

# **EXASCALE COMPUTING CHALLENGES AND OPPORTUNITIES**

Hearing of the House Science, Space, and Technology Committee

Subcommittee on Energy

U.S. House of Representatives

May 22, 2013

Dona L. Crawford, Associate Director for Computation

Lawrence Livermore National Laboratory

## **OPENING REMARKS**

Thank you, Madam Chair, Ranking Member Swalwell, and Members of the Committee. I am Dona Crawford, Associate Director of Computation at Lawrence Livermore National Laboratory (LLNL). I welcome the opportunity to provide my perspective on high performance computing (HPC). There are major opportunities and challenges associated with developing exascale computing, the next generation of HPC capability. At the same time, exascale computing is critically needed to support national security priorities, advance science and technology, and enable greater innovation in U.S. industry.

I applaud the Committee for its determined support to maintain U.S. leadership in HPC. The U.S. has benefited immensely from the investments the nation has made in developing and applying HPC capability. The U.S. has enjoyed a long record of success and unparalleled leadership in computing for many years. The world took notice, and many other nations started investing in HPC and now strive to challenge U.S. leadership in this critical arena. It is imperative that we continue to make smart investments of our limited resources to maintain U.S. HPC leadership.

In view of today's fiscal constraints, we must apply the lessons learned in our past successes to strategically target investments to make crucial early steps on the road to exascale-level computing. To cost effectively and efficiently maintain U.S. leadership in HPC, the nation must build upon and leverage programmatic and technical approaches that established the U.S. as the leader in innovative HPC systems over the past half-century. In particular, next generation HPC must be developed through an integrated partnership of the Department of Energy's (DOE) National Nuclear Security Administration (NNSA) and the Office of Science (SC). Both technical and fiscal responsibilities must be shared, taking advantage of the core capabilities of the partners—and working closely with industry. This includes balanced investments in both the ongoing core HPC computing programs and breakthroughs necessary to achieve exascale HPC. Leading-edge HPC is vital for U.S. national security and science missions and to advance U.S. economic competitiveness goals.

I thank the Committee for its support for the development of next generation supercomputing and its recognition that the dialogue must integrate NNSA and its laboratories in this effort.

## **NNSA's and LLNL's Role in and Reliance on Supercomputing**

HPC has been a core competency of the nation's nuclear weapons enterprise from the birth of HPC in the 1950s and has been essential to the nation's ability to develop and maintain a nuclear deterrent. LLNL is one of the DOE/NNSA's national security laboratories with responsibility for maintaining the safety, security, and effectiveness of our nation's strategic deterrent. HPC has played an increasingly important role in LLNL's nuclear deterrence mission since the cessation of nuclear testing over 20 years ago and the creation of the Stockpile Stewardship Program (SSP).

Current U.S. leadership in HPC is a direct result of the nation's investment in computational capability for the support of the SSP. After the U.S. decided to forego underground nuclear testing, DOE Defense Programs (NNSA's predecessor) embarked on a focused effort to develop vastly improved computational capabilities, along with advanced experimental capabilities, as a foundation for the SSP that would provide the scientific basis for maintaining the nation's nuclear deterrent without nuclear testing. This was and continues to be a grand technical challenge.

Nuclear weapons are extremely complex devices, with thousands of components made from a variety of materials that must work together seamlessly to produce a nuclear detonation. As they age, nuclear weapons are subject to environmental conditions that pose a number of challenges that affect performance of components and the weapon itself. Plastics can break down and give off potentially destructive gases, metals can corrode and weaken, and coatings can deteriorate. Some materials may change properties unpredictably in response to the high radiation fields, fluctuating temperatures, and other environments to which nuclear weapons are subject.

Congress acknowledged the challenges associated with maintaining the nuclear deterrent without testing by creating in the mid-1990s an initiative within the SSP to rapidly develop substantially more powerful computational, simulation, and modeling capabilities. At the time, DOE was not an HPC technology driver as it had been in earlier decades. If the Accelerated Strategic Computing (ASCI) Program had not been formed and aggressively funded, the HPC industry would have continued to evolve toward serving its consumer base, toward business and industry focused server solutions of relatively limited capability—and later toward gaming applications—and would not have enabled the success of the SSP.

ASCI was highly successful and has since evolved into the current program, called Advanced Simulation and Computing (ASC). The NNSA laboratories (LLNL, Los Alamos National Laboratory and Sandia National Laboratories) and our industry partners worked to develop computer architectures that would enable the laboratories to run large-scale, high-fidelity simulations integrating data from past underground nuclear tests and experimental capabilities to continue to assess and certify the safety, security, and reliability of the nation's nuclear deterrent. This effort spearheaded the revolutionary design shift in HPC computer architecture and applications development that occurred over the last two decades. ASCI and ASC in the SSP—and later Office of Science—pushed the extreme limits of what was possible, and this ultimately led the way for more competitive business and industrial related computing and simulation. The research and

development (R&D) has led to unprecedented advances in HPC and remarkably capable computing systems that now are becoming ubiquitous and are impacting scientific discovery and industrial competitiveness.

Time-urgent questions about the safety and reliability of the nuclear stockpile drove the DOE and its NNSA laboratories to invest in and develop supercomputers. These systems, developed by industry in response to national security demands, have dominated global HPC performance. In the last forty Top500 lists, the U.S. has held the top position 26 times, with NNSA HPC systems in 23 of these cases. Through investments in the ASCI and ASC programs, computing has become the single *integrating* element in assessment and certification of the U.S. nuclear weapons stockpile. As the global nuclear security landscape has evolved, these same computational tools are now being continuously applied to combating nuclear proliferation and bolstering counterterrorism—both nuclear and conventional.

The enormous success of ASCI/ASC has been a result of:

- sustained support for a Congressional initiative to develop HPC simulation as a pillar with the Stockpile Stewardship Program,
- leadership in DOE and NNSA driving the development of high-end computer architectures and associated simulation software designed for our unique mission requirements and targeted at specific challenges associated with the stockpile, and
- unprecedented level of cooperation between the national laboratories and industry to co-design/co-develop software and computational platforms required for the mission.

The importance of these three elements should not be underestimated. It is only through this initiative's combination of commitment, leadership, and cooperative R&D that U.S. computing made a revolutionary design shift in computer simulation required to ensure the safety, security, and reliability of the nuclear stockpile. As a result, computing capability used by NNSA's national laboratories increased over a million fold.

Each time the laboratories and their industry partners to develop a new generation of supercomputers, we discover new science that helps understand material phenomena and performance of the aging nuclear stockpile with a higher degree of accuracy. Today, the Sequoia machine at LLNL—a breakthrough ASC system with over 1.5 million processor units or cores, and 1.6 petabytes of memory—serves as a bridge between supercomputers of the past decade and exascale computers of the future. Indeed, it is arguably the first of the new era of daunting “many-core” computers. The machine's extraordinary capabilities are being used to improve models of weapons physics, particularly in the areas of hydrodynamics, radiation transport, and the properties of materials at extreme pressures and temperatures. In addition, Sequoia is able to run large suites of calculations designed to characterize uncertainties in weapon performance resulting from small variations in the weapon system and uncertainties in the physics models used.

Improved capabilities for uncertainty quantification (UQ) are essential for assessing the impact on performance of physical changes in aging weapons and for certifying stockpile Life Extension Programs (LEPs). Sequoia, with its 20-petaflop capability, can effectively address many stockpile issues through the use of UQ with two-dimensional (2-D)

applications. However, the system can only provide “entry-level” capabilities to run suites of three-dimensional (3-D) weapons physics simulations for UQ. It remains the role of exascale-class systems to address the full breadth of issues that will arise as the stockpile ages, as significant findings are identified, and as even more advanced safety and security features are added. In the future, 3-D predictive UQ analysis will become essential for essentially all aspects of work.

### **National Security Mission Need for Exascale**

With the planned modernization of the stockpile and simultaneous decrease in both its overall size and composition, advanced computing and simulation will play an increasingly critical role. A thousand-fold improvement over today’s modeling and simulation capability (exascale technology) is required over the long term to assure with confidence the safety, security, and performance of the nation’s nuclear stockpile. These more capable computers are needed to run large suites of high-fidelity simulations to fully map out the impact of uncertainties.

Nuclear weapons are engineered 3-D systems with complex materials that change over time as they age. Today, we do not have the computing power to simulate weapons performance in 3-D at the required resolution and incorporating detailed physics and age-aware material models. Additionally, we do not have the computing power to conduct the tens of thousands of high-resolution 3-D simulations to quantify the variation on weapon performance taking into account uncertainties in our modeling capabilities.

Through our success in developing and applying advanced HPC, we have resolved the energy balance problem—one of the physics issues remaining unresolved at the cessation of the underground nuclear test program. Today’s available technology allows us to simulate in two dimensions at high resolution and physics fidelity, or simulate in 3-D at moderate (not high) resolution, but today’s available technology does not enable our weapons specialists to simulate at high resolution and 3-D simultaneously. We are looking forward to HPC systems that are also capable of bringing to closure the grand challenge of modeling the physics phenomena of boost.

Examples of the role exascale computing will play in the continued maintenance of the stockpile include:

- **Warhead assessment and certification of smaller stockpiles:** As the stockpile decreases in size, the performance of individual weapons becomes increasingly important. Higher fidelity, 3-D simulation of warheads including detailed representation of initial conditions, engineering features, safety features, and security features are required to ensure the safety and performance of each weapon. These simulations will each require between 0.5-10 exaflops.
- **Material aging, compatibility, and acceptance of modern efficient manufacturing processes:** As weapons continue to age, the complexity of issues to be resolved could increase exponentially. To address potential material related issues in the life-extended U.S. stockpile, more detailed weapons science calculations will be required. Simulations at increased resolution that capture real materials (micro-structure,

interfaces, kinetics) as opposed to simulations based on simple models of bulk material properties require between 0.5-100 exaflops.

- **LEPs:** Confidence in the assessment and certification of future life-extended warheads will be informed by suites of high-resolution integrated weapons simulations to quantify and bound uncertainties in performance. Detailed uncertainty quantification (UQ) will require routinely running between 1,000 and 100,000 simulations per study in order to rapidly converge LEP design options. Each of these UQ suites and the confirmatory steps will require between 10 and 1000 exaflops.
- **Safety and Security:** Enhancing weapon safety and security to address the 21st-century threat environment (including non-state actors) is just one example of a potential LEP goal. The development and certification of even more advanced safety, security, and use control (surety) features that can be embedded in a nuclear warhead as part of an LEP may require at least 10-100 exaflops.
- **Boost:** Boost—the process of boosting the fission yield of weapon primaries—is key weapons performance, not well understood, and among the most challenging physical phenomena to model. Greater computational power is needed to apply the improved physics models developed within the ASC *Predictive Capability Framework* (PCF) in large ensembles of weapons simulations. Large ensembles are required for rigorous UQ explorations using those models. We estimate 10-100 exaflops is required for resolving the largest known uncertainty associated with boost.

### **Technological Challenges to Achieving Exascale Must be Overcome**

The development of exascale-class systems cannot be achieved through a straightforward refinement of today's technologies. Surmounting multiple technical issues will require sustained research and development and some key breakthroughs.

Succeeding generations of microprocessors, standard in computers of every scale, have grown faster by increasing the speed and shrinking the size of transistors, effectively packing more calculations into every unit of time and space on a computer. But now transistors are reaching a lower limit in size and an upper limit in speed. Although individual transistors could be pushed to run faster, speeding up millions of transistors on a microprocessor would drive energy demands and operational costs to unsupportable levels. To make exascale computing practical, the electrical power requirements must be reduced at least ten-fold per floating point operation. Without this reduction, exascale computers would need hundreds of megawatts—enough to power a small city—at an unacceptable cost of hundreds of millions of dollars a year to pay the electricity bill. Developing a low power system encompasses changes to every component of the HPC system: memory (e.g. stacked memory), networks (e.g. optics), and processors (including accelerators).

Exascale systems will be comprised of tens to hundreds of millions of components. Calculations run on these systems will require tolerance for component failure as well as the management of up to a billion separate, but coordinated threads of execution—in

short, the systems must become self-aware and must compensate, in real time, for failures.

In addition, memory and storage will be challenging. Science and nuclear stockpile applications require large amounts of memory per core. At exascale levels, data movement—not attaining a higher amount of flops—will be the performance bottleneck. This includes movement of data from memory to the processing unit as well as movement of data across the machine, from one processing unit to another processing unit. Current memory bandwidth and network technology is too slow, costly, and unreliable to support the millions of trillions of calculations per second required in an exascale machine. At the exascale level, more components mean higher failure rates. In addition, the faster the data moves, the more error-prone it becomes.

Finally, exascale computing will require new programming models that allow software developers to exploit unprecedented parallelism. Applications in the future may have to support upwards of a billion separate independently-executed instructions to efficiently use the hardware. This will require our scientists to find ways to break their problems into many more independent pieces than even today's largest computers support. They will have to re-think their entire solution approach to meet this challenge.

The U.S. has invested heavily over the last two decades to develop nuclear weapons simulation codes to maintain the deterrent. As we develop the next generation of supercomputers, it is of paramount importance that NNSA partners with SC to minimize disruption in the utility of current codes. If HPC technology were to evolve in a less controlled manner, there is the potential that our existing codes will become ineffective and possibly obsolete before we have time to rebuild them from scratch to operate efficiently on new architectures.

### **R&D Program Required to Attain Exascale HPC**

Attaining and harnessing exascale computing will require an integrated, decadal program of technology development and testing that balances technology push with applications pull. This requires the combination of a sustained federal investment, commitment to deliver hardware and software tuned for core missions, and strong partnerships between American industry and DOE laboratories tasked to define the exascale ecosystem. That ecosystem needs to support the combined national security, economic competitiveness, and science missions.

The overriding theme in an efficient and effective R&D program for exascale is co-design and co-development. Design and development are linked because mission application requirements influence computer architecture design and architecture technology constraints influence the formulation of algorithms and software in mission applications. To be effective in addressing stockpile mission requirements, future HPC systems must be capable of simulating the complex physics and material changes in a nuclear weapon. It is essential the NNSA laboratories are involved from the start in this daunting but necessarily seamless transition from today's systems to exascale. It is not sufficient to assume that next generation HPC computing developed for either mass

consumers or for science applications will meet national security mission needs. The U.S. government must invest in the technology development effort in a way that ensures an architecture that meets national mission requirements.

The mission drive of NNSA and the leveraging of NNSA's outstanding systems-engineering track record in delivering leading edge HPC will serve to focus technology development on the path to exascale and lead to an architecture that will serve our national security mission requirements. I endorse SC's early focus on long lead-time research and development in advanced technologies. I strongly support NNSA's continued focus on investments that support the design and the delivery of well-balanced and well-architected HPC systems used to meet mission requirements. Both agencies today are working together to maintain their current HPC capabilities that are of vital importance to meeting near-term mission deliverables for NNSA and SC. But these agencies have very limited resources for advanced technology efforts required to achieve exascale. Thus, a balanced alliance between SC and NNSA in the next three years, supported by Congress, in pursuing R&D of balanced and innovative systems is a cost effective strategy.

The HPC R&D program required to meet NNSA mission needs in the next three years includes:

- Development of a scalable design for advanced system architecture at 100-200 petaflops. This effort couples NNSA systems requirements with industry expertise and best technologies. By funding advanced systems architectures and scalability, development and design through this integrated research and system delivery program, we can minimize the potential drift of next generation hardware away from our mission application needs.
- Acquisition of 100-200 petaflop systems in the FY2016-2017 time frame will enable prototype builds on the path toward exascale and will deliver NNSA required mission capability. Through these interim prototype systems, potential new technologies can be evaluated in the context of new architectures that utilize them to solve the various challenges. To maximize cost effectiveness of this approach, NNSA laboratories and SC laboratories are combining forces. Argonne National Laboratory, Oak Ridge National Laboratory, and LLNL as one team and Los Alamos, Sandia, and Lawrence Berkeley National Laboratories as another are partnering to pursue system acquisitions in 2017 and 2015, respectively. For cost effective risk reduction, an exascale initiative should pursue a minimum of two technology tracks in future acquisitions.
- Co-development of new exascale algorithms, applications, tools and runtime environments are needed by programmers to achieve maximum sustained performance on exascale systems. Today's largest systems have greater than a million cores, while an exascale system is expected to require billion-way parallelism. To be able to harness these new hardware technologies and architectures for our nuclear weapons mission, we must immediately begin to develop software tools that enable scalability, programmability, fault tolerance and code portability. As described earlier, this kind of investment and

involvement in co-design is essential to ensure we do not have to rewrite millions of lines of computer code. This would add substantial cost and threaten NNSA's ability to maintain the deterrent.

### **Broad Benefits of Supercomputing Capability**

Over the last two decades, supercomputers have transformed the way the world conducts scientific research and enabled discovery and development across a broad set of disciplines from physics, chemistry, and bioscience to engineering. Simulation—the ability to virtually mimic physical phenomena with great accuracy—is now considered a peer to theory and experiment, and a pillar of the scientific method pioneered by Isaac Newton more than 300 years ago. HPC simulations have advanced medicine, energy, aviation, and manufacturing domains. In the 2008 U.S. Council on Competitiveness report “The New Secret Weapon,” the Council states “Supercomputing is part of the corporate arsenal to beat rivals by staying one step ahead of the innovation curve. It allows companies to design products and analyze data in ways once unimaginable.” Forefront HPC has moved from a tool developed and used by the national security laboratories like LLNL, to a critical tool for the U.S. science laboratories and Fortune 500 companies. The massive, complex simulations that run on today's HPC allow us to explore fields such as global food, water, and energy supplies, as well as tackle problems for which experiments are impractical, hazardous, or prohibitively expensive.

LLNL has certainly leveraged its high performance computing capability and applying it beyond the nuclear weapons program to other important mission area. New tools and expertise developed in other mission areas at LLNL can then be brought to bear on maintaining the nuclear deterrent. The multi-program utility of HPC capability and the joint benefits of applications to the weapons program and other mission areas are illustrated by three examples:

In 2010, the Department of Defense urgently tasked LLNL to develop an advanced conventional munition in record time. Based on the Mark 82 steel case form factor, LLNL combined a novel explosive design with a carbon fiber case that met the military need for a lightweight weapon that could deliver lethal effects with low collateral damage. Using LLNL's HPC simulation codes, originally developed in the nuclear weapons program, the team was able to accurately predict warhead performance under dynamic conditions and achieve desired strength properties by optimizing critical design features and tailoring fiber-composite winding patterns. This extensive use of HPC, allowing accurate design simulation, eliminated costly and lengthy iterated cycle of developing and testing prototypes. After only one set of proof tests, the BLU129/B advanced conventional munition was deployed in theater only 10 months after a Joint Urgent Operational Need was identified. The historic average of new munitions development is 4.5 years.

In the last year, Sequoia demonstrated its great scalability with a 3-D simulation of the human heart's electrophysiology. Using a code called Cardioid, created in a partnership between LLNL and IBM scientists, researchers are modeling the electrical signals

moving throughout the heart. Cardioid has the potential to be used to test drugs and medical devices, paving the way for tests on humans. The code, running in nearly real time across the 20-petaflop system (an astonishing 60 beats a minute) predicted an arrhythmia that was known to occur with the injection of a drug. The fact that a calculation of this complexity ran at 59% of peak on Sequoia is astonishing. Cardioid demonstrated outstanding scalability and time-to-solution over 1200 times faster than previous state of the art, and the simulation is performing within 12% of real-time. The work showed the promise of advanced computing, but it also demonstrated the extreme level of specificity and technical acuity required to achieve this result. The insights gained and techniques employed by the code team are proving useful to Sequoia's national security applications.

Through DOE economic competitiveness initiatives, a number of American HPC centers and laboratories—including Lawrence Livermore—are making large-scale HPC resources available to U.S. companies both large and small. In LLNL's California Energy Systems for the 21st Century (CES-21) initiative, the California Public Utilities Commission and state investor-owned utilities are collaborating with LLNL to improve and expand energy systems to meet future needs. The owners, operators, regulators, and a joint team of technical experts will use the nation's most advanced modeling, simulation, and analytical tools to gain unprecedented insight and generate new data—information that can reduce risk and inform solutions to issues facing 21st-century energy systems, such as renewable energy integration and use of smart-grid technology. CES-21 will benefit from LLNL's extensive experience in national security supercomputing. Our expertise will be utilized to perform realistic and verifiable tests of how utilities will need to operate in the future.

### **The U.S. Must Lead in High Performance Computing**

The U.S. is entering the 21st century as the global leader in HPC, with the vast majority of high-end computer systems produced worldwide using U.S. technologies. However, our leadership is not undisputed. U.S. leadership faces an unprecedented challenge, principally from China, but also other global players.

China is steadily increasing its investment and marshaling its technological capabilities and its state-owned industries to quickly develop next-generation supercomputers. China has indicated its intent to use HPC for oil exploration, aircraft design, and weapon system design. China's 5-year roadmap demonstrates its commitment to assume and maintain a leadership position in HPC. A report from China's Academy of Sciences has stated that China has fast-tracked the development of exascale computing with the intention of using indigenous technology to field an exascale system by the end of this decade. They are funding a smart, systematic strategy to develop three different approaches to key hardware technologies such as chips, operating systems, file systems, and networks.

This represents a dramatic change from just a few years ago. Over the years, from 1993 until November 2001, the Chinese had zero, one or two systems on the Top500 list. Press reports called it a "Sputnik moment" for the U.S. following the release in November 2011 of the Top500 list of the world's fastest supercomputers. For the first time, Chinese

supercomputers vaulted to the number one and three rankings on the list, displacing HPC systems at SC and NNSA labs. China has continued its focused investment in HPC and now has 72 systems in the Top500 list. This investment strategy has allowed them to surpass Japan and Europe in the numbers of systems on the Top500 list. Without dedicated U.S. investment, we risk ceding our leadership to China and eroding the U.S. HPC industrial base so important to our national security and overall economy.

Given our current global leadership position in HPC and our proven capability to innovate HPC simulation advances, the U.S. is positioned to continue HPC leadership if there is sustained R&D investment. The country that develops the key intellectual property for the next generation supercomputers will also control the high ground for standards-based decisions to be made over the coming decades: instruction sets, programming tools, I/O, visualization and protocols. This is a level of control that the U.S. has taken for granted due to its past domination in this arena. In addition, the low power-consuming technologies to be developed can be applied to the entire hierarchy of server offerings, down to the cell phone, providing almost unthinkable leverage and intellectual property value. Consider the position China could have in military operations if it controlled the low-power technology, next generation un-manned surveillance, and propagation of intelligence analysis to inform battlefield awareness in hand-held devices. Trusting U.S. nuclear weapon technology and even more sensitive national security technologies to Chinese-built HPC systems is untenable. We must invest in leadership HPC to support our national security interests and maintain a healthy U.S. HPC industrial base.

### **Comments on Draft Exascale Legislation**

I am very encouraged by the demonstration of bi-partisan and bicameral support for continued U.S. investment in HPC and efforts to develop exascale-capable computing. This is a critical capability that is essential to U.S. competitiveness in multiple fields. I commend Representative Hultgren for his leadership in moving forward an authorization bill for the exascale R&D effort. It is imperative that the U.S. embark on an R&D program to develop new technologies and computer architectures to support exascale computing. This needs to be a joint Office of Science (SC)-NNSA effort leveraging the strengths and expertise of SC and NNSA laboratories in close partnership with the U.S. HPC industry. Recognizing the jurisdictional boundaries of Congressional committees, I note that the draft bill only authorizes funding for SC participation in the exascale effort. It is imperative that SC and NNSA move in tandem in this R&D effort for cost effective leverage of resources across the Department and to ensure systems architecture to meet SC and NNSA mission requirements. Due to the technically challenging nature of developing exascale supporting technologies and computing capability, it is vitally important to ensure that there are competitive teams of SC and NNSA laboratories partnered with U.S. HPC industrial collaborators. Equally important is the development of an integrated strategy and program management plan. I believe the draft bill, by and large, addresses these critical elements of a successful exascale R&D effort and opens a door for further discussion on how to most effectively structure a joint SC-NNSA exascale R&D effort.

I look forward to working with the Committee to ensure the U.S. retains its HPC leadership.

### **Conclusion**

Although supercomputing may lack the glamour of the space race, U.S. leadership is critically needed to achieve key national priorities. Failing to maintain U.S. HPC leadership has consequences beyond national security, reaching much further and more broadly into our economic future. Supercomputers have become a differentiating tool for discovery and innovation, with profound impacts on science, national security, and industrial competitiveness. HPC at the exascale level will be a powerful lever to influence outcomes and foster prosperity and security as we face uncompromising competition in an uncertain world. In its 2008 report “The New Secret Weapon,” the U.S. Council on Competitiveness said, to “out compute is to out compete.” If we are to be partners in a world of global competition, I want us to come from a position of strength based on the best U.S. industry, academia, and the national laboratories have to offer. That is what put us and has kept us in the leadership role we enjoy today in supercomputing. It is imperative we now begin to push forward on the necessary technology to ensure a continued leadership position. The stakes are very high. A robust multi-year effort harnessing a partnership between DOE’s Office of Science and NNSA with industry is key to national security and science, which underpins competitiveness.

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