
# The Road to El Capitan

Developing Exascale Computing Power for National Security

# Contents




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-MI-2001343



*Acknowledgments*

**Authors:** Holly Auten and Meg Epperly

**Concept:** Bronis de Supinski

**Layout and Design:** Meg Epperly

**Proofreaders:** Holly Auten, Anashe Bandari, Kristen Howard, Deanna Willis

**Printing:** LLNL Print Plant

**El Capitan and Tuolumne Meadows Nature Photography:** Mike Carter

**Photography of Teams, Most Systems:** Garry McLeod

**Additional Photography:** Meg Epperly and Katrina Trujillo

***asc.llnl.gov***
***computing.llnl.gov***
***hpc.llnl.gov***


Lawrence Livermore National Laboratory

# Introduction

The road to the NNSA's first exascale-class supercomputer, El Capitan, began several years ago. This enormous undertaking of meticulous planning and strategic collaborations is already bearing fruit in a number of ways. For instance, close coordination with the DOE Exascale Computing Project, Office of Science Advanced Scientific Computing Research program, and the NNSA Advanced Simulation and Computing (ASC) program has paved the way for secure integration of national security applications with a robust, exascale-ready software stack.

This article series, written in 2023–24, spotlights Livermore Computing's efforts to stand up the NNSA's first exascale supercomputer. Named after Yosemite National Park's formidable rock formation, El Capitan came online in 2024 with the processing power of 2.79 exaflops, or 2.79 quintillion ($10^{18}$) calculations per second. The system will be used for predictive modeling and simulation in support of the stockpile stewardship program.

# It takes a village

At LLNL, completion of the enormous Exascale Computing Facility Modernization (ECFM) project—which finished on time in 2022 and under its $100 million budget—marked a major milestone in the preparations for El Capitan.

Partway through 2023, the road is becoming steeper both in activity and focus for the Livermore Computing (LC) program led by Terri Quinn, the LC Division led by Becky Springmeyer, and the Weapon Simulation and Computing (WSC) community led by Rob Neely. LC Chief Technology Officer Bronis de Supinski notes, "This is an exciting—and challenging—period with hardware deliveries, environment testing, and facilities completion. We've been connecting the 85-megawatt ECFM infrastructure to early access systems [EAS] since 2021, as well as configuring applications and system software to run on machines of this scale."

*People across the Lab are pulling in the same direction for what will be one of the best computing systems in the world.*

**Bronis de Supinski**
LC chief technology officer

Aerial view of the Exascale Facilities Modernization Project electrical yard and cooling towers

LC Facilities Team members Rich Herbert, Karl Mueller, Matt Williams, and Marquita Walker during the preparations for El Capitan

El Capitan will be inherently different from Frontier, its sibling system at Oak Ridge National Laboratory. The design combines Hewlett Packard Enterprise (HPE)/Cray's advanced Slingshot interconnection network alongside processors from Advanced Micro Devices Inc. (AMD). LLNL is also the first supercomputing center to use HPE's near-node local storage solution called Rabbits. de Supinski adds, "We'll use the ASC Tri-Lab Operating System Stack and the Livermore-developed Flux resource management software—and we have to make these work with HPE's software from day one."

According to de Supinski, significant efficiency gains in El Capitan's architecture will come from AMD's accelerated processing units, or APUs. He says, "All CPUs and GPUs will have equal access to memory in a truly unified memory space, so the challenges of moving data back and forth will be much lower. This memory architecture will increase our performance."

These preparatory activities are under the purview of the El Capitan Center of Excellence, a collaboration of developers and experts at the NNSA labs, HPE, and AMD. De Supinski explains, "We like to deploy early hardware because we can fix any issues ahead of the larger system deployment. Through both collaborative and ad hoc meetings with the vendors and other partners, and even a 'Rabbit hackathon,' we've made a lot of progress."

Within LLNL, siting El Capitan involves all of LC and a close partnership with WSC. De Supinski is quick to credit the work of others, stating,

"People across the Lab are pulling in the same direction for what will be one of the best computing systems in the world. Managing a new partnership with vendors is challenging, but Terri Quinn's experience combined with [WSC business manager] Kim Bosque's finance background really make the process so much easier. I'm proud of LC's fabulous hotline and tier-two customer support, collaboration with code teams, and respect for and from our users—not to mention our excellent facilities staff, system engineers, operators, and software and hardware experts. Working with such a great set of individuals makes doing anything worthwhile."

LC Facilities Team members Brad Davis, Brandon Hoang, Ralph Huddleston, Barbara Macchioni, and Jeff Bilke during the preparations for El Capitan

# All the moving parts

*The functionality that El Capitan is going to unlock for our users and for the programs is the most exciting aspect.*

**Adam Bertsch**
Integration Project lead

A supercomputer does not simply arrive in one piece ready to be plugged into a wall. The number and scale of components is staggering, to say nothing of the expertise required to assemble and connect them. How do the different types of processors work together? What type of floor will hold all that weight? How will the system draw electricity without crashing the power grid?

Then, even when the lights are blinking and the cooling system is humming, a researcher cannot simply press a button to begin running their scientific code. How will users log in? How will the code know which nodes to use, and when? Where will the massive amounts of simulation data go? How will administrators and users know if the system is working properly? What happens if another researcher has a completely different code and use case?

Questions like these are never far from the mind of Adam Bertsch, the Integration Project lead for El Capitan. He states, “My role is to bring all the various pieces together, both within Livermore Computing [LC] and with our vendor partners, to make sure that the system shows up, works, and gets accepted. My job is to make sure the facilities, hardware, and software teams have all the information about each other that they need—that they know what’s expected of them and that they’re ready.”

## Essential Infrastructure

Bertsch served as the integration lead for LLNL’s 125-petaflop Sierra supercomputer, which came online in 2018, but every HPC system is different. For instance, LC teams must learn the ins and outs of vendors’ products: Sierra has IBM CPUs and NVIDIA GPUs, and a Mellanox EDR InfiniBand interconnect, whereas El Capitan will have AMD APUs (accelerated processing units) with HPE/Cray’s advanced Slingshot interconnection network. Differences continue in the number of cores and nodes, operating system versions, and power requirements.

Furthermore, Bertsch points out, “During Sierra’s preparation and deployment, we didn’t have a really large infrastructure project going on for months.” LLNL’s massive Exascale Computing Facility Modernization (ECFM) project was a prerequisite for siting El Capitan. Bertsch continues, “ECFM was very successful. [Project manager] Anna Maria Bailey and her team did a really great job bringing new power and cooling capabilities to the Lab as well as to the building that will house El Capitan. But we’re not done. Now we’re installing and testing the system-specific infrastructure, including how it works with the early access systems [EAS].”

Aligning the hardware delivery schedule with facilities and infrastructure upgrades is just one facet of integration dependency. “On the software side, we’re responsible for a large chunk of the system software configuration and implementation,” Bertsch explains. “So, we need to make sure we have onsite all of the hardware that’s required to bring up the software. That way, the software will be ready when the final hardware is hooked up. We aim to deploy El Capitan as efficiently as possible.”

El Capitan Integrator Pythagoras Watson, Integration Lead Adam Bertsch, and LC Chief Technology Officer Bronis de Supinski


HPE contract workers responsible for the El Capitan installation

RZVernal, one of the El Capitan early access systems

The EAS machines play an important role in this readiness strategy. As smaller versions of El Capitan, the five systems—nicknamed RZNevada, RZVernal, Hetchy, Tenaya, and Tioga—have been in various stages of installation and testing since 2021. "We've had these systems onsite for a while," Bertsch notes. "We use them to confirm the hardware setup, software stack, and application portability."

## Early Success and Future Benefits

One integration success has been getting HPE's programming environment running on the Tri-Lab Operating System Stack (TOSS). Another is the extensive portability effort that enables scientific applications currently running on Sierra to run on the very different El Capitan EAS architecture. In addition, LLNL's next-generation resource management and job scheduling software, Flux, is already running on Tioga. Bertsch says, "Application teams were able to start running on these new platforms much more quickly and effectively than they were with Sierra. And when their code is combined with Flux, scientists will be able to get the most out of El Capitan."

As complicated as this exascale integration is, Bertsch is optimistic. He says, "The functionality that El Capitan is going to unlock for our users and for the programs is the most exciting aspect. The way people use computers to do their work will change." Additionally, he sees benefits for the LC team who keeps the supercomputers operational 24/7/365. "Because we're using TOSS, system administrators will be able to manage this new system the same way they manage our other systems. So, the transfer learning and onboarding will be much easier than it was with Sierra. When someone gets called in the middle of the night, they'll be able to respond with what they already know," Bertsch states. As for users who are accustomed to the Lab's current HPC systems, he adds, "We'll roll out Flux across the computing center, so researchers will be using the same tools they're used to."



# The right operating system

Supercomputers commissioned for the DOE's national laboratories do not exist in a vacuum. The entire process—procurement to design to build to installation to testing to acceptance to operation—unfolds with input from, and close coordination of, a number of stakeholders. Teams representing all of the interested agencies ensure each step is completed properly and transferred smoothly to the next step. Each decision is made with collaborative considerations: How will this system serve the DOE's mission? How will it enable multi-institutional teams to advance science? How will it push the leading edge of HPC technologies?

One such decision is about El Capitan's operating system (OS), and the requirements are steep. A suitable OS for the NNSA's first exascale machine needs to integrate fully with vendor software, function efficiently for users from multiple institutions, and streamline system administration. The chosen solution is the Advanced Simulation and Computing (ASC) program's Tri-Lab Operating System Stack (TOSS).

The "Tri-Labs" are the NNSA's national laboratories: Livermore, Los Alamos, and Sandia. Livermore Computing's (LC) Jim Foraker, Systems Software and Security Development Group leader, explains, "TOSS meets everyone's needs. It's a Red Hat Linux distribution created here at Livermore to provide a common base environment for the commodity clusters at all three labs. And other organizations around the DOE complex, as well as NASA, are starting to use TOSS."

*It's becoming more and more obvious to everyone that we really can do this.*

**Jim Foraker**

Systems Software and Security Group leader

## 'In Charge of Our Own Destiny'

According to Trent D'Hooge, who serves as the system software lead for El Capitan, the right OS was undoubtedly going to be a custom solution. But what was the best way to fulfill the requirements? He says, "We made the choice to be in charge of our own destiny. We have full lifecycle support for TOSS, including release management, software packaging, quality assurance, configuration management, and much more. Choosing TOSS means we can run El Capitan for at least ten years."

Foraker adds, "TOSS will simplify the user's experience and the administration side of El Capitan. The benefits are huge." The OS's track record and robust functionality gives LC's system software experts confidence in this decision.

## Future-Proofing

A sound decision still comes with challenges. TOSS needs to not only integrate with the AMD APUs, Hewlett Packard Enterprise (HPE)/Cray Slingshot interconnect, and Rabbit hardware, but also support system management and monitoring at El Capitan's scale. Accordingly, El Capitan's early access systems (EAS)—three of which rank in the top 200 HPC systems in the world—play a crucial role in software readiness. D'Hooge states, "The EAS machines are where we make sure all our system software supports the hardware that HPE provides. We've had success with Slingshot functionality already."

Each of the five EAS machines is already running TOSS ahead of El Capitan. Foraker notes that scaling up is difficult and complicated, but the smaller systems are an important proving ground. "We can do this," he says. "I think we're in a place now where it's becoming more and more obvious to everyone that we really can do this."

# Storage in the exascale era

Imagine the not-too-distant future. You're ready to run your 3D multiphysics application on the NNSA's first exascale supercomputer. El Capitan revs into action, processing more than a quintillion calculations per second thanks to the incredibly complex coordination between its sophisticated hardware and system software. Your code hits a snag and has to restart, but fortunately there's a checkpoint in place and it continues on almost seamlessly. Meanwhile, your application and its checkpoints are generating data, and as Livermore Computing (LC) computer scientist Brian Behlendorf observes, "That tremendous amount of data has to go somewhere."

"Somewhere" means a file system and the crucial I/O (input/output) process of storing and retrieving application data. Behlendorf, who serves as the I/O lead for El Capitan, explains that the machine is designed with two tiers of storage: a large-capacity tier alongside a near-node solution called a Rabbit. "Workloads are changing and stressing traditional file systems. The Lab's applications vary widely and require different types of storage for different computing paradigms," he notes. "El Capitan will have the global Lustre file system that's shared between all the nodes as well as HPE's [Hewlett Packard Enterprise] dynamically configurable Rabbits."

## Introducing the Rabbit Node

HPE's Rabbits are storage hardware installed directly into a supercomputer. By being physically located close to compute nodes, they can move data very quickly, which makes them ideal for an application's immediate storage needs. A Rabbit can be configured into its own Lustre

*Integrating Rabbits into the early access systems, and ultimately into El Capitan, is a huge co-design effort.*

**Brian Behlendorf**
I/O lead

Bronis de Supinski and Brian Behlendorf inspect one of the Rabbit node integrations

file system that performs better than the comparatively slower global file system, or as local block storage so the application can run as efficiently as possible. Furthermore, Rabbits are equipped with a processor and provide the capability to execute I/O and data analysis tools in containers on the Rabbits. Behlendorf states, "This Rabbit hardware is new, innovative, and flexible. I'm excited to build it from the ground up with LC and HPE."

When you run your application on El Capitan, a tightly choreographed I/O dance is happening. The application submits the job to the Flux resource manager, specifying the amount of storage needed. Flux determines which compute and storage resources to use, and sends that request to the HPE software governing the Rabbit. The HPE software, which runs in a Kubernetes container environment, configures the Rabbit file system for your application. When the application completes, the HPE software migrates your data from the Rabbit to the large-capacity tier, freeing up the Rabbit for the next job. If you need to re-run your application, the data can move back to the Rabbit. More complex workflows—such as those with machine learning algorithms or in situ data analysis—can take advantage of this hierarchical strategy of containers and local shared storage.

## Storing More than Application Data

In addition to application teams, LC is working with others to adapt user-level analysis tools for Rabbits. One pilot effort involves running the SCR (Scalable Checkpoint/Restart) software directly on a Rabbit node instead of on a compute node. For SCR project lead Kathryn Mohror, the Rabbits' potential for running these types of programs is just as exciting as the Rabbits' ability to handle application data. She explains, "Rabbits offer a flexible, temporary area for file sharing that's faster than the larger file system and offer the ability to run our tools on the Rabbit itself. What if we leveraged those capabilities for tools like SCR for checkpointing or Ascent for in situ data visualization?"

When SCR captures checkpointing data, it decides whether to reduce, move, or delete that data as the application continues. Currently, this process happens on the compute node, which can lead to interference because, as Mohror points out, "Everything SCR does competes for resources with the application running on the same node." If the SCR activities are moved to the Rabbit, the application could run faster without interference from SCR—or any similarly positioned software—getting in the way.

The strategy promises to be a win for the compute center, too. "With SCR or other data analysis tools running on a separate processor from the main application, they will have more opportunities to reduce the data. For example, we'll have the luxury of running a compression algorithm that might take a little longer. We can be smarter about how much data ends up on the large-capacity file system, which benefits everyone in the center," says Mohror.

## Production Readiness

LC installed the first Rabbit prototype into the Hetchy early access system (EAS) last year, and the I/O team has been kicking the metaphorical tires

*[With Rabbits] we can be smarter about how much data ends up on the large-capacity file system, which benefits everyone in the center.*

**Kathryn Mohror**
SCR Project lead


Bronis de Supinski with an early Rabbit node in front of Hetchy early access system



ever since. Three more EAS machines are in the process of receiving Rabbits for additional testing of the associated firmware, hardware, and software.

"Having Hetchy available for this work has been tremendously valuable, and now we're able to exercise the basic I/O functionality on it as well as begin prototyping the SCR integration," Behlendorf explains. "Long term, Hetchy will be used to stage updates for El Capitan, before new software gets to the big system. In the meantime, it's our primary development environment for testing integration and production-ready efforts."

Behlendorf praises the close coordination between HPE and LC. He says, "Integrating Rabbits into the EAS, and ultimately into El Capitan, is a huge co-design effort. Together we're making sure the software stack runs correctly in our environment and on HPE's hardware." As just two examples of this teamwork, James Corbett from LC's Workflow Enablement group spearheaded the building of new capabilities into Flux to accommodate Rabbits (see *A framework for complex workflows,* pp. 13–17), while Chris Morrone from LC's Systems Software and Security group led the effort to stand up the Kubernetes environment.

Behlendorf is confident this I/O strategy and Rabbit technology will be ready when El Capitan comes online next year, stating, "The core functionality is there. We'll scale it up."

# A framework for complex workflows

Livermore Computing (LC) has a long history of developing state-of-the-art resource scheduling software that allows users to run and manage their simulation codes on HPC clusters. However, Slurm and similar traditional technologies have not been able to keep up with the challenges of increasing system scales and interplays, such as those that occur between compute clusters and file systems.

Over a decade ago, LLNL led a collaborative effort to develop Slurm. "Perfectly reasonable choices made at the beginning of the Slurm project started to conflict with the way HPC technologies and use cases were evolving. Slurm's scheduler was designed and optimized for one node corresponding to one core, and it has a strictly two-level hierarchy for jobs and job steps. Additionally, Slurm limits how many jobs can run on a given system," explains computer scientist Tom Scogland. "Over time, we could envision, and had already begun developing, other technologies with more expressive resources."

The LLNL-developed Flux resource and job management framework aims to fill these gaps. Flux expands the scheduler's view beyond the single dimension of "nodes" and enables new resource types, schedulers, and services that can be deployed as data centers continue to evolve. Flux is easy to run, allows users to compose complex workflows, and supports very large numbers of concurrent jobs.

Flux has been a user-level tool for years, and now it's an effective cluster-level tool. "We are providing more capable resource management through hierarchical, multi-level management and scheduling schemes," says

*We are providing more capable resource management through hierarchical, multi-level management and scheduling schemes.*

**Becky Springmeyer**
LC Division leader

*[With Rabbits and Flux,] you'll have flexibility—whether you want node-like storage, your own personal mini-Lustre file system, or some combination of the two.*

**Tom Scogland**
Computer scientist

Becky Springmeyer, LC Division leader. "Users benefit from schedulers that have deeper knowledge of network, I/O, and power interconnections, along with the ability to dynamically shape running work. A major challenge in designing Flux was making sure its framework is general and extensible enough to support resources and use cases that are only now emerging in research."

One of these emerging use cases is El Capitan's novel node-level storage (see *Storage in the exascale era*, pp. 9–12). James Corbett of LC's Workflow Enablement Group is integrating Flux with Hewlett Packard Enterprise (HPE) hardware and software, including testing and writing the plugins and code that optimize Flux for the early access systems (EAS). Flux is the primary resource manager on four of El Capitan's EAS machines as well as LLNL's Corona supercomputer.

## Rabbits and Flux

Compared to traditional resource management software, Flux makes smarter placement decisions and offers greater flexibility and more opportunity for adaptation. These solutions help scientific researchers more effectively harness a supercomputer's power. For example, by considering a holistic view of a data center's input/output (I/O) bandwidth capability and utilization, Flux avoids the "perfect storm" of I/O operations that can occur when a naïve scheduler places I/O-intensive work without regard to I/O availability. This is especially important when using near-node storage like HPE's Rabbits.

Flux handles each batch job as a complete instance of the framework with all the functionality of the system-level resource manager. Because each job can act as a full Flux instance running hierarchically inside the primary instance, users can customize Flux at different nested levels within their jobs. For example, a user desiring to launch many small high-throughput jobs could submit a large, long-running parent job, then load a specialized scheduler—one that is streamlined for high throughput—inside it. Panning outward in scale, "parent" schedulers operating at a larger granularity can move resources between "child" jobs as bottlenecks occur and employ pluggable schedulers for resource types like Rabbits and other novel architectures.

Testing of Flux's integration with the Rabbits is under way on two EAS machines, and the sum is proving greater than the parts. "El Capitan users will be able to leverage the Rabbits not just as a network burst buffer or even a node burst buffer, but also as per-job configurable storage resources. You'll have flexibility—whether you want node-like storage, your own personal mini-Lustre file system, or some combination of the two," says Scogland, Flux's scheduling and research lead.

## Broader Benefits

Feedback from developers and users has significantly contributed to Flux's current features and its development roadmap. To ease the learning curve of moving to this new technology, LC's Ryan Day created wrapper functions that transform Slurm commands for Flux. Additionally, Corbett notes, "Besides being able to manage the Rabbits closely, users will like the Python libraries Flux offers for manipulating jobs. This feature is already popular among users who write workflows." Scogland adds, "The license on the Flux APIs [LGPL3] is attractive to other developers, which means that if you dynamically link to Flux, you don't have to use the same license. Even closed or proprietary software can link to Flux, and our active and growing API ecosystem can accommodate a range of use cases."

Winner of a 2021 R&D 100 Award and part of the DOE Exascale Computing Project's Software Technology portfolio, Flux is open source and available to HPC centers around the world via GitHub. After Flux's capabilities are fully demonstrated on El Capitan, LC and the development team will be primed for collaboration with the HPC and artificial intelligence

Close-up of a Rabbit node

communities, expanding Flux's scope to reach even more complex workflows and specialized hardware technologies.

The award-winning Flux team

# Modern compilers

Referring to El Capitan as an Advanced Technology System (ATS) almost sounds like an understatement. Livermore's ATS roster evolves as high-performance computing (HPC) architectures and capabilities evolve. This supercomputer lineup includes the petascale Sierra and Lassen machines along with El Capitan, its unclassified counterparts Tuolumne and RZAdams, and the early access systems already installed at LLNL. ATS machines are designed to take advantage of the industry's most cutting-edge components and attributes: processors, operating system, storage and memory, energy efficiency, and so on.

Perhaps one of the least publicized advancements in the exascale era is compiler technology. Compilers manage the incredibly tricky translation of human-programmable source code into machine-readable code (think 1s and 0s), and they optimize the latter so the former can run more quickly. As programming languages evolve, so too must compilers.

Progress is a double-edged sword, however. Computer scientist John Gyllenhaal, who leads Livermore's ATS compiler team, explains, "Programming languages aren't static. New standards, such as C++20 or OpenMP 5.2, add significant new features that make languages more powerful, expressive, and maintainable. But these standards also create a huge amount of work for compiler writers and support staff. There are a lot of opportunities for things to go wrong in a compiler." Gyllenhaal's team works closely with the vendors who develop and optimize the compilers for various ATS architectures, including El Capitan.

Most HPC vendors' compilers are based on LLVM, an open-source C++ compiler infrastructure created two decades ago at the University of Illinois. LLVM's widespread adoption promotes sustainable, distributed development across the organizations that use it, including Livermore. This community effort frees vendors to develop proprietary optimizations on top of LLVM's baseline code.

According to Gyllenhaal, this consolidation has tradeoffs. He states, "On all of our architectures, our historical aim was to accommodate at least two robust vendor-supported compilers, allowing code teams to switch between them if they had to wait for a bug fix on one. With most compilers now being LLVM-based, this approach has been slightly less effective. We've had to account for the LLVM baseline while addressing issues from vendor-specific optimizations."


John Gyllenhaal working on compiler debugging

An additional complexity facing the ATS compiler team is that, as Gyllenhaal states, "Large-scale codes often break compilers." Complex multiphysics codes, like those that will run on El Capitan in support of stockpile stewardship, rely on compilers in order to perform successfully in different HPC environments. Compiler-related bugs must be carefully managed on classified systems, especially when involving hardware vendors in the solutions. Furthermore, ATS compilers must work for other user groups and their codes—for instance, academic and NNSA partners.

As code teams begin to test drive Livermore's exascale systems, Gyllenhaal points to the coordinated compiler and runtime hardening efforts of his team, HPC vendor partners, and the Tri-Lab Center of Excellence (see *Collaboration is key,* pp. 27–29). "The primary vendor compilers provided for El Capitan's early access systems were new to us. Together we've found and fixed a number of issues," he says. "If our codes compile on the early access systems, we expect they'll compile on El Capitan."

# Packaging for everyone

The software required to run scientific codes on a supercomputer is, in a word, complicated. A massively parallel computing environment uses an operating system, compilers, schedulers, workflow managers, debuggers, and much more—all of which are customized for the hardware.

When a researcher wants to run their multiphysics code on a high-performance computing (HPC) system, the build and installation process has to identify the latest software versions needed and resolve the inevitable compatibility issues. Doing this manually is tedious and error-prone, and the reason why automated package management exists.

Perhaps no one is more prepared to address package management on El Capitan than the Spack team. Spack is an open-source package manager created at LLNL over a decade ago and now used widely throughout the HPC community. It was the packaging solution for the Department of Energy Exascale Computing Project, and now it's primed for the installation demands of exascale machines.

As popular as Spack is, however, not all scientific workloads use it. Computer scientist Greg Becker points out, "We need to be able to support El Capitan users regardless of their packaging solution—even if they aren't using a package manager at all and instead go the route of 'I install it all myself'."

This wider view of software installation, combined with lessons learned from previous supercomputer procurements, led to the formation of the Packaging Working Group with members from Lawrence Livermore, Argonne, and Oak Ridge national labs alongside El Capitan hardware vendors Hewlett Packard Enterprise (HPE) and Advanced Micro Devices

*We need to be able to support El Capitan users regardless of their packaging solution.*

**Greg Becker**
Packaging Working Group

Inc. (AMD). Becker continues, "The working group is the main conduit through which we talk to HPE and AMD about how our scientific software will interface with their system software. For example, we work closely with the HPE folks to make the custom Cray programming environment more usable and more like every other Linux system."

One major challenge has been managing the interaction between both vendors' system software. "Often packaging is simply a matter of system software working with user software," Becker explains. "In this case, we have relationships between the AMD software and HPE software, plus the relationships between our software and both of those. Some packages care a lot about exactly how the three-way connection is done."

LLNL Spack developers and external collaborators


According to Todd Gamblin, who created Spack and co-leads the working group, El Capitan's packaging solutions depend on the open-source community. He states, "Vendors don't want to reinvent software capabilities if they don't have to. We've built something that not only our scientists are asking for, but many of the vendors' other customers as well. We've even helped HPE increase their open-source contributions. Software projects such as Spack and Flux provide a way for the Lab to guide the future of HPC."

The group has also improved compatibility and usability by upgrading the Cray programming environment with MPI compiler wrappers and Linux-like installation features, as well as fine-tuning continuous integration (CI) processes that combine Cray CI with cloud-based CI solutions. Furthermore, they have implemented Spack's binary build caches, which speed up installations by using previous builds instead of building from source. Gamblin adds, "The working group has had some major victories. El Capitan is going to be an awesome machine. It's all very exciting."

# Prepping for performance

*It will be rewarding to see how all of our efforts contribute to the mission.*

**Olga Pearce**
Tools Working Group

Eavesdrop on enough conversations at the Livermore Computing Center and you'll repeatedly hear what it takes to "get application codes up and running on El Capitan." Just as lumber and nails alone don't make a house, high performance computing (HPC) hardware components require specialized software tools—and experts to wield them—in order to operate. Led by Matthew LeGendre, El Capitan's Tools Working Group focuses on correctness, debugging, and performance analysis tools.

Correctness software executes rigorous testing programs to confirm an application is running as intended, whereas debugging tools detect and fix errors that occur even when an application runs on the same hardware many times. LLNL teams use a range of performance analysis tools to monitor an HPC system's power management, throughput, resource utilization, and other metrics.

The Tools Working Group's mandate boils down to efficiency and accuracy. "We're concentrating on your code once it's running," explains computer scientist Olga Pearce. "If it's crashing or not producing correct answers, we use debugging and correctness tools. And when your code is working the way it should, and producing the correct answers, then you need performance tools to measure and monitor how fast it's running."

As with most aspects of the transition from petascale to exascale, and from CPUs to GPUs, these types of tools have undergone significant scale-up development. "At this point in the Sierra integration [2017–2018], we were concerned about the quality of debuggers," recalls LeGendre. "In the intervening years, industry and the HPC community have realized the need to debug code for GPUs. Debuggers are working well at this phase of siting El Capitan."

With more computing power at stake, the group continues to look closely at correctness and performance tools while sharing their expertise with El Capitan's chip manufacturer, Advanced Micro Devices Inc. (AMD). The supercomputer will have AMD's new integrated CPU/GPU processors, called APUs. Fittingly, the A in APU stands for *accelerated*. These processors will make up 100 percent of the machine's peak speed, providing enormous floating-point performance alongside increased fidelity and repeatability of calculations.

LeGendre notes, "Part of this collaboration is knowledge transfer. AMD is learning how to build performance tools, while we're learning how to use their processors." The Tools Working Group also leverages open-source software like HPCToolkit as well as contributions from El Capitan's interconnect and storage partner, Hewlett Packard Enterprise (HPE). "We are a pretty demanding customer," Pearce adds. "Our application teams want to understand the performance of their codes to a much finer degree than the other customers these vendors may have dealt with in the past."

User support is crucial in this new ecosystem. Application teams accustomed to Sierra's NVIDIA GPUs will need to learn how their codes work on El Capitan's AMD APUs. "Users will have different interfaces and workflows, but they'll find some familiarity within what these tools are doing," LeGendre says.

Both Pearce and LeGendre refer to El Capitan's simulation capacity as "a game changer." As the NNSA ramps up stockpile modernization efforts, the importance of investing in an exascale supercomputer remains front of mind for the teams working to bring the system online. Pearce states, "Scientists will be hard at work using El Capitan to respond to the Lab's mission, and it will be rewarding to see how all of our efforts contribute to that."

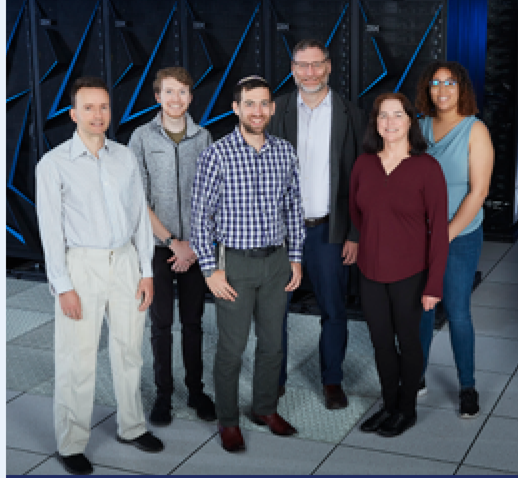

Members of the Tools Development Group

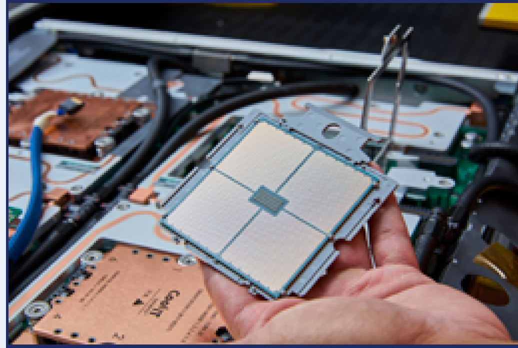

An AMD Instinct® MI300A APU above an open El Capitan blade

# Messaging and math

*What we're doing will enable El Capitan to handle a class of workloads that was never possible before.*

**Matt Leininger**
Messaging Working Group

The road to El Capitan isn't a single lane. It's more like a superhighway, with many specialized working groups solving an array of complicated challenges. While some of these pursuits may appear to travel in different lanes, they're all heading in the same direction. "Working groups allow LLNL to work shoulder to shoulder with vendors to influence the technology ending up in El Capitan. They wield a surprisingly large amount of influence," says Terri Quinn, Livermore Computing program lead.

Matt Leininger leads the Messaging Working Group, which is responsible for integrating the Hewlett Packard Enterprise (HPE) Slingshot interconnection network into the El Capitan ecosystem. Slingshot will enable large-scale calculations to be performed across many nodes. He explains, "The network interconnect is a critical technology for any of our high-performance computing [HPC] platforms. Its ability to efficiently move data between multiple processing nodes and storage permits a wide variety of complex simulations and scientific workloads—from traditional HPC to machine learning and data analytics."

Collaborating with Oak Ridge, Sandia, and Los Alamos national laboratories, the Messaging Working Group focuses on network performance optimization for these workloads, including the necessary nonrecurring engineering around messaging processes. Additionally, Leininger notes, "Our team includes vendor partners from HPE and AMD [Advanced Micro Devices Inc.]. HPE provides the Slingshot interconnect, and AMD provides the APUs [accelerated processing units], so both vendors are integral to solving issues with messaging."

Central to this effort are two software libraries related to the message passing interface (MPI) standard: the open-source Libfabric software and HPE's proprietary MPI solution. Message passing helps different programs and processes work together on parallel computing architectures. Leininger points out, "It's always good to support more than one MPI on an HPC system. This allows us to compare and contrast the two implementations and have an alternative for when bugs show up in one implementation."

To deploy the Slingshot, the group must manage the associated firmware, system software, automation processes, and more—all of which are first-generation technologies. El Capitan's early access systems installed at the Livermore Computing Center have been extremely useful in this regard. "Most of our early access systems are being used to test on a single node or small number of nodes, so the next step will be testing the messaging on more and more nodes as El Capitan is deployed," Leininger says. "What we're doing will enable El Capitan to handle a class of workloads that was never possible before."

Traveling beside the MPI communication lane on the El Capitan superhighway is a team dedicated to mathematical libraries. "These libraries are key to, and ubiquitous in, HPC modeling and simulation codes," states computational scientist Ramesh Pankajakshan. "The arrival of GPUs and novel architectures like El Capitan's have completely thrown a spanner in the works—in a good way. There is a lot more scope for making these libraries faster. It's an open research topic."

HPE and AMD have their own math libraries, and application teams may use other libraries such as hypre (linear solvers multigrid methods), LAPACK (linear algebra), SUNDIALS (nonlinear solvers), or PETSc (numerical solvers). According to Pankajakshan, who leads LLNL's contributions to El Capitan's Math Libraries and Machine Learning Working Group, fine-tuning math libraries for complex codes on GPU- and APU-based systems requires "a lot of back and forth." He adds, "We're unique among the vendors' customers because we have, say, applications with millions of lines of code written over 40 years, and we're using cutting-edge architectures."

*[Math] libraries are key to, and ubiquitous in, HPC modeling and simulation codes.*

**Ramesh Pankajakshan**
Computational scientist

El Capitan's memory and storage setup affects how math libraries execute calculations. Pankajakshan notes, "Previously, moving data back and forth was computationally expensive, so you had to pick an execution space and stay there. Now you can execute a calculation in its ideal execution space and switch as often as required. With El Capitan, in particular and APUs in general, there are lots of opportunities. It's always fun to work on the cutting edge."

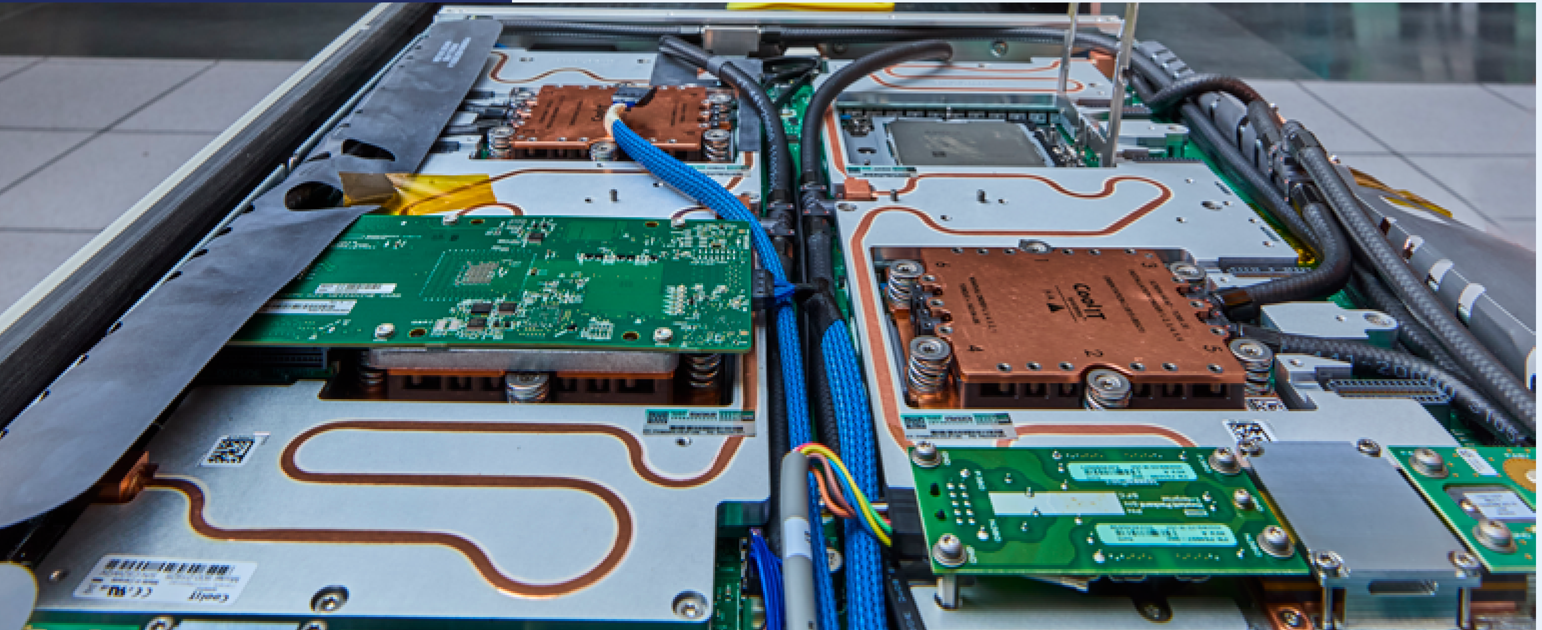

An El Capitan node

## Collaboration is key

Under the moniker Tri-Labs, high performance computing experts at Lawrence Livermore, Los Alamos, and Sandia national labs work closely on initiatives for the NNSA Advanced Simulation and Computing (ASC) program. For example, this team develops and deploys the custom Tri-Lab Operating System Stack (see *The right operating system,* pp. 7–9). Another area of collaboration is the Center of Excellence (CoE), which extends to El Capitan's hardware partners Hewlett Packard Enterprise (HPE) and Advanced Micro Devices Inc. (AMD).

The CoE is more than an administrative entity. Whereas working groups may include additional organizations with a stake in El Capitan's success, the CoE facilitates communication of requirements and pain points. "Our number one goal is that on Day 'zero', when this amazing system comes online, our programmatic applications from all three labs are ready to make efficient and productive use of it," says Judy Hill, Livermore's CoE lead. Anna Pietarila Graham from Los Alamos and James Elliot from Sandia round out the Tri-Lab's CoE leads.

Vendor and CoE staff work side by side with Tri-Lab teams. One key activity is the recurring hackathon, which brings together developers and hardware experts from each participating organization to fine-tune applications on El Capitan's early access systems. The CoE leads emphasize the value of hackathon interactions. "A core group of code teams attend all the New Mexico hackathons [at Sandia], and some travel to Livermore because they find it really useful to have not just the vendors there, but to have three days of dedicated time to work on their codes without distractions," notes Pietarila Graham.

*I compare [our challenges] to race car drivers versus regular drivers. The problems are of a different type and scale.*

**James Elliot, Sandia**
CoE lead

*A big part of the CoE's responsibilities is advocating for all future users.*

**Anna Pietarila Graham, LANL**
CoE lead

Between hackathons, the CoE conveys bugs and feature requests to HPE and AMD so that the overall software stack matures, scales, and stabilizes. "We help the vendors understand our challenges. From their point of view, our codes are massive and complex and written over many years, and their standard customers aren't doing the kind of science we are," Elliot explains. "We run into edge-case problems that vendors wouldn't normally encounter—the kinds of problems that arise when you push the limits. I compare it to race car drivers versus regular drivers. The problems are of a different type and scale."

This effort is a win–win: Tri-Lab expertise helps the vendors improve their R&D when these unique issues are resolved. For instance, issues with application correctness and performance reported through the CoE have led to vendor-released software updates to address those issues. Both vendors acknowledge this mutually beneficial experience.

"Our close collaboration with LLNL and AMD has been absolutely vital in designing and deploying a transformative system like El Capitan," states Gerald Kleyn, vice president, Customer Solutions, HPC & AI, HPE. "Our work together has spanned multiple years and has involved partnering on overall design, updating the LLNL facility, and ensuring application readiness through the CoE. This has enabled HPE to share our expertise and, in turn, learn from other world-class experts across a range of fields as we test new features and gain insights into priority applications that will secure better outcomes for our customers."

Nick Malaya, an AMD Fellow in High-Performance Computing, adds, "The collaboration with Livermore Computing and the Tri-Labs has been instrumental in ensuring El Capitan's hardware and software readiness. The CoE has assembled some of the world's foremost experts in computer architecture, computer science, and scientific computing. The team has been able to anticipate and address complex challenges in software porting and tuning, while also pushing the boundaries of what's possible in HPC."

Even with cutting-edge software and hardware, El Capitan's productivity will hinge on robust user support. Pietarila Graham states, "A big part of the CoE's responsibilities is advocating for all future users—not only making sure the applications will run, but that they'll run for regular users in addition to experienced code developers."

When the NNSA's first exascale supercomputer accepts its first user login and fires up its first application, the Tri-Lab collaboration will continue as the labs tackle their shared national security mission. Hill says, "I'm excited for those moments when El Capitan will surpass our expectations."

*I'm excited for those moments when El Capitan will surpass our expectations.*

**Judy Hill, LLNL**
CoE lead

El Capitan Center of Excellence hackathon

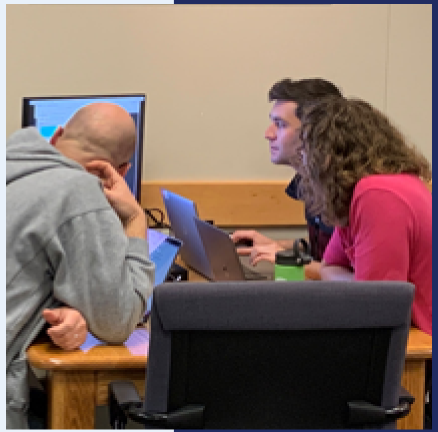

Technicians working on the final day of installation of El Capitan blades

# Industry investment

Perhaps no other square mile of California hums with as much activity and purpose as Lawrence Livermore's main campus. Amid the world-class experimental labs, specialized equipment, and manufacturing facilities, a workforce of more than 9,000 serves a national security mission through innovative R&D. One of the most exciting places these days is the Livermore Computing Center as the exascale El Capitan supercomputer and its unclassified siblings are brought online—a process that's considerably more complicated than unboxing a laptop.

The Lab's high performance computing (HPC) success depends in part on strategic relationships with U.S. commercial companies. From the early UNIVAC computers to today's massively parallel systems, computer scientists and software engineers have worked closely with industry experts to design and field each generation of machines.

Livermore Computing Chief Technology Officer Bronis de Supinski says, "Success at huge tasks cannot be achieved alone, whether you are an individual or a big organization. It takes partnership. This maxim is true for building supercomputers, and Livermore Computing's long history of strong, deep industry collaboration is born from our desire to succeed at that huge task."

CORAL-2 Quarterly Meeting

Hewlett Packard Enterprise (HPE) and Advanced Micro Devices Inc. (AMD) joined Livermore and the NNSA in the years-long El Capitan procurement and deployment. "The collaboration has worked extraordinarily well," states HPE system architect Chris Brady. "This is the best customer relationship I've been a part of. It's both pleasant and productive."

El Capitan's computational speed comes from AMD's unique MI300A accelerated processing units (APUs), which combine CPUs and GPUs for higher efficiency, higher resolution simulations. Debuting in El Capitan, the APUs are bolstered by HPE's Slingshot interconnect network, accelerator blades, and near-node storage components nicknamed "the Rabbits" (see *Storage in the exascale era,* pp. 9–12). HPE also provides the liquid-cooled cabinets housing all of this hardware in LLNL's primary machine room.

Terri Quinn, associate program director for Livermore Computing Systems and Environments, points out, "Our objective was to deliver to the NNSA the most capable computer possible within the given budget. To do this, we asked the vendors to build us a system that required them to push their technical capabilities hard. HPE and AMD have done this with El Capitan. They had the expertise and the courage to take risks."

Aiming for the exascale threshold was a challenge for everyone involved. "The scale is huge. Everything is much more difficult, whether we're upgrading firmware or figuring out how to run all the cables. For example, with smaller systems, five or six people can arrange and install the cables in a few days. This time, cabling took months," Brady notes.

"We had 13 project teams working on different aspects of delivery and 144 milestones along the way. There was incredible complexity in all the tandem work streams as well as the integration points between programming environments and system software," explains Gina Norling, former HPE engineering program manager and now AMD's senior engineering manager for the Center of Excellence HPC/AI.

*Livermore Computing's long history of strong, deep industry collaboration is born from our desire to succeed at that huge task.*

**Bronis de Supinski**
LLNL

*Every day there's a challenge you didn't see coming. Transparency plus time equals trust.*

**Brick Stephenson**
HPE

Chris Brady (HPE), Keith Shields (HPE), Adam Bertsch, (LLNL), Pythagoras Watson (LLNL) with the final El Capitan blade

Brick Stephenson, HPE director of program management for strategic customer engagements, adds, "The only way I can keep my head around the whole endeavor is to realize that I'm not the expert on any one thing. I don't have to know all the answers; I just need to know who to loop in. Every day there's a challenge you didn't see coming. Transparency plus time equals trust."

According to Keith Shields, HPE's senior director of strategic program management, the key to a large-scale deployment is managing the small-scale aspects. He advises, "Break it into physical and time-bound components. Where do the power drops go? Where does the plumbing go? Then it's not a scramble when the cabinets arrive."

In addition, HPE staff drew on lessons learned with AMD when building the exascale Frontier supercomputer at Oak Ridge National Laboratory, says hardware project manager Randy Law. "Aspects of El Capitan were familiar because of Frontier, and what we discovered during that project benefitted this one," he states.

As prominent contributors to the global HPC community, these industry partners share LLNL and NNSA's commitment to El Capitan's success. Stephenson, who brought five decades of computing experience to the project, states, "I look forward to hearing about the new scientific discoveries this system will enable." A relative HPC newcomer, field services engineer Jonathan Gutierrez notes, "This is a big job that takes lots of people. And assuming El Capitan debuts high on the Top500 list, I'll be able to say, 'I helped with that.' We hope our excitement carries over to the users."

The El Capitan project—which included siting three early access systems as well as the unclassified Tuolumne and RZAdams machines—highlights enduring connections between industry and national labs. Shields, who began his HPC career nearly 30 years ago as a software developer, routinely visited LLNL throughout the process. He notes, "I was born in Livermore and my dad was a Lab employee, so I've come full circle."

Stephenson, too, has Livermore ties. He was a field engineer at the Lab in the 1970s before moving to Cray, which was later acquired by HPE. Gutierrez worked at both Sandia National Laboratories campuses before pivoting to an HPC career just a few years ago. He will continue to work onsite at the Livermore Computing Center, helping to run the day-to-day operations. "Learning from experienced folks has been wonderful," Gutierrez says.

HPE technologists with the final blade of El Capitan

For many of LLNL's vendor colleagues, bringing El Capitan online has been a professional highpoint. Brady even put his retirement plans on hold to see the project through. "It's a really wonderful cap to my career," he states. "And frankly, after El Capitan, what else could be as exciting?" Norling draws motivation from the project's scale and difficulty, adding, "I have to pinch myself that I'm involved in this program. El Capitan has so many amazing technical aspects, but it comes down to the people. At the end of the day, it's those relationships and connecting folks to make a solution that's better than we could have ever produced on our own."

Livermore Computing staff and facilities team


WSC code teams

The last blade