

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

Quantum Computing Strategy 2026

S Hammond, J DuBois, F Graziani, Y Rosen, K Wendt, A Baczewski, M Descour, T Metodi, M Lang, D Richards, C Coffrin, S Eidenbenz, R Lowrie, A Sornborger, S Rudin, K Beck

October 2025



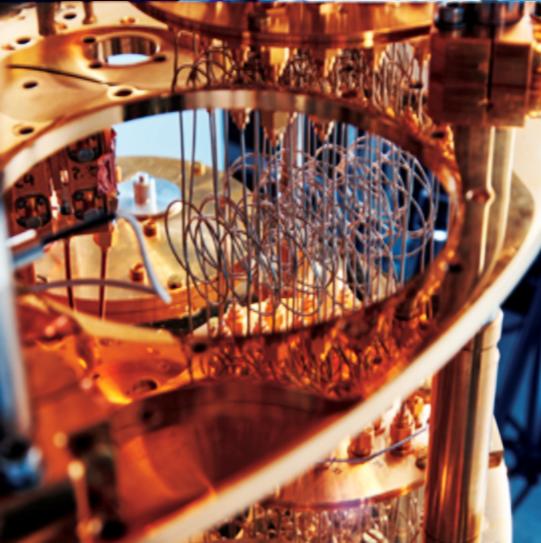
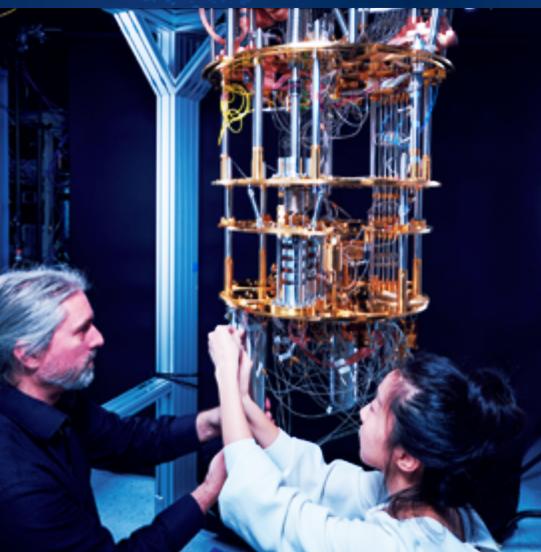
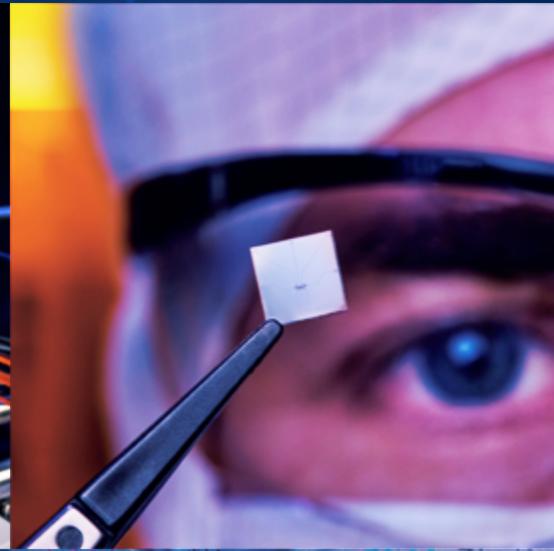
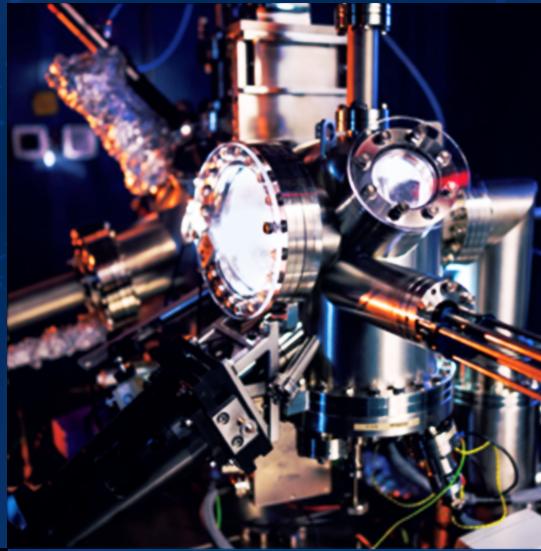
Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

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

ADVANCED SIMULATION AND COMPUTING

Quantum Computing Strategy 2026



LLNL-TR-2012045

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

For more information, contact the ASC Computing Program Team at dl-asc-computing@nnsa.doe.gov

On the cover

Top Center: A scanning tunneling microscope used to build electronic devices by placing individual phosphorus atoms into silicon. Devices that control the spin of individual atomic nuclei and use a single electron to couple those nuclei, are a candidate platform for quantum computing.

Top Right: Staff scientist Alex Abelson holds up a three-qubit superconducting device fabricated at LLNL's Center for Micro- and Nanotechnology.

Middle Left: Experimental physicists Sean O'Kelley and Yujin Cho mount new samples on one of the dilution refrigerators in LLNL's Quantum Design and Integration Testbed (QuDIT).

Middle Right: Theoretical scientists John Kallaugh and Ojas Parekh from Sandia National Laboratories find tasks in which quantum computers outperform classical computers.

Bottom Left: Close up of dilution refrigerator in LLNL's QuDIT used to test superconducting qubits and other supporting infrastructure for quantum computers at ultra-low temperatures.

Bottom Center: Electrical engineer Ray Haltli optimizes parameters before placing gold wire bonds on an ion trap at Sandia National Laboratories.

Advanced Simulation and Computing

Quantum Computing Strategy 2026

NNSA

Simon Hammond
Michael K. Lang*
David Richards**

LANL

Carleton Coffrin
Stephan Eidenbenz
Robert Lowrie
Andrew Sornborger
Sven Rudin

LLNL

Kristin Beck
Jonathan Dubois
Frank Graziani
Yaniv Rosen
Kyle Wendt

SNL

Andrew Baczewski
Michael Descour
Tzvetan Metodi

*on detail from LANL
** on detail from LLNL

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

Additional Contributors

NNSA

Sara Campbell
David Etim
Cheri Hautala-Bateman
K. Michael Lang
Thuc Hoang
Anthony Lewis
James Peltz
Stephen Rinehart
Robert Spencer

LANL

John Patchett

SNL
Robert Hoekstra

Table of Contents

Foreword	iv
Executive Summary	v
1 The Role of Quantum Computing in Stockpile Stewardship Science	1
1.1 Maturation Roadmap for a Production-Ready ASC Quantum Computer	2
2 Develop Capabilities in Mission-Relevant Quantum Applications	3
2.1 Simulating Materials in Extreme Conditions	3
2.2 Nuclear Dynamics	4
2.3 Partial Differential Equations	4
2.4 Uncertainty Quantification and the Quantum AI Nexus	5
3 Conduct R&D in Algorithms, Software, and Hardware.....	6
3.1 Algorithm Development and Optimization	6
3.2 Software Research Activities	7
3.3 Hardware Development Activities	8
4 Engage Vendors and Partners Through Testbeds and Co-design	8
5 Build Knowledge, Experience, and Workforce	10
6 Investment Priorities, Facilities, and Roadmap Phasing	11
7 Conclusions	12

This page is intentionally left blank.

Foreword

At the very heart of the NNSA's mission lies an on-going challenge to better understand, model, and predict the behavior of the physical world. For more than twenty-five years, the Advanced Simulation and Computing (ASC) Program has worked relentlessly in this field, helping to deliver an unparalleled virtual modeling and simulation capability, ensuring that the United States is able to commit to delivering a safe, secure, and reliable nuclear deterrent without the need to resort to underground testing.

The history of ASC is one of innovation, pioneering novel technologies for high-performance computing (HPC) and driving American technological leadership for decades. This foundational work is directly reflected in the nation's current standing, fielding the fastest supercomputers globally. A critical insight gained throughout this journey is that each exponential leap in computational power has invariably sparked a dramatic surge in scientific understanding and the demand for even greater computing resources.

The relentless demand for increased fidelity and accuracy of our models continues and forces us to look to emerging technologies as we continue to meet the needs of a complex and rapidly evolving national security mission. In this strategy, we lay out our vision for how quantum computing can augment our traditional HPC systems as a new resource to provide deeper, exquisite levels of scientific understanding. Put simply, quantum computing has the potential to provide insight into specific areas in a way which no other tool can match.

I thank the many contributors for their input in developing our quantum computing strategy, especially David Richards for shepherding the process and Meg Epperly for helping produce the final document with such care.

Our thoughts, bound within, will no doubt stimulate discussion, interest, and potentially debate. ASC cannot fulfill its mission or its vision alone—we invite collaborations and look forward to working with the innovative quantum computing industry, as we realize our desire to move computing into a new era.



Stephen A. Rinehart, PhD.

Assistant Deputy Administrator,
Office of Advanced Simulation & Computing and Institutional R&D

This page is intentionally left blank.

ASC Quantum Computing Strategy

Executive Summary

Quantum computing (QC) is a rapidly maturing technology with the potential for revolutionary impacts on stockpile stewardship science and national security. Recent developments in fault-tolerant architectures have compressed vendor roadmaps, and predictions of a production-ready quantum computer by the mid-2030s are becoming increasingly credible.

This strategy provides a roadmap for integrating QC into the Advanced Simulation and Computing (ASC) program by investing in four strategic focus areas:

1. Develop Capabilities in Mission-Relevant Quantum Applications:

ASC will prioritize developing quantum-ready applications in mission areas that have shown significant promise for quantum advantage, including simulations of materials in extreme environments, nuclear dynamics, solving linear and nonlinear partial differential equations, and uncertainty quantification. These applications directly support stockpile stewardship science and modernization objectives.

2. Conduct R&D in Algorithms, Software, and Hardware:

Sustained research into quantum algorithms, robust software tools, and quantum hardware is essential. ASC will develop efficient quantum algorithms; invest in quantum compilers, debuggers, and performance tools; and explore specialized quantum hardware tailored to NNSA's unique requirements.

3. Engage with Vendors and Partners:

Early and active collaboration with commercial quantum hardware vendors and academic partners is critical. Through testbeds, co-design agreements, and quantum demonstration facilities, ASC will influence hardware design, gain early access to emerging technologies, and ensure that quantum platforms evolve to meet mission needs.

4. Build Knowledge, Experience, and Workforce:

Expanding and upskilling the quantum-trained workforce is essential to long-term success. This includes hiring, internal training, university outreach, and postdoctoral support to ensure ASC maintains the expertise required to operate, program, and integrate quantum systems as they become available.

While quantum computing will never replace classical computing, it has the potential to solve certain problems with speed and accuracy that would be unachievable using any conceivable classical high-performance computing (HPC) system. By investing strategically in QC, ASC will help propel the emergent QC industry, maintain U.S. technological leadership, ensure mission readiness, and position itself to rapidly adopt quantum technologies as they mature.

This page is intentionally left blank.

1 The Role of Quantum Computing in Stockpile Stewardship Science

Quantum computing (QC) promises to deliver a variety of advantages over classical high-performance computing (HPC). Sufficiently powerful quantum computers are expected to solve certain problems much faster and/or with greater accuracy than any feasible classical computer. While QC will never replace classical computing (or even HPC) and current QC hardware is still in its infancy, ASC and NNSA will benefit from access to and participation in the development of gate-based quantum computers and related technologies. For ASC to remain competitive and provide research scientists with the latest technologies that can be utilized for NNSA mission execution, there must be strategic and sustained investment in anticipation of QC maturing to a point where it can realize revolutionary impacts for the unique problems of the NNSA.

QC has the potential to solve numerous problems in computational science and engineering, including many that are specific to the NNSA mission space. These include:

- High-accuracy simulations of materials in extreme conditions that are difficult, or sometimes impossible, to produce, even in multibillion-dollar experimental facilities.
- Detailed first-principles simulations of the nuclear dynamics of fission processes and reaction networks involving short-lived nuclei.
- Solutions of partial differential equations governing particle transport and diffusion, shock and radiation hydrodynamics, and a wide array of stockpile-specific engineering applications.
- Quantifying the sensitivities of these problems to both aleatoric and epistemic uncertainties.

The goal of this document is to lay out a strategy to ensure ASC is prepared for the transformative arrival of increasingly mature QC systems. This will require investing in four strategic focus areas:

1. Developing capabilities in mission-relevant quantum applications.
2. Conducting research and development (R&D) in algorithms, software, and hardware to enable ASC applications.
3. Engaging vendors and partners through testbeds and hardware/software co-design.
4. Building the knowledge, experience, and workforce necessary to procure, operate, and program QC systems.

By harnessing QC to perform simulations with speed and accuracy that would be unachievable using any conceivable classical HPC system, ASC can strengthen NNSA's responsive posture by providing on-demand, high-accuracy analytical capabilities for emerging stockpile challenges. QC will enhance our understanding of aging and performance issues, and higher-fidelity simulations will inform safety for longer-life weapons components and more precise life-extension decisions. Moreover, QC has broad U.S. competitiveness and national security implications beyond stockpile science, enabling new approaches in multiphysics/engineering simulations, uncertainty quantification, cryptography, and sensing. Quantum algorithms and quantum-enhanced sensors could also play a role in developing innovative tools to enhance arms control verification, nonproliferation, and incident response. Proactively investing in QC will ensure that NNSA maintains a technological edge and can respond adaptively to threats such as nuclear proliferation or terrorism with advanced computing analysis.

1.1 Maturation Roadmap for a Production-Ready ASC Quantum Computer

Vendors are broadly shifting to the development of fault-tolerant architectures and elements of quantum error correction have been demonstrated in all leading qubit technologies. The most aggressive public roadmaps suggest that systems with hundreds of logical qubits will be available by the end of the 2020s. While schedule risks remain and these roadmaps are extremely aggressive, a production-ready quantum computer by the mid-2030s is realistic. To avoid technological surprise, NNSA must proactively prepare to use increasingly capable QC platforms for its mission as they become available, while simultaneously adapting to unforeseen accelerations or delays in quantum hardware development. "There are only two ways to react to exponential change: Too early or too late."¹

A quantum operation (quop) is a single step in a quantum program. Larger and more complex programs require more quops, so the total number of quops a system can execute before errors affect the outcome is a useful metric for comparing different quantum computers and understanding their practical limits. Comparing systems on the basis of quops also emphasizes that the capability of a quantum computer is not determined only by the number of qubits it contains, but also by the qubit error rate.

Current estimates suggest that unambiguous quantum advantage in scientific computing applications is likely to be achieved with a teraquop quantum computer, i.e., a machine capable of performing 10^{12} quops without experiencing a fault that will

¹ Ethan Mollick, "Reshaping the tree: rebuilding organizations for AI," One Useful Thing (November 27, 2023). Accessed on July 13, 2025.

corrupt the result of a calculation. The only credible pathway to achieving teraquaop performance is with logical qubits and fault-tolerant operations. Figure 1 illustrates a roadmap toward this goal. While there is significant uncertainty in whether these scaling targets will be met, DARPA's Quantum Benchmarking Initiative is presently assessing vendor claims in coordination with the Department of Energy (DOE) and NNSA. NNSA can also reduce the risk of misalignment between future quantum hardware and mission-relevant use cases by helping vendors and the community with applications, benchmarking, and co-design.

Path to a Production-Ready Quantum Computer (PRQC)								
TODAY	2026+ →	Path to a Production-Ready Quantum Computer (PRQC)						
Transition out of NISQ era	Kiloquaop machines, est. '25-'26	Using "good" logical qubits,	Megaquaop machines, est. '27-'29	Using "very good" logical qubits	Teraquaop machines, est. '30-'33	Using "near-perfect" logical qubits	Petaquaop machines, beyond '33	Production-ready quantum computing for ASC
1-50 "minimal" logical qubits demonstrated: elements of error correction work, nuanced performance improvements	Ubiquitous demonstrations of "good" logical qubits with 2-4x improved performance	Universal logical operations won't be available, so demonstrations will focus on V+V of the hardware itself	1 "very good" logical qubit with $10^2\text{-}10^3\times$ improved performance	10-20 "very good" logical qubits enable universal logical operations and simulation mini-apps	100s of "near-perfect" logical qubits with $10^5\text{-}10^6\times$ improved performance	100s of "near-perfect" logical qubits enable beyond-classical exemplars	1,000s of "perfect" logical qubits with $10^8\text{-}10^9\times$ improved performance	Unambiguously super-classical scientific computing applications are possible
1 "good" logical qubit demonstrated: error correction improves performance 2x	Physical qubit quality improves a bit, becomes consistent at ~100 physical qubit scale – scaling further will lead to exponential improvement in performance	Physical operations will still be available for interesting noise-robust simulation tasks at the margin of quantum advantage	Physical qubits likely "good enough", logical scaling and integration becomes the biggest challenge – hockey-stick growth can begin	Capability benchmarking will become possible for testing subroutines in production-ready quantum algorithms	Marks the successful transition from a massive physics experiment to something closer to a traditional HPC system, no obstacles for continued scaling	The first complete science & engineering applications with unambiguous quantum advantage become possible	Fully mature quantum supercomputer capable of executing computational volumes associated with currently known challenge problems	Capable of mission-relevant simulations at scales comparable to classical HPC, but with improved accuracy

Figure 1. Development roadmap towards a production-ready quantum computer for ASC mission needs.

2 Develop Capabilities in Mission-Relevant Quantum Applications

The utility of a quantum computer will be measured by the types of problems that it can solve, especially those for which it confers an advantage over classical HPC resources. Several application areas in ASC's simulation mission have shown significant promise for quantum advantage. By developing quantum-ready applications for mission-relevant use cases, we will ensure QC efforts directly bolster NNSA's stockpile stewardship science and modernization goals.

2.1 Simulating Materials in Extreme Conditions

The ASC program is broadly concerned with the properties of materials in extreme thermodynamic conditions that are difficult and expensive to create and characterize. Recent advances in quantum algorithms, resource estimation, and architectures provide increasingly credible reasons to believe that sufficiently capable quantum computers could provide material models with accuracies that are classically inaccessible. Models of equation of state, phase diagrams, and transport properties (e.g., stopping powers, conductivities, and opacities), both in and out of local thermodynamic equilibrium (LTE), are all essential to the NNSA mission.

2.2 Nuclear Dynamics

Robust predictions for low-energy reactions involving the lightest nuclei (such as the neutron-induced breakup of ^6Li into ^3H and ^4He) are now becoming possible with advanced HPC architectures. Similar predictions at higher energies or for heavier nuclei are beyond the reach of foreseeable future classical HPC capabilities. As illustrated in Figure 2, quantum simulation of nuclear dynamics requires exponentially fewer resources on quantum computers relative to classical HPC. Quantum computing thus holds the promise of quantified, accurate, and predictive calculations of reaction, decay, and structural properties across the nuclear chart, including reactions on extremely short-lived nuclei that cannot be measured in the laboratory.

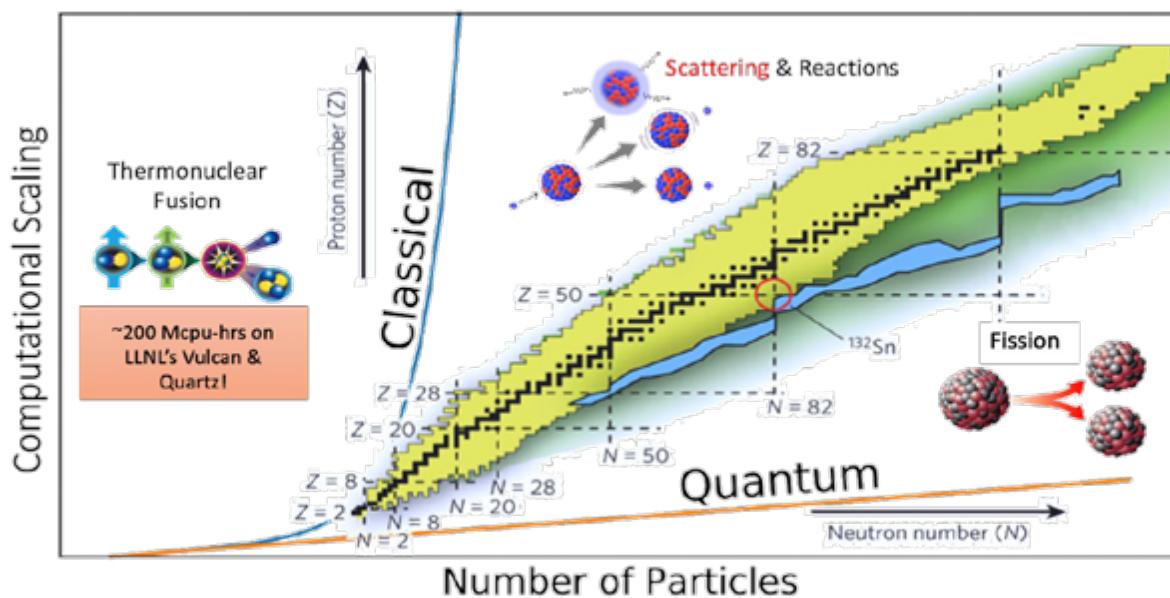


Figure 2. Illustration of the resources required to simulate nuclear dynamics on quantum computers compared to classical computers. QC enables simulations of heavier nuclei that would be inaccessible on classical computers.

2.3 Partial Differential Equations

The efficient solution of both linear and nonlinear partial differential equations (PDEs) on quantum computers holds promise for transformative advances across many scientific and engineering applications ranging from stockpile stewardship science to fundamental materials and nuclear science.

Linear PDEs

Solving large, sparse linear systems is a foundational capability for quantum-enhanced PDE solvers. Radiation diffusion and transport equations, which often dominate simulation run time, are typically linearized under semi-implicit schemes. A scalable quantum linear solver would thus directly accelerate radiation solvers, yielding significant mission impact. Moreover, quantum eigenvalue algorithms—already proving powerful in materials science—are being extended to nuclear criticality problems. Beyond providing key physical insights, these criticality calculations also serve as state-preparation primitives for time-dependent neutron transport simulations.

Nonlinear PDEs

Many high-fidelity physics models—shock hydrodynamics, radiation hydrodynamics, high-explosive detonation, turbulence, and coupled multi-physics systems—are intrinsically nonlinear. Though quantum algorithms for nonlinear PDEs are still nascent, initial results suggest exponential speedups for certain classes of problems. In addition, quantum-accelerated atomistic simulations often involve nonlinear differential equations at the microscopic scale and thus also benefit from advances in quantum algorithms.

2.4 Uncertainty Quantification and the Quantum AI Nexus

Quantum algorithms are expected to provide unique advantages for managing the high-dimensional spaces that commonly arise in uncertainty quantification (UQ) problems. For some conditions, classically inspired Monte Carlo sampling can realize quadratic quantum advantages. But quantum computers may also offer much larger advantages by taking an otherwise classically intractable approach and mapping high-dimensional classical dynamical evolutions onto Hamiltonian simulation instances. Much of the machinery developed for quantum simulation algorithms can be harnessed to deliver potentially exponential advantages in UQ. Given the use of UQ in many stockpile stewardship science applications, the prospective advantages of quantum algorithms for UQ could be particularly impactful for the ASC mission.

Akin to the emerging role of AI methods in uncertainty reduction in the ASC portfolio, the nexus of QC methods, HPC technologies, and AI has the potential to profoundly

influence workflows for uncertainty reduction. As an early example, emerging techniques to supplement machine-learned molecular dynamics potential functions for mesoscale materials simulations that rely on higher-quality DFT potentials could further benefit from higher-quality DFT potentials that are validated by first-principles, QC-based ground state calculations. Designing workflows that combine QC, AI, and HPC components will be a formidable, but extremely high-impact, challenge.

3 Conduct R&D in Algorithms, Software, and Hardware

NNSA must invest in internal R&D to help build essential components of the QC ecosystem. This includes designing and refining quantum algorithms tailored to ASC needs, developing software for integrating quantum co-processors with existing HPC workflows, and exploring custom hardware approaches. Investments in these three areas will also help ensure that ASC creates and retains the in-house technical expertise that will be needed to successfully evaluate and leverage QC hardware for our programmatic needs once it arrives.

3.1 Algorithm Development and Optimization

Sustained research into algorithms for quantum architectures will be essential for maintaining ASC leadership in simulation capabilities supporting stockpile stewardship science. Studying algorithms and subroutines extracted from applications will help produce resource-optimal implementations. For example, quantum phase estimation (QPE) is a ubiquitous algorithm used to calculate the eigenvalues of generic Hermitian matrices. However, there are many variants of the QPE algorithm, and implementations vary depending on the application. By tailoring the implementation of QPE and other algorithms to specific application use cases, we can ensure the availability of highly efficient building blocks as quantum hardware continues to improve.

Due to the range of possible quantum platforms that are or will become available, there is a need to pursue algorithmic and quantum control research for quantum simulation across a range of levels of abstraction. For future fault-tolerant quantum computers, research must focus on the development of more efficient algorithms with guaranteed upper and lower bounds on algorithmic complexity. In the near-term, analog quantum simulation methods are expected to provide novel insights into materials physics. For intermediate-term efforts, it is widely believed that there is an opportunity to demonstrate quantum advantage relative to the best classical algorithm for a given problem. The push toward quantum advantage is important both as a demonstration and as a means to push forward improvements in efficiency and fidelity of quantum algorithms leading towards fault tolerance.

3.2 Software Research Activities

HPC developers rely on a variety of software tools to build, debug, and analyze programs. A similar software stack will be needed to establish the bridge between quantum algorithms and quantum computers. Developers will need tools that can transform quantum algorithms to robust, efficient quantum circuits (i.e., quantum programs), and provide insight into program performance and correctness issues.

Code transformations and quantum compiling

Classical compilers routinely translate human-readable source code into highly optimized machine-level instructions. However, producing even a reasonably efficient quantum circuit from a problem description takes exponential time in the common case, limiting program scale to only a small number of qubits. Generalizing circuit generation to arbitrary problem sizes is a very active research area with many open problems. Quantum compilation needs significant research investment to match the convenience and capability we have come to take for granted in classical space.

Analysis tools

Classical simulation of quantum circuits can lead to better understanding of a quantum program's behavior, help establish a circuit's correctness, and improve developer productivity by reducing the need for access to precious quantum hardware during program development. By the very nature of quantum computation, classical simulation of quantum state evolution is a costly endeavor. There exist a few examples of quantum-circuit simulators that can exploit ASC capability-class supercomputers, but the potential for improvement remains.

Debuggers are invaluable for classical code development, so it seems likely that they will be equally necessary for quantum code development as algorithm complexity increases. Quantum debuggers are more challenging to construct than classical debuggers because a complete quantum state is not directly observable. It must be either statistically sampled or classically simulated, neither of which produces an accurate view of the state in a tractable length of time.

Gate compilation with quantum optimal control

On the current generation of quantum computers, accumulated errors can be reduced by minimizing the total circuit runtime. Using HPC resources enables custom gate compilation using pulse-level controls. These custom gates remove the overhead of compiling circuits into a fixed universal gate set. As a result, circuit design with gate decompositions chosen based on the problem or circuit of interest significantly improves effective results even without error correction.

3.3 Hardware Development Activities

ASC will prioritize QC hardware investments designed to ensure access to quantum computers that will be capable of solving problems in NNSA's unique problem space. This includes research in areas outside of mainstream academic and industrial research that would not be explored or funded otherwise.

Significant opportunities also exist to develop custom quantum hardware that emulates the complex dynamics observed in problems of interest to ASC. Specialized quantum hardware has already been demonstrated to mirror chaotic differential equations, nucleon scattering, and material behaviors. A focused effort on co-design with subject matter experts, low-level access to the hardware, and HPC-enabled control is likely to enable simulations that can achieve quantum advantage.

ASC has access to two on-site hardware testbeds, the Quantum Device and Integrated Testbed (QuDIT) facility based on superconducting technology and the Quantum Scientific Open User Testbed (QSCOUT) facility based on trapped atomic ions. These on-site testbeds offer ASC scientists full access to all elements of the computational functions, from the radio-frequency signals that control the quantum state, through calibration routines for mitigation of slow errors, to the measurement signals that interrogate the system state. This level of access allows scientists to tailor quantum simulations to their problem and explore low-level questions about how to best operate and optimize quantum devices for ASC applications. Commercial and cloud systems typically provide opaque interfaces that do not allow such explorations. On-site testbeds also provide opportunities to develop the specialized infrastructure required to house quantum computers and the classical computing resources needed to interface with and verify quantum computations.

ASC's classical HPC resources will also play a role in quantum hardware development. Modeling and simulation that elucidates the behavior of quantum hardware is essential to the development of better qubit technologies. HPC resources can also be used to investigate the implementation of optimal control protocols. Finally, large-scale simulations of noisy quantum circuits can be executed on these platforms as a part of the characterization, verification, and validation of commercial and NNSA lab-based quantum hardware.

4 Engage Vendors and Partners Through Testbeds and Co-design

Early engagement with multiple commercial vendors is critical to secure early access to technology and shape the next generation of quantum hardware before products

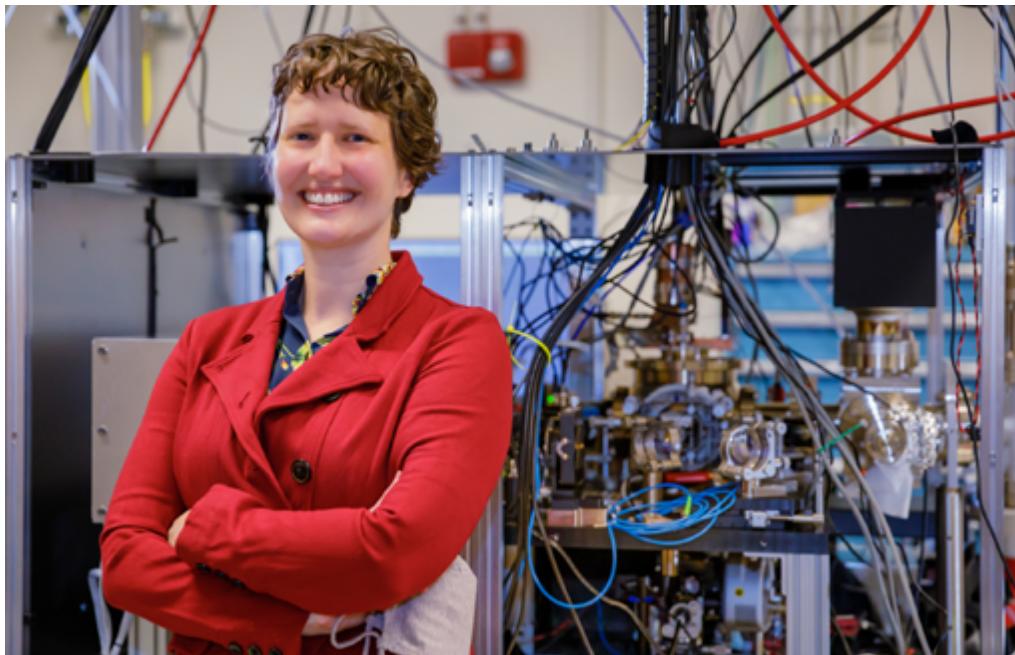


Figure 3. Susan Clark is the QSCOUT Principal Investigator at Sandia National Laboratories. The QSCOUT testbed is available to the research community for a range of QC applications.

mature to the point where changes are prohibitive. To this end, the ASC Tri-Labs (Sandia, Los Alamos, and Lawrence Livermore National Laboratories) have established strong working relationships with technical teams at leading vendors including (in no particular order) Quantinuum, IBM, Google, IonQ, QuEra, PsiQuantum, Xanadu, Rigetti, and D-Wave, as well as the cloud intermediaries AWS Braket and Azure Quantum. Because the level of maturity and the rate of advancement are highly variable among qubit technologies, there is substantial value in maintaining access to multiple vendors that span the space of credible qubit technologies.

Access to multiple credible qubit technologies may be realized through quantum demonstration facilities (QDFs), modeled after the successful Oak Ridge National Laboratory's Manufacturing Demonstration Facility (MDF). QDFs bridge the gap between testbeds like QSCOUT and QuDIT and any eventual ASC production-level QC facilities. QDFs provide industry with affordable and convenient access to infrastructure, tools, and expertise to facilitate rapid development of QC technologies. They also create platforms for NNSA scientists, modelers, and engineers to gain hands-on experience in fast-paced, iterative hardware-software development. Depending on the QDF location and supporting partnerships, experts from academic institutions could also contribute to QDF workflows. Finally, since QDFs support workforce and economic development at their locations, the cost of building and outfitting these facilities may be shared with state and local governments, amplifying the impact of ASC's investments in this strategy.

ASC has more than two decades of experience leading hardware/software co-design activities across the HPC ecosystem.² By using the same programs and techniques that have proven successful for classical HPC to influence QC, ASC will continue to strengthen our partnerships with vendors. Through QDFs, co-design agreements, and testbed programs, ASC can work alongside vendors to explore the integration of QC with classical HPC systems and tailor hardware features to NNSA mission requirements. This includes exploring opportunities for direct NNSA-funded hardware research in areas not pursued by commercial entities but critical to national security applications. ASC can also influence vendors' development direction by communicating our algorithmic goals via proxy applications and benchmarks.

5 Build Knowledge, Experience, and Workforce

Building a minimal QC research staff who are aware of the ASC mission space has taken nearly a decade at the Tri-Labs. Rebuilding this capability would be far more difficult now with the intense competition for quantum expertise from the commercial sector. ASC will expand our quantum-proficient workforce through hiring and retaining top talent, developing up-skilling programs for existing staff, and collaborating with DOE Office of Science and academia, ensuring we have the talent necessary to program and operate quantum machines. Targeted university outreach and support of postdoctoral research are also essential to train a new generation of quantum/HPC scientists and engineers to fill positions at the Tri-Labs.

ASC support for a broad spectrum of QC research activities will have the additional benefit of exposing staff to the science required to utilize quantum computers when application-scale machines exist. This is particularly valuable for communities at the laboratories that do not typically share their research externally. Having systems and hardware experts at the laboratories will allow the NNSA national laboratories to better understand resource requirements for NNSA applications. This is similar to historical HPC development, where the laboratories had to be heavily involved in the design of supercomputing resources to better achieve their goals.

6 Investment Priorities, Facilities, and Roadmap Phasing

Successful execution of this strategy will require a phased investment approach that balances urgency with preparedness and incorporates regular technology assessments to ensure investments align with the most current understanding of QC maturation. A sustained effort on algorithms and workforce coupled with increasing

²J. Ang, T. Hoang, S.M. Kelly, A. McPherson, and R. Neely, ASC Co-Design Strategy (2015, SAND2015-9821R).



Figure 4. The Quantum Computing Summer School (QCSS) at LANL has been giving top students research experiences in QC since 2019. The 2025 class of 20 students was selected from over 1000 applicants. Students spend 10 weeks at LANL over the summer working directly with research staff on cutting-edge problems in QC. The QCSS has resulted in more than 50 research papers, many of them in high-impact journals.

hardware investments as technology matures will ensure the technology is shaped to our requirements, and systems are deployed when needed. Delaying hardware acquisition, on the other hand, risks missing the window to influence design and could leave ASC unable to catch up without much greater expense.

Near-Term (1-4 years): Emphasize algorithm and application development, as well as prototyping using existing small-scale quantum processors and classical quantum simulators. Strengthen vendor co-design engagements with at least one and ideally many quantum hardware providers to influence system architecture early—even modest funding now can secure custom features or priority access that would be impossible to obtain later. Support ongoing hardware research, algorithmic prototyping, and growth of the system software ecosystem using on-site quantum testbeds. Establish QDFs to inform the design of a classified QC environment and ensure readiness when production hardware arrives. This includes dedicated R&D into the unique security, facility, and operational requirements for integrating quantum hardware into NNSA’s classified infrastructure.

Mid-Term (4-8 years): Procure at least one QC testbed tailored for NNSA needs (potentially a few hundred physical qubits, with a path to fault tolerance). Continue to expand both the application and system software ecosystem: develop the libraries, compilers, and tools to integrate the quantum testbed with ASC’s HPC-based workflows. Ramp up workforce training: use DOE QDFs and their resources, send staff to vendor labs or QC courses, recruit quantum specialists, and build internal QC communities of practice.

Long-Term (8+ years): Deploy increasingly more capable quantum systems and transition from experimental use to production science use cases within fully secure and accredited classified environments. At this stage, the investments in algorithms, software, and workforce will begin to be realized: NNSA will be able to run high-value calculations on the quantum accelerator, with results feeding directly into stockpile decisions. Continued engagement with industry will ensure that upgrades and next-generation hardware keep pace with the mission. Continued focus on workforce training will help QC become deeply embedded in ASC's simulation toolkit, working alongside exascale classical computers and AI workflows. It is expected that a large-scale fault-tolerant quantum computer (with hundreds of logical qubits) could perform hundreds of millions to billions of operations that dramatically accelerate R&D in materials science, chemistry, and optimization relevant to national security.

7 Conclusions

Quantum computing promises revolutionary impacts for stockpile stewardship science and national security, offering new approaches to solving mission-critical problems that are beyond the reach of classical high-performance computing. By strategically investing in quantum applications, algorithmic innovation, and software tools, ASC will help shape the evolution of quantum technologies. Early and sustained engagement with commercial vendors, academic partners, and internal testbeds will position ASC to influence hardware and software development, provide secure access to cutting-edge platforms, and integrate quantum solutions seamlessly into existing HPC workflows.

As QC matures, ASC's proactive approach will be critical to maintaining technological leadership and mission readiness in an era of rapid and unpredictable change. The phased investment roadmap outlined in this strategy provides the flexibility to adapt to breakthroughs and setbacks alike, while building a robust foundation of talent and infrastructure. By fostering innovation and collaboration across the Tri-Labs and with external partners, ASC will be prepared to harness a production-ready quantum computer as soon as it becomes available, driving forward the frontiers of science, engineering, and national security for decades to come.

Appendix A: Classified Addendum Available on Request

