload requirements to vendors.
- Started refreshing the list of procurement benchmarks, defined experiments for heterogeneous architectures.
- 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 FY25 and beyond funding. Focus on monitoring progress relative to proposed tasks and schedule, applicability, usability, and overall impact on LLNL applications.
- Provided feedback to Micron on their CXL testbed and evaluated their data mover interface.
- Used experimental CXL emulation testbed set up in LC Data Sciences Testbed to evaluate impact of using CXL memory for scientific applications
- Participated in the EES2 memory working group to characterize algorithm and application opportunities made possible by CXL. Report is in final draft form.
- In collaboration with the LBNL P38 Mosaic project, integrated the zhw floating point encoder into the Mosaic System on Chip framework. (Decoder is in progress).

Planned Activities in FY25:

- Continue development of Continuous benchmarking infrastructure to enable CI of procurement benchmarks across ASC systems and cloud resources.
- Complete development of the suite of procurement benchmarks, with reproducible multi-platform experiment definitions, performance analysis, and documentation, to represent the tri-lab ASC workload.
- Continue to participate in technical oversight of existing and new AMT contracts. Provide feedback on relative benefits of new memory technologies for LLNL applications.
- Continue CXL evaluations as hardware becomes available.
- Research into near memory/network computing with programmable gather/scatter near embedded ARM cores on NIC.
- Research into Content Addressable Memories and Associative Memories to accelerate similarity search.

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 FY24:

- Attended meetings with Micron and partner labs. Reviewed for DoE the reports, documentations, software, and hardware-architecture information provided by Micron as part of their engagement with the tri-labs. As part of this engagement, latest versions of (i) Transaction Engine (ii) Data Mover (iii) RTL for the accelerator's FPGA prototype have been installed and tested at the testbed "Foxcreek" node.
- Made Micron PCIe+CXL accelerator testbed and NDGSim simulator accessible on request to researchers at LANL for benchmarking. RTL for the hardware architecture is available, and a Xilinx Vivado license has been procured so as to compile and target the platform with custom IP inclusions as required.
- Successfully targeted Skywater's 130nm node open-source PDK with LANL-developed atomistic molecular dynamics accelerator architecture (MD-ASIC).

Planned Activities in FY25:

- Complete report on new hardware technologies (including AI hardware) algorithms, and codes using an agile / rapid response approach.
- Continue interactions with vendors and the tri-lab on advanced technologies (including AMT) providing timely feedback and access to workloads/benchmarks of interest.

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.

- Re-initiated relationships with P38 collaborators to investigate novel Government chip & system design capabilities. (Appendix J, target CSSE-1)
- Advanced the state of Molecular Dynamics strong-scaling on a new wafer-scale engine system. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Implement performance & functionality analysis workflows for custom chip design. (Appendix J, target CSSE-1)
- Outline and prototype connectivity between open-source chip design tools and SST to facilitate performance modeling and low-level chip design generation from a single source description. (Appendix J, target CSSE-1)

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 FY24:

- Management, execution, and verification of AMT FY21 and AMT FY22 Project contract deliverables.
- Release of AMT FY23 Project RFI and conducted Tri-lab evaluation of responses and preliminary selection.
- Begin SOW development and contract negotiations for selected partners. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Complete SOW negotiations and contract placement for selected partners. (Appendix J, target CSSE-1)
- Manage all rounds of AMT contract deliverables. (Appendix J, target CSSE-1)
- Reinvigorate coordination across Tri-labs for investigating advanced memory technologies. (Appendix J, target CSSE-1)

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 FY24:

- Evaluated the software stack including simulator and compiler toolchain of a vendor-provided dataflow accelerator and started initial mod/sim workload porting. (Appendix J, target CSSE-1)
- Investigated and tracked several potential future Vanguard system architectures to determine fit and viability in NNSA computing areas. (Appendix J, target CSSE-1)

- Generated an initial viable Kokkos back-end for vendor-provided dataflow accelerator API. (Appendix J, target CSSE-1)

Planned Activities in FY25:

- Deploy initial Vanguard II dataflow accelerator hardware and investigate system performance. (Appendix J, target CSSE-1)
- Improve the Kokkos back-end for dataflow accelerators and port more mission computing codes. (Appendix J, target CSSE-1)
- Develop initial framework for Vanguard III partnership. (Appendix J, target CSSE-1)

Projects for the Data Science and Artificial Intelligence Product (WBS 1.2.3.5.10)

The NNSA is one of the largest producers of scientific data in the world, developing large quantities of experimental, surveillance and design data as well as vast quantities of outputs from its exquisite modeling and simulation capabilities. Accessing, curating, and analyzing such vast volumes of data is intractable without scalable and capable data analytics capabilities. Similarly, artificial intelligence and machine learning has the potential to profoundly impact the way in which the NNSA approaches its science-based stockpile stewardship mission in the future. This work product develops scalable, high-performance data analysis and artificial intelligence methods which can be used in a demanding scientific domain. Activities cover a range of scopes including, but not limited to, the development of novel data science and artificial intelligence methods for exceptionally large datasets as well as the ability to ingest, curate, label and process such data for use in artificial intelligence and machine learning tools on ASC's computational platforms. In addition, dedicated AI platforms are procured, AI models trained, and AI focused R&D is performed, in this work product.

Tri-lab AI Infrastructure Project (LANL, LLNL, SNL)

The ASC program will procure dedicated AI resources and infrastructure to support the efficient and timely training and use of AI models. Additionally, the program will deploy inference-optimized resources, where applicable, to improve model access latencies. AI resources will be designed and optimized so that ASC AI teams working on various research, development and evaluation activities have sufficient scale and performance to prevent training and inference from being significant bottlenecks to program activities. The project will work with the AI vendor community to evaluate the feasibility of deploying AI resources in the appropriate computing environments to begin working on the near-term, actionable activities identified in the AI4ND Strategy. Activities will focus on applying AI models, tools, and techniques to reduce cost and schedule in the discovery, design optimization, manufacturing and certification, and deployment and surveillance phases of a nuclear warhead system.

Accomplishments in FY24:

- LLNL procured dedicated AI system for tri-lab use (Intel x86 CPUs with NVIDIA H100 GPUs) through the CTS-2 contract. (Appendix J, target CSSE-1.b)
- LANL procured a tri-lab dedicated AI system for deployment in the classified network. The system will be deployed in FY25. (Appendix J. target CSSE-1.b)
- SNL placed a contract with Cerebras to train a 70B parameter model of open, but ASC-relevant data sources.

- SNL placed a contact to purchase and deploy four wafers of a Cerebras system for installation at Sandia. SNL structured the contract with an expectation to expand the system to eight wafers in the future.

Planned Activities in FY25:

- Deployed dedicated AI system at LLNL on the restricted zone (RZ) network. (Appendix J, target CSSE-1.b) (LLNL)
- Transferring dedicated AI system at LLNL to classified network (SCF). (Appendix J, target CSSE-1.b) (LLNL)
- Deploy dedicated AI systems into the classified network at Los Alamos. (Appendix J, target CSSE-1.b) (LANL)
- Deploy Cerebras AI system into the restricted network at Sandia. (Appendix J, target CSSE-1.b) (SNL)
- Deploy and evaluate the performance of the Molecular Dynamics and HPCG mini-applications on the four-wafer Cerebras system installed at Sandia. (SNL)

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 FY24:

- LLNL contributed to the Tri-Lab LLM L2 Milestone, with an exploration of fine-tuning LLMs with information related to LLNL codes – both source code and documentation. Using open-source models, we were able to explore fine-tuning and inference on HPC hardware.
- The Hermit model for atomic physics calculations (including absorption, emissivity, and EOS for NLTE) has been fully integrated into Hydra for ICF calculations, with a workflow for retraining when out-of-scope inference calls are made during the simulation. These capabilities have been demonstrated on El Capitan EAS hardware.
- The ML-based Material Interface Reconstruction method has shown to be generalizable across multiple problems and resolutions, with higher overall accuracy than PLIC or Equi-Z. This algorithm is now being incorporated into a linking technique.

Planned Activities in FY25:

- HPC Support for AI Tools
    - Work with teams across the lab to understand how they are running AI workflows on LLNL HPC systems, what is working, and what needs to be improved. Develop a plan to address issues and make these teams more productive.
    - Profile AI workflow performance on AMD chips and determine how to make improvements as needed.
    - Support AI environments, including standardizes builds of tools, like PyTorch, across HPC systems, containers, etc.
    - Stand up the SCF PDS environment with NVIDIA GPUs, followed by the RZ environment when hardware is available, and support early test/dev with code teams.
- Integration of AI techniques into HPC code development (CSSE-3.d)
    - Demonstrate the Hermit model for atomic physics across a variety of input parameters, with a credibility assessment and retraining, on El Capitan hardware. (CSSE-2.a)
        - Compare Hydra results using Hermit and full-physics calculations for accuracy and performance.
        - Compare multiple approaches to the credibility assessment or the neural network.
    - Demonstrate the ML-based Material Interface Reconstruction algorithm on problems relevant to linking. Determine if additional functionality is needed to deploy an initial capability.
    - Continue exploration of LLMs for HPC code development, such as porting to accelerators and programming language translation
    - Evaluate neural network performance for multi-scale hydrodynamics coupling across parameter variations

Data Science Infrastructure (LANL)

The Data Science Infrastructure (DSI) develops capabilities for meta-data collection, storage and search project. The DSI project focuses on gathering meta-data from existing scientific activities, storing this meta-data and (pointers to) associated data in a database, and 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 FY24:

- Released DSI v1.0 with read, write, query/find, and data movement capabilities.
- Integrated DSI with the LANL HPC CONDUIT data movement capability on unclassified LANL HPC resources.
- Integrated DSI with National Security Data Solution (NSDS) via the CMF-NSDS API.

Planned Activities in FY25:

- Deliver DSI capabilities to support L1 Milestone.
- Continue integration of DSI with HPC data movement capabilities, especially CONDUIT on classified LANL HPC resources.
- Develop workflow patterns and examples for ASC AI/ML clients.
- Explore tri-lab data exchange standards and integration with LLNL and Sandia HPC and data management capabilities.

Advanced Machine Learning (LANL)

The Advanced Machine Learning project has mission-focused efforts in collaboration with PEM and V&V. 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 FY24:

- Evaluated and advanced intelligent aids for programming and software engineering in support of nuclear security.
- Developed and collected datasets that can be used to evaluate the performance of large language models for scientific computing tasks.
- Developed and evaluated techniques and models for translating code from Fortran to C++ and from CPU to GPU.
- Co-authored "Assessment of Data-Management Infrastructure Needs for Production Use of Advanced Machine Learning and Artificial Intelligence: Tri-Lab Level II Milestone (8554)", August 2023, https://www.osti.gov/biblio/2212844.
- Delivered database schemas and demonstrated metadata and data management for two ASC AI/ML projects: large language model (LLM)-based transpilers for Fortran-to-C++, and AI/ML model-generated strength model parameterization.

- Provided AI/ML workflow use cases and data infrastructure requirements to the Data Science Infrastructure (DSI) project.

Planned Activities in FY25:

- Write a paper evaluating large language models combined with few-shot learning for code translation.
- Build a dataset of synthetically generated code pairs that can be used to finetune a large language model to translate code from on language to another (targeting Fortran->C++ and Fortran->PyTorch).
- Develop a workflow for fine-tuning large language models on the aforementioned dataset using preference optimization techniques.
- Develop a tool to enable semantic search of ASC codes instead of string matching.
- Deliver data and metadata management for two ASC AI/ML project workflows such as large language model-based transpilers and radiograph inverse problems.

Artificial Intelligence for Nuclear Deterrence (SNL)

The Artificial Intelligence for Nuclear Deterrence project has mission-focused efforts in collaboration with all ASC sub-programs and in alignment with the AI4ND strategy to address AI/ML for discovery, design optimization; manufacturing and certification; and deployment and surveillance. Sandia's portfolio of projects addresses some of the key challenges identified in the AI4ND strategy that include, but are not limited to, physics-informed machine learning; verification, validation, uncertainty quantification, and AI trustworthiness; AI/ML for accelerated model development; application and development of ASC-relevant foundation models; and contributions to an enterprise-wide data infrastructure. In addition, Sandia intends to work closely with our ND program projects as well as production agencies through strategic collaborative efforts.

Accomplishments in FY24:

- Modified the Physics Informed Neural Network code to allow for utilization of experimental data and began developing machine-learned surrogate models for additively manufactured compression pads. (Appendix J, target CSSE-2)
- Completed report "A Maturity Model Framework for Scientific Machine Learning" that outlines the analysis criteria for each element of analysis essential for scientific machine learning credibility. Current research gaps and implementation challenges will be documented. (Appendix J, target CSSE-2)
- Incorporated a *Reduce for Simulation* tool into the Cubit 16.16 software release that uses ML to aid in preparing models for simulation, automatically identifying fasteners, holes, slots, and thin volumes. These tools were specifically designed to support structural dynamics (SD) and electromagnetics (EM) applications. (Appendix J, target CSSE-2)

- Made progress evaluating Kokkos-based code-generation using open-source and NVIDIA models using NVIDIA's NeMo GenAI software framework. Accomplishments include extensive quality metrics for C++ code generation, a Retrieval Augmented Generation pipeline to evaluate Sandia code repositories, and preliminary fine-tuned models trained from our own repositories. (Appendix J, target CSSE-2)
- Began deploying preliminary services for LLM capabilities for knowledge management with a particular emphasis on ND content. The ATLAS team now offers micro-services for text extraction, embeddings, cross-encoders, vector databases. (Appendix J, target CSSE-2)
- Developed hybrid and advanced RAG workflows with a focus on knowledge-transfer for exploding bridge wires and has started exploring knowledge graphs to provide reasoning in complex LLM queries. The IRIS/ATAS work provides a foundation for the deployment of LLM capabilities in sensitive and classified environments. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Develop Physics-Informed Machine Learning material models for solid mechanics, with emphasis on technologies for multi-fidelity fusion of experimental and simulation data. (Appendix J, CSSE-3.d)
- Perform R&D for inline credibility assessments of AI models. (Appendix J, target CSSE-2)
- Advance capabilities that use machine learning for Physics-Relevant Geometry Simplification in support of Sandia's Accelerated Model Development effort. (Appendix J, CSSE-3.d)
- Perform R&D to evaluate, develop, and deploy ASC-relevant foundation models for knowledge transfer, software modernization, and training on emerging AI architectures. (Appendix J, target CSSE-2)
- Establish guidelines for developing data-management plans for ASC projects with an expectation to register curated datasets into a common catalog. (Appendix J, CSSE-3.d)

Projects for the Compilers, Programming Models, and Runtimes Product (WBS 1.2.3.5.11)

The ASC program has developed many hundreds of millions of lines of source code since its inception. Such code can only be used efficiently on the computational platforms that ASC fields, if it is written to utilize stable, scalable, and performant programming models and if it can be compiled by high-performance optimizing compiler toolchains. As programming languages and standards have risen in complexity over the last two decades, the ASC program has increasingly needed to focus resources on ensuring that these critical tools work correctly and can optimize the sophisticated code required for stockpile stewardship. Once code is compiled, a high-performance and portable runtime is required to permit such code to run on ASC's high-performance computers. This work product covers development and R&D activities to ensure that ASC integrated codes have consistent and enduring access to stable and robust programming environments and compilation tools across its Advanced Technology, Commodity Technology, and Advanced Architecture Prototype systems.

Production Compiler Support and LLVM/Flang Development (LLNL)

The LLVM compiler ecosystem is the foundation for all modern HPC vendor compilers and provides an opportunity to influence, improve, and guide toolchains effectively. Through our development of LLVM, especially LLVM/Flang, we mitigate the risk of vendor compilers and, in the future, might allow for our own in-house compiler toolchain that is specialized for our codes and architectures. At the same time, this project will work with vendors to ensure their toolchains can be deployed on ASC machines to provide the users with a streamlined compilation experience that yields correct and performant codes. As problems arise, we work with vendors to find sustainable solutions that can be shared via LLVM to ensure features and experiences across compilers converge where appropriate.

Accomplishments in FY24:

- Installed and announced five AMD LLVM/new-flang point releases with El Capitan MPI support and AMD GPU support to allow Fortran ASC code teams to easily compile their codes and report issues specific to ASC codes. Several reproducers for blocking issues given to AMD that were fixed in later releases.
- AMD LLVM/new-flang's July 2024 release for the first time successfully compiled LLNL's Goulash GPU stress tests and yielded unexpectedly strong Fortran OpenMP GPU performance (achieving 82% of performance of the equivalent C++ OpenMP GPU tests).
- Helped AMD to integrate OpenMP offloading into LLVM/Flang such that SPMD kernels, including reductions, can utilize both AMD and NVIDIA devices fully. (Appendix J, target CSSE-2a)

- Implemented OpenMP SIMD parallelism for the GPU to allow for three dedicated levels of parallelism (user and Kokkos request). The OpenMP 6.0 “workdistribute” directive that allows offloading of Fortran array syntax has been prototyped. (Appendix J, target IC-1)

Planned Activities in FY25:

- Continue timely installation of AMD LLVM/new-flang releases with OpenMP GPU offloading support and continue to generate reproducers for ASC Fortran codes blocking issues and performance issues
- Identify and minimize performance differences between Flang-compiled code and established Fortran compilers, i.e., gfortran. Further improve the offloading support and performance. (Appendix J, target CSSE-2, IC-1)
- Develop compiler-based and ML-enhanced HPC/GPU centric heuristics and tooling to improve performance and provide actionable feedback to ASC users. (Appendix J, target CSSE-2, IC-1)
- Expand the LLVM/Offload project to target novel (ML) hardware and support existing languages such as CUDA and SYCL. This will provide interoperability between languages and portability across architectures to ASC codes. (Appendix J, target IC-1)

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.

Accomplishments in FY24:

- Flang saw significant growth in community engagement and functionality. This included a significant growth in testing functionality and core language features and bug fixes. Overall performance showed clear improvements.
- Added improved GPU-target code generation and runtime features for Kitsune. Observed up to 5-7x improvement over Kokkos and hand-coded CUDA for some use cases. Introduced advanced capabilities for dominance analysis. Initial reduction design complete and initial implementation complete for CPU targets with GPU support well underway.
- Continued work to explore mapping of computations (tasks) to heterogeneous architectures including modeling of code and hardware targets. Additional work exploring leveraging ML-based technologies for automatic mapping and modeling is underway.

Planned Activities in FY25:

- Increase Flang testing and evaluation, and complete initial deployment to the IC and PEM communities within ASC. Expand efforts to evaluate and improve the Flang tooling capabilities for various activities (e.g., Fortran-to-C++ code conversion). Continue outreach and engagement in support of the overall Flang and LLVM communities.
- Continue advanced compiler analysis and optimization for parallel constructs, including data-centric operations, advanced heterogenous architecture targets, mature and flexible reduction operations and code generation, and improvements to the toolchain for end-developer productivity. Broaden language support to include Fortran and subsets of Python for various IC, PEM, and other programmatic efforts.
- Continue exploration of auto-mapping and general task-based programming systems. Scope will include capabilities useful for general use cases (e.g., resiliency) and new use cases in machine learning and AI.

AMT Programming Models and Runtimes (SNL)

The DARMA project is exploring Asynchronous Many-Task (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.

The DARMA project supports the development of the DARMA toolkit, which includes: the core VT (Virtual Transport) runtime, a highly MPI-interoperable AMT interface; magistrate, a checkpointing and serialization library; the load balancing analysis framework (LBAF), a library implemented in Python for exploring new load balancing strategies and analyzing workloads in a quick prototyping environment; and VT-TV; a task visualization library for analyzing QOIs and visualizing them to better understand application workloads.

As part of our core activities, the DARMA team collaborates with application teams to understand their applications, integrate AMT tooling into their codebase to rebalance their workloads, and make their code more asynchronous and capable of overlapping work with communication to better utilize networks and achieve better performance.

Accomplishments in FY24:

- Developed a memory-aware distributed-memory load balancing algorithm capable of rebalancing tasks under tight memory constraints. We validated our implementation by reformulating the problem as a MILP and demonstrating that our fully distributed algorithm is within 2% of the optimal solution. (Appendix J, target CSSE-2)
- Effectively load balanced Gemma while keeping nearly 80% of each node's memory filled with the matrix using our distributed memory-aware load balancer. One example problem from vvtest with significant imbalance runs 2.3x faster during the matrix assembly. (Appendix J, target CSSE-2)
- Released several new versions of the DARMA toolkit with enhanced C++17 support and many bug fixes. The support of C++17 greatly improves the DARMA interfaces, making it much easier to create tasks and dispatch messages to handlers. (Appendix J, target CSSE-2)Developed a new visualization component of the DARMA toolkit (called VT-TV) in C++ to visualize tasks, communications, and shared memory blocks. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Enhance and generalize the machine-learning component required to predict task timings in Gemma to effectively load balance workloads. We plan to study and build a more general neural network model that can be applied across architectures. (Appendix J, target CSSE-2)
- Make our Gemma load balancing capability more production ready, testing it across more platforms and problems, and prepare it for merging onto their main branch. (Appendix J, target CSSE-2)
- Work toward making our new distributed load balancing algorithm effective for load balancing workloads with communication links and memory constraints for general problems using our reduced-order model. (Appendix J, target CSSE-2)
- Implement the VT-TV online load balancing capability so VT can be configured to visualize tasks and communications as an application is running. (Appendix J, target CSSE-2)
- Prepare the DARMA toolkit for C++20, which will enhance the user experience and code significantly. (Appendix J, target CSSE-2)

Kokkos (SNL)

The goal of the Kokkos Ecosystem is to provide performance portability for ASC application codes, allowing these codes to obtain~~ed~~ 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 FY24:

- Two significant Kokkos Ecosystem releases (4.2, 4.3) provided increased AMD GPU support in preparation for El Capitan (4.4 scheduled for summer). (Appendix J, target CSSE-2)
- C++26 Linear Algebra was voted into the C++26 standard. (Appendix J, target CSSE-2)
- Christian Trott worked with Todd Gamblin (LLNL) to establish the High Performance Software Foundation (HPSF). Kokkos joined HPSF as one of the inaugural members. (Appendix J, target CSSE-2)
- Implemented facilities in Kokkos for supporting host accessible memory for Grace-Hopper and to support unified memory for HIP AMD MI300. (Appendix J, target CSSE-2)
- Vanguard-II: Defined workplan with NextSilicon and evaluated their proposed NS-API extension for building a Kokkos backend, providing feedback throughout the year. Provided support for NextSilicon Kokkos backend. (Appendix J, target CSSE-2)
- Kokkos Kernels team delivered a novel Schur complement based approach for SPARC that improved solve phase performances by up to 3x for SPARC test

matrices. This has been integrated with SPARC and run on both MI250 and MI300 AMD GPUs. The improvements are particularly significant for GPU hardware since the Schur complement provides more parallelism to exploit. This will be particularly important in support of SPARC's future exascale milestone runs on El Capitan. (Appendix J, target CSSE-2)

- ODE integrators were released as part of Kokkos Kernels 4.3 release, (in support of Sierra ARIA). This included both explicit adaptive time integrators and implicit adaptive BDF integrators. The initial comparisons between ATS-2 and CTS-1 are promising, with 14-100x speedups for explicit time integrators and 4-6x speedups for the implicit time integrators. This is a capability delivered in support of the SIERRA Thermal team and their production GPU runs. (Appendix J, target CSSE-2)
- Kokkos Kernels team has started a new working group with AMD rocSPARSE team as part of El Capitan COE. Several performance issues have been identified, and several improvements to Kokkos Kernels rocSPARSE usage have been made. (Appendix J, target CSSE-2)

Planned Activities in FY25:

- Support ASC applications in deploying and optimizing their code base for El Capitan and CTS platforms. (Appendix J, target CSSE-2)
- Support vanguard program with preparing Kokkos for a potential platform deployment. (Appendix J, target CSSE-2)
- Evolve Kokkos interfaces to enable more generic programming via self-similar interfaces for execution policies. (Appendix J, target CSSE-2)
- Participate in the ISO C++ committee work to further the interests of Sandia and the HPC community in the evolution of the C++ standard. (Appendix J, target CSSE-2)
- Kokkos Kernels support applications readiness for El Capitan through performance profiling of application codes and optimization. (Appendix J, target CSSE-2)
- Kokkos Kernels new algorithm development for El Capitan and in general for all GPUs. (Appendix J, target CSSE-2)
- Development of new algorithms and programming techniques for new non-GPU platforms such as Vanguard and Cerebras. (Appendix J, target CSSE-2)
- Explore new machine learning approaches for developing a Kokkos-GNN (graph neural network) using Kokkos Kernels sparse matrix multiplication and auto differentiation work (for new machine learning work). (Appendix J, target CSSE-2)

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

This subprogram provides two critical enablers for the effective use of ASC tri-lab computing resources: 1) physical facility and operational support for reliable, cross-lab production computing and storage environments, and 2) a suite of user services. The scope of the facility operations includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities, and collaborations. FOUS also covers physical space, power and other utility infrastructure, and local area network (LAN)/wide area network (WAN) networking for local and remote access, as well as requisite system administration, cyber-security, and operations services for ongoing support and addressing system problems. Industrial and academic collaborations are an important part of this subprogram.

Accomplishments

ASC accomplishments from quarter 4, fiscal year 2023, and through quarter 3, fiscal year 2024, are reflected below for the FOUS subprogram.

- Took delivery of all El Capitan hardware and completed physical installation of the system, connection to power and cooling infrastructure, network cabling, and node installation. (LLNL)
- Continued stable production operation of Sierra (ATS-2) past the hardware warranty and vendor support period in support of tri-lab Advanced Technology Computing Campaigns. (LLNL)
- Delivered CTS-2 systems Bengal, RZHound, and Dane to production ASC customers (LLNL)
- Completed TOSS 3 to TOSS 4 conversions across the entire classified environment (LLNL)
- Provided 24x7x365 operational support for all production HPC resources and maintained consistently high system utilization levels and customer productivity. (LLNL)
- Provided 24x7 HPC operations of ASC systems and infrastructure to meet ATCC goals. (LANL)
- Preparation for start of construction SCC Electrical Upgrade design and procurement of long-lead infrastructure in preparation for ATS-5. (LANL)
- Achieved production readiness for Crossroads (ATS-3) through platform acceptance, contractual performance, and integration with a continuous support environment. (LANL)

- Production operation of Crossroads (ATS-3) and CTS-2 systems Rocinante and Tycho peaking with over 94% utilization and 99% system availability. (LANL)
- Began decommissioning activities for Trinity, Snow, Fire, Ice, and Cyclone. (LANL)
- Successfully deployed the CTS-2 cluster Amber into production completing FY24 ASC L2 milestone # 8640. (SNL)
- ASC ATS-3 Application Readiness system "Tachi" formally accepted. System released to users. Software environment expanded to meet code team needs. Major upgrade in coordination with vendor and LANL. (SNL)
- Completed additional 4MW power minor construction project to 725East data center. (SNL)
- ASC DevOps has successfully developed sustainable processes to easily deploy code environments across Sandia's ASC systems for both analysts and developers. This includes adoption of the ASC Unified Environment (AUE), improved build and test farm infrastructure, and improved metrics. The AUE is now being heavily used across many of the ASC code teams. This Sandia-wide collaboration across ASC is continuing and more efforts to increase efficiencies have been identified and are being executed. (SNL)
- Demonstrated deployment of LDMS as a common monitoring infrastructure in a hybrid configuration across on-prem production HPC systems and AWS cloud resources with run-time telemetry for all systems actively streamed to a common data analysis cluster sited at SNL. This demonstrates the ability for run time, locally accessed performance monitoring of applications in the cloud and bursting to the Cloud. (SNL)

Level 2 Milestone Descriptions

Milestone (ID#TBD): Converged HPC Center-Wide Data Analytics		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 8/31/25		
ASC WBS Subprogram: FOUS, CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC FOUS and CSSE programs		
Program Plan Target: Appendix J, targets FOUS-1, FOUS-5, and CSSE-1.a		
Description: This effort will build and demonstrate a converged data repository for center-wide monitoring. The repository will bring together multiple data sources to enable correlation and analysis across these historically disparate data sets. These correlations will enable insights into how the behavior and interactions of various disparate changes within the HPC center impact utilization and reliability. Data sources under consideration include benchmarking performance metrics, facilities operations data, and application-provided metrics.		
Completion Criteria: A program review will be conducted that shows how key data and metrics of interest to the ASC program were integrated into the central data repository, the capability for new analytics, and relevant documentation.		
Customer: LC staff, ASC customers, ATS system procurement teams		
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: LLNL FOUS, CSSE, and IC teams		

Milestone (ID#TBD): SCC Electrical Upgrade Commissioning Complete		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 9/30/25		
ASC WBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-3		
Description: Commissioning is complete for minor construction SCC Electrical Upgrade (SEU) project to prepare for increased facility electrical power distribution for ATS-5.		
Completion Criteria: Functional testing is complete for equipment of SCC Electrical Upgrade project.		
Customer: ASC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: FOUS		

Milestone (ID#TBD): El Dorado Standup and Deployment		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 9/30/25		
ASC WBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target FOUS-1.b		
Description: Standup El Dorado and integrate it into Sandia's IHPC network with the system environment sufficiently consistent El Capitan to begin use as an Application Readiness resource for codes to be deployed on El Capitan. Complete contract acceptance tests.		
Completion Criteria: El Dorado is functional and integrated onto Sandia's IHPC network. Basic operational software is functional. The configuration and environment are largely consistent with El Capitan, with the exception of site-specific considerations. User environment is sufficient for users to begin use. System is accepted per contract requirements.		
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		

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 FY24:

- Provided procurement support, contract management, and program planning. (Appendix J, targets FOUS-1 and FOUS-5)
- Managed the PSAAP III program for new Academic Alliance Centers and prepared for PSAAP IV request for proposals.
- Supported contracting and procurement for CORAL-2 and CTS-2 systems. (Appendix J, targets FOUS-1 and FOUS-5)

Planned Activities in FY25:

- Continue FY25 procurement support, contract management, and program planning. (Appendix J, targets FOUS-1 and FOUS-5)
- Continue to manage the PSAAP III program for Academic Alliance Centers, and prepare for PSAAP IV proposal evaluations.
- Continue contract management and procurement support for CORAL-2 and CTS-2 systems. (Appendix J, targets FOUS-1 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 FY24:

- Organized and hosted the ASC PI meeting.
- Provided LANL detailee support to HQ.
- Supported the programmatic needs of the PSAAP III & IV program.

Planned Activities in FY25:

- Provide LANL detailee support to HQ.
- Support Metropolis postdoc program.
- Support the programmatic needs of the PSAAP IV program.

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 FY24:

- Provided programmatic support for the PSAAP III and participated in preparation for the PSAAP IV proposal process.
- Participated in organization of the ASC PI Meeting that was held in May 2024 at LANL, partnering closely with HQ, LANL, and LLNL.
- Pursued partnership opportunities with DOE/NNSA and DOE/SC laboratories to support the National AI/ML strategy.

Planned Activities in FY25:

- Continue to support ASC integration and partnership with stockpile and modernization teams.
- Plan and organize Predictive Engineering Science Panel (PESP) review.
- Continue supporting ASC integration and partnership with stockpile and modernization teams.
- Continue to provide programmatic support for the PSAAP III program. Organize and host PSAAP IV Proposal Review, Selection, and Kick-off meetings.

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 FY24:

- Fully deployed center-wide monitoring solution based on Elastic Stack as a replacement for Splunk. (Appendix J, targets FOUS-1.b and FOUS-5)
- Migrated LC web and collaboration tools to containers to allow for rapid updates and reduced downtime.
- Completed phase 1 deployment of LC infrastructure container storage to provide fault tolerance, and disaster resiliency for building outages due to power or cooling failures. (Appendix J, targets FOUS-1.b and FOUS-5)
- Completed phase 1 deployment of LC's converged infrastructure hardware allowing for increased network speed and bandwidth, larger memory, and higher CPU based applications supporting LC infrastructure. (Appendix J, targets FOUS-1.b and FOUS-5)
- Deployed VAST NFS environment on SCF to allow for NFS high-speed scratch storage. (Appendix J, targets FOUS-1.b and FOUS-5)
- Deployed VNC replacement servers on the OCF to allow for higher capacity workloads enhancing user productivity and overall use of the service.
- Completed integration of LC OAUTH across all LC web services, replacing the commercial Atlassian Crowd authentication platform and allowing remote tri-lab users to utilize their local credentials to access LC's web services.
- Deployed a new Identity Management system, replacing a legacy commercial solution with a powerful and flexible role-based model for managing user accounts, groups, and access to resources.

- Deployed Starfish on the OCF to efficiently index and query metadata across all production file systems and archives including hundreds of petabytes of data and billions of files.

Planned Activities in FY25:

- Integrate El Capitan and Tuolumne CORAL-2 HPC systems into the production computing environment. (Appendix J, target FOUS-1.b)
- Complete phase 2 deployment of LC infrastructure container storage to provide fault tolerance, and disaster resiliency for building outages due to power or cooling failures. (Appendix J, targets FOUS-1.b and FOUS-5)
- Complete phase 2 deployment of LC's converged infrastructure hardware allowing for increased network speed and bandwidth, larger memory, and higher CPU based applications supporting LC infrastructure. (Appendix J, targets FOUS-1.a and FOUS-5)
- Deploy VNC replacement servers on the SCF 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 J, targets FOUS-1.b and FOUS-5)
- Deploy Starfish metadata indexing and query engine in the classified computing environment.
- Provide system integration, maintenance, and user support for dedicated AI system at LLNL procured through the Tri-lab AI Infrastructure Project. (Appendix J, target FOUS-5)
- Leverage ACI network architecture to directly integrate AI compute resources with the PDS container orchestration environment.
- Refresh PDS hardware environments for improved resiliency and performance.
- Enhance operational support posture for AI/ML/Data science services and resources. Perform a gap analysis and address identified risks in key focus areas including monitoring tools, change management processes, regression testing, service level agreements, training and documentation, and system availability.

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 FY24:

- Provided ongoing user support through hotline operations, documentation, and training while in minimal-normal operations. (Appendix J, targets FOUS-1 and FOUS-5)
- Developed documentation and training materials in the use of El Capitan/CORAL-2 and the CTS-2 class clusters. (Appendix J, targets FOUS-1 and FOUS-5)
- Provided support to the 9 ASC PSAAP 3 Alliance centers, successfully porting codes to TOSS 4 on the Quartz CTS-1 and Dane CTS-2 clusters. (Appendix J, target FOUS-5)

Planned Activities in FY25:

- Continue providing support through hotline operations, documentation, and training with a focus on assisting users porting to El Capitan/CORAL-2 and CTS-2 systems. (Appendix J, targets FOUS-1 and FOUS-5)
- Continue development of documentation and training materials in the use of El Capitan/CORAL-2 and the CTS-2 class clusters. (Appendix J, targets FOUS-1 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/CORAL-2 and CTS-2 architectures. (Appendix J, target FOUS-1)

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 FY24:

- Completion of Open Compute Facility (OCF) Network Re-architecture project – Migrated the remaining production workloads over to new IP address ranges for improved security segmentation.
- DisCom IHPC 100G Network Upgrade – Completed Phase-1 procurement of required hardware. (Appendix J, targets FOUS-1, FOUS-5)
- Procured and began testing of Apposite appliance to establish network telemetry and analytics capabilities for DisCom to enable 100G network tuning. (Appendix J, targets FOUS-1, FOUS-5)
- Completed integration of ATS clusters El Cap, Tuolumne, and rzAdams into Open Compute Facility (OCF) network. (Appendix J, target FOUS-1.b)
- Completed Phase-1 of Fiber Infrastructure Augmentation project. Increased fiber infrastructure capacity between Livermore Computing facilities (B451, B453, B654), allowing for continued growth and support of current & future projects.

Planned Activities in FY25:

- Begin planning and migration of Secure Compute Facility (SCF) Network Re-architecture project for improved security segmentation of critical services.
- Complete Phase-2 of Fiber Infrastructure Augmentation project. Increase fiber capacity within each Livermore Computing facility datacenter network locations, allowing for continued growth and support of current & future projects.
- DisCom IHPC 100G Network Upgrade – Complete Phase-2 and deploy new 100G encryptors at each tri-lab location.
- Complete project design and develop construction acquisition strategy for the B453 Sierra Retirement project.
- Purchase air handlers and switchgear for spare inventory to improve fault resiliency across the computing center.
- Execute site preparation projects in B-654 for upcoming system acquisitions.

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 FY24:

- Started decommissioning Trinity (ATS-1) and CTS-1 systems Snow, Fire, and Ice. (Appendix J, target FOUS-5)
- Operation of production CTS-1 Cyclone system. (Appendix J, target FOUS-5)
- Enhanced of utilization and availability of AT and CT systems. (Appendix J, targets FOUS-1.a and FOUS-5)

Planned Activities in FY25:

- Operation of production systems to include Rocinante, Chicoma, Tycho, and Crossroads. (Appendix J, targets FOUS-1.a and FOUS-5)
- GitLab Continuous Integration support on production platforms for ASC code-teams. (Appendix J, target FOUS-1.a)
- Explore and evaluate new artificial intelligence (AI4ND) cybersecurity tools and approaches for production ASC systems.
- Configure and operate new artificial intelligence (AI4ND) training and user systems.

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 FY24:

- Provided 24/7/365 Tier-1 support and operations of compute, network, and file system resources. (Appendix J, targets FOUS-1.a and FOUS-5)
- Performed continuous hardware maintenance, assessment of system life spans, and associated decommissioning activities for ASC computing equipment. (Appendix J, target FOUS-5)
- Supported onsite immersive training classroom for ASC users.
- Managed CREM from creation to destruction and disposition for storage systems supporting ASC simulations.

Planned Activities in FY25:

- Provide 24/7/365 Tier-1 support and operations of compute, network, and file system resources. (Appendix J, targets FOUS-1.a and FOUS-5)
- Support onsite immersive training classroom for ASC users.
- Continue developing Tier-1 support for ASC computing platforms.

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 FY24:

- Began facility preparation for hosting the ATS-5 system at the SCC. (Appendix J, target FOUS-3)
- Installed initial equipment for electrical operations (chiller plant economizer) and started design for cooling tower agitators.
- Operation of Crossroads and CTS systems for production use supporting ASC simulations. (Appendix J, targets FOUS-1.a and FOUS-5)

Planned Activities in FY25:

- Start construction for hosting the ATS-5 system at the SCC. (Appendix J, target FOUS-3)
- Operation of Crossroads and CTS systems for production use supporting ASC simulations. (Appendix J, targets FOUS-1.b and 5)
- Continue EPCU line item coordination. (Appendix J, target FOUS-3)
- Investigate increasing the efficiency of the facility warm-water cooling loop at the SCC (20 degree delta t).

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 FY24:

- Prototyped GUFI and supported production CONDUIT data management services to support ASC simulation workflows.
- Deployed new 15PB restricted enclave campaign storage.
- Supported the RCE working groups with parallel infrastructure required for tri-lab systems. (Appendix J, target FOUS-1.a)

Planned Activities in FY25:

- Continue to upgrade RCE capabilities for tri-lab users. (Appendix J, target FOUS-1.a)
- Deploy new filesystem infrastructure to expand capacity.
- Replace IB backbone with ethernet.
- Provide workflow and storage infrastructure for artificial intelligence (AI4ND).

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 FY24:

- Participated in RCE working groups for tools and techniques that improve ASC user productivity on tri-lab systems. (Appendix J, target FOUS-1.a)
- Continued service for production visualization of ASC simulation results supporting stockpile stewardship. (Appendix J, targets FOUS-1.a and FOUS-5)
- Managed and resolved of ASC user challenges on infrastructure and platforms to accomplish Annual Assessment Report and ASC milestones in other program elements. (Appendix J, targets FOUS-1.a and FOUS-5)

Planned Activities in FY25:

- Develop new continuous testing framework to avoid system regression and maintain system performance.
- Management and resolution of ASC user challenges on infrastructure and platforms to accomplish Annual Assessment Report and ASC milestones in other program elements. (Appendix J, targets FOUS-1.a and FOUS-5)

- Adopt to new artificial intelligence (AI4ND) user workflows.
- Explore AI/ML use cases to improve internal HPC operations.

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 FY24:

- Improved archival storage capabilities by importing existing NSRC data supporting ASC verification and validation.
- Continued development of production functionality for Marchive to enable next generation digital archive capabilities for ASC users.
- Provided top-tier diagnostic and troubleshooting support of archival storage for ASC users.

Planned Activities in FY25:

- Provide top-tier diagnostic and troubleshooting support of archival storage for ASC users.
- Formalize project management schedule and deliverables for Marchive development.

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 FY24:

- Successfully deployed the CTS-2 cluster Amber into production completing FY24 ASC L2 milestone # 8640. (Appendix J, target FOUS-5)
- Completed TOSS4 installation on legacy HPC systems in all four security environments. (Appendix J, target FOUS-5)
- Specified requirements and started procurement process for the 2nd ASC funded CTS-2 system. (Appendix J, target FOUS-5)
- Applied knowledge from ASC CTS-2 Amber cluster deployment to the deployment of non-ASC CTS-2 clusters including Flight/Condo (institutional), Hops (institutional GPU) and Dark Mead (NA-82). (Appendix J, target FOUS-5)
- Early deployment of the NA-82 funded 6SU CTS-2 based cluster (DarkMead) and associated supporting infrastructure. (Special Projects, Appendix J, target FOUS-5)
- Started retirement of legacy "Dark Bridge" TLCC2 HPC system (including EoL filesystems) and migrate user operations to next-generation platform. (Special Projects, Appendix J, target FOUS-5)
- Completed HPSS hardware/OS/software upgrades across all environments to update to RHEL8 and HPSS v9.3. (Appendix J, target FOUS-5)
- ASC ATS-3 Application Readiness system "Tachi" formally accepted. System released to users. Software environment expanded to meet code team needs. Major upgrade in coordination with vendor and LANL.
- Enhanced containerized workflow support on the Vanguard II testbed clusters via the Fuzzball workflow scheduler, in collaboration with CSSE and Ctrl-IQ. This work included codesign, development, and deployment of support for multi-container-image workflows, enhanced storage support and initial integration with a local WekaIO storage system, and improved operational capabilities based on our system experience. (Appendix J, targets FOUS-2, CSSE-1, and CSSE-2)
- Collaborated with LANL and SPARC teams to ensure SPARC performance and demonstrate system usability, as part of the Crossroads production readiness L2 milestone. (Appendix J, target FOUS-1.a)
- Released three new LDMS versions (4.4.[1-3]) into TOSS which provided new sampling capabilities for NVIDIA GPUs, AMD GPUs, and Slingshot switches, to address DOE architectures. These versions were deployed across the SRN CAPVIZ systems for monitoring systems and applications. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Enhanced TOSS LDMS packages to enable direct GMI-based (SNL configuration management) use from TOSS repositories, simplifying installation on SNL CTS systems. (Appendix J, target FOUS-5 and CSSE-2)
- Multiple successful LDMS and AppSysFusion education and outreach activities: LDMS User Group Conference (LDMSCON) held, and tutorials and/or

demonstrations of AppSysFusion analysis creation and dashboard use given to system administrators and code analysts. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)

- Spawned several best practices subgroups from the LDMS User Group to address current future needs of the DOE large-scale systems and other user base. These groups drove Ansible deployment strategies, YAML configuration development, among others, with plans and recommendations reported at the Cray User Group Meeting. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Delivered usage summaries, obtained via LDMS, of codes (e.g., GEMMA, ALEGRA, EMPIRE, LAMMPS) and SAW on HPC platforms, to code teams to enable usage understanding and explore potentials for efficiency improvements. (Appendix J, target FOUS-5)

Planned Activities in FY25:

- Explore a dual configuration of El Dorado as a hybrid Application Readiness Testbed and production resource. (Appendix J, FOUS-1, CSSE-1)
- In collaboration with the Vanguard program, stand up and deploy components of Spectra, the Vanguard 2 flagship platform. (Appendix J, FOUS-2, CSSE-1, CSSE-2)
- In collaboration with the CSSE program, stand up and deploy an advanced architecture testbed based on Cerebras' wafer-scale compute technology. (Appendix J, FOUS-2, CSSE-1, CSSE-2)
- Continue production support and/or decommissioning of advanced architecture testbed clusters and infrastructure including support for the Vanguard program, Application Readiness Testbeds (e.g., Tachi), AI/ML testbeds, CaaS/ADE testbeds, and other programs with advanced architecture needs. (Appendix J, FOUS-1, FOUS-2, CSSE-1, CSSE-2)
- Continue and advance support and design and installs of SNL's LDMS/AppSysFusion production capability and infrastructure. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Develop a standardized LDMS configuration and deployment model targeting multiple DOE computing sites' ATS and CTS systems and for new operations models (e.g., cloud), in collaboration with the LDMS User's Group and vendors. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Using LDMS/AppSysFusion data, continue and advance work with users and admins to improve diagnosis and performance and inform resource allocation decisions. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Deploy 2nd ASC funded CTS-2 systems with GPU and standard processors. (Appendix J, target FOUS-5)
    - Deployment of ASC CTS-2 SCN System (6SU) – "Nitro" (HPC)

- Deployment of ASC CTS-2 SCN GPU System (16 nodes) – "Cascade" (AI/ML and HPC)

- Deploy an integrated Starfish/FUSE solution for directed and/or automated data transfers to HPSS. (Appendix J, targets FOUS-5, and CSSE-2)
- Security: First Draft of HPC Security Overlay (for NIST 800-53r5) completed in FY24Q4. Second Draft with proposed expanded content completed in FY25Q1. Draft Publication and Comment Period completed by end of FY25 with possible finalized and published in FY25.
- Security: Rewrite SCN Security plan implementing NIST 800-223 guidance starting FY25Q2 for re-accreditation in FY26Q2. If available and authorized, use draft HPC Security Overlay (for NIST 800-53r5) for security control selection during re-accreditation.
- Continued operations and support of NSCC computing environment. Refine operations strategy around continued HPC capabilities in this environment. (Special Projects)
- Finalize Deployment of the NA-82 funded 6SU CTS-2 based cluster (DarkMead) and associated supporting infrastructure. (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)
- Support FASST (Frontier AI for Science, Security, and Technology) programmatic activities and data calls as needed (a place to capture our contributions to this effort).

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 FY24:

- Continued evaluating and deploying the latest versions of system configuration and management tools on testbeds and Vanguard II systems. Began discussions to expand those tools to new or planned architectures such as ARM, Grace-Hopper, RISC-V. (Appendix J, targets FOUS-1, FOUS-2, FOUS-5, CSSE-1, and CSSE-2)

- Developed and prototyped an HPC cluster on Amazon Web Services cloud infrastructure resulting in a reproducible deployment repository, initial benchmarks, and integration with local LDMS monitoring infrastructure (in conjunction with CSSE). (Appendix J, targets FOUS-5, and CSSE-2)
- Added Cloud-native Storage (supporting CSI protocol) alternative to the Vanguard II project in order to de-risk and provide a comparison / alternative to DDN Infinia. Continued evaluation of the latest versions of DDN Infinia. (Appendix J, targets FOUS-2, CSSE-1, and CSSE-2)
- Enhancements to LDMS including: a) credit-based flow control to bound user application data injection rates, b) fully-distributed resilience and load-balancing management layer, c) higher per-daemon processing performance, and c) improved in-line analysis capabilities. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Caliper and DARSHAN LDMS data connectors were integrated into their respective code bases. These send application profiling information during run time across the LDMS transport for run time evaluation (in collaboration with LLNL and ANL, respectively). (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Enhanced LDMS-Kokkos tools interactions to enable run time flow of PAPI performance counter information, collected at the Kokkos layer, across the LDMS transport for run time evaluation (in collaboration with CSSE). (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- New AppSysFusion analysis capabilities including an initial deployment of multithreaded time-series based analytics, anomaly detection with co-kurtosis, and DARSHAN summary analyses and dashboards. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Developed and deployed container-based LDMS services, to enable simplified standard installs. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Technical participation in the Computing and Information Sciences External Review Board in the areas of Advanced Workflow Orchestration, HPC-Cloud Convergence, and Autonomous Computing. (Appendix J, targets FOUS-1, FOUS-2, FOUS-5, CSSE-1, and CSSE-2)
- ASCR Award granted for the Software Tools Ecosystem Project (STEP), as one of the software sustainability projects, which includes LDMS leads as Co-PI's. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)

Planned Activities in FY25:

- Continue and advance vendor codesign explorations in advanced system management and functionalities, including specialized hardware, system provisioning, advanced job workflows, cloud-style computing, and novel storage solutions. (Appendix J, FOUS-1, FOUS-2, FOUS-5, CSSE-1, and CSSE-2)

- Internal development and deployment of novel operations models, including advanced HPC systems infrastructure tools and services, including configuration management and system provisioning tools, forward-looking job workflow management approaches and systems, and system observability tools, in order to support complex workflows, hybrid AI+HPC models, and autonomous computing. (Appendix J, FOUS-1, FOUS-2, FOUS-5, CSSE-1, and CSSE-2)
- Continue and advance LDMS/AppSysFusion capabilities, analyses (including AI), and infrastructure (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Develop a run time performance-tuning feedback infrastructure based on interoperability of LDMS+Kokkos and other tooling. (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Develop Computing-as-a-Service (CaaS) infrastructure for intelligently orchestrating HPC and AI jobs across multiple HPC and Cloud computing resources, integrating system health monitoring information, via LDMS, to route and map jobs efficiently. Demonstrate on the HAAPS Testbeds. (Collaboration with CSSE) (Appendix J, targets FOUS-1, FOUS-5, and CSSE-2)
- Deploy a multi-cluster SLURM to the SCN to enable more efficient use of HPC resources (Appendix J, target FOUS-5).

User Support (SNL)

The User Support project provides user support, development environments (e.g., DevOps infrastructure), and associated resources for SNL computing systems and tri-lab resources. User support activities focus on improving the productivity of the entire user community, local and remote, in utilizing the ASC HPC resources.

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

Accomplishments in FY24:

- Provided service desk, Tier-3 user support, and service desk improvements for SNL and tri-lab ASC computing systems. (Appendix J, targets FOUS-1, FOUS-2, and FOUS-5)
- Maintained and ensured developer tools (profilers, debuggers, and system performance and analysis) were available for ASC development teams and user workflows. (Appendix J, targets CSSE-2, ATDM-1, and IC-1)

Planned Activities in FY25:

- Provide service desk, Tier-3 user support, and service desk improvements for SNL and tri-lab ASC computing systems. (Appendix J, 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 J, 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 FY24:

- ASC DevOps has successfully developed sustainable processes to easily deploy code environments across Sandia's ASC systems for both analysts and developers. This includes adoption of the ASC Unified Environment (AUE), improved build and test farm infrastructure, and improved metrics. The AUE is now being heavily used across many of the ASC code teams. This Sandia-wide collaboration across ASC is continuing and more efforts to increase efficiencies have been identified and are being executed.
- Created initial criteria for code teams adopting ASC Unified Environment (AUE). (Appendix J, targets CSSE-2 and IC-1)
- Developed cost model for measuring increased efficiencies for adopting AUE. (Appendix J, targets CSSE-2 and IC-1)
- Delivered initial code usage data collection tool so consistent metrics across ASC can be collected. (Appendix J, targets CSSE-2 and IC-1)
- Implemented two new dedicated HPC system clusters (identical to Commodity Technology System (CTS-1) and CTS-2) for build-and-test continuous integration (CI). This increases compute capacity and replaces the decommissioned legacy standalone servers. This approach is ~30x cheaper per core than the former approach. Also, CI efforts are significantly improved by executing directly on an HPC cluster versus a surrogate compute server. (Appendix J, targets CSSE-2 and IC-1)

Planned Activities in FY25:

- Create observability platform where metrics for DevOps are collected for the ASC Unified Environment (AUE) product to measure processes over time and inform cost

models (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.
- Develop testing strategy for integration of ASC code teams with ASC Unified Environment (AUE) product (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.
- Implement a software bill of materials standard for improving security of the software being deployed (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.
- Leverage test harness creation discussed in Integrated Codes (IC) deliverable statement (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.
- Deliver resource managed Common Engineering Environment (CEE) build & test infrastructure which increases machine utilization and code team access to heterogenous architectures (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.
- Deliver comprehensive service management scaling plan (Appendix J, targets CSSE-2 and IC-1). Ties to CaaS/Platform Services Team.

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 FY24:

- Completed additional 4MW power minor construction project to 725East data center.
- Finalized design and began construction of cooling expansion minor construction project for 725East data center (planned completion is in FY25).
- Updated DisCom network service monitoring and fault isolation and created a standard operating procedure for carrier and site interactions and communications.

- Established network telemetry and analytics capabilities for DisCom to enable 100G network tuning.
- Replaced RedEagle COMSEC equipment ahead of December 2023 mandatory upgrade.
- Began isolating and reducing implicit access between SRN HPC and local (non-HPC) site resources.
- Updated iHPC 100G design (with feedback and lessons learned from FY23 S/RD upgrade) and completed upgrade.

Planned Activities in FY25:

- Complete construction of cooling expansion minor construction project for 725East data center.
- Improve isolation between HPC and local site resources (continuing from FY24).
- Upgrade SNL New Mexico DisCom black router hardware (continuing from FY24).
- Expand the SNL New Mexico iHPC IP (Internet Protocol) allocation to support El Dorado.
- Begin targeted IPv6 (Internet Protocol version 6) deployments to meet mandates.
- Complete the 100G iHPC COMSEC (Communications Security) upgrade for inter-site connectivity.
- Upgrade capacity on firewalls supporting iHPC.
- Research beyond 100G+ WAN (Wide Area Network) future state options for ASC DisCom.
- Expand SOP (standard operating procedure) documentation to cover troubleshooting common issues.
- Refresh and update WAN circuit documentation.
- Convert SNL New Mexico HPC network gateways to a highly available (HA) router configuration.
- Upgrade SNL (S/RD) HPC distribution layer hardware and optimize topology.
- Begin tuning maximum transmission units (MTU) between crypto tunnels and transport layers.
- Expand OHPC to CSRI (Computer Science Research Institute) in support of the Structural Simulation Toolkit project.
- Refresh obsolesced hardware at the network edge supporting HPC.
- Support SOC (Sandia Operations Center) network monitoring tool replacement project for local and remote resources.

- Evaluate BGP (Border Gateway Protocol) as an alternative dynamic routing protocol to OSPF (Open Shortest Path First) for path control across ASC DisCom.
- Craft a telemetry dashboard to highlight network utilization.

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 FY24:

- Maintained ES license support contract for FY24.
- Leveraged support contract, improved packaging and distribution of HDF5 and provide priority response to NNSA tri-lab user requests.

Planned Activities in FY25:

- Continue ES license support contract for FY25.
- Continue to leverage support contract, improve packaging and distribution of HDF5 and provide priority response to NNSA tri-lab user requests.

Projects for the Common Computing Environment Product (WBS 1.2.3.6.3)

The goal of the CCE product is to enable a common environment across the tri-labs that was initially deployed on the Tri-lab Linux Capacity Cluster (TLCC) systems. The scope of this product includes funded R&D projects to address gap areas identified by the tri-lab technical working groups.

The CCE working groups and projects focus on a common software stack, including but not limited to, OS software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues such as configuration management, licenses, WAN access, and multi-realm security. The CCE also coordinates the effort to provide a secure, robust, and performant remote computing environment for tri-lab users. This cross-lab remote computing environment provides users with access to HPC resources at the other two labs that is equivalent to the access at their own lab.

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 the El Capitan AT system as well as 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 FY24:

- Supported TOSS 3 on TLCC2 and CTS1/1+ systems with quarterly updates through the RHEL 7 End of Life in June 2024. (Appendix J, targets CSSE-2.b, FOUS-2, and FOUS-5)

- Continued support of TOSS 4 for CTS1/1+/2 and CORAL-2 systems with monthly updates. Released TOSS 4.7 (based on RHEL 8.9) and 4.8 (based on RHEL 8.10). Supported the integration of TOSS4 on CTS-2 clusters. (Appendix J, targets CSSE-2.a-c, FOUS-1.a, FOUS-2, and FOUS-5)
- Updated components of the TOSS buildfarm to enhance supply chain security and support RHEL 9 (Appendix J, targets CSSE-2.a-c, CSSE-2.e, FOUS-1.b, FOUS-2, and FOUS-5)
- Released TOSS 5.0, based on RHEL 9.4, to support testbeds, Advanced Technology, and other forward-looking systems. (Appendix J, targets CSSE-2.a-c, CSSE-2.e, FOUS-1.b, FOUS-2, and FOUS-5)
- Integrated three new LDMS versions (4.4.[1-3]) into TOSS which provided new sampling capabilities for NVIDIA GPUs, AMD GPUs, and Slingshot switches along with security, performance, and stability improvements. Also formalized better protocols and practices for incorporating LDMS into TOSS in the future.
- Organized and led the 2nd annual CCE All-Hands meeting at Sandia on April 2nd and 3rd with 43 attendees. The group planned several activities in order to address new challenges and opportunities presented by NNSA ASC HQ attendees.
- Continued research collaborations with six universities on HPC anomaly diagnosis and GPU bottleneck detection with machine learning, unsupervised clustering with time-series-based techniques, dynamic resource management, application instrumentation, and runtime I/O characterization leading to several publications and posters.
- Created a project plan for evaluating application progress during runtime and creating feedback mechanisms to resource managers to end hung processes and provide better runtime predictions.
- Broadened the common analysis framework with new analyses, I/O shim layers, and functional Jupyter Notebooks to make interoperability easier between sites.
- Sites continued to share lessons learned and tips for monitoring deployments and improvements, including preparing databases for Exascale monitoring sets, utilizing the Redfish stream from CSM, and making LDMS work within CSM.

Planned Activities in FY25:

- Continue support of TOSS4for CTS1/1+/2 systems and El Capitan. (Appendix J, targets CSSE-2.a-c, FOUS-1.b, FOUS-2, and FOUS-5)
- Continue development and support of TOSS5, based on RHEL 9, including support for CTS-2 systems as well as testbeds and El Capitan. (Appendix J, targets CSSE-2.a-c, FOUS-1.a, FOUS-2, and FOUS-5)
- Begin development of TOSS 6, expected to be based on RHEL 10, with a focus on supporting future platforms and management models. (Appendix J, targets CSSE-2.a-c, CSSE-2.e, 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, especially within the common analytics framework. (Appendix J, 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 J, targets CSSE-2, FOUS-1, and FOUS-5)
- Share new analysis techniques and monitoring approaches, form nucleation points for collaboration, and find tri-lab impact opportunities through periodic CCE meetings. (Appendix J, targets CSSE-2, FOUS-1, and FOUS-5)
- Explore and identify potential ASC-common processes for Open Source HPC software development and deployment and associated artifact release. Work with site-specific policy makers to explore the potential for tri-lab implementation. (In collaboration with Programming Environment Development/Support for tri-lab Systems) (Appendix J, 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 FY24:

- Continued to develop Pavilion2 features, test integrations, and dashboards.
- Managed support contracts for Tri Lab common tooling. HPCToolkit engaged vendors and ported to Crossroads, El Capitan, and CTS-2 systems. Trenza Survey, ParaTools TAU, and U. Wisconsin's DyninstAPI were also supported by contracts. (Appendix J, targets FOUS-5 and CSSE-2a)
- Designed and built prototyped Tri Lab-shared infrastructure for building programming environments, including a GitLab in AWS GovCloud. (Appendix J, target CSSE-2c)

Planned Activities in FY25:

- Continue Pavilion2 development in support of continuous testing, acceptance testing, performance testing and broader support for tri-lab resources.
- Tri-lab Container Environment study. Evaluate containers on Tri-lab systems with various MPI schemes to illuminate areas for future improvements and collaborations toward goal of running everywhere.
- Continue existing support contracts. HPCToolkit to harden support for El Capitan and extend support on Saphire Rapids and ARM systems. (Appendix J, targets FOUS-5 and CSSE-2a)
- Harden deployment of shared Gitlab for ASC PE. Deploy shared build and source caches for programming environments. (Appendix J, target CSSE-2.c)
- Explore and identify potential ASC-common processes for Open Source HPC software development and deployment and associated artifact release. Work with site-specific policy makers to explore the potential for tri-lab implementation. (In collaboration with System Software Deployment for Commodity Technology Systems) (Appendix J, targets CSSE-2, FOUS-1, and FOUS-5)

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 FY24:

- Continued to operate the tri-lab SARAPE process for all remote access account requests and improved the user interface experience and the remote site processing access and usability.
- Completed the SARAPE modernization which brought SARAPE to latest server software and provided new functionality.
- 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, including the split of Stockpile Sustainment broken down by system and P&EM separated from ASC.
- Upgraded the WCTool application's framework language to newer supported version.

Planned Activities in FY25:

- Explore opportunities for integration of Cyber Security Users Agreement (CSUA) training to be available in SARAPE and recorded in Training, Education, and Development System (TEDS).
- Integrate Common Engineering Environment Service Catalog/Shopping Cart portal (Nile) subscriptions/renewals with SARAPE, similar to Web Computer Account Request System (WebCARS).
- Continue to deliver quarterly utilization reports to DOE HQ and 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 design agency HPC simulation complex.

The RCE team is comprised of service providers from multiple HPC 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 (CI), authentication efficiencies, storage 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.

As a tri-lab endeavor with three notional endpoints, it is often useful to refer to a "senary-way" relationship when pertaining to new functionality and enhancements that are bi-directionally dependent between any two endpoints across each of the tri-labs. For many tri-lab wide RCE collaborative projects (e.g. 100Gbps DisCom, cross-site CI, et al.), a capability that has both server-side and client-side components can be considered fully deployed when it has achieved senary-way (i.e., 6-way) operational status.

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

Accomplishments in FY24:

- Began to tune to maximum efficacy the recently deployed classified DisCom 100Gbps infrastructure. Completed tri-lab procurement of Apposite Netropy devices for coordinated 4-site (LANL; LLNL; SNL, CA; and SNL, NM) full stack testing. LANL and LLNL tested black-side (unencrypted) lines between their endpoints to achieve near line rate. SNL, CA and SNL, NM black-side testing pending.
- Began the tri-lab endeavor of seeking security approvals and associated security plan updates in pursuit of classified cross-site CI capability modeling the FY23 unclassified restricted deployment of cross-site CI. This approval is currently now awaiting approval to operate (ATO) from NA-IM security authorizing officials. Once received, the tri-lab SMEs are ready to deploy the S/RD cross-site functionality for classified Gitlab instances.
- LLNL and SNL achieved a common understanding of the path forward for Kerberos trust between SNL RR (login nodes) and LLNL RZ to enable unprompted SSH CLI authentication. SNL to implement changes to IHPC authentication and then develop a MOU with LLNL.
- Investigated allowing remote desktop IP ranges across the IHPC network directly to RE/RR/RZ web server fronts (e.g., https protocol traffic). Focused on the LANL desktop to LLNL RZ web fronts as the first step in the ultimate senary-way honoring of this direct IHPC access. LLNL deployed firewall changes to accept incoming LANL desktop IP ranges. LANL currently undergoing security review such that outbound desktop web traffic can be routed through IHPC restricted unclassified network.
- Kickstarted a conversation concerning a "Tri-lab Data Runner" (previously Tri-lab Data View) Conduct initial study of a tri-lab file system (POSIX) and/or object store (S3 endpoint) storage capability, with physical and logical distribution across all three HPC centers presented as a singular namespace visible to users as the same data in all three locales bounded by design parameters of modest use cases such as small dataset transfers and data sharing functions. Current potential vendors include Hammerspace

and Vcinity, with recent tri-lab vendor discussions having been conducted with Vcinity.

- LANL successfully deployed the new Restricted Enclave File Transfer Agent cluster (RE-FTAs), providing a data management system in the IHPC.
- Tri-labs implemented a subset of data transfer tools in the restricted unclassified IHPC ubiquitously across sites, modeling same environment deployed in FY20 in the classified environment. Currently undergoing an IHPC SLA security approval update, after which transfer tooling will be green lit for deployment at all three IHPC restricted unclassified HPC centers.
- LANL removed the RE gateway dependency for inbound and outbound IHPC traffic gatekeeping purposes, favoring access to frontend nodes of clusters which is now accessible for inbound access from Sandia RR and LLNL RZ.
- Coordinated a procurement of tri-lab (four site) encryptors for the IHPC data transfer and visualization network to support encrypted 100Gbps speeds, which will represent an improvement of 10x over current throughput. This leverages underlying supporting line upgrades associated with the classified DisCom WAN 100Gbps upgrade of FY23.
- Deployed to production an LLNL RealVNC server configuration allowing LANL remote client Kerberos credential authentication coinciding with TOSS4 upgrades on same architecture.
- Tri-lab network SMEs coordinated over the course of months to diagnose and eventually successfully solve an IHPC 10Gbps Taclane encryptor problem that necessitated full endpoint reboot of LANL and LLNL encryptors every 7-10 days to avoid packet loss outages.
- LANL has developed training and documentation for user application containerization and usage based on CharlieCloud which has been successfully used for milestones supporting CCE tri-lab container portability efforts. This has enabled L2 milestone work for Sandia, LLNL, and LANL teams that have Crossroads user application container deliverables in FY24.
- LLNL has developed and hosted containers in CZ Gitlab which have been tailored for Crossroads and El Capitan ATS system portability inclusive of MPI library integration supporting CCE tri-lab container portability efforts. Pre-production tests have been successfully completed on those architectures with tri-lab participation from LANL SMEs.
- SNL collaborated with LLNL and deployed to production a RealVNC server in the SNL RR to support connections from remote desktop VNC clients at SNL, LLNL RZ, and LANL RE.
- LANL engaged with LLNL to evaluate RealVNC as track to use common toolset across sites. LANL has started its evaluation process of RealVNC.

Planned Activities in FY25:

- LANL and LLNL HPC centers to use recently deployed IHPC data transfer utilities to send approximately 180TB of digitized above ground test film covering several decades from LLNL to LANL as part of the larger multi-lab film archiving project. Such an exemplar use case is a good illustration of high bandwidth WAN data transfer capabilities between design agencies that may well serve future digital engineering/digital transformation uses.
- Continue to tune the classified DisCom 100Gbps deployments at the encrypted "red" layer at the four sites. The FY24 procurement of Apposite Netropy devices will be critical to this exercise. After this step, continue to debug up the software stack and note where any future investments need be made to maximize unclassified (IHPC) or classified DisCom high bandwidth data transfers.
- Continue tri-lab discussions on the Tri-lab Data Runner (TLDR) project, focusing on the shared WAN data view and transfer strategy. This encompasses further defining requirements and evaluating vendors.
- Engage with the NA-IM driven "ESN Hub" digital transformation project. Examine potential for integration and collaboration to effect change towards RCE's mission. Early indicators suggest potential for mutual advantages for DAs and potentially PAs alike.
- Transition to using the DOE OneID DUID as the authoritative identifier for translating user IDs across HPC centers for unclassified cross-site CI activities. Replaces current method of mapping one site's username to another site's equivalent username based on site-specific methods. Requires sending the DUID as part of the Jacamar cross-site payload from one site and performing the necessary LDAP translations at the receiving site handling the compute request. Achieving this milestone will illustrate leveraging disparate and distributed identity providing services such as have been true of the tri-labs for decades against an agreed upon common source of an authentication translational layer, thus illustrating an alternative to an untenable homogeneous federated identity providing service deployed at each

lab. We envision this distributed translational model to have future use cases beyond cross-site CI.

- Continue to work towards senary-way direct network routing from tri-lab institutional desktops without gateways or hops to remote-lab destination restricted unclassified web fronts (RE/RR/RZ web "store fronts").
- Complete senary way IHPC (unclassified restricted enclaves) high bandwidth data transfer tools (e.g., hsi2lanl, pftp2rzslic, pftp2ihpc, etc.)
- Add further usage information on SARAPE to clarify account processing and training requirements. LANL added HSPD-12 PIV authentication access for Sandia and LLNL for unclassified user documentation.
- 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.
- LLNL to further recent MPI enabled container portability achievements on El Capitan and Crossroads class ATS system by extending to "GA" status (production quality) inclusive of LC Hotline Technical Bulletin release and external documentation on hpc.llnl.gov. Further effort to also include MPI-enabled containers that are portable across CTS2 and ATS systems in cooperation with LANL and SNL.
- Implement a decided-on path based on final results of LANL's current evaluation of RealVNC in the manner deployed at SNL and LLNL.
- 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 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 FY24:

- Procured and installed two new CTS-2 systems for special purpose programs.
- Completed infrastructure improvements and modifications to support CTS-2 purchases, including installation of a new network and upgrade of another network.
- Supported effective use and maintenance of the Special Purpose Computing environment.

Planned Activities in FY25:

- Provide ongoing system maintenance and updates, software upgrades, and hardware repairs.
- Support customers in all aspects of the computing environment and maintain codes on Special Purpose Computing systems.
- Improve the 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 FY24:

- Configured new compute hardware for production use.
- Initiated construction of facility modifications in SCIF to host new Crossroads hardware.
- Began transition to new system management process to match Crossroads.

Planned Activities in FY25:

- Complete construction of facility modifications in SCIF to host Crossroads hardware.
- Install and replace storage system for Special Purpose computing.
- Move and install Crossroads hardware into newly upgraded SCIF.

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.

Accomplishments in FY24:

- Refer to WBS 1.2.3.6.2

Planned Activities in FY25:

- Refer to WBS 1.2.3.6.2

APPENDIX H: 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 FY24:

- Selected 40 fellows, a new record level for a CSGF incoming class, and provided benefits in STEM fields that use high performance computing to solve open and complex science/engineering problems.
- Held a successful annual CSGF Program Review that highlighted current and graduating CSGF fellows' research performed during their practicums and at their respective institutions in the past year.

Planned Activities in FY25:

- Collaborate with DOE Office of Science Advanced Scientific Computing Research (ASCR) in supporting a new cohort of fellows to be trained as next-generation leaders in computational science.
- Maintain the CSGF community of enthusiastic and committed doctoral students, alumni, and DOE/NNSA laboratory staff who together serve as a support system for the incoming and current fellows.

PSAAP III Centers

The following FY24 Accomplishments and FY25 Planned Activities are for the PSAAP III Centers, which are activities that started in September 2020.

University of Colorado at Boulder

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

Accomplishments in FY24:

- Verified elastostatic linear isotropic elasticity implementations in the Ratel-FEM, GEOS-MPM, LAMMPS-granular, and Tardigrade-MOOSE codes against an analytical solution for the quasistatic compression of a cylindrical sample.
- Completed the initial implementation of an implicit Material Point Method (iMPM) in Ratel, enabling basic finite-deformation simulations with limited material support.
- Established an Uncertainty Quantification (UQ)-guided material model calibration workflow for parameters required by the Ratel-FEM, GEOS-MPM, and LAMMPS-granular codes.
- Completed initial baseline simulations and convergence demonstration for the press manufacturing process using the LAMMPS-granular code.
- Performed experiments (at UT Dallas) and simulations (using Ratel-FEM and Tardigrade-MOOSE) of quasi-static unconfined compression of 5mm diameter IDOX-nitroplasticized-Estane cylinders, including full-puck 3D meshing with around 20 million elements (grains).
- Identified the importance of, and need to account for, under-resolved, small, tightly-packed grains in modeling the mechanical response of mock high explosives (HEs).

Planned Activities in FY25:

- Test a newly developed generative artificial intelligence (AI) approach for generating grain microstructures against experimental computed tomography (CT) measurements.
- Predict the thermomechanical behavior of pressed mock HE via a computational multiscale micromorphic framework.
- Quantify the variation in one or more processing variables (e.g., press temperature, pressure, or duration; grain size distribution, binder type, or formulation) on the mechanical behavior of specimen-scale mock HE
- Extend the implementation of implicit MPM in Ratel to support other constitutive models that encompass a wider range of materials, and enhance performance through improved algorithms in underlying libCEED and PETSc libraries.
- Perform experiments (at Army Research Laboratory and UT Dallas) and simulations (using GEOS-MPM) of dynamic compression of 5mm diam IDOX-nitro-Estane cylinders.

University of Illinois at Urbana-Champaign

Center for Exascale-Enabled Scramjet Design

Accomplishments in FY24:

- Organized the January 2024 NNSA-University Workshop on Exascale Simulation Technologies (NUWEST), where 61 PSAAP and NNSA Laboratory representatives shared ideas on technologies for facilitating exascale predictive science by highlighting available technologies, identifying challenges, and initiating further collaborations.
- Generalized the MIRGE-Com solver to allow quad/hex finite elements, with essentially the same user-facing drivers as for the previous triangle/tetrahedra mesh and without significantly disrupting the structure of the Loopy lazy evaluation path for code generation.
- Demonstrated portability and scalability of MIRGE-Com production code by running on laptops and workstations with AMD, Apple, and Intel processors, multi-core/node CPU clusters (LLNL Quartz), and AMD or NVIDIA GPU-accelerated supercomputers (LLNL Lassen and Tiago, and NCSA Delta), scaling to 2048 NVIDIA V100 GPUs on Lassen.
- For setting up multi-platform workflows, Parsl has been coupled with Globus Compute, providing a Jupyter Notebook driver capability for orchestrating HPC simulations within a larger workflow, which involves pre-processing, run(s), postprocessing, and possibly UQ sampling.

Planned Activities in FY25:

- Assess integrated predictive capability for a variable geometry scramjet with novel high-temperature composite combustor wall materials.
- Improve the code generation, transformation, and optimization (i.e., compilation) time of Loopy, to enable routine highly performant MIRGE-Com production simulations at scale.
- Complete predictive simulations for the supersonic ACT-II combustor experiments, expected to involve 100 million third-order elements for about one million time steps, using a reduced global ethylene–air combustion mechanism, a simple model for porous media transport, phenolic pyrolysis, and carbon oxidation in the wall material.
- 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 FY24:

- Completed the first probabilistic analysis of laser-induced ignition and compared Purdue experiments and HTR simulations with respect to the probability of successful ignition in a gas-phase model rocket combustor.

- Established a detailed verification roadmap coupled with a continuous integration / deployment regression suite for the Center's software framework.
- Deployed the first 3D ensemble browser to study the uncertainty of the ignition scenario and explore the sensitivity to initial conditions.
- Developed a new modeling strategy and completed the first simulations of a liquid-phase model rocket combustor, including spray and atomization of the liquid oxidizer.
- Constructed data-driven surrogates that represent the full system and enable acceleration of the reliability studies (within a multifidelity framework), and derived a data-driven physical model for multiphase flow modeling that can be embedded within the HTR solver framework to tackle multiphase ignition scenarios.
- Hosted a November 2023 workshop on multiphase flow modeling, co-organized with the University at Buffalo PSAAP Center. This workshop attracted 70 speakers and participants from PSAAP Centers and NNSA Laboratories, and about 20 participants from other government labs (NASA, ONR) and industry.

Planned Activities in FY25:

- Complete porting and optimization of HTR++ code to AMD GPU/APU-based systems, including Frontier and El Capitan.
- Perform HTR++ simulations to predict the reliability of laser-based ignition of cryogenic propellants in a liquid-phase model rocket combustor.
- Refine the multi-fidelity uncertainty quantification (UQ) strategy utilized to construct ignition maps.
- Compile and compare experimental ignition maps for the liquid-phase model rocket combustor.

University of Texas at Austin

Exascale Predictive Simulation of Inductively Coupled Plasma Torches

Accomplishments in FY24:

- Initial plasma torch simulations have been performed for a nitrogen feed gas, showing a good agreement with the measured exit temperature and correct qualitative changes relative to an argon feed gas.
- Evaluated the theoretical roofline performance and measured the performance, both bare-metal and containerized, of the Boltzmann Transport Equation (BTE) and Torch Plasma Simulator (TPS) solvers on NVIDIA and AMD GPU-based systems.
- Coupled the electron Boltzmann solver to multi-species plasma model, replacing electron continuity and energy equations.
- Identified and implemented numerical formulations that accelerate the modeling of transient effects during the start-up of torch simulations.

Planned Activities in FY25:

- Perform component validation experiments involving flow across a bow shock, as the ICP torch is exhausted into a vacuum chamber.
- Refine plasma models as guided by validation results.
- Perform multifidelity UQ to capture torch exit and blowout uncertainties induced by inflow condition and chemistry parameter uncertainty for nitrogen feed gas, and compare to experimental measurements
- Complete standalone, angle-adaptive, 3D radiative transfer equation (RTE) solver
- Additional development and Parla-based model implementations dictated by performance analysis of full-system ICP torch simulations

University at Buffalo

Center for Exascale Simulation of Hybrid Rocket Motors

Accomplishments in FY24:

- Developed preliminary GPU support for ABLATE's fluid solver, implementing a low-level library allowing Nvidia's CUDA compiler to interface with the framework.
- Validated with an automated approach models of methyl methacrylate (MMA, formed via polymethylmethacrylate or PMMA) and C32H66 combustion and the role of molecular O2 in ignition, demonstrating improved performance relative to prior published data and which perform closer to experimental data.
- Formulated a wall model and scaling law using direct numerical simulations (DNS) for turbulent boundary layers at various blowing rates.
- Used an experimental hybrid rocket motor assembly to generate PMMA ignition, pressure, and thrust data to validate against theoretical predicted performance.
- Implemented upgrades to slab burner assembly to allow characterization of additional flight regimes, allowing data collection and droplet measurements at higher flow rates more typical of common flight regimes, enabling improved uncertainty qualification and verification for fuel gasification through experimental means.

Planned Activities in FY25:

- Further develop the ABLATE and libCEED frameworks to enable chemical kinetics evaluation using GPU platforms to accelerate simulations at the high-performance computing (HPC) scale.
- Explore interface and droplet formation dynamics using direct volume-of-fluid (VOF) and interface reconstruction methods to better understand atomization process in hybrid rocket motors.
- Focus on reducing time and enhancing performance for large alkane mechanism generation, while improving models for paraffin wax fuel (C32H66) soot precursors and nucleation which may serve as training data for ChemTab.
- Improve turbulence and eddy modeling by leveraging machine learning to develop a wall model for reacting boundary layers with blowing.

- Generate higher-quality and more consistent data from the hybrid rocket motor assembly, reducing data losses and feedback noise to allow the quantification and evaluation of smaller thrust values to reduce uncertainty.

Massachusetts Institute of Technology

Center for the Exascale Simulation of Material Interfaces in Extreme Environments

Accomplishments in FY24:

- Conducted benchmark calculations and evaluations of Hf layers, developing Density Functional Theory (DFT) methodologies to improve accuracy of modeling porous materials and adsorption by focusing on accurate machine-learning interatomic potentials (MLIPs) barrier descriptions and fast molecular dynamics in StencilMD to achieve chemical accuracy "at scale", with improved performance in some scenarios.
- Supported improvements to Julia and OpenCilk/Kitsune to enable robust GPU and parallel compute capabilities, while integrating core packages and features which accelerate energy and force calculations for interatomic potentials (IAP) fitting, enabling both linear and neural network data reduction techniques.
- Finalized DFT approximations which reduce uncertainty and began developing AtomisticQoIs.jl and Cairn.jl, which will allow better modeling of equilibrium quantities of interest (QoIs) and active learning of atomistic modeling to statistically calibrate error and compare benchmarks.
- Began a new experimental validation campaign using microscopy techniques to enable atomic imaging of Hf-graphene oxide experiments which compare favorably to atomic-scale simulations.
- Iterated towards goal of many-query studies of hypersonic flow simulations on multiple fronts, with new procedures moving towards convergence of steady hypersonic flows which prevent solver failures.

Planned Activities in FY25:

- Extend hypersonic flow simulation work to support reacting flows, continuing promising work on inviscid flows, allowing increased physical complexity in future models capable of quantifying gas-surface interaction uncertainty and heat transfer conjugation.
- Refocus efforts on accurate predictive HF surface oxidation simulation, with the goal of advancing state of the art transition, oxidation, and microstructure problems.
- Use atom probe tomography (APT) and other methods as avenues to obtain an atomic-resolution dataset that can be directly compared with modelling.
- Continue development to implement GPU capabilities in the ML-POD package, enabling large-scale simulations of oxidation processes at high temperatures and pressures.

University of Maryland

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

Accomplishments in FY24:

- Submitted a paper on the development and application of the space-split sensitivity (S3) method to a turbulent flow and implemented S3 in the uPDE/LEA framework.
- Hosted a one-week Turbulence Summer School, with 46 on-site graduate students and 54 remote participants.
- Created a Reynolds-averaged Navier–Stokes (RANS) adjoint solver that enables field inversion of spatially varying model coefficients, improving the accuracy of the multi-fidelity sensitivity analysis (MFSA) approach.
- Submitted a paper on metric-conforming grid-generation and used it to model flow separation over a smooth ramp using Large Eddy Simulation (LES).
- The Light Exascale Application (LEA) forward solver demonstrated performance portability on the LLNL Quartz, Dane, Lassen, and Tioga systems, and scalability up to 2748 GPUs (Lassen).

Planned Activities in FY25:

- Define two common benchmark problems, channel flow and a non-canonical boundary layer problem, and apply and compare the different sensitivity algorithms that have been developed by this Center.
- Apply S3 as implemented in uPDE/LEA to simulate channel flows at higher Reynolds numbers and larger meshes, connecting to LES subgrid scales.
- Implement a metric-conforming grid generation algorithm for multi-block grids, utilizing GPU accelerators, and assess the numerical and modeling errors at resolution interfaces.

University of New Mexico

Center for Understandable, Performant Exascale Communication

Accomplishments in FY24:

- Developed and released the Beatnik proxy application for remapping and tree-based global communication, including a new cutoff-based solver that remaps between surface and spatial distributions.
- Simple performance studies using arbitrary datatypes for irregular data communication showed that stream and kernel triggering of communication is more beneficial than complicated application programming interfaces (APIs).
- Paper on “Quantifying and Modeling Irregular MPI Communication” accepted at CCGrid24.
- Developed new communication abstractions and primitives and tested them via careful assessment, benchmarking, and modeling.

Oregon State University

Center for Exascale Monte Carlo Neutron Transport

Accomplishments in FY24:

- The MC/DC code developed by this Center is now capable of utilizing GPUs effectively, modeling time-dependent and continuous energy problems, and scaling to large problem sizes.
- Co-developed (with MC/DC) the Python/Numba-based Harmonize framework to support GPU execution without requiring significant GPU-specific code.
- Made significant advancements in the iterative Quasi-Monte Carlo (iQMC) method, including a batched approach that now outperforms multigroup Monte Carlo in regions of low flux.
- Implemented variance deconvolution, an embedded uncertainty quantification method, in MC/DC, and demonstrated that it is consistently more accurate and efficient than the standard sampling-UQ approach for stochastic solvers.
- Defined final verification, validation, and exascale challenge problems.

Planned Activities in FY25:

- Utilize asynchronous scheduling to significantly increase thread residency, and thus performance, on AMD GPUs/APUs.
- Add additional functionality to MC/DC, including the k-eigenvalue method of nearby problems for solution verification and compressed sensing to reduce the number of tallies, and thus memory footprint.
- Complete batch parallelization and scaling studies for the variance deconvolution embedded UQ method.
- Use MC/DC to model the time-dependent small modular reactor challenge problem with UQ.

Appendix I: Construction and Capital Equipment

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

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY24**	**FY25**	**FY26**	**FY27**	**FY28**	**Future**	**Contact**
LLNL	Commodity Technology System (CTS)-2 Systems	Capital Equipment, MIE	N/A	70,000	30,000	20,000	20,000	0	0	0	0	Matthew L. Leininger 925-4224110 leininger4@llnl.gov
LLNL	AT System – El Capitan	Capital Equipment, MIE	N/A	600,000	446,000	96,000	58,000	0	0	0	0	Bronis de Supinski 925-422-1062 desupinski1@llnl.gov
LLNL	Unclassified ATS-4-like System, LLNL	Capital Equipment, MIE	N/A	17,000	4,000	3,000	3,000	3,000	3,000	1,000	0	Ned Bass 925-422-9389 bass6@llnl.gov
LLNL	Commodity Technology System (CTS)-3 Systems	Capital Equipment, MIE	N/A	0	0	0	0	20,000	20,000	20,000	40,000	Matthew L. Leininger 925-4224110 leininger4@llnl.gov
LLNL	AT System – ATS-6	Capital Equipment, MIE	N/A	250,000	0	0	0	2,000	8,000	70,000	170,000	Bronis de Supinski 925-422-1062 desupinski1@llnl.gov
LLNL	Bldg 453 Tuolumne Site Prep	Minor Construction	N/A	1,600	1,200	400	0	0	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 453 Sierra Retirement– G2 11072	Minor Construction	N/A	4,000	0	0	1,000	3,000	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 451 Power and Cooling Improvements/Modifications– G2 6109	Minor Construction	N/A	7,500	0	0	750	6,750	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	Bldg 453 ATS Power Improvements/Modifications– G2 6092	Minor Construction	N/A	4,000	0	0	0	4,000	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B-654 CTS Power and Cooling Modifications – G2 6111	Minor Construction	N/A	3,000	0	0	0	3,000	0	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov

				$ in Thousands								
Site	**Title/Description**	**Classification**	**Last CD Completed**	**TEC**	**Prior Years**	**FY24**	**FY25**	**FY26**	**FY27**	**FY28**	**Future**	**Contact**
LLNL	B453 CTS Power and Cooling Improvements/Modifications – G2 6099	Minor Construction	N/A	7,500	0	0	0	2,000	5,500	0	0	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B453 ATS-6 Site Prep – G2 4172	Minor Construction	N/A	22,000	0	0	0	0	0	0	22,000	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LLNL	B453 ATS Cooling Loop Improvements/Modifications – G2 6087	Minor Construction	N/A	3,000	0	0	0	0	0	0	3,000	Anna Maria Bailey 925-423-1288 bailey31@llnl.gov
LANL	SCC Electrical Upgrade	Minor Construction	N/A	20,821	20,821	0	0	0	0	0	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	ATS-5 Installation	Minor Construction	N/A	25,000	0	0	1,000	24,000	0	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	5,000	0	Jason Hick 505-667-4477 jhick@lanl.gov
LANL	Crossroads: Acquisition of Crossroads (ATS-3) system	Capital Equipment, MIE	N/A	115,000	103,000	6,000	6,000	0	0	0	0	Jim Lujan 505-665-0718 jewel@lanl.gov
LANL	ATS-5: Acquisition of ATS-5 system	Capital Equipment, MIE	N/A	250,000	0	2,000	15,000	100,000	95,500	20,500	17,000	Jim Lujan 505-665-0718 jewel@lanl.gov
SNL	Additional CT Systems	Capital Equipment, MIE	N/A	40,000			10,000	10,000	10,000	10,000	10,000	Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725E 4MW Power Expansion	Minor Construction	N/A	5,600	5,000	600						Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	725E Cooling Capacity Expansion	Minor Construction	N/A	14,700	100	13,500	1,100					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
SNL	725E Expansion	Minor Construction	N/A	5,000					5,000			Tom Klitsner 505-844-1901

Site	Title/Description	Classification	Last CD Completed	$ in Thousands TEC	Prior Years	FY24	FY25	FY26	FY27	FY28	Future	Contact
												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	39,000			500	4,500	34,000			Tom Klitsner 505-844-1901 tklitsn@sandia.gov
SNL	New Intermediate Data Center	Minor Construction		45,000					1,000	44,000		Tom Klitsner 505-844-1901 tklitsn@sandia.gov

* Project estimate is not complete. Funding profile serves as a placeholder.

Appendix J: FY25–FY29 Program Targets

This appendix lists the targets for each subprogram for the FY25–FY29 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 FY28, and intermediate work will support pegposts on the SCDS Modern Materials & Manufacturing and Future Deterrent strands.

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 FY25, and earlier work will support various SCDS pegposts in the Stockpile Sustainment strand.

V&V-3. Develop Engineering and Physics Common Model Frameworks with integrated V&V/UQ. The goal is adaptable, efficient, standardized credibility processes and frameworks to support engineering simulation workflows, sensitivity analysis, and integration into physics workflows as appropriate. Analyses including thermal, structural, aerodynamics, and hostile survivability simulations will be supported for the FY26 Combined Threat Environments Simulation SCDS pegpost. Develop integration between common model frameworks and data repositories to allow modern workflows to seamlessly integrate data into deterrent simulation capabilities and to ensure distinction between data used for calibration/training and that used for validation. The target is FY27.

V&V-4. Incorporate AI methods and validation techniques that enable agile predictions suitable for stockpile applications. The target is FY29, 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 FY26 to support subsequent Stockpile Sustainment and Future Deterrent SCDS pegposts.

Computational Systems and Software Environment

CSSE-1. Platform Acquisition and Deployment.

a. Procure and deploy El Capitan, including NRE activities. The target date for system acceptance is Q4FY24 with classified production service in Q2FY25.

b. Deploy CTS-2 platform across tri-labs and transition to production service. The target dates are FY22-FY25.

c. Procure and deploy Vanguard-II Advanced Architecture Prototype System. The target date is FY26.

d. 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.

e. Initiate planning for a future CTS procurement in FY25. Systems will be deployed at LANL, LLNL and SNL during FY26 – FY29.

CSSE-2. Software Environment Deployment.

a. Develop and deploy software and programming environments for El Capitan. The target date is FY25.

b. Deploy the Vanguard-II system with a functioning software stack composed of ATSE, TOSS, and FOUS CCE tools and packages. The target date is FY26.

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

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

e. Develop a new generation of software and programming environment for a future CTS procurement. The target date is FY26 – FY29.

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 FY25 – FY30.
b. Demonstrate applications of interest on BML-funded hardware, including initial evaluation of neuromorphic Intel Loihi testbed. The target date is FY25.
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. Production operation of Crossroads and El Capitan.

a. The Crossroads system will be operated in production starting in FY24.
b. El Capitan will be transitioned to users in FY25. The target date for installation and the beginning of operations is the end of FY24.

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 FY26.

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 FY27.

Appendix K: Codes

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

Appendix L: 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.