DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

LLNL-TR-2009298

Lawrence Livermore National Laboratory

El Capitan Readiness L2 Summary

M LeGendre, A Bertsch

August 2025

Disclaimer

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

El Capitan System Readiness

L2 Milestone Summary

July, 2025

Adam Bertsch, El Capitan Systems Integration Lead
Matthew LeGendre, CSSE Lead

Livermore Computing

Lawrence Livermore National Laboratory

LLNL-TR-2009298

NNSA National Nuclear Security Administration

Overview

This document is the written summary for the El Capitan System Readiness CSSE L2 Milestone from Lawrence Livermore National Laboratory. It briefly summarizes the efforts involved in preparing the El Capitan supercomputer for the local compute environment, and it describes how those efforts fulfill the requirements of ASC L2 #9235.

L2 Prose and Requirements

Milestone (ID# 9235): El Capitan System Readiness		
Level: 2	**Fiscal Year**: FY25	**DOE Area/Campaign**: ASC
Completion Date: 09/30/25		
ASC WBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Program Plan Target: Appendix J, target CSSE-1.a, CSSE-2.a		
Description: This milestone will certify that El Capitan is ready for integration into the classified computing environment. System hardware and software will be benchmarked and tested for correctness and reliability. Plans for system support and accessibility will be documented and communicated to the tri-lab community. At least three unclassified ASC applications will be demonstrated on the platform.		
Completion Criteria: Hardware deliveries from vendor to site are complete and on-site installation of the system by the vendor is substantially complete to the extent that is contractually required; contractual requirements for formal hardware acceptance have been substantially completed; system software needed for basic operation of the system is delivered, tested, and demonstrated to be operational; vendor has completed onsite benchmark and application performance testing and demonstration; and the system is ready to begin onsite integration into local computing environment. A tri-lab review will be conducted and the milestone results documented according to ASC milestone requirements.		
Customer: ASC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: CSSE, FOUS, IC, PEM, V&V, SNL and LANL		

In summary, the technical requirements for this L2 are:

- Hardware deliveries from vendor to site are complete.

- On-site installation of the system by the vendor is substantially complete.
- System software needed for basic operation of the system is delivered, tested, and demonstrated to be operational.
- Vendor has completed onsite benchmark and application performance testing and demonstration.
- Plans for system support and accessibility will be documented and communicated to the tri-lab community.
- At least three unclassified ASC applications will be demonstrated on the platform.
- The system is ready to begin onsite integration into local computing environment

And the administrative requirements are:

- A tri-lab review will be conducted.
- Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.

This document is part of the professional documentation and is the “written summary” that accompanies viewgraphs. The remainder of this document focuses on the delivery of the technical requirements.

System Readiness

This section summarizes the technical efforts that fulfill the requirements for the System Readiness L2.

Delivery

- ✓ Hardware deliveries from vendor to site are complete.
- ✓ On-site installation of the system by the vendor is substantially complete.

Figure 1: Delivery of El Capitan Compute Nodes began March 2024

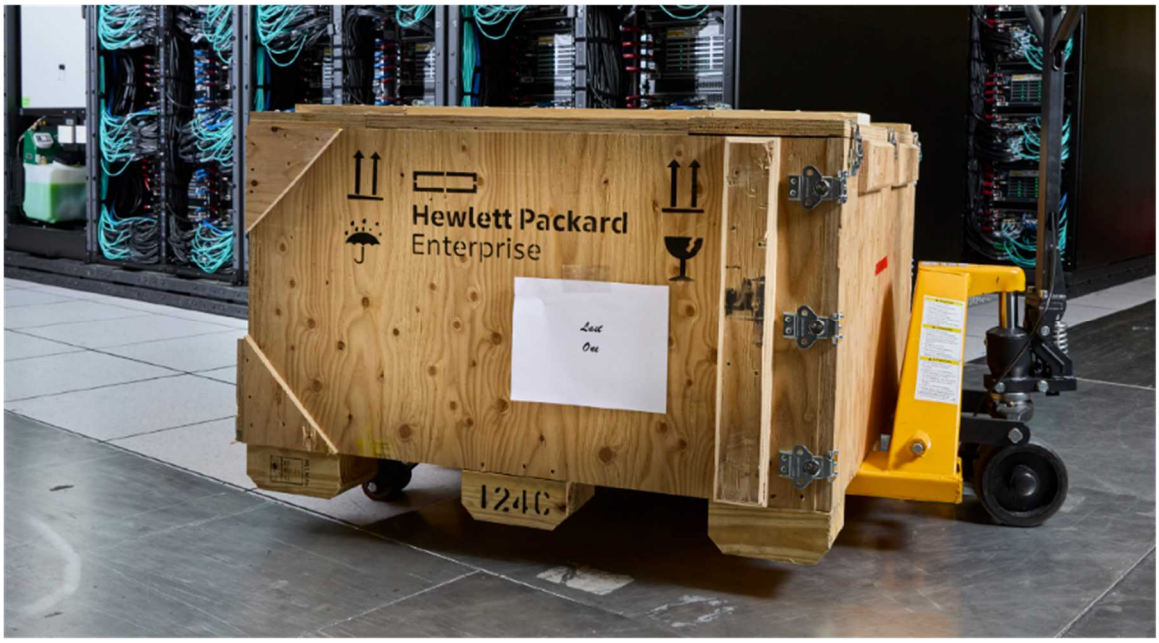

Figure 2: “Last One” Compute Blade Delivered May 2024

Figure 3: HPE Conducted System Installation

Figure 4: El Capitan was Accepted in October 2024

The delivery and acceptance of El Capitan hardware was an involved process taking much of calendar year 2024. Compute blades were delivered in Spring of 2024, which were installed by HPE staff through May. Early access to administrators began in June, and El Capitan was

subsequently made available to the first users on the CZ network. The final system was accepted in October of 2024.

System Software and Benchmarking

- ✓ System software needed for basic operation of the system is delivered, tested, and demonstrated to be operational.
- ✓ Vendor has completed onsite benchmark and application performance testing and demonstration.

TOSS

El Capitan is the first ATS system to use TOSS (Tri-Lab Operating System Stack) as its operating system. El Capitan uses TOSS-4, which is built around Redhat Enterprise Linux 8. The TOSS team integrated support for the MI300a APUs and HPE's Slingshot network into their distribution. They continue to work with RedHat and the community around MI300's to adapt system software for the HPC environment.

Programming Environment

HPE provides the Cray Programming Environment (CrayPE) with compilers, ROCm GPU runtimes, MPI, tools, and the utilities needed to build and maintain simulation codes on El Capitan. On top of CrayPE LLNL provides the TCE programming environment, which expands the available tools, compilers, GPU runtimes, and compiler wrappers.

The compiler wrappers from TCE provide users with a smoother experience linking ROCm and MPIs. The wrappers recognize and correct many common mistakes, and they provide bug workarounds at compile time. While many users find these compiler wrappers easier to use, they do hide the underlying CrayPE. Users are able to load modules that provide them with direct access to the CrayPE, which is ideal for reproducing bugs and is appropriate for applications already written to build across Cray environments.

Flux

El Capitan is the largest system yet to use Flux as its system resource manager. Flux provided support for heterogenous resource management needed to manage El Capitan's MI300 APUs and rabbit storage system. The Flux team scaled to El Capitan's full size, managing both large jobs and large numbers of jobs. They hardened the system to detect, drain, and manage hardware issues, and extended Flux with extended system administer interfaces needed for El Capitan.

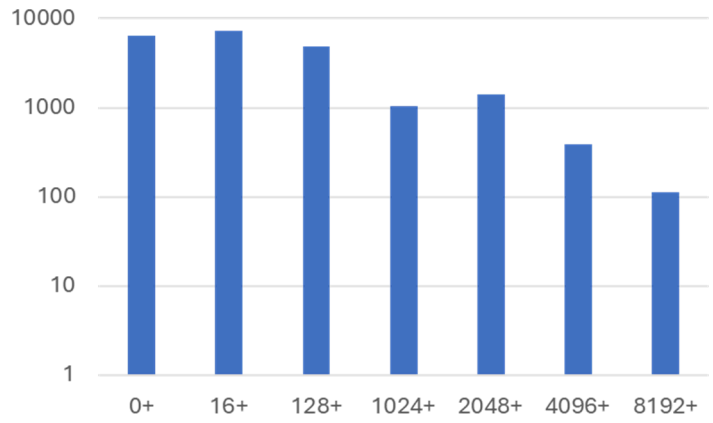

Figure 5: Number of jobs run by size before swing to SCF (1/1/25-3/16/25).

Spindle

Spindle is a tool to provide scalable job launch for applications. During launch at very large scales applications can overwhelm shared file systems searching for and loading libraries, python files, and other start-up files. This massive file system load not just slows down application startup, but also affects other users of the system who may just be trying to access files on the same file system.

Spindle provides a (mostly) transparent mechanism that identifies dependent libraries and python files, and it transfers them to a cache on node local storage. The application then loads the cached versions of libraries and does not overwhelm the shared file system. El Capitan is the first system to run Spindle by default, which required significant investments in hardening spindle for many application types, flux integration, and tuning spindle to assist AI-centric workflows.

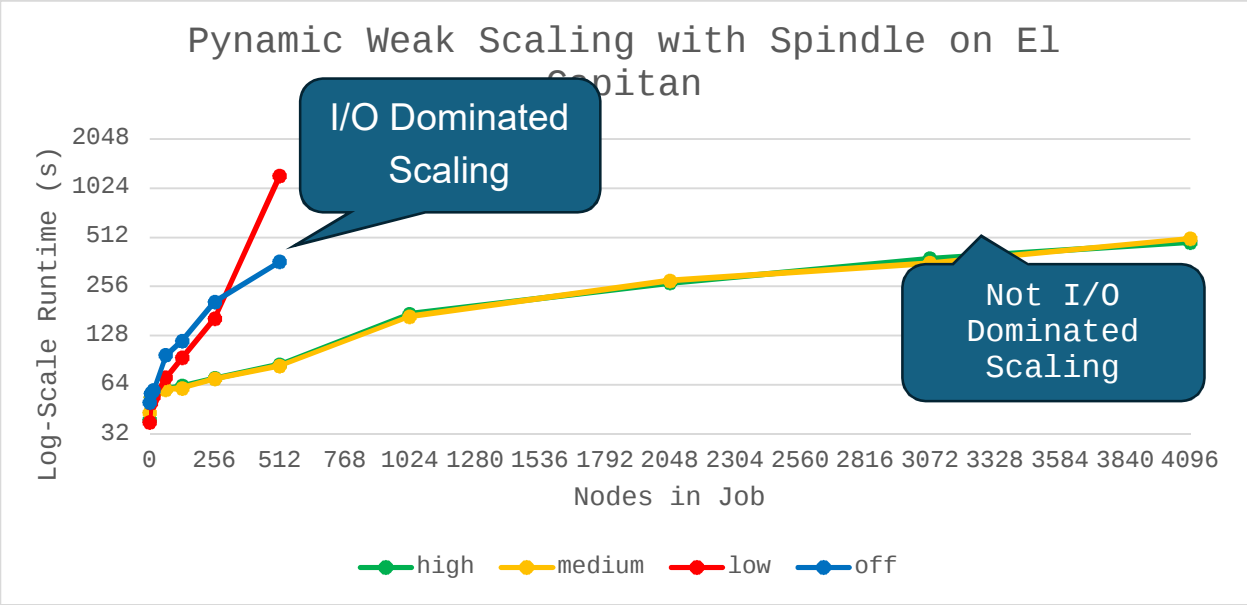

Benchmarks

A suite of benchmarks were run on El Capitan as part of system acceptance. Results of those benchmarks are not discussed in this document, but the list of benchmarks is:

- QMCPACK - Monte Carlo electronic structure
- HACC - Cosmology
- NEKbone - Navier-Stokes CFD
- LAMMPS - Molecular dynamics
- Quicksilver – Monte Carlo transport
- Kripke – Deterministic transport
- AMG Solve – Algebraic multigrid solver
- PENNANT - Hydrodynamics
- BDAS – Big data analytics
- LBANN – Artificial neural network
- Series of microbenchmarks covering MPI, memory, OpenMP, I/O, compilers, scalable python launch
- E3SM – Earth System Model
- Laghos – Hydrodynamics
- AMG Setup – Algebraic multigrid solver

Documentation

- ✓ Plans for system support and accessibility will be documented and communicated to the tri-lab community.

The El Capitan Center of Excellence prepared documentation for users. Non-public documentation is available in LC's confluence space and made available to users with El Capitan accounts. Public documentation is available on LC's public space at https://hpc.llnl.gov/documentation/user-guides/using-el-capitan-systems.

Highlights from this documentation include:

- Resource management with Flux
- Hardware Overview
- Compilers and Programming Environments
- Debugging and Performance Analysis Tools
- System Access

- System Usage Examples

El Capitan is using a new process for ATCC (Advanced Technology Computing Campaign), which aspires to provide a lower-overhead processes for regular customers to obtain system time. The new ATCC process has been discussed and agreed to by the ASC leadership.

Applications

✓ At least three unclassified ASC applications will be demonstrated on the platform.

The major unclassified ASC codes have run on El Capitan. Three highlights are:

Ares
The Ares code ran a 97.5B zone Convergent Rayleigh-Taylor calculation that they compared to Sierra. Using only 2048 nodes they experience a 4 times throughput improvement compared to Sierra's 4096 nodes.

Marbl
The Marbl team demonstrated a high-order ALE simulation of shock driven Kelvin-Helmholtz instability experiment on 2048 nodes.

Miranda
The Miranda team demonstrated an EMIT-LLNL high-resolution wind tunnel simulation running 4096 nodes.

Readiness

✓ The system is ready to begin onsite integration into local computing environment

El Capitan successfully transition to the classified compute environment in March, 2025. Classified early users are actively running on the system.