Reflections on Notable Events in High Performance Computing at LLNL Since 1995

An Account of LLNL's Return to Preeminence in HPC

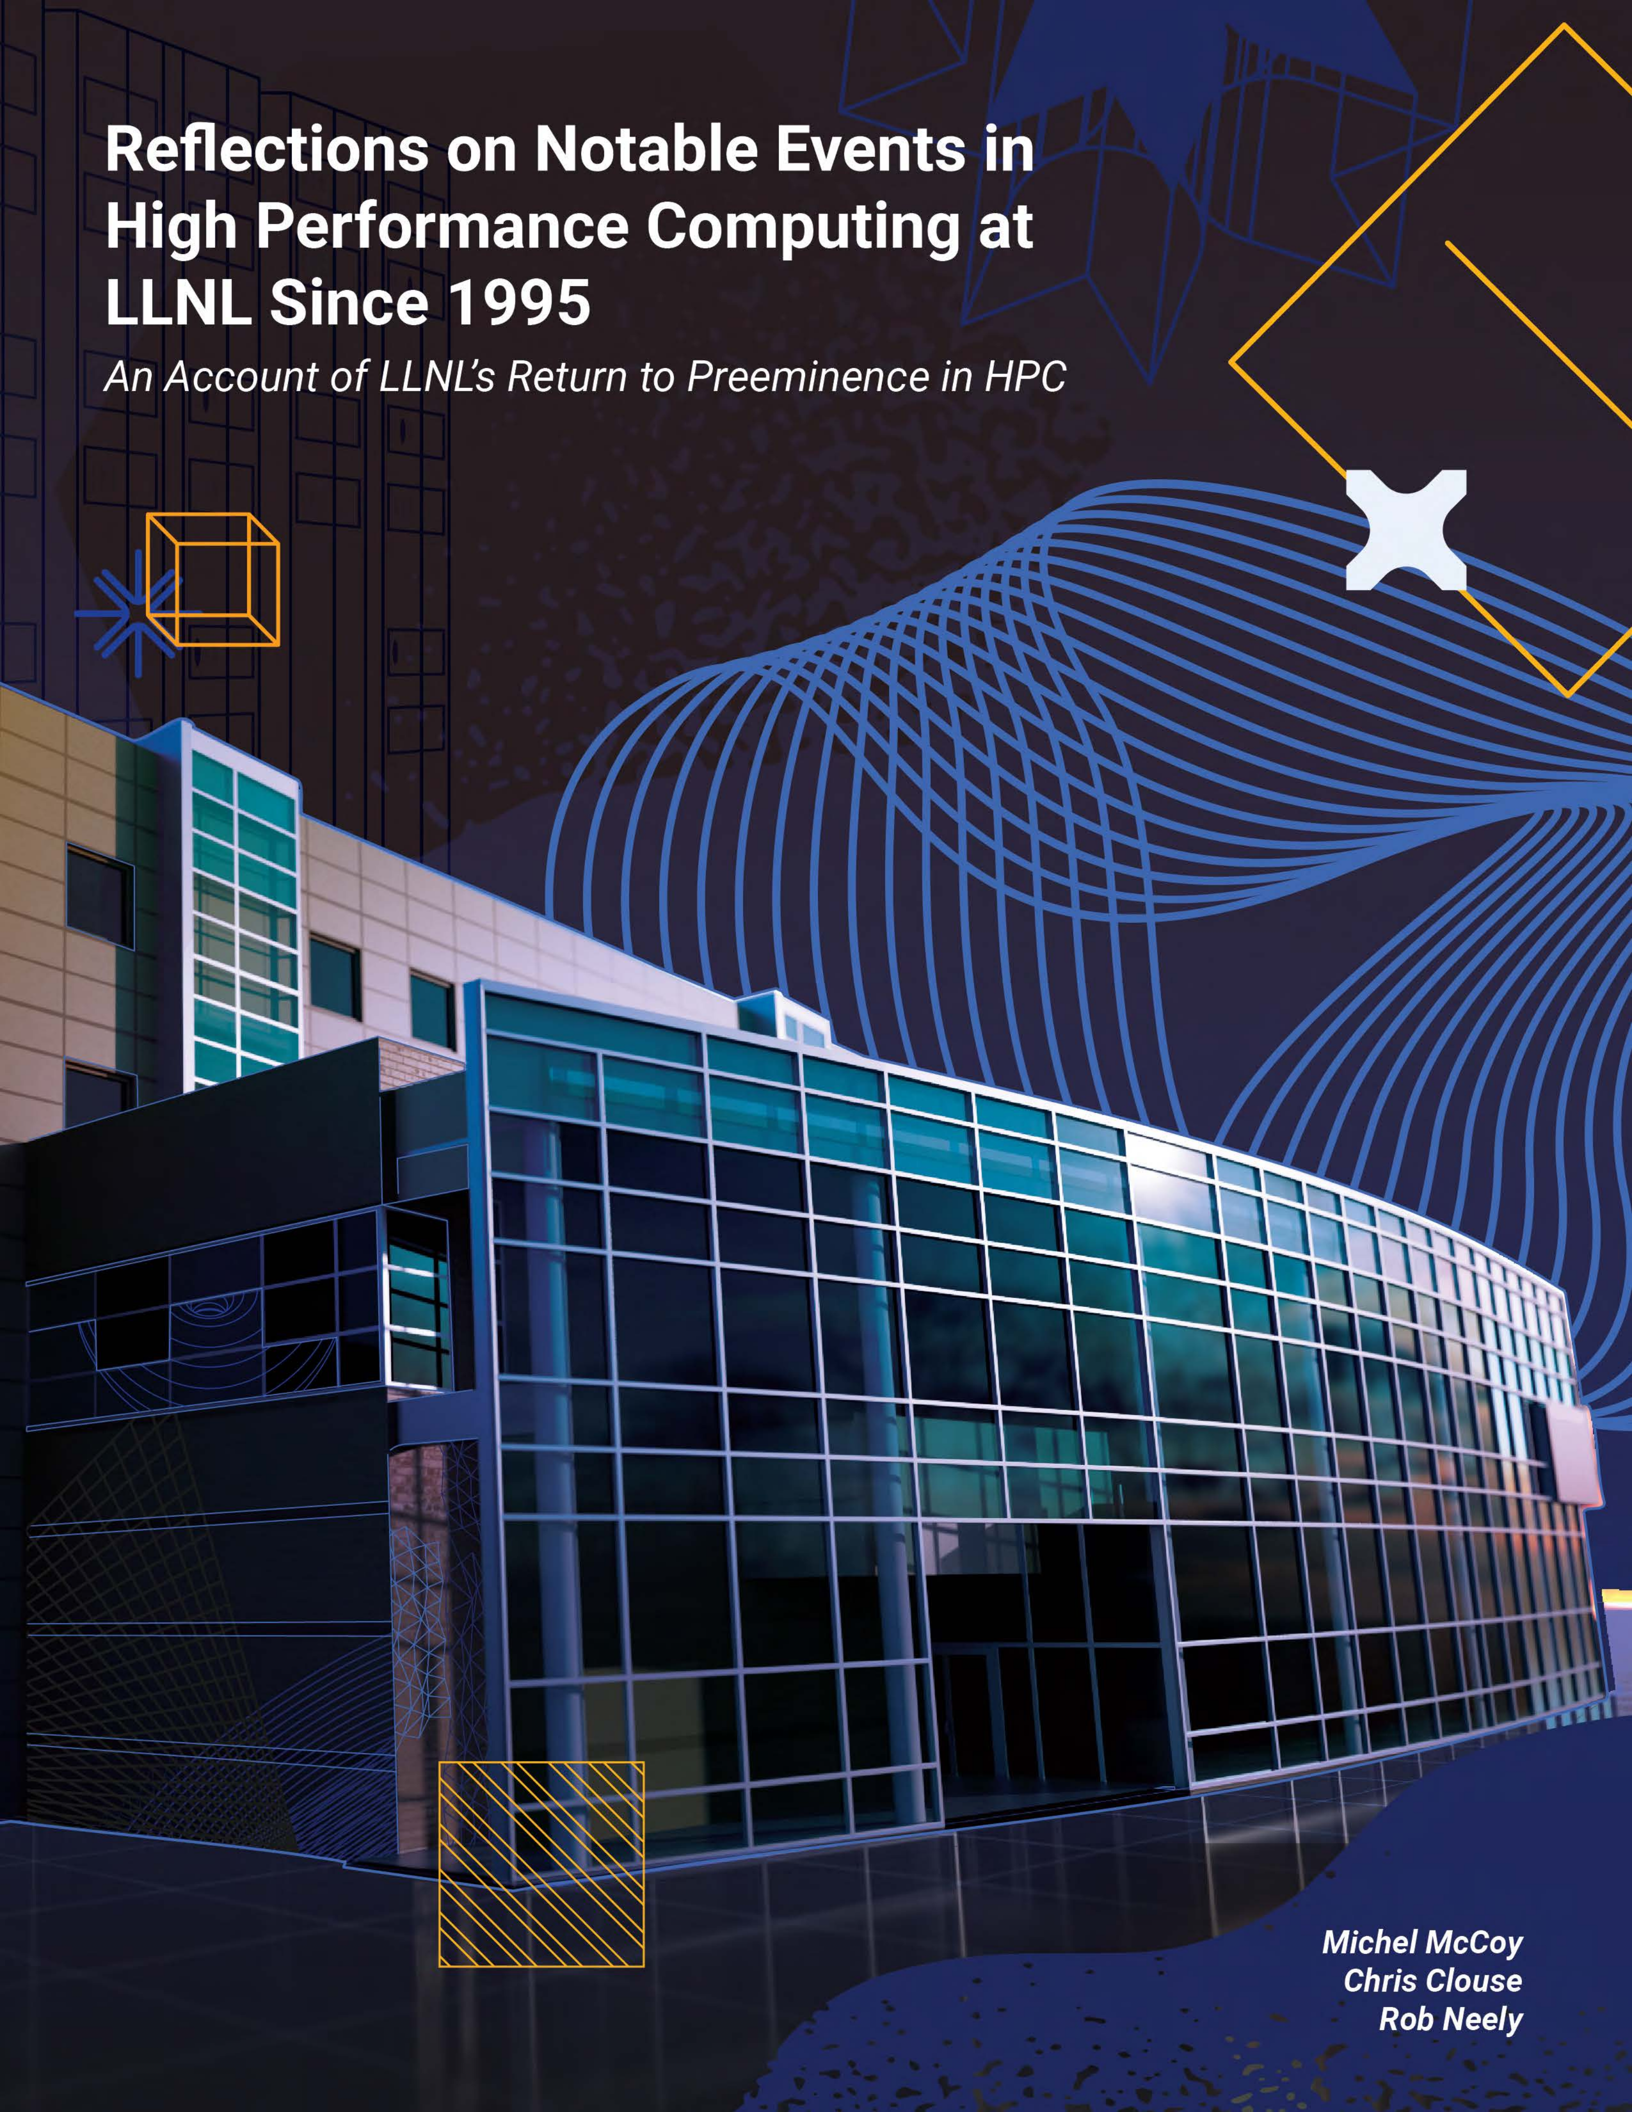

Michel McCoy
Chris Clouse
Rob Neely

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

LLNL-BOOK-2001620 TID | 25-81390

Managing Editor, Writer
Kristen Howard

Editorial Support
Margaret Davis

Creative Director
Alii Diaz

Art & Graphics
Heather Chandler
Jay Clouse
Jacob Long

Reflections on Notable Events in High Performance Computing at LLNL Since 1995

An Account of LLNL's Return to Preeminence in HPC

Michel McCoy, WSC Program Director, 2015–2020, and
ASC Program Director, 2002–2020

Chris Clouse, WSC Program Director and ASC Program Director,
2020–2022

Rob Neely, WSC Associate Director and ASC Program Director,
2022–present

ACKNOWLEDGEMENTS

Little of value can be accomplished without the help of others. For this account, the authors are indebted to George Zimmerman, Bob Tipton, Brian Pudliner, Becky Springmeyer, Randy Christensen, Tom Adams, and Mark Seager for their historical input. We thank Terri Quinn, Bronis de Supinski, Teresa Bailey, Brian Pudliner, Rob Rieben, and Jim Rathkopf for their comments and corrections to early drafts, which were essential for accuracy and reality checks. We salute Kristen Howard for her editorial advice and support.

The work of 30 years described here was the communal effort of an ecosystem of technical experts relying on the support of the NNSA ASC program, the LLNL Director's Office, and fertile partnerships with Los Alamos and Sandia national laboratories and the Office of Science laboratories at Argonne, Lawrence Berkeley, and Oak Ridge. We hope that those who listened and responded to our many supplications will take solace in what was accomplished.

I was fortunate to build a career at LLNL during a remarkable three decades of explosive growth and innovation in supercomputing for national security. When I joined LLNL in 1994 as a code developer in what was then called B-Program, I was aware of the Laboratory's storied history. Over time, in talking with colleagues and looking at documents, I gained a deep appreciation of LLNL's preeminent role in computing from the 1950s to the 1980s.

What I didn't know as a new hire is that I would have a front-row seat in an era characterized by a 25-million-fold increase in computational power. This astonishing achievement sprang from LLNL's culture of technical excellence, close partnerships, strong leadership, and willingness to embrace sometimes-risky change. Year after year, LLNL fielded the fastest systems in the world, developed powerful applications for scientific and national security, designed forward-looking facilities and infrastructure, and transformed the broader supercomputing ecosystem through open-source software development.

When I became the associate director for Weapon Simulation and Computing in 2022, one primary goal was to document this technological tour de force for posterity. I hoped my predecessors Chris Clouse and Michel McCoy might contribute to this effort in their retirement. Indeed they did, and what you are about to read is the result.

Initially, we thought we could tell the story in 25 or 30 pages. But when our first draft exceeded twice that length (though we had much more to share, combined with a great reluctance to cut), we decided to abandon strict page limits and focus on a comprehensive historical account.

Warm thanks to our many reviewers, who enriched the narrative with first-person insights and kept us accurate. And especially I offer heartfelt thanks to Mike and Chris, for their contributions to this project and especially their generous mentorship throughout my career—I owe my success to their guidance. Their influence on the WSC program has been bold and lasting. These pages chronicle but a portion of their legacy.

Rob Neely
Associate Director, Weapon Simulation and Computing
December 2024

Michel McCoy **Chris Clouse** **Rob Neely**

TABLE OF CONTENTS

1

INTRODUCTION

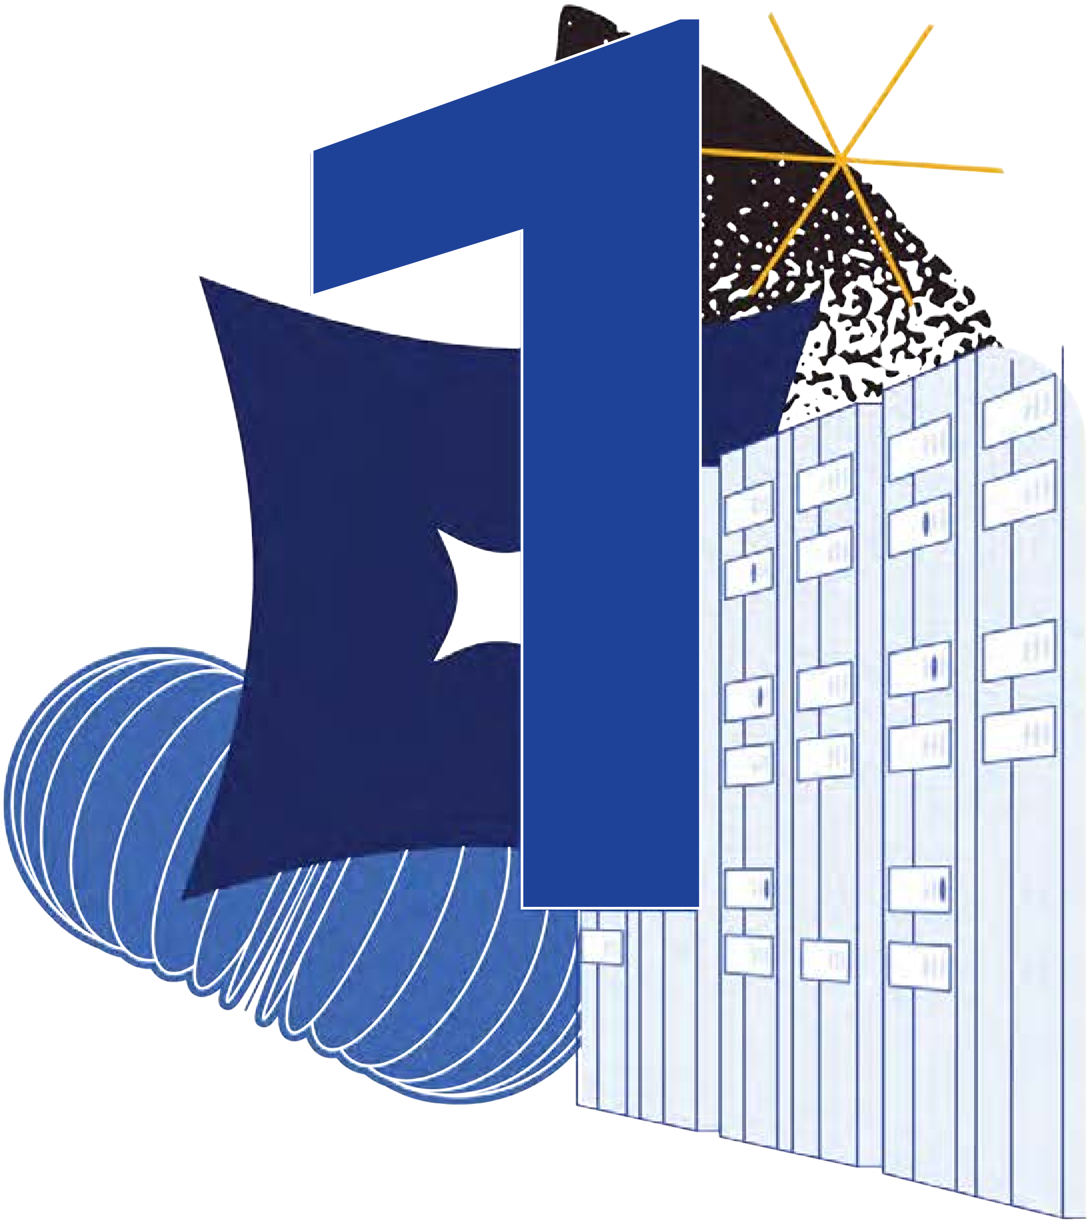

> *They say a good pundit never puts a number and a date in the same sentence. Nevertheless, the Accelerated Strategic Computing Initiative set a remarkable goal in 1995: to achieve an entry-level three-dimensional (3D) simulation of an integrated system on a 100-teraflop computer within a decade. ASCI succeeded with the Purple system at Lawrence Livermore National Laboratory in 2005. Perhaps more noteworthy is that a far more demanding capability, the routine use of 3D weapons codes at sufficient resolution for design and assessment, was realized some 15 years later. An initial capability demonstrated on Sierra circa 2019 is expected to become standard practice on El Capitan in 2025.*
>
> *In 2005, no pundit could have predicted when ASCI's 1995 goal might happen. But while intermediate goals had to be specified, LLNL kept its eyes on the prize. This suggests the importance of institutions that have the fortitude to persevere over a decade or more, despite the vagaries of the funding cycle and technical surprises. Predictive simulation is a journey with milestones but lacks a clearly demarcated destination. Routine 3D simulation for assessment is a major milestone in this journey. Such was the result not only of institutional perseverance, but of technical excellence in all relevant domains. The story is told in these pages.*

The evolution of computing capabilities and succession of world-class computing platforms sited at Lawrence Livermore National Laboratory (LLNL) since the inception of the Accelerated Strategic Computing Initiative (ASCI) has never been chronicled comprehensively until now. In these pages, we document many of LLNL's notable contributions to the global high-performance computing (HPC) landscape over the better part of three decades. Beyond this history, and far more important, we delineate the culture that made it happen, for the benefit of staff today.

In 1940, E.O. Lawrence spoke at Berkeley after having received a Nobel prize. He commented, "The day when the scientist, no matter how devoted, may make significant progress alone and without material help is past." If it is indeed true that complex programs are necessary to enable future advances, then while flashes of brilliance in one subdomain or another are always welcome, the key is extraordinary technical leaders who are free to make decisions regarding each component of a program. The role of management is to assure that overall integration and coordination continues and to address extraordinary fiscal, bureaucratic, or technical challenges that transcend a single area. However, the essence of Livermore's success lies in technical leadership and commitment from each subdomain, even as each recognizes their interdependence and adheres to the integrated plan. This model may seem obvious; but why is this combination of leadership and coordination so rare in human affairs? The Lab's approach in the simulation effort delivered ceaselessly to national security while catalyzing the evolution of the scientific method.

Livermore's HPC journey from terascale in the mid-1990s to exascale with the arrival of El Capitan in 2024 demonstrates an enhancement of computing capacity by roughly 50-million times over 30 years. This enduring advancement was realized through years of unnerving technical, fiscal, facility, and procedural challenges. Addressing setbacks requires courage, rapid invention, and seamless coordination among all the elements in a computing ecosystem. It has been said that what is often thought impossible is indeed possible, but usually takes longer. We learned through experience that occasional delays, while frustrating, left no option but to persevere with faith in the extraordinary imagination of those entrusted to deliver to national security.

The primary goal of this document is to give new generations of LLNL scientists and engineers a sense of how Livermore endured and overcame perplexing circumstances, so that as new challenges appear, they may be surmounted according to the historical context and character of the Laboratory. Persistence and perseverance are the virtues at the heart of what distinguishes great institutions from the mediocre.

We emphasize two themes throughout this narrative. First is the computing effort's contribution to national nuclear security by our assessment of the performance, safety, and reliability of the American nuclear deterrent. We accomplished this by developing increasingly predictive simulation in the absence of underground nuclear testing (UGT). The second, more esoteric theme, is LLNL's contributions to the evolution of the scientific method and HPC in general. Since the early 90s, the world has witnessed simulation joining theory and experiment to form a triad for scientific discovery.

LLNL's contributions to this evolution have been significant. These contributions were made possible by quick, strategic, and bold thinking to assure the computational health of the nuclear-weapons program as it accelerated under ASCI and the inclusion of all disciplines and programs at the Laboratory. In short, the touchstone was one lab or no lab.

LLNL pursued this unity in a sustainable and defendable way. It was arguably the first of the national-security labs to develop and execute a grand strategy for the computational empowerment of all scientists in the workplace. Rooted in multidisciplinary teams working on cutting-edge computational physics, computer science, and algorithmic innovation, the implementation of this vision catalyzed the evolution of the scientific method here. Many of the tools, methodologies, and processes foundational to today's HPC environments were largely byproducts of LLNL's efforts in this arena.

The work described was done under the aegis of the Department of Energy (DOE) National Nuclear Security Administration (NNSA) by the ASCI, and later the Advanced Simulation and Computing (ASC) programs that funded and managed the weapons computing components of the three NNSA laboratories. We are deeply indebted to the consistent support of NNSA leadership, who showed forbearance when things went awry and provided moral and financial support for promising Laboratory thrusts. This was characteristic of ASC headquarters (HQ) from 1995 to the present. Former directors Dimitri Kusnezov, Bob Meisner, Doug Wade, Mark Anderson, and current director Thuc Hoang were consistently responsive, even when LLNL requests were outside the bounds of current practice. Without their flexibility, little innovation would have been accomplished.

The Weapons Program principal directorate at LLNL directed ASCI and ASC efforts and provided essential moral support for the institutionalization of HPC. The Computation directorate provided the talent for Livermore Computing (LC), the Center for Applied Scientific Computing (CASC), and the Applications, Simulation, and Quality (ASQ) division, which has been integral to the mission. The directorate was led by Bill Lokke, followed by Dave Cooper, Dona Crawford, and today, Bruce Hendrickson. That all stakeholders at LLNL, both program and discipline, worked in step for the 30 years chronicled is a primary reason we can proudly point to the achievements described herein.

LLNL directors consistently advocated for computing—both ASC and institutional—during the good times, but especially when things were going seriously awry. These interventions are too many to list; suffice it to say that computing is a costly enterprise. Directors were relied upon to carefully assess which approach to take when fiscal challenges were severe in computing and other programs and priorities. It was a balancing act requiring cool but sympathetic temperaments. During the first days of ASCI (now ASC), director Bruce Tarter supported the creation of institutional computing and worked with IBM to solve early "near-death" challenges. Fast forwarding to today, directors Bill Goldstein and Kim Budil provided essential support in resolving complex Laboratory Directed Research and Development (LDRD) tax and lease-to-own issues. The ASC and institutional-computing programs accomplished the impossible only with the unwavering support of the directors.

The intent is not to provide a history of ASC, as there are well-written and complete sources that cover this topic.[1,2] Nor is this document intended to dwell on the significant accomplishments of our sister laboratories, Sandia and Los Alamos. Rather, this document is focused on LLNL's perspectives regarding some of its more prominent work in HPC. In this spirit, it is motivated by Bill Lokke's

---

[1] Larzelere, Alex R. II. 2009. "Delivering Insight: The History of the Accelerated Strategic Computing Initiative (ASCI)," LLNL document UCRL-TR-231286 (https://www.osti.gov/biblio/965460).

[2] Stevens, et al. 2020. "25 Years of Accomplishments in the Advanced Simulation & Computing Program." (https://www.energy.gov/nnsa/25-years-accomplishments-advanced-simulation-computing-program).

informative paper describing the early years of computing at LLNL,[3] as well as the wonderful series of interview-based stories collected by George Michael and available on the web.[4]

This document does not cover every significant effort in HPC at LLNL over the past three decades, but centers around major platform deliveries and attendant efforts in weapons applications, supporting software, and tools.

---

[3] Lokke, Bill. "Early Computing and its Impact on Lawrence Livermore National Laboratory." https://www.osti.gov/biblio/902225.

[4] Michael, George A. "Stories of the Development of Large-Scale Scientific Computing at Lawrence Livermore National Laboratory." https://www.computer-history.info.

2

HISTORICAL CONTEXT

Historical Context

a. Status of the Weapons Computing Program Circa 1995

From the time that LLNL was founded in 1952 until the late 1980s, the trajectory of computing at LLNL was aggressive, supporting a growing nuclear-weapons design and engineering program. With the collapse of the Soviet Union in 1991 and subsequent Strategic Arms Reduction Treaty (START), followed by the cessation of American UGT in 1992 and the 1996 Comprehensive Nuclear Test Ban Treaty, many in the nation demanded a "peace dividend." This meant redirecting funding from defense (including the nuclear enterprise) to pressing national needs elsewhere in the national budget. Major decreases in funding were made to the weapons laboratories—LLNL, Los Alamos National Laboratory (LANL), and Sandia National Laboratory (SNL). Today, the country is confronting the reality that this dividend was at best a loan that must now be repaid. In 1991, many wanted to believe that superpower confrontation was over. It wasn't. As Mark Twain may have said, "History does not repeat itself, but sometimes it rhymes." In the final chapter, we consider where LLNL might lead in today's environment.

In 1989, the number of LC staff was 225. By 1992, full-time employees were reduced to 132. While fewer operators hanging tapes for calculation restarts were needed, owing to technological innovations (which mitigated some of the effects of reduced manpower), the 40 percent staff reduction left the LC under extreme stress. At the same time, funding for new platforms slowed markedly. As the adage goes, knowledge is power. The nation's knowledge and expertise in the most awesome weapons ever known resided at the three NNSA labs, which were now threatened with severe attrition. On one occasion, Dave Nowak, the first ASCI executive, showed Laboratory senior management a viewgraph indicating that LLNL's computational capability was slightly below that of Finland. This was not to mock Finland, but to hammer home how far nuclear-weapons computing had declined. It took marketing talent to make this point vividly.

A stopgap solution was to fund the labs to work with industry in technology-transfer programs, or cooperative research and development agreements (CRADAs). The theory was that the tri-laboratories' scientific and computational expertise could be put to great advantage by American industry. The CRADAs were funded by the weapons program and participating industrial partners. This holding pattern reached an apex around 1994, and in cases like the SuperCRADA between Cray, LANL, and LLNL, was helpful in developing both important partnerships with industry and local expertise in parallel computing. However, the key benefit of the CRADA effort was to create a little bit of time for the country to invent and develop a strategic and sustainable plan to support nuclear security. The result was the Stockpile Stewardship Program (SSP), which was established to maintain the reliability of nuclear weapons long past their projected lifetimes. Weapons originally designed to be replaced at regular intervals of approximately every 25 years now had to be maintained much longer, with quantified lifetimes for component replacement and refurbishment. Key elements of the SSP were the development of more-predictive computing capabilities through the ASCI program, alongside experimental facilities to understand key aspects of weapon design to validate the codes.

b. Parallel Computing at LLNL before 1995

Cray vector-based computers dominated the production-computing landscape from the 1980s to the mid-1990s. It was becoming increasingly obvious, however, that the microprocessor revolution could be disruptive and had to be explored. LLNL's exploration was jumpstarted by researchers making constructive cases for funding, either to the weapons program itself or the institution.

The Massively Parallel Computing Initiative

In October 1989, the LDRD office funded the Massively Parallel Computing Initiative (MPCI). Led by Eugene D. Brooks III, the three-year initiative explored the harnessing of large numbers of Brooks's "killer micros" to Laboratory computer applications. The term was coined for a memorable and oft-quoted talk Brooks gave at Supercomputing 1989, "Attack of the Killer Micros." What followed was the acquisition of Livermore's first massively parallel computer, a 64-node BBN-ACI TC-2000 machine, and then an upgrade to a 128-node system. Scientists from across the Laboratory's technical directorates were funded to develop codes in their area of expertise and address software challenges. Results were first published in 1991 in a compendium of work from plasma physics to sedimentation modeling.[5] A very substantial Multiprogrammatic and Institutional Computing (M&IC) effort coalesced as ASCI was

---

[5] Brooks, E D, et al. 1992. "The Attack of the Killer Micros." MPCI Yearly Report, UCRL-ID-107022.

gathering speed at LLNL, and it depended critically on two things—the Laboratory's experience from MPCI and the confidence this generated in the Director's Office that the Lab could field (and the various disciplines would benefit from) generally accessible large-scale parallel computers.

Building on the MPCI with an eye toward classified missions, Dale Nielsen published an internal LLNL technical report in 1991, UCRL-ID-108228, "General Purpose Parallel Supercomputing," that was an early prompt for the weapons program to change direction toward parallel computing and away from the vector-based systems that dominated the landscape.

The Numerical Test Site

By late 1992, research ideas around parallel computing and the early seeds of what would become ASCI started to take shape, as articulated in an LLNL proposal called the "Numerical Test Site"—a not-so-subtle allusion to the Nevada Test Site, where underground nuclear testing had recently ceased. Championed by director John Nuckolls and fleshed out by Randy Christensen, David Nowak, and Eugene Brooks (among others) was a bold and prescient vision of how computing could play a seminal role in the future support of the stockpile. Subsequent pitches of the proposal called it the "Numerical Test Facility," as George Miller, head of the Lab's Council for National Security (CNS), was uncomfortable with implying equivalence between the two NTSs.

Christensen's presentation of the idea to Lab senior management represented perhaps the first official request that the Lab take advantage of parallel computing to benefit the weapons program, which was seeking a new strategy in the wake of UGT cessation. Many NTS elements were later found in early planning documents for ASCI, as management felt strongly that despite LLNL's leadership, a unified front by all three labs was required for the vision to take root. Ultimately, NTS did get funded, once Vic Reis stepped into the leadership of defense programs at DOE HQ in 1993 and cemented the vision of science-based stockpile stewardship with ASCI a key element.

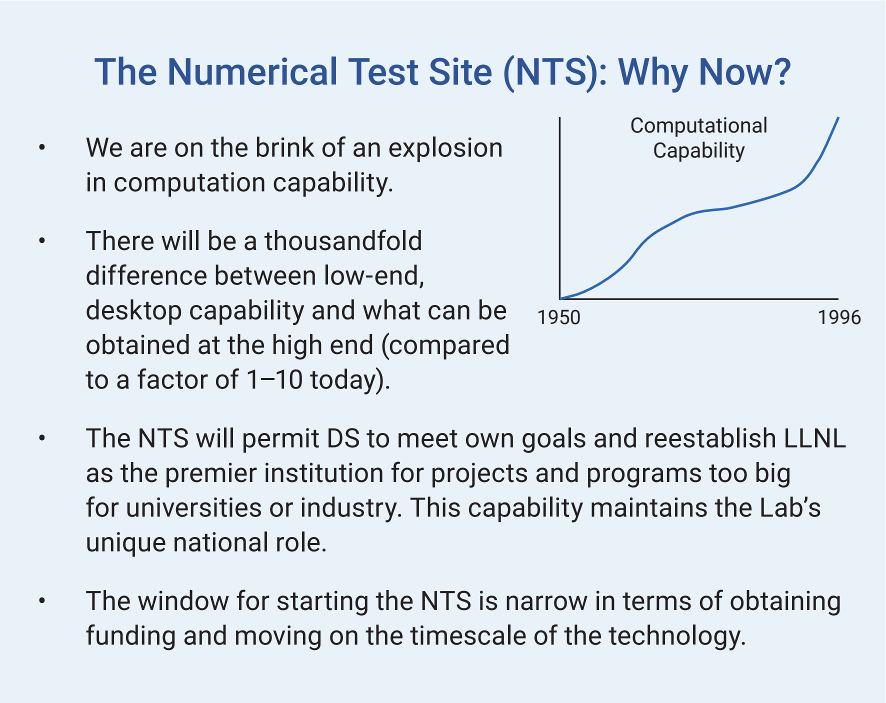

The Numerical Test Site (NTS): Why Now?

- We are on the brink of an explosion in computation capability.
- There will be a thousandfold difference between low-end, desktop capability and what can be obtained at the high end (compared to a factor of 1–10 today).
- The NTS will permit DS to meet own goals and reestablish LLNL as the premier institution for projects and programs too big for universities or industry. This capability maintains the Lab's unique national role.
- The window for starting the NTS is narrow in terms of obtaining funding and moving on the timescale of the technology.

Figure 1. Introductory slide from a presentation by Randy Christensen to the National Security Council pitching the "Numerical Test Site" (late 1992).

The LC's Meiko CS-2

Now keen on getting serious about parallel computing, the weapons program began to focus on the acquisition of a Meiko Compute Surface-2 (CS-2) system. The Meiko architecture was viewed as a compelling transition machine, as it required distributed-memory parallel processing but relied on vectorization for performance. Significant investment in tuning applications for vector architectures had been made in previous decades, and this system offered the hope of leveraging that work. When this $18M procurement was announced circa 1995, there was some political turbulence regarding security and national competitiveness (Meiko was a subsidiary of British-owned Meiko World). This controversy was managed by the Laboratory but foreshadowed the issues and challenges (occasionally associated with approvals) seen by nearly every system sited from then on. Mark Seager, a principal in selecting and integrating the Meiko system, conveyed the story thus:

> The Meiko was delivered with no shortage of hardware and software issues. This was typical of experiences with serial #1 of a new generation. It was particularly true when working on the system build, delivery, and integration of a large-scale, massively parallel system from a small business with strong foreign ties.

CS-2 was a second-generation system based around SuperSPARC and Fujitsu μVP vector processors. On the HW side, the Solaris kernel and firmware in the motherboard needed complex debugging to get the SPARC (two per node) processor talking to the Fujitsu VPUs (two per board). In addition, the LC found out, much to its consternation, that the Fujitsu VPU had major buffer limitations, resulting in terrible performance on gather, scatter operations. Unfortunately, by the time these performance issues were discovered, the LC had migrated the system from the open to the secure side, because LLNL's A and B divisions needed the system for code porting and were anxious for access. That made debugging more problematic, as most of the Meiko engineers were UK citizens. So while the LC had some Meiko engineers onsite, they had to be escorted on the Secure Computing Facility computer floor, which was difficult for all.

The LC system was the largest Meiko ever built, with 224 processors, and was installed in B113 for the weapons program in 1994. The machine peaked at 40.3 gigaflops/second.

Despite the birth pains, Meiko went on to be a very productive system. The weapons program at LLNL used the platform to investigate distributed-memory processing, with modest parallelization capabilities implemented in some production codes.

The machine ran up against the capability of ASCI Blue Pacific in 1998. Nevertheless, it continued to serve and was retired around 2002.

c. The State of Parallel-Computing Applications

In the early 1990s, simulation tools provided some guidance and set certain expectations, but heavily relied on UGTs as a means of calibration. If simulation results did not match experimental, the tools could be calibrated to give a good match. This gave designers some degree of confidence that calibrated tools could be used to calculate the performance of devices similar, in both design and testing conditions, to those tested. With the end of UGT, the continuing evolution of the nuclear stockpile due to aging, and the changing threat environments in which a device must operate, calibrating simulation tools was no longer a viable option. This meant that simulation tools had to be more predictive, primarily by means of improved models and physical data. Eliminating calibration knobs through the implementation of improved models could be achieved only with significant increases in resolution, which were required for model accuracy, and the ability to truly model 3D effects. This translated to problem sizes that needed to be thousands to millions of times greater than the state of the art in the early 90s.

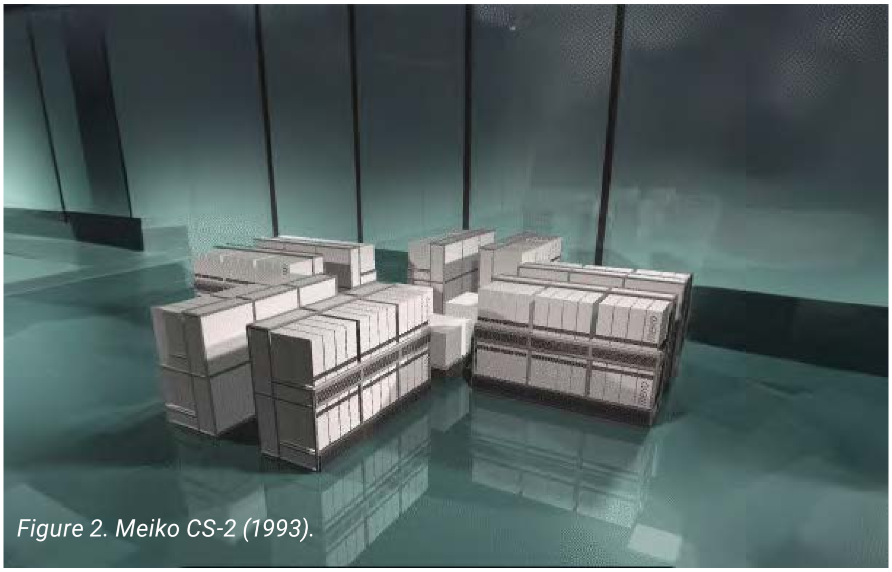
Figure 2. Meiko CS-2 (1993).

In addition, the physical data critical to the accuracy of a calculation, such as material strength and nuclear cross-sections, was mostly limited to data that could be collected in Lab experiments under conditions covering only a small portion of the phase space of a nuclear weapon. Improvements in our experimental capabilities in the stockpile-stewardship era would extend the validity of some of this data, but much of the phase space could be reached only through first-principles simulation tools. These specialized science codes modeled materials at the atomistic level under extreme conditions and would require millions of times more computing power. One-dimensional (1D) simulation tools dominated most of the history of the nuclear-weapons program, while two-dimensional (2D) tools became prevalent through the 1980s and early 1990s. However, the issues faced by an aging stockpile and the threat environment created by an adversary were inherently 3D in nature. Three-dimensional capabilities were essential in lieu of nuclear testing. The only viable path that could accommodate the increase in problem size and need for more computing power was embracing the notion of massively parallel computing. In the early 90s, parallelism was limited to vector processing on Cray computers, where memory was shared among less than a dozen vector processors. The need for higher resolution also drove a greater demand for implicit algorithms unconstrained by the small timesteps associated with highly resolved zoning. These algorithms, along with their global communication needs, would have to be adapted to massively parallel machines without the global memory that Cray computers provided.

d. The State of Unclassified HPC at LLNL in 1995

The weapons program had funded LLNL HPC environments since the Laboratory was established; that computing was essential to the mission was understood from the beginning. Lokke noted that Edward Teller contracted for the first machine before the Lab doors opened. Certainly, all the large systems were procured (and computing environments developed) by that program. Given the nature of the work, the platforms ran classified on their own protected network. While unneeded systems were moved[6] to the unclassified network from time to time, there was no notion of a robust, sustained, dedicated unclassified-computing presence at LLNL.

Scientists not working directly for the program who needed time for their research could get access to classified systems by collaborating with designers and others with programmatic time allocations, and these partnerships could be fruitful. Many scientists did not enjoy ready access, however—for instance, computational biologists, or even scientists working on needed advances in computing physical data relevant to the weapons program. Many lost interest in mainframe access and bought SUN workstations, thereby joining the workstation diaspora and bidding the LC adieu.

With ASCI on the horizon, everything changed. The Director's Office had been keenly aware that LLNL could not thrive with "have and have not" computational citizens. One program was about to compute at the multiple teraflops, and soon petaflop, capability, thinking 3D, using mammoth memory, and implementing expensive advanced-physics models. Meanwhile other programs were settling for Sun workstations or time on NERSC[7] machines, where proposal success was never guaranteed, and allocations were often viewed as insufficient. This created a dysfunctional sociological environment that at minimum inhibited computational crosstalk among programs. There could be no computational lingua franca. There could be no sharing of software tools.

The decision to fix this was made by the Computation associate director (AD), Bill Lokke, with the concurrence of LLNL director Bruce Tarter. The lucky recipient of the guidance to work this was the newly appointed LC director, Michel McCoy (the previous director, Randy Christensen, had been promoted to ASCI deputy under David Nowak, the first ASCI executive at LLNL). Miller's directions were clear: one program (i.e., the weapons program) could not support the work of other programs. This was a legal constraint. Each program had to carry its own weight. However, the institution could help, just as it funds the library and campus roads, which everyone uses regardless of programmatic affiliation. The goal was to invent a sustainable and sociologically acceptable model that would pass muster at LLNL and at NNSA.

Figure 3. Cover of Science & Technology Review *magazine featuring LLNL supercomputing capabilities.*

After a couple of failed attempts at a solution, Eugene Brooks (who was kindly helping McCoy and had witnessed his struggles with the CNS) had an idea. Why couldn't the institution procure the computer, and each program buy into that system? One-hundred thousand dollars would buy 10 percent of the cycles on the $1M system until it was retired. The institution would cover all other costs at the center, including staffing and maintenance. This design would undercut the cost

---

[6] Today, moving a classified system to unclassified use is largely prohibited due to heightened security awareness. In the 1990s, the simpler memory architecture of say, the Control Data Corporation (CDC) 6600, made this less an issue.

[7] NERSC was the National Energy Resource Supercomputing Center, originally housed at LLNL and moved to Lawrence Berkeley National Laboratory in 1995.

of workstations while providing the lure of capability calculations where no workstation dared to go. The approach was briefed to the CNS and approved. Convincing the LLNL legal office and the NNSA Livermore Field Office (LFO) required careful arguments, but eventually succeeded.

There was then the minor matter of convincing the Laboratory diaspora, “Get in, the water’s great!” After many meetings replete with recounting of past grievances, a reunion was achieved by creating the Institutional Computing Executive Group (ICEG). Modeled after the NERSC Executive Group, the ICEG comprised key users across the participating laboratories. Experience at NERSC was paying off. The group issued annual report cards on progress and M&IC responsiveness to the Director’s Office, while goodwill and patience came from the ICEG as M&IC fired up. There wasn’t much initial trust—that had to be earned. By 2001, M&IC was fully woven into the computing fabric of LLNL. The original charter remains relevant and fresh today.

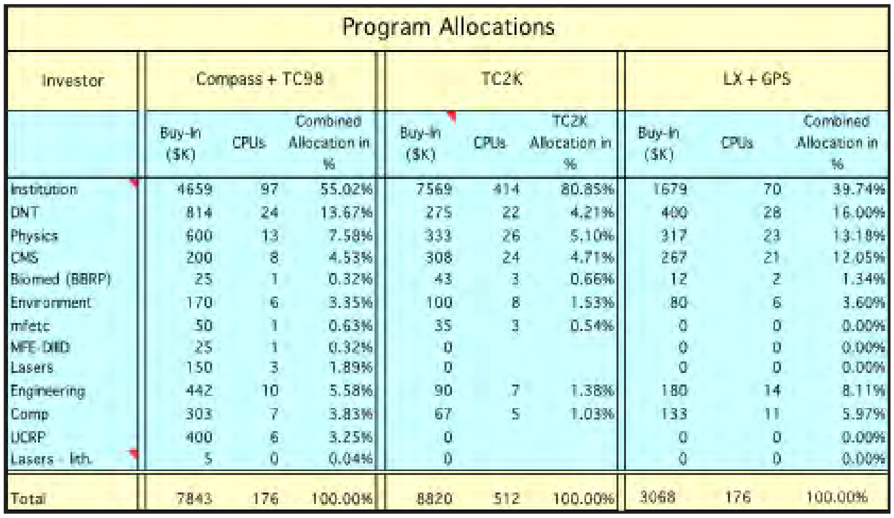

Program Allocations

Investor	Compass + TC98			TC2K			LX + GPS		
	Buy-In ($K)	CPUs	Combined Allocation in %	Buy-In ($K)	CPUs	TC2K Allocation in %	Buy-In ($K)	CPUs	Combined Allocation in %
Institution	4659	97	55.02%	7569	414	80.85%	1679	70	39.74%
DNT	814	24	13.67%	275	22	4.21%	400	28	16.00%
Physics	600	13	7.58%	333	26	5.10%	317	23	13.18%
CMS	200	8	4.53%	308	24	4.71%	267	21	12.05%
Biomed (BBRP)	25	1	0.32%	43	3	0.66%	12	2	1.34%
Environment	170	6	3.35%	100	8	1.53%	80	6	3.60%
mfetc	50	1	0.63%	35	3	0.54%	0	0	0.00%
MFE-DIID	25	1	0.32%	0			0	0	0.00%
Lasers	150	3	1.89%	0			0	0	0.00%
Engineering	442	10	5.58%	90	7	1.38%	180	14	8.11%
Comp	303	7	3.83%	67	5	1.03%	133	11	5.97%
UCRP	400	6	3.25%	0			0	0	0.00%
Lasers - lth.	5	0	0.04%	0			0	0	0.00%
Total	7843	176	100.00%	8820	512	100.00%	3068	176	100.00%

Figure 4. shows a table from a FY00 request to the Director’s Office for additional funding, demonstrating that the investment model proposed to CNS seemed to be working.

Environments across classified and unclassified domains are almost identical today, as are programming models, computers, and tools. Exchanges across classification domains are now common and an essential part of development. These are some of the most satisfying examples of how simulation supports scientific discovery via this coordination. Before the Office of Science received major funding around 2005 for its leadership systems, the M&IC computing environment was comparable to that at NERSC, serving scientists at all DOE labs. While all this seems natural today, at the Lab in 1996 it wasn’t remotely the case. Though initially eyebrows were raised at the sister labs regarding legalities and costs, the necessity of doing something permanent became glaringly obvious. In time, SNL and LANL built sustainable unclassified environments, and each lab then developed its own defendable funding and governance approach for the unclassified environment. In this effort, however, LLNL was the trailblazer.[8]

The ICEG Charter

To provide LLNL scientists from all LLNL programs and research areas with access to a sound, responsive and first class unclassified, capacity computing infrastructure, and to strive, beyond this, to provide the most promising efforts with ample access to the most powerful computing platforms to foster breakthrough science.

The desire ultimately is to enhance LLNL’s reputation as broad-based science laboratory as distinguished from a tightly focused programmatic laboratory.

An excellent encapsulation of the vision at that time, presented nearly a decade later to the Director’s Office, is diagrammed below. We immodestly asserted, “LLNL is an acknowledged leader in simulation world-wide because of its institutional vision and decade-old, coordinated strategy.”

---

[8] https://str.llnl.gov/content/pages/past-issues-pdfs/2001.10.pdf

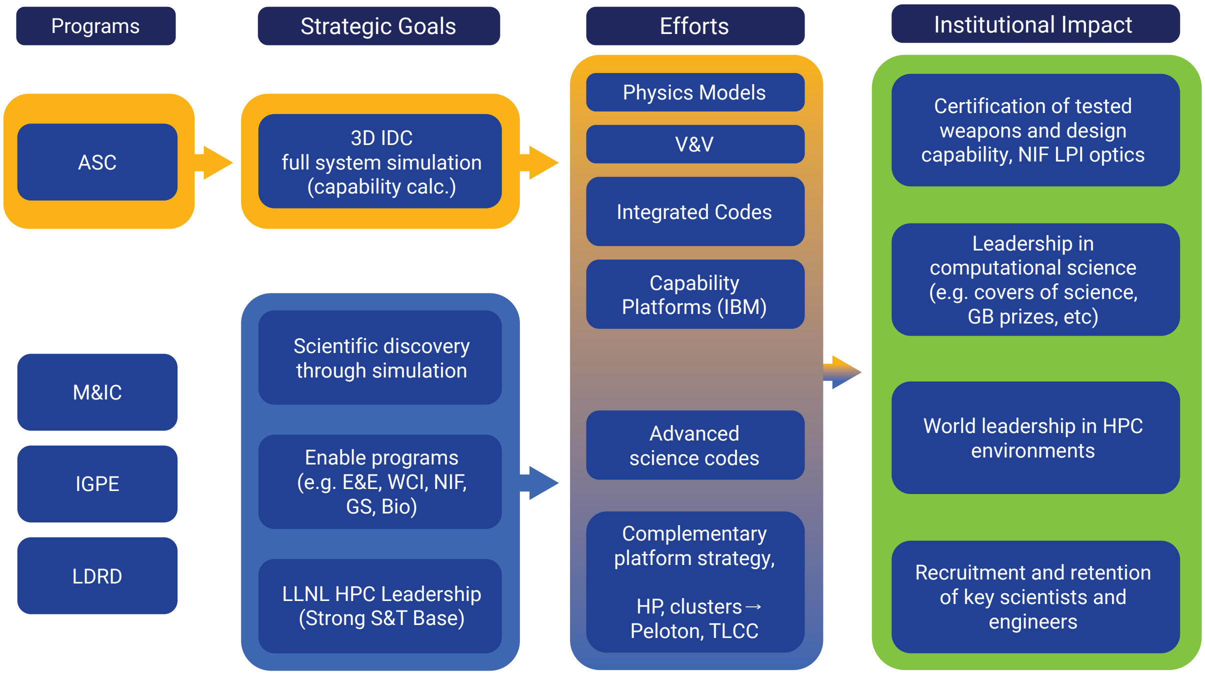

Figure 5. Diagram highlighting the synergy or cooperation between unclassified funding sources.

e. The Departure of NERSC and Genesis of CASC

Today, the Office of Science funds three major sites: the leadership sites at Oak Ridge and Argonne national laboratories (ORNL and ANL) and the National Energy Resource Supercomputing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL). NERSC preceded the others by three decades. The NERSC was established to provide computing power to magnetic fusion energy (MFE) researchers across the country. Many at LLNL are unaware that the NERSC, originally called the Controlled Thermonuclear Research Computing Center (CTRCC), was sited at LLNL using a CDC 6600 retired by the weapons-program computing center, Livermore Computing. The NERSC soon moved into a new building, B451. It was the first Office of Science -funded site to serve remote customers with satellite connections; that it was born at a defense-programs laboratory, not an Office of Science lab, was due to the presence of the LC. In short, there was operational knowledge and leverage to be had that existed solely at national nuclear-security laboratories.

CTRCC served MFE researchers across the country, beginning in 1974. It was a bold effort executed exceptionally by John Killeen. Figure 6 captures the original systems staff, including Killeen's deputies Hans Bruijnes (back row under the sign), Dieter Fuss (center in plaid shirt in front of Hans), and Bruce Griffing (third from right). Killeen and business manager John Fitzgerald are not shown.

At the behest of Hans Bruijnes, members of the small computational-physics group also embedded in the center under Art Mirin. Killeen had insisted on creating this group to stay in close touch with research. Latencies were programmed into

Figure 6. Early CTRCC (NERSC precursor) staff with the new CDC 7600 in 1975. Courtesy of NERSC 40th-anniversary calendar.

their computer connections to equilibrate their frustration with that of compatriots across the country who were dealing with long-distance satellite connections and other interfaces. Since CTRCC was the first DOE site offering remote computing via satellite, the difficulties faced by users dealing with roughly one-second latencies were not well understood, and this was a way to learn more. Additionally, since CTRCC provided a time-sharing system, users could increase their priorities (thereby burning allocations faster) to get quicker turnaround. Cycles were scarce, and priority battles among users were frequent. No one wanted to wait too long for a job to finish or, even more important, for a debugger to reach a breakpoint.

CTRCC's computational-physics group had an advantage. Their priority changes did not suffer satellite latency, so they were able to fine-tune priorities much more quickly and efficiently than colleagues at, say, Princeton Plasma Physics Laboratory or General Atomics in San Diego. Hans gleefully listened to and ignored the howls of the computational-physics group at the artificial equality imposed on them—not only did they have to suffer the frustration of latencies, but they lost their advantage in the priority bidding wars. With 40 years of hindsight in which to cool down, it is now possible to acknowledge begrudgingly that Hans was right. To this day, LC prides itself on its service to remote customers, including those at the Tri labs, and remains committed to making computing at LC feel as local as possible to remote users through the Remote Computing Enablement (RCE) effort led by Todd Heer.

It was inevitable that some people would have a problem with "thermonuclear" in the CTRCC's name, and it was soon changed the National Magnetic-Fusion Energy Computer Center (NMFECC). Later, as it began serving Office of Science scientists from all disciplines, it was renamed NERSC. The center thrived at LLNL, procuring multiple world-class Crays and running a time-sharing operating system like that at the LC, using many of the tools and libraries developed there.

At the time the ASCI was approved, the Office of Science decided to recompete NERSC, with LLNL and LBNL the main contenders.

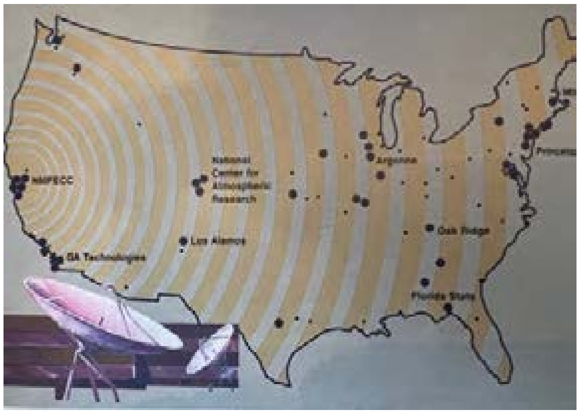

Figure 7. High-level depiction of the satellite network and customer sites served by the CTRCC.

In 1995, the decision was made to move NERSC to LBNL. While it's understandable that the Office of Science desired to place an important program like NERSC at one of its own laboratories, LLNL was deeply disturbed to lose the program. There was little expertise at LBNL for running NERSC, and it had to be built rapidly. There was hope at LBNL that most of the LLNL NERSC employees would migrate 45 miles away to Berkeley.

LLNL was even less happy about seeing valuable staff leave than about losing a program. It was a difficult time for everyone involved—even gut wrenching for those who had been with NERSC since inception. Shortly after the decision to move, senior members of the Director's Office, including George Miller and Bill Lokke, drove to NERSC in B451 and assured staff that those who stayed would have meaningful jobs at LLNL. The NERSC staff were in shock. Over half the staff and the deputy director of NERSC, Michel McCoy,

Figure 8. NERSC staff with the Cray-2 at NERSC in Livermore. The first Cray–2 was installed at the NMFECC (now NERSC) in 1985. Dubbed "Bubbles" for its heat-exchanger water tank and liquid cooling, this was the world's fastest computer; yet by 2011, Linpack tests showed the iPad 2 could rival a four-processor Cray-2 in processing speed.

elected to stay at LLNL, while a substantial number went to LBNL, including the director, Bill McCurdy. A number who left to join LBNL later returned. LLNL was particularly happy to welcome back John Fitzgerald, among others.

In what must have been one of the more ironic moments in LLNL computing history, the timing was almost perfect for the Lab. With new ASCI funding coming to LC and more money flowing in than people to spend it on (arguably a first in DOE history), dozens of seasoned staff were ready to be recruited. Many joined the LC, and as it grew, the weapons AD, Mike Anastasio, named McCoy the LC director. He was happy to accept and looked forward to the opportunity to make amends by doing better.

The staff who migrated to LBNL included John Bell and his world-class computational scientists and mathematicians at the Center for Computational Sciences and Engineering. The center had not enjoyed close connections with the weapons program, as the model for moving its products into the weapons program was neither robust nor overthought. This modus operandi was criticized as "throwing stuff over the fence." One staff member, Steve Ashby, proposed a center at LLNL that would have much better integration and connections to the weapons program while maintaining productive connections to the Office of Science. This would be a balancing act requiring considerable persuasion and salesmanship, at which Ashby excelled. He received great encouragement from ASCI leaders David Nowak and Randy Christensen and worked closely with McCoy to develop a working model for the new Center for Applied Scientific Computing (CASC). Ashby's renowned organizational skills soon bore fruit. CASC settled into an umbrella organization under McCoy that included the LC, and CASC, and what was then called Lab Net, the backbone for unclassified Laboratory networks.[9]

Over the years, Ashby built a powerhouse at LLNL. Though he eventually moved on to become director of Pacific Northwest National Laboratory (PNNL), his legacy thrives even today in the CASC. Overcoming earlier disconnects required constant fine-tuning and close working relationships between weapons-program developers in CASC and developers in the Weapons directorate. Ashby took direct responsibility for the codes, underscoring the importance of culture in a program's success.

CASC was housed in the former NERSC offices along with LC's development environment and information-management and graphics groups. CASC and LC managers partnered to build strong CASC groups in tools, visualization, and scientific data management, and forged extensive collaborations between developers and researchers. LC had a deep focus on operating systems, compilers, visualization, data management, and key aspects of computing at scale, such as debugging, performance, parallel processing, and architectures. CASC hired doctorate-level researchers in each of these areas, and the LC program integrated them into projects within the ASC Computational Systems and Software Environment (CSSE). Strong LC/CASC collaborations in open-source systems software and tools live on today in R&D 100 -winning efforts such as Flux, SCR, Spack, and STAT. LC's chief technology officer, Bronis de Supinski, and key research software architect, Todd Gamblin, both started as CASC researchers and brought their knowledge, vision, and connections to LC.

Figure 9. Randy Christensen.

Researchers continue to affirm CASC's commitment to this collaborative model through innovative, research-based models. In turn, LC brings in complementary operational knowledge that balances early research prototypes with production-quality software rollouts. Their unique blend of researchers, developers, systems architects, and operational specialists has delivered landmark machines and a catalog of open-source software used around the world.

f. ASCI: Challenges and Opportunities

ASCI was a unique and bold applications-driven initiative to develop computational tools for assessing the performance, safety, and reliability of the American nuclear stockpile by means of

---

[9] It was a much simpler world; not until the Wen Ho Lee incident at LANL did the labs respond fully to the necessity of integrated and well-funded computer- and network security, both classified and unclassified.

increasingly predictive simulation. It had four major goals:

1. Develop 3D weapons codes with improved physical models
2. Accelerate platform performance by twice beyond what was possible by merely leveraging Moore's Law
3. Leverage industrial R&D and investment in computing and foster partnerships with academia
4. Within a decade, run a full-system calculation running at sufficient resolution to demonstrate an entry-level 3D capability

Randy Christensen was primarily responsible for the analysis that suggested entry-level 3D capability would require a 100-teraflop computer. The April 1998 issue[10] of *S&TR* magazine featured a discussion by Christensen on ASCI's 6b gt. march to 100 teraflops. This was to be managed through well-defined intermediate goals with rigorous reviews. For computers, the first system (which had been in procurement before the initiation of ASCI) was SNL's Intel 1.8-teraflop ASCI Red. The systems that followed were 3, 10, and 30 teraflops, culminating in 100 teraflops within a decade. Code development was subjected to intense reviews and scrutiny. These were the "burn-code reviews" chaired by Kim Molvig of MIT. The first of these demanded 1,000-processor 3D calculations of the performance of the primary (the nuclear trigger of a two-stage thermonuclear weapon). Given that the program had struggled to run 2D calculations on a few dozen processors, the leap demanded was colossal. Physics improvement was minimized in favor of the forced march to 3D. Reviews of other ASCI components were also demanded, including problem-solving environments (PSEs) and visualization-corridor development.

Many informed scientists, including some leading the charge as key members of the code teams, were skeptical that computing-based stewardship was sustainable. Substituting computational experiments for UGT (along with nonnuclear experiments and information from new stockpile-stewardship facilities like NIF) seemed naïve and even perilous given that not even the infrastructure for such bold simulations existed. ASCI had to overcome deep disbelief by showing rapid initial progress—and any early failures could sink the ship. That meant sailing a bit before it was completely seaworthy, meanwhile proving it could at least float. The initiative was aided greatly by the relentless DOE deputy assistant for research, development, and simulation for defense programs, Gil Weigand, and strongly supported by Vic Reis, the leader of what is today NNSA.[11] Weigand and Reis recognized the necessity of speed and quality but wisely understood that it was sometimes okay for quality to catch up a little later.

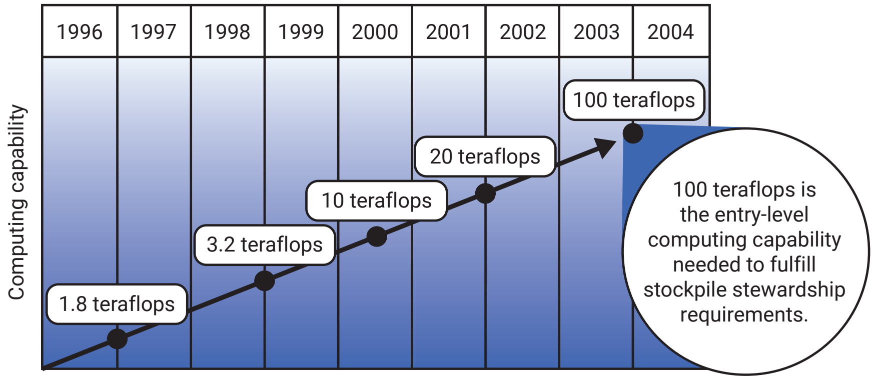

Figure 10. Computing-capacity graph from 1998 S&TR magazine article showing teraflops needed to fulfill original ASCI vision of entry-level 3D simulation of a nuclear system.

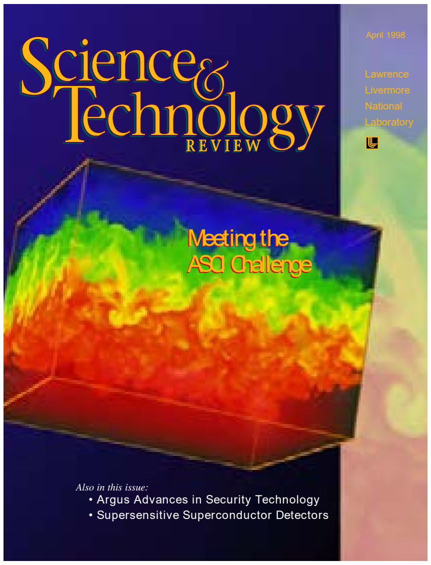

Figure 11. S&TR cover, April 1998, featuring ASCI article.

Figure 12 documents an early ASCI moment in 1996 at the Supercomputing Conference (SC). At bottom left is Vic Reis (the assistant secretary for defense programs in the DOE, equivalent

---

[10] https://str.llnl.gov/content/pages/past-issues-pdfs/1998.04.pdf

[11] Before 2000, Defense Programs was part of the DOE. In 2000, NNSA was created as a semiautonomous agency within the DOE.

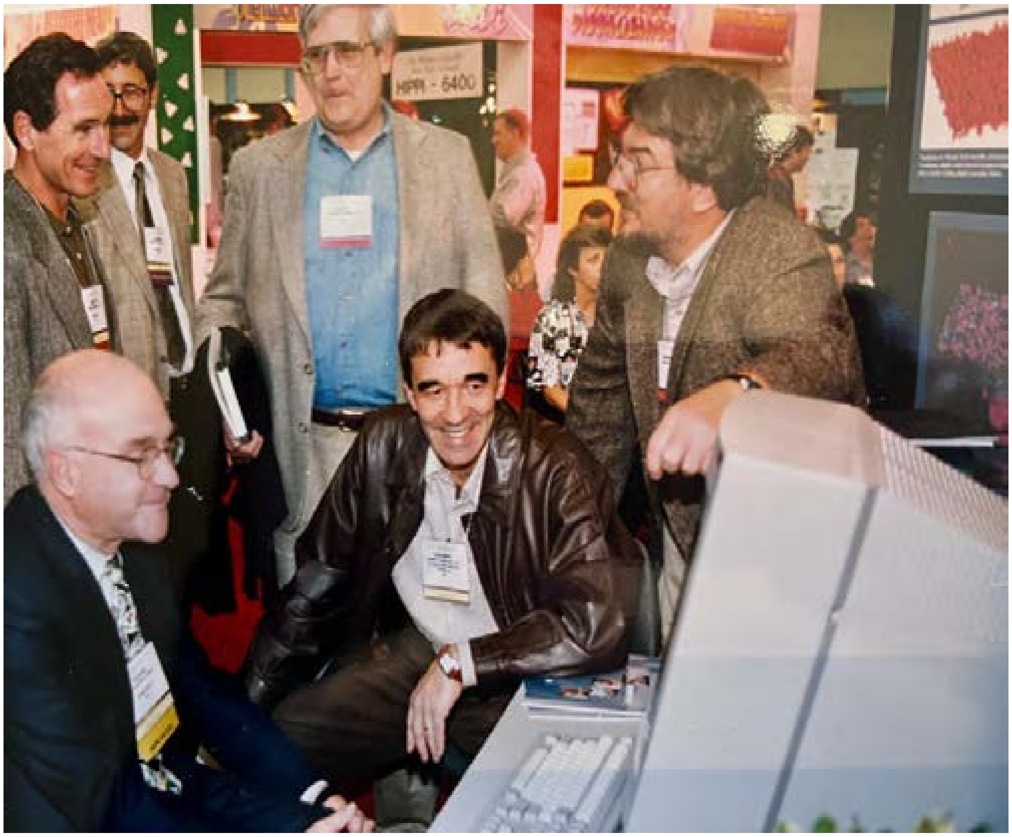

Figure 12. Clockwise from bottom left: Reis, McCoy, Art Mirin, Miller, Anastasio, and Nowak.

to NA-1 today, and an HQ force behind stockpile stewardship, of which ASCI was a part). Clockwise from Reis are Michel McCoy (head of the LC), Art Mirin (a CASC computational mathematician), George Miller (head of the Council for National Security), Mike Anastasio (LLNL AD of defense and nuclear technologies [DNT]), and, seated, Dave Nowak, the first LLNL ASCI executive.

To be candid, ASCI was an opportunity to grow funding rapidly until the 100-teraflop system was sited and in service. There was apprehension that interest in Congress and even at NNSA would wane under competing pressures. Future funding cuts, however, came off a very robust peak, guaranteeing a viable computing effort well beyond the decade allotted. Once funding started to flow, the challenge was to build up as fast as possible while delivering from day one.

3

ASCI BLUE PACIFIC

a. Launching the Blue Procurements
b. Early Facility Wake-Up Calls
c. Deliveries and Integration

ASCI Blue Pacific

a. Launching the Blue Procurements

After ASCI Red at SNL, LANL and LLNL hoped to site the Blue procurement system as the first fully ASCI-led competitive procurement. Gil Weigand secured sufficient funding to site fraternal twins, one at each laboratory, by means of a common request for proposal (RFP), with the two top bidders invited to negotiate a contract at one lab or the other. ASCI Blue Pacific was to be sited at LLNL, and ASCI Blue Mountain at LANL.

The top bidders were IBM and Silicon Graphics (SGI). Deciding which computer would go where was a problem. LANL asked for priority in choosing and indicated a preference for the SGI offer. At LLNL, the principals, including Nowak, Christensen, Dave Cooper (the Computation AD), George Miller, Mark Seager and Michel McCoy, held a short meeting. Miller asked, “Which solution do you prefer?” and the only strong vote for IBM came from Cooper, primarily based on his perception of IBM commitment and financial strength. The others had no strong preference—both systems were very high risk, with obvious shortcomings. Seager and McCoy were hoping for ICEG approval to use M&IC funds to forge a third path with Digital Equipment to create more future options for ASCI and M&IC. In the end, Cooper’s views were prescient, as LLNL did over $1B in business with IBM over the next 20 years. While the road was often rocky, IBM’s commitment and ability to deliver (financially and technically) were an enormous factor in LLNL’s long-term success. IBM’s seminal contributions in that partnership must be acknowledged.

Figure 13. Blue Pacific (1998): 3.8 teraflops peak performance.

As a part of the contract for ASCI Blue and with the support of Gil Weigand at HQ, Dave Nowak asked IBM to add an option for a follow-on platform at ~10 teraflops. This was an excellent strategic move, as it created an easy path for a subsequent system without requiring an additional competitive procurement and invoking the inevitable question of which lab would get it. Since Nowak understood that speed was of the essence to HQ, it was easy for Weigand to select LLNL for the follow-on system, namely, ASCI White.

b. Early Facility Wake-Up Calls

In the 90s, the LC was sited in B113, which featured a computer room used for decades by the weapons program. After the IBM Blue Pacific system went into contract, LC leaders—including Bruce Griffing (formerly NERSC), Barbara Atkinson, Seager, and McCoy—were shocked to get word from the facility manager, Hal Nida, that wedging the new system onto the floor would essentially max out the building’s power and cooling. Furthermore, there would be little electrical or mechanical redundancy in the building to

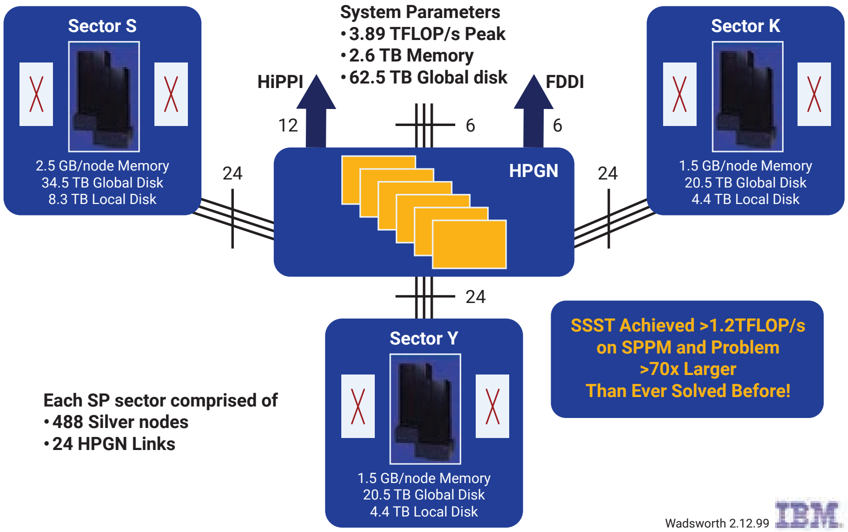

Figure 14. ASCI Blue Pacific system parameters.

support resilience. It is a bit unfair to say that the floor was primordial, but clearly its time had passed; and it was a terrible shock to McCoy and Seager, who had been blissfully unaware that they were boxed in. The problem could be solved in two ways: the easier was to find an interim solution for the next platform; the harder was to propose a permanent solution.

c. Deliveries and Integration

The contract specified that the Blue Pacific machine was to feature the most advanced IBM HPC node (codenamed Night Hawk 1), which was still under development. The timetable for delivery was very tight, given extreme pressure applied by ASCI HQ. Promises had been made to Congress and generous funding had been provided. Clearly, continuation of this generosity would be far easier to obtain if the program met its initial milestones on time and on budget. The linchpin was the computer, as the codes needed to be further developed on the target platform and then run at 1,000 processors to completion.

The first sign of trouble was IBM's struggle with the development of the node. It was soon obvious that there would be a slip. HQ was aware of the issue, but it would have been politically damaging to accept a slip so early in the program. As part of a technology-overview discussion with IBM, a more mature but less powerful node emerged as a candidate from the PowerPC line. The difficulty was that the processor was too slow to get to 3 teraflops using the largest IBM switch available. It was proposed to build three separate systems at 1.3 teraflops each and bind the system together using a high-performance gateway node (HPGN). One should take "high performance" on advisement: it would provide insufficient bandwidth to run weapons calculations across all three sectors. Thankfully, the program was not ready for full machine scale anyway. After all, the target for the first milestone was one thousand processors.

HQ found the alternative solution acceptable, but it was a very close thing. Had LLNL and IBM not cooperated intimately to identify a potential solution, ASCI could have taken a serious political hit at this early, vulnerable stage.

The ASCI Blue Pacific hardware that emerged was an IBM RS/6000 scalable parallel (SP) system using IBM PowerPC 604e processors. The system featured 1,464 nodes, each with four processors. IBM also delivered a separate, unclassified platform housing an additional 352 nodes containing 1,408 processors with 950 peak gigaflops processing speed, 524 gigabytes of memory, 20 terabytes of global disk, and 3.5 terabytes of local disk memory. This delivery

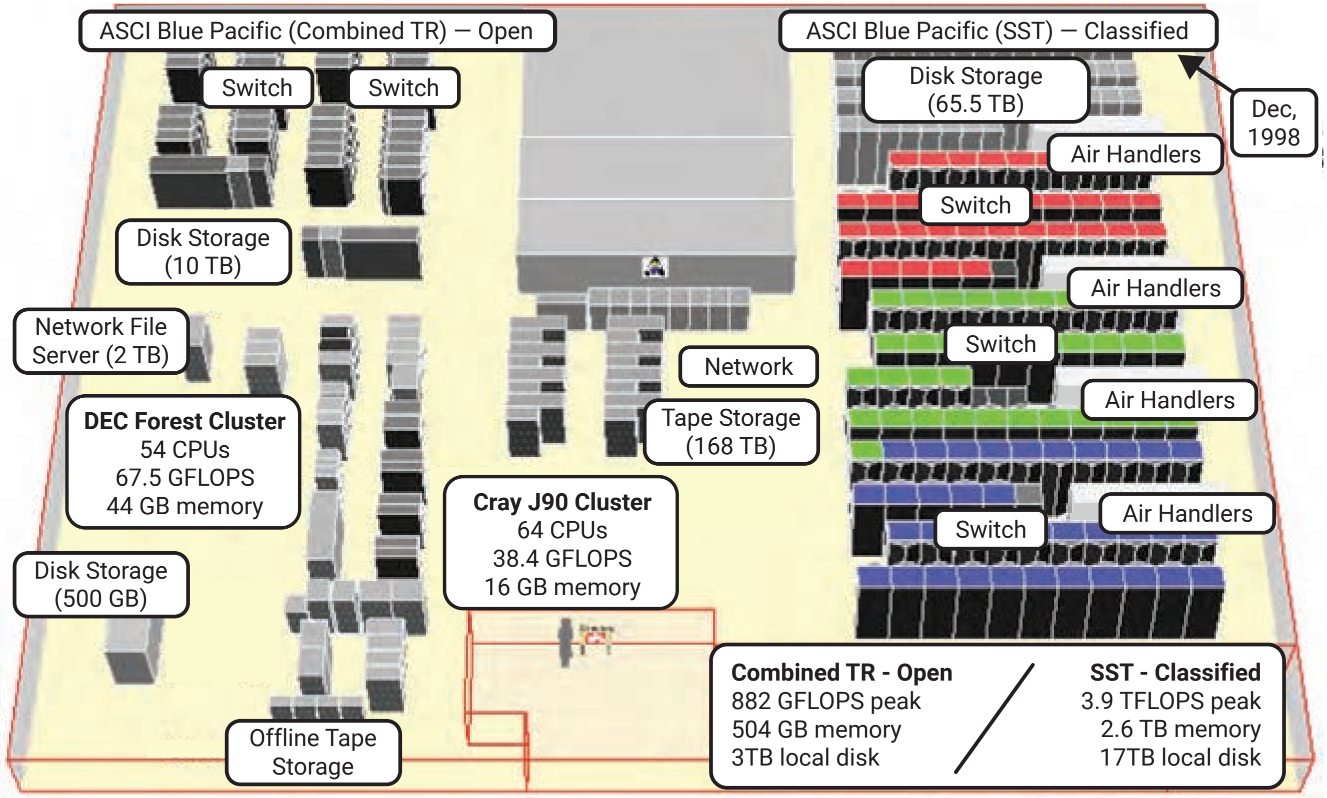

Figure 15. Drawing of the Sustained-Stewardship Teraflop (SST) system on the floor in B113 after delivery.

followed the successful LLNL strategy of acquiring classified and unclassified systems with comparable programming models, computer hardware, and software tools. Sticking with the blue theme, the unclassified system was named Sky, with individual sectors S, K, and Y.

Integration of the computing environment was both a vision and perennial goal. While the sophistication of integration improved with experience and technological advances, at any given time everything had to fit within some vision, or frankly it could be argued we didn't know what we were doing. Figure 16 depicts an early idea of the integrated environment, involving all systems on the floor, down to the working scientist.

This was to be the last time LLNL would use B113 for HPC. Ironically, while this system alternative was arguably a step back from the tightly integrated target in the original contract, the machine had a higher peak speed of 3.9 teraflops. The HPGN was able to provide sufficient bandwidth between the three sectors to run Linpack and achieve a respectable Top500 ranking of number two at 2.144 teraflops, behind the 2.379-teraflop SNL ASCI Red system. Given the bullet that LLNL had avoided, the Lab gladly took the standing. More important, the weapons program was able to use the system effectively to meet the first burn-code milestone.

In fact, every major system procured at the Lab had at least one near-death experience. What neutralized these threats and created opportunities for success was the tight technical partnership between LLNL and IBM and the willingness of both partners to acknowledge the other's pain, bend when the wind blew, and refuse to give up. Perseverance and personal ownership were key. Either we made it right or we failed—no one else would be blamed. IBM was a huge factor in these successes, because setbacks were expensive to rectify. But in the choice between failure or delivery, IBM always stepped up.

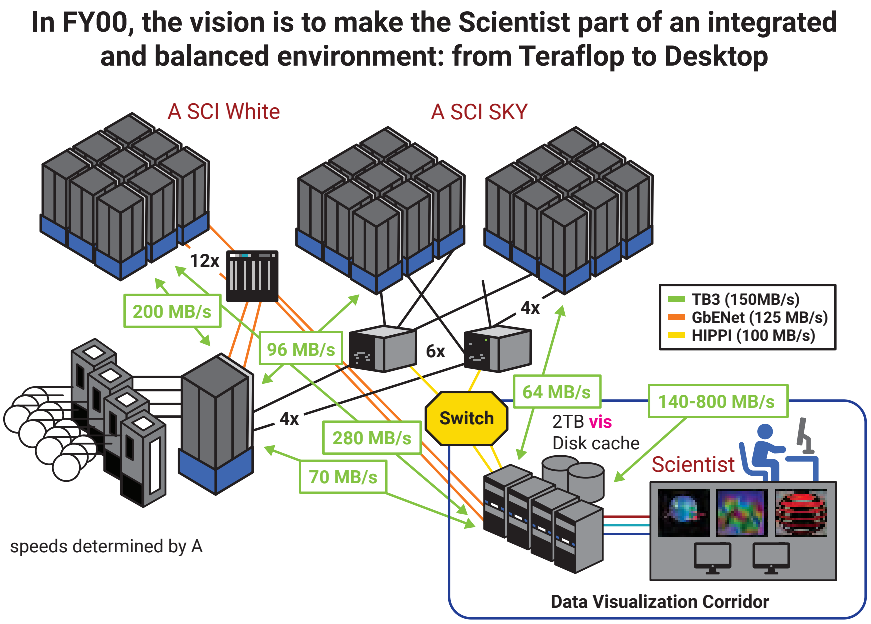

In FY00, the vision is to make the Scientist part of an integrated and balanced environment: from Teraflop to Desktop

Figure 16. Vision of an integrated computing environment.

4

APPLICATIONS DEVELOPMENT IN THE EARLY ASCI ERA

a. Burn-Code Reviews

b. M&IC and the Compass Cluster

Applications Development in the Early ASCI Era

The launch of the ASCI program placed emphasis on developing a new generation of simulation tools that were 3D capable and built to be massively parallel from the ground up. As noted in subsequent chapters, the term “massive” evolved with time.

In the summer of 1996, LLNL director Anastasio chartered a DNT simulation-code strategy committee to develop a coordinated weapons-simulation approach across organizations to meet stockpile-stewardship needs. The committee was chaired by Randy Christensen, with members Charles McMillan, Dale Nielsen, Peter Raboin, Tom Weaver, Charles Westbrook, and George Zimmerman. The multipart strategy codified LLNL’s applications approach in the ASCI era and included the following high-level recommendations:

1. Define a portfolio of new, next-generation simulation codes.
2. Meet the near-term needs of the Stockpile Stewardship and Management Program (SSMP) while new codes are developed.
3. Ensure that needs in key supporting areas, such as material properties, problem generation and visualization, and R&D are met.
4. Measure and reduce uncertainties in simulation codes.
5. Organize the Lab’s code-development effort.

With the delivery of ASCI Blue Pacific in 1998, the goal was to scale across a thousand processors. Two new code-development projects were launched at LLNL to meet the goals of ASCI: one in A Program, whose mission was to design and support secondary physics in a two-stage nuclear weapon, and the other in B Program, which designed and supported primary physics in a two-stage weapon. Both these projects included classified and unclassified simulation tools, but they took different computer-science approaches. They broke away from the Fortran language, a long-established staple of scientific programming, in favor of more modern languages. The B Program project chose C as its primary programming language, while the A project chose C++.

C was widely used in industry, but LLNL led the scientific community in its use for production-level scientific applications, beginning in 1987 with the CALE code written by Bob Tipton. Tipton was requested by the head of B Program code development, Randy Christensen (later an LC leader and subsequently ASCI deputy under Nowak), to develop a code for both the large mainframe supercomputers and Sun 3 workstations, enabling use of the workstations for problem generation. The notion of code portability introduced by LLNL would be a significant theme carried over to the ASCI program to accommodate both capability and capacity computers. The Cray supercomputers that made their appearance in the mid 70s were delivered as bare hardware; it was up to LANL and LLNL to put operating systems and language compilers on them. LLNL developed the Livermore Time-Sharing System (LTSS) with the Cray operating system and Civic as the Fortran compiler. By the early 80s, LLNL’s Fortran codes ran at 40–80 megaflops on the Crays by making extensive use of vectorization. The CDC Star 100, Cray’s primary competitor in the late 70s, was achieving less than 10 megaflops with similar codes.

LLNL’s success with its homegrown software stack on the Crays encouraged LANL’s adoption of LTSS (renamed CTSS) and Civic. By the mid 80s, the Lab developed a hybrid C compiler that included a Civic backend, essentially allowing C to perform as well as Fortran on the Cray, with CALE achieving roughly 60 megaflops. These achievements let B Program’s ASCI team feel reasonably confident in their choice of C as a performant language.

More conservative thought in B Program was reflected in other

Figure 17. Cray 1 supercomputer.

choices, such as straightforward extensions of proven algorithms to 3D. By contrast, A Program embraced a language that had not yet demonstrated performance comparable to Fortran or C, but offered considerably more flexibility for the future. Especially in the early years of the project, new features of the C++ language standard were not well supported and could be wildly inconsistent across compilers. Moreover, the algorithms chosen, while offering the potential for increased flexibility, had not been proven in any production-level simulation tool. Python was used in an infrastructure that tied the C++ physics packages together and allowed the user freedom to customize algorithms and interrogate data during a run. This provided flexibility well appreciated by many users, especially those experienced in the homegrown Yorick and BASIS languages implemented in George Zimmerman's Lasnex code. These were used by many ICF designers as a similar steering interface, data interrogator, and selective loader of packages to produce a program. The A Program code team wanted to move away from homegrown packages maintained by one or two developers, but the evolving state of C++ made it difficult to maintain Python interfaces that exposed the compiled C++ objects.

Another difference in the two approaches was reflected in personnel. B Program mixed new code developers with seasoned veterans respected in the design community. Gary Carlson was designated the first project lead of their ASCI effort, largely because of his experience leading the existing B Program production code, with Tom MacAbee co-leading the effort. Once the new project was well established, the reins were handed to Brian Pudliner, an up-and-coming code developer heavily involved in the development of new ASCI code. Partly owing to the cutting-edge approaches taken, A Program was staffed almost entirely with relatively new developers, including its original project leader, Doug Miller. Later, Mike Zika took over as the project was starting to be used in production. Notwithstanding Zika's efforts and significant leadership skills, a lack of veterans with established relationships among designers, along with handicaps in performance, slowed the design community's early adoption of the A Program's ASCI effort to replace legacy capabilities.

Another large code project, ALE3D, was also developed within B Program at the time. Started by Richard Sharp in 1987, ALE3D brought together the 3D capabilities of Dyna3D with the new arbitrary Lagrange–Eulerian (ALE) techniques pioneered by Bob Tipton in CALE. In the period between the end of UGT and the start of ASCI, ALE3D found new life as part of a CRADA between LLNL, SNL, and Alcoa Aluminum for metal-forming analysis. One goal of the CRADA was to develop a parallel version of the code, building on the success of Meiko and T3D and foreshadowing ASCI. That parallel version was prototyped by a student named Ed Luke and picked up by a new hire in 1994, Rob Neely, who oversaw a long transition from Fortran to C (and eventually C++). In the late 90s, the ALE3D team consisted of many Lab leaders in what is today the Strategic Deterrence (SD) directorate, among them Brad Wallin, Ivan Otero, Juliana Hsu, Scott Futral, and Rose McCallen. While the Alcoa CRADA largely fell to the wayside once ASCI started, ALE3D lived on and has become one of the most widely used codes from SD's Weapon Simulation and Computing (WSC) program, whether in or outside the Lab. ALE3D has a particularly large presence in the DOD, where it excels in conventional-weapon modeling, blast effects, and simulations that require advanced-material modeling.

Meanwhile, CASC began developing open-source software packages using combinations of programmatic, LDRD, and ASCR funding that eventually underpinned the major codes. The need for better implicit algorithmic capabilities that were not timestep constrained and could take advantage of massively parallel computers spurred Rob Falgout to begin developing Hypre—a library of linear solvers shareable among all simulation codes at the Lab. SAMRAI was also one of CASC's earliest projects, started by researchers Rich Hornung and Scott Kohn after LLNL's primary adaptive mesh confinement (AMR) team decamped to LBNL with the departure of NERSC.

a. Burn-Code Reviews

Gil Weigand of the ASCI program commissioned the creation of an external burn-code review committee to scrutinize LLNL and LANL efforts annually to measure progress toward ASCI goals. The committee, composed of knowledgeable academics and experts from the labs, was chaired by Kim Molvig from MIT. The first major milestone of the ASCI program was to demonstrate a 3D primary simulation of a nuclear weapon by the end of fiscal year (FY) 1999. At only four years from the establishment of the ASCI program, this was an extremely ambitious goal, given the typical 10 years to develop a new production code. The accelerated timeline was important for maintaining congressional support and served as a major challenge to Lab code-development resources. During this period, the development of new physics capabilities and model improvements supporting the weapons program was essentially halted in favor of developing new codes that initially would only replicate capabilities in existing production codes—except for moving to 3D and running much larger calculations that could span thousands of processors.

The annual reviews helped maintain pressure to produce results as quickly as possible, culminating in B Program in 1999, when the code team engaged in round-the-clock monitoring of weeks-long simulations to meet the ASCI milestone. The team was the only team among the two labs that succeeded, allowing the ASCI program to claim its first major success. It is important to note that while the 1999 milestone was achieved with most of the basic physics in place, the new ASCI code was still a long way from having all the physics capabilities of existing production codes. The team had much work to complete, including the replication of existing 2D capabilities. So much emphasis had been placed on the early success of 3D that 2D capability was put on the back burner.

Figure 18. First-ever primary simulation, B Division code team, December 1999. Front row (L to R): Becky Darlington, Janine Taylor, Brian Pudliner, Tom McAbee, B.I. Jun, Gary Carlson, Greg Greenman, Shawn Dawson, Jeff Grandy. Back row (L to R): Mike Collette, Jeremy Meredith, Bill Oliver, Chris Clouse, Grant Bazan, Ivan Otero, Tom Adams, Rich Procassini, Chris Hendrickson, Frank Graziani.

b. M&IC and the Compass Cluster

The M&IC center was born in FY97 with both programmatic co-investment and institutional support, as intimated in the 1996 "Report of the Working Group on Institutional Computing." This group was populated by distinguished scientists from all directorates. The broad goal of M&IC was to provide first-rate resources for production use to programs and the institution. There were thirteen members of the original ICEG, many of whom had a storied history at LLNL; two others went on to lead other DOE or NNSA laboratories. Their names are given below to emphasize the quality of governance that benefited M&IC throughout the early years.

- **Steve Ashby**—Computation
- **Eugene Brooks**—Physics and Space Technologies
- **Randy Christensen**—ASCI
- **Ron Cohen**—Energy Programs
- **Charles McMillan**—Defense and Nuclear Technologies (DNT), B Division
- **Doug Rotman**—Environmental Programs
- **Tomas Diaz de la Rubia**—Chemistry and Materials Science
- **John Fitzgerald**—University Relations Program (formerly a NERSC deputy)
- **Dave Hardin**—DNT, A Division
- **Steve Langer**—DNT, X Division
- **John Lindl**—Laser Programs
- **Rob Sharpe**—Engineering
- **Tom Slezak**—Biology and Biotechnology Research

Eugene Brooks, who authored the first report card to the Institution wrote,

> Members of the LLNL research staff who depend upon computing activities are finding that they get much more work done at the M&I center than they can get done on work group-funded resources, and the scope of the applications (execution time/memory size) is larger than can be run on group-operated resources.

The LC had made progress in halting the workstation diaspora.

What was the solution that earned encouraging remarks from Brooks, the man who coined the term "killer micros?" Recall that the two ASCI Blue solutions were from SGI and IBM. Seager, the LC architect, was concerned that there might be insufficient bidders for major procurements in the future and identified Digital Equipment Corporation (DEC) as a likely contender that should be kept in the game. LC consequently fielded a cluster composed of eight shared-memory multiprocessors, named the Compass cluster because the microprocessors were arranged in a rough circle. Compass consisted of 80 processors, 56 gigabytes of memory, and over a terabyte of disk. The peak performance was over 70 gigaflops in aggregate. The DEC Alpha microprocessor used on these symmetric multiprocessors (SMPs) was the fastest available. On average, the per-processor memory was roughly 700 megabytes, substantially larger than typically seen in workstations.

DEC later found itself in financial difficulties and was purchased by Compaq, which was subsequently absorbed into HP. LANL went on to procure the 30 teraflops-per-second ASCI Q system (originally bid by Compaq) from HP. This machine was subsequently limited to 20 teraflops for budgetary and technical reasons, but it can be said that Seager's persistence bore fruit by keeping another vendor in the game. This is just one early example of how M&IC activity provided positive feedback to ASCI.

5

ASCI WHITE (CY2000)

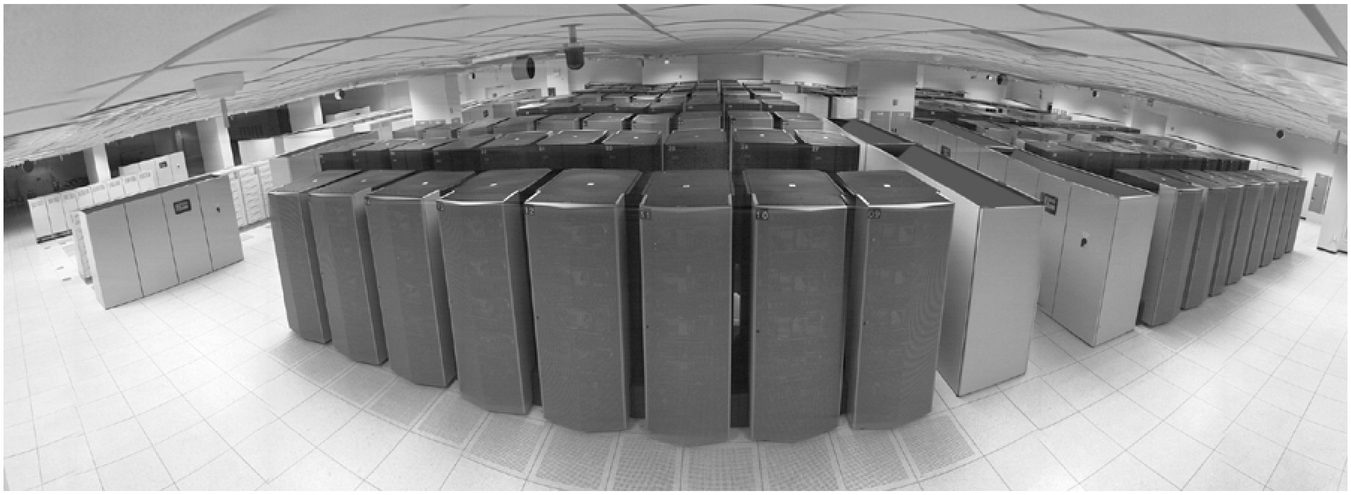

Figure 19. ASCI White (2001), featuring 12.3 teraflops peak performance.

ASCI White (CY2000)

As discussed previously, ASCI HQ authorized LLNL to exercise an option to deliver a 10-teraflops-per-second system to LLNL in 2000. The B113 facility was inadequate to house this large a system, so the old B451 NERSC facility was retrofitted with additional space, power, and cooling—a $15M investment.

While inadequate for the long term, this improvement was a lifesaving placeholder, as facility options at LLNL were extremely limited. As shown in Figure 21, the system nearly filled the 20,000 square foot machine room.

ASCI White went into classified service after its dedication on August 15, 2001. Architecturally, it was a computer cluster based on IBM's commercial RS/6000 SP node. It was a second generation of the node and was originally intended for Blue Pacific some years before, but did not meet the tight ASCI program schedule. White consisted of 512 interconnected nodes, with each node containing sixteen 375 MHz IBM POWER3-II processors. This totaled 8,192 processors for a peak speed of 12.288 teraflops per second (exceeding the 10 teraflops-per-second requirement). The system featured six terabytes of memory and 160 terabytes of disk storage and required three megawatts (MW) to power.

After initial installation, a long integration period followed to burn in the nodes and harden the software stack, making it suitable for production use. The inevitable hurdles that accompany first-of-a-kind machines struck from an ominous direction. The nodes would fail at random, making long-running calculations unfeasible. LLNL was unable, therefore, to accept the machine. The system, at around $100M, was on IBM's books awaiting resolution. On the other hand, LLNL and the ASCI program were losing precious time and were quietly fearful that the system would never be useful. It required weeks of deep forensics before IBM engineers identified faults in nodal interposers[12] that could randomly reveal themselves.

Figure 20. The 18-inch subfloor was raised to 30 inches and 8,000 square feet was added to the south, as shown.

---

[12] Interposers are compression-mounted interconnect devices that allow signals to pass quickly between boards or sockets.

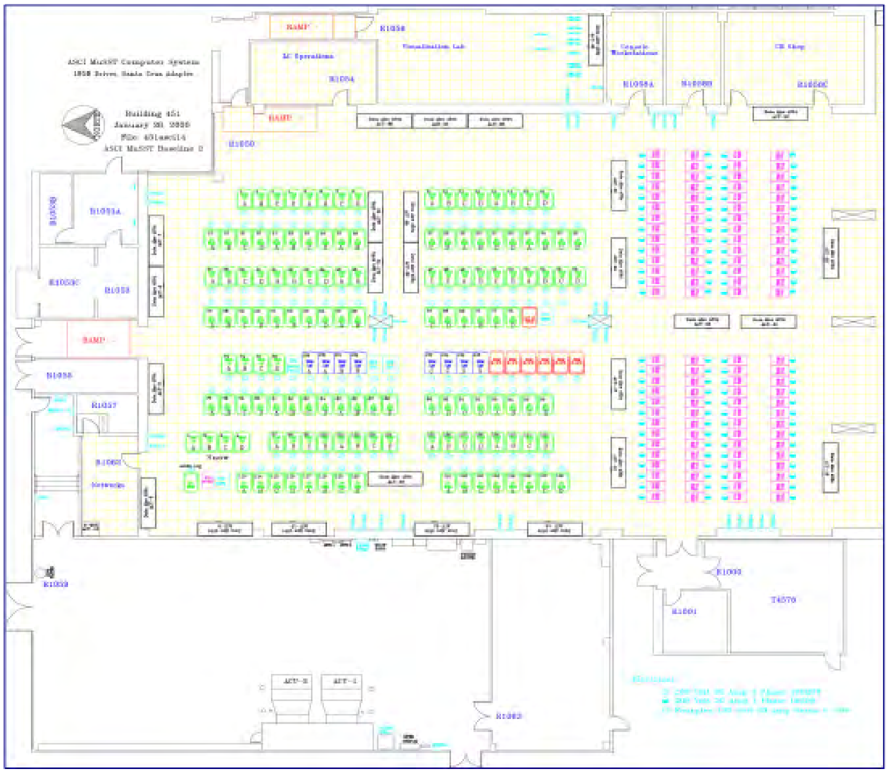

Figure 21. View of B451 machine-room floor layout.

LLNL and IBM made repeated attempts to find expeditious interposer solutions, but no quick fix guaranteed a fully reliable production machine. Finally, an executive vice president at IBM flew to LLNL to meet with director Bruce Tarter, the Computation AD, Dave Cooper, and ASCI-program leadership. After introductions in the director's meeting room on the fifth floor of B111, the IBM executive announced, "Now I will speak." He then relayed his decision to replace all the interposers across the entire machine—this time, with interposers that met the newly defined specifications. This was very far from a no-cost solution for IBM, as it meant flying out engineers to dismantle every node across the machine, which took weeks to complete. This was a defining moment in the LLNL/IBM relationship, representing as robust a commitment to integrity as to partnership.

In November 2000, ASCI White was number one in the Top500 list for the fastest supercomputer in the world, with 4.9 teraflops-per-second Linpack performance. Once fully assembled in June 2001, it reached 7.226 teraflops per second. At that time, 215 of the top 500 computers were from IBM. In June 2002, ASCI White was surpassed by the Japanese NEC Earth Simulator, a $600M vector system, which took the world by storm—a dinosaur swallowing a mammal. ASCI White was retired in 2006.

a. Progress in Weapons Applications and Science

ASCI White brought with it the next two major computing milestones in the ASCI program, again reviewed annually by the ASCI burn-code committee. The first was a 3D secondary simulation to be completed by the end of 2000. While the B Program code team had successfully completed the 1999 3D primary simulation, the A Program team was on the hook for the 2000 milestone. Many problems were encountered, including obstacles mentioned earlier, such as the immaturity of the algorithms employed and the heavy use of C++, with its relatively immature compiler. Message-passing interface (MPI) scaling on the machine was poor and continued to be a major focus and collaborative effort between IBM and LLNL; this wasn't fully resolved until the arrival of Purple in 2005. Given these issues, LLNL was not able to meet the milestone deadline. In the end, LANL completed the simulation in time to ensure ASCI continued to deliver as a national program.

Eventually, the A Program code fixed many of its C++ and algorithmic concerns. The production code has been used for over 20 years, and the high-risk technological aspects of the project resulted in many lasting contributions to the field of computational science. The use and production hardening of Python as a driver language, for example, was pioneered by this project, and staff contributed to the development of NumPy, which is used broadly. Within LLNL, this project ensured C++ compilers continued to advance and become production quality. Most major software-development projects rely on C++, and it has become the current language of choice within LLNL's ASC program. Finally, within A Program, computationally efficient algorithms and numerical methods were established for key areas of physics previously assumed intractable. These numerical methods became the default in LLNL programs.

The milestone for 2001 was intended as a culmination of the previous year's work, which was demonstrating a 3D full-system (primary and secondary) simulation. All LANL and LLNL code teams were to participate

ASCI White Integration Schedule

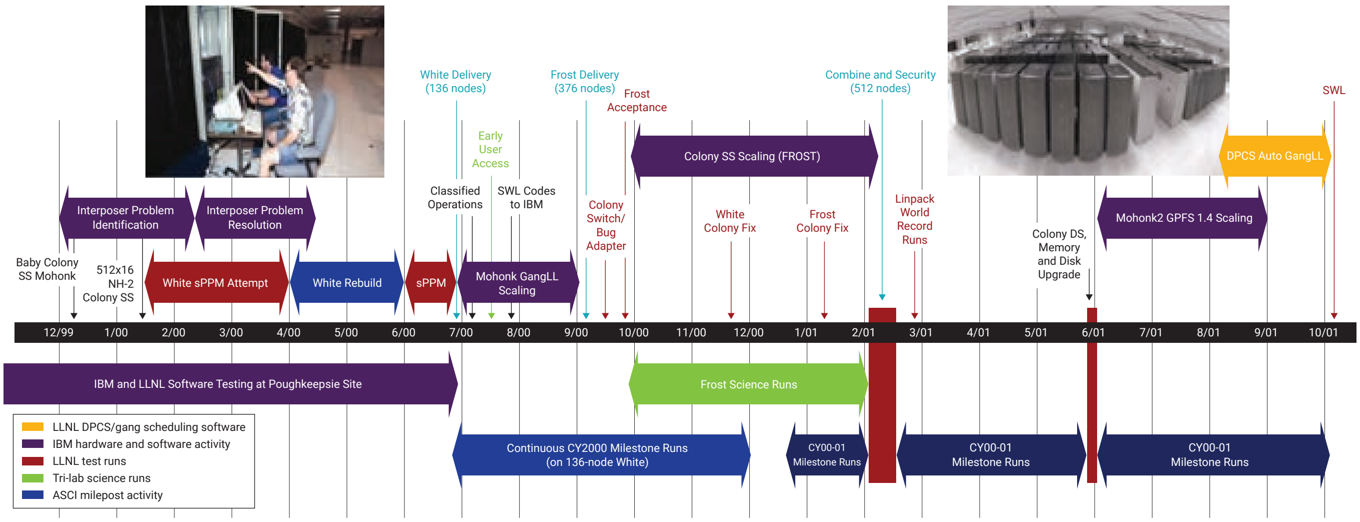

Figure 22. The complexity of ASCI White's system integration is suggested by this high-level schedule from a review at the time.

in the milestone. To get around the MPI scaling problem, the B Program code team adopted a threading strategy, running four threads per processor. They were the first team to achieve the 2001 ASCI milestone, running the simulation 39 days on 1024 processors. Implementing a threading strategy that only served its purpose on ASCI White was a Herculean effort spanning many months. But once IBM solved the MPI scaling problems, the threading approach was abandoned. In later years, when threading once again appeared to provide some advantage, a more efficient, coarse-grained threading approach was adopted.

b. Earth Simulator Disrupts the Firmament

A discussion of ASCI White would be incomplete without a rendition of the impact of the Japanese Earth Simulator (ES) on ASCI. White was a 12.3 teraflops-per-second peak system consuming about 3MW with a cost (for the platform itself) of less than $100M. For about a year it was the number one system on the Top500. In 2002, however, the Japanese sited the first in a series of climate-modeling computers, the NEC SX-6 vector-processing ES at 40 teraflops-per-second peak, achieving an extremely efficient 36.5 teraflops on the Linpack benchmark—almost five times faster than ASCI White.

The cost of this project was about $600M, including facility construction. The international response to the ES was very positive, and the keynote speech at the SC that year was given by the head of the project. The ES featured very expensive and high-power-consuming, non-commodity vector technology that ASCI had abandoned, believing a transition towards commodity processors was the only long-term, economically sustainable solution for HPC.

The U.S. political blowback to the ES approached hysteria. ASCI was criticized for having failed to assure American preeminence in computing after the huge investment in national-security computing. The damage was so significant that a Senate staffer came to LLNL for a day to investigate, grilling the managers (McCoy and Seager in particular) about the failure. Obvious and compelling explanations (at least, to the Lab) about cost performance and energy efficiency were largely ignored; nothing seemed to damp the umbrage in Washington and the public at large. To make matters worse, some leaders of other government programs argued they could do computer leadership better than ASCI. It was a purgatorial period. ASCI faced a political crisis, and the only resolution could be through a political response.

This came less than two years later with the first quadrant (16 racks) of the 360 teraflops/second Blue Gene/L computer. In November 2004, a 16-rack system (each rack with 1,024 compute nodes) took first place in the Top500 list, with a LINPACK performance of 70.72 teraflops per second. This was about double the performance of ES and involved only a quarter of the final machine. Hence, this number was essentially achieved by an approximately $20M system. Blue Gene/L was designed with simple "light" processors best suited to computations that could be spread out over much of the machine without requiring significant communication, like molecular-dynamics calculations in physics and biology. The idea was to study the architecture and eventually improve and expand the architecture's application reach, as LLNL had recognized early that low power consumption would be key in the future.

In short, Blue Gene/L did not have the application breadth of the ES, which was built to run multi-package physics calculations. While this should have been relevant to critics, politically it didn't matter. Superficiality had worked against LLNL originally but was now working on the Lab's behalf. This lesson—that politics, even if based on factoids, can make or break a valid strategy—was a lesson ASCI management at LLNL never forgot.

The criticism subsided, but there was permanent scarring in terms of budgets. The irony was that Blue Gene/L had been hanging in the balance under budget pressures, but survived thanks to the vision and courage of the ASC HQ program leader, Dimitri Kusnezov. Kusnezov had faith LLNL could deliver a knockout blow while exploring a potentially fertile path for national-security computing.

6

UNIX AND LINUX STRATEGY FOR CAPACITY COMPUTING

Unix and Linux Strategy for Capacity Computing

ASCI's vision orbited around building 3D codes to model full nuclear systems. It did not focus overly on less-intensive stewardship computing needs and certainly was not fixated on the unclassified needs of a multidisciplinary laboratory with multiple programs. The difficulty was that the bulk of the work going on at the labs was 2D in nature, not requiring the full resources of a capability computer. Yet packing smaller runs on a capability computer undermined the expeditious scheduling of larger jobs. In short, the weapons and unclassified research programs needed well-planned access to both capacity and capability. Given this dilemma, would it not be possible to procure a nimbler and cheaper solution for capacity needs?

There was a second conundrum. LC had been looking for a risk-mitigating and liberating strategy to avoid vendor-proprietary software and hardware solutions. Vendors frequently rejected LC's proposed operating system (OS) enhancement requests or were slow to deliver bug fixes. In addition, LC's in-house software-developer expertise was wasted working around limitations in the vendor's OS. In 1999, Charlie McMillan (then the B Program code-group leader, later director of LANL) asked for a meeting with LC management. He had been exploring Linux on his desktop and suggested aggressively investigating the potential of open source to reduce costs and enhance efficiency at the LC. This was well-received, and the McMillan meeting put commodity-cluster, open-source exploration on a high-priority path.

What followed was the development of a comprehensive scalable-systems Linux cluster strategy. This included building clusters from commodity components, designing hardware and software for easy manageability, and leveraging the open-source software model to the extent possible. This approach enabled the LC to supplement the base Linux distribution with in-house development expertise and vendor partnerships, allowing the development of best-of-breed software. Advances included a robust, scalable cluster management toolset, an efficient, scalable resource manager (SLURM) to exploit maximum utilization of resources, the Lustre parallel file system, and a high-performance interconnect (initially Quadrics Elan3). The LC developed a multitiered software-support model (see Figure 23) that included LC system administrators, onsite developers, vendor partners, and the open-source community.

All of this didn't self-coordinate: it required forward-looking and insightful management by the LC Linux group, which, for instance, leveraged years of investment in the ASCI PathForward program (2000). PathForward "consists of multiple partnerships with computer companies to develop and accelerate technologies that are expected to either not be in the current business plans of computer manufacturers or not be available in either the time frame or the scale required by ASCI." This was manna from heaven for the computing centers at each laboratory. Each brought forth ideas for investments, which were vigorously debated at ASCI Tri-lab meetings, because there were more ideas (good and bad) than money. Eventually, the labs focused on a subset of investments that all (or at least, most) found acceptable. Among noteworthy accomplishments were improvements in the Quadrics switch in the 2001 timeframe, which later fed into major ASCI platforms, clusters, and investments in the Lustre open-source file system, greatly helping assure the viability of large clusters.

The capacity-hardware solution began with the concept of the scalable unit, originally implemented at SNL and borrowed by LC: namely, the design of a complete computing system that could be integrated by plugin into additional scalable units to build larger systems, like building with Legos.

Open-source software infrastructure developed at LC has a widespread impact on global HPC

- SLURM - resource manager originated at LLNL and is currently deployed on most Top500 machines
- TOSS - Tri-Lab Operating System, led by LLNL, provides a common OS for all NNSA commodity computers
- HPSS - High Performance Storage System was the first hierarchical storage management capability, now used at over 40 global HPC centers
- SPACK - manager for package builds and installs on HPC systems

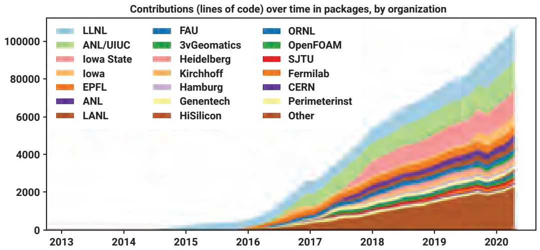

Figure 23. Some examples of LLNL's contributions to the open-source community.

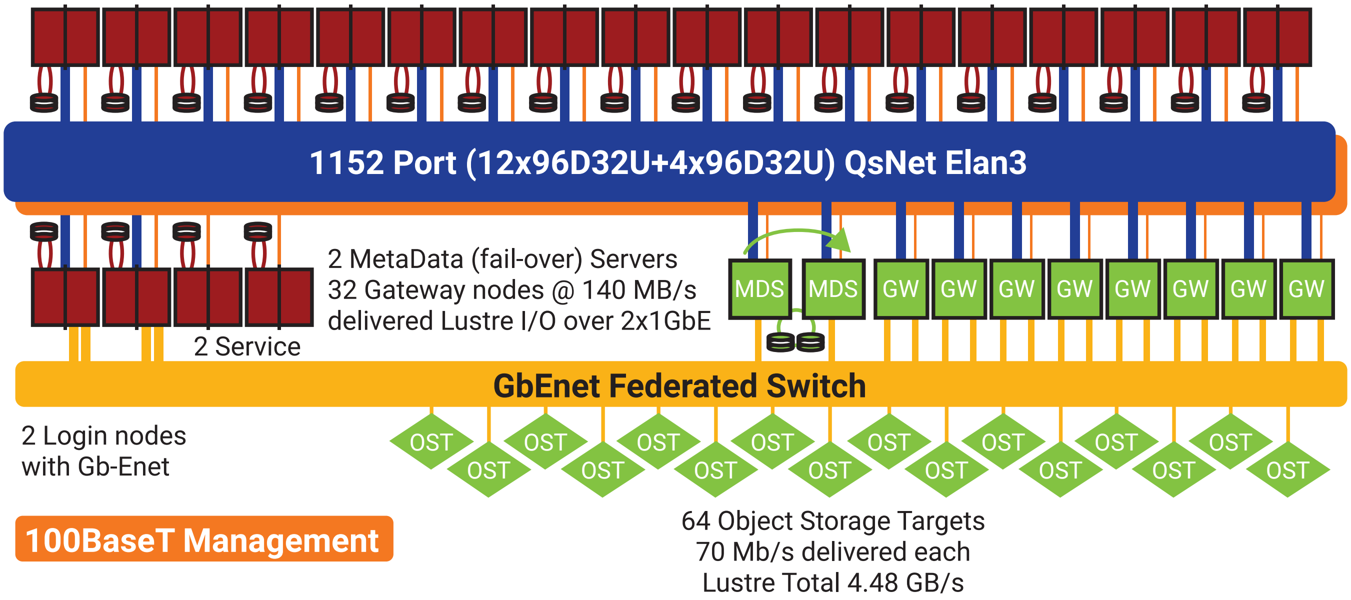

Figure 24. An example of the Multiprogrammatic Capability Resource (MCR), a system sited for M&IC in 2002.

To achieve all this required the leverage of M&IC and ASCI at LLNL. Building software toolsets was expensive and well outside the purview of M&IC. ASCI, with its core foundation in capability, was not funding capacity clusters, though it was the natural place to develop appropriate software solutions. LC leveraged ASCI PathForward[13] investments with vendors, such as the Lustre file system, as well as ASCI funding for the local R&D of essential toolsets like SLURM. With the software funded, the LC used M&IC funding to build unclassified systems at scale to convince skeptical audiences. In short, M&IC leveraged ASCI R&D effectively to build highly functional, low-cost computational systems and provided ASCI with real data about potential future investments. A serious and recognized risk was disappointing M&IC customers with undeveloped service, but such was the close attention at the LC that this risk was usually recognized and dealt with early—which is not to say that improvements were seamless. There were times, as in the deployment of Lustre, that the LC had to retreat from its most aggressive multisystem ambitions.

A lengthy sequence of primarily institutional investments was prosecuted, each larger and more sophisticated than its predecessor. Three were of great importance:

- 2002: the Multiprogrammatic Capability Resource (MCR), a mammoth 1152-node cluster composed of dual-processor, 2.4 gigahertz Intel Xeons. The MCR's configuration included the first production implementation of the Lustre parallel file system. It achieved number three on the Top500 list in June 2003, a stunning result at the time.
- 2003: The M&IC Thunder Linux cluster, a product of California Digital, Quadrics, and Intel. Thunder was a powerful Linux supercomputer—a 4,096 Itanium2 processor-based cluster delivering 19.94 teraflops of sustained performance made it the most powerful in North America. It was also the largest Itanium 2-processor deployment and the largest implementation of Quadrics's low-latency interconnect technology. These technologies allowed Thunder to achieve record cluster efficiency of 86.9 percent.
- 2006: Appro's Peloton Opteron/Infiniband Linux cluster procurement in June. Peloton clusters were built in 5.5-teraflop scalable units (SUs) of approximately 144 nodes using AMD dual-core Opterons with eight central processing units (CPUs) per node. The six Peloton systems represented a mix of resources for ASC and M&IC capability and capacity.

The architecture of Thunder is depicted in Figures 25. It was mammoth for its time and generated a huge peak at very low cost while providing functionality to run integrated physics codes. It was nearing time for a reckoning

---

[13] "ASCI PathForward: To 30 TFLOPS and Beyond", https://doi.org/10.1109/4434.678783

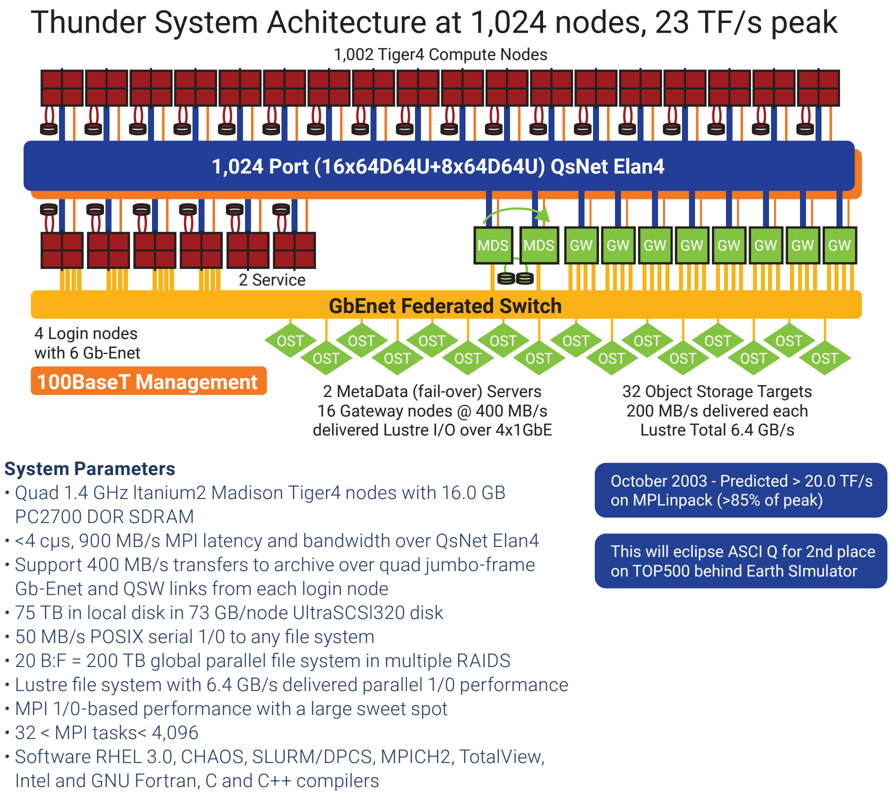

System Parameters

- Quad 1.4 GHz Itanium2 Madison Tiger4 nodes with 16.0 GB PC2700 DOR SDRAM
- <4 cμs, 900 MB/s MPI latency and bandwidth over QsNet Elan4
- Support 400 MB/s transfers to archive over quad jumbo-frame Gb-Enet and QSW links from each login node
- 75 TB in local disk in 73 GB/node UltraSCSl320 disk
- 50 MB/s POSIX serial 1/0 to any file system
- 20 B:F = 200 TB global parallel file system in multiple RAIDS
- Lustre file system with 6.4 GB/s delivered parallel 1/0 performance
- MPI 1/0-based performance with a large sweet spot
- 32 < MPI tasks< 4,096
- Software RHEL 3.0, CHAOS, SLURM/DPCS, MPICH2, TotalView, Intel and GNU Fortran, C and C++ compilers

October 2003 - Predicted > 20.0 TF/s on MPLinpack (>85% of peak)

This will eclipse ASCI Q for 2nd place on TOP500 behind Earth SImulator

Figure 25. Diagram of Thunder system architecture with parameters, from a 2003 LC presentation.

and an adjustment to the ASC platform model.

ASC HQ leader Bob Meisner, impressed by the successes of Thunder and even more so by the turnkey Peloton systems, dropped his concerns about procuring other than capability systems. The argument was obvious: use cheaper systems for capacity to allow capability calculations to breathe on systems designed for such use. This argument conveyed to Congress that ASC was still focused on capability computing and was simply looking for greater efficiency. In 2007, HQ decided to invest in Tri-Lab Commodity Clusters, or TLCC, later known as Commodity Technology Systems (CTS). This became a separate ASCI hardware-investment line. Deviation for the sake of efficiency and innovation can be dangerous, but the messaging to Congress via Meisner was managed effectively.

HQ selected the LC CHAOS software stack for TLCC and CTS. Deemed the most evolved and tested across the Tri-lab, CHAOS was renamed TOSS (for Tri-lab Operating System). From this point, the three labs worked together on RFPs for procurements, the procurements themselves, and the software repository, while much testing and development continued at LLNL.

The contribution of the LC's Matt Leininger, who guided the transition from TLCC to CTS on behalf of ASC, was enormous. CTS procurements had increased technical and contracting risks because these procurements offered more architectural options to serve multiple NNSA organizations, complicating system-software processes and adding to the contract-management burden. Also added was the risk of embarrassing failures that would be visible near and far. With Matt taking the lead, LLNL's technical and contracting team designed an improved contracting approach to manage the increased technical risks.

Matt's vision and influence were seen also in the Hyperion project, in which vendors and the LC coinvested in experimental systems and shared access. As a leader in driving the scalability of Linux clusters, LLNL was obligated to participate and coinvest in the development and testing of key capabilities at scale. Hyperion was built to address this need. Matt was integral to the development and management of Hyperion's vendor partnerships and in coupling these efforts to LLNL's R&D capabilities. Selected outcomes were higher-performing commodity interconnects, Lustre at-scale testing, and early evaluation of data-intensive and cloud architectures. Hyperion and a follow-on project using the so-called Catalyst system revolutionized cluster computing in fundamental ways by providing critical software and hardware components for a highly scalable simulation and data-intensive environment.

This was obviously a complex and challenging evolutionary process. The extremely tight coordination within the LC and between M&IC and ASCI was a high card during this period—a force multiplier for LLNL and ultimately a great benefit to the Tri-lab community. We are reminded that initiative at the labs, which may often skirt the margins of orthodoxy, has greatly contributed to the success of the labs and HQ goals for national security. It was a combination of vision, internal coordination, and benign, robust funding, both from LLNL (the institution) and from ASCI at HQ, that made this leap possible.

7

FACILITY CHALLENGES AND SOLUTIONS

Facility Challenges and Solutions

The homing of the ASCI Blue system in B113 was the end of the line for the facility. With ASCI White in planning, it was imperative to identify a new site in existing Laboratory space. Building 451—the old NERSC building—was the best candidate and was duly upgraded. The real problem was accommodating whatever might come next.

At the time, little was understood about the mammoth power and cooling requirements of the future. Prudently, the LC insisted on flexibility and expandability, especially since it had no clear understanding of future needs.

One afternoon in 1997, Seager and McCoy stood at a whiteboard in B113 and scoped out the bare-bones requirements for a long-term facility, using Seager's wristwatch calculator to crudely extrapolate power and space requirements:

- High levels of power and cooling available as needed were required to support up to two systems with power requirements equivalent to of two 2004-class 100-teraflop systems—estimated that day to be around 15 MW.
- A guarantee was needed that any ASCI-scale system would have sufficient high-quality floor space and high probability that two systems with requirements like those of two 2003/4 100 teraflops-per-second computers could be sited simultaneously. This required about 50,000 square feet of unobstructed machine-area space (with at least a four-foot underfloor). In Seager's famous words, "Pillars are a tax on the program."

Expandability was key, but not explicitly in the requirements. The LC wanted to avoid any suggestion that it was designing to a limited budget (less than $100M) to avoid additional governmental scrutiny, which

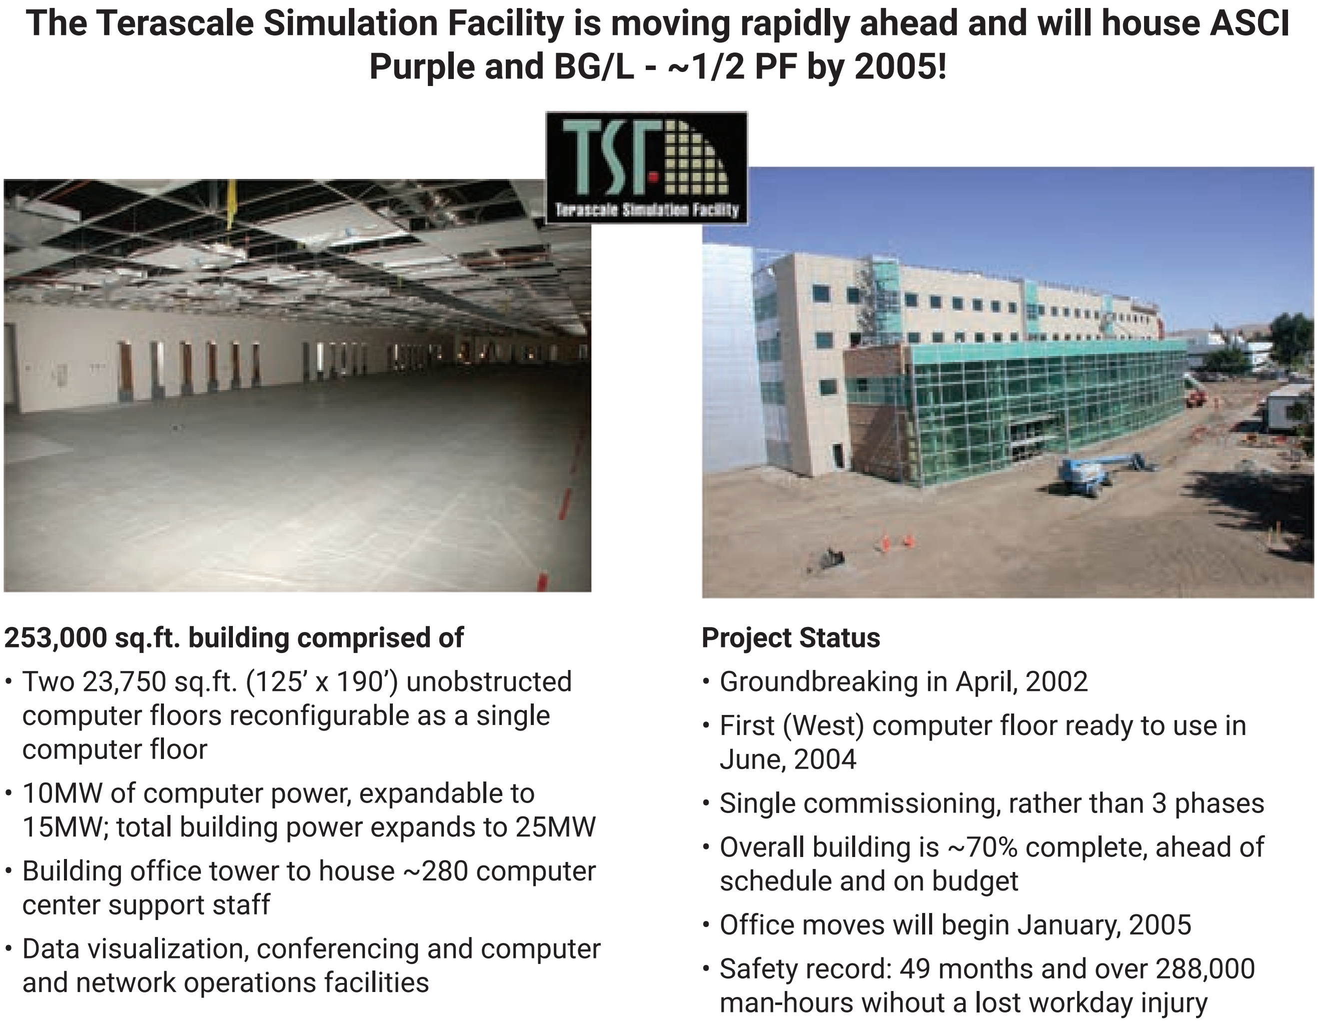

The Terascale Simulation Facility is moving rapidly ahead and will house ASCI Purple and BG/L - ~1/2 PF by 2005!

253,000 sq.ft. building comprised of

- Two 23,750 sq.ft. (125' x 190') unobstructed computer floors reconfigurable as a single computer floor
- 10MW of computer power, expandable to 15MW; total building power expands to 25MW
- Building office tower to house ~280 computer center support staff
- Data visualization, conferencing and computer and network operations facilities

Project Status

- Groundbreaking in April, 2002
- First (West) computer floor ready to use in June, 2004
- Single commissioning, rather than 3 phases
- Overall building is ~70% complete, ahead of schedule and on budget
- Office moves will begin January, 2005
- Safety record: 49 months and over 288,000 man-hours wihout a lost workday injury

Figure 26. Viewgraph showing progress made for the new B453 facility, circa 2004.

could add delays and potentially scuttle the effort. Nonetheless, LC management understood the need for expandability for future chilled water, floor space, electrical power, and mechanical infrastructure, including water cooling at some point. This meant the site could not be hemmed in by other buildings or strangled by buried utilities. Detailed exploration of the Laboratory for suitable sites was frustrating and lengthy, as NNSA required a thorough survey to rule out any existing facility that could be modified at lower cost than building from scratch. The LC wanted a solution that would be relevant in 40 years, not an expeditious and ultimately futile kludge generating an obsolete facility within 5–10 years. The process was respected while unyielding persistence won the day. The B453 site was identified for new construction—it was near enough to B451 (the old NERSC building and home of CASC and ASCI White) to encourage a campus atmosphere, but far enough from the gravitational center of the Lab (e.g., B111, B113, and B132) to avoid buried infrastructure that made the addition of electrical and mechanical substations and pipelines very difficult. The nearby open parking lot to the west of B451 would also be ideal for an electrical or mechanical substation if ever required.

The first computer room was to be available in August 2002, well in time for the 100-teraflop system in 2004. The LC went for the bleachers in terms of space, knowing that adding floorspace later would be extremely difficult, especially with systems running on the floor. As Seager joked, "We follow More's Law—more is better."

In every aspect of reviews, internal design, and construction, Barbara Atkinson of the LC took the lead, aided immeasurably by Anna Maria Bailey's visionary mechanical and electrical expertise. That the building has been so successful for so long and remains current today was due in the greatest part to their insights and communication skills regarding requirements. In addition, Sam Brinker, assigned by the local NNSA office to provide oversight during construction, helped immensely in overcoming technical and bureaucratic challenges, including an HQ-mandated floorspace downgrade of about 8,000 square feet, bringing it to 48,000 square feet. In the years ahead, the critical need was for power and cooling, but this was unknowable in 1997.


Figure 27. Building 453.

A viewgraph from a presentation to the institution is in Figure 26. The sense of excitement at the time comes through even today.

The LC's steadfast insistence on expandability proved prescient. In preparation for exascale, B453 was upgraded to 85 MW, with commensurate cooling made possible by an expansive utility yard where a parking lot used to be (see "Exascale Computing Facility Modernization" section below). Today, B453 is perhaps the most capable facility in the DOE complex.

If there was anything to regret, it was the project name, "Terascale Simulation Facility," which unfortunately stuck as the building's informal name. In pulling together an initial request to HQ, McCoy had suggested this name for the project because it felt expeditious and consistent with ASCI ambitions to go to 100 teraflops, and he hoped it would loosen the purse strings of the gatekeeper, Gil Weigand. The name provided a source of unending embarrassment, however, and as McCoy endured jokes from colleagues as the world blew past terascale, the name was quietly dropped. Today it is merely "B453." If this was the worst outcome of the project, the LC was very lucky. As Bruce Goodwin was wont to say when things went better than expected, "It is better to be lucky than good."

8

100-TERAFLOP PURPLE AND BLUE GENE/L (BG/L) CHANGE THE WAY SCIENCE IS DONE

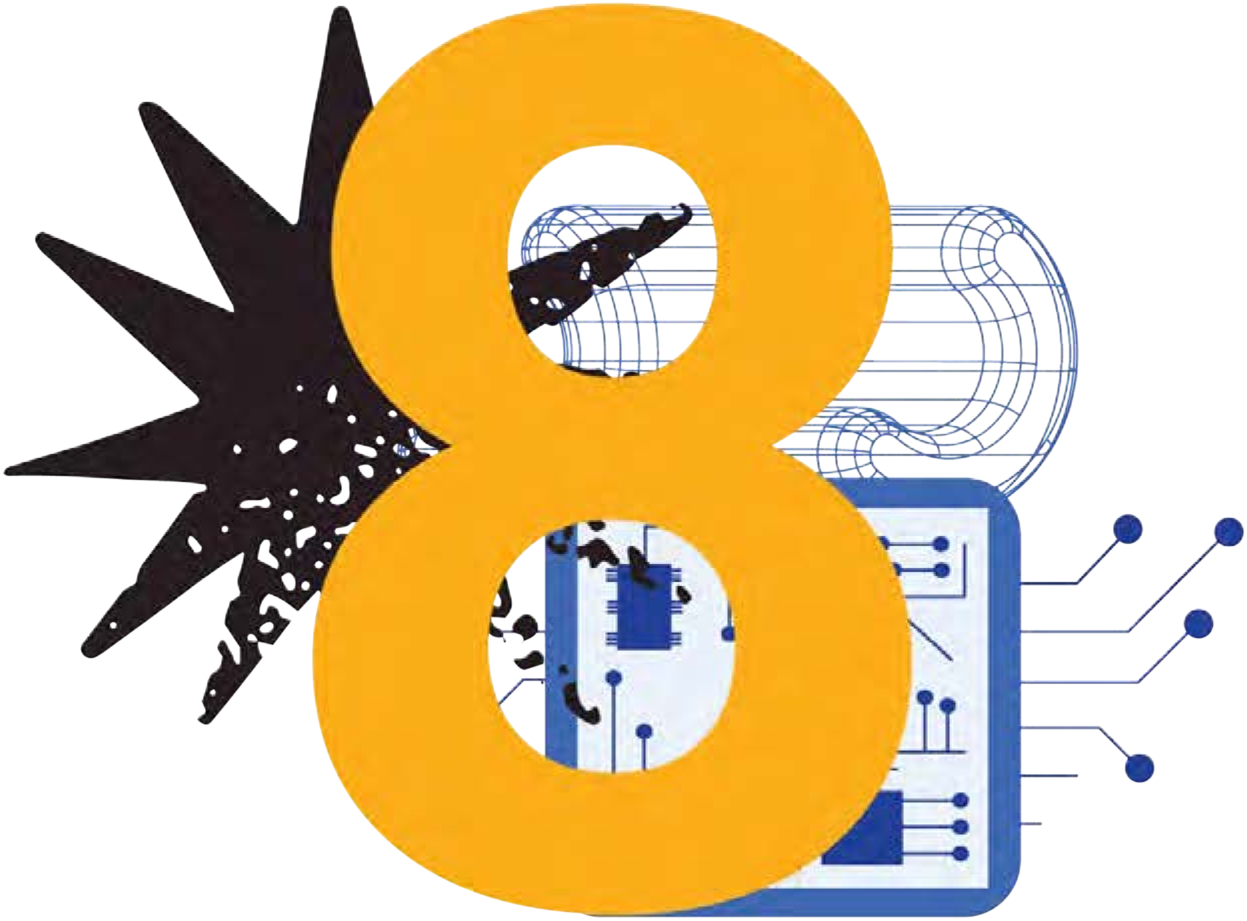

100-Teraflop Purple and Blue Gene/L (BG/L) Change the Way Science is Done

a. Purple RFP and the Genesis of BG/L

Even today it is a challenge to think of another procurement with the lasting reach of the Purple procurement. This process produced two major systems now sited on B453 floors: the ~100-teraflop Purple computer and the 360 teraflops BG/L system. For these acquisitions, the LC invented a procurement model for first-of-a-kind, advanced systems that is widely used today across the DOE. The process provided for two very different platforms that changed the way NNSA organized its weapons-science program and provided compelling evidence that simulation was changing the nature of scientific discovery. In simultaneously siting two orthogonally different systems of enormous capability, the LC accomplished a tour de force in operations that set a high standard of professionalism.

The Purple procurement was to produce the last system win in the ASCI march towards the 100-teraflop platform for running full-system 3D simulations of a nuclear device at entry-level resolutions. This was to be a proof of principle that computers could use, as they were refined, to enable certification of the U.S. nuclear stockpile for the foreseeable future without new UGTs. The initial budget at $290M was set by HQ through a Moore's Law extrapolation from the cost of the 30-teraflop LANL procurement. Although this budget was based on just one data point, there were no other defendable approaches for estimating the cost of a new platform well ahead of its procurement.

While siting White and planning Purple, LC management was increasingly concerned that power consumption could be a limiting factor for future systems and was exploring alternative technologies at lower power-consumption levels. At that time, accelerators, while common today, were immature and lacked the capability to run 64-bit calculations routinely.

In 2001, LC systems architect Mark Seager and the first ASCI leader at LLNL, Dave Nowak, returned from a visit to IBM in New York. Seager argued that a simulation machine could be built from a generalization of the IBM massively parallel cellular-processor architecture designed for protein-folding simulations, a biological process. IBM's Alan Gara had been thinking about an extension of the design to a more general-purpose computer, and this caught Nowak's attention. Seager was soon a major advocate and influence. What came from this was a modest R&D award to IBM from NNSA. The subsequent design was sufficiently provocative that the LC envisioned building a large, unprecedented machine based on this work. The problem was twofold:

(1) Where would the money for this system come from, and what would be the procurement mechanism? IBM obviously would not refine a design and build a nonproduction system without a substantial procurement.

2) Would NNSA risk authorizing the procurement of a large system with an application reach that might have a very limited effect on weapon assessment?

To answer the first question, LC conducted a study of costs about a year after the original budget was proposed and suspected a ~$60M slice could be taken from the $290M without affecting the 100-teraflop goal. Taking a page out of Nowak's playbook, McCoy suggested that a sentence be added to the RFP package instructing the bidders that if the vendor had other compelling technologies to offer in addition to their 100-teraflop bid, they were invited to do so. This was briefed beforehand to Dimitri Kusnezov at HQ, who was supportive.

The second question could be answered only by detailed studies of what this architecture could do for weapons science. This set in motion studies demonstrating that the machine could be a powerful, low-cost, materials-modeling molecular-dynamics (MD) platform that would offload science work from the 100-teraflop platform. Materials-science (MS) modeling had often taken a backseat to weapons design, so this argument scratched an itch.

The machine was novel and risky, and no such weighty decision could be made in a political vacuum. All stakeholders had to weigh in, including those representing potentially competing interests at other DOE labs. Because failure would reflect on NNSA HQ, not just LLNL, leadership and support were needed at high levels in Washington. As an example of potential headwind, it was obvious that the machine (to be named Blue Gene/L, with "L" for "light") would eclipse the Linpack score of the ES supercomputers deployed at Yokohama Institute of Earth Sciences. How would those hoping to build their computing budgets by capitalizing on a politically expeditious tailwind—generated by the perceived failure of ASCI to maintain leadership—react to moving forward and thereby calming the waters? Even so, the reviews went well, chiefly because the preparation had been so comprehensive it was impossible to dismiss. Without the strong support of the ASCI office under Kusnezov, however, the idea would have remained just that—another good idea among so many others on the scrap heap of aspirations dispatched by political headwinds or budgetary stress.

The LLNL machine was the first example of ASC advanced-architecture procurements. It had been assumed that while machines would become very large, their designs would be conventional or, more precisely, attached to the vendor's product line, even if improved for ASCI needs through additional R&D funding. The energy-consumption crunch and the path first lit by BG/L informed HQ that it would be necessary to countenance riskier machines to forge an affordable future. Risk-taking decisions at LLNL combined with support from HQ made a lasting difference.

While the direct impact on programmatic use cases was not large for BG/L, it proved a real wake-up call for our programmatic codes, uncovering numerous algorithmic issues with extreme (~million processor) MPI scaling. The small amount of memory per processor (0.5 gigabyte initially) also proved a major challenge and forced a complete rethinking of how arrays are stored and how the large, third-party libraries that provided much of the fundamental data are stored and accessed. Cooperative sharing of memory between libraries and host-code processes had to be implemented, and the distribution of data across processors was needed in many cases. Tackling these challenges under BG/L helped prepare the LLNL code teams for the arrival of Sequoia in 2012, a third-generation BG architecture. BG/L was also an opportunity for basic science codes with simpler parallelization strategies to demonstrate the power of BG/L in running simulations of unheard-of size, providing unique insights into the basic properties of matter.

b. The 100-Teraflop ASC Purple Contract

Purple was targeted to deploy a 100-teraflop system as a capability platform. Several conditions made this deployment an extraordinarily high-risk, high-reward undertaking: very aggressive schedules, highly complex technology development, and a volatile budgetary cycle. As discussed earlier, the budget of $290M proved sufficient to acquire and deploy both the 100-teraflop system and an unproved but potentially game-changing 360-teraflop BG/L technology featuring none of the safeguards of a commercial standard product (e.g., a maintenance and service structure, not to mention system software). The budget for the latter was slightly more than $60M, including the file system, leaving approximately $230M for the Purple 100-teraflop system itself.

As IBM originally proposed, the basic Purple servers (64 PE H nodes) consisted of highly reliable general-purpose computers. Much of the infrastructure on these machines was ideal for serving the time-sensitive transactional needs of businesses requiring extremely high reliability. This attribute was not high on ASCI's list of priorities, since jobs could easily be restarted. LBNL's NERSC computing leaders under Bill McCurdy and Horst Simon were also shopping for a major system with IBM and were similarly inclined to look for a more economical, lightweight SMP. Both labs communicated their concerns to IBM and ultimately collaborated on a lightweight node structure through direct engagement between Lab application developers and IBM chip designers. The result (the 8 PE IH node) was an economical solution suitable to scientific and technical computing and responsive to national-security mission needs.

Of these 8-way nodes, 1530 were connected using a high-performance switch called Federation. The system thus contained 12,288 POWER5 microprocessors in total with 50 terabytes of total memory and 2 petabytes of total disk storage. The computer consumed 7.5 MW, including cooling. These requirements in themselves justified the new ~$91M B453 building where both Purple and BG/L were integrated, essentially simultaneously, with Purple in the west machine area and BG/L in the east. Purple had a theoretical peak speed less than the ASCI target of 100 teraflops at 92.78 teraflops and achieved a very respectable 75.26 teraflops on the Linpack benchmark, taking the number four slot at SC in November 2006.

How LLNL, ASC HQ under Kusnezov, and IBM ultimately succeeded is an interesting backstory, in which the relationship between M&IC and the LLNL ASC program was enormously helpful. As discussed earlier, the LC sited the M&IC cluster Thunder, taking ranking number two on the Top500 list in June 2004. Let's pause for a moment to take this in: an M&IC (institutional) computer costing less than $20M took second place worldwide. It also ran at about one fifth the peak of what ASCI wanted to achieve as its endpoint. The integration was not turnkey and lacked vendor-backed guarantees. But LC, working with various component vendors, had demonstrated that clusters of relatively inexpensive Linux-based nodes could provide world-class compute capability.

Thunder had been conceived primarily by Seager, working with the system vendor, switch vendor, Linux support vendors, and LC operations and tools experts. Seager imagined a design for a

4X Thunder system and drew up a ~90-teraflop system that he believed could be built for around $80M. This raised doubt on why NNSA should pay $230M for the IBM Purple system. To add to the complexity, pressure on the ASC budget was increasing, and Kusnezov was asking the labs where cuts might be possible. What followed was a renegotiation of the Purple solution with IBM, forsaking the original (64 PE H-node) system for the lightweight 8PE IH nodes described above. This was expected to lower the cost of the Purple system and so relieve budgetary pressure that would have placed the BG/L procurement at risk. After concessions on both sides, a Purple system emerged for less than $150M. A casualty was the 100-teraflop target—the machine ended up at just under 92.78 teraflops. LLNL could have gotten to 100, but it would have meant jeopardizing other weapons work, and it was clear that cutting back on an important but less-visible effort to preserve in its entirety a publicly visible project would provoke, at minimum, an examination of LLNL priorities.

Another casualty of this redirection was the schedule. The original Purple demonstration date, per contract, was Dec 2004. Reengineering the solution would add an additional six months to the projected demo date and consume any schedule buffer, exposing the entire project to outside scrutiny should the June 2005 date erode further. Complex technology development—the manufacturing of more than 1500 new SMP nodes in three months—proved a bridge too far. However, because the most senior IBM executive with oversight interest had given personal assurances to LLNL director Mike Anastasio, the development organization created a dual-network environment (one at LLNL, one at IBM). This allowed necessary and rigorous testing to proceed without incurring a two-month scheduling setback from system breakdown, shipment across country, and installation at the Lab. This decision put greater than $10M of hardware at risk of being scrapped. But in the event, the June 2005 demonstration was preserved.

The existence of Thunder as a proof of principle most likely spared BG/L from the chopping block. The leveraging of two healthy computing efforts at LLNL—classified weapons and unclassified M&IC—far exceeded the benefits anticipated.

Figure 28. Purple (2005): 93 teraflops-per-second peak performance.

IBM was less than overjoyed at seeing millions of dollars taken from its Purple contract. This was especially true on the power-server side of the house. Their dissatisfaction was conveyed to Director Anastasio, Dona Crawford (the Computation AD) and Michel McCoy when they traveled to Yorktown to visit senior executives at IBM after the dust settled. IBM had a point. The original $290M budget for a 100-teraflop system had been used to buy two computers from IBM, one of which was not even a commercial standard product. IBM willingly undertook this extraordinary challenge and could well have created room in its original bid for the 100-teraflop computer to make BG/L possible. IBM likely viewed this as a partnership package deal. This was not the Yorktown meeting McCoy was hoping for in encouraging his director and AD to visit IBM, and he was downcast after the meeting. On the flight home, Anastasio kindly took the time to tell him not to be too troubled: business is business, even between good partners like IBM and LLNL. Indeed, LLNL and IBM partnered happily on multiple subsequent machines.

As a final anecdote, for the November 2006 Top500 list, SNL had just upgraded its Cray Red Storm computer to 127 peak, 101.4 Linpack using its 2.4 gigahertz Opterons. This machine then scored number two on the Top500 list (ahead of Purple at number four). An IBM Watson BG/L system scored number three, with the LLNL BG/L was first. It would have been embarrassing for LLNL if the laboratory selected to site the ASCI 100-teraflop platform were eclipsed by a low-cost upgrade from a sister lab at a politically conspicuous point in time. Red Storm featured a high peak but was less balanced overall for designer workloads than the Purple system. As demonstrated previously, subtleties do not change political pressures, and politics is frustrating to those on the wrong end of its ridicule. BG/L held a towering lead over all contenders, however, and held the number one position on the Top500 for over three-and-a-half years, rendering LLNL an undisputed number one by that metric. From this perspective, it could be argued that Red Storm was a midpoint between architectural exploration for speed and full-on weapons-code functionality, and not a low-cost improvement to Purple. But it took BG/L to make those arguments easy. In production, Purple was lauded by weapons designers for its memory bandwidth, capable file system, aggregate memory, and reliability.

c. The BG/L Contract: A New Procurement Model for DOE

As complex and trying as negotiations for the 100-teraflop platform were, the negotiations for BG/L proved more influential, changing the way future procurements were conducted for advanced-architecture and leadership systems and enjoying broad adoption across major DOE HPC sites.

The complicated issue during BG/L contract negotiations was that while Mark Seager was asking IBM to agree to strict performance metrics for BG/L, IBM lawyers were unwilling to sign them. The machine was not yet a product—it was highly experimental, as the company had not sufficiently developed the design. Nor was there ample testing to support high confidence that taxing requirements would be met. This led to an impasse. During a painful telephone negotiation among Seager, Nowak, and McCoy in Seager's office, Dave was inspired to suggest, Why not initially define targets for the contract and then harden the deliverables when more was known? Either both sides would agree to move ahead, or they would walk away owing each other nothing. IBM soon got back to LLNL with a tentative yes.

With help from Gary Ward, Seager set to work refining the idea with IBM and making it workable from the perspective of NNSA procurement regulations. What eventually emerged was that, once targets were defined, LLNL was to fund an R&D contract to help in the design and building of a small test system. This system could be used to test the targets and assess the effects of missing some. Some targets might be exceeded, some not met; but in aggregate, they would know whether the machine was good enough to build. This would be decided at a fraught go/no-go meeting. In a less flexible laboratory (or under stiffer NNSA oversight), this idea would have died in its crib. But miracles do happen. All future LLNL procurements of advanced architectures, Blue Gene/Q (BG/Q-Sequoia), Sierra, and El Capitan followed this model.

d. Early Calculations on Purple Catalyze the NNSA Capability Framework

Shortly after Purple made its way to the classified partition, Kusnezov, desiring the demonstration of some eye-catching achievements with ASCI's latest investment, asked the LLNL code teams to "take the machine" and do something to catch people's attention. After some deliberation, the B Program team settled on a deep dive into the thermonuclear-burn process in a primary. Limited understanding of this process had resulted in conservative choices made over the years to ensure the high reliability of our weapons systems. A deeper understanding of how thermonuclear burn works in a primary could lead to more efficient designs and substantial cost savings in maintaining the stockpile. The code team, led by Brian Pudliner, used Purple to conduct extraordinary 2D simulations at previously unheard-of resolutions, providing data for insights into the thermonuclear-burn process. The results were sufficiently intriguing.

NNSA launched a major national initiative, the Thermonuclear Burn Initiative (later renamed the National Boost Initiative, or Boost), to further investigate this

Figure 29. Blue Gene/L (2005): 590 teraflops-per-second peak performance.

process. The project was led by Frank Graziani, an LLNL physicist with extensive experience in advanced simulation. Boost stimulated broader discussions between NNSA and the labs about the nature of the overall weapons-science effort, and NNSA subsequently created the Predictive Capability Framework to focus on key unknowns impeding predictive simulation—that is, simulation with quantified uncertainties.

As Boost work progressed, 3D simulations demonstrated that some processes demonstrated in 2D had a far more limited effect in 3D. While this was a disappointment to those who thought they were onto something, the Purple machine nonetheless provided an undeniable vindication of the ASCI program and NNSA, demonstrating that ASCI's initial insistence on 3D simulation was prescient. Some phenomena were indeed highly 3D in nature, and 2D results in those cases could point to nonphysical explanations. In addition, it was apparent that the spatial resolution would have to increase substantially to decrease uncertainties. Predictive simulation demanded correct physical models and sufficient resolution in a 3D framework. Furthermore, a large number of 3D calculations were needed to quantify uncertainties. All of this flowed from the Purple experience.

e. The Lasting Impact of BG/L

Before going on to some of the early calculations on BG/L, it is useful to provide a description of the global influence of generations of BG systems. In 2020, LLNL and IBM won the SC20 Test of Time prize for a 2002 paper that was the first peer-reviewed overview article to disclose details of BG/L, including nodes, system packaging, and software support. As noted in *HPCWire*, "The machine was predicted to be at least 15 times faster, 15 times more power efficient and consume about 50 times less space per computation than the fastest supercomputers that existed at the time."

As LC program leader Terri Quinn noted at the time of the award,

> *The Blue Gene/L effort was born out of an urgent need to deliver more powerful but affordable computers for the National Nuclear Security Administration's national security work and IBM's willingness to help. This was a calculated risk by both parties that paid off. The line of Blue Gene systems, beginning with BG/L and up to the recently retired Sequoia, were marvels, and set us on the path for exascale-class applications and systems.*

BG/L remained the fastest supercomputer in the world until June 2008, peaking at 596 teraflops after a series of upgrades.

Two major unclassified calculations on BG/L piqued global interest. First, LLNL researchers under Fred Streitz won the 2005 Gordon Bell Prize for peak performance for a simulation of the solidification of tantalum at extreme temperatures and pressure. This was the first scientific application to exceed 100 teraflops of sustained performance. This early achievement of sustained performance required Herculean efforts. Fred Streitz wrote, "Sometimes all-night calls with Jim Sexton back in NY, [dealing with] the jury-rigged system whereby the login nodes would buzz my pager if the machine crashed in the middle of the night, the craziness of getting the weird 'double hummer' configuration to work, etc."

The MD calculation depicted atoms undergoing high pressure and consequently solidifying in small islands within the molten material. When the work was complete, Streitz prepared a video presentation, "Powers of Ten," that depicted scenarios for scaling each calculation up—the first with the few atoms possible on previous computers, then with each subsequent run having 10 times more atoms than the previous calculation. To the eye, each calculation seemed reasonable on its own. But the team was able to prove only that the instantaneous formation of solidified islands was physically correct at the point of reaching 16.4 billion atoms. In other words, for the first time in computing, sufficient atoms were included to observe the detailed solidification process.

Shortly thereafter, Streitz wrote,

> *Bruce Goodwin was hosting High Commissioner Bernard Bigot from CEA for a visit shortly after SC05. He, you [McCoy], and I tag-teamed a presentation to Bigot and his entourage that ended with the "Powers of Ten" movie, as I proved that using all of BG/L, we could capture solidification for the first time without any spurious size effects. We were finally simulating a system that was "large enough." He (Bigot) turned to his advisor and said, "But this changes everything!"*

---

[14] "An Overview of the Blue Gene/L Supercomputer," https://doi.org/10.1109/SC.2002.10017

[15] https://www.youtube.com/watch?v=7QWEH-RkxWY

CEA would go on to order a BG/L of their own shortly thereafter, starting a flood of purchases that cemented Blue Gene as *the* supercomputer of the decade. Bigot had quickly understood that we were at last at a turning point for simulation.

As *HPCWire* noted,

> *The following year, a team including LLNL scientists won the Gordon Bell Prize for a large-scale electronic structure simulation of the heavy metal molybdenum using Blue Gene/L. And in 2007, a joint LLNL/IBM team won another Gordon Bell award for a first-of-a-kind simulation of Kelvin-Helmholtz (KH) instability in molten metals on the system.*

A close look at the 9.2 B MD calculation in Figure 30, which can be envisioned like the effects of wind blowing over the sea, depicts the formation of waves with small secondary instabilities forming on the waves themselves. The wave formation was only possible using MD as opposed to differential equations, which require seeding perturbations to initiate instability. This advance demonstrates how it often requires extreme capability and new techniques to squeeze out the physics. Figure 30 depicts a similar, but 3D, KH calculation; but alas, it was beyond the convenience of BG/L to run the calculation long enough to observe the formation of the secondary instabilities.

HPCWire reported Streitz's comment at the time:

> *The Blue Gene/L machine really set the stage for what massive parallelism could accomplish. It seems quaint now to think about people questioning whether you could use 100,000 processors, but you look at modern machines that have millions and millions of processor units and no one says, "oh you can't use that," because you know what? We can. It was quite a remarkable time, and just an enormous amount of fun. And through all this time it stands as one of the things I'm most proud of, that we were a part of making this little bit of history.*

HPCWire went on to write,

> *For nearly a decade the Blue Gene/L series won multiple Top500 awards and Gordon Bell Prizes, including finalists, and served as a vehicle for many research publications," said Test of Time Award chair Amanda Randles, a former LLNL computational scientist and later assistant professor in biomedical sciences at Duke University. "In addition, Blue Gene/L was a precursor of the importance of energy efficiency before it was a recognized problem in the community. It's now a dominant constraint in the design of HPC architectures. This paper has had a tremendous and ongoing impact on the design of subsequent supercomputers."*

In September 2009, President Obama recognized IBM and the Blue Gene family of supercomputers with the National Medal of Technology and Innovation, the country's most prestigious award for technological achievement. IBM's senior vice president, John Kelly, invited Seager and McCoy to the October 2009 gala in Washington, DC, in acknowledgment of LLNL's role.

Figure 31. Mark Seager with Bernie Alder (at top right) and with Alan Gara (bottom right).

BlueGene class simulation provided critically needed insights into material behavior for Stockpile Stewardship

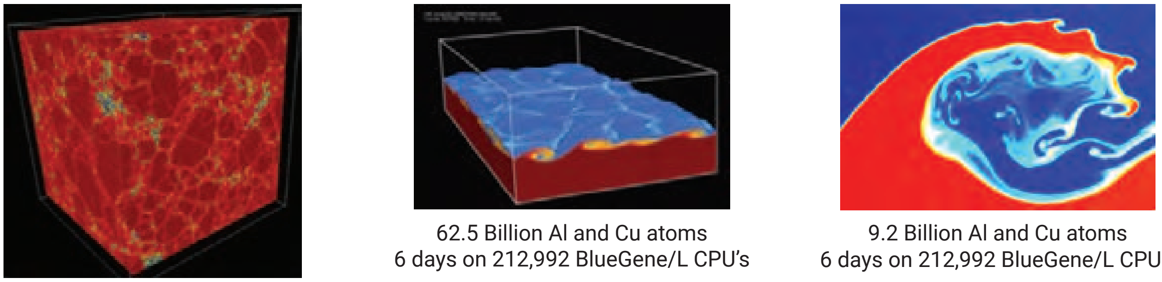

16,384,000 Ta atoms
6 days on 131,072 BlueGene/L CPU's

62.5 Billion Al and Cu atoms
6 days on 212,992 BlueGene/L CPU's

9.2 Billion Al and Cu atoms
6 days on 212,992 BlueGene/L CPU's

Current simulations are limited by simulation size and resolution

Figure 30. The 9.2 B MD calculation depicts the formation of waves with small secondary instabilities forming on the waves themselves.

Alan Gara was honored as the chief IBM designer and force behind BG/L, and LLNL's Bernie Alder was recognized for his lifetime work on (indeed invention of) molecular dynamics.

With this BG/L effort, LLNL found itself at the forefront of changes in scientific discovery. Partly because of the success of the work, LANL moved ahead with the Roadrunner computer in partnership with IBM, using a graphics accelerator called the STI (Sony, Toshiba, IBM) chip. This machine was the first to sustain a petaflop. While using this machine for weapons work proved unsustainably taxing, it contributed to the ASC decision to move to an advanced-architecture line of computers that shifted procurements from base vendor products (albeit adjusted) to riskier systems that promised faster advances in simulation, a benefit sorely needed by nuclear security.

9

SEQUOIA

a. Sequoia's Strengths and Weaknesses
b. Sequoia Simulation Environment and Integration
c. Unclassified Science, Gordon Bell Submissions, and Other Recognitions

Sequoia

In 2011, IBM delivered Sequoia, a BG/Q system and the third and final IBM Blue Gene series computer. Sequoia culminated the low-power-consuming, massively parallel computing approach explored by IBM with partners LLNL and ANL. The name was selected from a rich set of California attractions, setting a trend for future systems. BG systems were not in IBM's formal product line; had DOE not participated, IBM would not have considered such a major project.

While LLNL led with the first-generation BG/L system, ANL procured the largest second-generation system, Blue Gene/P (BG/P). LLNL procured a system half its size, Dawn, employed primarily as a code-development platform for Sequoia. Later, ANL procured a 10-petaflop BG/Q system, while LLNL went on to the 20-petaflop Sequoia. The partnership with ANL under ANL AD Rick Stevens, had been fertile. Communication across the labs had been very useful in bringing machines into production faster, developing system software, and organizing workshops for the user community.

Sequoia was a mammoth machine, peaking at 20 petaflops with 1.6M cores and 1.6 PB of memory. The peak performance of just one cabinet of Sequoia was equivalent to the peak of the Purple machine. The interconnect was not a switch, but a 5D torus. Remarkably, the machine required only 8 MW of power and nearly doubled the processing capability of the previous number one computer, the Japanese K computer, at a third less power. Sequoia was installed in B453. The Laboratory was naturally pleased with its first-place machine, as reflected on the cover of *Science & Technology Review*. Shortly after, a companion one-quarter the size of Sequoia was installed: a Vulcan BG/Q unclassified M&IC computer. The LC had been interested in forging computing partnerships with industry (which was a major initial driver of LLNL's open campus), with Vulcan a target platform for collaborative code development. This was the inspiration for the system's name.

Figure 32. Dawn (2009) peaked at 500 teraflop performance.

Figure 33. July/August 2013 cover of S&TR *magazine featuring Sequoia.*

The choice of a BG/Q system for the unclassified environment was a departure from the LC's original strategy, which was to use the M&IC environment to test alternative architectures to the ASCI systems. At this point, the weapons program and LC management felt more confidence in their choices and less need to reduce risk by dividing bets. With commodity Linux systems moving into the mainstream in both environments, it was expeditious to have identical architectures across classification domains and made life simpler for scientists working in both environments. It was also consistent with creating a computational lingua franca across the Lab, by which all programs and scientists could enjoy the same

---

[16] Livermore Valley Open Campus is an LLNL program to encourage collaboration with industry and academia by providing open facilities at the east of the Livermore site. Badges and clearances are not required.

high-quality tools and cooperate for the betterment of all missions. It would allow scientists trained in the unclassified environment to move efficiently into the weapons program.

Sequoia initially took number one on the Top500 list in June 2012, with a remarkable 16.32 petaflops out of 20 peak (over 80 percent efficiency). However, for many months in 2012, Sequoia and Vulcan had sat in the unclassified environment undergoing integration and early science runs. Sequoia featured 16 rows and Vulcan had four. Sometime after June, with Vulcan production hardening largely completed, LC management decided to integrate the two systems to create a single 25-petaflops platform before decoupling and moving Sequoia into the classified arena. This fusion was to make possible the largest unclassified problems on a 25-petaflop system (some of these calculations are discussed below), with the knowledge that the combined machine would certainly retain the first place for LLNL on the Linpack benchmark for the Top500 list of November 2013.

Figure 34. Picture of Sequoia on the machine floor in B453.

The number achieved by the combined system (over 20 petaflops) would indeed have kept LLNL at number one. But as the ASC program leader, McCoy was unsettled by the possibility of breaking rules, so he checked with a Top500 leader at LBNL, Horst Simon. Horst confirmed that the number would not be accepted, because LLNL had publicly stated that these would be two separate systems during their permanent production use. Although this was disappointing to McCoy and the LC, the logic was sound and they had the consolation that it would be applied consistently to all submissions. Without the uplift from Vulcan, Sequoia placed second on the list, just a hair below Titan, the new 27-petaflop peak ORNL system, at 17.6 petaflops.

In addition, but not commonly internalized, LLNL maintained two preeminent environments—classified and unclassified—at a staff size comparable to other major sites with only one major environment. This efficiency was achieved through common architectures and software stacks. LLNL's development of highly functional, world-class environments in two security domains was remarkable and the Lab's innovation continues to this day.

a. Sequoia's Strengths and Weaknesses

The Boost initiative and the Predictive Capability Framework (PCF) focused on critical issues impeding predictive simulation. Given the anticipated strengths and weaknesses of the Sequoia system, the computer was tasked in 2009 with addressing PCF requirements in the realms of uncertainty quantification (UQ) and MS, as related to boost and certification. Consequently, the acceptance criteria for Sequoia included:

1. Achievement of 12 to 24 times Purple's throughput for integrated weapons calculations related to UQ (stretch goal: much greater than 24 times)
2. Achievement of 20 times BG/L (stretch goal: 50 times) on an MS effort

Note, there was no serious requirement for running full-system (integrated multiphysics) calculations across the entire machine. Instead, the integrated design codes (IDCs) would run in parallel as a suite of smaller UQ calculations. The retired BG/L machine and Dawn had helped the LLNL code teams prepare for the massively parallel scaling needed given the relatively small amount of memory per processor. A major deliverable for the Sequoia machine and the codes using the machine was to serve as a testbed for massively parallel MPI scaling. During the Sequoia era, LLNL demonstrated the scalability of deterministic-transport-sweep algorithms to 1.5 million MPI tasks. This was supported by the development of theoretical parallel-performance models, which indicated that a family of these algorithms can be made scalable beyond Sequoia with reasonable changes to existing code bases—a significant finding. The results from the theoretical parallel-performance models resulted in guiding principles for future deterministic-transport porting efforts on Sierra.

Significant collaborative research among the Tri-labs and universities emerged from the Sequoia MPI scaling challenge. While all the integrated multiphysics codes struggled with reduced memory per core and scalability challenges, the A Program codes were at a disadvantage with this architecture because of heavy memory needs for the physics of secondaries.

This is not to imply, though, that B Program applications were unhampered by memory restrictions. Indeed, Sequoia never really made its mark as a capability machine for the weapons program (i.e., as a machine that could run integrated multiphysics at problem sizes no previous machine had achieved). However, as the acceptance criteria highlighted, Sequoia did prove adept at running large ensembles of smaller (e.g., 2D) calculations. Towards the end of its service, Sequoia was running targeted 3D ensembles to free up resources on commodity machines. Thus, this machine did meet the first goal and demonstrated the value of a system optimized for the throughput of ensembles. Likewise, some basic physics applications such as molecular dynamics were able to demonstrate capability-class calculations, meeting the second goal. The scale of large multiphysics applications in the weapons program was limited more by memory and communications costs, however, than by number of processors. This was more a statement of practicality than a hard limit. Very large problems could be scaled across the machine, but communication costs coupled with relatively slow processors made such calculations largely impractical.

Sequoia nevertheless played a substantial part in the LLNL weapons program as a throughput machine that could run thousands of moderately sized calculations simultaneously. This boosted the adoption of more rigorous validation and verification (V&V) and UQ methodologies, as Sequoia made such work much more practical. The ability to launch and manage thousands of runs simultaneously did not come free. LC staff at LLNL worked hard to expand the capabilities of their job scheduler to handle such loads, making LLNL computers unique in their ability to schedule and manage multiple submissions of large suites of simulations that simply overwhelmed other job schedulers. The code teams at LLNL invested heavily in software such as the UQ Pipeline, Maestro, PSUADE and Merlin that could generate a suite of simulations, submit them to the scheduler, monitor progress, and process output. LLNL was a leader in incorporating V&V methodologies in its annual assessments and development activities.

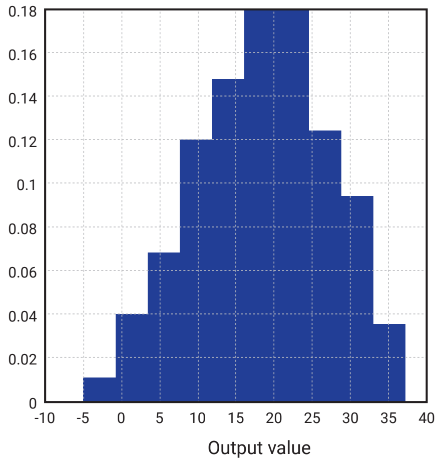

Figure 35. Probability of an output value as inputs are varied over their uncertainties.

Though Figure 36 (prepared by McCoy for a review) may seem naïve today, the key to getting support was often a one-viewgraph synopsis of the vision. This graphic conveyed that the series of systems delivered to LLNL were part of an enduring idea, even if that idea was not fully formed in 2005.

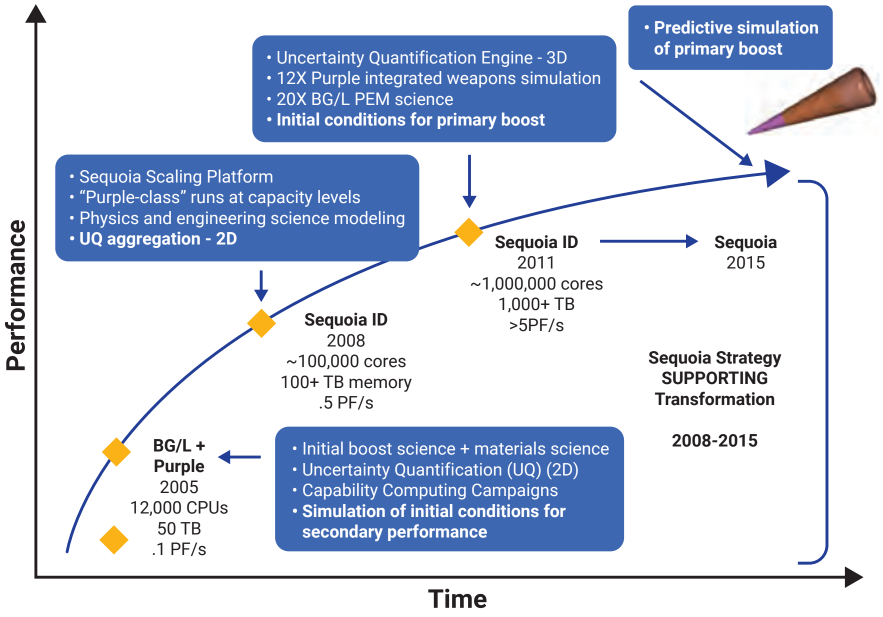

Figure 36. An envisioning of the steps to predictive simulation of primary performance.

b. Sequoia Simulation Environment and Integration

In a 2009 talk to the Predictive Science Panel (PSP), an academic-review board that gauged progress in predictive simulation, Mark Seager commented on the three scalability challenges that Sequoia had to meet—for hardware, software, and applications—and put his finger on what had made LLNL preeminent in computing: an emphasis on integrated management across all components required for success.

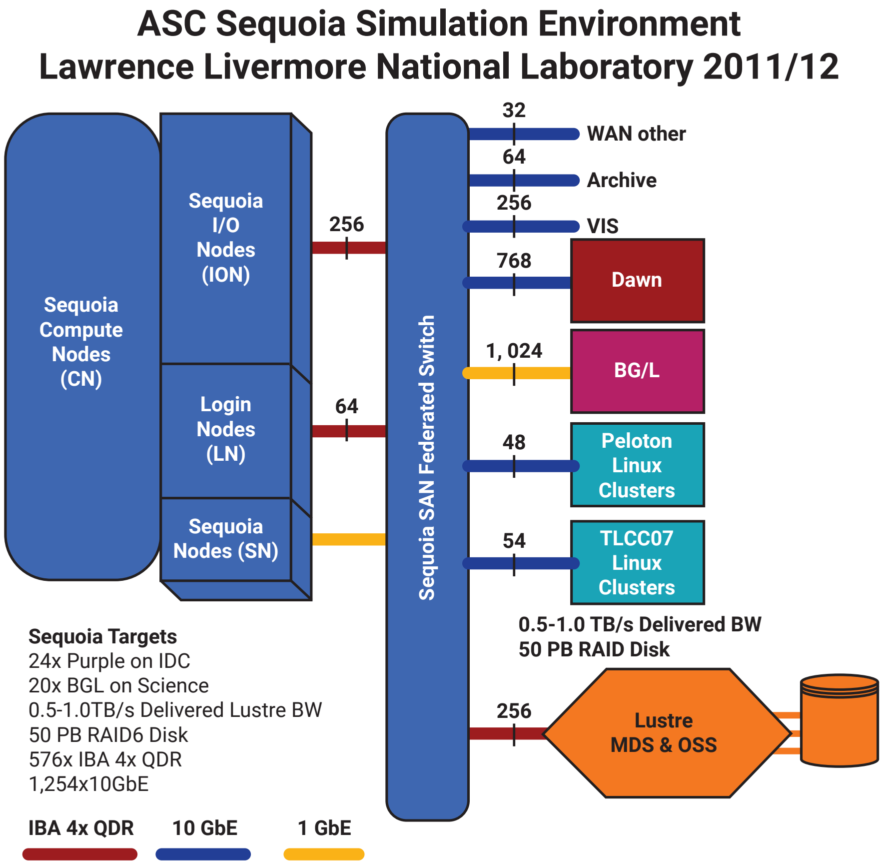

Figure 37. The proposed classified LC local-area network (LAN) highlighted in Seager's 2009 talk to the PSP.

The vision for a common file system (Lustre) included connectivity to Linux clusters and other major systems. The same was planned on the unclassified side. From this perspective, users could view their files as being in the LC "cloud" well before the cloud concept was commonplace. This notion was aggressive, if not revolutionary, for 2009. There were major setbacks in 2012 and 2013, however, in making it all work, and the LC back off some from the integrated-cloud vision as they worked with file-system islands local to computers to improve performance. Even so, the spirit of innovation and attendant risk was there.

Without going into detail on Seager's three challenges, the complexity of the challenge and required integration must be noted. From the application-programming perspective, MPI parallelism at the top level involved static allocation of MPI tasks to nodes and sets of cores and threads, effectively absorbing multiple cores and threads in MPI tasks, supporting multiple languages (C, C++, and Fortran03), and allowing different physics packages to express node concurrency in different ways. Figure 38 shows a high-level rendition of the Sequoia software stack from the applications perspective.

Seager, who had done much to bring the LC to this point, moved to Intel just weeks before Sequoia was integrated, and the LC made some rapid reassignments in responsibility. Terri Quinn became the program leader, with Robin Goldstone and Kim Cupps taking on additional LC roles. Kim took point on integration and did masterfully, despite some major problems flagged in red in Figure 39.

Sequoia provided two near-death experiences (one more than usual). The first was bent pins on the node cards, which could have caused tens of millions of dollars in loss for IBM had they not found a way to repair them. Luckily, the partnership was not put to that test.

Improper installation of water-cooling infrastructure was another near-death. Technicians noted impure, reddish water flowing through Sequoia racks, forcing an immediate shutdown of the system to forestall irreparable damage. It was soon determined that the water-pump vendor had over-sprayed the pumps with a red paint that was flaking off as water flowed. It took weeks to flush and filter the entire machine (see Figure 40). Once flushed clean, Sequoia could move on. These are the kinds of things that turn a technical manager's head gray. How could such a colossal failure be explained to Laboratory and HQ management?

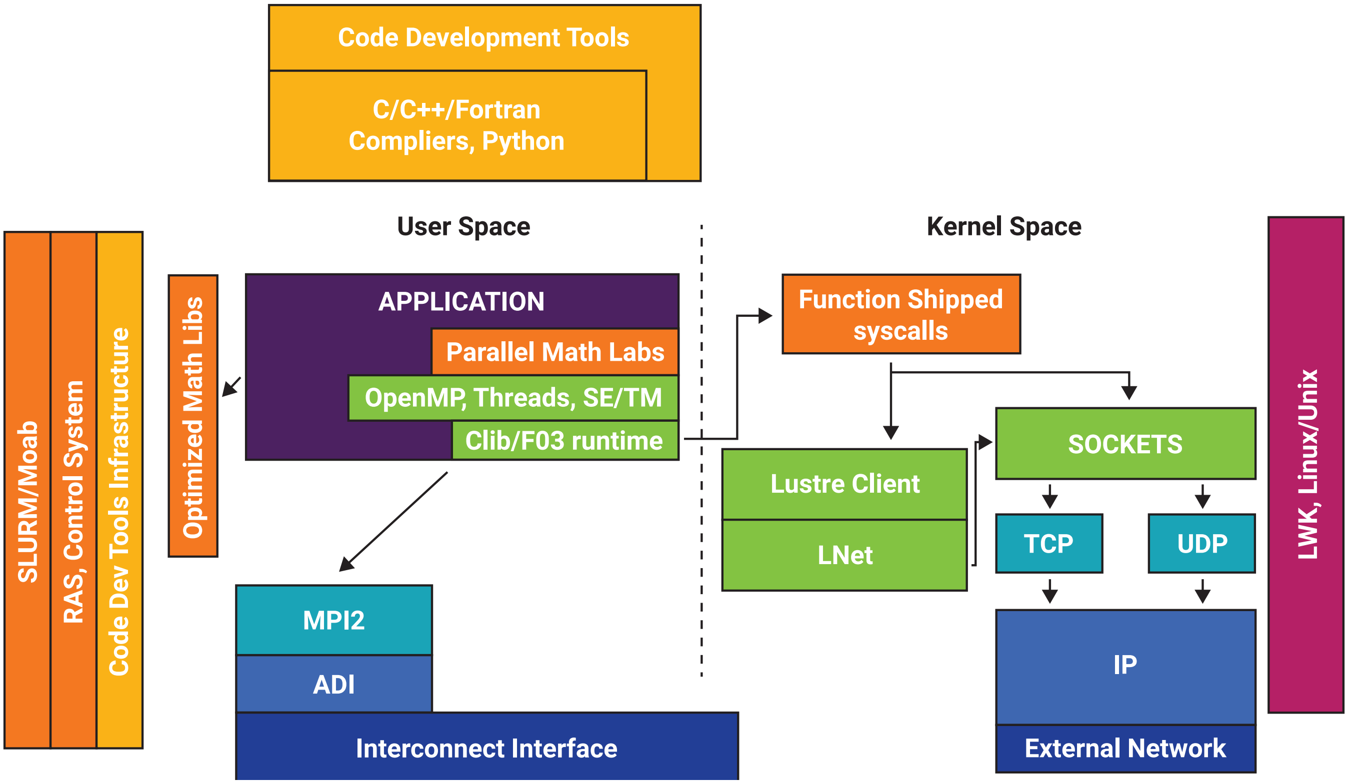

Figure 38. The Sequoia (Blue Gene/Q) software stack as deployed in the LC.

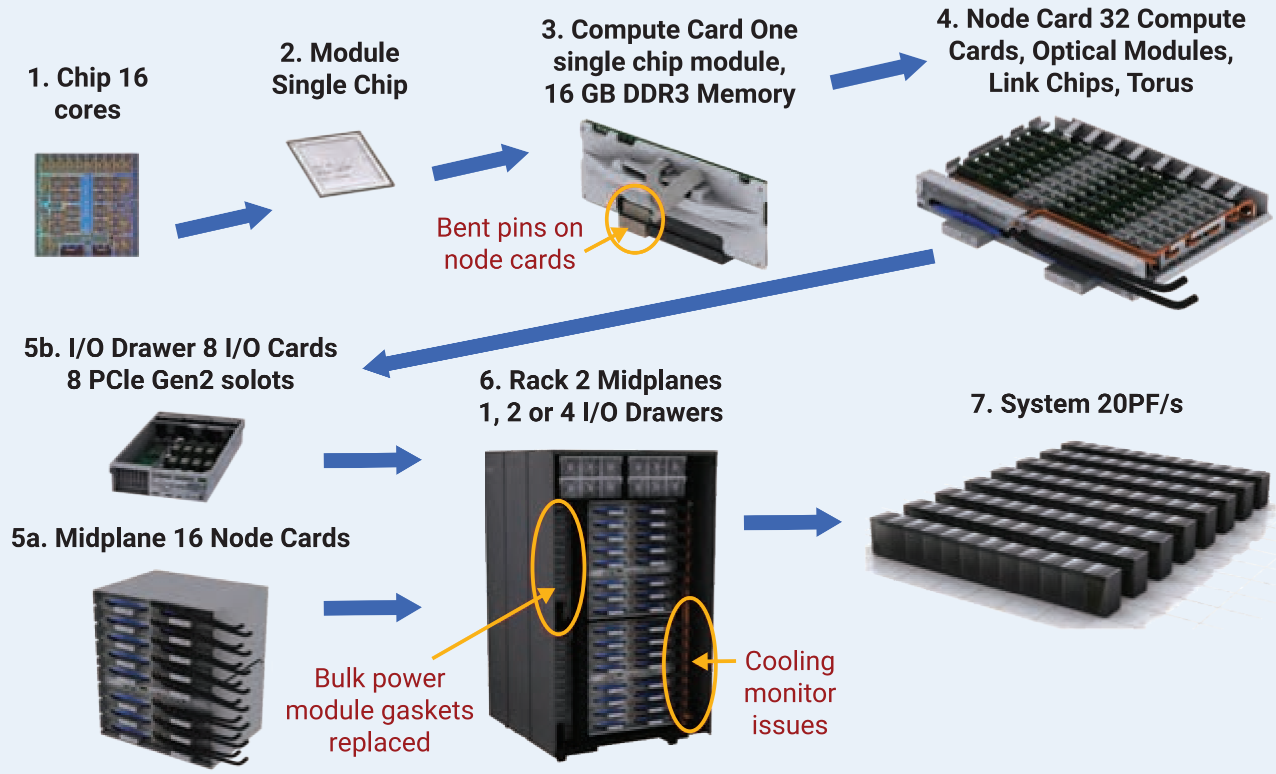

Figure 39. An excerpt from Kim Cupp's presentation on difficulties in the Sequoia stand-up.

Figure 40. (Left) The red-painted water pump before repairs; (right) the cleaned-up pump.

c. Unclassified Science, Gordon Bell Submissions, and Other Recognitions

The Sequoia and Vulcan integration period was an opportunity to run problems that used the combined peak of the two machines (25 petaflops). The challenge for scientists was exploring what kind of science could be done with this unprecedented computing power. Multiple codes had access to the machines during this period, with an emphasis on seismology, laser–plasma interactions for NIF, MD, and cardiology. One of these codes (Cardioid) was a finalist for the Gordon Bell Prize.

Cardioid (Streitz et al.) was intended for electrochemical human-heart simulation. One reason for LLNL's selecting this code was its intuitive appeal to the public, something that mission-specific codes like MD lacked. Using a highly scalable code developed in partnership with IBM, LLNL modeled the electrical signals travelling from cell to cell, triggering them to contract. The code achieved better than 50 percent efficiency (10 petaflops sustained)—an amazing level given the complexity of the model. Earlier modeling of the heart was limited by the speed and memory of smaller machines and permitted the modeling of only a few heartbeats. Cardioid took this to another level, modeling hundreds of heartbeats during a single run and as efficiency increased, in real time, that is, 60 beats per minute. With this improvement, the team was able to induce arrhythmia from the simulated injection of a drug—another first.

The Blue Gene series continued to get attention. In November 2012, *Popular Mechanics* listed the IBM Sequoia in its eighth annual Breakthrough Awards ("celebrating world-changing ideas, innovative products"). The ceremony in Manhattan was attended by IBM leaders and LLNL's Weapons Program PAD, Bruce Goodwin, and Michel McCoy. As seen in the "Men in Black" image in Figure 41, *Popular Mechanics* advertised Sequoia as "outflopping the world." Elon Musk was the major recipient (for SpaceX and Tesla), and spoke at the ceremony, as did Goodwin.

As emphasized previously, the Blue Gene series opened the doors to acceptance of simulation as a peer to theory and experiment in scientific discovery. It was beginning to be possible to think of using computers for scientific prediction, beyond their previous use as tools for gaining insight. With the ability to launch large ensembles, simulations could now be characterized with quantified uncertainties much more routinely. This put them on a more solid footing as a pillar of scientific discovery and guided researchers in choosing experiments that might help constrain simulation uncertainties.

Publications that relied on simulation for scientific advances were increasing rapidly. LLNL's emphasis on unclassified computing had borne fruit. The examples in Figure 42, taken from a 2010 presentation by Brian Carnes, the deputy of ASCI and M&IC leader, support E.O. Lawrence's observation, "The day

Figure 41. The "Men in Black" photo featured in Popular Mechanics coverage of LLNL's Sequoia.

when the scientist, no matter how devoted, can make significant progress alone and without material help is past.”

Our state-of-the-art simulations appear on the covers of leading scientific journals

> “Access to the M&IC Grand Challenge program has enabled predictive and systematic simulation with respect to length scales, dimensionality, and interface effects without having to compromise on the level of theory used.”

Figure 42. Some journal covers featuring LLNL simulations.

10

WEAPONS-CODE EVOLUTION

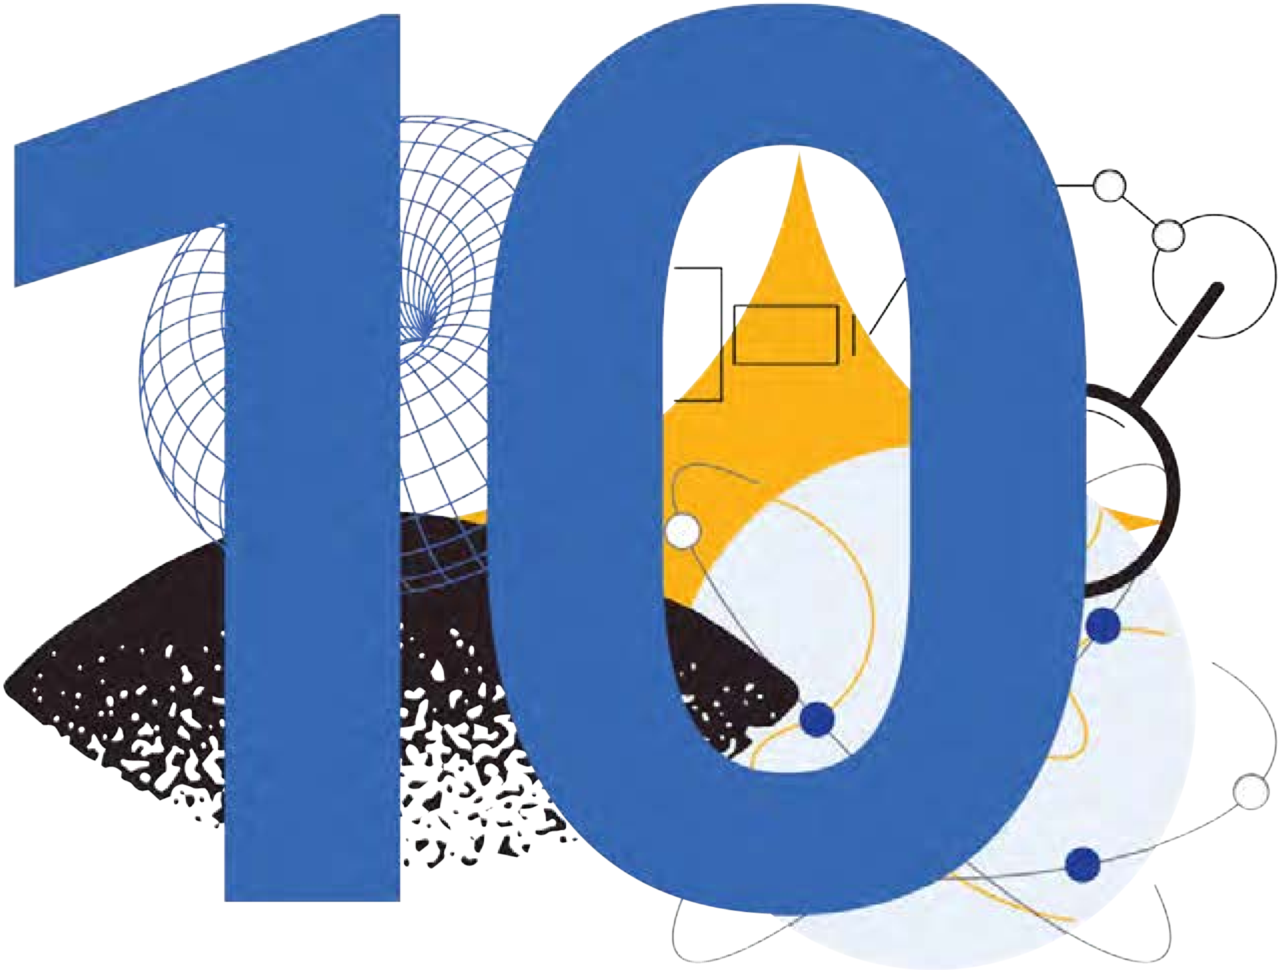

a. Reorganization and Centralization

b. Launching the Next-Generation Code Effort

Weapons-Code Evolution

a. Reorganization and Centralization

Despite overlap in the simulation capabilities needed for A and B programs, each had its own independent code development effort until 2007, when the decision was made to consolidate management of the two code efforts under a single leader. Jim Rathkopf had been leading the A Program effort at the time and was the natural choice to take over, since Tom Adams, who had succeeded Charlie McMillan at B Program, was recently retired. The creation of a single leader for code development was essentially a compromise. Rathkopf still reported to both the A Program leader (initially Charlie Verdon, then Des Pilkington), and the B Program leader, Mike Dunning, as well as the ASC program executive, Michel McCoy. During this period, little *real* progress was made in unifying code development in A and B programs.

The situation improved in 2015 when the program was reorganized by Charlie Verdon, PAD for the weapons program (then named Weapons Complex and Integration, now called Strategic Deterrence). Charlie united A and B design into a single program under Weapon Physics and Design (WPD) and created a new program, Weapon Simulation and Computing (WSC), under which all code-development activities were consolidated. McCoy, the first program director for WSC, selected Chris Clouse as his associate program director for computational physics (CP).

While one of the great strengths of LLNL was the tight relationship between weapons-code development and hardware procurement and development, that relationship relied primarily on personalities—individuals who maintained the highly responsive structure of the LC to meet the simulation needs of the weapons program. Verdon's reorganization of the weapons program essentially codified that relationship by placing the LC and weapons simulation development under one program, WSC.

b. Launching the Next-Generation Code Effort

The year before the reorganization, Clouse, the acting leader for A and B code development after Rathkopf departed for HQ in Washington DC, launched a new code effort: a next-generation (Next-Gen) code specifically architected for anticipated changes in advanced-computing platforms. These changes were largely motivated by impending power constraints in the industry. Purple, for example, provided roughly 100 teraflops of compute capability, using roughly 4 MW of power. A 20-petaflop machine like Sequoia would clearly not be affordable at 20 times the electrical power of Purple. Thus, IBM took the many-lightweight-processors approach embodied in the Blue Gene series, though other architectures were also contenders. The use of graphics-processing units (GPUs) eventually proved a better option, although forays into that heterogeneous computing (like LANL's Roadrunner machine, which took the top spot on the Top500 in 2008) were insufficiently evolved for production weapons work. Roadrunner nevertheless provided valuable lessons that informed the first truly successful GPU-based platform for the weapons program, Sierra.

Regardless of approach, computationally heavy kernels with minimal communication were clearly a winning strategy. These new architectures had lots of local flops, but communication remained expensive. With that in mind, Clouse chose Rob Rieben as the project lead for the Next-Gen code effort. Rieben and collaborators in CASC had been developing new hydro algorithms under LDRD funding that took a very different algorithmic approach. Rather than tackling problems with low-order schemes and simply increasing resolution when higher fidelity was needed, Rieben used high-order schemes that were more computationally intensive, but required many fewer zones for equivalent fidelity, resulting in higher flop/byte ratios. This seemed a good match for anticipated future architectures.

LLNL's Next-Gen effort spurred ASC HQ to lobby for funding a Next-Gen effort across the NNSA complex. Initially, Doug Wade, the acting head of ASC, had rewarded LLNL extra money in the budget for their forward thinking, but was looking for more money for the entire ASC Tri-lab program. The message ASC HQ conveyed was that our weapon-simulation capabilities were in danger of becoming obsolete on anticipated computing platforms. This warning resonated with Congress, and Wade secured funding for a new program element within ASC—the Advanced Technology Development and Mitigation (ATDM) program. With ATDM funding, LANL and SNL began to spin up their own Next-Gen efforts.

Several potential hurdles for the Next-Gen strategy remained. The first was that the next advanced platform was unknown. Certainly, a code running on a GPU-based machine would be optimized differently from a code running on many lightweight processors—and those were just two of the potential architectures of the future. This problem was tackled by two

LLNL computer scientists in the weapons program, Rich Hornung and Jeff Keasler, who developed a programming model called RAJA. The idea was simple: replace for-loop constructs with similar syntax to be expanded at compile time into machine-specific optimizations, usually with no changes to the loop body itself

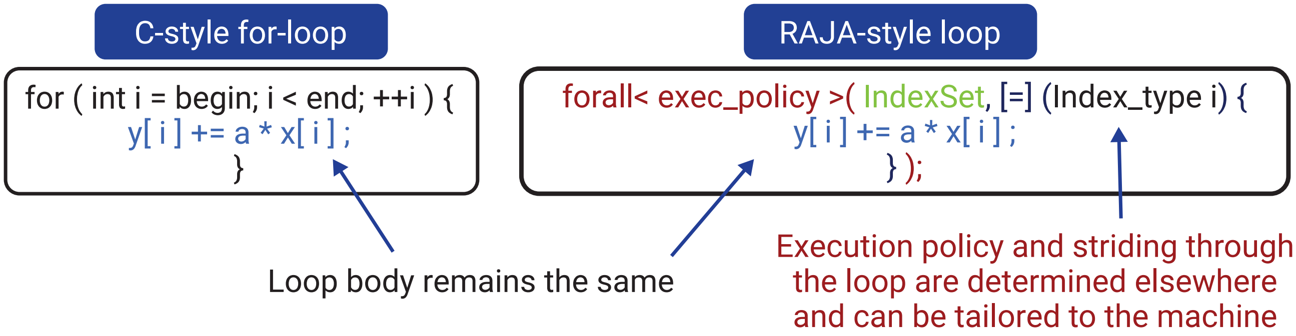

C-style for-loop

```
for ( int i = begin; i < end; ++i ) {
    y[ i ] += a * x[ i ] ;
}
```

RAJA-style loop

```
forall< exec_policy >( IndexSet, [=] (Index_type i) {
    y[ i ] += a * x[ i ] ;
} );
```

Loop body remains the same

Execution policy and striding through the loop are determined elsewhere and can be tailored to the machine

Figure 43. A RAJA-style loop.

The simplicity of the model and its deliberately limited scope allowed relatively quick adoption by the various codes in WSC, since it could be retrofitted into existing codes in a straightforward manner. Work on the optimization of loops could then be done in RAJA for the benefit of all host codes. SNL was concurrently developing a similar approach in Kokkos, but with a much broader scope, including memory layouts and host-device memory motion, making it initially more intrusive to retrofit into existing codes. Both RAJA and Kokkos were major contributions to the exascale effort for which the Office of Science eventually got funding, and which led to the Exascale Computing Project (ECP). Given NNSA's early start on Next-Gen efforts with ATDM, many ECP projects ended up relying on RAJA or Kokkos, while others used the evolving OpenMP standard or nascent SYCL standard promoted by Intel.

The second hurdle was development redundancy. The new code workload was added to the continuing development and maintenance of existing production codes, leaving resources easily overwhelmed. Initial ATDM funding helped stave off severe constraints, but ATDM had a limited lifetime, and the increasing complexity of codes and architectures was certain to take a toll.

After a bout of strategic planning in 2015, Clouse embraced a modular development approach. Major physics packages were to be modular and shareable among code projects, expanding a practice often followed in A and B code teams on a small scale. This included hydrodynamics, which historically was the hub of a code around which other physics accreted. Infrastructure was also made modular and reusable among code projects, and Rich Hornung became the first project leader for shared computer-science components. A key aspect of this approach was the creation of a datastore or central storage for all mesh-field variables. Each physics or CS package could access datastore arrays and perform its operations to update state variables in each zone. CS packages could process the data for output or other uses, updating the datastore and making the data accessible to the next package. Highly successful packages like the radiation-transport capabilities developed by Paul Nowak were retrofitted for modularity. New capabilities were developed, such as Brandon Morgan's RANSBox, which modularized commonly used turbulent-mix models.

Another aspect brought to the forefront in Clouse's reorganization was the creation of a workflow project headed by CASC's Dan Laney. This project recognized that most end-user time is spent on problem set-up and post-processing results. Improving this workflow became a major emphasis in WSC. A significant product from the workflow project was a code-

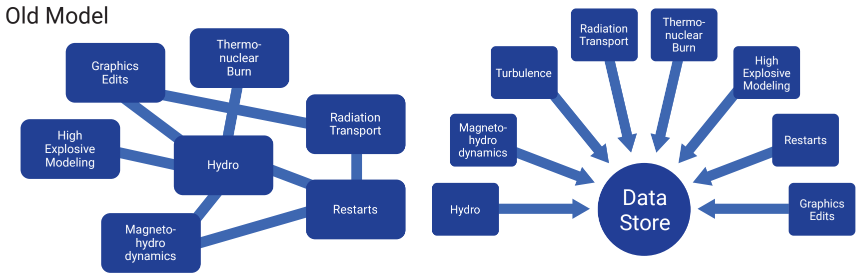

Figure 44. Clouse's new model (right) increased modularity with 40 percent of large integrated codes consisting of "infrastructure," increased the modularity of physics packages, and centralized mesh-data information.

agnostic library for generating geometry setups, named C2C, with a format developed at LANL. The use of C2C, together with the modification of LLNL's problem-generation project, PMESH, to read C2C syntax, enabled Rieben's Next-Gen code to provide simple and fast problem generation from the input of other codes—a significant advantage in gaining code adoption.

While the benefits generally outweighed the disadvantages of modular code development, there are several negatives associated with the approach. Writing reusable code requires more initial effort, so longer development times can be expected in achieving the same capability as development within a single code. Modular code also requires greater collaboration and coordination across code projects to agree on standards and avoid potential conflicts. Code builds can become quite complex, as each module has its own needs for external libraries. This last problem was partly addressed in 2013 by CASC's Todd Gamblin at the inception of the Spack project. Spack is a package-management tool to simplify the building of complex simulation tools with numerous dependencies on HPC systems. Its open-source format allows contributions from users across the globe to record dependency information and location on various platforms. Leveraged by the community, this data can reduce build times on a new HPC platform from weeks to hours. Spack has grown rapidly in popularity and is now used worldwide, with over 1,200 contributors and thousands of users.

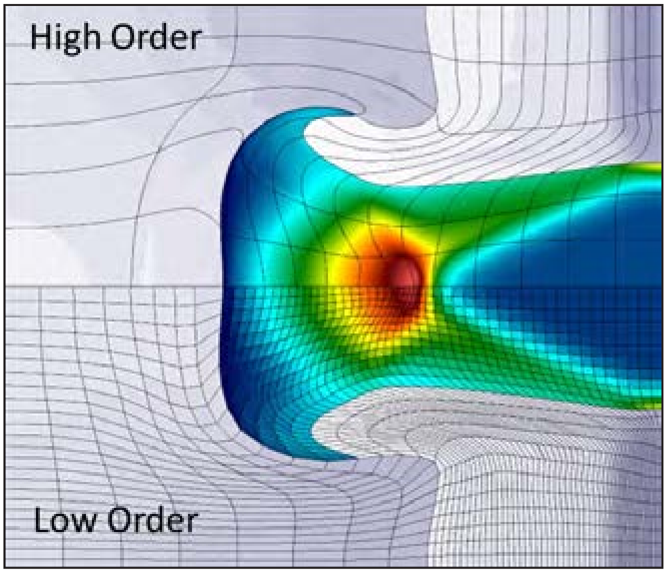

Figure 45. Contrast between a low-order and high-order mesh. Note the curvature in the mesh element boundaries allowed by high-order discretization.

The genesis of GPUs for computing at LLNL (as told by Becky Springmeyer)

In the early days of ASCI, visualization-project leader Randy Frank led a team of researchers, developers, and operational experts, notably visualization architect Dale Southard, now at NVIDIA. LC had deployed very expensive SGI visualization servers. Dale and Randy took note of developers working on the open-source Clustered High-Availability Operating System (CHAOS), derived from RedHat, and decided to build a GPU cluster running CHAOS. When MCR, the first large commodity cluster, was deployed in 2004, it was complemented by the Production Visualization Cluster (PVC) with NVIDIA GPUs. The inexpensive commodity visualization clusters replaced the SGIs.

Vis team members programmed on the visualization clusters using OpenGL. They collaborated on open-source visualization and data-analysis solutions. Randy Frank led a contract for Stanford through the ASCI program to fund Pat Hanrahan and students. The collaboration included the development of the R&D 100 award-winning Chromium parallel OpenGL API as part of the software for the visualization clusters. One of the students was Ian Buck, who worked on a streaming-programming model called Brook. At the 2016 NVIDIA GPU technical conference, Buck, now an NVIDIA employee, presented a talk, “From Brook to CUDA.” NVIDIA had released CUDA, and the race was on to leverage GPUs for more general-purpose applications (GPGPUs).

As ASCI visualization funding waned, the team dispersed to multiple programs and companies. Those who stayed leveraged their skills in GPGPU programming on an array of projects. Some went to the Persistics team to prototype algorithms, test concepts, and explore scalability for system design. They delivered scalable algorithms running on heterogeneous architectures to solve problems in wide-area persistent video surveillance. Fluent in graphics-programming environments, they were on the bleeding edge of GPGPU programming and involved in interdisciplinary projects that demonstrated potential new roles for GPUs. One of the first was a project on hardware-accelerated simulated radiography, published in IEEE Visualization in 2005 by Dan Laney and others.

Rob Rieben was one scientist with the foresight to appreciate GPUs. Around 2010, he observed demonstrations of GPUs solving equations relevant to scientific simulations. About the first, a Playstation 3 demo of guiding a toy duck in water, Rieben noted, “The water physics was simple, but rendered in real time, which made it clear that the NVIDIA GPU inside the PS3 could solve equations very fast.” MFEM project leader Tzanio Kolev brought to Rob’s attention a real-time hydrodynamics simulation and visualization by Ben Bergen at LANL. They recruited Tingxing “Tim” Dong from UT Knoxville, who was working at the Innovative Computing Laboratory, to intern at LLNL for two summers and get an early version of Blast running on the Edge visualization cluster with NVIDIA Tesla M2050 GPUs. Rieben showed Brian Pudliner a real-time demo of GPU Blast versus CPU Blast: “I hit enter at the same time for both versions, so it was like a race. It was clear that the GPU code was running much faster. Brian said something like, ‘Oh man, I want some GPUs.’ And well, we all know what happened after that.” Rich Hornung and Jeff Keasler, working within the Weapons Simulation and Computing program, developed the core of the RAJA portability suite that enabled most of our codes to run on multiple GPU architectures.

LC’s deployment of Edge in 2011 with 2,592 Intel Westmere processors, 412 NVIDIA M2050s, and 20TB of memory represented a decade of evolution in accelerating Linux visualization clusters. Edge enabled the porting of several application codes to GPU architectures. Southard, now at NVIDIA, sought scientists willing to partner with NVIDIA experts. Rieben, Kolev, and Dong were the first. Another application was a project on lattice quantum chromodynamics that used Edge to perform novel calculations. The researchers estimated that equivalent calculations on homogeneous clusters would have taken at least three times longer.

Edge was replaced by Surface in 2014, followed in 2018 by the Pascal visualization cluster still in service in 2024. Many code teams ported codes to GPUs on Edge and Surface as LC brought in GPU training. In 2014, the Sierra early-access systems arrived. LC ran several centers of excellence for Sierra to partner with IBM in preparing code teams on a much larger scale.

Fast forward to the Sierra dedication ceremony at LLNL. In his remarks, Ian Buck recalled the LC–Stanford visualization collaboration and how it helped him develop Brook. Buck was hired at NVIDIA after graduation to help develop CUDA. NVIDIA CUDA experts sited at LC were instrumental in helping code teams port and tune their codes on Sierra.

GPUs have remained a core element of High Performance Computing. Most notably, El Capitan has AMD MI300 accelerated processing units (APUs). This brings us full circle, from combining CHAOS and graphics cards to combining TOSS, AMD APUs, a Slingshot network, and an innovative I/O system with near-node local storage on El Capitan. Through two decades, the ASC program not only anticipated, but led the way in ensuring that NNSA codes leverage commodity technology (including GPUs) coupled with open-source software to ensure success with a variety of vendor and academic partners. Randy Frank noted, “I think LLNL has played an important role in the evolution of GPUs and the transition to GPGPU simulation environments. Starting from a basic need to transition to more economically scalable visualization solutions to early exploitation of programmable GPU logic one can trace a line of research and development projects that helped shape and realize the notion of the modern GPU-based computing environment.

11

ASC SIERRA

Figure 46. Sierra (2018): 125 petaflops peak performance.

ASC Sierra

Sierra realized the ASCI vision of modeling nuclear devices in 3D with high fidelity. While the Purple machine provided a proof of concept for running in 3D at entry-level resolutions, Sierra made high-fidelity calculations routine. Sierra enabled investigation into the next steps in capturing sub-grid physics that go well beyond the initial goals of ASCI. This mammoth accomplishment is not widely appreciated. The weapons-code teams had to refactor their codes to run on accelerators based on GPUs, an enormous undertaking that continues to bear fruit today. How this effort achieved a step change in simulation capability is outlined below.

a. The Genesis of the Enduring Partnership with Office of Science

Touched upon earlier was the spontaneous partnership among LLNL, ANL, and IBM in developing the Blue Gene series. It made sense for the two labs to coordinate closely—not to partner formally in a procurement, but to support each other in separate quests.

When the time came to begin planning the ~100 petaflops+ LLNL Sierra system, ANL's Rick Stevens, at a dinner with McCoy in Washington, DC, suggested that the Office of Science ASCR program was considering formalizing the grassroots partnership at the two labs. They recognized how effective Blue Gene interactions had been and concluded that a formal agreement would make it even better, and ASCR and ASC HQ agreed. This was an adroit move on DOE's part. For one thing, if the labs did common procurements, the dollar total would be more dramatic, attracting the attention of vendors. Microprocessor and system design were driven by forces far richer than those within DOE, so any possible aggregation of funds was advantageous—advanced architectures required major vendor R&D investment, and the prize had to be worth the effort. Additionally, the labs and DOE saw exascale computing just over the horizon. HQ anticipated that Congress would be supportive only if DOE had an integrated partnership rather than Balkanized camps. Moreover, while the Office of Science had very strong support for funding systems, the NNSA mission introduced a national-security driver into funding decisions. Even if NNSA were challenged fiscally, the national-security arguments for exascale computing were uniquely resonant.

In suggesting a common procurement, Stevens made the point that HQ wanted ORNL to join the partnership with ANL and LLNL. He was unsure what McCoy's reaction would be, as ORNL and LLNL had been quietly competing for number one on the Top500 (as noted earlier, Sequoia plus Vulcan would have grabbed the top spot on the list, but this gambit was disallowed by Top500 leadership). Stevens braced for a possible volcanic eruption across the table; but to McCoy (after a moment's thought to get his bearings) it was a golden opportunity to move the procurement from "love to have" to "let's get it done." He knew evidence of the three Office of Science and NNSA labs cheerfully working together would please DOE (and Congress) and generate impressive momentum for the procurement. McCoy asked only that the solicitation originate from LLNL, and this was readily accepted.

What came of this was a productive, multiyear collaboration among ORNL, ANL, and LLNL (dubbed CORAL). The LLNL–ORNL relationship blossomed and grew close as the two labs selected the same IBM–Nvidia solution for the 2018 system (Sierra and Summit, respectively) and for exascale systems with Cray (later Hewlett Packard Enterprises, HPE). ASCR required that ORNL and ANL select different vendors; ORNL had the inside track for GPU-based solutions, as the Lab had selected these before working with Cray, and its user base had an affinity for GPUs. The relationships among the leaders at ANL (Rick Stevens), ORNL (Jeff Nichols), and LLNL (McCoy) remained steadfast through these major procurements and the behind-the-scenes slog to generate congressional and executive commitment to exascale.

b. The Sierra Procurement

A common solicitation must originate from one lab, and LLNL was given the responsibility. To understand this choice, a look at ASCR computing history is warranted. Before 2004, NERSC at LBNL was the primary computing site for DOE researchers. The ASCR office at HQ and the labs without large systems, in particular ORNL and ANL, wanted to establish a powerful presence in HPC at their labs. There was congressional support for this idea, notably by Senator Alexander of Tennessee. In addition, the lingering scars from the Earth Simulator humiliation inspired arguments for DOE Office of Science investments. HQ ordered a solicitation and competition to determine the leadership sites. ORNL emerged as the primary winner, with ANL a partner. The emphasis for these sites would be leadership (read risky) architectures, and the plan was to provide capability computing for a moderate number of key computational-science teams to advance the scientific method at scale. NERSC would remain primarily a capacity center serving large numbers of scientists. The winning HPC vendors were also determined, and ORNL's was Cray. From that time forward, until the CORAL partnership emerged, ASCR sites never underwent a full-blown, complex competition, as the NNSA sites had received routinely. With HQ's support, ORNL managed to make the case that each new system was an upgrade. While this strategy reduced the uncertainties, risks, costs, and delays of a procurement, it did not provide ORNL with all the experience necessary to do a procurement, especially at this scale. LLNL made a second argument. The procurement model for advanced architectures developed at LLNL for the Blue Gene system and used at LLNL for later procurements was perfect for the experimental, risky

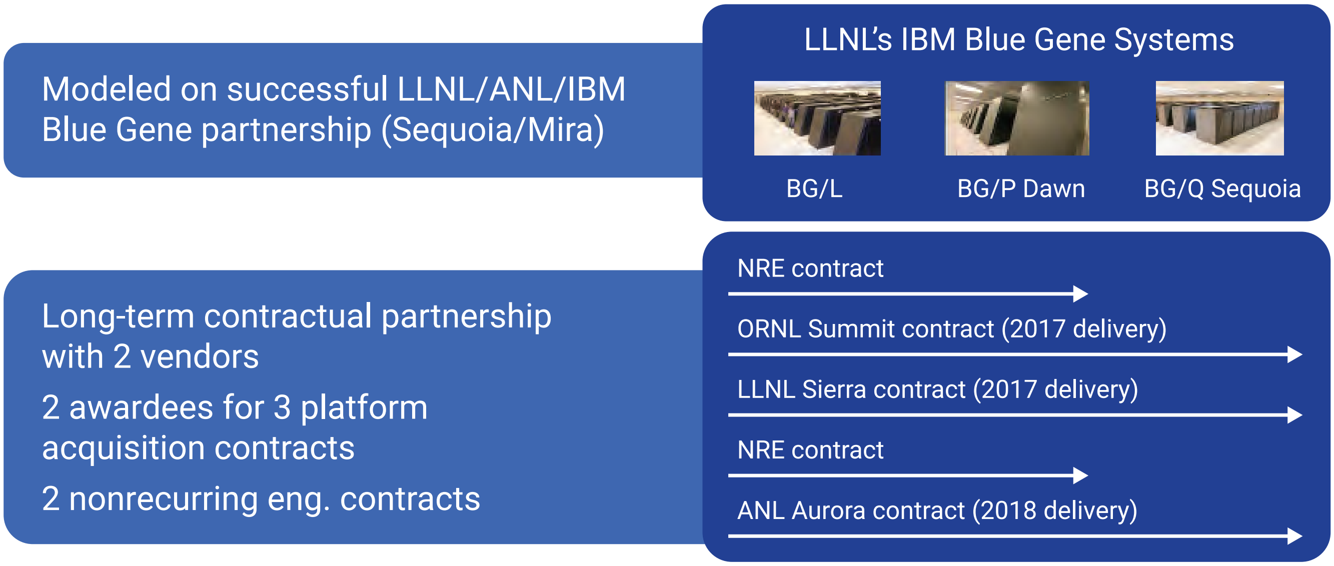

Figure 47. Representation of the LLNL–ANL–IBM Blue Gene partnership for Sequoia and Mira systems.

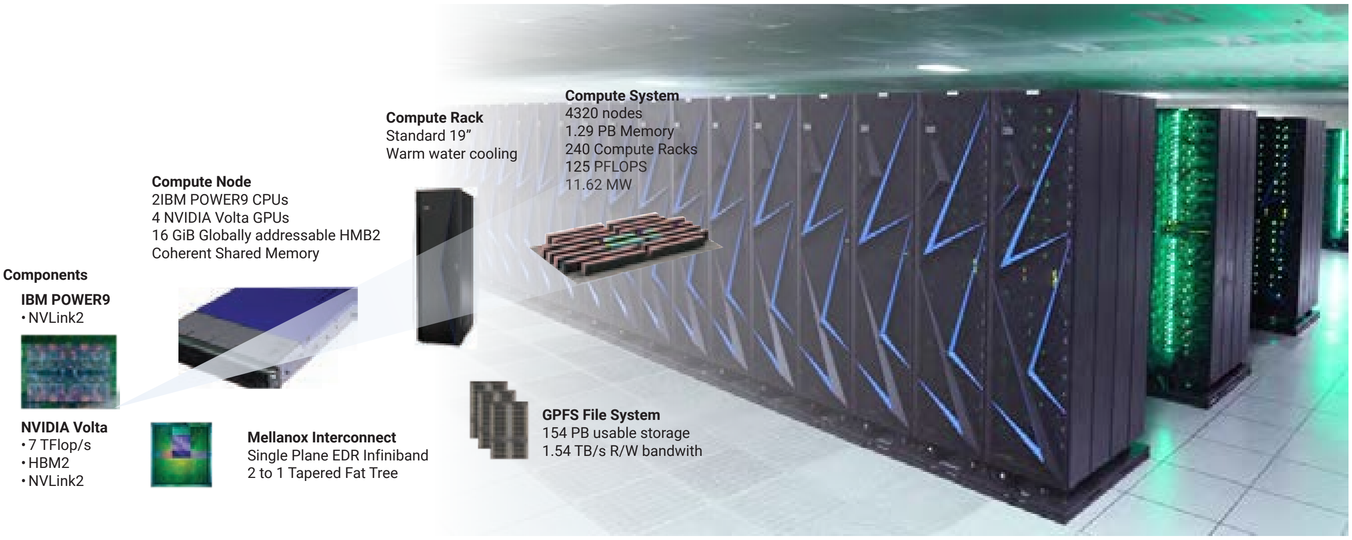

Figure 48. Sierra architecture as installed (left) and view of Sierra (right).

systems to be procured by ANL and ORNL. But this kind of procurement was complex and tricky.

The complex multi-lab RFP that emerged under Bronis de Supinski, Robin Goldstone, Matt Leininger, and manager Terri Quinn reflected the partnership with ORNL and ANL, incorporating their requirements. Provision was made for nonrecurring engineering (NRE) contracts in the Blue Gene model. Ultimately, LLNL (Sierra) and ORNL (Summit) selected an IBM system with Nvidia GPUs, and ANL (Aurora) went with an Intel system. R&D contracts followed as well. The Sierra system had a peak performance of 125 petaflops, featuring 4,320 nodes interconnected by a 2:1 tapered Mellanox Infiniband fat-tree topology. Each Sierra node featured two IBM Power 9 processors and four NVIDIA Tesla V100 GPUs. There were 320 gigabytes of fast memory per node, spread across 256 gigabytes of DDR4 and 64 gigabytes HBM2. The ORNL system had a higher peak, having selected three Voltas per Power 9 (as opposed to two at LLNL) to balance the machine for weapons codes. As it turned out, the two-Volta configuration was a sweet spot for throughput of the communication subsystem.

c. Applications Confront and Overcome Architectural Turbulence

Looking back, LLNL’s experience with the Blue Gene series provided many lessons. Managing scalability and memory constraints continued to be issues in future architectures, as was maximizing those machines for throughput, at which they excelled. But LLNL also recognized that the low-powered, many-core-processor approach was not a winning strategy for capability computing, and shied away from Intel’s Xeon Phi series, which was eventually discontinued. IBM also abandoned the low-power, many-core-processor approach and in 2014 responded to the CORAL RFP with an intriguing architecture that married IBM Power9 processors to Nvidia GPUs, as previously described.

LLNL and ORNL were sufficiently impressed with IBM’s work in showing the projected performance of their proposed architecture on CORAL benchmark applications that they chose the IBM/Nvidia hybrid architecture. This was a major gamble: while GPUs showed some promise in running limited applications, LLNL had not successfully demonstrated a large multiphysics application on GPUs. Even the demanding benchmark applications did not adequately address the challenges presented by these huge, multipackage codes. Yet LLNL felt the potential payoff was worth the risk. CPU-based HPC systems were running out of runway to continue their decades-long improvements based on increased clock speeds and massive processor counts and were projected to provide modest gains in application performance relative to Sequoia. But IBM’s proposed architecture might lead to more than an order-of-magnitude increase in capability if the daunting hurdles of optimizing LLNL’s scientific codes for GPUs could be overcome.

The work that went into modifying code to optimize for GPUs was more taxing than the preparation for any previous machine in the ASC era. For a modest two- to four-times improvement in computational power, it might not have been worth the investment—but the LLNL code teams were shooting much higher. An estimated 25 percent of all coding efforts over four years was spent on code optimizations preparing for Sierra. As discussed later, RAJA was becoming widely adopted by the code teams and made the job of optimizing for Sierra easier while also avoiding the dreaded lock-in of vendor-proprietary models that make portability nearly impossible. Much of the work, however, was in the modification of core algorithms, which fortunately had a positive effect on code performance across other platforms, including CTS, which ran most daily design work. Optimization for GPU performance generally involved data layouts and orders of operation that translated to improved efficiency on CPUs as well. Over the years, many algorithms were re-derived and adapted to map onto GPUs. The code teams embarked on independent efforts to expand, redesign, and retrofit their code bases, and as the individual code teams advanced, new capability was added to RAJA based on what the teams learned. Thus, the new technology became sharable across applications. The full value of RAJA's performance portability will play out as new systems with other nonstandard programming environments are deployed, including the anticipated El Capitan and AMD's ROCm and HIP programming models.

One major Sierra technology advance that made it practical for general-purpose computing was coherent memory between the GPU and CPU. This freed the programmer from moving data back and forth between the two memory spaces. Data movement was a major limitation on Roadrunner with the Sony, Toshiba, IBM, or STI chip used as an accelerator and resulted in LANL code that was never integrated into final products once Roadrunner was decommissioned. While this unified virtual memory never performed well enough to rely on continuous back-and-forth data motion between CPU and GPU memory, codes could rely on it to page-in memory to the GPU during problem or package startup behind the scenes, with little or no explicit programming. Umpire and CHAI products, which provided efficient memory management, were developed alongside RAJA to round out a suite of performance-portable solutions that applications could adopt in whole or part, keeping with the overall philosophy of allowing projects to retrofit portability as they saw fit.

Figure 49. Cover of S&TR magazine featuring Sierra, March 2017.

To assist with this disruptive transition, the CORAL procurement asked vendors to bid NRE that would fund a vendor partnership with the lab teams. Thus, the Sierra Center of Excellence[17] (COE) was born, uniting lab code teams and experts from IBM and Nvidia to work on application preparation. Because LLNL's M&IC program intended to purchase a smaller unclassified version, a similar institutional COE was formed for codes outside the core weapons program.

Steady progress ensued as code teams tackled the long poles in the tent, only to uncover the next limiting factor. Crucial to this work was the all-in attitude and collaboration of the code teams


S&TR March 2017

Supercomputing Advancements

A Center of
Excellence
Prepares for
SIERRA

IN November 2014, then Secretary of Energy Ernest Moniz announced a partnership involving IBM, NVIDIA, and Mellanox to design and deliver high-performance computing (HPC) systems for Lawrence Livermore and Oak Ridge national laboratories. (See *S&TR*, March 2015, pp. 11–15.) The Livermore system, Sierra, will be the latest in a series of leading-edge Advanced Simulation and Computing (ASC) Program supercomputers, whose predecessors include Sequoia and BlueGene/L, Purple, White, and Blue Pacific. As such, Sierra will be expected to help solve the most demanding computational challenges faced by the National Nuclear Security Administration's (NNSA's) ASC Program in furthering its stockpile stewardship mission.

Lawrence Livermore National Laboratory 5

S&TR March 2017

Application developers and their industry partners are working to achieve both performance and cross-platform portability as they ready science applications for the arrival of Livermore's next flagship supercomputer.

4 Lawrence Livermore National Laboratory

Figure 50. A feature article on the preparation for Sierra (str.llnl.gov/march-2017).

---

[17] "Application Modernization at LLNL and the Sierra Center of Excellence", IEEE CiSE vol 19 issue 5, https://doi.org/10.1109/MCSE.2017.3421556

and libraries. The new modular-development strategy embraced several years before proved useful as optimizations gained in each physics or computer science (CS) package were realized across multiple code projects. Large multiphysics weapons applications eventually saw 10 times or better improvements in turnaround, compared to an equivalent number of nodes on the CTS machines. Similar speedups were seen in comparing a Sierra node's CPU performance with its GPU performance. The largest speedups were with codes that ported nearly all the heavy computing to GPUs, reserving CPUs mainly for managing work distribution and small serial sections of computing.

The real wow factor for Sierra, however, was not the ability to run CTS-sized problems 10 times faster, but its running of capability calculations of unprecedented size in 3D. Three-dimensional simulations that would have taken weeks on any other machine could be turned around in hours or days on Sierra, and scientists were able to run problems at resolutions never attempted. Sierra thus established itself as a quantum leap forward in capability. That one could run moderately sized 3D calculations on a couple percent of the machine also meant it could be used as a throughput engine, allowing quick turnaround for 3D ensembles. This provided a totally new avenue for the weapons program, which had previously run large validation suites in 2D or even 1D because of the severe limitations of computers and codes written for earlier, less capable computers. The key was 3D, as GPUs were very good at processing massive data quickly; but struggles with efficiency surfaced if the pipelines were not kept full. Three-dimensional problems generally met that requirement, but not 2D, which still represented much weapons work. Speedups for 2D problems were modest—two to three times, typically—but they could be relegated to commodity capacity machines, demonstrating again that LLNL's development of commodity technology was prescient.

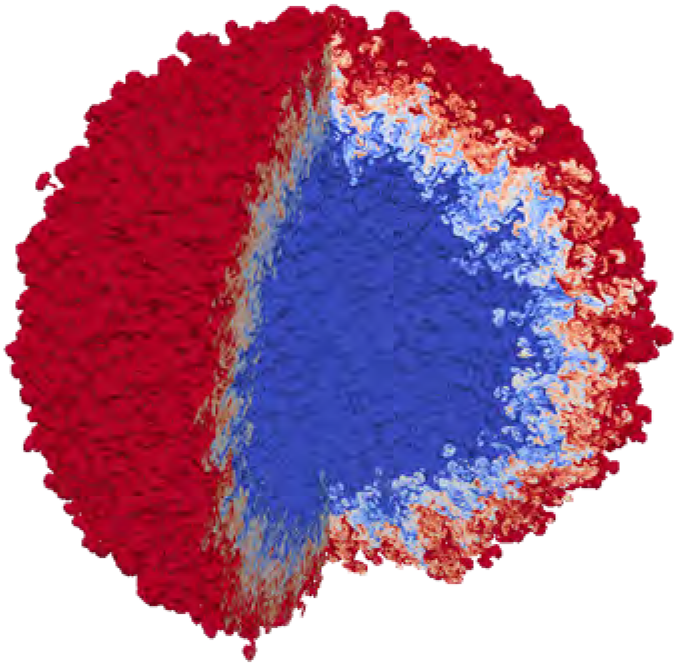

Figure 51. A 97.8-billion-element simulation run on Sierra showing 3D growth of an ICF-inspired Rayleigh-Taylor instability.

In sum, the original ASCI vision to run large 3D problems routinely was finally realized in Sierra. As we have emphasized throughout, what initially seems impossible is in most cases possible, but takes longer. In this case, the 1996 vision was essentially accomplished by 2020.

Early in Sierra's integration, weapons-code development leader Brian Pudliner projected,

> *Design and analysis in LLNL's weapons program until today relied primarily on 2D approximations because 3D simulations could not be turned around quickly enough to make them a useful routine design tool. Sierra's architecture, which is expected to bring speed-ups on the order of 10X for many of our 3D applications, will be able to process these crucial simulations efficiently, changing the way weapons designers work by making the use of 3D routine.*

The Next-Gen code project leader, Rob Rieben, added,

> *We are also developing next-generation simulation codes for inertial confinement fusion and nuclear weapons analysis that employ high-order, compute-intensive algorithms that maximize the amount of computing done for each piece of data retrieved from memory. These schemes are very robust and should significantly improve the overall analysis workflow for users. These advanced simulation tools enabled by Sierra will improve throughput along two axes: faster turnaround and less user intervention.*

It would be foolish to say that we are the end of history and nothing more is to be gained by pursuing better machines, physics models, artificial intelligence (AI), and resolution. It is indisputable, however, that LLNL can envision and accomplish great things by the relentless perseverance at the heart of the Lab. From 1996 to 2020 was a span of 24 years. Building NIF and achieving ignition took even longer. But LLNL's great advances, like NIF ignition, 3D multipackage advanced-physics models, and modern nuclear weapons took extraordinary perseverance and were miracles of human creativity. Sierra and the applications that run on it are a superlative example of what LLNL can deliver to the country.

Sierra and Summit, ORNL's sister platform, took second and first, respectively, in the Top500 list of 2018. The list was soon dominated by GPU-based architectures, making it clear that GPUs had won out over competing architectures, at least for the time being, and gave the most computational bang for the energy buck.

While the focus of this retrospective has been platforms and the work necessary to make them perform, it is important to keep sight of the bigger picture, that is, how increasing computer performance fits with meeting weapons-program needs. Running today's problems faster or in 3D or at higher resolution are all key in delivering on the stockpile-stewardship mission. But without investment in better physics understanding and models and more accurate and performant algorithms, the gains from compute power would be far more limited. Just as moving hardware forward takes years of investment, improved models and algorithms are long-term projects that may take years to mature to the point of reliability in mission-critical applications.

Throughout the progression from ASCI's early days to exascale today, investments in making simulation tools faster and more predictive have reflected a careful balancing act. Because better physics models and more accurate and complex algorithms frequently require more computing power to be practical, advances in models, algorithms, and computing power go hand in hand. The decades of code work on new models and algorithms, the continuing tight integration of the ASC Physics and Engineering Models (PEM) subprogram into our code and platform development, and the efforts to improve performance on the hardware have combined to render the machines on the floor of primary importance for mission applications.

Figure 52. LLNL researcher Arthur Rodgers with a 3D seismic-simulation image.

d. The Third Leg of the Scientific Method: Seismology and Predictive Biology

Seismology

The advancement of nuclear weapons codes to routine 3D are just one example of the growing capability of the discovery triad's third leg. Seismic modeling is another powerful example.

The motion of a strong earthquake near a fault is highly variable and poorly constrained by limited empirical data. Supercomputers enable the simulation of earthquake movement to investigate hazards and risks to buildings and infrastructure before damaging events occur. The large scale (about 100 km) and fine detail (about 10 m) of high-frequency seismic waves greater than 5 Hz require today's most powerful computers.

Using SW4-RAJA, a 3D seismic-simulation code ported to the GPU hardware on Sierra, LLNL researchers increased the resolution of earthquake simulations to span most frequencies of engineering interest on regional domains, and rapid throughput enables the sampling of various rupture scenarios and subsurface models. As Arthur Rodgers wrote regarding Hayward-fault modeling, "Sierra's thousands of GPU-accelerated nodes allow SW4-RAJA to compute earthquake ground motions with hundreds of billions of grid points in shorter run times so we can resolve high-frequency waves and investigate different rupture scenarios or earth models."

Some SW4 work can be seen in the "Exploring Earthquakes" exhibit at the California Academy of Sciences, San Francisco.

Predictive Biology

The importance of biosystems to national security and global health was made clear by the COVID-19 pandemic. The interaction of biological systems like viruses and humans is complex. Predicting behaviors and designing interventions relies on large quantities of data integrated with complex models. HPC is the integrating element that enables these models and their applications. The convergence of high-fidelity simulation and AI-based models with automated experiments is enabling a new generation of predictive-biology models and molecular-design systems.

The Generative Unconstrained Intelligent Drug Engineering (GUIDE) program for rapid monoclonal-antibody design provides a good example. LLNL scientists have demonstrated the computational design of antibodies to target future variants of viruses like SARS-CoV-2. These designs are based on predictive forecasting models for viral mutations, generative AI models that propose new antibody sequences and structures, and mechanistic-simulation and machine-learning (ML) models that predict their properties. It is now possible to design new therapeutics in weeks, not years, and get ahead of biothreats.

The GPU-accelerated architectures of Sierra and Lassen are a good match for the computational needs of biodesign applications. In a single design-optimization cycle, we can run millions of molecular-dynamics calculations to evaluate free energy for an antibody binding to a viral protein target. We can also train and evaluate ML-based property-prediction models for human-likeness of the antibody structure, as well as manufacturability properties like thermal stability. These operations are efficiently accelerated using GPU architectures.

The GUIDE team demonstrated this approach concretely in responding to a rapid-response antibody-design request in December 2021. The newly emerged Omicron variant of the SARS-CoV-2 virus had rendered all available antibody therapies ineffective. The spike protein on the virus had mutated sufficiently that existing therapeutics could not bond and neutralize the virus. The GUIDE government sponsor asked LLNL to run a rapid redesign of an existing commercial antibody. Allocating a significant fraction of Sierra, the team produced a complete design optimization in less than three weeks. The top molecular designs from that process were experimentally validated and shown to restore the full potency of the antibody therapeutic.[18]

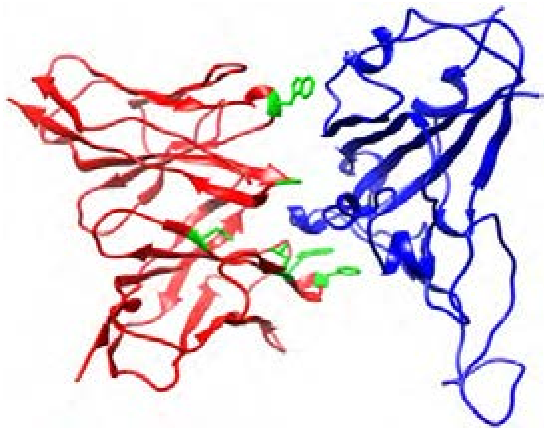

Figure 53. The receptor-binding domain of a computationally designed antibody (red) engaging a SARS-CoV-2 spike protein region (blue). 20M CPU core hours and more than 1M GPU hours were required to reach a design validated to bind and neutralize the live virus.

Current work is just scratching the surface of potential biological-design applications. New methods for training and inference from extremely large-protein foundation models are increasing performance and reducing uncertainties in molecular design. Coupled with new biotechnologies for producing trillion-token experimental datasets for protein–protein interactions and powered by next-generation computer architectures, these models will vastly extend our capabilities for protein- and biosystem engineering. From cancer to neurological disease, application areas will be transformed by these methods. We're addressing the simplest classes of targets in these initial efforts. New methods to model and modulate the full biological regulatory and metabolic network of cells will ultimately be possible under advances in computing, AI, and automated laboratories.

---

[18] Desautels et al. https://www.nature.com/articles/s41586-024-07385-1).

12

EXASCALE AND EL CAPITAN

a. DOE's Quest
b. The El Capitan Procurement: Trials and Triumphs
c. The RFP, Selection, and Aftermath
d. Exascale Computing Facility Modernization

Figure 54. El Capitan (2023): greater than 2 exaflops peak performance.

Exascale and El Capitan

a. DOE's Quest

It is beyond the scope of this document to recount all the travail involved in building congressional and administration support for the $4B exascale initiative and NNSA funding for ASC's procurement of a ~$600M computer. The bigger the numbers, the greater the scrutiny and processes and marketing and supplication required, conducted in an impenetrable fog of uncertainty. NNSA computing budgets were strained and had been in general decline since the time of the Purple acquisition at the end of the ASCI period. To get to exascale, ASC budgets needed to grow very substantially in the face of enormous investments to rebuild the complex. The rendition below provides context to understand LLNL's contributions over a period spanning a decade.

In 2010, Steve Koonin, the DOE undersecretary for science (responsible for coordinating and overseeing DOE research), became aware of the problem in retaining U.S. leadership in computing and contemplated how DOE might keep its leadership by focusing on exascale. Chinese number-one systems, despite their inefficiencies with real applications, were becoming a public challenge to continued American supremacy. David Dean, an ORNL detailee working for Koonin as a senior adviser, was tasked with penciling out how such an initiative might be organized. He developed a blue-sky planning exercise for a DOE exascale initiative, including viewgraphs that were widely shared. In hindsight, it is remarkable how well Dean captured what was needed and how it might be done.

Later in 2010, ASC and ASCR asked four Office of Science labs (ANL, ORNL, LBNL, and BNL) and the three NNSA labs (LLNL, LANL, and SNL) to form a steering committee to propose an exascale initiative. The committee was informed by workshops sponsored by the Office of Science to gather requirements and by technical meetings of lab computing leaders and visits from potential vendors and technology providers. The meetings were generally held near the Denver airport to even out the pain of travel. In 2012, the DOE undersecretary, Steven Chu, submitted the "DOE Exascale Strategy," as requested by Congress. This was a high-level outline of a plan to provide capabilities roughly a thousand times more capable than current systems, within a reasonable power envelope, and applications co-designed with the computer technology to be developed.

On July 29, 2015, President Obama established by executive order the National Strategic Computing Initiative (NSCI)[19] to maximize the benefits of HPC for U.S. economic competitiveness, scientific discovery, and national security. DOE was responsible for executing the Exascale Computing Initiative (ECI). By 2016, while copious new funding was short, major programmatic efforts at Office of

---

[19] https://www.federalregister.gov/documents/2015/08/03/2015-19183/creating-a-national-strategic-computing-initiative.

Science labs were reprogrammed to focus on developing an exascale technology and applications. ASC secured modest supplementary funding for code development and CS (discussed below). A memorandum of agreement (MOA) was signed among the three key NNSA and three Office of Science laboratories describing how the initiative—now a formal DOE project—would be managed among them. The MOA included Paul Messina from ANL as the initial project director and a board of directors, chaired for much of this period by LLNL director Bill Goldstein.

The scope of the project encompassed four focus areas: application development, software technology, hardware technology, and exascale systems. The new Exascale Computing Program (ECP) would be managed according to DOE order 413.3B, Program and Project Management for the Acquisition of Capital Assets, tailored for the ECP.

Order 413.3B is a DOE process intended for standard capital acquisitions with clearly defined milestones; it was not well suited for an ostensible research project. Recognizing that the ECI was anything but standard and alarmed at the reporting overhead demanded by 413.3B, ECP and the HQ program offices modified the process into something manageable for the technical teams while maintaining a defendable structure. Having succeeded Messina as project leader, Doug Kothe (later the ORNL AD for computing and then SNL chief research officer) focused primary attention on technical output while convincing overseers in the administration and Congress that dollars were spent responsibly.

During this period, ASC saw fewer funding plus-ups (budget increases) than Office of Science, given the overall stress on NNSA budgets. Moreover, some annual national-security deliverables could not be dropped by ASC to reprogram dollars into ECP deliverables. This meant that NNSA's nominal dollar contribution was $1.4B, as opposed to ASCR's $2.5B. Discussions that focused on applications development and CS tools from CASC, LC, and ASQ amplified this point. The number of LLNL CS tool deliverables to the project that were ultimately picked up and used by the broader community was a clear indication, however, of the broad appeal and high quality these products embodied. The national-security codes were also expertly aligned with mission needs and the approaching architectural solution. In short, NNSA contributed heavily to the areas it most cared about, and LLNL software-tool development was prominent among these.

b. The El Capitan Procurement: Trials and Triumphs

Funding Challenges

The CORAL2 exascale procurement RFP was issued from ORNL, as agreed by the three labs (as mentioned earlier, the CORAL1 process that delivered Sierra was issued from LLNL under Terri Quinn and Bronis de Supinski of the LC). While Quinn and Supinski provided technical and administrative leadership, specialist Gary Ward enabled most of the procurements listed in these pages. Every procurement raises complex issues and may invite novel strategies that are viewed as unorthodox up the chain. Rather than retreating under pushback, Ward worked cooperatively with LC and officials at all levels to find an approach that worked. The Blue Gene procurement model, used later for most large acquisitions, is an example of Ward's expertise. His work with officialdom to reach accommodations led to a viable model for the acquisition of advanced architectures.

For the CORAL2 procurement, the first tiny detail was funding a $600M system. NNSA budgets were under pressure. This obviously affected the ASC HQ bottom line, which had been on a downward trend for a decade. Nonetheless, at an ASC principal investigators (PI) meeting in Monterey, Doug Wade, who oversaw ASC HQ, and his deputy Thuc Huong, acceded to McCoy's increasingly desperate arguments to get the procurement process started with the CD-0 requirements document. Given the pressures at NNSA HQ, this was not an easy decision,

To add to the uncertainty and stress, a copy of the ECP president's budget for FY18 showed a planning number of $500M for an ASCR exascale machine in FY21 and $250M for an ASC machine in FY19. Even with a year's lag, $250M would be insufficient to procure an exascale-class system. The FY19 Future-Years Nuclear Security Program (FYNSP) budget conveyed the message that NNSA would not field an exascale computer.

Then seeming magic happened. On March 26, 2018, Rick Perry, the secretary of energy, visited LLNL. He had heard that LLNL would not field an exascale system, and when he got to the Lab, asked if this were really the case. From his perspective, the exascale initiative had been advertised as vital for U.S. leadership in science and nuclear security; how could it be that the nuclear-weapons program that essentially invented HPC wasn't interested in an exascale system? The Lab responded that the FY19 FYNSP for ASC precluded the hamstringing of critical weapons deliverables to site a computer. At the end of his visit, Perry asked LLNL to prepare an ASC budget that would make delivery of such a system possible. Director Bill Goldstein asked ASC program leader McCoy to prepare it, and McCoy calculated the annual plus-ups needed to cover a $600M (not $500M!) system over the FY19 FYNSP period. Thuc Hoang, later to lead the ASC program from HQ, provided input and concurrence with McCoy's calculation; the total over-target request was about $500M over the five-year FYNSP.

In April 2018, the LLNL Director's Office responded to Secretary Perry, and ASC budgets were adjusted upwards for a time —not to the full $500M, but sufficient to request proposals for exascale-class technology commensurate with ASCR numbers. While pressures on ASC budgets continued, technical and COVID issues intruded, affecting the supply chain and technology delivery. The planned deployment of the system was delayed until FY24, relieving some pressure on the ASC budget to pay the total cost on a quicker timescale (systems are paid by lease to own, or LTO. Like mortgages, LTOs allow payment over multiple years so the burden can be spread over time). It is very likely that if Perry had not raised the issue, the new LLNL system would have been less than exascale class.

c. The RFP, Selection, and Aftermath

ORNL issued an RFP in early 2018. Bronis de Supinski, Robin Goldstone, Matt Leininger, and their manager, Terri Quinn, provided extensive LLNL input to the RFP, and LLNL was confident its needs were reflected well. By midyear, four creditable responses were received, and preliminary decisions were made. LLNL and ORNL selected the same vendor (Cray) and ANL would most likely work with Intel. By collaborating, LLNL and ORNL could share the cost of the NRE contract. In the coming months, Cray was absorbed into HPE.

Requirements in the RFP were guided by key performance parameters (KPPs) as seen in Figure 55.

LLNL wanted and expected to do much better than 10 to 12 times' improvement in performance but was also pretty sure that power consumption would tend towards the threshold limit of 40MW. One mitigating factor was that the LLNL machine would be deployed a year after the ORNL system. From LLNL's perspective, the Laboratory warranted improved performance for each dollar spent as a result of this differential in time to delivery.

For the first time since ASCI began, LLNL declined to go with IBM for a major system. This was not an easy decision; the Lab had flourished during their long partnership and, indeed, friendship. But evaluation processes are formal things, leaving little room for sentiment, and the reviewers and ultimately management were committed to disinterested decisions based on their expertise, judgement, and what they had read. After this setback, IBM backed away from the HPC arena, removing a giant from the constellation of vendors willing to bid in the future. The absorption of Cray into HPE removed an additional bidder, leaving only HPE and Intel. How this will affect future NNSA and Office of Science procurements is yet to be seen, but it does not portend well. LC strategies to address this problem depend on its control over the software stacks it manages, which largely liberate the Lab from vendor software.

Cray's winning proposal contained multiple component options, including nodes, from a variety of potential subcontractors. This meant that the labs had leverage to compete further for the nodes within the envelope of the vendor, Cray. In addition, since LLNL was going a year later than ORNL, it could and did argue for improved technology

Requirements	Threshold	Objective
Performance improvement for benchmark IDC over Sierra	5x	10x
Performance improvement for full system science and benchmark IDC over Sierra	6x	12x
Performance improvement for large ensemble/throughput simulations over Sierra	6x	12x
Aggregate memory addressable from codes	4 PB	10 PB
Mean time between application failure due to system fault requiring user or administrator action	4 days	6 days
Maximum power consumption (system and peripherals)	40 MW	20 MW

Figure 55. Table of requirements, thresholds, and targets in design documents presented to HQ.

and performance. All this took considerable time and effort, but in the end, ORNL went with an AMD node modestly enhanced from the MI200 bid (originally called the MI250). This had been offered to LLNL as the improvement requested, but both AMD and LLNL thought the additional year might offer greater improvements. After considerable work, AMD advanced the delivery timeline of their next-generation MI300A, which offered a single instance of high-bandwidth memory across GPUs and CPUs. The product was to become the first realization of the chiplet-based accelerated processing unit (APU) vision imagined in the earliest days of the FastForward[20] program, over 10 years before. This created space for ORNL to argue successfully for the MI250.

El Capitan is based on the Cray Shasta system, the newest line of Cray supercomputers, (now offered by HPE). Included was a system-software stack to operate and manage the system. Much of this software was new and represented a new architectural approach. As El Capitan was not the first DOE system to take delivery of the Cray Shasta system, LLNL observed experiences at other sites with the new software (e.g., ORNL). Lengthy ordeals materialized in testing and stabilizing systems. With the inevitable hardware delays around the MI300, El Capitan's schedule contingency had been reduced, and the Lab could ill afford major delays from integrating the vendor software stack on the MI300 architecture. Fortunately, LLNL was in a unique position to address this risk with a mature solution, TOSS, and had decided to use TOSS on El Capitan about three years before delivery of the system. While this risk-mitigation decision put LLNL and the experienced TOSS team in control, it also put the Lab to a highly visible test of putting its money where its mouth was. TOSS is running on El Capitan early test systems and is currently being installed on El Capitan. This strategy will certainly benefit LLNL in the post-exascale era. The ability to take control of the system-software stack allows the Lab to consider a wide array of vendors for future systems and may prove essential to being first to support novel architectures.

The flexibility this capability provides to LLNL is hard to overstate. One need only consider that IBM is no longer is building large systems, and Cray has been absorbed into HPE. There is a vanishingly small set of potential vendors for full-feature systems, but a site with a software stack takes enormous pressure off potential vendors who can provide hardware, but not proven software. This capability expands the bidding space to LLNL's advantage.

As noted earlier, with each system procurement comes a life-threatening crisis. For El Capitan, this was not technical, despite delays and glitches in the MI300A. It was related to NNSA approval of the financial structuring. ASC's plan was to use staged financing to cover portions of the costs for the acceptance-contract milestones. LLNL began the process early on, as this was not the Lab's first procurement battle and staff knew anything and everything could go wrong. Though LLNL had used lease-to-own financing for two decades, there were two new challenges to overcome: the magnitude of the financing, which was raising eyebrows, and a new DOE approval process accompanied by approval delays that threatened the schedule. Extraordinary efforts by LLNL's legal, financial, and contracting personnel, strong support from LLNL's senior managers, including the Director's Office, and NNSA's ASC program managers ultimately secured the approvals. It should be noted that though HPE was seriously concerned, its management stood by LLNL, confident that ASC and the Lab would persevere. Honest communication among the parties was essential to preserving the partnerships. Few battles are decisive into the infinite future. The LTO had been in place ever since John Fitzgerald of the LLNL NERSC had tried it out for Cray procurements. Over the years, multiple LTOs gave ASC and M&IC the ability to fund large systems, even if not every dollar was available in the year of acceptance. Sometimes it took a little persuading; but as time went on, the process became routine—until it wasn't. Complacency can be fatal in a highly regulated environment where moving forward requires an array of approvals from people who are distant from direct responsibility to deliver and new to the responsibility of oversight. Occasionally, one is vividly reminded of this reality.

d. Exascale Computing Facility Modernization

Much has been said here about the intention of creating a B453 that could adapt or expand for future needs under a supplementary, rather than replacement, approach. Ponder for a moment if the building design were insufficiently flexible: the cost of placing a structure elsewhere, with adequate power and cooling for exascale and beyond, would have doomed LLNL computing. A $350M new-facility

---

[20] https://asc.llnl.gov/exascale/fast-forward.

line item would have sunk the ship. The B453 site had been selected partly because there was a large parking lot to the west where substations could be erected. And there was 48,000 square feet of basement space, ample to install additional electrical and mechanical equipment.

By 2014, long before budgets had solidified for El Capitan, LC facility leaders under Anna Maria Bailey understood that power and cooling to the site would be insufficient to site two exascale systems simultaneously. LLNL gathered input from multiple HPC vendors regarding projected weight, power, and cooling requirements for the next two planned LLNL systems, El Capitan in 2023 and a follow-on in 2028. With this information as a basis, ASC program managers at LLNL requested NNSA support for a line item called the Exascale Computing Facility Modernization (ECFM) project. This was the Lab's highest-priority line-item request. In response, the NNSA ASC program office defined five objectives for its computing environment, as listed in Table 1.

The capacity to support multiple platforms at once offered NNSA two advantages. First, the implementation of a new Advanced Technology System (ATS) takes roughly a year. To minimize operational disruption, the old system and the new must coexist for about 18 months. Second, should NNSA choose or need to consolidate at one ATS site, it would be imperative that the systems operate simultaneously in a robust and efficient facility during the lengthy transition period. In this scenario, it is possible that two such systems (possibly with varying architectures) could coexist for long periods.

Figure 56. The electrical switchyard and cooling towers developed for the ECFM program, a few hundred feet west of B453.

A line item and funding within the ASC budget were required. The need was obvious to LLNL management, and the how to do the project was equally well understood by facility management. But NNSA construction and line-item leadership at HQ were not as sanguine regarding the Lab's abilities. Hence, the LC was subjected to the rigors of

Requirements	Threshold	Objective
Adequate square footage to handle exascale systems and their environment	15,000 ft2	24,000 ft2
Power capacity to meet demand from exascale systems	85 MW*	110 MW*
Increase water capacity to implement innovative mechanical liquid-cooling solutions	18000 tons	25000 tons
Accommodate heavier racks (4' x 4')	315 lb./ft2	500 lb./ft2
Ensure sustainable and energy-efficient facility solutions are implemented	PUE** = 1.08	PUE** = 1.05

Table 1. Key facility requirements for Exascale Computing Facility Modernization.

*Does not include power for file systems, which is small compared to the computer itself

**PUE = power-usage effectiveness—a measure of how much energy is used by computing equipment (in contrast to cooling and other overhead). PUE = (total facility energy) / (computing-equipment energy) complying with NNSA processes. After the smoke cleared, the project was completed in late 2021, ahead of schedule and below the $100M budget. It was an extraordinary achievement, especially considering that much of the work was done during the COVID shutdown, and a testimony to the prescience and professionalism of Anna Maria Bailey and her team.

13

SUMMARY OF MAJOR OUTCOMES FROM THE MASSIVELY PARALLEL ERA AT LLNL

Summary of Major Outcomes from the Massively Parallel Era at LLNL

This account begins not with a rendition of ASCI achievements, but with the questions and challenges faced at LLNL (and all weapons labs) before the ASCI period. These included transitioning from vector computing to a new era of "killer micros" without any real understanding of how these disruptive tools could be harnessed. Indeed, a clear winner in the message-passing arena was not yet identified. Many of LLNL's grassroots explorations are highlighted here. A new era was dawning, with what seemed to be gigantic numbers of low-cost processers working together on complex sets of coupled partial differential equations. This era reached its apex with the first exascale computers, and El Capitan may well be the crowning achievement of that period.

Today, in 2024, what comes next is speculation. How the labs speculate and the actions they take to explore are important. The future will likely be a more-complex era that builds on everything described here but adds cognitive and "intelligent" tools, exploiting increasingly sophisticated cloud infrastructure (possibly relying on massive data and increasingly heterogeneous hardware) that might provide an epic improvement in predictivity.

Hopefully, this tour through history, including the dark period preceding ASCI, demonstrates LLNL's genetic ability to face uncertainty and prevail through the technical brilliance of its scientists and their dogged perseverance. Let us summarize what the last 30 years produced in HPC at the Lab. Much of this product is remarkable, and to our credit, some was even planned. But most was unimagined in 1995 and invented as we went along. With a nod to Bruce Goodwin's comment that it is better to be lucky than good, we just got very lucky—even if we might also argue that the harder we worked, the luckier we got. Louis Pasteur is often credited with observing, "Chance favors the prepared mind."

1. **Original ASCI goal** The original ASCI goal of entry-level 3D full-system calculations was achieved at the end of the ASCI period. Such calculations were run on ASCI Purple and reflected a forced march to 3D, which, despite its drawbacks, was necessary if only because it appealed to the country and provided a goal that people could understand. We now know this was just the end of the beginning.

2. **Rite of passage for 3D weapons codes** Routinely resolved 3D weapons simulation was achieved on Sierra around 2020, and very substantial refinements will be seen on El Capitan. Critical to this success was the evolution of LLNL weapons codes to adapt to new architectures and maximize their capability. The original ASCI vision was at last achieved, with LLNL codes at the forefront. The magnitude of this advance can be fully understood only by the computational physicists, computer scientists, engineers, and designers who wrote the codes married to the complex computer architectures, tested and validated them, and now use them in anger on next-generation weapons systems.

3. **Institutional computing** The Director's Office understood early on that ASCI could not operate in a vacuum at LLNL. All scientists and engineers had to be enfranchised with access to world-class computing if it were to be one lab and not a Balkanized federation of winners and losers. The M&IC model was therefore put in place. This document highlights the fruition of this idea. LLNL would not have achieved its stature in computing today had it not been a groundbreaker for NNSA in this area.

4. **Mission focus** The enormous success of LLNL computing is largely a consequence of a fundamental and relentless focus on the Laboratory mission. This commitment began with the foundational national-security mission and extends to the full range of Laboratory science and technology programs. An essential element is continuing, robust, two-way communications with the user community, in which immediate feedback (positive and negative) is sought from the users, information is gathered about future mission needs, and open communication of candidly bad news (such as near-death experiences that require mitigation) and good news (like increased system capabilities that might require a modified system or user software to access). Decisions—even high-risk decisions—were always focused on developing and delivering capabilities to meet mission needs. Basic research, whether in computer architecture, system software, applied mathematics, or other, was guided by long-term mission needs. The history of LLNL computing demonstrates that decisions made long ago almost invariably yielded essential mission capabilities in the future.

5. **Multidisciplinary application teams** Through a combination of culture and organization, ASC application teams have always thrived in diverse disciplines. Computational physicists, engineers, mathematicians, and other specialists have ensured cutting-edge and robust algorithms lie at the core of our capabilities, and computer scientists and software engineers have developed scalable, robust, and maintainable applications that are increasingly proof against technological surprise. Over time, titles blur as team members learn from one another and venture out of their lane. That's when the magic of multidisciplinary teaming happens.

6. **Advanced architectures: getting real about power and cooling** With Blue Gene/L, LLNL sited the first major system focused on making extreme parallelism affordable. The whole NNSA advanced-architecture systems approach had its spark here. Along the way, LLNL invented a procurement model that made it possible for vendors to bid on the leading edge, rather than accepting conservative products built for less demanding customers and then tweaking them. This procurement model is good for the vendors, good for us, and good for the country. It has been used by major DOE laboratories, not just NNSA labs, and the exascale systems now running in this country are products of this advance.

7. **Building infrastructure while avoiding obsolescence** Understanding early that flexibility was critical, LLNL designed B453 with a generous future in mind. That building now sits astride 85MW of power and 28,000 tons of liquid cooling, thanks to ECFM planning, and will serve the nation for decades. There is no other site in the DOE this integrated and flexible: LLNL can and will be able to site two exascale systems simultaneously.

8. **Production Linux clusters** LLNL jumped on the Linux cluster idea and made such capacity systems production ready. While many labs experimented with Linux clusters, the Lab approached the problem systematically, made clusters practical for daily mission workloads, and built increasingly complex systems over time. LLNL borrowed the scalable-unit idea from SNL and built a software stack and scalable computers that were so successful that HQ adopted these as one of its two major procurement solutions (commodity clusters and capability systems). The LLNL SW stack (CHAOS) was adopted by the Tri-lab as the baseline OS. Subsequently, capability systems could be used for purposes like 3D integrated simulations and 3D UQ, rather than competing with capacity runs for access.

9. **SW tools development** LLNL is recognized for its contributions to open-source HPC software development, spanning system software, computational mathematics, build tools, data management and visualization, performance portability, performance analysis, workflows, and application infrastructure. The mission-critical needs of NNSA drove the local development of widely adopted, practical solutions to exascale problems.

10. **Top500 leadership** LLNL has taken a nationally recognized role in siting systems, as confirmed by the Top500. LLNL systems have frequently taken the number one position, and the number of computing frames on the list at any given time has been remarkable and unique. At one point, a capacity system, Thunder, took second place in the Top500. LLNL leadership generates global credibility; and credibility can be taken to the bank when need arises.

Finally, as discussed in the historical background, LLNL pitched for a greater role for computing in the years between the end of UGTs and the start of ASCI, actively promoting the Numerical Test Site vision. Nevertheless, the ASCI program did not self-generate and appear out of the ether at the behest of some genie. It took grassroots work at all three labs and HQ to clear the fog and define the path. Today, we are clearing the fog away to view the next era. The ideas discussed in the following section might set the stage for tomorrow's adventure.

14

EPILOGUE–LOOKING FORWARD

Epilogue—Looking Forward

This document captures a generational effort in computing at LLNL. Beginning with the rise of parallel computing and journeying through low-power lightweight cores before pivoting to heterogeneous computing and GPUs, the work culminated with the delivery of an exascale system in 2024—a 50-millionfold increase in peak floating-point performance. The Lab rode an exponential curve where each system wasn't merely incrementally faster, but replaced each predecessor's capability with 5–25 times the compute power of before. It has been said that a factor of ten is not just a quantitative improvement, but a qualitative improvement too, because new things can be done. LLNL kept pace with hardware advances, developing impressive, complex applications and software so that these platforms were true scientific apparatus. Combined with the myriad NNSA experimental facilities used to validate our codes, LLNL's science-based stockpile stewardship has met or exceeded all expectations. Skeptics in ASCI's early years who asserted an altered or refurbished weapon would never be fielded without confirmatory underground testing have largely fallen silent. We are confidently and judiciously forging a new era of stockpile modernization.

We are at an inflection point between familiar challenges and new opportunities. Below are several topics that will likely define our future in computing and have already influenced our evolving strategy at LLNL.

Post-Exascale

With the advent of exascale, we are reaching a brass-ring capability touted for years: the ability to do large ensembles of high-fidelity, 3D, full-system studies as part of the daily workflow of a designer. One question we must always ask is how much is good enough. With each new system procured, we have struggled to predict the new discoveries and insights that might ensue. History shows our exceptional track record, insights, cost savings to the complex in each new generation, and significant payback for our relatively modest investments made in computing. Under increasingly dynamic global threats, engagement scenarios, and economic strains on defense budgets, it's difficult to imagine that future increases in computing capability will not justify their cost. Assuming both continuing need and looming technical hurdles for the entire computing industry, will we even have the option of pursuing exponential growth, or have we truly run the course and must settle for incremental gains in delivered performance?

Many of the same challenges we rose to a decade ago—power requirements, levels of parallelism, system complexity, cost—remain daunting as we think about what it will take to carry on to 10, 20, 50+ exaflop systems. Some past challenges in exascale will look trivial compared to the changes required to continue leveraging industry trends—whether they be low-precision floating point, non-von Neumann architectures, or the sheer cost of modern GPUs' skewing the cost–benefit of riding that wave. The 2023 National Academies of Science report[1] on post-exascale computing argues that NNSA has demonstrated that mission needs require that we lead in HPC development. But it also laid bare the daunting challenges ahead. We're right on the heels of a decadal push to exascale that has left code teams yearning for stability and a chance to focus on physics and algorithmic advances. To continue bending the curve to pursue leadership systems based on raw computing power, several key outcomes must hold true: 1) the mission need clearly demands this technology, 2) the disruption of pursuing new architectures is worth the gains, 3) there is strong support in government for DOE to pursue global HPC leadership, 4) staffing budgets and research portfolios support this direction, and 5) the risk reward outweighs alternatives, such as assigning limited resources into algorithmic advances and other critical areas.

Data Management

NNSA and LLNL have data in abundance, but it may not be accorded the importance it deserves. This is changing as stockpile modernization drives the weapons enterprise to integrate in ways not seen since 1980s Cold War weapons production. The need to employ data as an integrating element is captured in NNSA's digital-transformation efforts, which are nascent but gaining traction. The initial goals are complete digitization and sharing of small-scale data files and tools among the labs and production agencies. Beyond that, digital transformation will lay the groundwork for wrangling data as the basis of artificial-intelligence

---

[21] "Charting a Path in a Shifting Technical and Geopolitical Landscape: Post-Exascale Computing for the National Nuclear Security Administration." https://nap.nationalacademies.org/read/26916/chapter/1https://nap.nationalacademies.org/catalog/26916/charting-a-path-in-a-shifting-technical-and-geopolitical-landscape.

advances and decision support. This is where digital transformation meets HPC. We will pursue means to integrate LC's modeling and simulation (mod/sim) capabilities with the broader burgeoning ecosystem across the enterprise. Success will allow us to integrate simulation results with data from experimental facilities, modern manufacturing equipment, and even legacy underground-test data in ways that enable digital twins, AI, and decision support. Data management is the foundation for all this.

Artificial Intelligence

Looking at broad industry trends, large-scale computing is getting more attention than ever—not in pursuit of traditional modeling and simulation, but for AI. Because AI is entirely contingent on managing and harnessing large amounts of data, concentrated work in data management is a necessary precursor. There is ample evidence that AI will transform the Lab's mission; smaller projects have already demonstrated the potential, and we imagine harnessing our data into trained models with speed and insight that traditional modeling and simulation can't achieve. ASC is starting to generate funding opportunities to take AI research into the realm of our core mission in areas like material discovery, design optimization, surrogate models leading to effective zettascale, automated code generation and translation, digital twins, design for manufacturability, and large language models trained and fine-tuned on the Lab's classified corpora of documents and data. This just scratches the surface; each year brings new developments and mission potential. Currently, the Frontier for Science, Security, and Technology (FASST) initiative is gaining support as an aspirational large DOE initiative and hopefully the starting point for a big advance in LLNL computing history.

Cloud Computing

Related to the above trends, cloud computing is proving as disruptive as promised. Hyperscalers like AWS, Google, and Microsoft are filling out the bottom end of the market with cheap access to web servers on demand and fielding large HPC systems dedicated to AI training at scales unachievable to anyone without traditional skills (like DOE's) and deep pockets.

Many think of "cloud computing" as outsourcing hardware needs to an off-premises data center. This approach will not displace the need for a world-class HPC center like the LC anytime soon. The cost is simply too prohibitive for the benefits gained when computing at our scale. However, the cloud offers other advantages and is the source of much hardware and software innovation. It's foolish to ignore it based on cost comparisons alone.

The question mustn't be one of pitting traditional on-premises HPC against off-premises cloud. We must use both approaches to get the best of both worlds. Cloud offers ease of use and a range of compelling software services that shoot across the bow of the traditional HPC center. LLNL is working with cloud partners to the benefit of both parties. We are learning how to bring "cloudiness" into our data center and map a vision by which users burst easily into the cloud for specific parts of less-compute-intensive workflows; our cloud partners are learning how to modify their offerings for large-scale, parallel scientific applications. Meanwhile, hyperscalers are increasingly developing new chips—this is where much of silicon innovation is happening today. This hardware is available only in proprietary data centers, however, so access will require LC and cloud providers to rethink their procurement models for a win–win scenario.

While many questions remain on how cloud computing will play out, the one clear answer is it can't be ignored. We can't beat them; but we can influence them; and we must join them in identifying what future HPC leadership looks like.

Algorithmic Innovation

A hallmark of LLNL application- and system-software development has been success in creating new methods from whole cloth and moving them through the research pipeline into production use. Whether it's new computational physics and math methods, system software and tools, or the myriad software technologies in between, few places on the planet do scientific computing as well as the Lab. Just as exascale took great resources in the past decade, AI is poised to do the same. We must make room for top experts in the field to come here for research, development, and hardening of research ideas into production software.

Software and Productivity

Finally, we must address the persistent need to make our users more productive through increasingly robust software and easy-to-use integrated applications. HPC and mod/sim at the labs have been fairly characterized as difficult to use—certainly relative to most commercial offerings, though the payoff is cutting-edge algorithms, unparalleled scaling, and performance at the fingertips of those who negotiate the learning curve. But there's no reason we shouldn't have both. We must prioritize user productivity without sacrificing the ability to drill down and make the codes sing, as they say. Some answers lie in the topics named above. AI code generation will increase developer productivity by automating tedious tasks; cloud software will inspire service-based approaches; and workflows will become increasingly modular and automated under a common ecosystem of analysis tools. Many open-source tools aim at improving user productivity—for example, Jupyter, a web-based interactive-computing platform used by LLNL Next-Gen simulation tools. The computing ecosystem is much more diverse than at the beginning of the ASCI era. LLNL's continued national leadership in computing will depend not only on local innovation, but on leveraging external capabilities.

a. The Next Five Years

We have presented a spectrum of promising technologies, any of which could become the foundation of a national-security initiative in modeling and simulation. The ASCI effort exploited the extreme parallelism arising from commodity processors, driven by the need for 3D simulation in the absence of integrated underground tests. Many possibilities in software and hardware were explored before a likely path (MPI and clusters) emerged as the key candidate. Today, the field of potential technologies is at least as open and unexplored as in 1990, so making plausible bets is even more formidable. Nevertheless, it is imperative to initiate organized but relatively modest investigations, because in the end, the labs present possibilities for the nation's consideration—not the other way around. These possibilities need to be convincing, and for that there is nothing better than proof of principle. It is up to us.

The programmatic driver is also shifting today, building on the deep understanding our stockpile-stewardship tools have given us, but refocusing on speed and agility across the entire weapon lifecycle, from material discovery to design, manufacturing, and surveillance. With the arrival of a third nuclear peer in China and with Russian aggression reviving the specter of tactical nuclear weapons, the U.S. will likely rethink its strategic stance. It might find, for example, that more options in perhaps fewer fielded weapons fills a deterrence gap. The labs may be called upon without notice to offer alternative nuclear-weapons designs with properties and missions inaccessible to the current stockpile. While we can't select new design properties now, we can think about achieving the high-level design of a device in mere months, for detailed design later, and thereby save years. Several large LDRD projects, including those of Brian Spears in 2018 and Jon Belof in 2020, have explored rapid design through machine learning and AI.

One enormous advantage right around the corner is El Capitan, a colossal simulation and AI engine that will be world class for at least half a decade. The LC also possesses immense secondary storage capabilities. The development of rapid-design capability will surely draw on the ingredients of data management and AI discussed above. Skunkworks with rapid design and assessment capability as a goal could begin with modest national-security and LDRD investments using El Capitan and exceptionally challenging unclassified applications as workbenches. This would call for collaboration with the plants and other labs. Of course, rapid-design capability would be of interest to many industries with product goals far different from ours, but the potential for fertile collaboration is obvious. National nuclear security's growing needs for rapid design and assessment would attract industrial players seeking a committed partner. Failing to move forward today with some target technology, whether rapid design or other, would squander a rare opportunity afforded by AI and El Capitan.

b. Parting Thoughts

Our advanced computers should be considered our grandest experimental facilities—designed both to confirm what we think we know and lead to new insights and discoveries. Computing with the best systems and codes in the world is much like having access to the James Webb space telescope. We aim it at something and have a pretty good idea that we're going to see familiar things in better detail than ever before, confirming or denying our theories and models. But we sometimes look far and deep enough with a new instrument to surprise ourselves and discover things we didn't know. Discovery unleashes the creative scientific mind to make predictions of what phenomena mean and what may lie beyond, driving the pursuit of the next more-powerful instrument that we didn't dream we needed.

We anticipate that El Capitan and succeeding systems will lead to discoveries and insights otherwise impossible. Rapid design is a condign focal point for such an effort, comparable to a focal point for exploring a galaxy 13 billion light-years away. We wish to reinvigorate discussion of HPC and confirm that our computing leadership is something we cannot cede. We must not insist that every new computer justify its procurement by expressing rigid predictions of what it will allow us to do.

Declining to swing for the fences is the biggest risk. We trust that, commensurate with the achievements captured in this document, the next three decades of computing at LLNL will tell an extraordinary story.

15

ACRONYMS

2D	two-dimensional
3D	three-dimensional
AD	associate director
AI	artificial intelligence
AMR	adaptive mesh confinement
ANL	Argonne National Laboratory
APU	accelerated processing unit
ASC	Advanced Simulation and Computing program
ASCI	Accelerated Strategic Computing Initiative
ASCR	Advanced Simulation and Computing Research
ASQ	Applications, Simulation, and Quality [Division]
ATDM	Advanced Technology Development and Mitigation
ATS	Advanced Technology System
B	billion
BG/L	Blue Gene/L
BG/P	Blue Gene/P
BG/Q	Blue Gene/Q
CASC	Center for Applied Scientific Computing
CDC	Control Data Corporation
CHAOS	Clustered High-Availability Operating System
CNS	Council for National Security
COE	center of excellence
CORAL	Collaboration of ORNL, ANL, and LLNL
CPU	central processing unit
CRADA	cooperative research and development agreement
CS-2	Compute Surface-2
CTRCC	Controlled Thermonuclear Research Computing Center
CTS	Commodity Technology Systems
DEC	Digital Equipment Corporation
DOD	Department of Defense
DOE	Department of Energy
ECFM	Exascale Computing Facility Modernization
ECI	Exascale Computing Initiative
ECP	Exascale Computing Project
ES	Earth Simulator
FY	fiscal year
FYNSP	Future-Years Nuclear Security Program
GPU	graphics processor unit
HPC	high-performance computing
HPGN	high-performance gateway node
HQ	headquarters
HW	hardware
ICEG	Institutional Computing Executive Group
IDC	integrated design code
KH	Kelvin–Helmholtz
KPP	key performance parameter
LAN	local-area network
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LC	Livermore Computing
LDRD	Laboratory Directed Research and Development
LFO	Livermore Field Office
LLNL	Lawrence Livermore National Laboratory
LTSS	Livermore Time-Sharing System
M	million

M&IC	Multiprogrammatic and Institutional Computing
MCR	Multiprogrammatic Capability Resource
MD	molecular dynamics
MIT	Massachusetts Institute of Technology
MOA	memorandum of agreement
MPCI	Massively Parallel Computing Initiative
MPI	Message Passing Interface
MW	megawatts
NERSC	National Energy Research Scientific Computing Center
NIF	National Ignition Facility
NMFECC	National Magnetic Fusion Energy Computer Center
NNSA	National Nuclear Security Administration
NRE	non-recurring engineering
NSCI	National Strategic Computing Initiative
OCF	Open Computing Facility
ORNL	Oak Ridge National Laboratory
OS	operating system
PCF	predictive capability framework
PCR	parallel-capacity resource
PI	principal investigator
PNNL	Pacific Northwest National Laboratory
PSE	problem-solving environments
R&D	research and development
RFP	request for proposal
SC	Supercomputing Conference
SD	Strategic Deterrence
SCF	Secure Computing Facility
SGI	Silicon Graphics
SMP	symmetric multiprocessors
SNL	Sandia National Laboratory
SP	scalable parallel
SRD	secret restricted data
SSP	Stockpile-Stewardship Program
SST	Sustained Stewardship Teraflop
START	Strategic Arms Reduction Treaty
SU	scalable units
TOSS	Tri-lab Operating System
TLCC	Tri-lab Commodity Clusters
UGTs	underground nuclear tests
UQ	uncertainty quantification
U.S.	United States
V&V	verification and validation
VPU	vision processing unit
WCI	Weapons Complex and Integration
WPD	Weapon Physics and Design
WSC	Weapon Simulation and Computing

Lawrence Livermore
National Laboratory