LA-UR-14-20080
Approved for public release; distribution is unlimited.

Title: Without Testing: Stockpile Stewardship in the Second Nuclear Age

Author(s): Martz, Joseph C.

Intended for: University Lecture, Stanford

Issued: 2014-01-07

Without Testing: Stockpile Stewardship in the Second Nuclear Age
—Joseph C. Martz

Joseph C. Martz, technical staff member at Los Alamos since the 1980s, is an expert on plutonium aging and co-developer of ARIES, a system for safely recovering the plutonium from dismantled excess nuclear weapons. He also led the Enhanced Surveillance program for nuclear weapons within the Stockpile Stewardship Program.

Stockpile stewardship is a topic dear to my heart. I've been fascinated by it, and I've lived it—mostly on the technical side but also on the policy side from 2009 to 2010 at Stanford University as a visiting scholar and the inaugural William J. Perry Fellow. At Stanford I worked with Perry, former secretary of defense, and Sig Hecker, former Los Alamos Lab director (1986–1997), looking at nuclear deterrence, nuclear policy, and stockpile stewardship and at where all this was headed.

The Nuclear World Changes

In my career, the most consequential period with respect to nuclear weapons were the years from 1989 to 1992. Three very important things happened during those years, and they led to profound changes in U.S. nuclear policy. First, we had the fall of the Soviet Union, presaged by the fall of the Berlin Wall in 1989. The dissolution of the USSR came on December 25, 1991, and that changed everything. The Cold War and its nuclear arms race were over, making an anachronism of MAD [Mutual Assured Destruction], the policy whereby, to deter nuclear war, the United States and the Soviet Union each deployed enough nuclear weapons to ensure the complete destruction of the other.

Second, in 1989, the government halted work at the Rocky Flats Plant, outside of Denver, Colorado, where plutonium pits [the nuclear triggers for thermonuclear weapons] were produced. That turned out to be a seminal moment in the history of the nuclear weapons complex because, frankly, with very little forethought, it ended our ability to produce new weapons and it effectively shut down the entire nuclear weapons production complex! Over the next 10 years, more than 50 percent of the historic nuclear weapons complex was shuttered forever.

Third, the Soviet Union had proposed a moratorium on nuclear testing and conducted its last test on October 24, 1990. In 1992 President George H. W. Bush endorsed a U.S. testing moratorium. On September 23, 1992, "Divider" was the last nuclear test conducted by the United States. The era of nuclear testing was over.

Any one of those changes would have radically altered how the Lab carried out its national security mission, but all three events together put the Lab in unprecedented territory: instead of *designing and engineering* weapons for the nuclear stockpile, it would now *maintain* the stockpile. But the cessation of nuclear testing was the loss of the most important tool the weapons designers had used for 50 years to develop nuclear weapons and ensure that the stockpile was safe, secure, and reliable.

Also lost was the means that, along with nuclear testing, had developed and maintained the skills of weapons designers: the continued design and production of new, upgraded nuclear weapons. Maintaining the designers' skills was vital because although the Cold War was won, shifting global politics presaged new national security needs—and perhaps new nuclear weapons requirements to meet them (such as enhanced security measures in a post-9/11 world). So, not only had we closed the factories and put a moratorium on testing, but we had also agreed not to develop new weapons. Together, these required managing an aging stockpile, while remaining agile in the face of changing national security needs.

Inventing "Science-Based" Stewardship

After the collapse of the Soviet Union, President Bill Clinton commissioned the first Nuclear Posture Review to examine the role of nuclear weapons in a post-Soviet world. This review (and every review since) reaffirmed the continued need for U.S. nuclear deterrence, while also recognizing the changing conditions and constraints in the global security environment. For itself and its allies, the United States would continue to maintain its nuclear stockpile, and nuclear deterrence would remain a central element of our supreme national security posture. But that presented the nuclear weapons laboratories with a huge challenge. How could the nuclear weapons labs ensure that nuclear weapons remained safe, secure, and reliable in the absence of nuclear testing?

The question was particularly important because the weapons were going to enter configurations that we had no experience with; that is, because we weren't continuing production, the weapons we had would, by default, age beyond their design life. Could we and our allies rely on these complicated weapons in spite of their aging? We would have to understand how age affected the weapons' performance, safety, and security and do that without any further nuclear testing.

This also meant finding new ways to train the next generation of designers without the live tests the first generation had used.

Rethinking Mission "How To's"

Clearly, we had to rethink the entire problem of meeting our national security mission. Leading that process was Assistant Secretary for Defense Programs Vic Reis in the Department of Energy [DOE], assisted by Los Alamos National Laboratory Director Hecker.

Hecker, with Reis's guidance, convened technical experts from across the DOE weapons complex, and what the experts came up with was the realization that maintaining the stockpile would require an approach that was the complete inverse of the one used during testing. I'll explain what that means.

Nuclear testing was a wonderful tool. It was also the world's biggest shortcut. It meant that we didn't have to understand all the details of a nuclear weapon and how it functioned. During the nuclear testing era, we got away with knowing just enough about how things worked and how materials behaved to configure a device and make a pretty good guess as to how it would perform. We then took it to Nevada, buried it in the ground, and detonated it to see if it worked. It usually did, but sometimes it didn't, and we didn't always understand why. Basically, we solved the problem of building safety, security, and reliability into a weapon from the top down: if the full device worked, its components must be working. So, we froze the design at this point and did our best to build systems that exactly replicated what we had tested.

In the post-testing era, we realized that without the top-down approach, we would have to piece together how nuclear weapons function from the bottom up—that is, gather all the basic science pieces underlying the behavior of each of the weapons' different materials and physical processes and use that information to calculate how the complete nuclear weapon would function (see "Nuclear Weapons 101—How They Work").

We quickly realized that the best way to do this was to represent all this basic science as a series of mathematical models and then to integrate all those models, along with copious amounts of data on physical properties, into a huge computer calculation that would accurately predict a weapon's performance.

To be sure the calculation was accurate, we would validate it by comparing its results with the results of past nuclear tests [over 1,000 tests for the United States], as well as with results from newly conducted "integrated" experiments. Integrated experiments reproduce in the real world some portion of what the computer calculation predicts—for example, how some configuration of materials in a warhead behaves when hit by shock waves during detonation. We adopted this process to put real-world checks and balances on our virtual-world calculations and predictions.

This was the bottom-up approach. It would enable weapons designers to make technically sound judgments about a weapon's performance without any new nuclear tests. In 1994, shortly after its conception, we named this approach "science-based stockpile stewardship," now the Stockpile Stewardship Program. A colleague of mine, Jas Mercer-Smith, had a good line about this. He used to say stockpile stewardship is all about doing nuclear testing in a computer. It's just damn hard on the computer!

Finding the Fundamental Science Pieces

We realized in the early 1990s that not only were computer calculations of weapon performance going to take a level of computing power that didn't exist at the time but that the basic science pieces for building those computer calculations were also missing.

One of the missing science pieces was an understanding of many of the properties of the materials used in weapons—for example, the strength and compressibility of many of the materials within the "physics package"— and how those properties changed under extreme pressures, temperatures, forces, and accelerations—especially after the materials had aged.

In the nuclear testing era, we'd never thoroughly characterized the properties of the materials that went into the weapons—we hadn't needed to because the weapons were tested and regularly replaced. This lack of characterization was especially evident for the most important material in the weapon: plutonium.

For example, we didn't understand the details of how the hollow plutonium sphere [the "pit" inside the primary of a nuclear weapon] gets compressed when its surface is hit from all around by a strong high-explosive shock wave. The pressure from the shock wave causes the plutonium not only to implode (move inward) but also to get denser because the atoms in the plutonium are forced closer together [i.e., compressed]. But how much pressure causes how much compression, or increase in density?

We needed to put that quantitative information into our computer codes so they could predict accurately exactly when, during implosion, the subcritical pit would reach a supercritical configuration needed to sustain a fission chain reaction. But we didn't have accurate data to give us that quantitative information. Since we didn't know this, we certainly couldn't describe how decades of aging change plutonium's ability to compress either. In fact, we didn't even know whether its compressibility, strength, and metallurgical stability actually *would* be affected by aging. So one of the first things we had to do in stewardship was build the tools and facilities needed for measuring these types of material properties in plutonium and in other key weapon materials.

Figure caption: *The gas gun (artillery designed for use in a laboratory) contained in this glovebox at the Laboratory's Plutonium Facility fires projectiles at speeds ranging from about 200 to 4,000 miles per hour into plutonium samples. Researchers measure how much the plutonium gets compressed by the resulting shock waves. Photo by Paul Moniz (*Actinide Research Quarterly, *4th quarter 2001)*

In 1997 I moved from the group that was charged with examining pits and plutonium and asked to start a program to study aging in *all* the materials within the weapon. We called this Enhanced Surveillance. Initially, Enhanced Surveillance was a $7-million program at Los Alamos, but within a few years, it grew to five times that size. By 1999 Los Alamos had 40 science projects devoted to learning how the various materials would age and how that aging would affect a weapon's performance. Another 100 such projects were being conducted at other labs and sites—a clear sign of just how concerned we all were with the problems of aging stockpile materials. Some of the country's best chemists, engineers, and materials scientists became focused on aging nuclear weapons, and the success of their work formed a key basis for the Stockpile Stewardship Program.

Another missing science piece was a detailed quantitative understanding of the physical processes that go on during a nuclear detonation. For example, the plutonium pit's implosion, which must crush this hollow metal sphere into a supercritical configuration, requires a precise sequence of processes that must work together perfectly: first the detonators fire, then the high explosives ignite, and the shock waves from the ignited explosives collapse the plutonium inward with the right timing and symmetry for the pit to go supercritical. The shock waves break through the plutonium into the gas at the pit's center, where fusion takes place, providing extra neutrons that boost the fission chain reaction in the plutonium shell. As it turns out, knowledge about much of the physics of these processes was incomplete, especially for the later processes associated with the boost process.

Through experience and the nuclear tests of the testing era, we had gotten this sequence to work. And the processes had been partially measured and modeled, but never to the degree that would make us confident that a bottom-up calculation would be predictive—that is, would provide an accurate picture of exactly how all the processes fit together into a working whole.

For example, when the shock wave breaks into the center of the pit, bits of plutonium are ejected into the hot deuterium-tritium (D-T) gas at the center. The D-T gas must become sufficiently hot to begin undergoing fusion, just like the hydrogen in the sun. But the effect of plutonium "ejecta" mixing with the D-T gas is just like the effect of an ice cube being dropped into your coffee: things cool down, which can dramatically affect the boost process. But by how much? How big are the bits of plutonium ejecta, and how many are there? How fast do these bits of plutonium mix with the D-T gas, and how does the rate of mixing affect the boost process? Will ejecta from an aged pit be different from those produced by a fresh pit and have a different effect on the boosting process? These were open questions.

To do the basic science experiments needed to improve our understanding of these processes and convert that understanding into mathematical models for high-resolution 3D computer calculations, the stewardship program provided for a number of new research programs and for new facilities to be built at Los Alamos, Lawrence Livermore, and Sandia National Laboratories. By the year 2000, DOE had established nearly a dozen "campaigns" to address these science issues. These campaigns have made tremendous progress in filling in the gaps in the myriad physics and materials issues of relevance to weapon assessment, and they continue to make advancements to this day.

Accelerated Strategic Computing Initiative

One important campaign was about investing in powerful new scientific computing capabilities—advanced, fast supercomputers and new computer codes—to perform the huge calculations that would include all the new fundamental science and data from new experiments. To assess how a weapon worked, we would take all the new data on materials properties, combine those data with the physics we learned, wrap all that into new weapons computer codes, and have the codes step through a detonation piece by piece. A code would mock up the nuclear weapon virtually, first in two dimensions and ultimately in three, using millions of pixels to model the exact shapes of weapon components. And the code would track the changes in each pixel for many tiny increments of time to accurately simulate the detonation.

When we started this approach in the mid-1990s, a full-system bottom-up calculation—from the detonation of high explosives to the final energy release of the entire warhead—would have had to run for years to reach completion at our newly desired levels of detail. We knew if we could accelerate the rate at which computer power was increasing, in 10 years we could reduce that running time from years to months. It was important to get the required computing power up and running as soon as possible because most of the weapons designers with testing experience would be retiring over the next decade or two. The designers' real-life testing experiences would be critical to evaluating the accuracy of the computer models we hoped to generate.

We were running to beat the clock, so DOE created ASCI, the Accelerated Strategic Computing Initiative. Under that initiative, DOE and the computer industry began producing computational platforms at an accelerated pace, and those computers started to break records in terms of their capabilities. For example, we reached one thousand trillion calculations per second in 2008 with the Roadrunner supercomputer, a milestone that was widely considered impossible in the 1990s. And all this was done to enable the massive calculations that were needed for modeling a nuclear weapon.

Integrated-Test Facilities

But it wasn't enough to have the fundamental data in the new codes and to run the new codes in a timely way on the new supercomputers. We also needed to validate that the predictions from the new codes were right, so we brought to bear the third major investment for stewardship, namely, facilities where weapons designers, old and new, could do integrated tests. The real-world results from those integrated tests would provide a check on what the codes predicted for the same phenomena.

The most common type of integrated test, both today and in the testing era, is the hydrodynamic test, or hydrotest, a nonnuclear test in which the replica of the first stage of a nuclear weapon, the primary, undergoes implosion and the implosion is imaged by x-rays. These implosion experiments are called *hydro*tests because, at the high pressures attained during implosion, the materials flow like liquids such as water. To keep the hydrotest nonnuclear, a surrogate metal is used in place of plutonium.

Hydrotests at DARHT

At Los Alamos the most important integrated test facility is DARHT [pronounced "dart"], the Dual-Axis Radiographic Hydrodynamic Test facility. At DARHT the hydrotest of a primary occurs inside a sealed, steel test vessel, and two powerful x-ray machines set at a 90-degree angle to each other take simultaneous x-rays of the implosion process, giving us two views at one instant. One of the machines takes a single image, and the other captures a four-image sequence, thereby making a kind of short "movie."

DARHT's images have unprecedented resolution, and we can compare them with our calculations to check whether the calculations simulated the implosion correctly. These capabilities make DARHT a unique experimental facility and critical to the stewardship program because having integrated experimental data for direct comparison is absolutely key to validating our codes and calculations. Lawrence Livermore also developed an important integrated testing tool at the same time—the National Ignition Facility, NIF. While DARHT concentrates on hydrodynamics and understanding the implosion stage, NIF studies later elements of weapon function that are also important to model and understand.

Nevada Nuclear Security Site

Because we don't use plutonium at DARHT, we needed to build a facility at the Nevada Test Site [now the Nevada Nuclear Security Site] where we could do integrated tests involving plutonium and high explosives. By both Executive and Congressional order, such experiments would have to be subcritical; that is, any experiment in which we are using explosives to dynamically compress plutonium must never produce a critical mass.

A shaft called U1a had been dug there, about 970 feet below the desert surface, for a nuclear test that never took place. We went down the shaft and expanded it into a very sophisticated plutonium-testing facility with advanced diagnostics. For example, we built a miniature version of DARHT nearly 1,000 feet underground to take x-ray pictures of plutonium during implosion. Recently we have added another diagnostic tool, the PDV (photon Doppler velocimetry), which uses dozens of laser beams reflecting off the *inside* of the pit's surface to continuously track the velocities of different sections during implosion. Quite literally, we built a state-of-the-art dynamic testing lab for nuclear materials in a mine! Our first subcritical experiment (subcrit) was Rebound, conducted on July 2, 1997, in which we compared new with aged plutonium. Since then we've conducted dozens of subcrits that have provided invaluable data for our understanding and assessment of weapon function.

Proton Radiography

DARHT's x-rays let us take up to four images of the implosion of a surrogate weapon primary. Because it uses strong x-rays, DARHT is very good at imaging dense materials like metals. Many of the materials in the weapon aren't dense; they're relatively lightweight, like explosives, foams, and cushions. When DARHT is tuned to look for the movement of heavy metals, it can't easily image the movement of shocks in things like the lighter high explosives. This problem has been known for many years, and some very clever Lab scientists figured out that protons—the nuclei of hydrogen atoms—would make an excellent probe to image these lighter materials.

Hence, the science of proton radiography, pRad, was born. The pRad facility, an outgrowth of the Los Alamos Neutron Scattering Center (LANSCE), uses protons to take more images, and at a higher contrast, of many of the materials in the physics package. Proton radiography is especially well suited to the study of shock-wave movement inside the explosives themselves. Very short pulses of protons, accelerated to over 80 percent the speed of light, can penetrate these materials and create a sequence of 10 or more 2D "movie" images of, say, a detonation travelling through high explosives at 17,000 miles per hour!

With pRad, we can measure the effects of aging on high explosives. We can also use it to reveal how ejecta form when a detonation breaks through a metal surface. In addition, pRad experiments are used to develop and test the physics models that go into the codes and to interrogate some versions of scaled-down integrated experiments.

Different Slants on Stewardship

Now the United States has all the parts in place to piece together an assessment of nuclear weapon safety and functionality. Three key advancements—new fundamental materials and physics experiments, advances in supercomputing, and next-generation integrated test facilities—have made it possible to successfully ensure a safe and reliable stockpile without the need to conduct any additional nuclear tests.

However, while the United States took the lead in using a science-based stewardship approach, other countries deployed different solutions for stewardship of their stockpiles. The Russians, for example, have maintained a production-based stewardship program. They didn't close their production factories. They now simply replace their aged weapons by manufacturing new ones. The Russians can still produce thousands of new warheads a year if they desire.

Weapon Autopsies

An important element of stewardship is a surveillance program to monitor the aging of weapons in the stockpile. Each year several weapons of each type are returned from the Navy and Air Force. Most of these are nondestructively examined and returned to the military. A small fraction of these weapons undergo destructive evaluation. In essence, we perform an autopsy on them. The weapon is disassembled into its components, and those components are returned to the production agency for evaluation. The pits are returned to Los Alamos, where we cut them open for detailed examination. Plutonium is extracted and subjected to a variety of tests to look for aging or birth defects [flaws created during original manufacture]. Measurements from the tests are compared with the manufacturing records for that specific unit, and we note changes that may have resulted from aging.

One issue we've had is the incompleteness of the initial manufacturing records. We didn't record all the data we'd like to have to fully assess aging in the weapon. A large part of our activities in Enhanced Surveillance involved reconstructing time-zero and birth data for many of these units so we could compare those data with our surveillance measurements.

During surveillance operations, if we found a deviation from specifications, we would report this as a "significant finding notification," or SFN. The designers would evaluate each SFN, and if they felt it required further assessment, they would elevate the notification to a "significant finding investigation," or SFI. From 1995 to 2005, between Los Alamos, Livermore, and Sandia National Laboratories, 156 SFIs were opened and investigated. Of these, 75 were determined to be "nonactionable," meaning that the investigation and assessment revealed no reason to expect an impact to safety, security, or performance.

The remaining 81 SFIs were deemed to be actionable. In this case, a component or material was changed, often as part of a scheduled refurbishment process, or a change was made to the certification of the weapon, usually through a limitation in the environmental storage, deployment, or military requirements. Of the 81 actionable SFIs, 18 arose in the nuclear components for which Los Alamos and Livermore are responsible, including 7 SFIs related to pits.

So, we've done lots of surveillance on all of our weapon types, including so much surveillance on the W88 warhead and its pits that the Navy said we were taking too many warheads out of the stockpile. That's one reason why Los Alamos, in lieu of Rocky Flats, was asked to gear up a new pit-production capability in the late 1990s: to replenish the Navy's inventory.

This became the pit rebuild program, active from 1997 to 2010. In addition to resupplying the Navy with pits for the W88, there were two other important goals for pit rebuild. One was to capture manufacturing technologies previously used at Rocky Flats and to develop replacement technologies for those processes that couldn't be replicated at Los Alamos' Plutonium Facility, known as TA-55. The other was to demonstrate that we could use the tools of stewardship to certify the newly rebuilt pits, along with certain new production methods.

Pit rebuild wasn't the only activity to require new production methods. Other programs began to look at how certain weapon characteristics could be improved to reflect post–Cold War stockpile needs such as enhanced safety and security. One of these activities was the Reliable Replacement Warhead (RRW) study, which illustrated many of the challenges of stewardship. The story of RRW is one with which I'm intimately familiar.

Rebooting Aged Weapons

In order to appreciate the goals of the RRW study, it's important to understand the design goals and characteristics of the Cold War–era stockpile, the stockpile we still have today. Weapons designed during the Cold War placed a premium on military characteristics designed to deter a specific adversary, the Soviet Union. One of the design goals was to stretch limited plutonium inventories as far as possible in order to build the greatest number of weapons from the limited supply of this strategic material. We were in an arms race with the Soviet Union, and every gram of plutonium mattered. If you could reduce the amount of plutonium in a weapon, you could build a few more weapons for less money.

At the same time, we wanted to optimize the yield-to-weight ratio in warheads going onto missiles; we wanted the biggest bang for the least amount of plutonium and in the smallest and lightest warhead package. This allowed us to place multiple warheads on a single missile, expanding the target set and enhancing Cold War deterrence. This was especially true for warheads on strategic missiles, where weight was at an absolute premium. Optimizing yield to weight was everything in our designs, and the history of Cold War design at the national laboratories was one of exceptional success. We're very proud of the fact that we did, indeed, build very lightweight, compact, and powerful nuclear weapons. These weapons helped win the Cold War.

But the success of our designs didn't come for free. The price for optimizing the yield-to-weight ratio was reduced margin and increased complexity—in some cases we made these warheads *very* complicated. We also didn't leave much room for error in the performance of these designs. The margin was quite low. In these Cold War weapon designs, even something small going wrong can affect the weapon's performance. We often compare weapon designs to sports car designs. Cold War weapons were much like Ferraris: complex and lightweight, with high performance but little margin for error and with costly build and maintenance requirements. And man, oh man—did it take nuclear testing to make sure those things would work! Not just one nuclear test but in some cases up to a dozen! We needed that many to confirm that the highly optimized designs would work under all kinds of environmental and combat conditions. We tweaked these designs between tests to ensure they were operating as we intended, given their tiny margins for error.

Remember that during the Cold War we never thought we would cease nuclear testing or that we would lose our production capability. We designed weapons to stay in the stockpile for 10 to maybe 20 years, certainly not 50 or 70 . . . or 100 years. New production and new designs had always replaced older weapons in the stockpile. But all this changed with the period from 1989 to 1992.

The result is that the age of our weapons today requires us to eventually refurbish and “life extend” each warhead. This refurbishment is the work of the Life Extension Programs (LEPs). The LEPs are designed to refurbish, modify, update, or replace components to ensure that the weapons remain safe, secure, and reliable for an additional 20 to 30 years. The LEPs were executed first for the W87 (a Livermore design) and then for the W76 (from Los Alamos). The B61 bomb (also a Los Alamos design) is next. Eventually, all the weapon types will be “rebooted” in this fashion.

The Reliable Replacement Warhead

Back to our discussion of the RRW. In conjunction with the LEPs, we also considered an option to design a new warhead with modular components that would work across different weapon systems. These warheads would be easier and less expensive to both build and disassemble for maintenance and surveillance. And they would be less dependent on a lot of exotic materials and fabricating processes. They would have the very best safety features and the kinds of security features that would render them unusable unless and until the president authorized their use. The prior tradeoffs made to maximize yield to weight would be relaxed, and the space and weight this freed up would be used to add safety and security features. And it’s important that these designs would NOT have new military yields or delivery means. They would functionally replace existing systems in the stockpile, using existing delivery systems, and would have no new military capability. Perhaps the biggest challenge was to build a new warhead in which we had so much confidence that the military would accept it and deploy it without nuclear testing.

All of those requirements became the design goals of the RRW program. As good fortune would have it, I was asked to head the RRW team for Los Alamos and Sandia/New Mexico. Perhaps the most rewarding 18 months of my career were from May 2005 to December 2006 when we conceived the RRW design. The Los Alamos-Sandia team came up with ideas that were absolutely ingenious. The team’s design was entirely “next generation” in safety, security, and reliability, and it could have been made ready for cost-effective production in about 48 months. It could have been built in about a third of the time needed for a new conventional weapon.

The RRW, unfortunately, had opponents concerned that it was too much like a "new weapon" and would send the wrong messages from a proliferation standpoint. RRW was cancelled in 2009, although many of the ideas that came out of the RRW have made their way into LEPs. Far from being a one-time, cancelled effort, the RRW provided ideas and thinking that are starting to pay real dividends as we look at improving margin, safety, and security in existing weapons. They have become an integral part of our LEPs. I can think of many examples in which ideas conceived during RRW have been adapted into LEPS.

In many ways, the RRW program represented a mature stage of success for stockpile stewardship. All of the understanding, tools, data, and computations came together to show just how much we had learned in the prior 20 years. The fact that we could offer a new-generation design to meet 21st-century needs and have that design accepted by the military is a signature moment in our program. It wouldn't have been possible without the many stockpile stewardship advances and investments earlier.

The Silent Sentinels

Some people say nuclear weapons are fading from consciousness and aren't that important anymore. In my mind nothing could be further from the truth. A few years ago, a member of a congressional committee looking at funding for the nuclear weapons program asked a general, "But General, what role do nuclear weapons have? Why do we need them anymore?" The general straightened up, looked that congressman in the eye, and said, "Mr. Congressman, nuclear weapons function every day. They are our silent sentinels, reminding everyone that this country has the ultimate means of reprisal . . . . [T]hose that would wish ill to the United States must always calculate and have second thoughts when contemplating an act against us." That's deterrence.

At a deep level, I don't like the fact that we have to threaten retaliation to maintain peace. However, that's the contradiction of deterrence. And from 1945 until today, it hasn't failed. It still operates. Norris Bradbury [the Laboratory's second director] used to call new staff members into his office for a short discussion. He would start by saying, "If the products of our work are ever again used in anger, then we will have failed in our mission. We don't build nuclear weapons to kill people. We build nuclear weapons to buy time for our political leaders to find a better way." Bradbury understood the contradiction of nuclear weapons—that we retain these objects of awful destruction in order to preserve the ultimate peace.

Nuclear weapons have a destructive power that current generations have never witnessed. They've never seen a nuclear weapon tested; it's only in the abstract that they can appreciate the awful power of these creations. Harold Agnew [the Laboratory's third director] proposed that once in every generation a nuclear weapon be detonated above ground, with world leaders required to witness it and see its sheer and size and power for themselves. If each generation of leaders did this, they would surely never use a nuclear weapon.

A Better Way

Is there a better way that Norris Bradbury alluded to? Can we achieve the benefit of deterrence while lessening the risks? These questions were a few that I examined during my time at Stanford University as the inaugural William J. Perry Fellow. I came to understand that the work of Los Alamos and the laboratories was growing in importance as the country and the world strove to reduce the sheer numbers of nuclear weapons. Indeed, could we find a roadmap to Bradbury's vision of a better way?

I've come to believe that as stewardship has moved forward, there's been a new kind of payoff. As we become really good at understanding *how* nuclear weapons work and more confident that we can, with agility, reconstitute an arsenal to respond to new threats, that capability becomes a growing part of the deterrent. This is the future. The weapons that we designed at Los Alamos are not the sole protector of our security. The work itself—the science and engineering—is also part of the deterrent. Elements of this strategy, a *capability-based deterrent*, have been adopted as part of the most recent Nuclear Posture Review, conducted by President Obama. (See "Reconstitution as Deterrence," J. Martz, *Actinide Research Quarterly*, May 2011, pp. 1–9.)

The most important element of stockpile stewardship and a capability-based deterrent is the *people*. I've been a witness to innovations that astound me to this day. The clever ideas, the dedication, and the work ethic of Los Alamos staff are extraordinary. There have been achievements here in support of stewardship that the previous generation couldn't have imagined. My greatest pride comes from interacting with hundreds of exceptional scientists and engineers at the Laboratory. The challenge in capability-based deterrence is ensuring that this capability is *agile*, that we can respond to world developments with sufficient agility so that no one doubts our ability to overwhelm and defeat an adversary.

In a very real sense, our deterrent will evolve so that it's not just the products of our work—the nuclear systems we design and maintain—but also our work itself that will become the protector of our security. In this vision, the Laboratory is more important than it ever was, and that's where we're headed.

Sidebar: Nuclear Weapons 101—How They Work

Nuclear weapons are complex devices operating at the extremes of physics, chemistry, and materials science. The temperature, pressure, velocity, density, and energy produced in a nuclear detonation are essentially unprecedented in human experience. Furthermore, the need to ensure the safety and security of nuclear devices leads to a great paradox: The weapon must be designed to ensure that the weapon's exceptional destructive power *does not* manifest itself when not desired but *always does* when required. And all the components that produce both results must be designed to fit within a volume and mass of material smaller than many kitchen appliances.

A nuclear detonation can be viewed as a series of cascading, compounding events, each of which helps amplify energy production for use in the next main stage. A modern thermonuclear weapon has two main stages: the primary and the secondary. The primary is essentially a fission bomb that releases energy from a runaway fission chain reaction. That energy reaches the secondary, setting it off. The fuel in the secondary undergoes both thermonuclear fusion and fission to release hundreds to thousands of times more energy than a fission bomb does.

Detonation of the modern thermonuclear weapon begins with a small electrical signal to a detonator in the primary, a signal that is scrupulously controlled to ensure it is transmitted only when it is certain that a detonation is desired. This signal fires a number of detonators, and each detonator ignites a small charge of explosive. These small charges ignite the primary's main charge of explosives. The symmetrical detonation of this main charge is essential, as it compresses a core of fissile material—material capable of undergoing nuclear fission when hit by neutrons—into a supercritical mass. Plutonium and uranium are the fissile materials most often used to make cores. This core of material is contained within a component called a "pit." When the core is compressed to a supercritical mass, a runaway fission chain reaction begins in the plutonium or uranium, generating tremendous amounts of energy very rapidly.

In modern designs, the core is a spherical shell that surrounds a small quantity of tritium and deuterium, heavy forms of hydrogen. When sufficiently compressed and heated by the fissioning shell, the deuterium and tritium can fuse. This fusion reaction generates copious quantities of additional neutrons. The extra neutrons "boost" additional fissions in the surrounding shell, increasing the primary's yield.

The energy from the primary is manifested as radiation (x-rays and gamma rays), neutrons, and fragments of the fissioned nuclei. Much of this energy is rapidly deposited in the materials that make up the weapon, heating them to temperatures exceeding the temperature of the sun. In modern, two-stage thermonuclear weapons, the primary is held within a metal containment shell called a "radiation case." This component contains and reflects the radiation from the primary onto the secondary, which contains both fission and fusion fuels. The tremendous amount of energy absorbed by the secondary creates a crushing shock wave that compresses the secondary into a state that produces vast amounts of further fission and fusion.

The yield from a secondary greatly exceeds what the primary can create, and this energy manifests itself as radiation, neutrons, and fission products, just as in the primary. The secondary's radiation is also absorbed by the materials that make up the weapon, heating them to temperatures so extreme that they re-radiate x-rays into the environment immediately around the weapon. In an atmospheric detonation, these x-rays are absorbed by the air, creating a fireball and a tremendous shock wave, the source of the direct damage from a nuclear explosion. Other effects of the nuclear detonation include direct radiation, both x-rays and neutrons, as well as nuclear fallout in the form of fission products.

Safety and security, collectively known as "surety," are paramount in nuclear weapon design as well. The requirement to avoid unintended or accidental production of yield from nuclear weapons is a long-standing one, as is a desire to mitigate the consequences of various accidents. These needs have been considered from the beginning of nuclear weapon design activities. They are formally stated in the "MC"(military characteristics) documents associated with each weapon. These documents describe the range of environmental conditions a weapon may experience, including accident scenarios, as well as the weapon's desired response to each environment, such as no production of nuclear yield and containment of radioactive materials, when possible.

One of the most challenging requirements for a weapon is "one-point safety." This requirement states that if the explosives in a weapon were detonated at the "most vulnerable point," there would be less than a 1 in 1,000,000 chance of a nuclear yield greater than the equivalent of four pounds of TNT. The most vulnerable point is actually quite challenging to calculate or determine. A considerable number of nuclear tests were conducted to prove or determine the one-point safety characteristics of given nuclear designs, including the location of the most vulnerable point within the explosive charge. Today, considerable effort is expended on advanced codes that can calculate accident scenarios instead of normal performance. These "safety" codes can certify one-point safety.

The need to counter post-9/11 terrorist scenarios has increased the desired requirements for nuclear weapons. The goal is to deny weapon yield to a foe, even a determined foe who knows the weapon's design details and is willing to sacrifice his life to achieve this yield. This is an exceptionally challenging objective. Both of the proposed Reliable Replacement Warhead designs contained advanced features that met this objective.