Scaling El Capitan:
Exascale Computing on a Mission

Judy Hill
Lead, El Capitan Center of Excellence

2025 Argonne Training Program on Extreme-Scale Computing
July 28, 2025

Prepared by LLNL under Contract DE-AC52-07NA27344.

This is a critical time for the nation's nuclear deterrent

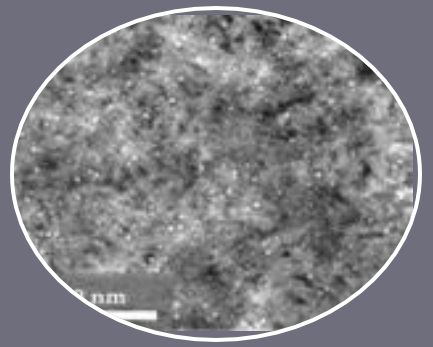

An aging stockpile stressing our assessment capabilities

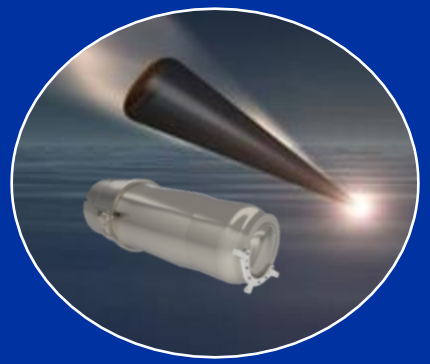

Challenging stockpile modernization efforts under way

Production complex undergoing major modernization

A workforce that is new to the nuclear security enterprise

NNSA's stockpile stewardship mission relies on HPC

Surveillance and assessment of aging weapons

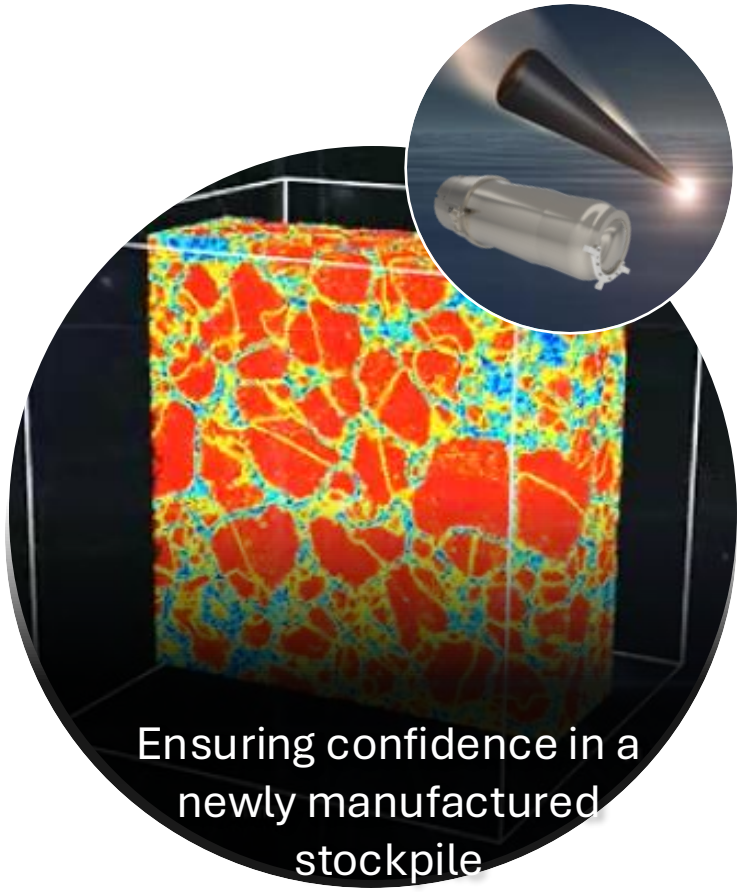

Ensuring confidence in a newly manufactured stockpile

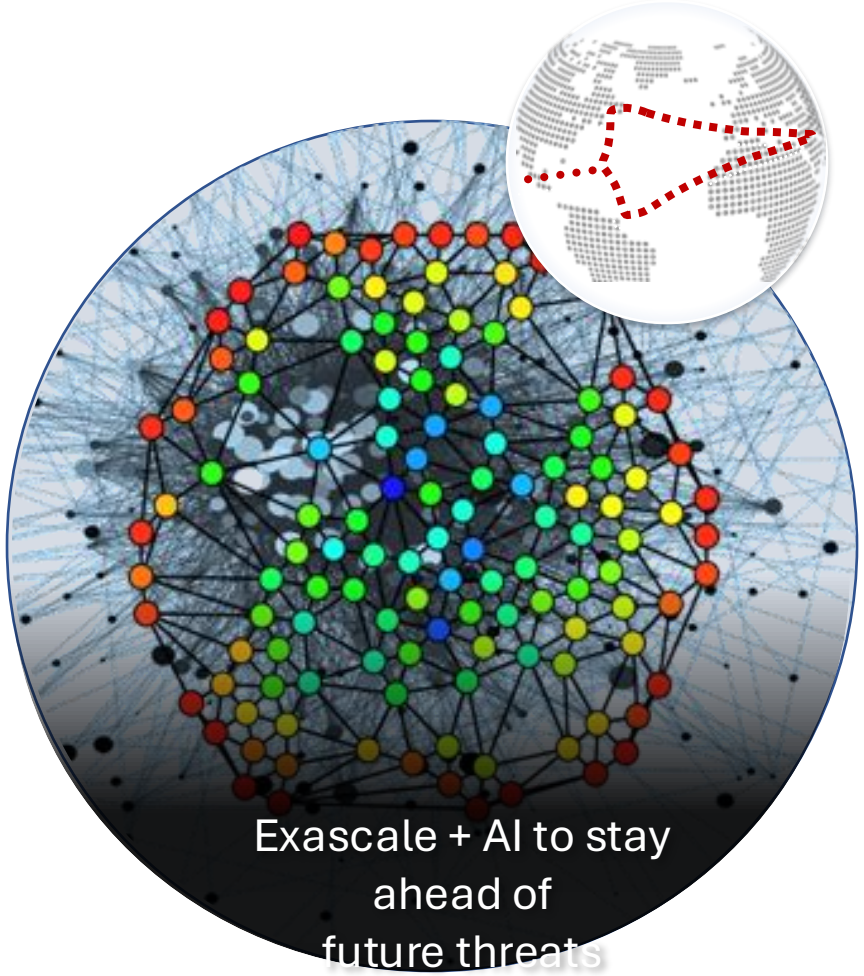

Exascale + AI to stay ahead of future threats

NNSA Exascale computing requirements are driven along three main axes

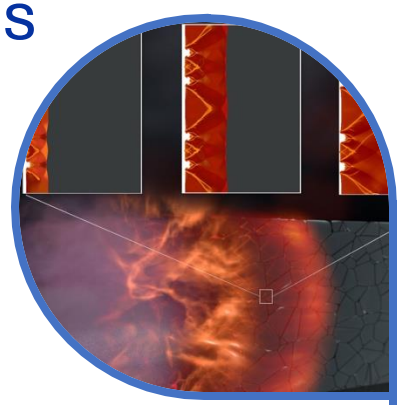

1 Improved physics and algorithms

Cumulative result: Improved accuracy and confidence in simulation results requires unprecedented levels of compute (exascale)

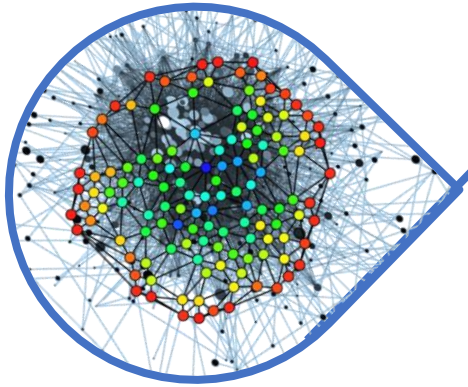

2 Large ensembles for UQ and AI training

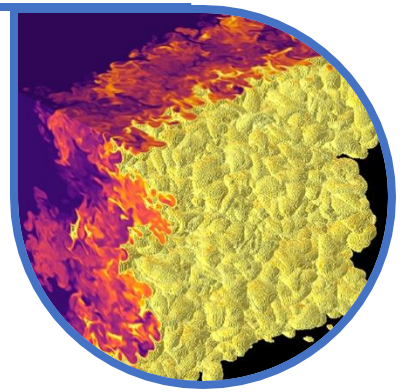

3 Higher resolution in 3D

World’s fastest supercomputer

EL CAPITAN
ADVANCED SIMULATION & COMPUTING

Highlights

- Sustained performance of 1.742 exaFLOPs
- First exascale supercomputer for the NNSA
- Peak speed of 2.79 exaFLOPs
- Built by Hewlett Packard Enterprise with Cray EX architecture
- AMD Instinct MI300A APUs optimized for AI workloads
- HPE Slingshot interconnect and near node local storage technology
- Unprecedented modeling & simulation for national security
- Unclassified sibling system Tuolumne will be used for open science

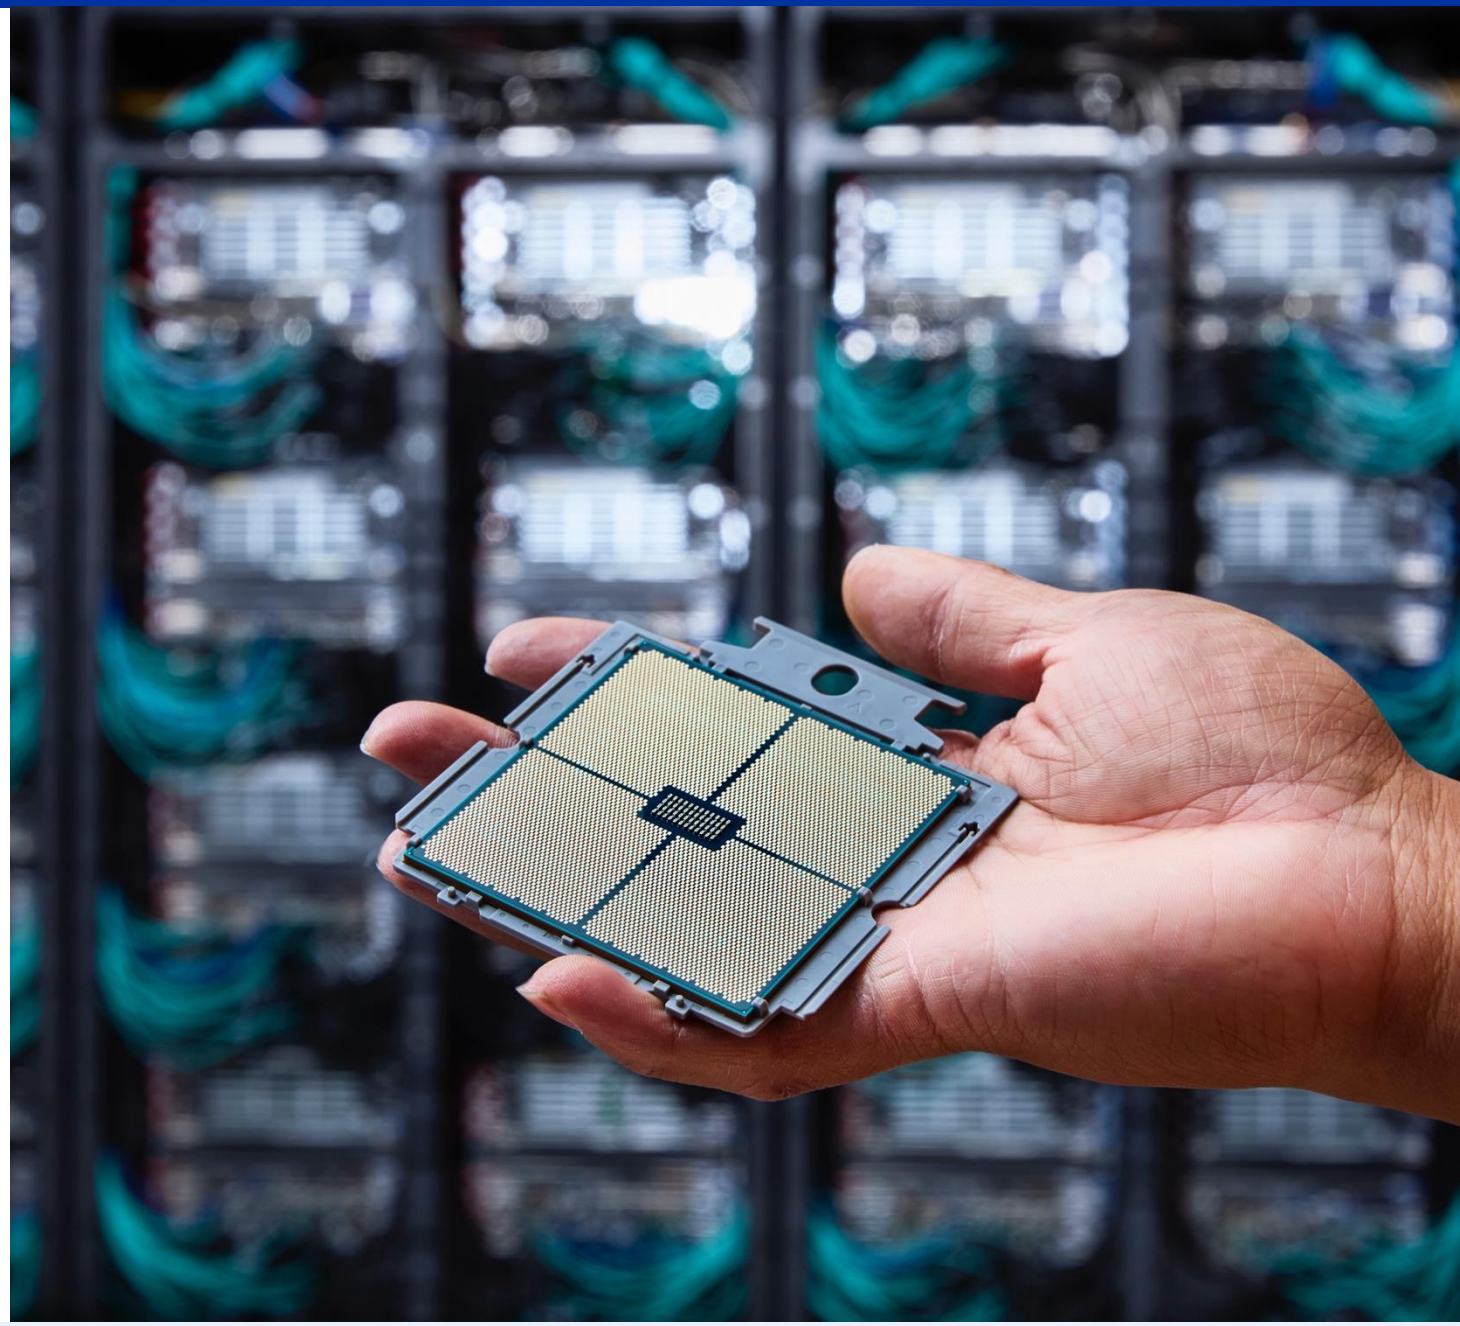

El Capitan is the product of several co-designed innovations

- Major infrastructure upgrade at LLNL to position the lab for the future
- AMD's MI300A, world's first data center APU, directly addresses multiple challenges
- El Capitan includes an innovative near node local storage solution
- Is using LLNL developed software stack enhanced with HPE software

El Capitan is the product of several co-designed innovations

- **Major infrastructure upgrade at LLNL to position the lab for the future**
- AMD's MI300A, world's first data center APU, directly addresses multiple challenges
- El Capitan includes an innovative near node local storage solution
- Is using LLNL developed software stack enhanced with HPE software

HPC exascale infrastructure demands are unprecedented and drove utility upgrades serving HPC at LLNL

Exascale Computing Facility Modernization (ECFM) Project

Status: Project is complete

Construction timeline: 3/2020 to 5/2022

Objective: enable LLNL to operate two exascale class systems simultaneously

ECFM Highlights

- No structural upgrades needed
  - Existing facility had ample square footage with 48,000 SF and structural integrity up to 625 lbs/SF
- Cooling scaled to 28,000 tons with new 18,000 ton cooling tower
  - Loop extended avoiding chillers
- Electrical supply upgraded from 45 MW to 85 MW
  - Capacity = 1771 Watts/SF
  - Dynamic monitoring and control systems to ensure seamless 24/7 HPC operations
  - Two electric utilities tied in parallel at 115kV

Exascale Computing Facility Modernization Project: 3/2020 – 5/2022

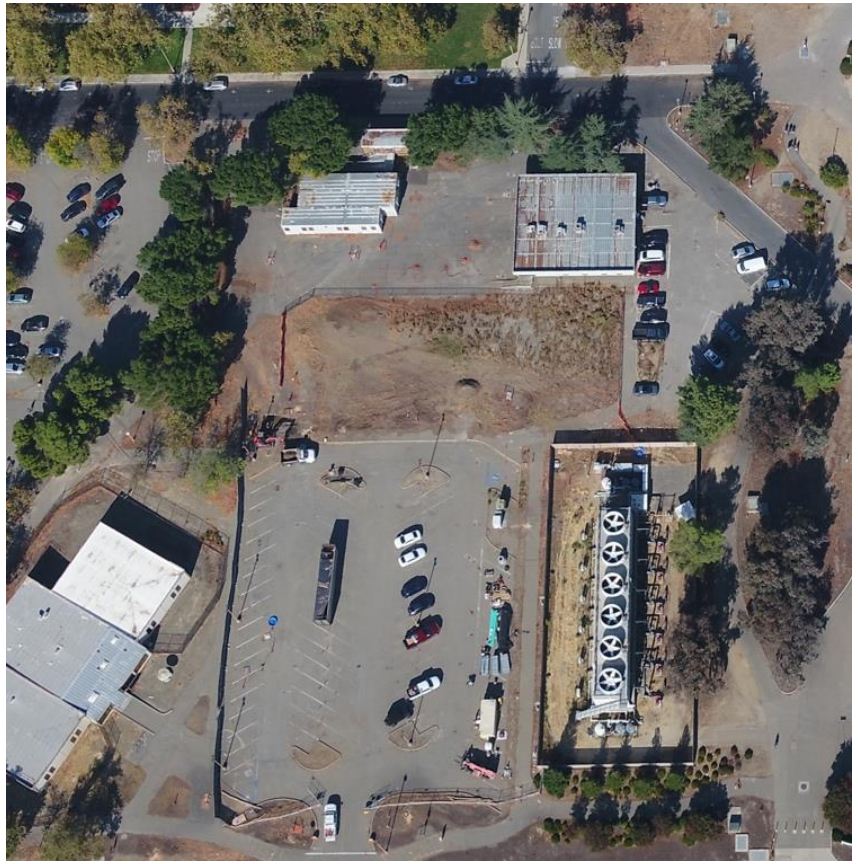

Before

After

ECFM team illustrates scale of 18,000 ton cooling towers

Fielding a supercomputer often requires a facilities construction project in addition to siting the machine

El Capitan is the product of several co-designed innovations

- Major infrastructure upgrade at LLNL to position the lab for the future
- **AMD's MI300A, world's first data center APU, directly addresses multiple challenges**
- El Capitan includes an innovative near node local storage solution
- Is using LLNL developed software stack enhanced with HPE software

AMD's MI300A, world's first data center APU, directly addresses multiple exascale challenges

AMD MI300A APU

- 4th Gen AMD Infinity Architecture: AMD CDNA™ 3 and EPYC™ CPU "Zen 4" together
  - CPU and GPU cores share a unified on-package pool of memory
- Groundbreaking 3D modular (chiplet) packaging
  - CPU | GPU | Cache | HBM
  - 24 Zen4 CPU cores, 146B transistors, 128GB HBM3
- Designed for leadership memory bandwidth and application latency
- APU architecture designed for power savings
  - compared to discrete implementation

MI300A offers 3D CPU+GPU integration for next-level efficiency

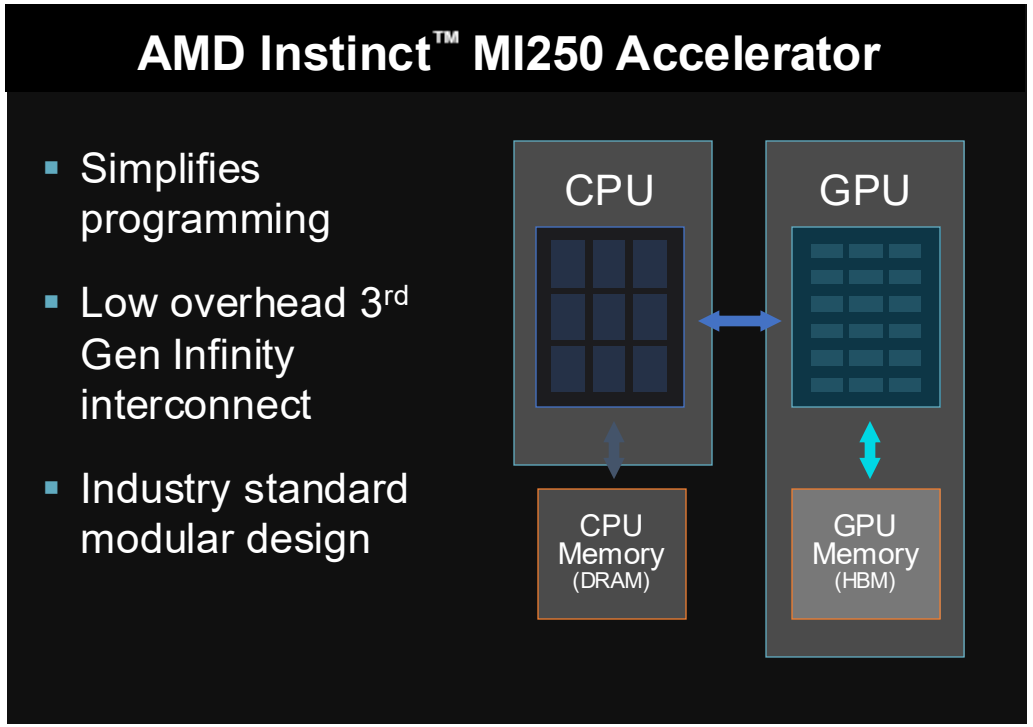

(Frontier)

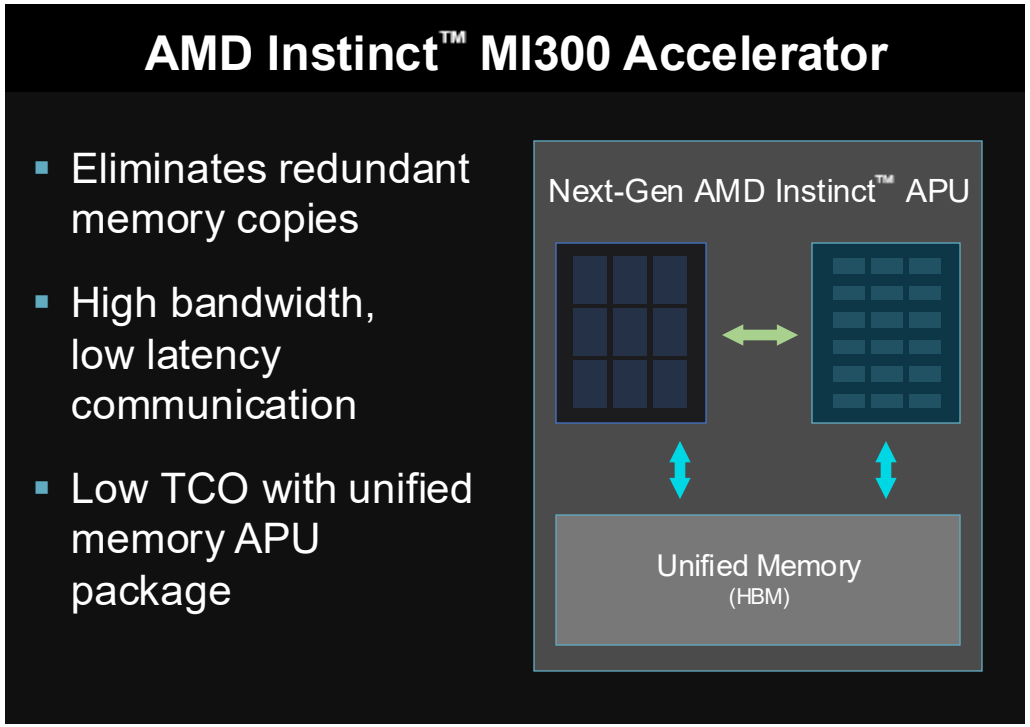

(El Capitan)

The MI300A design was informed by the needs of TriLab multiphysics applications

El Capitan is the product of several co-designed innovations

- Major infrastructure upgrade at LLNL to position the lab for the future
- AMD's MI300A, world's first data center APU, directly addresses multiple challenges
- **El Capitan includes an innovative near node local storage solution**
- Is using LLNL developed software stack enhanced with HPE software

The El Capitan Rabbit Architecture:

A novel, flexible near-node local storage

- One Rabbit blade per compute chassis
- Each Rabbit houses 18 SSDs (16+2 spare) with PCIe connections to every compute node via PCIe switch (*Rabbit-S*)
  - 2TB capacity per compute node
- Each Rabbit contains 1 AMD EPYC CPU (*Rabbit-P*)
- Rabbit blades are connected to the high-speed interconnect

Compute nodes on El Capitan can access storage in different ways with different performance characteristics

- Permanent Lustre
  - System-wide parallel file system that operates in the same way that parallel file systems do today

---

- Ephemeral Lustre on Rabbit
  - Shared Lustre file system instance on Rabbit storage private to a job
  - Duration of file system is the life of a single job
  - Can span multiple Rabbits
- Persistent Lustre on Rabbit
  - Shared Lustre file system instance on Rabbit storage private to a group, shared across jobs
  - Duration of file system is independent of a job lifetime; duration dependent on LC policies
- Node-local storage on Rabbit
  - Every compute node in a job has a node-local file system with fast PCIe connection (similar to Sierra); storage is private to a job
  - Duration of file systems is the life of a single job

El Capitan is the product of several co-designed innovations

- Major infrastructure upgrade at LLNL to position the lab for the future
- AMD's MI300A, world's first data center APU, directly addresses multiple challenges
- El Capitan includes an innovative near node local storage solution
- **Is using LLNL developed software stack enhanced with HPE software**

El Capitan is the first ATS to use TOSS and TCE in production

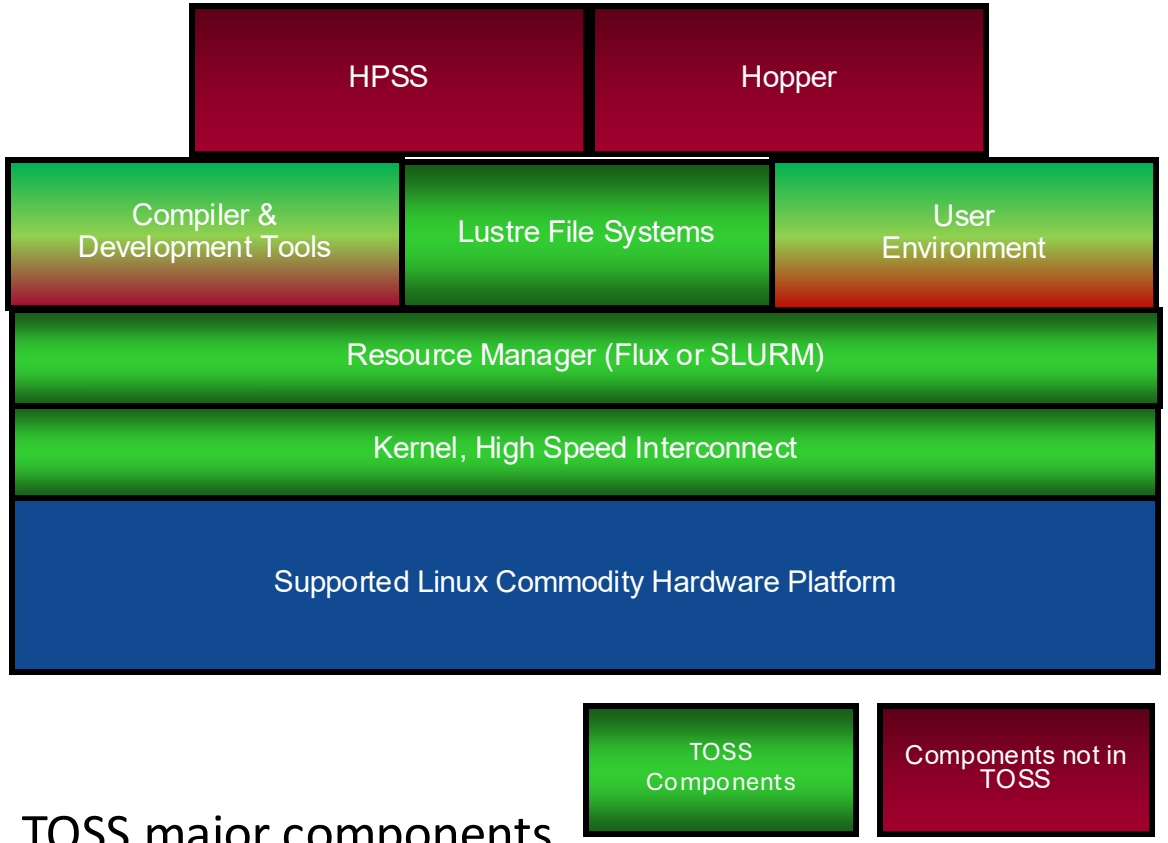

- TOSS major components
  - The OS – LLNL's Linux distribution based on RHEL
  - Resource Manager (SLURM or Flux)
  - Lustre

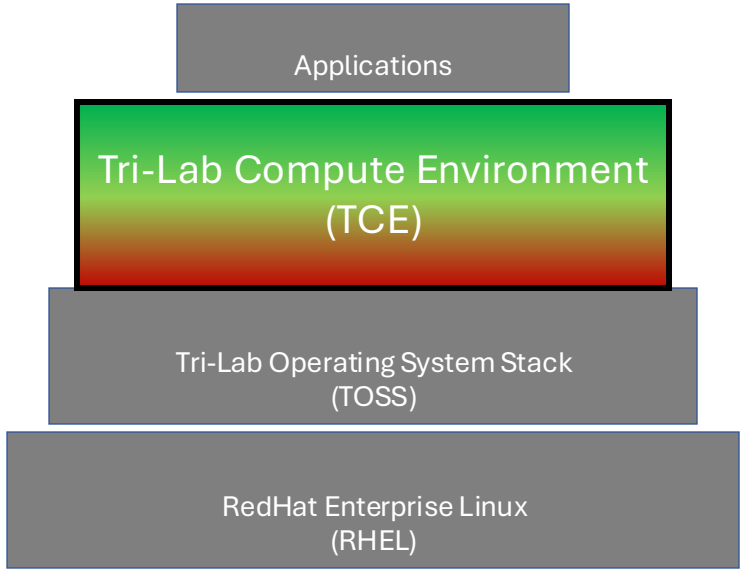

- The Tri-Lab Compute Environment (TCE) is an application development environment (DE)
  - Compilers (Intel, PGI, GNU, …)
  - MPI (MVAPICH, OpenMPI, …)
  - Debuggers (TotalView, Allinea)
  - Performance Tools

El Capitan is the first ATS to use Flux as its scheduler

Flux offers a new scheduling model to meet modern-day workflow challenges.

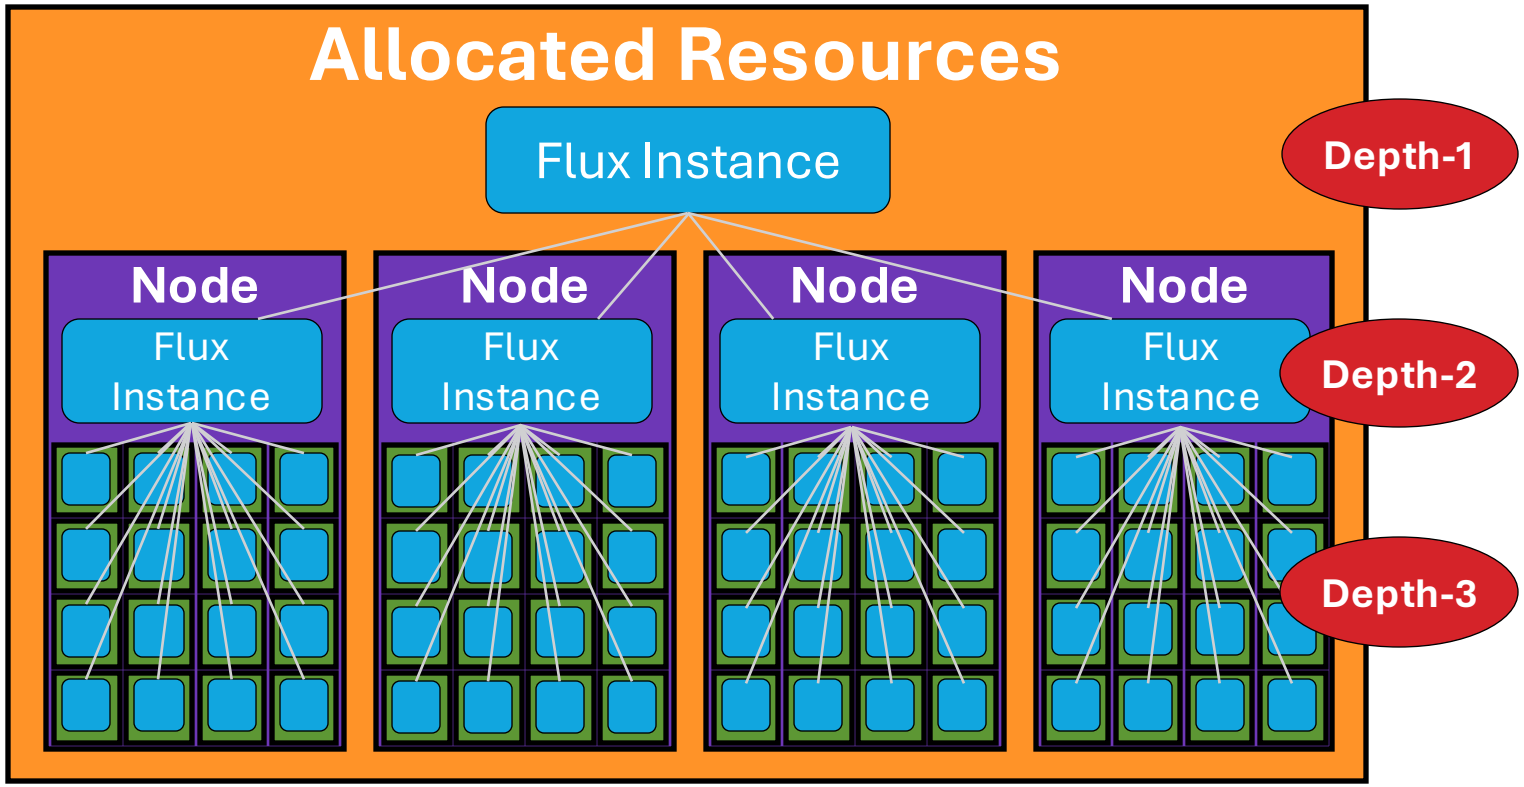

- Full workflow-enablement support
  - Hierarchical resource subdivision
    - Workflows that program to Flux as their top level scheduler will be able to use the same interface to subdivide a single node
  - Sub-resource manager per subdivision with specialization
    - Can separately allocate GPUs and CPUs to jobs
    - Can be extended to to oversubscribe resources
    - Can use mpibind to associate CPUs with GPUs
- Rich, well-defined interfaces
  - Facilitate communications and coordination among tasks within a workflow

How did we prepare for the world's fastest computer?

Since it's inception,

LLNL has deployed **world-leading** supercomputers

1953	1960s	1970s	1980s	1990s	2000s	2010s	2018	Today
 UNIVAC 1	 CDC 3600 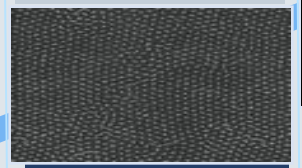 Pioneering simulations of particle tracking	 CDC 7600 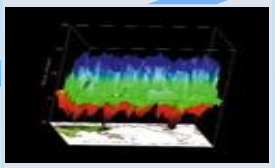 Ozone mixing models	 CRAY 1  Dynamics in three dimensions	 ASCI Blue-Pacific 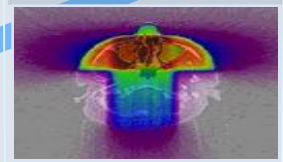 Helping medical community plan radiation treatment 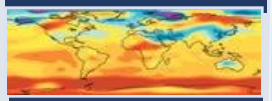 Global climate modeling	 Blue Gene 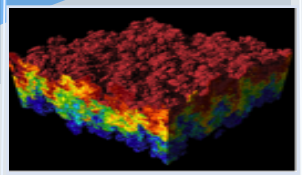 Breakthrough visualizations of mixing fluids  Unprecedented dislocation dynamics simulations	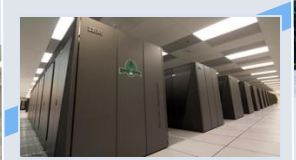 Sequoia 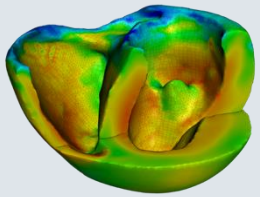 World's most detailed simulation of the human heart in action	 Sierra 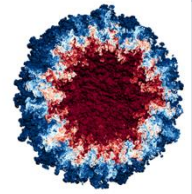 Highly resolved simulation of Rayleigh-Taylor instability 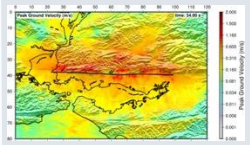 Fast, highly resolved earthquake simulations	 El Capitan

What are we and will we learn?

Each generation of our supercomputers has enabled new discoveries in a plethora of scientific fields

Embracing GPU-based computing with Sierra paid off with big performance increases

Ares, RT Mixing
13x speedup

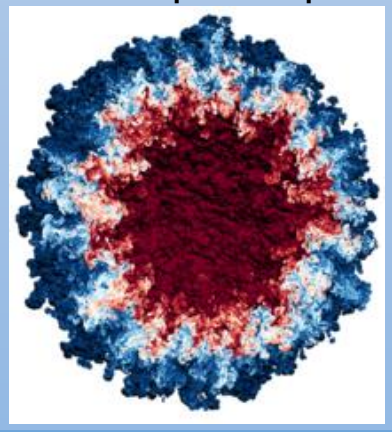

Ardra, Reactor Safety
16x speedup

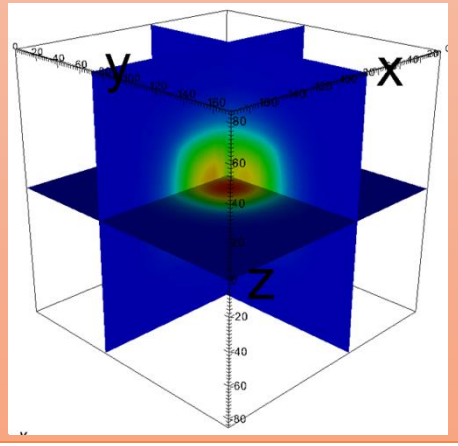

ALE3D, Shaped Charge
8x speedup

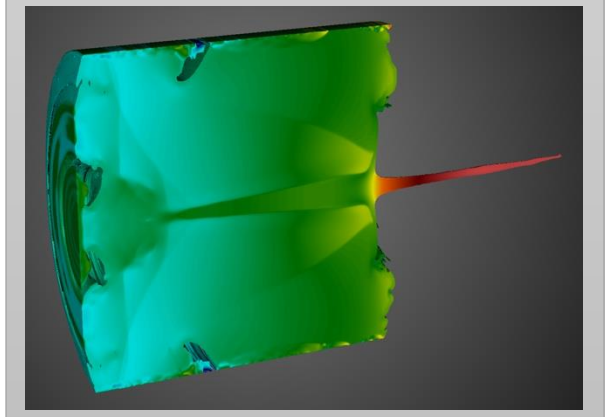

Kull/Teton
Radiating Sphere
7x speedup

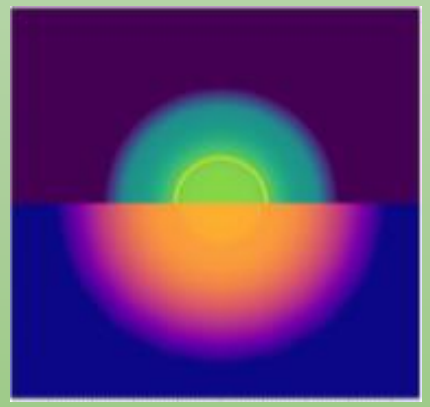

SW4, Hayward Fault, 28x speedup

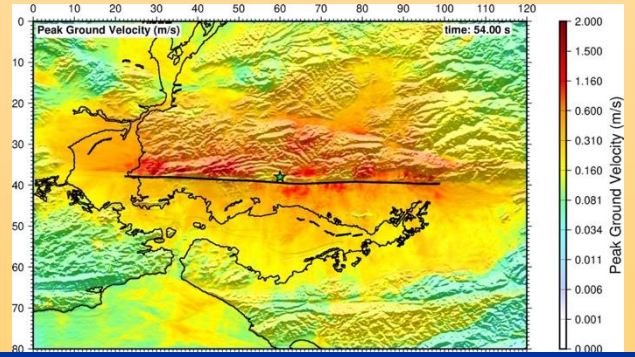

Lessons learned from Sierra and earlier machines helped us prepare for El Capitan

Strategic investments in RAJA portability abstractions for Sierra are now paying dividends

	Sierra (V100)	El Capitan (MI300A)	Speed up
Flop Rate (Linpack)	5.5 TF	39 TF	7.1x
Memory Bandwidth (Stream Triad)	850 GB/s	3600 GB/s	4.2x

Raja kernel performance: MI300A vs V100

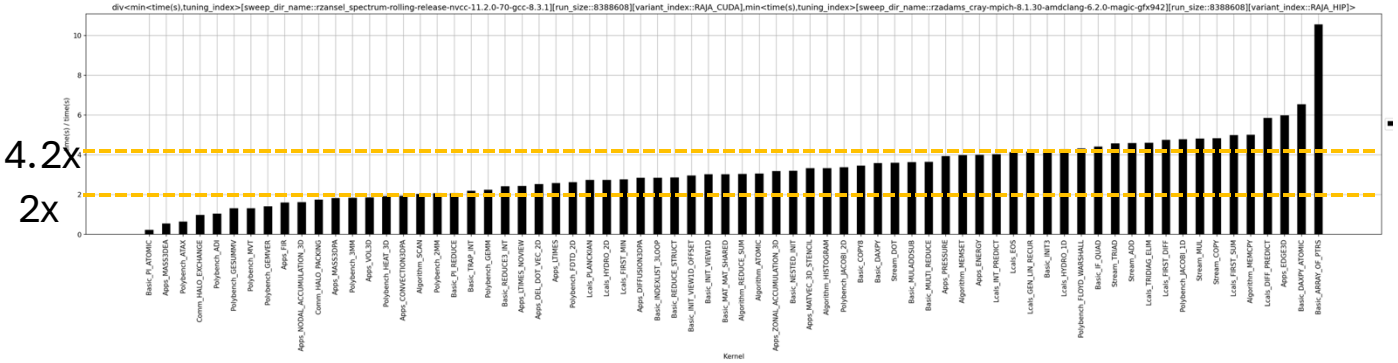

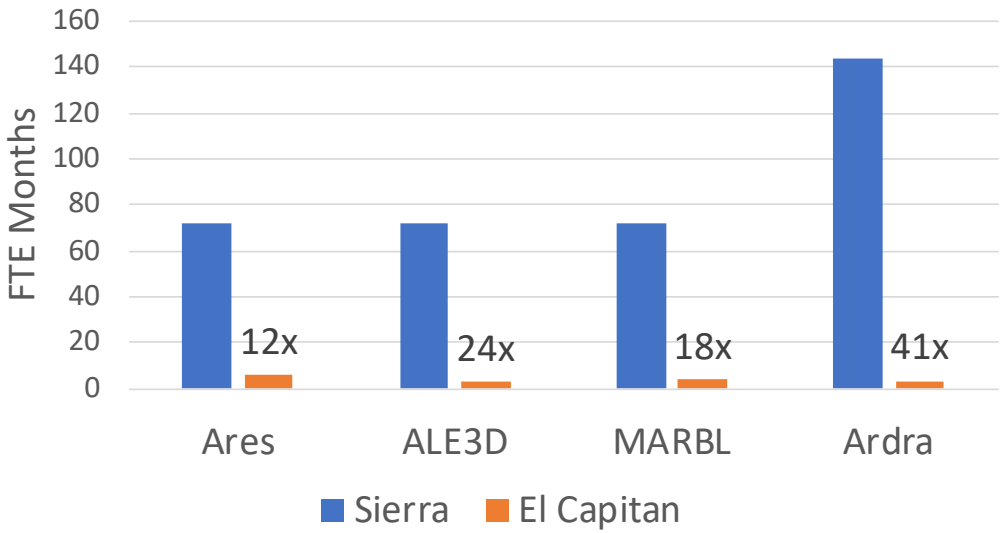

RAJA

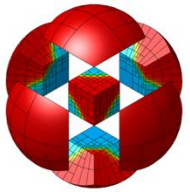

MFEM

bLT

Umpire

RAJA abstractions significantly reduced porting time to El Capitan without impacting performance. RAJA kernels are also a good target for performance optimization.

The El Cap Center of Excellence has been preparing Tri-Lab codes for El Capitan

- A close partnership between the DOE NNSA Laboratories, HPE and AMD
- Joint work plans, information sharing, and collaboration mechanisms
- Dedicated HPE and AMD staff **working alongside lab code teams**
- Vendors providing NDA information and **early access to hardware** and software

Forming a Center of Excellence has become a recognized best practice for large DOE system procurements

Our newest code, MARBL, with advanced algorithms, is critical to simulating fusion ignition

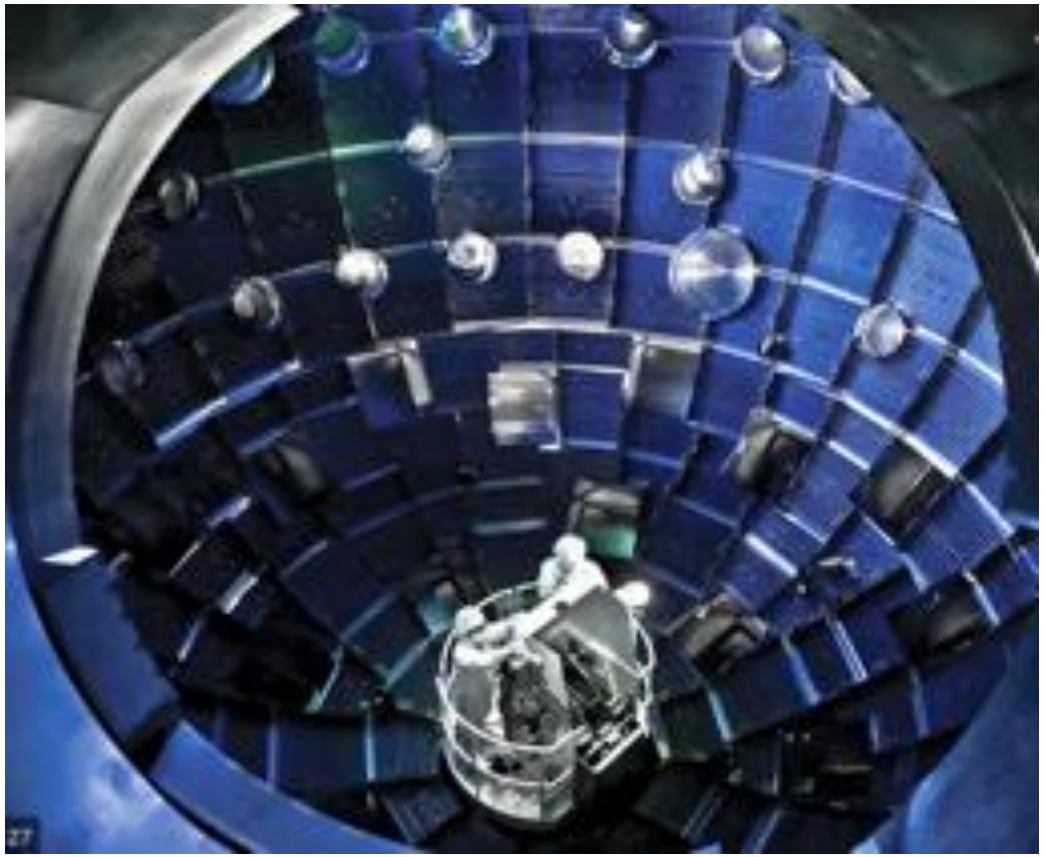

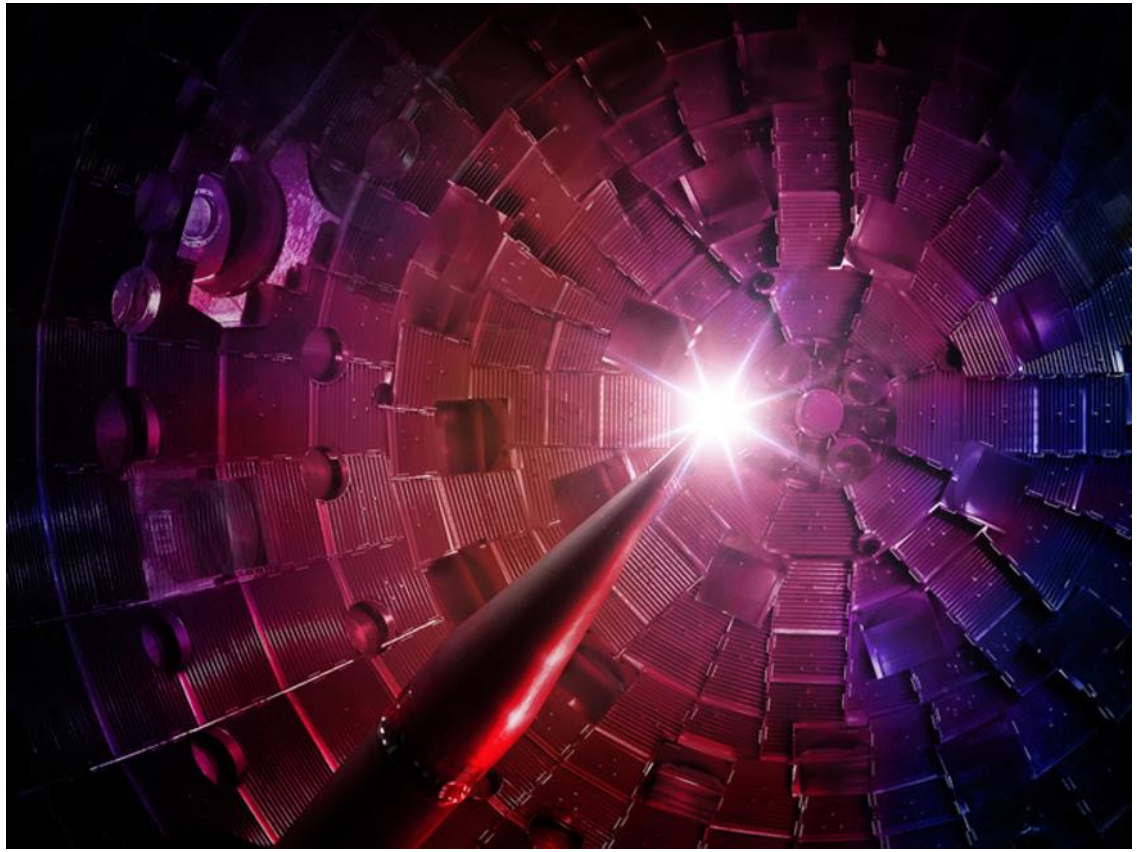

MARBL has been preparing for El Capitan for many years in collaboration with the COE

Hotspot (Pre-imploded ICF capsule benchmark) – High order ALE Rad Hydro with TN Burn

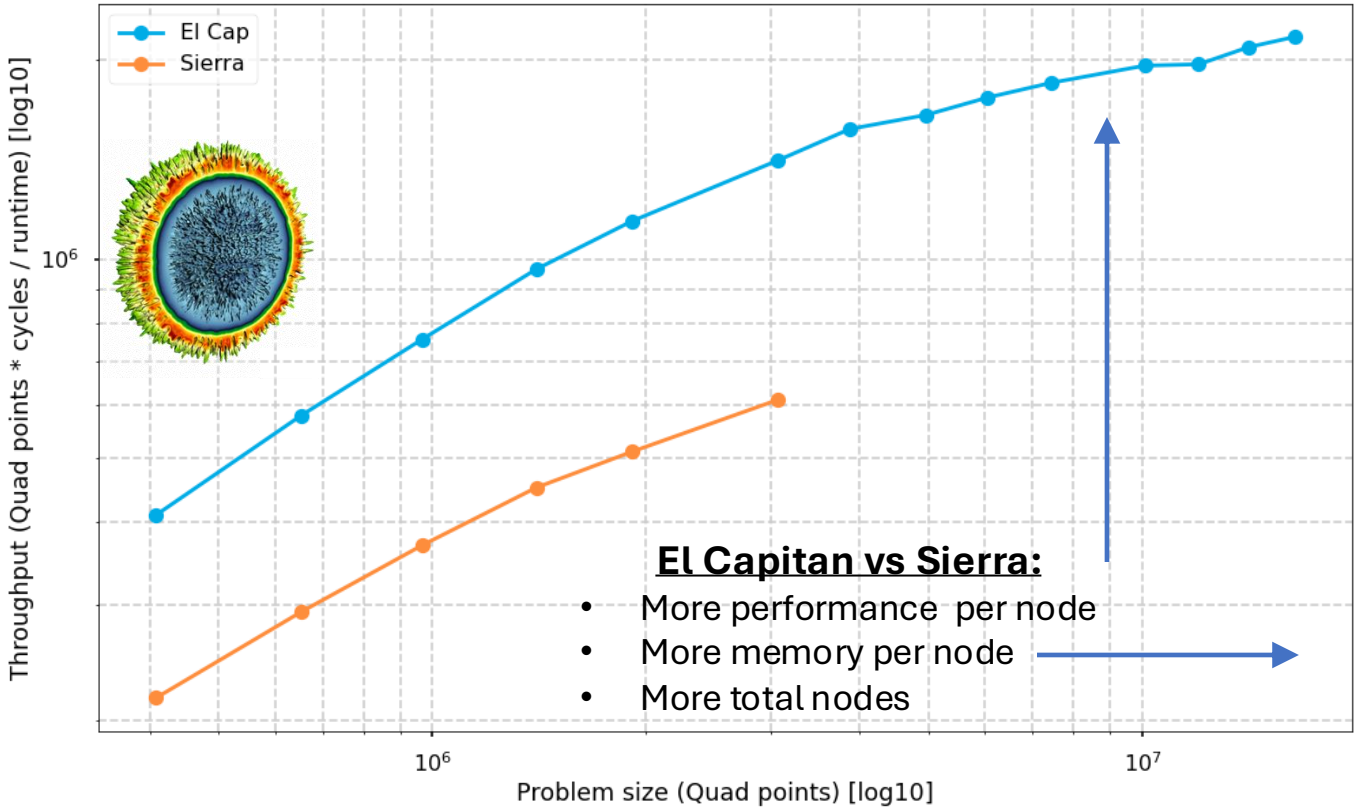

MARBL	Sierra (s)	El Cap (s)	Single Node Speedup
Total Physics	259.07	110.15	2.35X
Radiation Diffusion	228.13	96.41	2.37X
Thermal Conduction	19.34	8.89	2.18X
Hydro	8.42	3.05	2.76X

RAJA Umpire mfem Caliper

Slide courtesy R. Rieben

El Capitan is the right machine at the right time to study high fidelity mod-sim, coupled with ML/AI workflows

→ Project ICECap is leveraging exascale computing and AI to automate millions of simulations and discover new robust ignition designs

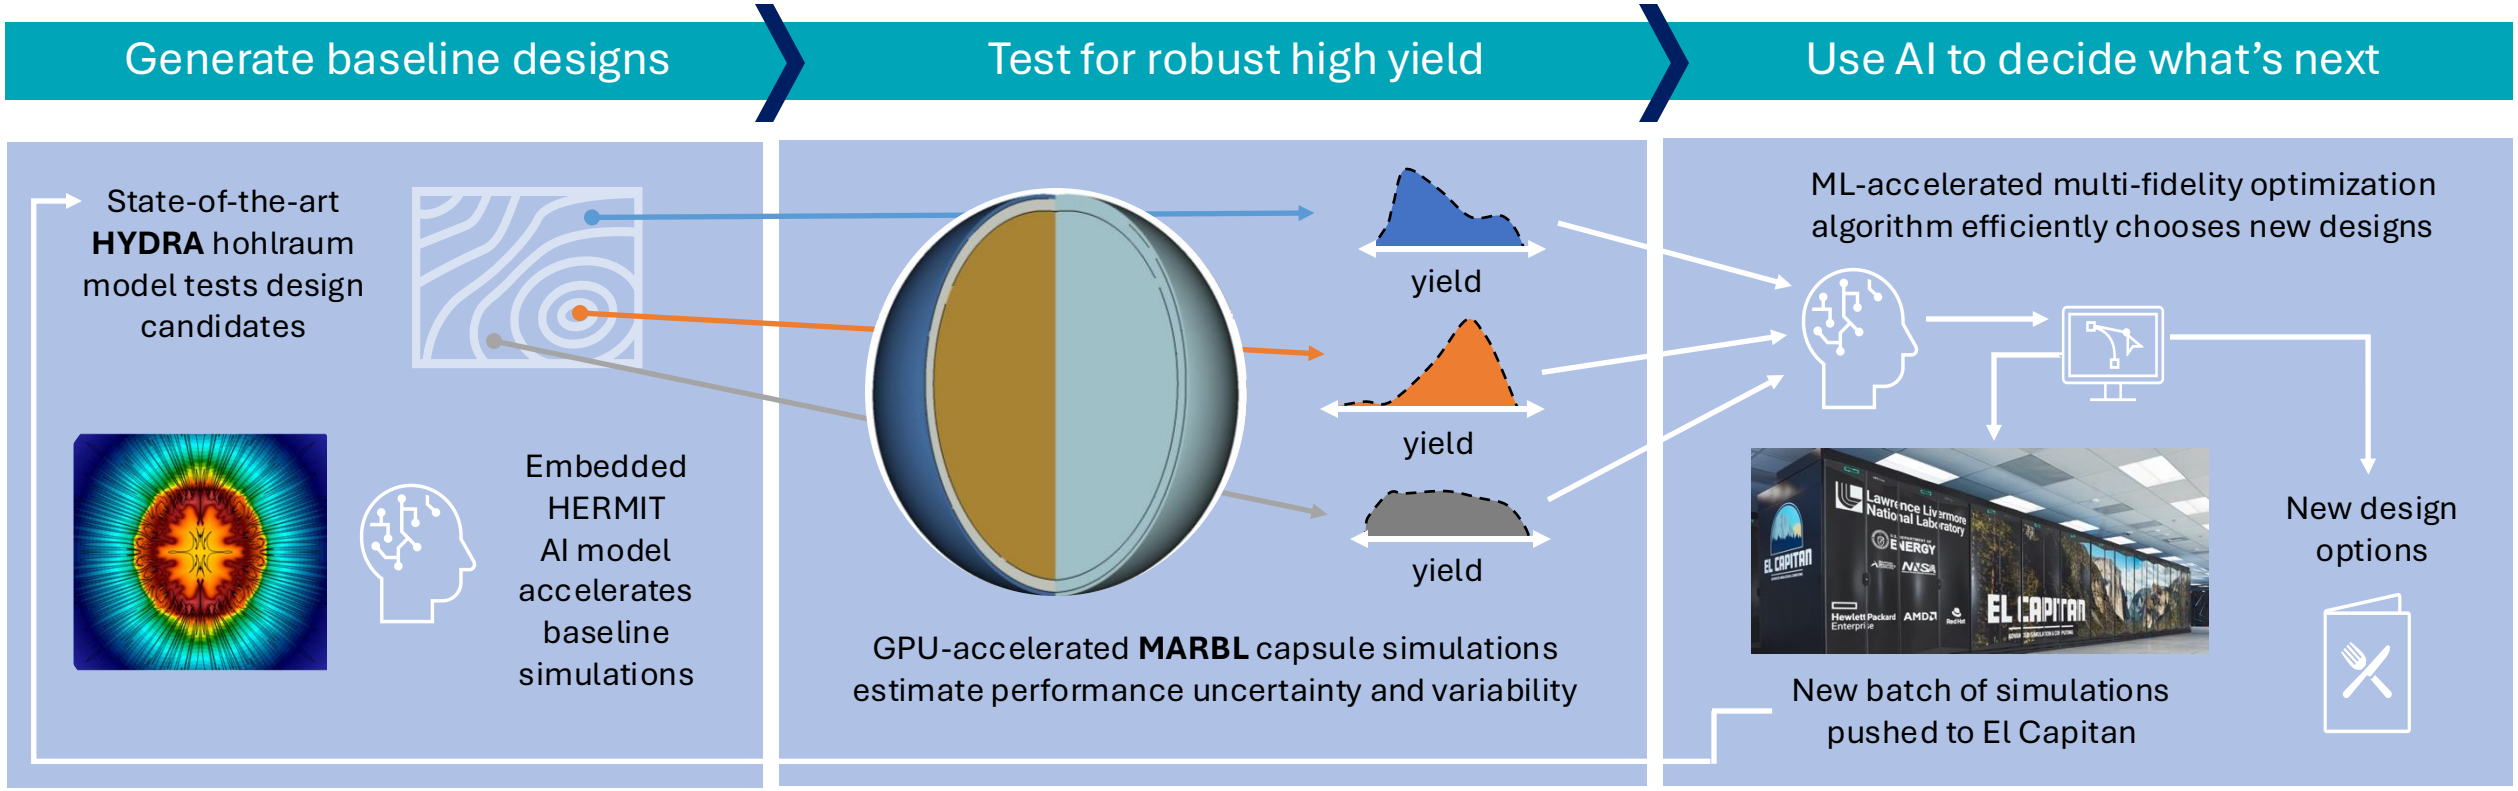

Slide courtesy R. Rieben

The GUIDE Program is Accelerating the Design of Novel Antibodies Using Supercomputing

HPC-Enabled AI Models and Simulations Help GUIDE to rapidly search the massive antibody design space of 10^{40} possible designs!

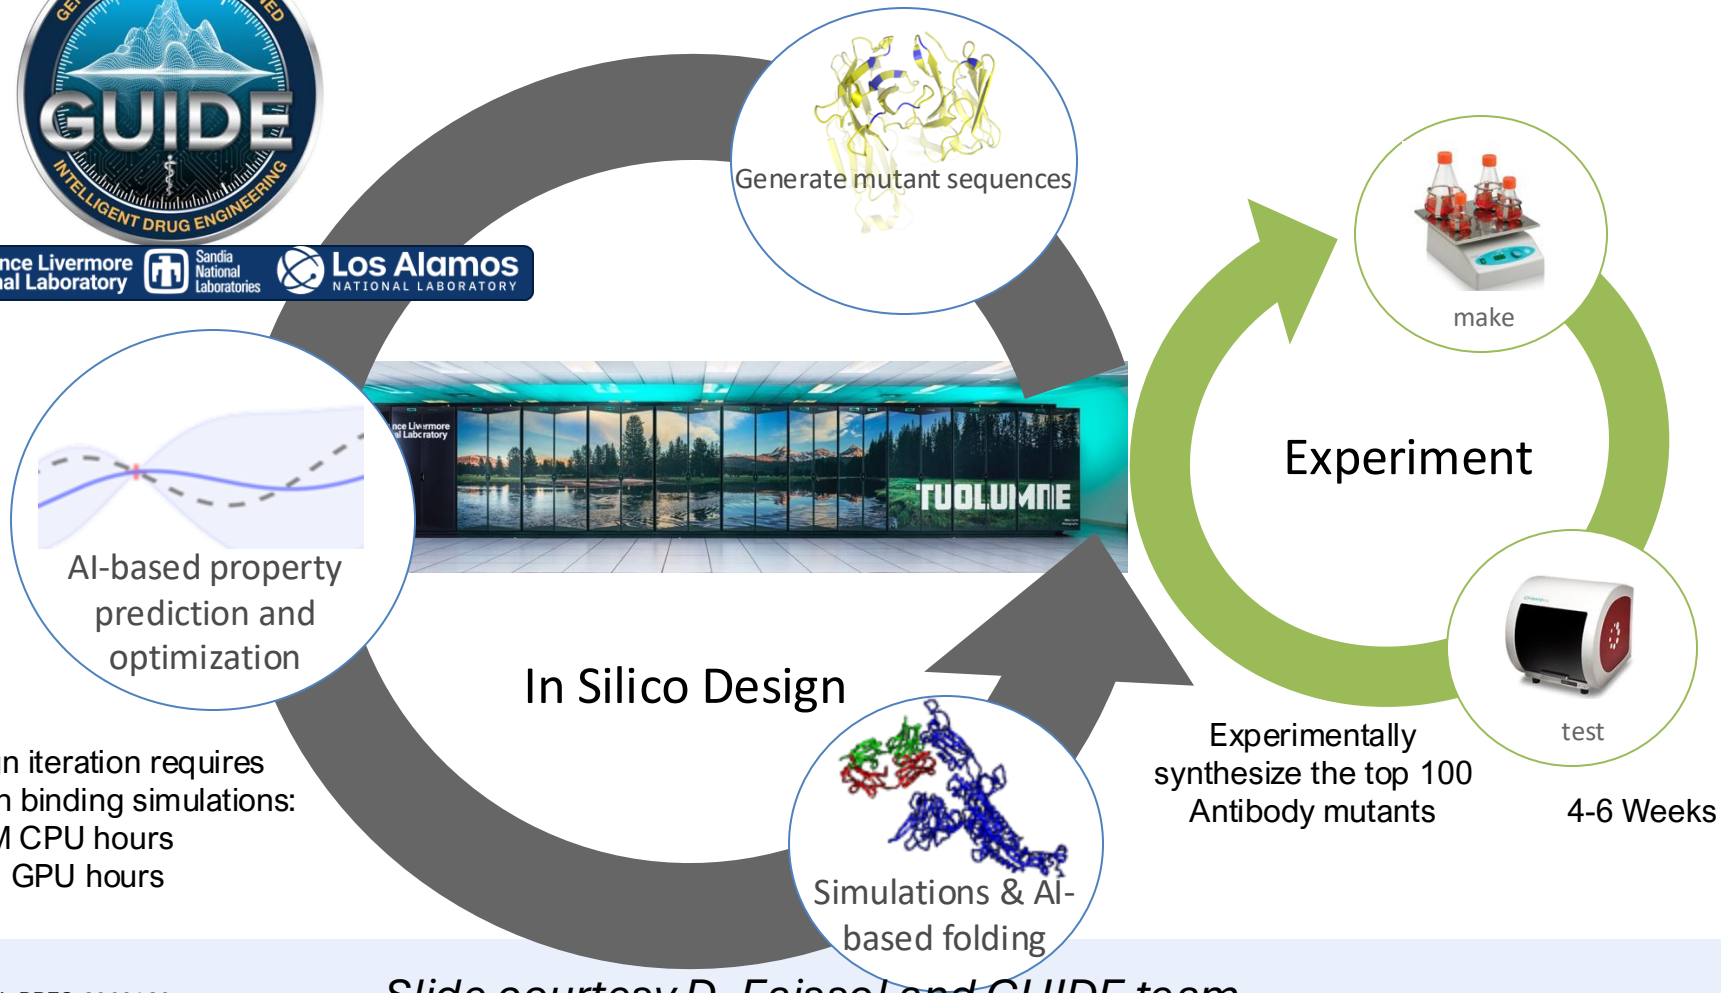

Slide courtesy D. Faissol and GUIDE team

The GUIDE computational platform has demonstrated success in antibody design and optimization across multiple viral families

Neutralizing antibodies for SARS-CoV-1 that do not bind SARS-CoV-2

Mutate antibody sequence to maximize affinity to SARS-CoV-2

6 mutations on antibody enabled SARS-CoV-2 neutralization

SARS-CoV-1 antibody

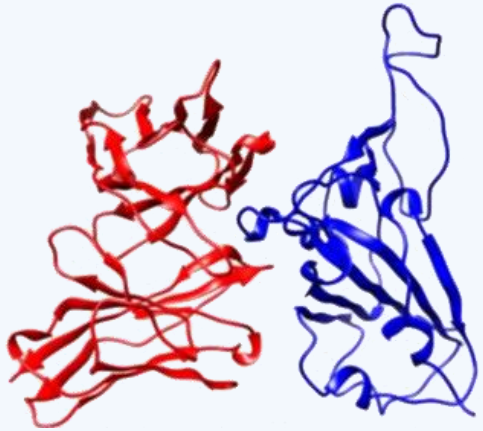

10^{40} possible antibody mutants!

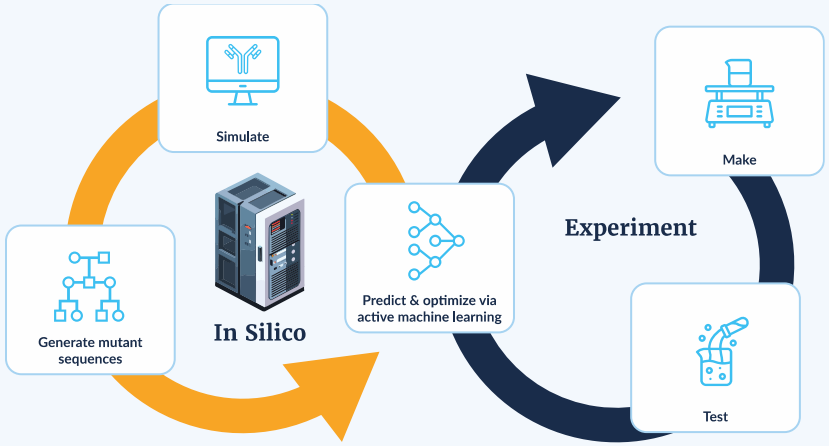

5 million binding simulations
20M CPU hours
3M GPU hours

100 Antibody mutants experimentally tested

Mutated Antibody

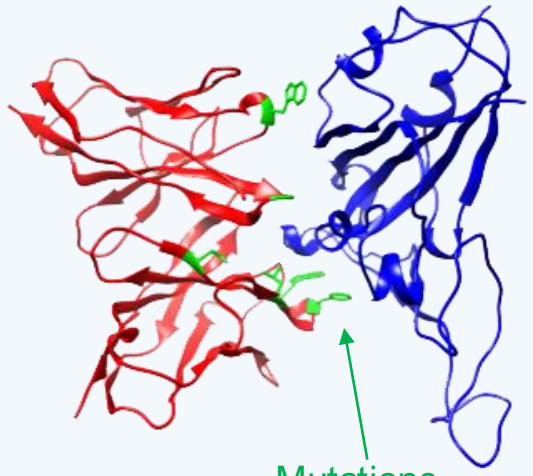

“Hero” Sierra calculations can now be run as ensemble members on El Capitan

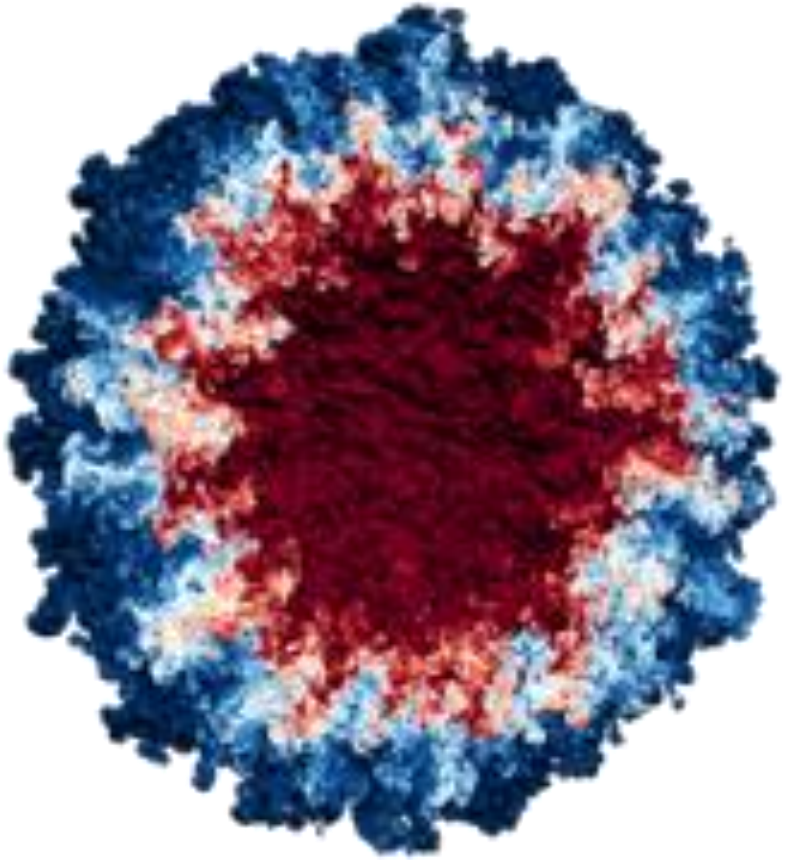

- A 2018 97.5B zone Convergent Rayleigh-Taylor calculation that we used to test out Sierra
- **This was the largest problem that could be run on Sierra**
- On El Capitan, we see significantly better performance per node and run at 2x the speed on half the nodes

Platform	Node Count	Time per Cycle (s)	Throughput Speedup
Sierra	4096	0.47	1
El Cap	2048	0.24	3.95

Information courtesy B. Ryujin, J. Burmark

El Capitan is enabling unprecedented scales

Multiple codes are scaling out to over 1 Trillion zones

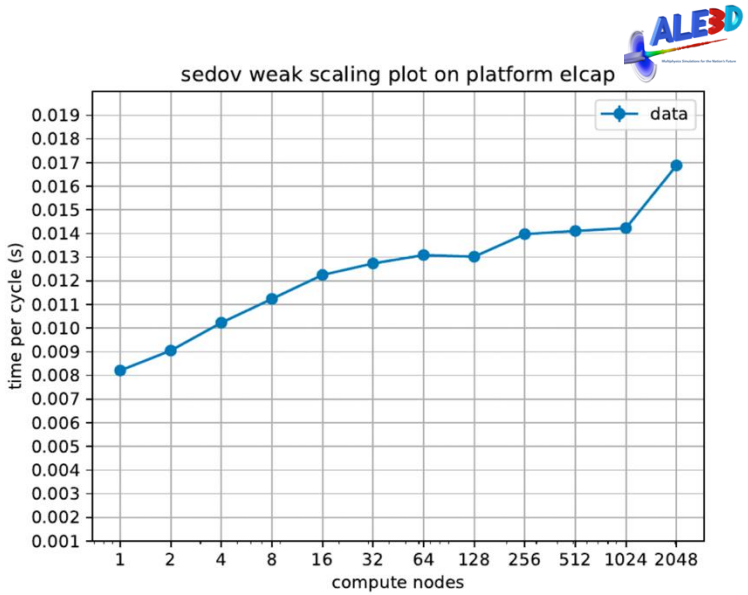

Hydro scaling in ALE3D with GPU aware MPI (on node GPU kernel reductions currently disabled due to correctness bug)

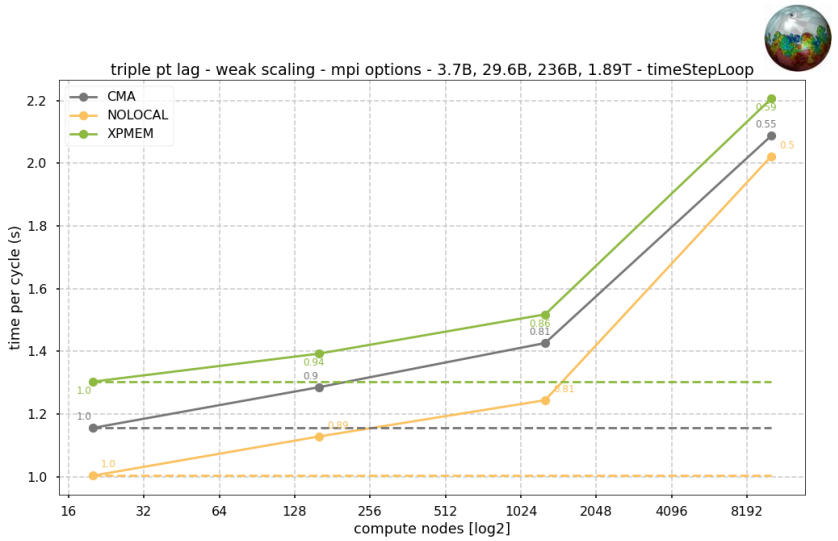

Hydro scaling in MARBL with various MPI options. 10k results are probably a known point-to-point performance issue

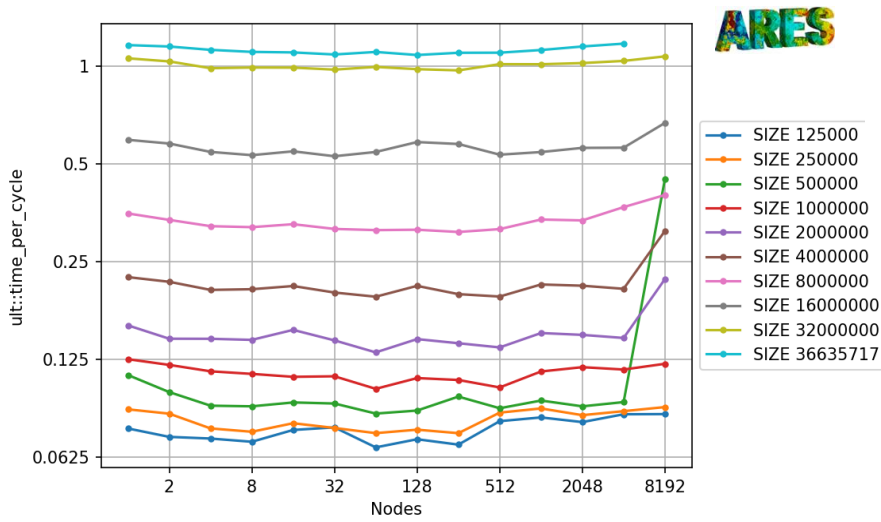

Hydro scaling in Ares with XPMEM looks good. 8k result outlier likely result of a bad node

Information courtesy B. Ryujin, J. Burmark, T. Stitt, P. Robinson

El Capitan is enabling unprecedented scales

Multiple codes are scaling out to over 1 Trillion zones

ARES

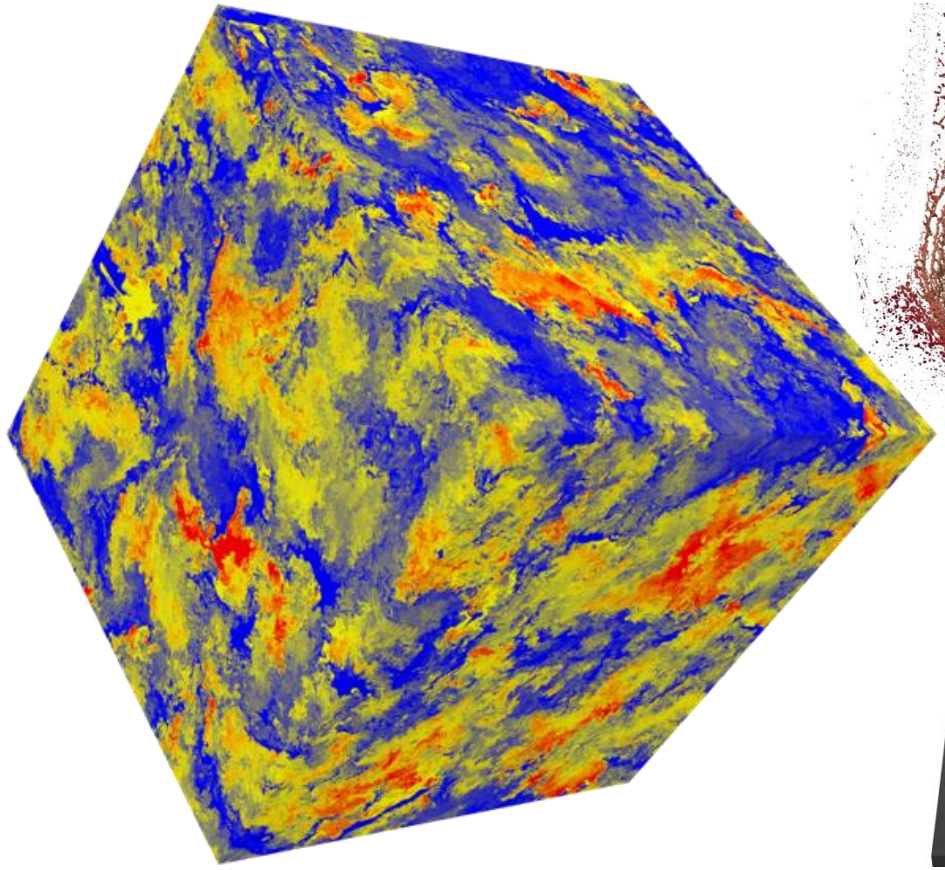

68.9B zones on 512 nodes
three component mixing

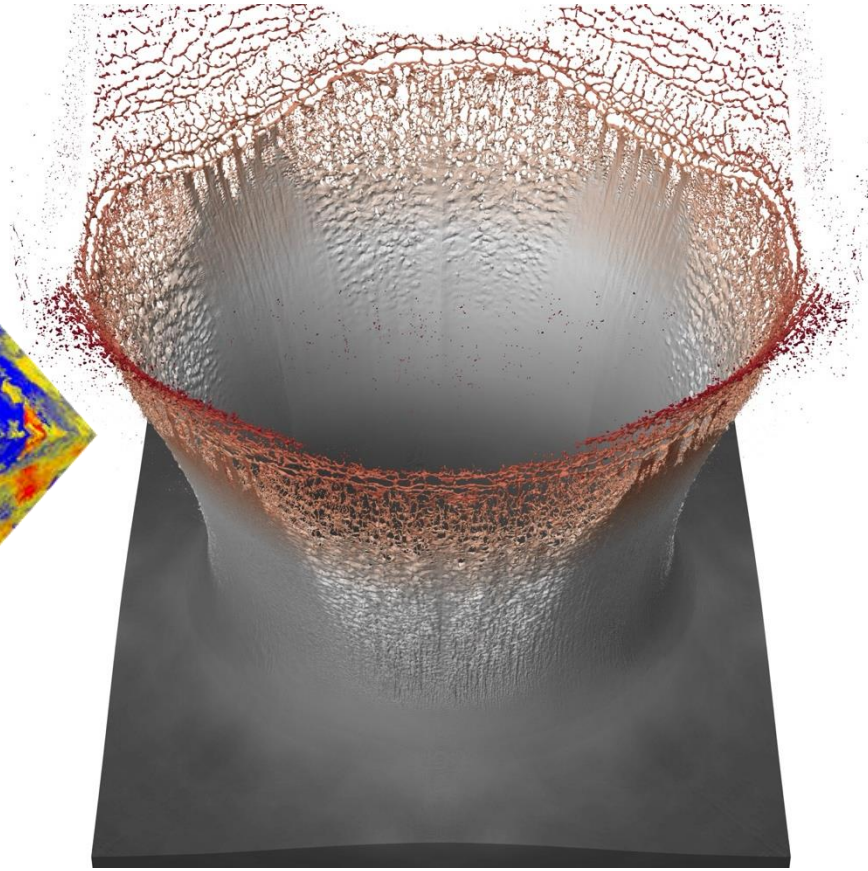

250B zones on 2048 nodes
hypersonic tin particle impact on tin surface

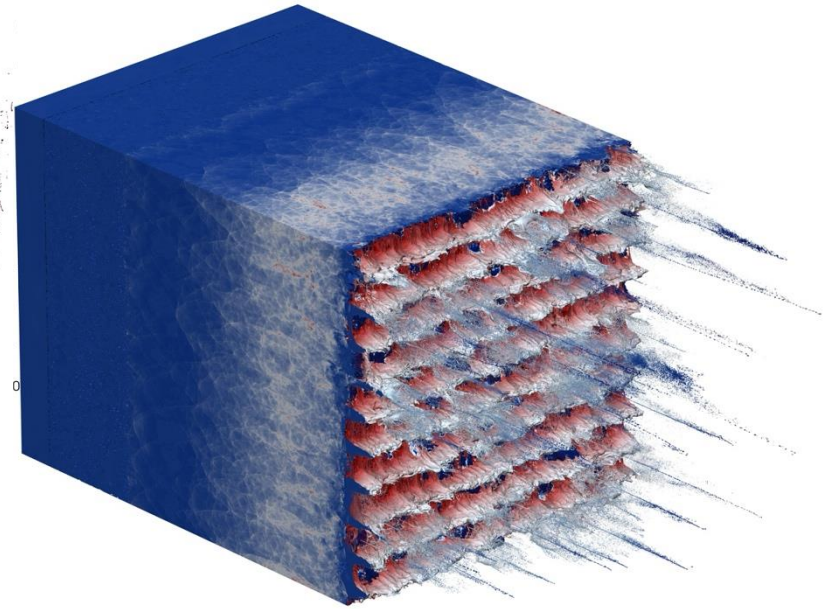

139B zones on 2048 nodes
shocked tin surface undergoing
shallow bubble collapse

El Capitan and Tuolumne are enabling new scientific discovery

Sequoia
20 petaFLOPS

Sierra
125 petaFLOPS

Tuolumne
288 petaFLOPS

El Capitan
2.790 exaFLOPS

National Security

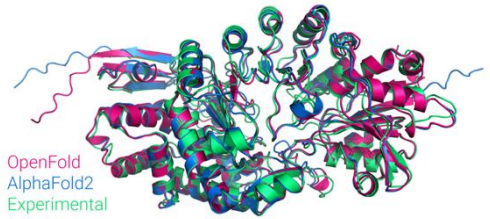

Artificial Intelligence

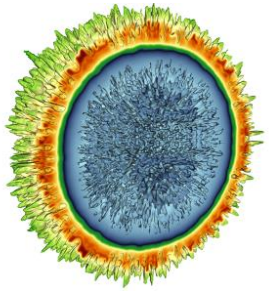

Fusion Research

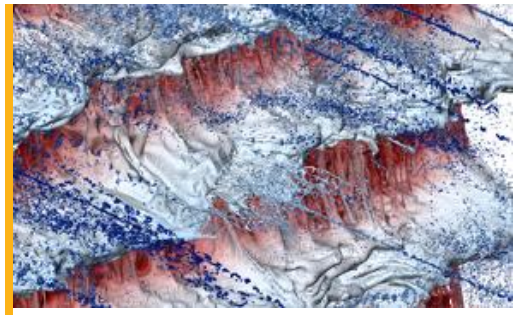

Materials Science

Questions?

Thanks to: T. Bailey, J. Burmark, B. de Supinski, D. Faissol, R. Pankajakshan, T. Quinn, D. Richards, R. Rieben, P. Robinson, B. Ryujin, T. Stitt
and many others in LLNL's WSC/CP, LC, and at AMD and HPE