ASC™

Advanced Simulation and Computing
COMPUTING STRATEGY

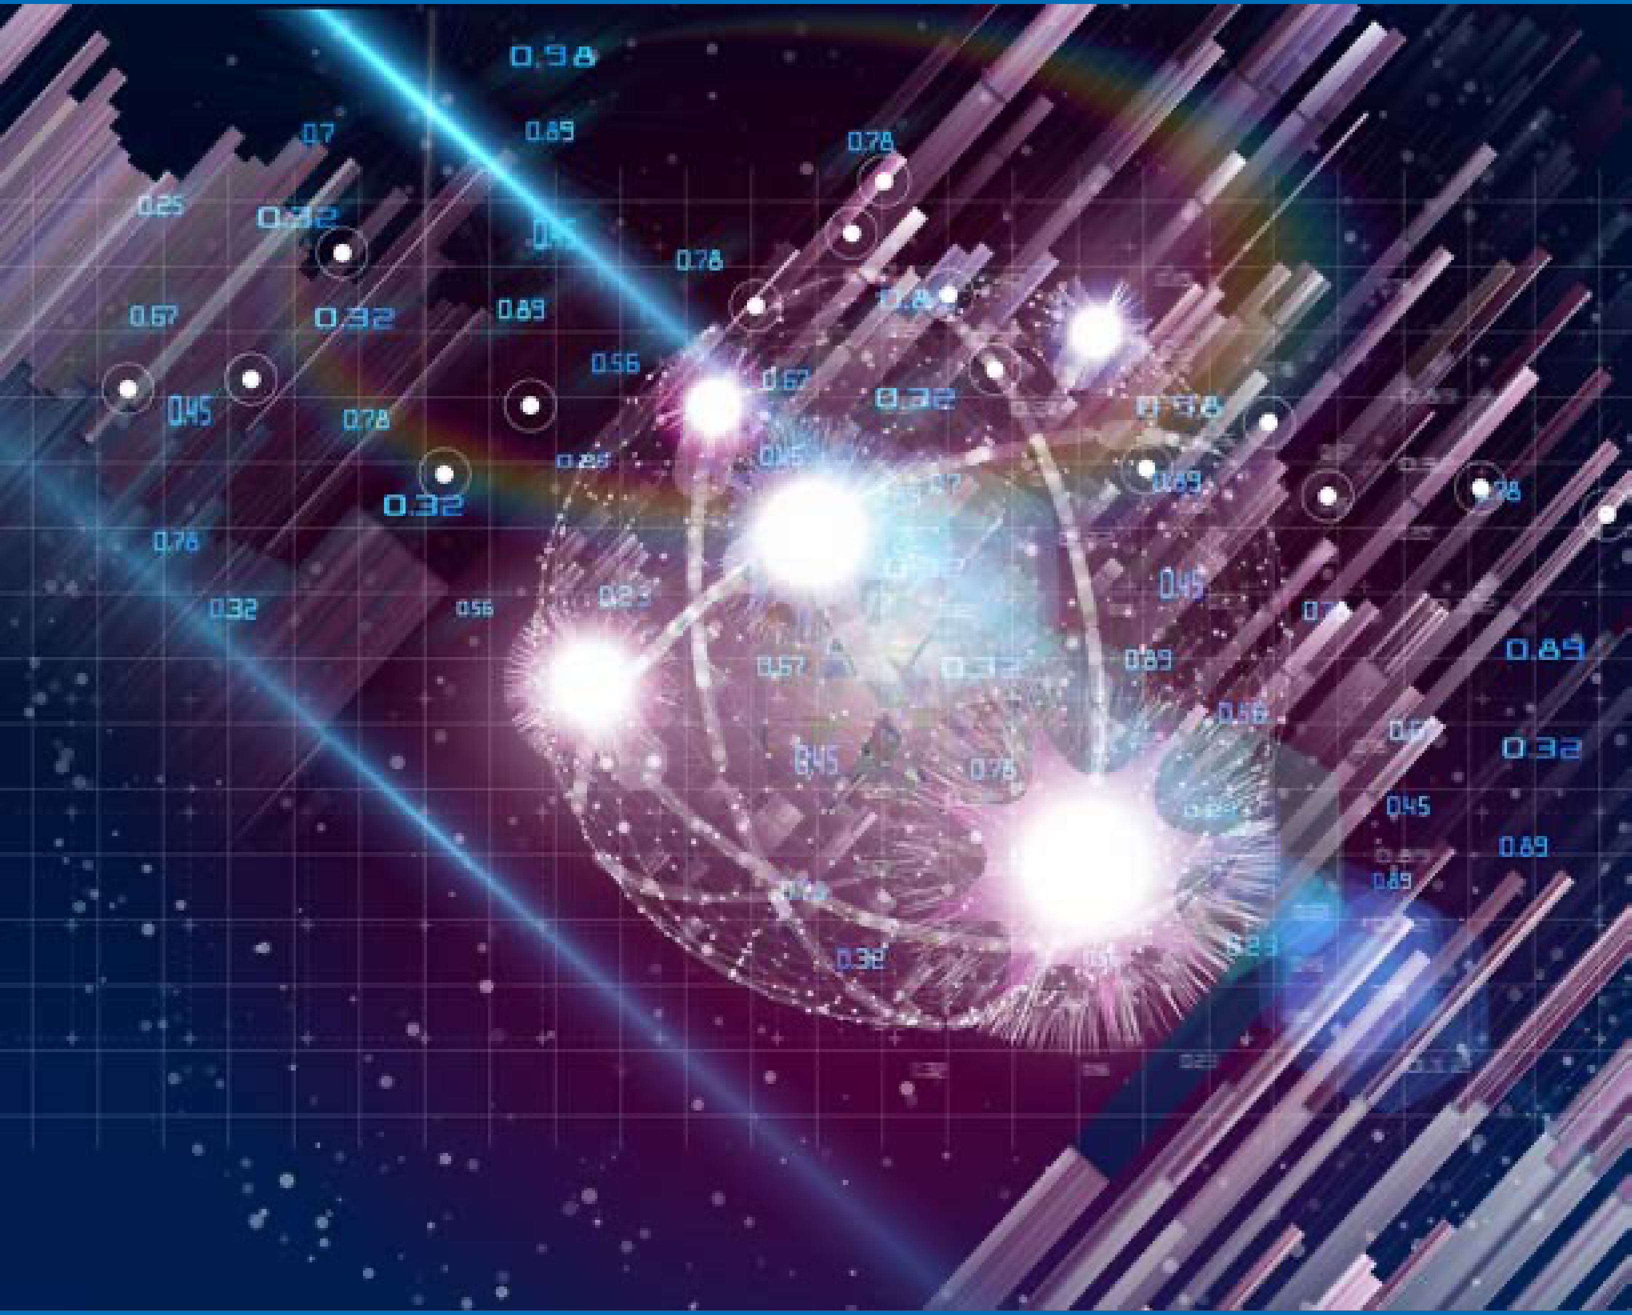

NNSA
National Nuclear Security Administration

LA-UR-22-25073

Issued by Los Alamos National Laboratory for NNSA Office of Advanced Simulation and Computing & Institutional Research and Development Programs, NA-114

For more information, contact Thuc Hoang at Thuc.Hoang@nnsa.doe.gov.

Advanced Simulation and Computing

Computing Strategy

NNSA

Thuc T. Hoang
K. Michael Lang

LLNL

Matthew P. Legendre
J. Robert Neely

LANL

Michael K. Lang
James W. Lujan

SNL

Robert J. Hoekstra
Stephen T. Monk

July 2022

A Publication of the Office of Advanced Simulation and Computing
& Institutional Research and Development Programs (NA-114)
NNSA Office of Defense Programs

ACKNOWLEDGMENTS

This ASC Computing Strategy document is the fruits of the labor of many members of the ASC program at NNSA Headquarters and the three NNSA national laboratories. Listed below are the contributors who have helped the writing team with their technical insights, feedback, and notetaking service, all of which made the document as comprehensive as it is. I would like to especially acknowledge and thank my co-authors, most of whom have been working side by side with me in the ASC program for many years, for their invaluable service in producing this document as an update of the 2013 ASC Computing Strategy. Finally, I would also like to thank Cristina Olds and Marisa Lamb (LANL) for their critical support as the technical editors of the document.

Thuc Hoang

Director, Office of Advanced Simulation and Computing
& Institutional Research and Development Programs

NNSA

Srini Arunajatesan (SNL detailee)
Eric Chisolm (LANL detailee)
Kevin Elzie
David Etim
Simon Hammond
Garry Kuhn (Leidos)
Anthony Lewis
Tina Macaluso (Leidos)
Robert Meisner (Leidos)
James Peltz
Michelle Quirk
Erich Rummel (Leidos)
David Stevens (LLNL detailee)
Emily Simpson (Leidos)

LLNL

Ned Bass
Todd Heer
Robin Goldstone
Jeff Long
John Allen
Todd Gamblin

LANL

Patrick McCormick
Cristina Olds
Bradley Settlemyer
Galen Shipman

SNL

Ann Gentile
Tom Klitsner
John Naegle
Aron Warren

CONTENTS

FOREWORD

The mission of the National Nuclear Security Administration Office of Defense Programs (DP) is to maintain a safe, secure, and effective nuclear weapons stockpile for the United States. Essential to this mission is the Advanced Simulation and Computing (ASC) program's high-performance simulation and computing capabilities that inform critical DP stockpile stewardship decisions. The ASC program's success relies on its ability to provide to the nuclear security enterprise the necessary simulation tools, which serve to enable broad and deep knowledge of the individual processes involved in a nuclear weapons explosion as well as a comprehensive understanding of the complex interactions among these processes. By continually developing and deploying the credible, science-based simulation tools that enable NNSA to certify the current and future stockpile, ASC provides the required confidence in the Nation's nuclear deterrent.

To ensure a reliable nuclear deterrent in the face of a changing geopolitical climate, NNSA must imbue its Stockpile Stewardship, Management, and Modernization programs with responsiveness and agility qualities to stay ahead of emerging threats. The ASC program will be called upon to contribute to the development of new weapons options that could potentially be very different from those in the current stockpile. To enable the nuclear security enterprise to field these options with agility, the ASC program will be required to extend the application of its tools to improve the efficiency of the NNSA production complex. In the near future, the modeling of weapons systems beyond the stockpile's lifetime and simulation of production processes to enable the NNSA production complex to be more efficient will represent a significant growth in the scope of the program, in addition to programmatic commitments.

As the program's scope increases to meet the expanding DP mission requirements, ASC must also contend with major, disruptive computing architecture changes. The ongoing challenges associated with potential decreases in application performance and increasingly complex facility equipment and management practices, as a result of the forthcoming architecture changes, must be addressed proactively in order to meet NNSA's current and future requirements.

The 2022 ASC Computing Strategy addresses these challenges and provides an updated approach to the development and procurement of high-performance computing platforms through a balanced portfolio of systems and technology investments to provide a stable, production-level computing service while tracking progress of the computing industry. Partnering with industry and fielding advanced architecture testbeds and prototypes to keep pace with the technology changes will allow the program to anticipate and prepare for potential disruptions. Sustained, tightly coordinated co-design collaborations with industry partners and other U.S. agencies will enable the optimal use of these technologies to meet stockpile mission requirements.

In summary, this Computing Strategy outlines how ASC will endeavor to stay responsive and agile, in the midst of fast-paced technological changes and priority shifts, to ensure that current and future generations of ASC simulation capabilities and resources will continue to underwrite our nation's nuclear deterrent. ASC will diligently execute its programmatic commitments according to the new NNSA slogan of "Innovate. Collaborate. Deliver." as it aligns perfectly with ASC operating principles since its beginning in 1995.

Thuc T. Hoang

Director, Office of Advanced Simulation and Computing
& Institutional Research and Development Programs

NNSA Office of Defense Programs

EXECUTIVE SUMMARY

As an update to the May 2013 *Computing Strategy*, the 2022 *ASC Computing Strategy* will provide key principles and high-level details describing how the ASC program will execute its plans for research, development, deployment and operation of its high-performance computing systems, user environment, and facility infrastructure for the 2020–2030 decade. While the document is very detailed and comprehensive about the ASC computing environment planned for this decade, the presented information is meant to convey the strategies, and not actual implementations, that the program will strive to execute according to the currently known budget plans and laboratory staff availability.

Below are key guiding principles ASC will follow to execute its new computing strategy:

- Provide a stable, production-level computing service for the tri-lab[1] user community in support of current and future NNSA nuclear deterrent missions;
- Modernize and maintain its tri-lab computing network, as well as local and remote access infrastructure;
- Maintain a balanced portfolio of investments in advanced research, development and deployment of leadership-class capability and capacity systems, architectural testbeds, and prototype platforms; and
- Collaborate with partners in other Department of Energy (DOE) programs, other U.S. federal agencies and international high-performance computing (HPC) entities to share lessons learned and engage in meaningful collaborations, where appropriate.[2]

The ASC computing service, both classified and unclassified, is an irreplaceable resource to the nuclear security enterprise that needs to be continually re-capitalized and upgraded over time. The environment within which ASC operates during this decade, within NNSA and externally, will be more complex than ever due to increasing scope and additional budget constraints. Co-design collaborations with industry and close engagement with other U.S. agencies, in addition to international partners, will continue to be important mechanisms for ASC to directly influence and track next-generation architectural and associated application performance issues.

For ASC to execute its computing strategy of deploying planned HPC systems on a scheduled timeframe with a stable, production-service user environment, it requires a highly technical and sustained interdisciplinary workforce at the NNSA laboratories and Headquarters to ensure delivery of required solutions for the many technological challenges that lie ahead, in support of the nuclear security missions.

[1] Otherwise known as NNSA laboratories: Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), and Sandia National Laboratories (SNL).
[2] See Appendix A for updated Lessons Learned.

1.0 INTRODUCTION

The Advanced Simulation and Computing (ASC) program continues to be a cornerstone of the National Nuclear Security Administration (NNSA) Stockpile Stewardship Program (SSP). The ASC program underpins the nuclear deterrent by providing simulation capabilities and computational resources to support the entire weapons lifecycle from concept and design to dismantlement. In doing so, the program helps sustain U.S. leadership in high-performance computing by fielding top-ranked computer systems and cultivating the workforce needed to support the simulations and the computing ecosystem.

The NNSA mission drivers will evolve over the next decade and beyond. The 2018 *National Defense Strategy* and the 2018 *Nuclear Posture Review Report* call for NNSA to expand its focus beyond sustaining and modernizing the enduring stockpile to also maintain readiness to introduce new options as needs emerge. The 2022 *NNSA Strategic Vision* emphasizes, "we will sustain the current stockpile, undertake comprehensive weapons modernization, recapitalize our nuclear weapons infrastructure, and strengthen our cutting-edge ST&E (science, technology, and engineering) capabilities." Additionally, the NNSA Office of Research, Development, Test, and Evaluation's (RDT&E) Beacon Strategy emphasizes that the RDT&E programs, including ASC, must lead in three areas:

1. Fostering an environment of innovation aimed at improving the stockpile, serving as the proving grounds for new ideas;
2. Applying RDT&E capabilities to enable a more responsive and efficient production complex; and
3. Investing in the next generation to lead the nuclear security enterprise by equipping them with the knowledge, judgement, and state-of-the-art scientific and computing capabilities to achieve the mission.

Developing options for the future through innovation requires ASC tools to be capable of credible application beyond the realm of the current stockpile, putting greater emphasis on predictive capabilities underwritten by strong validation and uncertainty quantification. As ASC tools are deployed for the production complex, new methods, novel models, and production-level support for new workflows are required. Taken together, these requirements represent a significant expansion of the scope of the ASC program. Covering this expanded scope is essential to ensuring that the ASC program continues to support the nuclear deterrent.

The computing activities discussed in this 2022 *Computing Strategy* document are funded by the following program elements:

- Computational Systems and Software Environment (CSSE)
- Facility Operations and User Support (FOUS)
- Advanced Technology Development and Mitigation (ATDM)'s Next-Generation Architecture and Software Development (ASD) product group (until end of FY 2023)

The CSSE subprogram procures and integrates the computing systems needed for weapons simulations. 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 input/output (I/O), storage and networking, and post-processing visualization and data analysis tools. In this subprogram, ASC will continue to pursue advanced research and development (R&D) in next-generation computing technologies and also embark on research investigations of Beyond Moore's Law to include quantum, neuromorphic, and non-complementary metal-oxide-semiconductor (CMOS)-based computing techniques. The 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, cybersecurity, 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. Lastly, the ATDM ASD product group funds the computer science projects that pursue long-term computing goals relevant to both exascale computing and the broad national security missions of the NNSA.

Evolving user requirements during this decade have major implications for the ASC computing ecosystem which the program must continue supporting. Simulations that support stockpile certification and modernization, address Significant Finding Investigations (SFIs), and investigate safety scenarios use large ensembles of multiphysics, three-dimensional calculations with large data movement requirements. This requires a robust, stable ecosystem that can efficiently support numerous large calculations and workflows to minimize turnaround times. The emphasis on improved predictivity for future systems translates to increased numbers of high-fidelity and/or first-principle scientific simulations, along with the integrated multiphysics simulations with potentially large data storage requirements. Furthermore, as these simulations explore the increased use of machine learned models, machine learning (ML)-specific hardware will be investigated carefully to realize the huge potential performance gains.

The use of ASC tools to support the modeling of production processes means that the codes, workflows, and computing ecosystem must support rapid exploration of manufacturing process options. As the NNSA complex explores additive manufacturing techniques, the need to understand the effects of microstructure variability on component behavior could reveal the need for large-scale, first-principle calculations. In all these cases, as new methods, applications, pre- and post-processing approaches, and tools are adopted, the computing ecosystem serving the user community must remain robust and capable of supporting the increasing quantity and variety of workloads.

There are numerous challenges in supporting the demands and increased workload from the user community. NNSA mission requirements and fiscal constraints dictate the tempo and means by which the ASC program acquires planned computational systems. However, external entities such as cloud service providers and data centers, along with industry's response to the end of Moore's Law and Dennard scaling, strongly affect the cadence of commercial computing progress. In particular, vendors who are driven by business and consumer markets focus heavily on graphics processing units (GPUs) and data-intensive computing rather than on high-accuracy numerical simulation. Additionally, the current rapid emergence of machine learning has caused HPC system vendors to veer away from the needs of scientific computations, which historically rely on double-precision floating point performance and non-local data accesses, to dense, single- or half-precision floating-point computing. NNSA laboratory activities funded by ASC, Laboratory-Directed Research and Development (LDRD) programs, and DOE Office of Science (SC) are all looking into mixed-precision computing techniques to maximize performance on lower-precision hardware that vendors are producing for the data-intensive market.

The explosion of the aforementioned new research trends in consumer computing has created a demand for programmers and computer scientists that threatens the talent pipelines upon which the laboratories depend. The ASC program and NNSA laboratories put a high priority on the development of a highly specialized workforce. The laboratories are actively seeking, hiring, developing, and retaining the talent needed to work on the NNSA missions. The unique challenge of the NNSA missions and long-time leadership in HPC are major factors for drawing top talent and young staff to the NNSA laboratories. ASC supports the laboratories' workforce development efforts by its continued funding of the ASC Predictive Science Academic Alliance Program (PSAAP) and the DOE Computational Science Graduate Fellowship, a jointly funded effort with DOE SC. Forging stronger partnerships between NNSA Headquarters and the NNSA laboratories is key to enabling the RDT&E programs to increase return on investment from various DP RDT&E academic programs, including the NNSA Stewardship Science Graduate Fellowship, the Laboratory Residency Graduate Fellowship, the Joint Program in High Energy Density Laboratory Plasmas, the Minority Serving Institution Partnership Program, and the Tribal Education Partnership Program.

Adapting commercial computing technology options for NNSA applications and developing or upgrading simulation tools that can take advantage of the more modern and powerful computing capabilities are challenges that require a coherent, executable platform strategy, accompanied by a simulation development strategy, to ensure stable, production-level computing support to the ASC user base. The ASC program must help to ensure an industrial supply base in the long term and participate in the technology change process to have the computing technology/components available as needed to sustain throughput for the DP mission.

The ASC program also recognizes that both the use of existing open-source software packages and the development and maintenance of new open-source software packages fosters the creation of like-minded communities of software users, developers, and researchers that leverage and extend state-of-the-art ASC software technologies. However, taking advantage of open-source development does not remove the need for in-house talent, so ASC must also renew its efforts with the NNSA laboratories to maintain healthy recruitment and development of the staff needed to execute the mission.

This 2022 *ASC Computing Strategy* concentrates on how ASC will provide the computational infrastructure required by present and future user requirements, bounded by the program's budgetary constraints. The computing infrastructure described herein is a complex tri-lab environment that integrates many types of hardware and software products. Whenever possible, ASC utilizes products from commercial vendors and the open-source software community. However, when the required technology is not available from these sources, ASC invests in internal laboratory research and development (R&D) and non-recurring engineering (NRE) activities with system vendors and subcontractors to close the gaps. Maintaining a specialized workforce at the NNSA laboratories to engage with vendors effectively and strategically to both build and sustain the crucial HPC ecosystem for the nuclear security enterprise is also a key part of the ASC program. In addition to strategic investments, ASC will continue to identify areas to consolidate lab-specific activities into a common, tri-lab environment to better maximize the value of its investments and provide enhanced fail-over capabilities, should a disaster or natural emergency occur at a NNSA laboratory.

2.0 Platforms

For the ASC program, an HPC platform is an integrated system of hardware and software that comprehensively provides the required computing environment (classified and/or unclassified) in which a weapon analyst or designer can run simulations and analyze results. It is not just a computer; it is a host of hardware and software components (e.g., compute and login nodes, networks, file systems, long-term storage, operating systems, compilers, numerical libraries, developer tools, etc.), which are often developed independently from one another by component vendors and deployed by the HPC system integrator. For the ASC platform acquisitions, it is cost prohibitive for the program and resource prohibitive for the vendors to provide all the necessary system software and user tools on the procured systems that the ASC users will need to run their application codes. Therefore, the ASC program is required to expend additional resources to develop the software and tools in-house (at NNSA laboratories) or acquire the needed, supplementary software packages either through third-party vendors or academia. The latter step is a risk mitigation to allow the program to have as much control as possible over its sustainment of a persistent, common computing environment across the laboratories.

In recent years, system integrators and hardware vendors have been embracing the concepts of throughput computing, deep memory hierarchies, general-purpose GPUs (GP-GPUs), accelerators, large-scale networking, and computer architectures that were novel or not feasible for scientific computing five to seven years ago. At the same time, market forces have shifted vendor attention away from large HPC systems to machine learning, cloud computing, data analytics, and mobile computing, especially in the consumer space. The HPC market constitutes only a tiny fraction of the primary computing market—a fact that limits the influence of the HPC community on the direction of future computing technology. ASC continues to face formidable challenges in making strategic investments for its applied research and platform procurement portfolios, while balancing an increasingly large computing scope with a decreasing buying power.

2.1 Principles

The challenge of producing and deploying necessary computing resources for this decade requires an updated approach to the 2013 *ASC Computing Strategy* to cover the following principles:

- ASC will provide a stable, production-level computing service for the tri-lab user community to ensure current and future nuclear deterrent missions of the nuclear security enterprise are met. The focus areas for this decade will be code efficiency, performance portability, and user productivity.
- ASC will maintain a balanced portfolio of systems with advanced technologies, commodity technologies, and advanced-architecture prototypes. The focus will continue to be on transformative co-design[3] of the application requirements on computing architectures and software, and vice versa.
- ASC will partner with industry to keep pace with the fast-moving, disruptive technology changes that are not inherently aligned with scientific computing needs. The focus will be continuing advanced architecture investigations, relating to both hardware and software issues, with the goal of achieving application optimization for performance and efficiency.

2.2 Platform Classes

The ASC platform strategy is built around deploying a set of platforms that strike a careful balance among delivering reliable production cycles, pushing the boundaries of current technology, and looking beyond the horizon to what is coming next. This approach to balanced risks has served ASC well over the last 25 years and is revisited with each Computing Strategy publication. In this strategy document, we retain the Commodity Technology and Advanced Technology Systems, and add the third element called Advanced Architecture Prototype Systems to formally recognize the SNL Vanguard program, which was established in 2017. The Vanguard program is responsible for keeping ASC at the leading edge of the computing ecosystem by exploring promising new technologies that are emerging as possible future forerunners in the continuous race to increased performance and productivity. All these system categories are focused on ultimately delivering production-level computing cycles to the ASC user community.

- **Commodity Technology (CT) Systems:** These are workhorse, production-service, general-purpose systems that provide stable compute cycles to the NNSA design and analysis community. These will run the Tri-lab Operating System Stack (TOSS) and persistent, common software environment, with annually exercised capabilities for back-up, data recovery, and remote mission continuation in case any of the NNSA laboratory HPC centers become unavailable for an extended period.

[3] Ang et al., ASC Co-Design Strategy (Sandia National Laboratories, 2016).

- **Advanced Technology (AT) Systems:** These are leading-edge architectures that can solve the most demanding simulations for our stewardship mission. They incorporate newer technologies that push the limits of the ASC program in terms of facility requirements, software infrastructure, and applications. They will have NRE investments prior to system deployment to support acceleration and maturation of critical vendor technologies, for both software and hardware, for a solid, production-ready platform to meet the ASC code needs. The experiences and lessons learned from each AT system deployment will inform ASC of possible future scaling and performance capabilities, as well as issues in advance of what will show up in the workhorse CT systems.
- **Advanced Architecture Prototype (AAP) Systems:** These are node-level testbeds, system-level prototypes, and pre-commercial hardware/scaled-up systems. Investments will be made to evaluate and accelerate the development of emerging technologies, which have the potential to dramatically accelerate mission workloads, to increase their viability for future large-scale production platforms. The goal of these projects is to reduce the risk of deploying unproven technologies by identifying gaps in the hardware and software ecosystem and making focused investments to address these gaps when moving from small-scale testbeds to potentially large-scale, production-service systems.

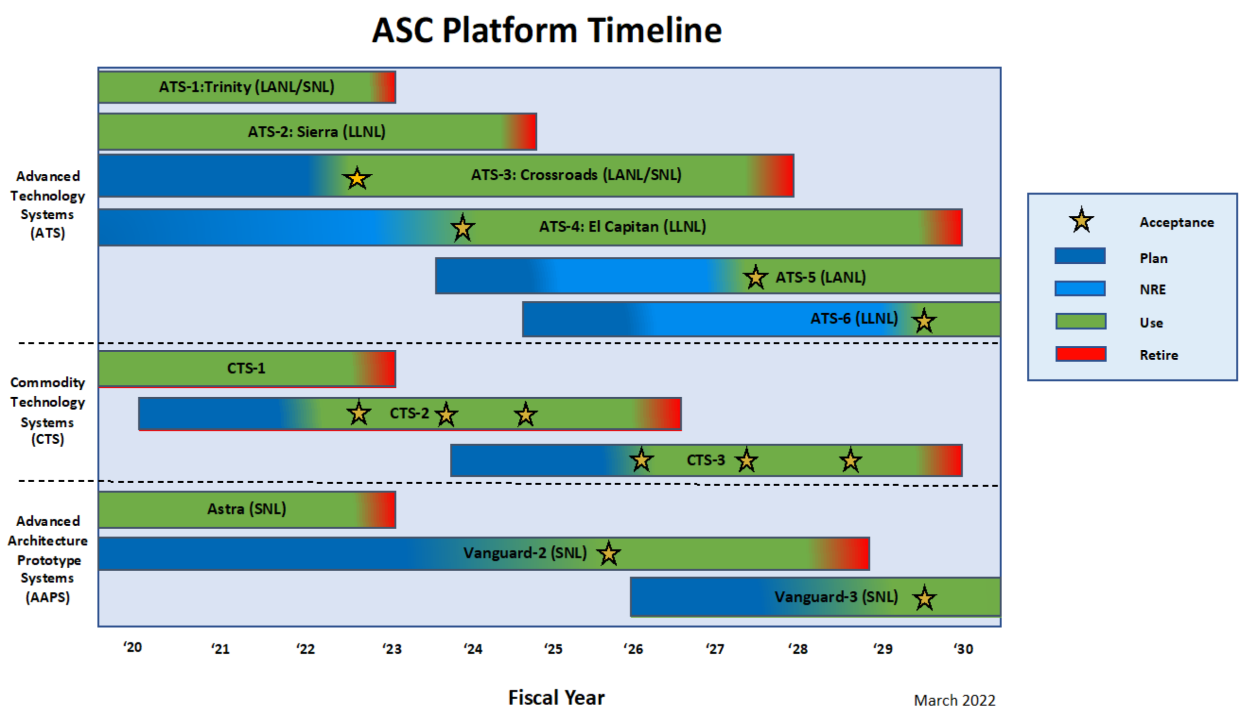

Figure 1. ASC Platform Acquisition Plan for FY 2020–FY 2030. AT systems will alternate between the NNSA laboratories in California and New Mexico roughly every three years. The cadence for new generations of CT systems will be every four years and the cadence for Advanced Architecture Prototype systems is every three to four years.

A key element to prepare for AT-class systems is to deploy early-access hardware at the laboratories to allow porting of applications, analysis of code performance, and maturation the vendor software stacks. Early-access systems are included as part of the overall AT system acquisition strategy and include all three labs' requirements to ensure all ASC simulation codes have the needed lead time and resources to prepare for advanced technology systems.

To address the performance, energy, and resilience challenges for ASC applications, ASC will direct its NRE and industry R&D investments towards advanced developments of new AT system technologies that will support highly complex, multiphysics applications, some of which still employ asynchronous message passing and ensemble-calculation techniques. Recent NRE investments by ASC have focused on improvements to programming environments and math libraries, which increase the ability of mission applications to take greater advantage of the hardware performance. Whereas NRE is an explicit part of an AT system procurement, other industry R&D investments are pursued outside of the NRE activities with the intent to drive technologies to a maturity level beyond which they can be considered in multiple future procurements. Since ASC has limited funding for NRE and technology developments, prioritization will be given to investment in advanced architecture technologies that help preserve the value of and enhance the performance of the ASC weapons code base.

2.3 Expectations and Characterizations

Each platform class comes with a unique set of minimum required capabilities that will guide the ASC HPC system acquisition strategy during this decade.

2.3.1 Commodity Technology Systems

The CT systems will continue to

- Provide a common, persistent tri-lab computing environment;
- Meet mission requirements for simulations that do not require the extreme scale of AT systems;
- Operate at full-system availability and stability, including minimal delays between system delivery and general availability;
- Maximize value in acquisition and operation; and
- Minimize code development costs.

Figure 2. CT systems are supercomputers built from commodity parts. The pictured Magma system at Lawrence Livermore National Laboratory has 1,544 Intel Cascade Lake CPUs and achieves 5 petaFLOPS. (LLNL)

Given these expectations, certain platform characteristics naturally arise. The CT systems will predominantly consist of commodity-level hardware and software, employing computing technologies that the ASC integrated design codes (IDCs) can utilize without necessitating disruptive code adaptations. However, customization of specific hardware or software technologies may still be warranted if a relatively small investment would bring a large benefit to the user community.

To further support reliable service to the users, CT systems are designed and provisioned to maximize availability to the end users and are operated as a production-level resource. Disruptions due to system maintenance activities are kept to a minimum. If problems should arise during system maintenance activities, the entire platform will not be rendered unavailable and will be returned to its production state as quickly as possible while the problems are resolved.

2.3.2 Advanced Technology Systems

The AT systems will continue to

- Meet mission-support computing needs for the most challenging engineering and physics simulations required for predictive capability, supporting both increased fidelity and quantification of uncertainties;
- Incorporate new computing technologies that provide benefits (e.g., scale, speed, stability, energy efficiency, programmability, manageability), which are not presently available in existing commodity offerings; and
- Cultivate and develop pathfinding technologies that could appear in future CT systems.

The benefits of new computing technologies will be measured from the perspective of the ASC code developers and end users. This means that a faster processor or specialized architecture may not exclusively be the appropriate technology for the users' mission computing needs. Other architecture capabilities may be more beneficial to ASC users, such as the following:

- Increased balanced memory capacity and bandwidth,
- Integrated memory architecture,
- Higher I/O bandwidth, and
- Improved network performance.

AT systems represent significant investments for the ASC program in terms of computing capabilities. As such, the amount of preparation, planning, and technical project management necessary to deliver cutting-edge technologies is considerable and requires greater lead time than that for CT systems. Also, to maximize the opportunity for success for code team readiness at the time of system deployment, ASC invests in early-access systems and tight collaborations with AT system integrators and technology providers well in advance of the final system being deployed.

Mechanisms such as a Center of Excellence (COE) facilitate the transition of applications to new system architectures by pairing expertise from system vendor partners with the NNSA laboratory application developers to focus efforts on porting and optimizing applications and system software in concert with novel hardware. This has significantly reduced the time between acceptance of a system and its productive use by NNSA mission applications and is now a standard practice for the AT system procurement and deployment process.

Once an AT system becomes generally available to the tri-lab user community on the classified network, time on the machine is reserved through a lightweight proposal and cycle allocation process, called the Advanced Technology Computing Campaign (ATCC), that allows programmatic tri-lab priorities to dictate usage of that tri-lab AT system resource.

Figure 3. Crossroads will be the next ASC ATS platform deployed at LANL in 2022. (LANL)

2.3.3 Advanced Architecture Prototype Systems

The AAP system platforms are a new class of systems in the program, focusing on the following:

- Evaluation of new architectural and system software technologies with regards to viability and impact on mission applications,
- Deployment of new technologies at scale to allow evaluation with mission applications in the classified environment, and
- Cultivation and development of pathfinding technologies that could appear in future AT and/or CT systems.

In particular, new technologies expected to impact future mission workloads include the following:

- Specialized accelerators focused on kernels in mission applications,
- Tightly integrated heterogeneous solutions,
- Machine learning/AI accelerators, and
- Highly elastic, cloud-like solutions, utilizing technologies such as containers.

AAP systems are not required to focus on delivering general-availability, production-level service, although they can be transitioned to the classified, production-computing environment if they prove to be of high value to the mission workloads. These systems are to create an opportunity to take high-risk, high-reward technologies and develop and mature them in collaboration with vendors and the R&D community. When successful, these technologies will then be a high priority for maturation for future AT and CT systems.

Figure 4. Vanguard Astra, ASC's first at-scale advanced architecture prototype system and the first petascale Arm-based supercomputer.

The ASC program has a long history of making strategic investments that foster the development of advanced computing technologies necessary to improve the scalability and performance of HPC platforms for scientific computing in general and for NNSA simulation needs specifically.

Since 1995, when ASC's predecessor program Accelerated Strategic Computing Initiative (ASCI) was established, the program's investments in advanced technology development and NRE have supported both hardware and software efforts to accelerate the availability and increase the scalability of technologies that the U.S. computer industry would not otherwise provide. The computing ecosystem is advancing in multiple interesting, complex directions and at a pace not seen before due to commercial demands in the data center, cloud computing, and mobile computing sectors. One of ASC's guiding principles is to collaborate with vendors in co-design of computing technologies to benefit from this innovation wave and coax these technologies in directions suitable for advancing our stockpile mission and the HPC market. Recent expansion of technology paths for computing, such as machine learning/artificial intelligence (ML/AI), quantum computing (QC), and other Beyond Moore's Law designs, significantly diverge from current HPC architectures. The ASC program is increasing strategic investment in these areas to drive forward innovative computing designs that leverage these new opportunities in the HPC industry.

While ASC HPC systems have utilized highly customized system interconnects to maximize scalability and performance of large-scale mission workloads, nearly all the recent systems have utilized commodity processors. In the future, the program is exploring even more customized approaches to processors, accelerators, and networks to better match with the workloads and maximize code efficiency. Close collaboration between ASC and its sister HPC program in DOE SC, the Advanced Scientific Computing Research (ASCR) program, along with other federal agencies, is intended to magnify the impact of these efforts across the federal government and the U.S. technology sector.

As a part of the increased emphasis on design exploration, the ASC program has expanded the scope of these research activities and defined a strategic lifecycle of R&D investments. This lifecycle is shown in Figure 5.

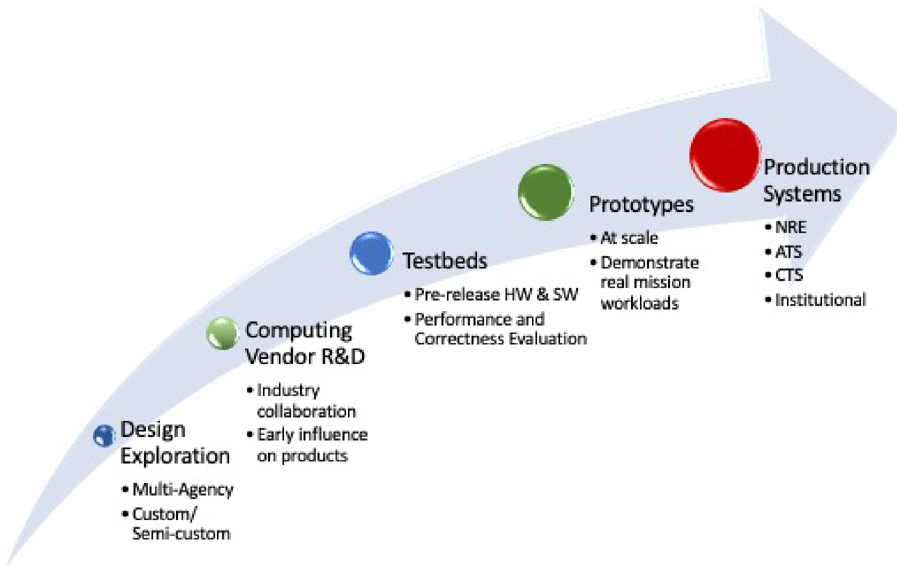

Figure 5. ASC investment pipeline to accelerate the availability of advanced architectural features for production systems.

ASC will apply investments in areas unique to the HPC user community, provide early access to promising new technologies to allow the ASC codes to adapt, and provide NNSA computer and computational scientists a direct pathway to have co-design dialogues with the HPC industry collaborators through these strategic activities:

- *Design Exploration*: Changes in the computing industry, especially in the design and fabrication of semiconductor devices, have made customized designs significantly more feasible and affordable for the federal agencies. ASC has invested in simulation capabilities for computing hardware that are being used to evaluate the potential benefits of new designs.
- *Computing Vendor R&D Partnerships*: In the joint DOE Exascale Computing Initiative that began in 2016, ASC and ASCR together invested over $250M in three-year PathForward R&D contracts with U.S. computing vendors to accelerate the development of new computing technologies to meet DOE mission work. These vendors (AMD, Cray, HPE, IBM, Intel, and NVIDIA) also matched over 40% of total project funds for an additional $150M of investment. These contracts concluded successfully at the end of 2020 and the ASC program will make additional R&D investments, albeit at lower funding levels, to accelerate maturation of needed computing technologies for its NNSA missions.
- *Node- and Rack-Scale Testbeds*: The NNSA laboratories regularly procure small-scale testbeds of early-access hardware from computing vendors. This allows the program to both prepare the software teams to utilize this hardware in the future and deliver feedback to the vendors on issues with the technology that must be resolved before it can become a viable option for a CT or AT system procurement.
- *At-Scale Prototypes*: This is the newest part of the research lifecycle for the program embodied in the AAPS Platform Class (reference 3.2.3). Prototypes are system deployments “at scale” that are intermediate-sized systems allowing both the hardware and software to be exercised on real mission workloads. The SNL Vanguard program successfully deployed the first ASC prototype system, named Astra, which is an HPE system and was the first petascale Arm-based supercomputer in 2018.

- *Non-recurring Engineering (NRE)*: NRE contracts are part of the first phase of AT system procurements. NRE is used to fund additions to or customization of technologies that the general marketplace may not support due to their unique characteristics favoring NNSA workloads. It is also used to accelerate the timeline for development of key technologies. Within the first year of or two before an AT system is deployed, these contracts are utilized to complete final maturation of crucial technologies which will increase the performance of and reduce engineering risks associated with the system.

Additionally, development of co-design capabilities within the ATDM (Advanced Technology Development and Mitigation), CSSE (Computational Systems and Software Environment), and FOUS (Facility Operations and User Support) subprograms will identify multiple options for new capabilities that can and should be implemented in future computing technologies. We intend to quantify the beneficial impact of adopting these changes while reducing the barriers and risks to industry adoption of innovative advances. R&D partnerships with vendor partners are a key element of this strategy; active interaction between the industry partners and the ASC co-design projects is required to enable laboratory scientists and engineers to have early performance data and gain insight into impacts of the advanced architectures on the ASC codes. The ASC co-design approach is detailed further in Section 4.

3.0 Computing Environment Ecosystem

When deploying an HPC platform, it is important to consider the environment in which those incoming AT and CT systems will operate. A successful ASC platform must integrate within an ecosystem of workflow processes, user tools suites, storage systems, and communication devices. Additionally, the facilities infrastructure and operational support services are critical to a successful deployment. This environment is constantly changing. Support for heterogeneous computing elements is currently taking significant efforts and HPC systems look to be more heterogeneous in the future. This complexity is reflected in all levels of the software environment ASC needs to support but is especially evident in software development, debugging tools, and the programming environments, all of which will require strategic partnerships with system and independent software vendors.

Due to the comprehensive and broad landscape of the ASC computing environment discussed in this section, the following subsections 3.1–3.3 will have a challenge/strategy theme provided for a short summary.

3.1 User Environment and Workflows

There are many types of workflows; two of the most prevalent are *developer workflows* and *scientist/user workflows*. While variations exist, the typical developer and user workflows are shown in Figure 6. In a user workflow, users develop inputs via meshing tools using common models; manage simulation runs or ensembles of these runs; and then focus on data analysis, visualization, and data storage. The developer workflow is similar and shown on the right of Figure 6. Ultimately, the ASC environment needs to be persistent and consistent, to the extent possible, among all currently deployed and future ASC platforms as different-sized problem sets are required to run interchangeably on either AT or CT systems.

Challenge: In the current ASC environment, solutions often require orchestration of multiple applications, services, and platforms. Accelerated, shorter time-to-solution is key to user productivity.

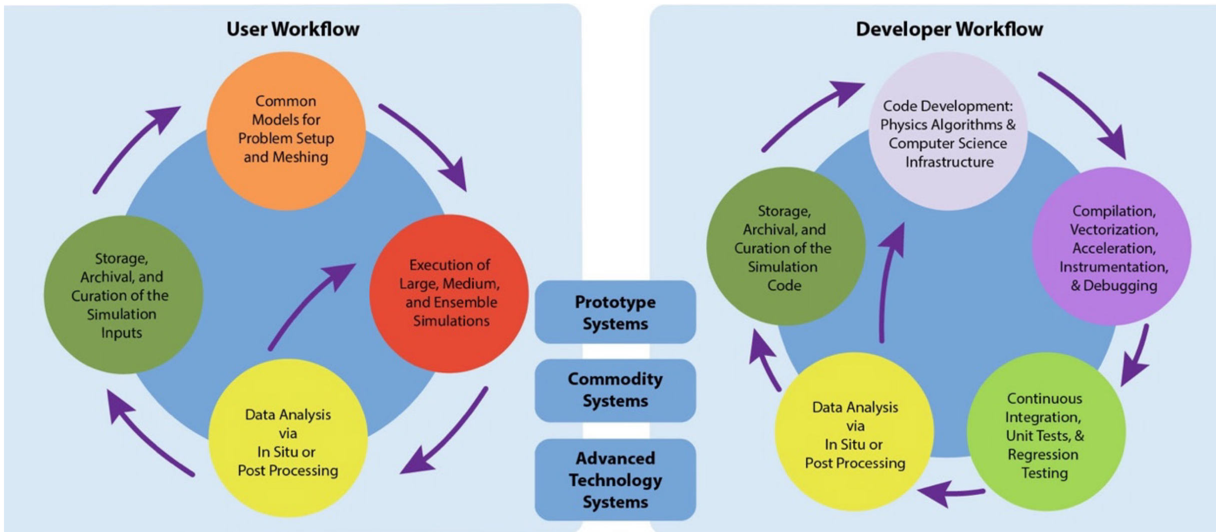

Figure 6. Typical ASC User and Developer Workflows. The ASC computing environment is delivered by the CSSE, ATDM, and FOUS sub-programs and provides the tools and resources to support complex workflows on the three types of supercomputer platforms.

Developers are concerned with rapid evaluation of code changes and improvements on the variety of computing resources that ASC provides for the tri-lab community. They rely on a software development best practice of *continuous integration*, where code changes are automatically tested on various computing platforms with a set of problems that exercise the breadth of functionality of the simulation applications.

User or scientist workflows start with the question (or problem) of stockpile importance being posed, then proceed through problem setups, simulations of the parameter space, and in situ or post-processing of the results for analysis, all as part of producing the answer to the question. This workflow involves multiple storage systems for management of data, with prioritized scheduling of the individual application jobs that make up the overall simulation and many avenues of analysis. These jobs can be of various sizes: many small ones consuming one or two nodes of a supercomputer or large jobs saturating an entire HPC platform over multiple months.

Another type of emerging workflow integrates machine learning into the simulation loop (see Figure 7). In one example of an ML-coupled workflow, ML techniques are used on simulation data to pick a few prime areas for expensive full-simulation runs, thereby reducing the parameter space to be investigated.

Many workflows benefit from a common computing environment provided by the diverse computing resources at the three NNSA laboratories. Best practices are leveraged for technologies to support these workflows, such as containers, data labeling, provenance, reproducibility, trustworthiness, composability, data curation, and archival. Performance instrumentation of these workflows is important to enable workflow optimization efforts. ASC's goal is to greatly reduce turn-around time of user workflows over the next few years. For the ASC environment, it would be ideal to achieve the same ease of use for small job workflows, parameterized ensemble runs, and full-system runs for the user community.

Key to supporting workflows are the integration of data and analysis at large scale and the movement of data between workflow components. Hence, traditionally input/output services have been key to

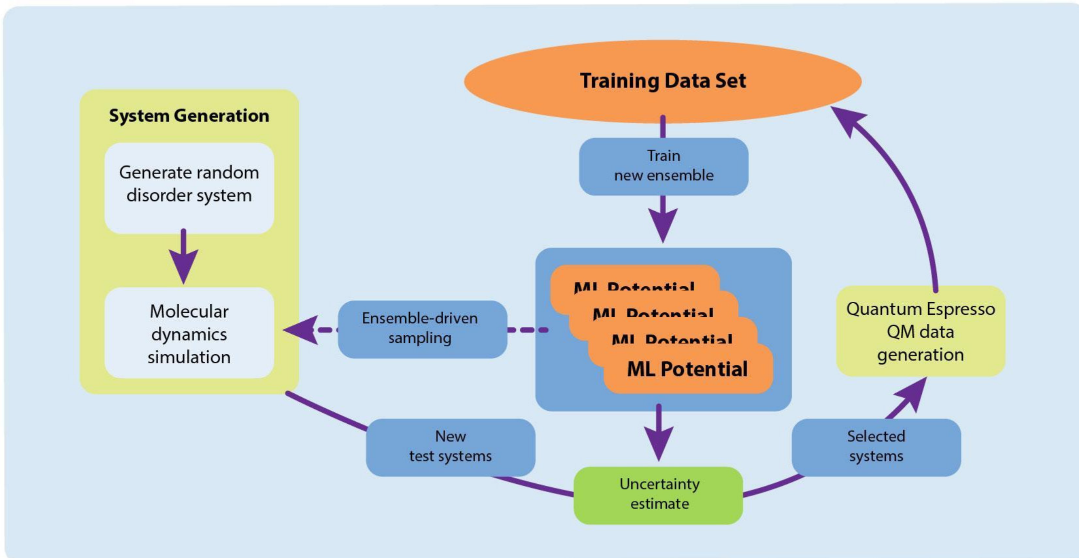

Figure 7. A hybrid workflow involving a molecular dynamics simulation coupled with machine learning.

workflow capability and robustness, and these services have centered around storage technologies and file system abstractions. At present and moving into the future, new node-based persistent storage capabilities, such as non-volatile memory, are beginning to change this paradigm, and future workflow capabilities are likely to start bypassing file system approaches in favor of direct data transfer between components.

Strategy: Workflow tools and technologies designed to orchestrate and manage simulations that encompass multiple applications, platforms, and services will be deployed to greatly enhance developer and end-user productivity. A focus on greatly decreasing the turnaround time of end-to-end workflows is critical.

3.2 Software Environment

The ASC program has historically increased its tri-lab computing capacity by both deploying higher node counts and by leveraging increasingly powerful individual nodes. While both trends are likely to continue, most future growth will come from increasingly powerful nodes. Accelerators, such as GPUs, are shifting the degree of parallelism on a single node from the realms of 100s of parallel-processing units into the realms of 10,000s of parallel-processing units. In addition, new software technologies, such as machine learning, are driving further shifts in ASC applications and changing how application problems/questions are asked. These compounding shifts are requiring the ASC software environment to be more flexible, better integrated, more diverse, and highly efficient, while maintaining the fundamental capabilities needed to deliver on ASC commitments to the users.

Challenge: ASC will enhance its world-class software environment to provide the most capable technologies to users and developers supporting the ASC program's national security missions. Hardware platforms are not useful to end users without a functioning software stack.

ASC's software ecosystem is built on an overlapping combination of open-source, commercial, and ASC-developed software. The ASC program acts like a conductor over a diverse software stack, orchestrating various pieces of software into world-class supercomputers. For example, it works with hardware vendors who may know the best way to write low-level software for their devices but do not understand how ASC HPC workloads will stress them. It co-designs software such as schedulers and tools with system integrators and vendors. It leverages and contributes to open-source software to target HPC, and ASC takes a leadership role to develop software that meets its needs and the needs of the broader HPC community where appropriate.

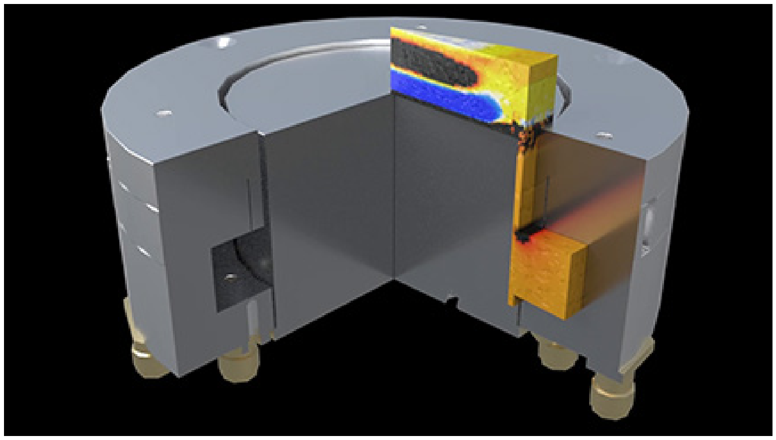

Figure 8. EMPIRE simulations are used to simulate the radiation/plasma environment near the detonation of a nuclear weapon to ensure the survivability of U.S. assets. Shown is the rendered validation experimental CAD geometry as fielded on the Z pulsed power X-ray facility with an overlay of a plasma generated in an EMPIRE simulation. The extreme-scale computational resources of Trinity, Astra, and Sierra are necessary to model the complex dynamics generated by the X-ray generation from the Z facility. This validation experiment has a simple geometry elucidating the basic physics of the X-ray environment, which includes photoelectron emission from surfaces and subsequent ionization of the gas. (SNL)

3.2.1 Programming Models

Programming models provide abstractions that allow developers to write high-level codes that focus on the physics and mathematics of the problems they are solving, which are automatically mapped into efficient levels of parallelism on hardware. Maintaining a leadership role in programming model R&D is critical to ASC as the NNSA weapon applications will need to be performant in the rapidly changing hardware ecosystem. The ASC program has made significant investments in updating existing programming models, such as OpenMP, and researching new ones, such as RAJA and Kokkos, for GPUs and accelerator-driven computing. But significant work remains. Programming models need continuous porting to new hardware and additional research is needed when hardware undergoes fundamental shifts. When certain algorithms do not map well to existing programming models, innovative research in new abstractions and implementations may also be required.

3.2.2 Development Tools

Tools are a force multiplier in software development. Debuggers and performance analysis tools help developers hunt through millions of lines of code to identify bugs and performance issues. Continuous integration tools, which provide constant testing, can identify new problems soon after they are written. The ASC program leverages both sets of tools, from the broader software development community and from the vendors. However, tools frequently need to be adapted to the ASC program's unique environments and programming models. Typical debuggers and performance analysis tools are not designed to run at the scale of ASC systems. Supporting HPC programming models is typically not a priority for tool developers. Further, ASC's trailblazing efforts to migrate onto new architectures sometimes means the hardware arrives before tools are production ready. The ASC program addresses these challenges with software R&D and investments, such as funding a performance tool team to add support for an important programming model or architecture. Sometimes these challenges need fundamental research such as building debuggers that can scale to thousands of nodes.

3.2.3 System Software

System software such as operating systems, schedulers, compilers, and file systems underlie all of the ASC computing infrastructure. The broader community's investment in system software dwarfs what the ASC program, or even the HPC industry, would be able to solely support. This provides opportunities and capabilities that were not available in the ASC program's early days when it wrote much of its own system software but also creates new challenges in influencing the broader community. For vendor-provided software, the ASC program leverages large system procurements to focus their development on HPC. Open-source system software is particularly valuable, as the ASC program can cultivate local expertise and help build broader developer communities to advance the software. However, significant challenges in system software remain.

The users' ability to store data on file systems is falling behind their ability to compute data, which is being driven by higher levels of node performance and data-intensive computing jobs and services to support machine learning. While there may be some further gains with respect to increased I/O parallelism, ultimately hybridized solutions will need to be explored and implemented. New approaches will likely be required for "scratch space" (temporary storage space) while calculations are running and for the backup and recovery systems. These problems cannot be solved with new hardware solutions alone. Sophisticated software packages will be needed to make the hardware accessible and high performing.

3.2.4 Software Integration

To a greater extent than other computing disciplines, HPC relies on the ability to integrate applications, libraries, tools, and system software written in a wide range of programming languages. Application developers rely on hundreds of dependency libraries, both internally developed and open source. System software also relies on hundreds of additional packages, and machine learning stacks are their own equally complex ecosystem. With the advent of containers, the line is blurring between maintaining a traditional system software stack like TOSS and maintaining an application's own dependencies. Moreover, programmatic codes must reuse and share dependency libraries to leverage new features and avoid duplication of efforts.

Integrating all of the aforementioned software is a daunting task, and it is not always the case that open-source developers or vendors have optimized their stacks for bleeding-edge HPC systems. Thus, ASC develops and maintains tools such as Spack that enable software integration across all levels of the software stack and ensure that open-source software can be easily built and integrated on HPC systems by facility staff, user support teams, application teams, and users of ASC HPC resources.

The development and adoption of open-source software is a major productivity multiplier for ASC, as it allows for a broad community to contribute to improvements to open-source software originating from the labs, as well as adoption of software developed and maintained elsewhere. The use of open-source software is not without risk however, as it places ASC teams at risk of discontinued development and maintenance exposure to others' poor software quality practices or security vulnerabilities. These risks are well understood and managed as part of a software quality assurance process. Likewise, tools like Spack further help ensure that only well-tested and verified versions of externally developed software are integrated into a final production code.

3.2.5 Data Analysis and Visualization

Higher mesh resolution, increased data dimensionality, and finer time increments stress current visualization tools and supporting infrastructure, such as I/O. New techniques such as in-situ and in-transit data analytics offer promise for relaxing the infrastructure requirements imposed by post-processing visualization. Another dimension of data analysis is the ability to compare sets of related data. This is crucial for code validation, as well as a thorough understanding of the physics.

3.2.6 Math Libraries and Application Frameworks

Many ASC simulation applications are designed around common mathematical and computational foundations. Rather than rebuild these foundations for each application, ASC has factored its codes to rely on math libraries and application frameworks that can be shared across applications. For common mathematical operations the ASC program relies on externally developed software. For example, most hardware vendors provide math libraries for common linear algebra operations, which usually have unmatched performance on their native hardware. ASC also maintains a leading R&D program to develop its own application frameworks to meet its more specific needs. For example, the Hypre library and Trilinos framework, which were both developed in the early days of ASC, are still heavily used within the DOE community. R&D continues to produce new application frameworks such as the MFEM library, a powerful high-order methods solver that has been developed by LLNL.

Strategy: ASC develops and deploys a required software stack, which is developed internally by the tri-labs and by external partners such as system or third-party vendors, in order to support the productive and efficient development and execution of its mission applications on the broad range of deployed platforms.

3.3 Facility Operations and User Support (FOUS)

Challenge: Increasingly powerful and complex HPC platforms place significant demands upon supporting facility infrastructure and operations.

Facilities and operations are the support foundations of the ASC computing infrastructure, providing space, power, cooling, operations, and systems monitoring to increasingly complex and densely packaged computing and storage systems. The current computing-associated demands, such as power, have already resulted in the need for changes in the facility hardware and in management practices. These will only increase as systems increase in heterogeneity and complexity, as ASC is increasingly called upon to support ML applications with more irregular resource demands than MPI[4]-based applications, along with additional data management and storage services requirements. Additional changes will also be required as ASC must increasingly embrace remote computing. The ultimate goal of successful operations is maximizing the availability of computing resources to the end users with minimum expense and effort required of the computing center. Additionally, due to the important missions the ASC codes support, the NNSA labs are required to implement a data back-up and recovery strategy.

3.3.1 Networks

The networking community at the NNSA laboratories has recognized that the existing 15-year-old 10Gb/s DisCom wide area network infrastructure is insufficient to support the classified remote computing needs for the upcoming platforms. As part of an ASC tri-lab effort, ASC is upgrading DisCom to support 100Gb/s. This effort will require the development and deployment of a new architecture that supports the

[4] Message Passing Interface (MPI) is a standardized and portable message-passing standard designed to function on parallel computing architectures.

100Gb/s MACSec Type 1 encryptors. Over the next few years, new tools will be developed and deployed to support effective data transfers at this performance level. The local infrastructure at each lab will be upgraded to enable an end-to-end 100Gb/s data path.

The ASC tri-lab community is also working closely together to improve the user experience of the unclassified collaboration environment. This involves network improvements, better authentication mechanisms, and more consistent services for the unclassified environments operated between the labs.

3.3.2 Data Management

Data management is a growing challenge for ASC. Larger and faster machines are producing unprecedented amounts of data that need to be labeled, tracked, shared, properly stored, and eventually retrieved. Further complicating data management are new multi-tiered storage systems where users can store data everywhere from fast and ephemeral on-node storage, to databases, to medium-term shared file systems, to long-term and slow-rate archival storage. Previously, data management was done with humans-in-the-loop processes, where individuals or teams had their own favored processes for data management. ASC is researching better ways to consolidate and automate these processes with workflow tools and methodologies that will keep data safe and organized for future use.

3.3.3 Continuous Integration and Release

To ensure code correctness, test suites are run on testbeds and production systems (AT, CT, and AAP systems). Hundreds of thousands of tests per night per code suite per platform are typical. Determining additional resource requirements to ensure throughput is difficult because the tests increase in number and complexity, and the test runs compete with production runs which can affect performance due to competition for shared platform resources (e.g., network, file system, etc.). System configurations, performance analyses, and operational policies must continue to be developed to ensure the throughput of the nightly tests, as well as the production workload.

Additionally, ASC code dependencies frequently include DOE-SC-funded mathematical libraries and third-party software, whose code repositories may reside at remote locations. System configurations to enable remote code checkouts while ensuring security postures must be developed and deployed.

3.3.4 Container Technology

Container technology, also known as just "containers," is a method to package an application so it can be run with its dependencies, isolated from other processes. The major public cloud computing providers, including Amazon Web Services, Microsoft Azure, and Google Cloud Platform, have fully embraced container technology. As the programming environment grows more complex and as ASC codes are increasingly reliant on additional software, containers can ensure a user-defined and user-maintained environment that can persist through code revisions and enable performance and result comparisons. The use of containers may ease the complexity of supporting continuous integration since much or all of the supporting software would be part of the containerized environment. Further, containers may facilitate code development on platforms different from the target architecture.

However, support of containers places new challenges on system operations and user support:

- *Operating System and Programming Environment Maintenance* – With more of the supporting software moved into containers, responsibility for content shifts from system support roles to the users. More container-specific interaction is expected between container developers, current programming environments, system administration teams, and user support teams, to ensure correctness and diagnostic support.

- *Launch and Runtime* – Infrastructure for launching containers is currently under investigation. There are several competing container technologies that must be concurrently supported. New methodologies for performance analysis, tuning, and monitoring of applications in containers will need to be developed.
- *Containers for Management* – Containers for the system management plane are being investigated since they hold the promise of increasing scalability, reliability, and resilience. This will, however, add complexity to the overall management processes of ASC resources.
- *Provenance and Security* – As more of the supporting software comes from user-accessed sources, security concerns increase. These are further exacerbated for multitenant allocation schemes. Shared resource access and container attacks are ongoing technical challenges.

3.3.5 Resource Management on Heterogeneous Systems

The need to support increasingly diverse applications and associated algorithms (see Section 4.1) on an increased diversity of hardware can most be efficiently accomplished through a co-design approach (see Section 4.0). This increases the complexity of system operations. A rising challenge is the need to manage systems consisting of heterogeneous hardware throughout the system. New operational methodologies will need to be developed to ensure the correct match of applications, and even phases within an application, to this variable hardware. These include mechanisms for maintaining and launching multiple concurrent architecture-specific software images, and more complex and dynamic scheduling and allocation logic.

Increasingly powerful and resource-heavy components, such as GPUs, may be largely idle when applications not designed to utilize these resources are executing. Particularly within a single platform, these resources will require more complexity in their use and operation since there is no single rule of thumb as in the past with traditional applications and hardware. For example, ML-based applications have more dynamic and less balanced resource demands than historical HPC MPI-based applications. Scheduling multistage workflows, and even multiple phases within a single application, can result in wasted resources if allocations are based solely on the requirements for the most resource-intensive phase.

Node-sharing of applications (a.k.a., multitenancy) can provide benefits but has cybersecurity challenges that need to be addressed. Scheduling mechanisms and policies that support allocation changes over applications' and workloads' lifetimes will need to be supported. At a meta-level, cost-benefit analyses to assess the feasibility of allowing a site's computations to take advantage of unused resources at another site as a form of "surge capacity" will need to be assessed.

3.3.6 Agility of Computing Capacity

Currently, ASC computing resources consistently run at or near full capacity. In the absence of increased resources, improving overall throughput will require more intelligent use of resources. Development of analyses and operational methodologies to enable this is complicated by anticipated changes in ASC compute hardware, applications, and workflow modes. The ability of a center to support dynamic resource demands is termed "agile capacity" or "resource elasticity." The model employed by cloud computing service providers is highly elastic, and the ASC program continues to explore and incorporate lessons learned from the cloud computing community to improve ASC's ability to support mission needs. In particular, fine-grained node resource management (node-sharing by applications) and containers for increased portability are being evaluated. To meet current and future mission demands, the program will need to maximize the flexibility, availability, and elasticity of the computing resources and minimize the challenges of portability of workflows across a wide variety of heterogeneous resources.

3.3.7 AI for Operations

AI/ML approaches, utilizing system monitoring and application performance data, are being investigated to provide support for resource management and diagnostics, enabling "smart facilities." These have the potential to reduce the human cost of maintaining a machine and reduce time spent in degraded conditions or downtime. Targets for operational analytics include detection of application resource under-utilization and performance-impacting contention and resultant mitigation, detection of abnormal or degraded application performance, and attribution of performance variation to code changes or system conditions, among others. In support of this work, ASC must motivate computer vendors to expose instrumentation data. Analyses can be informed by co-design insights relating to performance and architectural features.

3.3.8 Facilities Constraints

Facilities enable optimal computing services when planned in concert with systems. Facilities are properly viewed as an extension of the computing platform and, together with the hardware and software components, make up the ASC integrated computing environment. The NNSA computing facilities and HPC systems are interdependent, and the ASC program must make investments in both facilities and computing systems to fully serve mission needs. These investments often have to meet power, cooling, size/floor load, and public utility constraints; however, by gathering electrical, mechanical, and structural requirements early in the planning process NNSA laboratories can minimize constraints that would be otherwise placed on system procurements.

Not all system requirements can be precisely known up front; thus, flexibility, scalability, capability, reliability, and sustainability must be designed into every solution. High-level details on power, cooling, and structural needs can be determined years in advance by regularly engaging with system vendors about their product roadmaps, while detailed system facility requirements are often not available until as late as one year prior to delivery. Facility investments therefore require both near-term and long-term strategic plans.

With the slowing of Moore's Law, the density of ASC compute platforms (both AT and CT systems) is growing rapidly. Increasing power density in HPC systems is an industry trend that continues to drive demand for higher rack weight/floor load and complex liquid cooling solutions in the near term. Most new HPC systems feature unique node packaging, internal cooling mechanisms, and rack layouts to meet performance requirements; facility-level cooling solutions must be tailored to match system specifications.

Notwithstanding increased power density, to meet future, longer-term, increasing mission computing requirements, the NNSA laboratories will need to increase the number or size of systems deployed at any given time. As this set of requirements grows, other potential constraints include energy and water conservation requirements, as well as public utility infrastructure capacity and load limitations, all of which require long-term planning to coordinate with local utility providers. These constraints will possibly impact the design and architecture of systems that the ASC program must accommodate in existing and future HPC facilities.

The cost of facility upgrades, including cooling distribution units, plumbing, pumps, electrical upgrades, and other infrastructure, must be budgeted. Lead time to design the solutions, procure the equipment, and perform the site construction must be factored into project timelines. Dependencies between systems and facilities have become increasingly intertwined. Successfully delivering HPC capabilities at the scales required by the ASC program therefore depends vitally on close collaboration between facilities, codes, and systems in all phases of project planning and execution.

Strategy: ASC has a balanced investment approach that ensures underlying facility and networking infrastructure fully supports and integrates its major hardware and software investments.

3.4 Cross-Site Computing

In order to develop a more tri-lab approach to managing ASC compute resources to effectively meet user demands, the program must align infrastructure across the tri-labs and reduce policy limitations to ensure that the remote HPC user experience is as close as possible to the local user experience. There are many challenges in a multisite HPC complex for which each site handles its own identity management domains, often differing at a given site between the institution and the HPC center.

Specific difficulties particular to cross-site computing include the following:

- Reducing the number of gateway interconnects required to get to a desired destination;
- Marshalling network zone security models in like fashion and allowing for cross-site credential trust, thereby reducing the need for users to keep multiple tokens from multiple institutions and empowering a better security understanding of where codes can run in remote HPC environments;
- Increasing wide area network (WAN) bandwidth, allowing remote users to avoid a significant penalty for using tools like graphical profilers and debuggers;
- Extending a simplified identity management infrastructure capable of addressing the needs of an increasing number of new web applications while maintaining support for legacy software;
- Offering simplicity by coalescing on a similar suite of tools for common activities like data transfer between sites, inclusive of a common lexicon of storage access points;
- Permitting session-persistence for remote environments that have a desktop “look and feel”;
- Mitigating the disparity in short- and long-term data storage policies across the complex; and
- Enabling tri-lab continuous integration, permitting users to launch remote HPC jobs from their local repositories.

Complicating these cross-site challenges is the reality that there exists a narrower solution space available to solve problems for classified computing, given that similar infrastructure may not be viable in both classified and unclassified environments due to security reasons. Novel solutions that are often presented as options in the open computing, or unclassified, environment are often not an option in the classified environment. Running two solution sets on the same problem for both classified and unclassified computing is particularly taxing on available tri-lab resources.

With increased capabilities for cross-site access and the deployment of multisite collaborative tools, it becomes crucial to ensure need-to-know access is enforced not only for HPC computing environments, but also for each site’s corporate/business data. Assessing and limiting the scope of access can be a lengthy and resource consuming process. Achieving progress on these fronts will require ongoing complex-wide communication across the NNSA enterprise.

3.5 Cloud Computing

Over the last decade, the market for HPC scientific simulation has been dwarfed by dynamic, new markets that include cloud computing and AI. These high-margin, high-volume markets demand solutions that are orthogonal to simulation’s traditional needs and have outcompeted them for vendor attention. The NNSA labs constantly evaluate the question “Can ASC mod/sim workloads run in the cloud through independent service providers?” The answer is still a resolute “no” because of the high premium costs primarily due to the 80–98% system utilization and diverse data management and storage services demanded by the ASC codes and their users. However, ASC must continue to track the evolution of HPC cloud technologies to be able to respond in an informed fashion. While cloud providers have evolved their HPC offerings, barriers remain. These barriers include the lack of cloud computing facilities being operated at the required security level, insufficient interconnect support for the largest HPC platforms, and underperforming hardware libraries, in addition to cloud MPI offerings not being as high performing as MPI libraries on HPC systems. However, portions of the ASC workload can be run on cloud-like environments and ASC is expected to integrate many cloud-like features into its flagship HPC installations in the coming years. This will leave the ASC program well positioned for when its most intensive workloads could be transitioned to a future cloud environment.

Before that shift occurs, cloud technologies such as containers and micro services environments can enhance scientific discovery workflows on existing and next-generation systems. Recently, private cloud or on-premises cloud services have been deployed and are enjoying early successes. The next challenge becomes how best to integrate cloud technologies alongside large HPC systems and within HPC systems. Currently, HPC network technologies do not support the isolation required to run cloud services in a performant way and the ASC storage model (globally mounted file systems) makes secure isolation difficult within nodes. Integrating cloud-native virtualization and resource management tools with HPC resource management tools will be crucial for ensuring that software developed in industry can be reused.

Regardless of whether the ASC program will ultimately be able to use the cloud for production workloads, the cloud will increasingly become the way that the rest of the world leverages parallel computing resources. Cloud providers continue to improve their HPC offerings, and the widespread industry adoption of ML/AI workloads puts the focus of cloud vendors much closer to that of traditional HPC users. For ASC software to gain adoption and influence the industry, and in order for DOE to build communities around its open-source projects, lab code developers must ensure that the ASC libraries, applications, and workflows can run well on both HPC and cloud resources. ASC can facilitate this by ensuring that the software stack, particularly key components like resource schedulers (e.g., Flux) and software integration tools (e.g., Spack), can work in both the cloud and on HPC resources.

3.6 Cybersecurity

Increasing cybersecurity requirements primarily written to target corporate resources can affect HPC ability to respond to R&D needs across classified and unclassified networks. Currently, each tri-lab site is positioned differently in their ability to insulate from increasing cybersecurity requirements and such differences can impede progress in multisite collaborations. ASC will continue to push for integrating the latest cybersecurity technologies and processes into our HPC ecosystem. This includes the hallmarks of a strong cybersecurity posture, such as strong multifactor authentication, identity management, multilayer network protections, and stringent monitoring. The focus will be enhancement of sharing of the most capable systems across multiple security domains while simultaneously supporting both the ASC user community at the tri-labs and the ASC partners from academia and the private sector.

3.7 Software Quality Assurance

Software developed under the ASC program ranges from research software, designed to rapidly explore new approaches to simulation, to large multiphysics production codes that are key to informing decisions on the reliability and safety of the nuclear stockpile. The NNSA labs employ a risk-based, graded approach that ensures ASC research software is not overburdened with undue process, while the weapons production codes are held to a higher standard of verification and validation and development practices supporting that standard as they are required to be in compliance with DOE Order 414.1D and NAP-24A. Robust software engineering is particularly important to the long-term health of ASC codes and rigorous software quality assurance (SQA) practices are in place at the national laboratories to monitor, maintain, and continuously improve those assets. SQA tools are needed for static analysis of source code (for defects), measuring test coverage, and developing software quality metrics, along with automated and reproducible software development practices that ensure software quality is maintained, improved, and continuous as developers transition their involvement in a given project over its lifetime. Furthermore, ASC laboratory documents and implementation plans such as Software Quality Assurance Plans, Software Test Plans, Software Configuration Management Plans, and Disaster Recovery Plans are reviewed regularly and updated as needed.

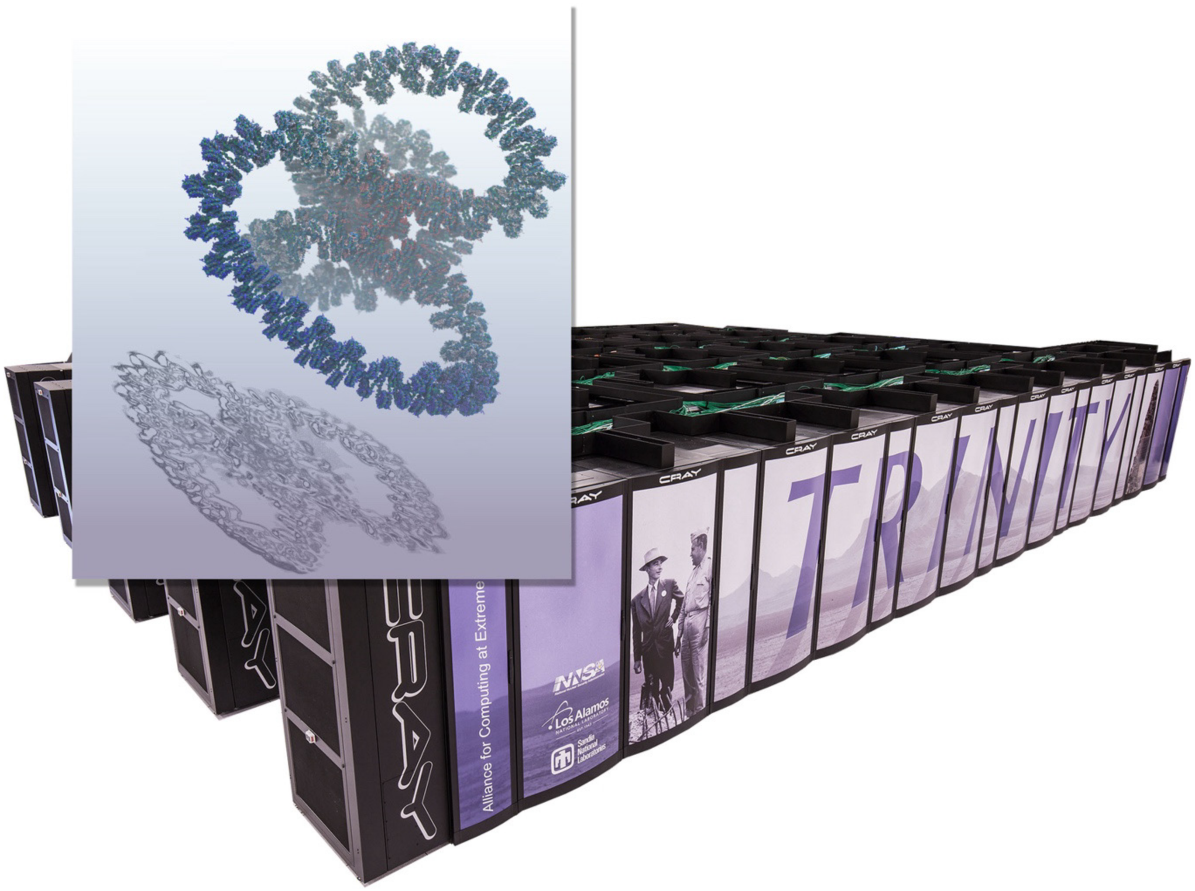

Figure 9. The first billion atom simulation in biology was performed on Trinity to help biomedical researchers determine the effect of DNA compaction on gene expression. Shown in atomistic detail is the GATA4 gene locus, important for stomach and colon cancer. The DNA of the gene is wrapped around 427 molecular spools (histones). Understanding the response of DNA compaction to viral infections, such as COVID-19, may aid in characterizing subtle effects, missed by current diagnostics. (LANL)

4.0 CO-DESIGN APPROACH

The ASC program has embraced a co-design process to evolve and transition codes into the exascale era and beyond. The co-design process enables an integrated approach to computational challenges through tight collaborations between hardware, software, and application specialists. Broadly speaking, co-design is a process of end-to-end optimization—from a code's fundamental physics and numerical methods, its algorithms and data structures, interactions with other processes in the overall workflow, to how that all maps onto the hardware and system software that the code will ultimately run on. This process of co-design is implemented through formation and management of collaborative, multidisciplinary teams that include DOE computational and computer scientists, and hardware and software vendor representatives. These teams conduct their investigations through a tightly coupled cycle of application, algorithm and system software development, performance modeling, hardware simulation, and hardware design. Through the co-design process, optimal design tradeoffs are identified for hardware, run-time environment, parallel file systems, physics algorithms and packages, and programming models. The relationships within the team are fluid and agile, reorganizing and realigning themselves as necessary throughout the lifetime of the effort.

The NNSA tri-lab co-design portfolio was instituted to investigate a diverse set of technical areas such as architecture-aware algorithms, performance portability, programming models, system software, hardware architectures, resiliency, workflows, and power management. The projects are driven by the multidisciplinary nature of co-design and draw participation and influence from all of the ASC program elements.

High-impact focal points for these investigations are the following:

- Algorithmic exploration,
- Proxy application development,
- HPC architectural simulation,
- Advanced architecture testbeds and prototypes, and
- Future hardware trends.

Characteristics that impact performance should be understood as early as possible in the analysis and design of new computers. Furthermore, it is often the case that there are multiple ways to design and implement the algorithms used in an application, and the choice can have a dramatic impact on application performance. Hardware architectural choices, if properly exposed through the system software to the application developers, can likewise have a tremendous impact on performance.

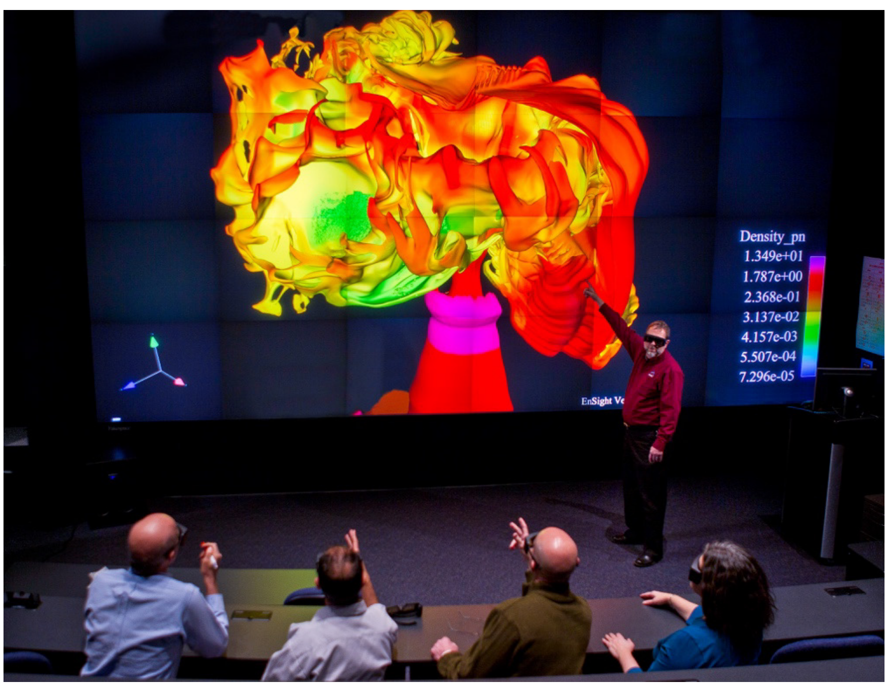

Figure 10. ASC scientists use the Strategic Computing Complex powerwall to demonstrate the 3D xRage simulation of a laser experiment to generate a laboratory-scaled version of an astrophysical jet. (LANL)

4.1 Algorithmic Exploration

For many years, the program was able to focus algorithmic development on bulk-synchronous MPI approaches with limited parallelism within the node. Processors on the node had limited core counts and message passing approaches were viable even within the node. Most of the algorithmic innovation was associated with "scale-out" of the algorithms with the goal being to approach near perfect scalability for very large node counts and MPI processes (approaching millions). This landscape has drastically changed for more recent ASC platform procurements. At the very least, the nodes have much higher core-count processors which require shared memory parallel algorithms to take advantage of the greater parallelism. To add even greater complexity, many of the new systems incorporate GPU-based accelerators that require work to be off-loaded from the CPUs and have much more complex memory subsystems, which create data movement challenges and incorporate much higher degrees of parallelism (10s to 100s of thousands of threads). Future systems will likely have even higher levels of heterogeneity as a diverse set of application/algorithm specific accelerators are made available. User expectation in improved physics and new features has not been reduced to accommodate the harnessing of the speed of these new platforms. This requires that the accuracy and mathematical flexibility of needed algorithms improves. Fortunately, the arrival of new platforms provides a unique opportunity to pursue these activities.

Key areas of algorithmic research that address the above challenge include the following:

- High-order and high-order/low-order (HOLO) hybrid solution methods
- Extreme levels of shared memory parallelism
- Mixed-precision and compression approaches
- Reproducibility
- Data-centric algorithms
- Multiscale physics coupling
- Implicit or semi-implicit methods
- Concurrent multiphysics operators

See Appendix B for more details on the algorithmic topics.

4.2 Proxy Application Development

Application performance is determined by a combination of many choices, such as the hardware platform, runtime environment, languages and compilers, algorithm choice and implementation, plus other considerations. In this complex environment, the use of mini-applications, skeleton applications, and kernels is an important approach for rapidly exploring the parameter space of all these choices. Furthermore, the use of proxy applications or small, self-contained software products enrich the interaction among application, library, system software and hardware architecture designers and developers by providing explicit functioning software and concrete performance results that lead to detailed, focused discussions of design trade-offs, algorithm choices, and runtime performance issues.

Application performance analysis has historically taken advantage of two important properties of some applications. Firstly, although an ASC application may have millions of lines of source code, performance is sometimes dominated by a relatively small subset. Secondly, for the remaining code these applications often contain many physics models that are mathematically distinct but have very similar performance characteristics. Mini applications take advantage of these two application properties by encapsulating only the most important computational operations and consolidating physics capabilities that have similar performance profiles. The large-scale application developers, who are tasked with designing and developing the proxy application, guide the decisions, resulting in a code that is a small fraction of the original application size yet still captures the primary performance behavior.

While the current approach to proxy application development has proven quite valuable, it does have its deficiencies. First, the engineered simplicity of a proxy application can lure a researcher into making modifications for purposes of making performance trade-off decisions that are wholly inappropriate for the application it represents. In addition, it is not uncommon for an algorithmic choice to be well suited to a specific architecture but not another. This may force a significant rethinking of a proxy application to adopt a new algorithmic approach, (e.g., higher-order methods with greater computational intensity and parallelism for GP-GPUs). Second, as applications become increasingly complex or coupled together in intricate workflows, proxy applications must be expanded to study how to best take advantage of advances in HPC hardware and software that exploits resource allocation and scheduling, execution asynchrony, extreme shared memory parallelism, and custom acceleration. This may be manifested in a larger number of representative computations, multiple motifs of execution, more complex data movement, or all of the above. Distilling these applications and workflows to comprehensively capture their interactions will require new approaches to proxy application development to balance the higher level of complexity with the desire for a compact code that facilitates rapid design cycles and improved communication among co-design stakeholders.

4.3 HPC Architectural Simulation

Architectural simulator capabilities are important tools to enable co-design to close the loop back to computer architects and hardware component designers. Without this capability, co-design threatens to retreat to "business as usual," in which new HPC systems are procured and ASC code teams, algorithm developers, and system software developers are then tasked with extracting the best performance possible from the HPC systems. The intent behind HPC architectural simulation is to obtain quantitative data to guide the technology development and design of all elements of the integrated HPC system.

HPC industry partners have a tradition of using simulators to analyze and model processors, interconnection networks, and other features of their proprietary designs. Some simulation capabilities are cycle-accurate and highly proprietary. To the extent that the ASC program can access and use these simulators, or provide proxy applications to drive these proprietary simulators, important quantitative data can be obtained to inform the co-design process. Processor models can be integrated with memory subsystem and network interface models to provide a node-level model. At a lower level of fidelity, it may be acceptable to give up on cycle-accurate processor models in order to enable simulations of entire HPC systems that integrate node-level models with models of the interconnect networks and to include analysis into holistic properties, such as energy consumption, thermal stresses, and system reliability. Again, proxy applications can be used to drive HPC architectural simulators to provide quantitative data to the co-design process.

4.4 Advanced Architectural Testbeds and Prototypes

Looking forward over the rest of this decade, it is apparent that revolutionary technologies will be required, particularly if one considers the end of Dennard scaling and the approaching end of Moore's Law. These factors have helped drive the adoption of many-core CPUs and GPUs as primary HPC technologies. Over the coming years further advances are expected such as tighter integration of CPU and GPU technologies, improved memory technologies, and an increasing number of algorithm-specific accelerators technologies, largely in response to the needs of AI/ML workloads. ASC will adopt revolutionary technology changes, up to and including rewriting our applications, if that is what is required to meet mission needs. The original ASCI program made sustained and substantial investments to transform the NNSA application codes from vector- to MPI-based computing to create the current ASC application portfolio. In much the same way, the program now must continue to undertake the challenge of transforming the NNSA application codes to incorporate advances spanning programming models, runtime systems, and new hardware technologies.

It is critical to have a diverse set of experimental architecture testbeds to guide technology investment decisions. As a community, access to and experience with these experimental architecture testbeds will allow ASC to build upon our leadership in HPC and co-design process, increase our agility to changes in hardware, develop advanced methodologies for using the new hardware, and inform future programming model changes. Perhaps more importantly, this experience provides a foundation for decision makers to determine future HPC platform investments required to meet mission deliverables.

4.5 Future HPC Hardware Trends

As mentioned in the previous sections, the HPC community has seen a drastic shift in hardware architectures that both create challenges for our mission work and opportunities to realize massive performance gains beyond the Moore's Law curve.

Examples include the following:

- Special-purpose accelerators for key algorithmic kernels
- Memory systems optimized for sparse data algorithms
- Photonics-based networks for disaggregated architectures
- AI/ML-focused architectures
- Increasing heterogeneity, both inter- and intra-node
- Computation embedded within the data path (storage, memory, network)
- Quantum
- Neuromorphic

These trends present unique opportunities and challenges for ASC and reinforces the critical importance of co-design with the vendor ecosystem. Co-design of new algorithms with special purpose accelerators and new memory technologies could significantly improve performance across a number of our applications. Evaluating and assessing the suitability of AI/ML-focused architectures for both data-intensive and compute-intensive workloads will be important as the number of startup and established companies operating in this space increases. Increasing heterogeneity presents new challenges and opportunities for platform designs optimized for a diverse ASC simulation workload. Computational resources available throughout the data path could impact all layers of the software stack and require changes to programming systems to effectively use these resources. ASC is investing in Beyond Moore's Law technologies, such as quantum and neuromorphic architectures, to understand and develop algorithms and applications that have direct applicability to mission problems to prepare as these technologies mature in the near future. Maintaining ASC leadership in HPC will require active engagement in each of these areas.

5.0 External Partnerships

The execution of this ASC Computing Strategy requires frequent and close communication, alignment, and partnerships with many organizations and institutions. These include U.S. computer companies, universities, and the DOE/SC ASCR program.

5.1 Active Engagement with the U.S. Computer Industry

The ASC program has been extremely successful in leveraging commodity computer technology to create systems without incurring the tremendous costs of fully custom designs. However, commodity computing, especially in the consumer markets, is rapidly moving away from the types of designs the program had successfully leveraged in the past. The scientific HPC community, and NNSA in particular, must maintain close collaborations with computer vendors to ensure that features essential to its needs are continuously identified, integrated, and supported in future products.

While close collaboration with computer vendors will mitigate the risk of a catastrophic loss of HPC capability, ASC simulation codes are unable to exploit enhanced hardware features without employing new programming approaches. This represents a significant challenge for ASC application developers, as mission drivers for new capabilities in simulation codes already stretch programmatic resources. The ASC program must make hard decisions on programming approaches that meet its current needs and for the future, while managing the timing of any technological disruptions.

The jointly funded PathForward projects by ASC and ASCR were a recent example of tremendous R&D investment (~$400M) in the U.S. computer industry. These projects span the majority of the influential companies and they have focused on accelerating the availability of high-impact technologies in the three forthcoming DOE exascale systems.

5.2 Training and Education of a New Generation of Scientists and Engineers

The ASC Predictive Science Academic Alliance Program (PSAAP) has been a critical pipeline for the NNSA laboratories' next generation of scientists and engineers. The PSAAP II Centers, which completed their six-year activities in 2021, explored predictive science for multiphysics applications that use future generations of advanced computer architectures. The new PSAAP III[5] Centers have the additional focus on AI/ML technologies and their potential impact on mission capabilities. The ASC academic alliance program, spanning more than a dozen major universities over more than 20 years, exposes young graduate students and post-doctoral candidates to the power of interdisciplinary, large-scale scientific simulation and computation as they work on the Centers' overarching, unclassified application problems and on research projects during their requisite NNSA lab visits. ASC not only funds research in disciplines of high interest to the weapons program, but also makes enormous unclassified computing resources available to the Centers' research staff, post-doctoral candidates, students, and their mentors involved in these activities.

5.3 Strategic Partnership with DOE/SC ASCR

As described in Section 5.1, the ASC program has been partnering with ASCR in the DOE Exascale Computing Project (ECP)[6] since 2016 to leverage both programs' core talents and fund a set of collaborative hardware and software R&D projects with industry and the DOE laboratories. The funded ECP focus areas are Application Development, Software Technologies, and Hardware and Integration. Although ECP is scheduled to conclude at the end of FY2023, the tremendous benefits and impacts of ECP successes will be long lasting and ASC will endeavor to continue supporting key exascale technologies for its mission needs into the future.

Another equally successful collaboration with ASCR is in procurements of ASC AT systems with ASCR leadership-class systems. A major benefit of this approach has been the opportunity to leverage the NRE investments of critical technologies to benefit both programs and also the tremendous technical expertise of the involved staff at the three DOE SC and three NNSA laboratories. The Collaboration of Oak Ridge, Argonne, and Livermore (CORAL) was a joint procurement activity among these three DOE laboratories launched in 2014 to build state-of-the-art high-performance computing technologies that are essential for supporting DOE national security and science missions. The collaboration has been procuring leadership-class computing systems for two national labs (Summit system at ORNL and Sierra system at LLNL via

[5] Predictive Science Academic Alliance Program, https://psaap.llnl.gov/
[6] DOE Exascale Computing Project, https://www.exascaleproject.org/

the CORAL-1 procurement) and will deploy the next two exascale systems for DOE via the CORAL-2 procurement (Frontier system at ORNL at end of 2021 and El Capitan system at LLNL in 2023).

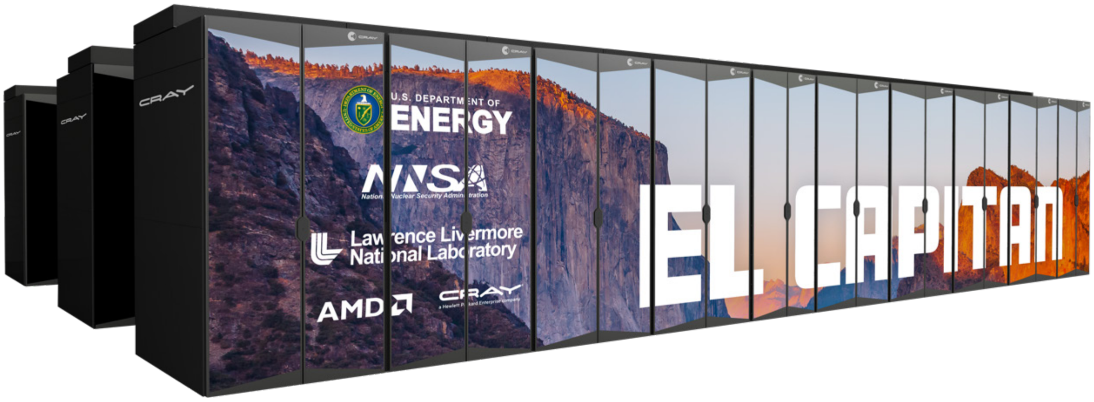

Figure 11. El Capitan will be sited at Lawrence Livermore in 2023 and will be NNSA's first exascale system servicing Stockpile Stewardship and Management Programs. (LLNL)

5.4 Collaboration with NNSA Production Sites

Historically, the NNSA production sites at Kansas City, Pantex, and Y-12 have made very little use of HPC capabilities. This has changed over recent years. To modernize and improve their mission work, these sites have embraced the use of HPC for modeling and simulation. To support these activities, the ASC program has made HPC resources available at the tri-labs, helped the production sites procure significant capacity systems of their own, and supported the production sites' use of a range of ASC engineering codes, data analysis, and visualization tools. In addition, the new ASC Production Simulation Initiative (PSI) will engage with the NNSA production agencies to develop and deploy more modern simulation tools to help accelerate the design, manufacturing, and qualification of nuclear and non-nuclear components.

Figure 12. Sierra is an IBM supercomputer sited at Lawrence Livermore. Containing 17,280 NVIDIA GPUs and 8,640 IBM P9 CPUs, Sierra is capable of 125 petaFLOPS. (LLNL)

6.0 SUMMARY

Successful stockpile stewardship is fundamentally dependent upon simulations and analyses of extraordinarily complex devices and physical and production processes in order to advance our scientific understanding and inform critical weapon system decisions. The ASC program is responsible for ensuring that all necessary scientific expertise, high-performance computing platforms, and software infrastructure are available to meet current and future stockpile stewardship needs and challenges. To accomplish this goal, the program must provide a stable, production-level computer service for the tri-lab user community, modernize, and maintain its tri-lab computing network, and maintain a balanced portfolio of investments in advanced research, development, and deployment of leadership-class HPC platforms. The overriding objective is to maximize user and developer productivity, while simultaneously enabling scientific code improvements and enhanced confidence in simulations of device and system performance beyond the realm of the current stockpile.

As the NNSA mission drivers evolve over the next decade and beyond, the ASC program must remain responsive and agile to ensure that current and future generations of ASC simulation capabilities and resources will continue to underwrite our Nation's nuclear deterrent. The work scope outlined in this ASC Computing Strategy for the upcoming decade is very ambitious and daunting. However, with the very capable and experienced staff at the NNSA laboratories and HQ working diligently on the program priorities, while navigating carefully through the usual and continual budget uncertainties, ASC will continue to deliver on its commitments to the nuclear security missions as it has in the past two decades.

7.0 ACRONYMS

AAP	advanced architecture prototype
AI	Artificial Intelligence
API	Application Programming Interface
ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
ASCR	Advanced Scientific Computing Research (DOE Office of Science)
AT	Advanced Technology
ATCC	Advanced Technology Computing Campaign
ATDM	Advanced Technology Development and Mitigation
CORAL	Collaboration of Oak Ridge, Argonne, and Livermore
COTS	Commercial off the shelf
CPU	Central Processing Unit
CSSE	Computational Systems and Software Environment
CT	Commodity Technology
DAM	Defense Applications and Modeling
DOE	Department of Energy
DOE SC	Department of Energy Office of Science
FLOPS	Floating-point operations per second
FOUS	Facility Operations and User Support
GP-GPU	General Purpose Graphics Processing Unit
GPU	Graphics Processing Unit
HPC	High Performance Computing
IC	Integrated Codes
I/O	Input/Output
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LLNL	Lawrence Livermore National Laboratory
ML	Machine Learning
MPI	Message Passing Interface
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NRE	Non-Recurring Engineering
NSE	Nuclear Security Enterprise
PEM	Physics and Engineering Models
PSAAP	Predictive Science Academic Alliance Program
RAS	Reliability, Availability, and Serviceability
RCE	Remote Computing Enablement
R&D	Research & Development
RDT&E	Research, Development, Test, & Evaluation
SCDS	Stewardship Capability Delivery Schedule
SFI	Significant Finding Investigators
SNL	Sandia National Laboratories
SSP	Stockpile Stewardship Program
TOSS	Tri-lab Operating System Stack
TLCC	Tri-lab Linux Capacity Cluster
UQ	Uncertainty Quantification

V&V	Verification & Validation
WAN	Wide Area Network

Appendix A. Updated Lessons Learned

Lesson 1: A stable, modestly changing computing environment significantly increases the productivity of code developers, designers, and analysts.

Our highest-end simulation needs have often required new technical features only provided by the most advanced vendor-provided solutions, which are often developed by partnering between the computer scientists and users at the NNSA laboratories and the computer vendors. To implement these solutions, ASC recognized the balance between short-term user productivity disruptions and the long-term productivity increases that result from more capable codes running on more powerful systems. To implement new technologies with minimum disruption to users, ASC historically focused on an explicit message passing programming model that allowed our applications to evolve independently from the characteristics of a particular generation of high-end computers. However, the most recent generation of platforms requires much greater parallelism at the node level due to many-core CPUs and GP-GPU accelerators. This has significantly increased the complexity of ASC applications and the program has reacted by investing in on-node programming model abstractions (e.g., OpenMP, Kokkos, RAJA, and FleCSI) to reduce the cost of implementation and porting of these applications across disparate architectures. The result of these investments has allowed us to appreciably advance our science throughput on pre-exascale platforms (e.g., Trinity, Sierra, and Crossroads) and then El Capitan, NNSA's first exascale system, while future-proofing our applications against inevitable disruptions coming in post-exascale architectures.

Lesson 2: The weapons application workload benefits from a mix of computer systems available to match cost performance to problem needs; the program needs to leverage and influence the wave of new, more specialized architectures showing up in the marketplace.

In recent years, the marketplace has evolved to a state in which high-end, commodity-based systems met the needs of many ASC capacity problems at substantially less cost. In the first 10 years of the ASCI/ASC program, the focus was on commercial-off-the-shelf (COTS) processors, but the inter-processor communications fabric and software were customized because no alternatives existed. Now COTS providers have expanded to include the communications fabric and systems software including the Linux operating system, as well as a variety of open-source software for debugging and performance tools, system monitoring and control, job scheduling, and file systems. ASC investments contributed to this evolution of COTS technology. This has made possible the acquisition of CT systems, which can handle a substantial fraction of the workload at a significantly reduced cost. However, accelerator-based architectures are becoming more common in DOE's leadership-class platforms. As applications port to and optimize for these architectures, we will see larger fractions of the CT platforms transition to these architectures as well to maximize throughput for our mission.

Lesson 3: Investing in and influencing the development of market-based technologies for ASC supercomputers has proven to be a successful strategy for balancing system costs and progress in scientific computations.

From its inception, the ASC program decided to work with the computing industry to leverage its business models to build supercomputers for scientific applications. A benefit from this partnership is that commodity-based solutions provide an evolutionary path for applications, ensuring that code investments can cost-efficiently carry over to future generations. Another benefit is that although market-based supercomputing platforms are expensive they are still more affordable than custom-built architectures, and vendors are able to build, test, and deliver them in a relatively short period of time. Furthermore, given the low HPC sales volume from the scientific community, these business-based solutions leverage a much larger market and provide a stable basis for producing ongoing generations of supercomputers.

With the expansion of more specialized technologies such as GP-GPU accelerators, the program made successful investments in pre-production R&D through such programs as FastForward, DesignForward, and PathForward. The joint R&D efforts with the Office of Science and the U.S. computer vendors have pulled in and optimized new technologies for the benefit of DOE's supercomputer procurements and played a major role in the successful procurements for the forthcoming three DOE exascale-class systems.

Lesson 4: Bringing leading-edge systems to a production level is both a time- and resource-consuming process that requires a strong partnership between the laboratories and vendors.

To meet requirements and to ensure that needed petascale computers would exist in the future, the ASC program procured systems that accelerated the business plans of its vendor partners. This resulted in both an invigorated HPC industry and a series of "serial-number-1" systems that were acquired two to four years in advance of market offerings from a cross section of the industry. Over the past two decades, the ASC approach fostered competition and brought systems to market that would not have existed otherwise. Such systems provided a means to explore problem spaces previously not possible, but the application of such systems to production work introduced unforeseen problems in hardware and software reliability and system features. The tri-lab system integration teams worked closely with the ASC applications groups and system vendors to ensure that when the applications uncovered bugs in hardware or software, the issues were dealt with quickly and the solutions were implemented in a practical manner. Often understated, this was an essential ingredient for success in an advanced development environment. A new construct, the Center of Excellence (COE), has become a best practice of ASC's supercomputer procurements. This tri-lab and vendor collaborative effort is a multiyear focal point for preparing our applications prior to final system delivery to run effectively from the start and to dramatically improve productivity of the system early in its life.

Lesson 5: Innovative architectural approaches provide significant future capabilities even though use of the advanced technology/architecture may be confined, in the early stages, to a subset of the important physics simulations for which ASC is responsible.

The AT systems, though higher-risk endeavors, have provided significant returns to ASC in terms of illuminating innovative architectural features. For example, the IBM BlueGene/L (BG/L) and Intel Xeon Phi architectures drove on-node parallelism for performance and the IBM Power/NVIDIA GPU architecture has pushed this requirement even further for ASC multiphysics applications. ASC's use and investigation of the architectural innovations in AT systems with industrial partners ensures that the NNSA user community understands how to use the next-generation computers to solve stockpile stewardship problems and demonstrates that future HPC systems are suitable for these problems. Strategic investments in AT systems, which may begin as only applicable to a subset of stockpile issues but later expand in scope to become powerful general-purpose production engines, has been realized over multiple generations of AT procurements and is an effective and productive model for ASC success. A number of benefits are realized by simultaneously fielding a new AT system while executing day-to-day production computing responsibilities, including the following:

- Learning how to write algorithms and code for future advanced architectures,
- Gaining insight into how ASC applications can be made to work efficiently on new and possibly revolutionary architectures,
- Porting the large weapons application codes in advance of the production phase, and
- Training users and system personnel in the use of the AT system.

Lesson 6: Computing-at-a-distance continues to be an integral capability for NNSA, critical for resolving the most complex weapons system problems across the complex.

Advanced, powerful ASC computers, whether deployed at LANL, LLNL, or SNL, are tri-lab resources with major cycle allocations provided for each laboratory. This is a successful usage governance model that enables scientists at each of the laboratories to compute effectively from their home laboratory on the most powerful systems available within the NNSA complex with new systems arriving on a regular cadence. This model's success is enabled by robust classified networking resources and by data assessment tools that can be run in a variety of ways to fit particular programmatic requirements and platforms. There still exist bottlenecks in the tri-lab ASC computing infrastructure and the program continues to invest in improvements to throughput and usability of this infrastructure to reduce those barriers. In the last few years, the ASC Remote Computing Enablement (RCE) project has made considerable contributions toward the end goal of enabling remote HPC user experience as close as possible to the local user experience to maximize productive utilization of computing resources across the NNSA HPC simulation complex. ASC will continue to use the proven, successful tri-lab usage governance model to allocate future AT system cycles.

Lesson 7: U.S. competitiveness and leadership in high-end computing, enabled by government investment and industry commitment, are necessary for progress in science-based stockpile stewardship.

The NNSA Stockpile Stewardship Program could not have succeeded without a sufficiently healthy high-end computing industry. NNSA's unique nuclear security mission has required computing performance beyond the capabilities normally available in the commercial marketplace. Designing and delivering state-of-the-art supercomputers that meet stockpile stewardship needs within practical cost and schedule constraints have been possible through the efficiency and innovation of a healthy and competitive industry. Government support, through the ASC program and a few other U.S. agencies, helped to engender this competitiveness and leadership through both competitive procurements and a commitment to share some NRE costs for development of capabilities needed to scale the systems upward to meet our mission computing needs. The common denominator for continued U.S. competitiveness remains to be successful, long-term partnerships between government agencies and industry.

Lesson 8: ASC's integrated program approach and strong mission needs have been key to its success, and ATDM provided a strong cross-cutting element that strengthened our ability to deliver on exascale.

From its inception, ASC has been an integrated program (see Appendix C for program structure), with an approach that has been critical to unifying the goals of the program around the ultimate mission goal of science-based stockpile stewardship. A natural tension exists between users' insatiable need for more compute power versus the need to minimize computing disruptions to code developers and users. The ASC program has balanced both needs by continuing to lead the world in deploying usable, production-ready platforms to deliver increasingly predictive science in support of the U.S. national security mission. The deep partnerships with computer vendors, discussed at length throughout this document, are the unwritten additional program element of ASC, which together provide a blueprint for success by incentivizing a unified and cohesive approach to the mission.

The addition of the ATDM program element in 2014 provided a way to address the disruptions expected from exascale computing and a clean-sheet approach to next-generation code development unencumbered by production demands. This offered an additional integrating element that will have lasting effects after ATDM ends in 2023, when its scope will be transitioned primarily to the Integrated Codes (IC) and CSSE program elements. ATDM taught the ASC community many of the same lessons that co-design did—that success is more likely when different parts of the ecosystem understand each other's needs and capabilities. Cross-cutting efforts within each laboratory and across the tri-labs will remain a critical element of the ASC culture.

APPENDIX B. ALGORITHMIC RESEARCH CHALLENGES

- *High-order and High-order/Low-order (HOLO) hybrid solution methods*: High-order methods have shown promise in increasing both algorithmic robustness as well as reducing memory traffic per floating-point operation. HOLO hybrid solution methods have shown significant speedups in application runtime while preserving solution quality by reducing the need to use a high-order method for each step in the simulation. Future work in algorithmic advances that enable asynchronous execution of the high-order and low-order operators could result in even greater improvements on future systems.
- *Extreme levels of shared memory parallelism*: Co-design of new algorithms, programming models, system software, and underlying hardware will be required to address the extreme levels of shared memory parallelism available on modern CPU and GPU architectures. The choice of algorithm, data structures, and hardware level optimizations can have significant performance ramifications for shared memory parallelism.
- *Mixed-precision and compression approaches*: Mixed-precision approaches will be required to realize the next step function in performance for many applications on a variety of GPU and ML-optimized architectures where performance of fused, multi-precision operations can be 10x or higher than traditional double-precision operations. Likewise, lossy compression of floating point is emerging as a way to reduce memory capacity and bandwidth, with similar algorithmic tradeoffs to low precision. Co-design of numerical algorithms alongside a variety of new hardware accelerators will be increasingly important as the diversity of these hardware accelerators continues to grow.
- *Reproducibility*: ASC codes will continue to evolve to leverage higher degrees of asynchrony spanning algorithms, programming systems, runtimes, and the underlying HPC hardware. Mitigating the impact to scientific reproducibility through numerical robustness of algorithms, new software technologies, and analysis tools to enable data inter-comparison will be required.
- *Data-centric algorithms*: A large number of ASC applications are data-intensive such as in-situ analysis, interactive data-analysis, V&V, and more recently in the area of AI/ML. Coupled with advances in programming systems, system software, and hardware advances (such as compute resources along the entire data path), the need to significantly reduce the time to insight for these workloads opens new opportunities for co-design with data-centric algorithms.
- *Multiscale physics coupling*: Multiscale physics coupling offers the opportunity to capture finer-scale physics resulting in a higher-fidelity simulation (e.g., granular length scales in materials). These applications can significantly benefit from asynchronous execution, algorithm specific hardware acceleration, and heterogeneous system architecture. Co-design of new approaches to multi-scale physics coupling with asynchronous programming systems will enable these higher-fidelity approaches on future HPC systems.
- *Implicit or semi-implicit methods*: Implicit or semi-implicit methods are appropriate for longer timescale problems (e.g., modeling additive manufacturing processes) where explicit methods suffer from excessive time steps that make simulation intractable. Research in new discretization techniques and matrix solvers are required to make these techniques robust to widely variable material stiffness, and even shock discontinuities.
- *Concurrent multiphysics operators*: Most multiphysics codes are designed so that physics packages (operators) are run in a prescribed sequence within the application's time step loop on a common set of resources. Allowing operators to run concurrently on separate partitions of nodes could allow an extra degree of parallelism, utilize nodes that best map to the algorithm, (e.g., CPUs vs. GPUs, high-bandwidth memory [HBM] vs. dynamic random-access memory [DRAM]), and scale appropriately to the needs of the package.

Appendix C. ASC National Work Breakdown Structure (as of 2022)