Twenty Questions with the NNSA

D. Kusnezov

Dr. Dimitri Kusnezov is Director of the Office of Research and Development for National Security Science and Technology.

SciDAC Review: *Your supercomputing needs seem insatiable. How much is enough?*

Dr. Kusnezov: How much is enough really depends upon the questions you are asking. For example, if I were to ask you, when will you have done enough fitness training? then I am sure you would consider what it is you are training for. If you knew that you were training for an important event but you did not know what event it was, or even when it was, then you would have to establish a flexible and versatile fitness program that covered the most likely occurrences, but also took into account the time you had available to train. The same is true with the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing (ASC) Program and computing. If you were to ask, what does it take to answer the majority of the stockpile questions today? then you get to where we are today. Today we supplement the expert judgment of the remaining weapons designers still in the complex with the current state of the art in our understanding as captured in our simulations of nuclear behavior. Now if you were to ask, what will it take to answer a detailed stockpile question in about 15 years time? then we predict that about 100 exaflop/s will be needed. At that point we are in the regime where all our experts with real test experience are long gone, and our approach to certifying the stockpile has significantly matured to help fill that void. In 15 years it is an entirely different game than today. Beyond that, requirements are much more uncertain and will depend upon the state of our capabilities and the type of questions being asked at that time.

As an aside, think about our nuclear weapons that are sitting around for decades. We know from reliability theory that even if every component of a nuclear weapon did not age, when you put the system together, they still have aging issues. What does that mean? When you think of ourselves, even if our cells never aged, we would still not live forever. That is because we face risks from accidents to disease, and our lifetime is often defined by events that are independent of our age. With weapons it is the same. Regardless of how pristine we keep parts of the systems, they still have a lifetime. For us, we use preventive methods to find early signs of medical issues so we can seek treatment early enough to have a positive outcome. For weapons, we conduct continuous analysis through supercomputer simulation and experiment to try and understand the issues we find. In either case, constantly improving methods are critical—giving more advance time, providing more mitigation approaches, and finding solutions/cures—for what we detect. Supercomputers are the only tools that can be used to reconcile problems we find through our virtual approach.

In 1995 your goal was 100 teraflop/s and now it sounds like it is 100 exaflop/s. What has changed?

For one, the world has changed. What we worry about today and the set of challenges we face has evolved significantly. The goal in 1995 was to establish the foundations of a Stockpile Stewardship Program, which was dependent upon simulation without the ability to validate via underground nuclear tests. It was anticipated that this would be a 10 year journey that demanded the utilization of thousands of processors working in unison, not just in performing calculations, but moving, storing, and visualizing the data produced. Your own personal computers have shown a similar scale of improvement over that period. In addition the goal at the time was to move to simulation software that allowed the resolution in three dimensions. The goals were all successfully accomplished and have enabled us to do groundbreaking calculations, which have provided us insights into the greater complexity of the problems we are dealing with, the discovery aspect. Based upon this increased understanding, we have defined the goals of the next phase of the Program, which is to drive to an enhanced predictive capability. This requires further significant improvements in the scientific basis of the simulation tools and at least a further order of magnitude improvement in resolution.

Do nuclear weapons continue to drive supercomputing almost two decades after the United States stopped testing?

Whatever one thinks about nuclear weapons, there are three things we must be prepared to address: (1) what we must do with our weapons today, to keep them safe, secure, and reliable; (2) how we understand what others are doing; and (3) how we anticipate the future. Weapons undergo continuous screening, as we look for issues and irregularities that might help the custodians seek advanced treatment. Everything from moving them around to taking them apart requires complex simulations to ensure safe handling. We need to be right because these activities directly impact the people who work with them. We must also be prepared to analyze what other state and non-state actors are doing. Here again, accurate simulations of scenarios is important because it can impact policy options. Finally, to be truly responsive as a nation to nuclear security issues that emerge, we need to invest in the research to ensure we apply our best scientific talent to these problems so that in the event we receive unpleasant news, we can take appropriate action based on credible technical analysis.

Over the past decade, you have learned to test nuclear weapons virtually. Why not just run these simulations from now on?

I wish that were the case, because then we could focus more effort on other pressing national issues. Although the integrated and specialty codes we have developed and applied up to this point are the best we can make them, created in a partnership between our code developers and the nuclear design community, there are still phenomenologically determined parameters, calibrated to the underground test base. Today that is okay because we still maintain experts who understand the nuances of the phenomenology as it pertains to how we tested and diagnosed nuclear explosions. But not too far from now, that will no longer be the case, and all our expert phenomenologists will be gone. As we move into the future and face slightly altered configurations and aged material, we must be able to predict behavior outside the as-tested design space. When we were testing, we designed and shot devices within the space we could simulate with our one and two-dimensional codes. In those days we controlled the design space. Today, Mother Nature controls the design space and we must respond by developing the capabilities and the tools to meet her challenges. This clearly means deeper science, greater understanding, and three dimensions. Aging breaks symmetries and drives material behavior; today's design and analysis communities must be prepared to answer questions we cannot bound.

One of your stated goals is "predictive capability." What does that mean?

The word "predictive" may be a bit strong. It should be understood by contrast with how we used simulation in the past, the so-called baseline models. The baseline models are a legacy of our historical approach to nuclear weapons design. At that time, weapons were really validated through testing, and simulation was heavily calibrated to the measured outputs so that it could be used as an interpolative tool for the next test. Many underlying scientific issues, that are non-equilibrium, complex dynamical processes, were simply absorbed into the calibration. And why not? A detailed understanding did not matter, and furthermore, the tests provided integral data that could not shed light on the microscopic physics, materials science, and chemistry behaviors. Today, our nuclear test history is becoming increasingly irrelevant to the stockpile as we swap out aging parts and update materials as needed to maintain them. In order to understand whether a new foam or component will perform as expected under nuclear conditions—that is, stellar temperatures and pressures—we need new scientific rigor in the computer codes. That will push us well beyond our baseline approach which to date has been the workhorse of design and analysis. A predictive capability enables accurate simulations of device behavior outside the parameter space spanned by the underground test data. Such a capability enables us to simulate behavior, to predict responses and performance beyond the range of test data, the last of which were collected 17 years ago. It is the only place we can go from here, and is a driver for exascale computing.

Supercomputers are expensive. How do you justify the cost?

For those of us in the business of national security, it's really not hard. In the 1950s, people were kept awake at night worrying about nuclear threats. Here we are almost 50 years later—most people have stopped worrying. How much is that worth? However, from a purely business perspective, the leading edge supercomputers are typically in the $100 million to $200 million range. Taking the support costs and the power bills, you need to factor a similar amount for its lifecycle of five years or so. Part of the calculation has to include how you plan to use such a system. It might be that if you have many users, it is cheaper to get a collection of smaller systems. So knowing what it is for and who will use it is important. For us, I think about the forthcoming critical national security decisions and issues in the nuclear enterprise. These might be high-leverage decisions which are informed through the simulations. Doing simulations does not guarantee you will save money, but it will provide a sounder basis for the decisions that have to be made. In the past we have had a number of hundred million to billion dollar level decisions we have been able to avoid by informing ourselves this way. So the benefits in terms of cost-avoidance are clear. There are far more subtle benefits as well. So again, how much is your security worth?

Can you save money by buying the last-generation supercomputer?

Not really, not for the job we have, and not for the commitment we have made. I do recognize that last-generation supercomputers are noticeably cheaper than leading-edge machines. And ASC does take advantage of that and buys some computers where the established technology can deliver workhorse machines where very high availability is a key requirement, the so-called capacity computing requirement. However, my science programs need high-end production platforms, capability computers, where the largest weapons calculations can be performed. These machines tend to be the current-generation technology and a single calculation may utilize all, or a significant fraction, of a computer, sometimes for timescales of a year or so. Finally since both supercomputer technology and nuclear weapons simulation requirements are growing and, I believe, will continue to do so, ASC also invests in advanced architectures, like Blue Gene, Roadrunner, and Sequoia. These machines assist the program in being ready to rapidly deploy and exploit the next generations of supercomputing technology.

What is so unique about the weapons science community?

I would say the focus of a powerful mission and the quality and breadth of the science that is being done to support it. As an academic, I only became aware of the scientific scale of this problem when I entered government. I have been enormously impressed by the level of scientific excellence and the level of creativity in this community. What is remarkable is how it takes a full complement of laboratories' infrastructure to enable such a coherent and long-term mission. You may think of the nuclear weapons program as an applied physics and engineering enterprise, using available tools to answer well-defined questions. That's not the case at all. A much better picture is to understand this community as a body of talented scientists applying the results of fundamental explorations to national priorities, many of which happen to be time urgent. The work done by this community spans a broad spectrum of scientific disciplines—from atomic and nuclear physics to the high-speed flow of liquefied metals, to applied mathematics, algorithms, computer architects, and code developers. The insights that the basic research provides lead to a deeper understanding of the various steps along the path from detonation through implosion to final yield. This is a rewarding grand challenge to confront, as well as an important service to the nation. Much of the knowledge cross-fertilizes national security issues well beyond the scope of our stockpile. I find that the tools we develop find important and time-urgent applications, from the satellite shoot-down that the President authorized in February 2008 to the analysis of the failure modes of the Space Shuttle Columbia that led to the accident in 2003. Without computing and these experts, none of this would be possible.

For NNSA, our mission responsibility requires the largest and the most powerful computers to develop and run the simulation tools that are used to inform decisions in national nuclear security.

With national support for nuclear weapons waning, how will you retain and recruit the next generation of experts?

About halfway between New York City and Boston, along Interstate 95, there is the small town of Mystic, Connecticut. It seeks to recreate the infrastructure that was needed to maintain ships and shipbuilding in the 19th century. There, the old skills are preserved. Today in the nuclear weapons program, we work with technologies that are decades old, looking at materials that are no longer manufactured, electronics from the 1970s and 1980s, and a host of dated issues. Imagine the growing contrast we will have in another decade. I worry that if we don't update the mission portfolio, we will be building our own Mystic seaport, an enclave of researchers working disconnected from the day-to-day national security issues, supporting a historical technology. It will not be a lure for our brightest scientists and engineers. The only way out that I see is to update the portfolio of research of our labs that has currency in today's national security landscape. Resetting the problem set we defined at the end of the Cold War to one that is relevant in today's world is critical to remaining vital in our core nuclear mission.

With the number of nuclear weapons declining, will ASC computer needs also drop?

In my opinion, I believe a strong case can be made for an inverse correlation. As the number of nuclear weapons decreases, our focus is on our confidence in the remaining, smaller number of both targeted and untargeted weapons. We have to be able to assure policy makers that they can rely on what they have determined they will keep in the stockpile. Our scientific understanding must be better, our simulation capability up to the task of ensuring the credibility of our information, and, in the end, our uncertainties in performance smaller. Short of replacement or significant rebuilding, the margins are what the margins are. All we can do is try our hardest to calculate the error bars in our prediction of failure modes and provide that information to our customer, the Department of Defense. That means scientifically-based simulations and analyses, and those must rely on large-scale, full-system calculations on the most capable computers. I will point out, as an aside, that even if we had no stockpile, we must be able to analyze devices that state and non-state actors have the potential to field and—I hope never—attribute an exploded device anywhere in the world to its source.

What do you view as the role of the Department of Energy (DOE) in supercomputing, and NNSA in particular?

Although historically DOE, and the weapons program in particular, has been the driver for the big machines, our primary role is one of a consumer. Broad brush, I see the department as having two essential components: one of "closed" national security science, and one of a commitment to pushing the forefront of "open" science. Both are critical to this nation and complementary. Supercomputers are the tools that enable these DOE science missions, and their complexity and long-term planning have been a core competence of the Department and its predecessors throughout the history of supercomputing and its roots in the nuclear weapons program. I believe that DOE is the federal centerpiece for leadership in supercomputing, and that these tools should be broadly applied to the growing demands in both the national security arena and open scientific discovery.

For NNSA, our mission responsibility requires the largest and the most powerful computers to develop and run the simulation tools that are used to inform decisions in national nuclear security. It's not just about the weapons; we worry about issues of proliferation, detection, counter-terrorism, and a collection of adjacent issues where the computers classified at increasing levels of security can be turned to pressing problems. The program is faced with what are probably the most complex nonlinear problems that have ever been represented on a computer. To drive the uncertainties down to levels that meet the needs of the Department of Defense and to assure the President with confidence that the testing moratorium can continue together drive our demand for the supercomputers with systems and interfaces that make them useable by the community responsible for national nuclear security. This is our role and we need to be sure we have the tools to be able to do it into the future.

Are national laboratories the right homes for supercomputing?

For us, it's the only home. There are two main reasons. First, both physical and cyber security are things we all worry about. Today we can get LoJack on our laptops to track them if they are stolen, and worms, viruses, and pervasive threats are somewhat mitigated by software tools. But despite all our care, systems are still hacked and information lost. When you think of the work we do that keeps our big computer networks humming day and night, we really cannot afford any loss of information. It is not just guards and fences; it is top cyber teams and carefully constructed networks.

Second, dependably secure and available supercomputing to do simulations of nuclear weapons is also the backbone of the nuclear weapons certification and analysis process. The NNSA national security labs (Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Sandia National Laboratories) have been drivers for the supercomputing industry since the dawn of the nuclear age. At the laboratories reside the multidisciplinary science and technology that lead to innovation and enable us to address a broad suite of problems of a scale that do require the most powerful computers. This potent combination of talent and mission-need is the reason that the laboratories have been and will continue to be the right home for the next generation of national security supercomputing.

How are you responding to the Secretary of Energy's 2008 vision for the broader national security role of NNSA?

NNSA is marching forward to put flesh on the bones of this vision. We are working with teams at the NNSA laboratories to develop a portfolio of potential areas of applied research where this agency can contribute to challenges that confront our national security. An example of immediate interest is the ability to do credible forensics on the debris from an exploded nuclear weapon anywhere in the world with the goal of attributing the source of the material and the provenance of the design. This would provide policy makers with the information to enable an informed choice of actions to be taken to respond. These teams are working on the identification of other agencies with responsibility for aspects of national security with whom we might partner. We are looking to the next budget cycle of 2011–2015 as the next opportunity to create a wedge to support this expanded mission. The expertise that has been developed at the weapons laboratories in national security science and attendant classified work supports this vision of a broader science and technology mission for NNSA. In 2008, the Secretary of Energy signed off on this expanded mission and NNSA is working with other agencies to begin implementation.

Being a national security program, do you feel that you have any role in supporting U.S. competitiveness?

A role, sure; a focus, no. I have been committed to understanding how money we spend can simultaneously help the United States gain the edge internationally. Working with a number of outside organizations, first and foremost the Council on Competitiveness, I try to keep an eye on both near-term opportunities and longer-term possibilities. One of the problems I see is that supporting U.S. competitiveness is not simple. With respect to the edge supercomputing can give, no federal agency has such a mandate to support this area. Further, in my experience, it is typically when companies have had their backs against the wall that they are willing to change their research and development (R&D) from more costly empiricism to predictive, virtual-based simulations. When this has happened, there have been notable benefits. We do push the envelope of simulation and these modern techniques are available for all to use. Cross-fertilization is important whenever opportunity arises. Similarly, we drive computing architectures and numerical libraries for those architectures, which are available for all to use. I view this as a rising tide that lifts all boats and, in this way, does improve the U.S. position in the global economy. This nation is known for its innovation and the ASC program contributes mightily to the general knowledge of computational science. A specific effort that had a very good outcome is the partnership between Sandia National Laboratories and The Goodyear Tire & Rubber Company to develop simulation capability to design the next generation of tires. As a result of the simulations, Goodyear was able to reduce their time-to-market for consumer tire products from three years to one year, at a significantly reduced cost, saving of the order of $100 million per year.

Is it a DOE responsibility to ensure a vital U.S. supercomputing industry?

If history is any guide, probably so. In 1993, at the beginning of the ASCI program (the initiative that preceded the ASC program), an identified goal of the program was to ensure the viability of U.S. supercomputing, since at that time the nuclear weapons program had backed away from acquiring the latest supercomputers and the high-end component of the industry was failing. ASCI turned around that trend and kick-started a rejuvenation in supercomputing. Over the past 15 years we have invested in and worked hand-in-hand with the supercomputing industry to ensure their continuing health. This has taken the industry from a few billion floating point operations per second in 1992 to computers that are a million times larger than that today. These investments have been essential for the success of the ASC program and have been of enormous benefit to U.S. industry as well.

A few specifics are worth discussing further. It is well known that ASC has a long history of leveraging long-term partnerships with industry to accelerate computing technology development. What is not well understood about this process is that, in the partnerships, we have had the opportunity to collaborate with our high-performance computing (HPC) partners to develop, scale, and deliver systems that are more suited to our applications. Examples of this are the ASC PathForward (now FastForward) program element and the Advanced Architecture development efforts that gave rise to Roadrunner and Blue Gene (BG). (The development of BG/P and Sequoia were undertaken in partnership with the Office of Science.) In PathForward, among many other activities, we collaborated with industry to develop and scale five generations of interconnect technology (HiPPI, Colony, Federation, Elan4, and InfiniBand). Over several recent TOP500 lists, approximately 40% of the systems in the TOP500 list directly benefited from ASC PathForward and Advanced Architecture investments.

We are committed to continuing this historic engagement of industry in these areas. In addition, we are developing new and innovative ways of partnering with industry to develop, and more importantly, scale commodity-off-the-shelf technologies for HPC and the broader community's use. The Hyperion partnership with 10 leading HPC industrial partners at Livermore recently won a Federal Computer Week 100 award for this innovative approach to joint development. We also have a keen attention to lowering the Total-Cost-of-Ownership (TCO) of delivered systems. We pioneered this approach with the Tri-Laboratory Linux Capacity Cluster (TLCC) procurement that delivered over 28 Scalable Units (SU) to five sites in 15 different clusters. Two fundamental enablers of this strategy are the concept of building multiple clusters of widely varying sizes out of a highly-replicated SU, and common hardware and system software stack across multiple sites. With this strategy, we lowered our TCO by more than 50% over previous deployments of large Linux clusters. We will continue innovating in the Linux cluster arena for capacity systems in the future.

TCO improvements have been obtained from the huge improvements of flop/s per Watt focus of our advanced architecture approaches. For example, with Sequoia at 6.0 MW for 20 petaflop/s, this represents a 200 times improvement over the landmark Purple supercomputer and over four times improvement over other systems of similar capability based on COTS for delivery in 2011–2012. With the cost of power these days, both in dollars, but also in terms of carbon footprint, these are huge savings.

Will the Top500 list, which ranks the world's fastest computers, ever flatten out?

Ever is a long time, but in the near term, I anticipate a continued positive slope. Ever since the invention of computers the supercomputer industry has managed to maintain an impressive growth rate of the most powerful computer in the world. Initially this growth realized a three order of magnitude increase in computational speed every 15–16 years, but for the last 15 years or so, coinciding with the introduction of the ASC Program, the factor 1,000 growth has been every 10–11 years. During this time supercomputer manufacturers have managed to overcome a wide variety of anticipated barriers to the continual increase. Therefore, although there are certainly some very significant challenges to be overcome to increase computational power a further three to six orders of magnitude, the expectation is that these challenges will be overcome. However the more recent rapid growth may not be maintained without continued investment at a scale commensurate with the challenge.

How and when will you get to an exaflop?

We need it, so we better plan for it and continue down this path with as much energy and certainty as we can muster. NNSA already has envisioned several key simulation requirements that demand exascale computing. It is already anticipated that the delivery of supercomputing beyond a petaflop will require a complex computing environment exploiting large numbers of computing nodes, each with large number of cores, and possibly also including specialized computing processors like those used in Roadrunner and video game consoles. It is expected that, on current supercomputing growth rates, an exaflop computer will not be available until the 2018 timeframe, and that efficiently exploiting these computers will be extremely challenging. ASC is, therefore, currently investing in early delivery vehicles of some of these technologies, for example, Roadrunner and Sequoia, to allow the Program the opportunity to start to address the likely challenges, in both the computer environments and the application of both system and scientific software.

NNSA is making some significant investments in petascale platforms in order to pave the way towards exascale. Today's and future petascale systems all have several common attributes: unprecedented levels of parallelism; innovative architectures to dramatically improve the flop/s per Watt; and innovative approaches to reliability. To a lesser extent these systems are also addressing the problems associated with the fact that memory bandwidth, and recently also capacity, are not keeping up with processor improvements. Two major trends in the industry are being leveraged in these systems: homogeneous systems with low power, system-on-a-chip (SOC) designs, leveraging the embedded and gadget markets; and heterogeneous systems with attached coprocessors, leveraging the gaming and information processing markets.

Fielding a workable state-of-the-art supercomputer system requires the cooperative work of a community of computer scientists, computational scientists, engineers, and technologists.

This delicate balance happens in advancing a petascale strategy designed around rapidly advancing the state-of-the-art while also delivering production systems for the nuclear security mission. An advanced architecture system, Roadrunner, deployed in 2008 with 1.3 petaflop/s peak, is exploring the applicability of heterogeneous architectures utilizing accelerators to the ASC workload. The Sequoia system in late 2011 with 20 petaflop/s peak is the third-generation Blue Gene system based on a homogenous SOC architecture and is now evolving that advanced architecture into the next generation production system. Prior to delivering Sequoia, the 2010 Zia system at 1 petaflop/s peak will attack another fraction of the workload as described earlier. Both the Roadrunner and Sequoia systems dramatically improve the flop/s per Watt over traditional architectures.

The primary challenge that we will face with Roadrunner and Sequoia is extracting much higher levels of parallelism out of millions of parallel lines of existing codes. The Roadrunner approach is in the research phase and is utilizing the optimized library approach. For Sequoia, we have extended the "Livermore model" programming model to include multiple styles of shared memory parallelism within the existing high-level message passing (MPI) parallelism. By carefully allowing the nesting of

OpenMP (loop level parallelism), Pthreads (threads and lock style parallelism), and innovative approaches such as transactional memory and speculative execution in each MPI task, we can offer multiple packages the opportunity to utilize the most applicable style of parallelism, while minimizing the operating system and application programming complexity. In addition, to dramatically improve overall system reliability, Sequoia takes a hierarchical hardware and software architecture where both the hardware and software on the highly-replicated components—the compute nodes—are made extremely simple, and more complicated system functions for Input/Output (I/O) and external networking are forwarded to much fewer I/O and Login Nodes. This hierarchal system model has delivered highly-reliable systems in the first two generations of Blue Gene (BG/L and BG/P).

Can NNSA do this alone?

If we have no other choice, then yes. Fielding a workable state-of-the-art supercomputer system requires the cooperative work of a community of computer scientists, computational scientists, engineers, and technologists. The systems are of a complexity and size that no single group of people can do all of the work required to effectively field a single machine, let alone do the research to ensure that we can use the hardware and software advances on future architectures. From advanced architectures to storage systems to graphics to debuggers, the ASC program (and many other users of HPC) have benefited from the contributions by academia and vendor R&D. We are members of the computing community and certainly make significant technical contributions, but we also participate in the larger community in other ways. For example, ASC funds a group of universities, as an outreach to the world of unclassified research, in the Predictive Science Academic Alliance (PSAAP) program to establish validated, large-scale, multidisciplinary, simulation-based predictive science as a major academic and applied research program. These centers are unique in their approach to solving a single large-scale multidisciplinary problem on a university campus. Our centers at Caltech, Stanford, University of Utah, University of Chicago, and University of Illinois–Urbana-Champaign were recognized earlier this year for a decade of leadership in helping define the field of computational science, as well as for contributions to U.S. competitiveness and the security of the country. These centers highlight how long-term strategic partnerships can be made to effectively leverage the country's diverse talent pool.

How can all of DOE work together to achieve these common goals?

We are doing it already. We and the Office of Science have many common interests in computing and other scientific areas. The Office of Science laboratories are key players in developing tools to make high-performance computing systems more usable and efficient and we are developing ways to formalize our interactions. We have formed two institutes, the Institute for Advanced Architectures with Sandia and Oak Ridge, and the BALL (Berkeley, Argonne, Lawrence Livermore) institute to capitalize on the expertise across the complex in advanced systems and computational sciences. We also are engaged in the Scientific Discovery through Advanced Computing (SciDAC) program to capitalize on the Office of Science investments in new ideas advanced by academia and other laboratories.

In terms of philosophies, I believe that across DOE we have recognized that success is based on intelligent leveraging of the mammoth investments of others in computer technology. For instance: foundries and base silicon technology investments represent about a $10 billion/year investment; Microprocessor design, compilers, operating systems is perhaps another $5 billion; what could be called HPC system design, interconnect development might be $1 billion/year; and ASC or Leadership Class systems within DOE might be together $200 million/year. So, we are at the tip of the pyramid, and we better be really, really smart in how we leverage the base.

What we together need to do is to make sure that the base of the pyramid (the Commercial Off-the-Shelf world) keeps pace with the vertically integrated systems we need to procure. We should focus on issues like hierarchical system hardware and software and on overall system architectural issues. DOE can lead this, and we should lead it.

An example, mentioned earlier, of what we are doing today is the Blue Gene R&D contract with IBM, cost shared 50/50 by the Office of Science and us. This contract has made possible both the BG/P and BG/Q technologies and followed the model set by the original BG/L R&D contract between NNSA and IBM. As a result of this effort, the Sequoia system is coming together for NNSA, a large BG/P system is now sited at Argonne, and a smaller system at an NNSA lab. Others have leveraged our efforts, with a 1 petaflop/s BG/P system to be sited in Germany.

As a final example, the two institutes that are being established across NNSA and Office of Science laboratories are a big step. What needs to be added to the discussion, of course, is that the issue is funding. In order to get the attention of the vendors, substantial investments (like the BG investment) are necessary. Even with both the Office of Science and NNSA's Defense Programs working together and sharing the cost, funding two institutes plus the industry investment, is expensive. However, without such enlightened investment, we will be at the mercy of the base of the pyramid, and the base does not worry much about the tip, where we live.

What keeps you up at night?

First inaction, then wrong actions. It takes years to change direction in this simulation enterprise, finding or training the right expertise, developing the integrated codes with the relevant physics, and validating and verifying their correctness. In other words, if the needs of national security change, there will be an inherent lead time to come up to speed and ready the capabilities to be able to respond. So, I often ask myself the question: are we solving the right set of problems today; are our investments in the science, technology, and engineering aligned with what we think the world will be like in the future? How do those of us at the Federal level, who help steer the development of the essential national security capabilities for the country, anticipate correctly what the future threats will be, and how do we make the right investments in hardware, codes, and people that will be crucial five to ten years from now?

Thank you for taking the time to answer our questions. ●