

# Introducing El Capitan

El Capitan: The World's Fastest Supercomputer Is Revolutionizing National Security and Sci...



One of the world's most powerful and energy-efficient supercomputers helps scientists safeguard the U.S. nuclear stockpile.

Each new generation of technological advancement in high-performance computing (HPC) is named according to the number of calculations the most sophisticated machines can perform. Welcome to the exascale era, where a quintillion floating-point operations per second—an exaflop—is the new computing threshold.

Numerically, exascale is a billion billion, expressed as  $10^{18}$ . This level of processing power is staggering to comprehend.

A person can calculate a basic equation in about one second. Imagine everyone in the world calculating their own equations every second for a year. In four years, they would still collectively fall short of an exaflop. An exascale supercomputer can do in an instant what billions of people cannot.

Thanks to the Laboratory's long history of HPC leadership, the promise and potency of exascale computing have arrived on the Livermore campus. The El Capitan supercomputer, named for Yosemite National Park's distinctive rock formation, is fully operational and already delivering unprecedented results. As powerful as Livermore's 125-petaflop ( $10^{15}$ ) Sierra supercomputer is, El Capitan's peak performance of 2.79 exaflops is more than 22 times faster. In fact, it has been benchmarked as the fastest supercomputer in the world.

The scale of this new system is commensurate with its responsibilities. El Capitan is not the Department of Energy's

(DOE's) first exascale system—Oak Ridge and Argonne national laboratories respectively deployed the Frontier (2022) and Aurora (2023) supercomputers—but it is the first for the National Nuclear Security Administration (NNSA), which is tasked with ensuring the safety, security, and reliability of the nation's nuclear deterrent. In the decades since the Comprehensive Nuclear Test Ban Treaty, HPC has become an essential part of modern stockpile stewardship.

Indeed, complex multiphysics simulations are more important now than ever. The higher the fidelity and resolution, the higher the confidence in stockpile reliability and performance. “Aging materials and parts may not behave exactly as expected. Whether or not this is an issue requires detailed modeling and accurate simulations,” explains Rob Neely, Livermore's associate director for Weapon Simulation and Computing. And as global adversaries grow their nuclear capabilities, he notes, “The U.S. may need to quickly field a weapon to fill a gap in our deterrent. In that case, simulation and modeling will be key to rapidly converging to a viable design.”

Amid this evolving threat landscape, NNSA's Advanced Simulation and Computing (ASC) Program focuses on simulation-based analysis of nuclear weapon functionality, combining the efforts of Lawrence Livermore, Los Alamos, and Sandia national laboratories into the Tri-Lab collaboration. Researchers at all three laboratories run simulations on El Capitan. (See box below.)


## Exascale Excellence

Deploying a supercomputer is a complicated undertaking. The process includes many overlapping steps and milestones—from procurement and design to installation and testing—over several years. “Our team includes experts in system administration, high-performance computing (HPC) tools, hardware, software, storage, user support, facilities, and security alongside vendors who provide processors and other components as an integrated system,” states El Capitan integration project lead Adam Bertsch. “Despite challenges along the way, we always come together in the spirit of partnership and complete our goals.”

To communicate requirements, resolve issues, and manage expectations, the partnership with Advanced Micro Devices, Inc. (AMD) and Hewlett Packard Enterprise (HPE) included several technical working groups. Vendors and laboratory staff worked side by side to coordinate every detail, such as fine-tuning features and troubleshooting bugs.

Formed to focus on applications, the El Capitan Center of Excellence (COE) featured close interactions among the Lawrence Livermore, Los Alamos, and Sandia national laboratories' collaborative code teams. “We've built a history of successful COEs. For El Capitan, AMD and HPE were new organizations for the Laboratory to work closely with, so collaboration was the key to success,” recalls Judy Hill, Livermore's COE lead. Rob Neely, Livermore's associate director for Weapon Simulation and Computing, agrees, “The COE has been a true success story of how we can work with our vendor partners to prepare our applications for a first-of-a-kind system far in advance of its deployment.”

Working group interactions have improved the vendors' products, such as system software adjustments that arose from El Capitan's unique purpose and specifications. The collaborations also strengthen the Laboratory's reputation as a trailblazer in the field. Hill explains, “Because of our subject matter expertise in ‘all things HPC’ spanning disciplines and organizations, Livermore is one of very few places in the world that



can deploy and use a system as complex as El Capitan.”

## Accelerated Architecture

Sierra’s architecture combined central processing units (CPUs) and graphics processing units (GPUs) and was the first of Livermore’s heterogeneous systems to lean on GPUs for improved computing performance. (See *S&TR*, August 2020, **The Sierra Era**.) El Capitan’s architecture upgrades this paradigm with accelerated processing units (APUs), a tight integration of CPUs and GPUs in a single package that brings exceptional efficiency and fidelity to 3D simulations. Built by Advanced Micro Devices, Inc. (AMD), the MI300A series APUs are also primed for AI-assisted data analysis of exascale simulations.

El Capitan is the first supercomputer in the world to use these APUs, four of which make up a compute node. Livermore Computing (LC) Chief Technology Officer Bronis de Supinski says, “Discrete CPUs and GPUs have separate physical memory systems for each processor type, which need significantly more programming to move data between them. APUs eliminate the need for data transfers between memory systems.”

In addition to advanced processors, exascale applications rely on an intricate circuitry network to manage the operations happening in the blink of an eye. Hewlett Packard Enterprise’s (HPE’s) high-speed Slingshot interconnection network routes computational traffic across the entire El Capitan system. A hierarchical connectivity pattern known as dragonfly topology further enhances the network, resulting in low latency and high bandwidth.

Every large-scale multiphysics simulation generates huge amounts of data—not only from the application code itself but also from the system checkpoints and input/output (I/O) operations that store and retrieve data. El Capitan’s architecture includes both a high-capacity storage tier and a high-bandwidth local tier so that applications can read and write data in different ways.

Livermore is the first supercomputing center to use HPE’s new near-node local storage modules, which are installed alongside El Capitan’s compute nodes. Aptly nicknamed “the Rabbits,” these components move data quickly and immediately while the application runs.

When the application completes, data migrates from the Rabbits to the larger, longer term Lustre file system. Data can move back to the Rabbits when the application runs again. Furthermore, the Rabbits have their own processors that run data analysis tools and software containers instead of running those operations on adjacent compute nodes. (See *S&TR*, March 2024, **Solving at the Speed of Exascale**.)

This highly efficient combination of storage solutions opens the door for more complex simulation workflows such as those with AI algorithms or in situ data analysis. de Supinski remarks, “The Rabbits shorten the distance between compute nodes and storage, cutting down on I/O time without affecting the network or interfering with computational progress.”

## Custom Upgrades



Livermore staff and vendor partners worked closely to “unbox” El Capitan.

Just as a laptop or smartphone has an operating system, so too does a supercomputer. However, unlike personal computing devices, no off-the-shelf program can manage an HPC system's hardware components, streamline its system administration, or accommodate users' needs. An HPC-optimized solution is the only solution, so nearly two decades ago, the ASC Program backed development of the Tri-Lab Operating System Stack (TOSS). Based on Red Hat Enterprise Linux, TOSS is used on commodity computing clusters at all three laboratories.

El Capitan is running the fourth version of TOSS, specially upgraded to integrate with the vendor software that helps run the APUs, Slingshot, and Rabbits. This version is adapted for exascale system management and monitoring, yet is flexible enough for use on other Livermore systems. "With TOSS, we bring our whole software ecosystem onto the platform, leveraging our experience from many years of installing and debugging the environment across different architectures," notes LC Division Leader Becky Springmeyer. "Our system administrators are expert TOSS users and developers, and our users do not need to learn yet another proprietary operating system."

In a second departure from household computers, a supercomputer does not merely plug into a wall outlet. During El Capitan's early planning stages, Livermore's existing computing facilities were not ready for exascale infrastructure demands—but now they are. Completed in 2022, the Laboratory's massive Exascale Computing Facility Modernization (ECFM) project laid the literal foundation for the electrical substation, transformers, and cooling towers necessary for an exascale system. (See *S&TR*, July/August 2022, **Charging Up and Rolling Out.**)



In preparation for El Capitan, upgrades to Livermore's primary machine room included scaled-up, energy-efficient cooling infrastructure. Water flows through the green and blue pipes without the use of chillers.

The ECFM project finished ahead of schedule, under budget, and without any safety incidents. The project team future-proofed this 85-megawatt infrastructure so that a second exascale supercomputer can join El Capitan in the years ahead. Occupying over 550 square meters of the primary machine room floor, El Capitan requires about 35 of those megawatts, and its energy efficiency design earned it a spot in the top 20 on the Green500 list of energy-efficient supercomputers.

## Scaled-Up Software

Various software programs, libraries, and frameworks are needed to guide scientific application codes from start to finish—analogue to different parts of an airplane providing lift, aerodynamics, navigation, and thrust. As HPC hardware rises to meet the exascale challenge, so too must the software. Livermore computer scientists and developers have long been working on a software portfolio that anticipates El Capitan. (See *S&TR*, February 2021, **The Exascale Software Portfolio.**)

Springmeyer explains, "Creating and supporting the software necessary for scientific computing workflows requires immense coordination across teams, including those in vendor organizations who deliver network and systems software. This ecosystem draws on decades of designing and developing HPC software."

These exascale-ready software tools—many developed at Livermore—are highly specialized and optimized for managing all steps of an application code's workflow. For instance, the Spack package manager installs the interdependent software packages that enable the code to begin running. (See *S&TR*, March 2023, **Expediting Research with Spack.**) The Flux resource manager schedules each application run for processing on the supercomputer's nodes. (See *S&TR*, July 2022, **Optimizing Workflow with Flux.**) Mathematical libraries such

as SUNDIALS and MFEM provide nonlinear solvers and finite element discretizations, respectively, to perform a simulation's numerical operations. (See *S&TR*, February 2021, [The Exascale Software Portfolio](#).) Data visualization software such as Ascent supports real-time simulation data analysis, and a host of performance tools help researchers evaluate and optimize simulation performance. Other software solutions tackle extreme-scale power management, code compiling, I/O operations, checkpointing, debugging, and data compression.

## Making the Move

Perhaps the trickiest exascale software challenge is portability. An application that runs on Sierra will not automatically run on El Capitan, a system with a significantly different computing architecture. To address this disparity, Livermore's RAJA Portability Suite provides software abstractions that insulate the source code from implementation details associated with specific programming models and hardware architectures. The Laboratory's application teams invested heavily in this portability suite when preparing codes for Sierra, then built on that investment for El Capitan with much less code disruption. "Applications that adopted RAJA were able to run on El Capitan hardware with relatively little additional effort compared to Sierra, and could focus on performance tuning more quickly," states RAJA Portability Suite project lead Rich Hornung.

This portability solution minimizes code rewriting when a new HPC system arrives—a scalability boon for the Laboratory's mission-critical multiphysics applications, many of which were in development long before the rise of GPU-based architectures. Adapting to El Capitan's APUs, the ARES team used RAJA's benchmarking features to optimize performance. ARES models high-explosives, materials physics, inertial confinement fusion (ICF), pulsed power, and other laser-driven experiments. "The larger capacity per processor allows us to run more physics integrations within the same simulation, and we've been able to scale performance better than ever before," explains Jason Burmark, a computer scientist working on both ARES and the RAJA Portability Suite.


Application development cannot grind to a halt while addressing portability issues. Since 2015, the MARBL code has simulated hydrodynamics processes such as burning plasma, radiation diffusion, and compressible flows, so its portability to exascale architecture was a priority. Taking advantage of greater computing capacity meant refactoring the numerical algorithms that underpin MARBL's physics capabilities.

Computational physicist and MARBL project lead Rob Rieben states, "The ever-increasing scale of systems like Sierra and El Capitan has led to multiple surprises. We always discover something when we push our code to new resolutions and scales."

For Livermore's application teams, scaling codes to run on each successive HPC generation—irrespective of hardware manufacturer—is a high-effort, high-reward endeavor. According to Rieben, "Transforming MARBL from being a CPU-only code to running on Sierra's GPUs, and then on El Capitan, was a big development effort requiring careful collaboration between computer scientists, computational physicists, and applied mathematicians." Here again the lessons from Sierra paid off. Rieben adds, "The work spent preparing for Sierra and the development of the RAJA Portability Suite greatly accelerated the process of porting our code base to AMD's APUs."

## Proving Grounds

As with previous supercomputer procurements, neither the Laboratory nor its vendor partners wait for full system



This 3D, shock-driven Kelvin-Helmholtz fluid experiment was simulated with the MARBL multiphysics code. The simulation consists of 13 billion high-order quadrature points and ran on 256 nodes of El Capitan.

deployment to verify that everything is operating correctly. Since 2021, Livermore teams have been working on a series of early access systems (EAS), which are smaller versions of El Capitan technology with the same operating system and similar architecture but a fraction of the processing power. The five EAS machines are put through a gauntlet of software installation, hardware configuration, and electrical and cooling connections. Tri-Lab users run real applications on these test beds to identify bugs or surprises.

For instance, ARES ran on the EAS machines so the team could benchmark and adjust the code's performance before moving to large-scale runs on El Capitan. Among the Laboratory's proprietary application codes, ARES often makes the first foray onto new systems at Livermore. Burmark notes, "We set this expectation because of our zeal to provide a performant, scalable capability for national security. Being first sometimes means we encounter challenges before other teams, so we work with the COE and other partners to find solutions."

Scientists and developers have also been test-driving El Capitan's unclassified sibling systems. (See the box below.) For example, Livermore's MARBL code team first tested peak processing performance on the new APU-based RZAdams machine. Early application runs demonstrated that full-scale El Capitan runs were within reach.

These smaller systems also serve as test beds for the physical aspects of installing a supercomputer of El Capitan's scale. Each cabinet—consisting of compute nodes connected onto boards, which are then combined into 64 blades per cabinet—arrived at the Laboratory on a convoy of semitrucks. Facility staff carefully wheeled each cabinet onto the loading dock, then into the freight elevator, then onto the machine room floor. The EAS machines, Tuolumne, RZAdams, and finally El Capitan came together node by node, cable by cable. "Monumental describes not just El Capitan's power but also the level of effort it takes to field a machine like this and ensure it performs at the expected level," confirms Springmeyer.



El Capitan integration project lead Adam Bertsch poses with the system's final blade (consisting of two compute nodes) prior to installation. Team members signed the back of the blade to commemorate the milestone.

## Sibling Systems

The enormous El Capitan granite monolith rises over 900 meters above Yosemite Valley in eastern California. Fifty miles away at the other end of Yosemite National Park lies Tuolumne Meadows—a diverse ecosystem of alpine vegetation and hundreds of animal species where the Tuolumne River winds amid glaciated granite domes. This rich environment is the namesake for the El Capitan supercomputer's unclassified counterpart, also sited at Livermore.

The Tuolumne system shares El Capitan's



accelerated processing unit (APU) architecture on a smaller scale. With a peak performance of nearly 300 petaflops, this supercomputer is the new workhorse system for the Laboratory's unclassified research in energy security, astrophysics, disease therapeutics, and more. "Tuolumne represents a huge increase in computational capability over the petascale Lassen system and is more powerful than even Sierra," states Judy Hill, who leads the El Capitan Center of Excellence at Livermore.



An additional unclassified system, RZAdams, supports work in inertial confinement fusion, high-energy-density physics, and conventional weapons, among other areas. Altogether, the APU-based systems expand the range of scientific challenges Livermore scientists can undertake. Hill points out, "With this evolution in computing, our scientific simulations now have higher resolution as well as better, more detailed physics."

## Technology Runway

The DOE's Exascale Computing Project (ECP) laid the runway for the technology now deployed with El Capitan. The multiyear collaboration took up the challenge of creating a production-ready hardware and software ecosystem that could be fully integrated with scientific application codes running on exascale systems. According to Terri Quinn, associate program director for LC Systems and Environments, "DOE partnered with industry to deliver exascale-class computing much earlier than what industry would do on its own. El Capitan is evidence of the success of this approach, and I have no doubt we'll see amazing new discoveries over its lifetime."


From 2016 to 2024, ECP brought together hundreds of researchers from DOE's Office of Science, NNSA, and all 17 DOE laboratories to conquer the technological challenges of exascale computing. The scope and stakes were high. "A system such as El Capitan would not have been available in this timeframe without ECP's investments in U.S. companies to accelerate exascale technologies," says Quinn, who led ECP hardware integration efforts.

ECP successfully concluded under the leadership of Lori Diachin, Livermore's principal deputy director of the Computing Principal Directorate. Today, the project's goals are realized with three exascale supercomputers serving the nation's science and security missions. Quinn adds, "DOE has maintained a leadership position in scientific and technical computing, as has NNSA in using computational science for national security."

## High-Fidelity Science


To grasp the exascale leap in computing capability, Neely offers the James Webb Space Telescope as an analogy. "Astronomers can see known parts of the universe in much more exquisite detail, which translates into better insights. They can also make discoveries they didn't expect, such as new distant galaxies or supernovae," he notes. "Similarly, El Capitan will give us more detailed understanding of the weapons behavior we've been studying for decades, while also providing insights to possibly discover new and unexpected phenomena."

With El Capitan, researchers can move beyond ensembles of 1D and 2D simulations to larger ensembles of 2D and 3D simulations.



Run on 64 nodes of the Tuolumne supercomputer for the Inertial Confinement on El Capitan project, this MAREL simulation of the N210808 burning plasma National Ignition Facility shot is used to prepare the code for larger scale runs on El Capitan.

Additionally, simulations can run with higher spatiotemporal fidelity that exposes more detail—and therefore higher accuracy—in the physical phenomena crucial to stockpile science, such as explosive detonation, fluid flow, material behavior under extreme conditions, and radiation transport. Consequently, scientists can better quantify a simulation’s uncertainty and error.



For extreme-scale scientific workflows to be feasible, accuracy demands acceleration. El Capitan’s unique architecture exploits parallel computing in new ways, decreasing the time to run expensive simulations that in turn reduce costly experimental trial and error. Rieben points out, “Exascale computing gives us a huge boost in the rate of scientific discovery and throughput we can achieve, leading to new insights and better predictive capabilities. The simulations we’re running on Tuolumne and El Capitan are the highest resolutions we’ve ever run for MARBL.”

## The AI Boost

The eras of exascale computing and AI overlap at the Laboratory. AI technologies do not replace scientists and engineers but add efficiency to experiments and simulations. Neely points out, “While exascale simulations give us novel insights into the details of weapons behavior, we also need the ability to move more quickly. AI is an increasingly important tool in our toolkit.”

Livermore researchers integrate AI models and techniques into scientific workflows in an approach called cognitive simulation. (See *S&TR*, September 2022, **Cognitive Simulation Supercharges Scientific Research**.) AI models can automate repetitive tasks, identify complex patterns, and accelerate predictive and inference capabilities. According to computational physicist Luc Peterson, “We can teach an AI system to search for answers to scientific questions. The model won’t be biased, and it can rigorously find solutions we might have dismissed. AI frees us to think at a higher level apart from the daily grind of running experiments.”

AI consumes significant computing resources, so an exascale system is a research windfall. “El Capitan gives us the ability to explore cognitive simulation in the weapons program like we’ve never been able to do before,” adds Neely.

## Energy Experiments Go Exascale

Since Livermore’s historic fusion ignition achievement in 2022, scientists at the National Ignition Facility (NIF) have repeated the multi-megajoule feat several times, edging closer to the twin goals of understanding the physics of nuclear weapons and developing a clean fusion energy source. HPC simulations have long contributed to these experiments, and El Capitan offers powerful new technology in the Laboratory’s quest to advance ICF applications.

A groundbreaking project called ICECap—Inertial Confinement on ElCapitan—combines an AI-based workflow with large-scale multiphysics simulations and ICF experimental data. AI algorithms extrapolate important features from massive data sets, then make predictions based on those features. The ICECap team expands the ICF design space by boosting these algorithms with supercomputers.

“Can a million AI-powered simulations find the next generation of NIF experiment designs? That’s the question we’re answering,” states Peterson, the project’s principal investigator. “ICECap combines AI, HPC, multiphysics modeling, and pencil-to-paper science to tackle something seemingly impossible.” With data from only a few hundred simulations, the workflow ran on an EAS machine and produced a new design for the hohlraum, which holds the ICF target capsule.

## Full Potential

As Livermore leads the NNSA into the exascale age, those closest to El Capitan’s operations look toward the next

HPC horizon. During her 40-year Livermore career, Springmeyer has seen computing power evolve from megascale ( $10^6$ ) to exascale—a trillion-fold increase in performance. She reflects, “I’m struck by how quickly the time flew from when I used a CDC-7600 machine to installation of El Capitan and Tuolumne. The next generation of engineers, operators, system administrators, developers, and business teams will repeat this evolution again and again. This is an exciting time at the Laboratory.”

El Capitan’s impact on Livermore’s mission cannot be overstated, with many experts anticipating future scientific discoveries. “Science isn’t just doing an experiment and proving a hypothesis once. We must be able to repeat it under varying conditions so we can trust the answer. Exascale computing provides more trust in results and transforms the way scientists think about their work,” says de Supinski. Quinn adds, “I’m eager to see what our highly talented users will accomplish when they leverage the full potential of El Capitan.”

—*Holly Auten*

For further information contact Becky Springmeyer (925) 423-0794 ([springmeyer1@llnl.gov](mailto:springmeyer1@llnl.gov)).