

[Back](#)

Lawrence Livermore National Laboratory's El Capitan verified as world's fastest supercomputer



[\(Download Image\)](#)

With a peak performance of 2.79 exaFLOPS, El Capitan comprises more than 11,000 compute nodes and provides Lawrence Livermore National Laboratory with a flagship machine 22 times more powerful than its previous fastest supercomputer, Sierra. (Photos: Garry McLeod/LLNL)

Lawrence Livermore National Laboratory (LLNL), in collaboration with the **National Nuclear Security Administration** (NNSA), **Hewlett Packard Enterprise** (HPE) and **AMD**, have officially unveiled El Capitan as the world's most powerful supercomputer and first exascale system dedicated to national security.

Verified at 1.742 exaFLOPs (1.742 quintillion calculations per second) on the High Performance Linpack — the standard benchmark used by the **Top500** organization to evaluate supercomputing performance — El Capitan is the fastest computing system ever benchmarked. The system has a total peak performance of 2.79 exaFLOPs. The **Top500 list** was released at the 2024 Supercomputing Conference (SC24) in Atlanta.

As NNSA's first exascale supercomputer, El Capitan is a premier resource for the NNSA Tri-Labs — LLNL, and Los Alamos and Sandia National Laboratories — to advance nuclear weapon science and scientific discovery, providing the vast computational power necessary to ensure the safety, security and reliability of the nation's nuclear deterrent without nuclear testing. This state-of-the-art system marks a monumental leap forward in HPC, enabling unprecedented modeling and simulation capability essential for NNSA's Stockpile Stewardship Program that certifies the U.S. nuclear stockpile, and other critical nuclear security missions such as nonproliferation and counterterrorism.

“El Capitan's introduction continues the capability advancement needed to sustain our stockpile without returning to explosive nuclear testing. This computational capability, backed by decades of data, expertise and code development is the heart of science-based stockpile stewardship,” said Jill Hruby, Department of Energy (DOE) under secretary for nuclear security and NNSA administrator. “We will continue to invest in the technological and scientific infrastructure necessary to underpin the nuclear security enterprise.”



Designed to optimize the convergence of AI and high performance computing, the MI300AInstinct APUs deliver unmatched computational performance, energy efficiency and reliability, and are well-suited for work in support of AI-driven workloads that will impact national nuclear security, as well as efforts in fusion energy, climate research and drug discovery.

NNSA Tri-Lab scientists will utilize El Capitan's speed and unparalleled capabilities to further advance NNSA's core mission of maintaining an aging stockpile while simultaneously pursuing weapon modernization such as the W87-1 and W93 warheads currently under development. El Capitan will realize a multi-decade goal to model weapon performance and safety in high-fidelity resolution with quantified uncertainties using codes that have been developed and tuned for exascale over the past decade. It also will be used to model advanced high-energy-density physics experiments, such as inertial confinement fusion, the complex dynamics of ballistic reentry and detailed understanding of material behavior under extreme conditions.

El Capitan also will support novel new AI-based workflows to address emerging challenges in the NNSA mission, including material discovery, design optimization, advanced manufacturing, digital twins and intelligent AI assistants trained on classified data. Advances in these national security capabilities will impact the broader mission of the DOE and the scientific community at large, including clean energy, climate science, seismic modeling and building a more efficient and agile enterprise based on advanced computing.

“This tremendous accomplishment, years in the making and the result of tireless efforts by hundreds of dedicated employees in this large collaborative team, is a testament to the Laboratory’s leadership in driving scientific discovery. It continues a legacy of supercomputing excellence that spans more than 70 years,” said LLNL Lab Director Kim Budil. “El Capitan’s extraordinary computing capabilities will allow us to tackle complex challenges that were previously out of reach. We are proud to lead this achievement in partnership with industry, and advance science in ways that will benefit society and the nation as a whole.”



Lawrence Livermore National Laboratory technician Ty Nguyen works on installation and maintenance of El Capitan’s compute blades.

With more than 11,000 compute nodes and 5.4375 petabytes of total memory, El Capitan represents a more than 20-fold peak increase in computing performance over LLNL's previous most powerful system, Sierra, which has a peak performance of 125 petaFLOPs (PF). Complex, high-resolution 3D simulations that would take weeks or months on Sierra will be done in just hours or days on El Capitan, leading to previously unimaginable insights, according to LLNL experts.

"We expect El Capitan to make yesterday's 'hero runs' of large-scale 3D models commonplace, allowing us to analyze components of the stockpile in greater detail and with more precision than ever before," said LLNL's Weapon Simulation and Computing Program Director Rob Neely. "El Capitan will enable us to simulate entire weapons systems, incorporating various real-world factors such as materials, manufacturing imperfections and environmental variables, meaning more accurate predictive capabilities and better-informed decision-making for NNSA and the Stockpile Stewardship Program."

Built on the HPE Cray Supercomputing EX system specifically designed for exascale computing, El Capitan is a product of numerous cutting-edge innovations in HPC. It features HPE's direct liquid-cool leadership-class solutions, including the HPE Slingshot interconnect, alongside custom-built, ultra-fast, near-node local storage tiered to a global Lustre file system that's shared between all the compute nodes. It capitalizes on an end-to-end approach encompassing everything from system architecture and data storage to networking and software. El Capitan is powered by the AMD Instinct MI300A Accelerated Processing Units (APUs), which combine CPU cores, graphics processing unit (GPU) cores and high-bandwidth memory into a single shared package.

El Capitan marks the third exascale machine HPE has deployed for the Department of Energy, extending a longstanding supercomputing legacy.

"El Capitan is the results of years of engineering and innovation, and this achievement is a testament to our years-long collaboration with NNSA, LLNL and AMD as we push the boundaries of what's possible in computing," said HPE's Senior Vice President and General Manager of HPC & AI Infrastructure Solutions Trish Damkroger. "Built to address tomorrow's

challenges today, we can't wait to see the incredible discoveries that will come from this machine, propelling our society forward while overcoming almost any obstacle. El Capitan is truly the pinnacle of supercomputing."

Designed to optimize the convergence of AI and HPC, El Capitan's AMD Instinct MI300A APUs deliver unmatched computational performance, energy efficiency and reliability, and are well-suited for work in support of modeling and simulation workloads that will impact national nuclear security, as well as efforts in fusion energy, climate research and drug discovery.

"At AMD, we push the boundaries of technology, and our partnership with NNSA, LLNL and HPE on El Capitan has been a tremendous opportunity to do just that," said Forrest Norrod, executive vice president and general manager of the Data Center Solutions Group at AMD. "The AMD Instinct MI300A APUs are truly the heart of El Capitan's extraordinary capabilities, delivering a level of performance and efficiency that was unimaginable just a few years ago. With the strength of AMD processors, El Capitan is poised to deliver unprecedented performance and enable groundbreaking discoveries in science and national security. We're excited to see what breakthroughs it will spark in the future."



El Capitan's unclassified "sister" system, Tuolumne, will support projects in open science, including energy security, climate change, computational

biology/drug discovery and other areas of public interest.

In addition to the classified El Capitan, its companion system Tuolumne — built with the same architecture and components as El Capitan, but roughly 1/10th its size — reached No. 10 on the Top500 list. As an unclassified system, the 288-petaFLOP Tuolumne (288.8 PF peak) is funded by multiple NNSA programs and LLNL's Institutional Computing program. The machine will be essential for advancing open science at the Tri-Labs, such as modeling the impacts of climate change, performing high-resolution earthquake modeling, exploring advanced manufacturing and energy sustainability and advancing computational biology supporting other agencies through Strategic Partnership Programs.

Together, El Capitan and Tuolumne will provide NNSA scientists with an unmatched ability to handle large-scale data sets and complex simulations, accelerate research and provide solutions to some of the most pressing problems facing the world today.

To see more information about El Capitan, [**click here**](#).

Nov. 18, 2024

CONTACT



Jeremy Thomas
thomas244@llnl.gov
(925) 422-5539

RELATED LINKS

National Nuclear Security Administration

Hewlett Packard Enterprise

Top500

TAGS

HPC, Simulation, and Data Science

Climate

Energy

Nuclear deterrence

Computing

Strategic Deterrence

Industry Collaborations

ASC


FEATURED ARTICLES



LLNL and Starris sign agreement, schedule conference talk for Aug. 13



Livermore teams honored with Defense Programs Awards of Excellence



LLNL pushes frontier of fusion target design with AI