

# Description of Facilities and Resources

Livermore Computing | Lawrence Livermore National Laboratory (LLNL)



LLNL computational scientists are supported by Livermore Computing (LC), which delivers a balanced high performance computing (HPC) environment with constantly evolving hardware resources and a wealth of HPC expertise in porting, running, and tuning real-world, large-scale applications. Currently, LC delivers over 3 exaflops of compute power, massive shared parallel file systems, powerful data analysis platforms, and archival storage capable of storing hundreds of petabytes of data. This balanced hardware environment supports key collaborations between LLNL application developers and LC experts. This involves the creation, production use, performance monitoring, and analysis of results of HPC parallel applications in a wide variety of scientific disciplines.

In addition to general access for development and debugging, researchers access allocations on these resources through several competitive internal programs. With these allocations, science and technology directorates sustain and enhance institutional capabilities to help ensure mission successes. Such efforts benefit the entire Laboratory by developing the expertise of the next generation of scientists and engineers who will, in turn, use these capabilities for continued growth of Laboratory programs. Furthermore, researchers with Strategic Partnership Projects projects and external collaborators needing HPC access can purchase cycles on the Multi-programmatic and Institutional Computing systems.

Major systems include the 2.88-exaflops El Capitan and its sibling systems; the 288-petaflops unclassified Tuolumne with a combined file system capacity of 401 pebibytes; and the 32-petaflops RZAdams; several El Capitan Early Access Systems, including RZVernal at 6.9 petaflops, Tioga at 5.8 petaflops, and Tenaya at 4.4 petaflops; the 5-petaflops Magma; and additional large multi-core, multi-socket Linux clusters with a variety of processor types. Our El Capitan systems feature HPE/Cray Slingshot interconnect technology and contain AMD MI300A APU processors.

We also have several cutting-edge AI powerhouses: three stand-alone SambaNova systems and the Cerebras wafer-scale engine.

Computational scientists may also take advantage of several testbeds for evaluating next generation hardware and software. Researchers use these testbeds to investigate hardware advances in areas such as multi-core processors, networking technologies, I/O, GPUs, memory, and power-aware HPC (via a dedicated power lab), as well as investigations of software technologies. In addition, LC hosts production collaboration environments that facilitate the sharing of scientific data among international research groups, including the Green Data Oasis.

Several facilities house the simulation infrastructure at LLNL. The largest (LEED-certified) simulation facility offers 48,000ft<sup>2</sup> and 85 MW of power for systems and peripherals, as well as additional power for the associated 28,000-ton cooling system. Engineering and facilities staff maintain it in a physically secure environment.

The balanced LLNL simulation environment includes Lustre multi-cluster file systems, High Performance Storage System (HPSS)-based archival resources, a high-speed Ethernet and InfiniBand networking infrastructure, advanced visualization resources, and a rich tool environment, as described on page two.



## LC RESOURCES

### High Performance Archival Storage

LC provides high-performance archival storage services via the HPSS. A world-class array of hardware integrated beneath HPSS includes disk arrays, tape subsystems, mover nodes, networks, robotics, and hundreds of petabytes of media.

### File Systems

HPC systems require high performance storage, which LC provides in multiple forms: open-source Lustre and the all-flash VAST. These systems are mounted across multiple compute clusters and deliver high-performance global access to data.

### LC Networking

LC's simulation environment includes a software-defined, 100G high-speed datacenter Ethernet networking infrastructure, as well as InfiniBand SANs. LC employs InfiniBand, Omni-Path, and HPE Slingshot fabrics for high-speed interconnects.

### Containerization

LC has deployed a selection of on-prem-cloud, science-focused platform capabilities to enable persistent data services, HPC workflow tools, LLMs, and machine learning frameworks. These services will include container management and orchestration, and an object storage system.

### Visualization Facilities

LC runs two large visualization clusters—the CPU-only Tron on the SCF, with a peak speed of 434 teraflops, and the hybrid-architecture Matrix on the OCF, with an initial peak speed of 4 petaflops—as well as visualization partitions on our ATS machines. LC supports RealVNC on login nodes and NICE DCV on batch nodes, and operates several visualization theaters, ranging from auditoriums with PowerWalls to smaller collaboration spaces.

### HPC Tool Environment

LC provides a stable, usable, leading-edge parallel application development environment that significantly increases the productivity of applications developers by enabling better scalable performance and enhancing application reliability. The tool environment includes high-performance compilers, debuggers, analyzers, editors, and locally developed custom libraries and application packages for software development. LC also provides a range of web-hosted tools: GitLab for repository hosting and continuous integration; Orbit for analytics in Python; and Confluence and JIRA for collaboration. LC enables users to exploit emerging technologies while making software development within the compute center as easy as it is in the cloud.

## PRIMARY PRODUCTION COMPUTING PLATFORMS

### El Capitan, Tuolumne, and RZAdams

At 2.79 exaflops peak, with 43,808 AMD MI300A APUs and HPE Slingshot interconnects, El Capitan is one of the most powerful supercomputers in the world. Its sibling system, Tuolumne, at 288 petaflops, is completely dedicated to unclassified computing. RZAdams, at 32 petaflops, shares their architecture.

### El Capitan Early Access Systems

Our 3 early access testbeds—RZVernal (6.9 petaflops), Tioga (5.8 petaflops), and Tenaya (4.4 petaflops)—combine 1 HPC- and AI-optimized 3<sup>rd</sup> generation AMD EPYC 64-core CPU with 4 purpose-built AMD Instinct MI250X GPUs.

### Bengal, Dane, and RZHound

The 10.7-petaflops Dane, 7.9-petaflops Bengal, and 2.7-petaflops RZHound from Dell contain Intel Sapphire Rapids CPUs and Cornelis interconnects.

### Magma

A 5.3-petaflops Penguin Computing Relion system, comprised of 752 nodes with Intel Xeon Platinum 9242 processors, is included. The cluster has 293 terabytes of memory, liquid cooling, and an Intel Omni-Path interconnect.

### Jade

Jade, a Penguin Computing system, has 2,688 nodes, 343 terabytes of memory, and a peak speed of 3.2 petaflops. It uses Intel Broadwell processors. RZTopaz is a smaller 748-node cluster with the same architecture as Jade.

For a complete list of nearly 20 production compute platforms supported by LC, see the [Livermore Computing Systems Summary](#).

## INSTITUTIONAL RESOURCES AND FACILITIES

LLNL researchers benefit from an institutional IT infrastructure that provides desktop support and experts in server technologies. The latter includes virtualization expertise, which has been applied to provide multiple operating systems on shared resources and to create a wide variety of virtual machines to leverage resources across LLNL. Networking service is also provided by an enterprise team. Connections into LLNL include ESnet and a wide variety of programmatic networks.