OAK RIDGE National Laboratory

Beyond Frontier – Thoughts on the OLCF's Post-Exascale System

Scott Atchley

Chief Technology Officer, OLCF

SOS26

Cocoa Beach, FL

ORNL is managed by UT-Battelle LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Overview

- OLCF's Mission
- DOE's Integrated Research Infrastructure
- Changing Advanced Computing Ecosystem
- OLCF-6's RFP adapts to the new Ecosystem
- How the SOS Community Can Help

OLCF's Mission

OLCF was Established in Response to the 2004 American Supercomputing Leadership Act

Department of Energy High-End Computing Revitalization Act of 2004 (Public Law 108-423):

The Secretary of Energy, acting through the Office of Science, shall

- Establish and operate Leadership Systems Facilities
- Provide access [to Leadership Systems Facilities] on a competitive, merit-reviewed basis **to researchers in U.S. industry, institutions of higher education, national laboratories and other Federal agencies.**

Amended: American Super Computing Leadership Act of 2017 (P.L. 115-246]

Executive Order -- Creating a National Strategic Computing Initiative, July 29, 2015:

Mandated that DOE be the lead agency to develop a capable exascale program

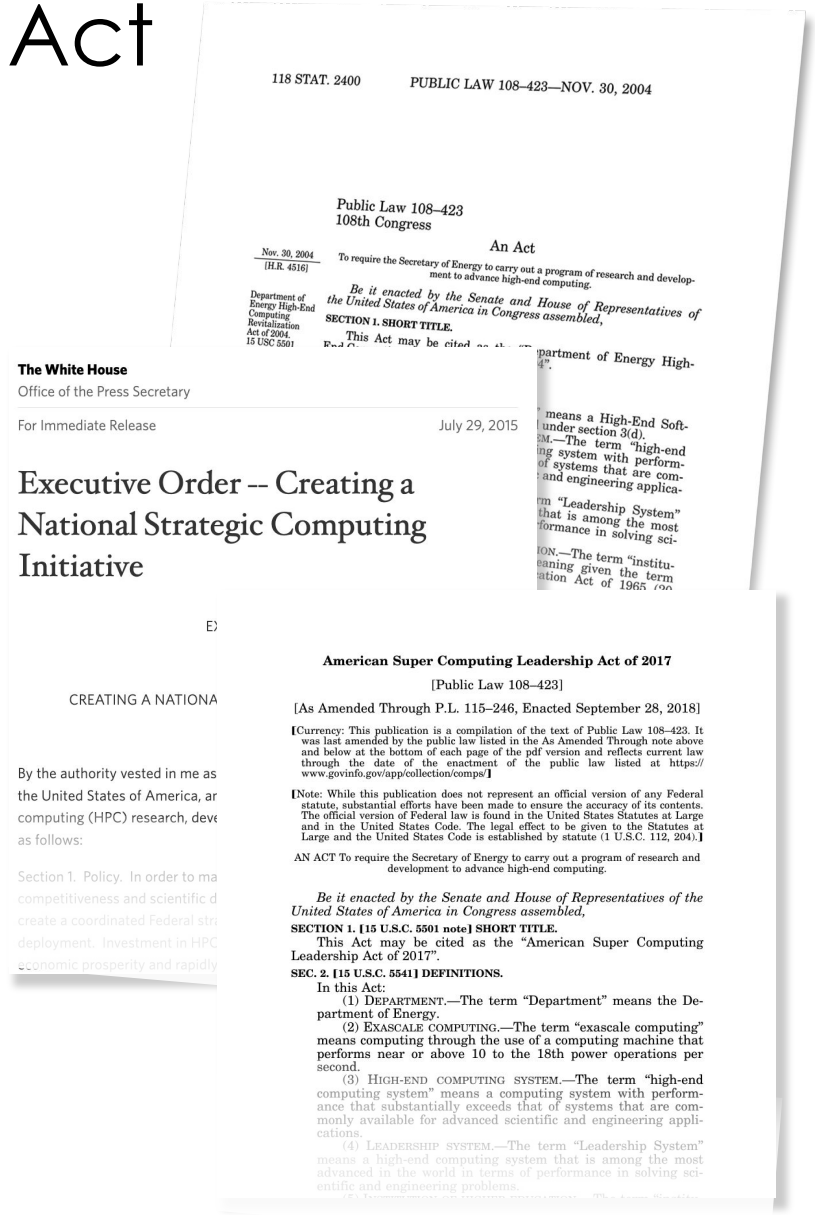
118 STAT. 2400 PUBLIC LAW 108–423—NOV. 30, 2004

Public Law 108–423
108th Congress

An Act

Nov. 30, 2004 [H.R. 4516]

To require the Secretary of Energy to carry out a program of research and development to advance high-end computing.

Department of Energy High-End Computing Revitalization Act of 2004. 15 USC 5501

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. SHORT TITLE.

This Act may be cited as the "Department of Energy High-End Computing Revitalization Act of 2004".

The White House
Office of the Press Secretary

For Immediate Release July 29, 2015

Executive Order -- Creating a National Strategic Computing Initiative

CREATING A NATIONA

By the authority vested in me as the United States of America, ar computing (HPC) research, deve as follows:

Section 1. Policy. In order to ma competitiveness and scientific d create a coordinated Federal str deployment. Investment in HPC economic prosperity and rapidly

American Super Computing Leadership Act of 2017

[Public Law 108–423]

[As Amended Through P.L. 115–246, Enacted September 28, 2018]

[Currency: This publication is a compilation of the text of Public Law 108–423. It was last amended by the public law listed in the As Amended Through note above and below at the bottom of each page of the pdf version and reflects current law through the date of the enactment of the public law listed at https://www.govinfo.gov/app/collection/comps/]

[Note: While this publication does not represent an official version of any Federal statute, substantial efforts have been made to ensure the accuracy of its contents. The official version of Federal law is found in the United States Statutes at Large and in the United States Code. The legal effect to be given to the Statutes at Large and the United States Code is established by statute (1 U.S.C. 112, 204).]

AN ACT To require the Secretary of Energy to carry out a program of research and development to advance high-end computing.

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

SECTION 1. [15 U.S.C. 5501 note] SHORT TITLE.

This Act may be cited as the "American Super Computing Leadership Act of 2017".

SEC. 2. [15 U.S.C. 5541] DEFINITIONS.

In this Act:

(1) DEPARTMENT.—The term "Department" means the Department of Energy.

(2) EXASCALE COMPUTING.—The term "exascale computing" means computing through the use of a computing machine that performs near or above 10 to the 18th power operations per second.

(3) HIGH-END COMPUTING SYSTEM.—The term "high-end computing system" means a computing system with performance that substantially exceeds that of systems that are commonly available for advanced scientific and engineering applications.

(4) LEADERSHIP SYSTEM.—The term "Leadership System" means a high-end computing system that is among the most advanced in the world in terms of performance in solving scientific and engineering problems.

OLCF by the Numbers

Mission: Provide Leadership Computing **to researchers in U.S. industry, institutions of higher education, national laboratories and other Federal agencies.**

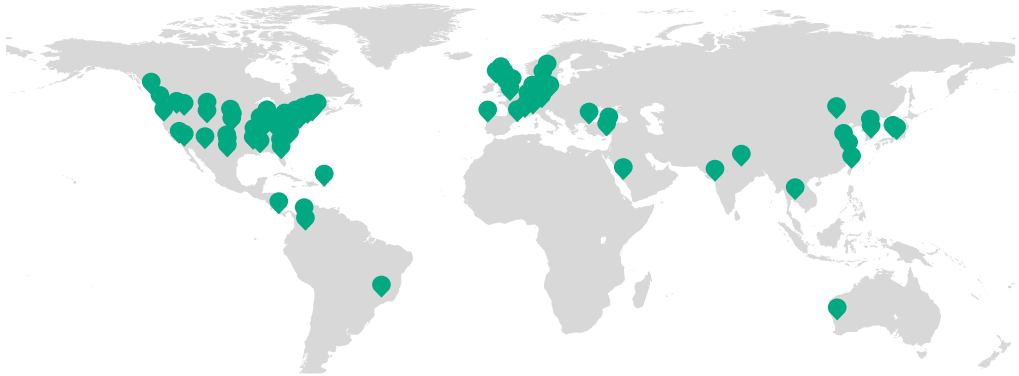

50% Universities

40% Federal Agencies

10% Industry

Users access OLCF through 3 competitive programs (Hours)

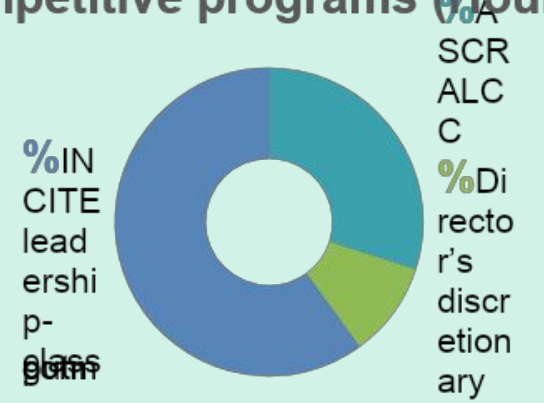

1,744 annual users
262 institutions

5,800 confirmed peer-reviewed publications since 2012
12 Gordon Bell Prizes

Effective & numerous channels for user engagement

High annual survey scores for user engagement

DOE Office of Science Programs

Basic Energy Research

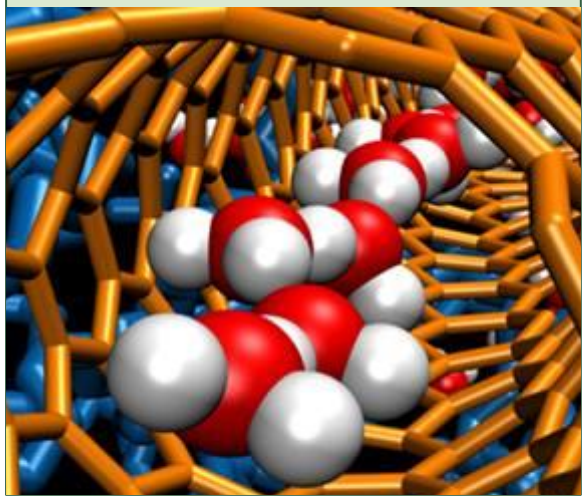

Biological & Environmental

Fusion Energy Sciences

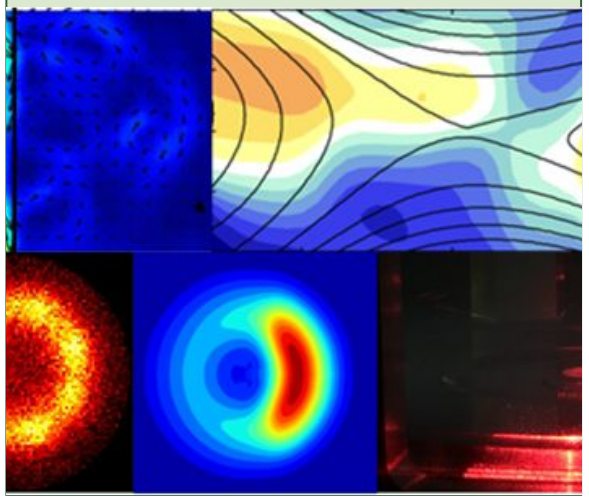

High Energy Physics

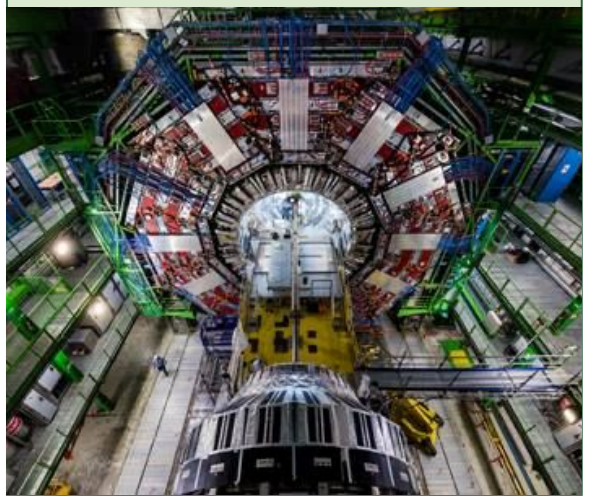

Nuclear Physics

Accelerator R&D and Prod

Isotope R&D and Prod

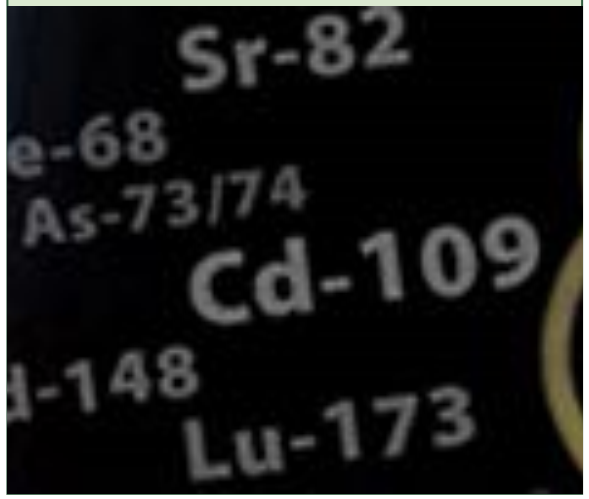

Strong Partnerships with other Federal Agencies

NOAA

- NOAA paid funds in to fill out the GPUs in Titan
- Collaborated on container technologies that we ultimately deployed on Summit, Frontier, and the NOAA systems

VA & NIH

- Main drivers for building the CITADEL capability
- End-to end AI camp aigns from discovery to population level deployment
- Driving Federated Learning capabilities

Air Force

- Experience learned on deploying AFW systems helped the OCLF to accelerate Frontier timeline
- AF investment in NVIDIA GPUs provides OLCF with access to current and emerging GPUs

NASA

- NASA has received significant computing time through INCITE projects
- NASA applications have benefited from OLCF partnerships
- The ability to host ITAR data allows NASA to bring new problems

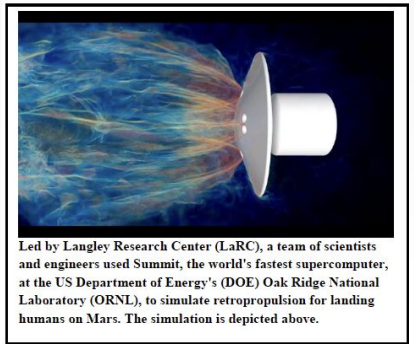

Led by Langley Research Center (LaRC), a team of scientists and engineers used Summit, the world's fastest supercomputer, at the US Department of Energy's (DOE) Oak Ridge National Laboratory (ORNL), to simulate retropropulsion for landing humans on Mars. The simulation is depicted above.

OLCF Industrial Partnerships Program Innovates

- The OLCF supports small to large Industrial partners
- Supported **175 unique industrial projects** over the last 5 years with notable outcomes
- OLCF **pioneered the Industry Partnership User Agreement**
- Industry was strong driver for the development of **OLCF's secure computing infrastructure known as Citadel**
- Strong engagement with the **ECP Industry Agency Council (IAC)** and plans are in the works to continue those engagements into the future

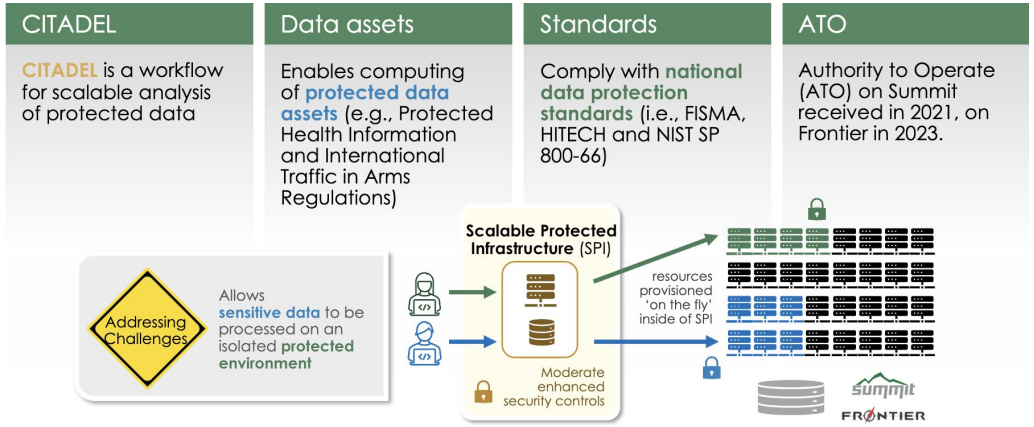

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Nuclear Energy

Leadership computing enables better engineered, safer vessels that can accelerate next-gen reactor deployment, including small modular reactors (SMRs)

Titan

Detailed snapshot of how neutrons move and interact (neutronics) and fuel at reactor startup

Summit

Multiple detailed snapshots of neutronics for **multiple points** in the reactor lifetime

Frontier

High-fidelity modeling of a whole reactor over **its** lifetime, **coupled to** thermal hydraulics

OLCF-6

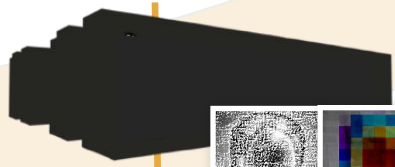

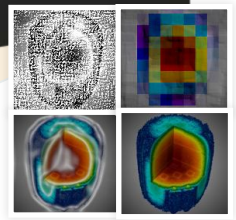

Multiple high-fidelity whole-reactor models **to** accelerate the design and regulatory cycle

OLCF-7

Optimized design at high fidelity for deployment

ML-based multi-physics emulators to render the geometry design space tractable

Climate Change

Leadership computing creates a high-resolution climate laboratory and full-scale digital twin of the Earth to improve forecasts and protect national infrastructure

OLCF-7

Numerical laboratories capable of testin mitigations

OLCF-6

Frontier

Summit

Titan

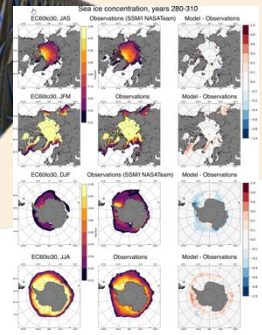

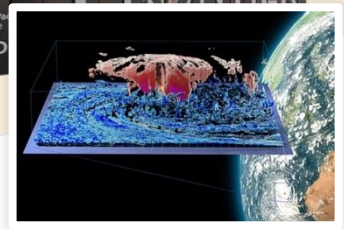

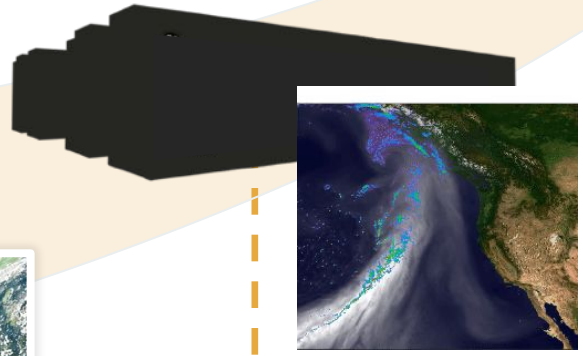

High-resolution, coupled climate models

Higher resolution, including physics of clouds and sea ice; **near-decadal** timescales

Cloud-resolving, 40-year climate studies at rates exceeding **one** simulated year per day

Virtual climate laboratory to explore multiple impacts and effects (plants, clouds, urban heat models, etc.)

ML emulators reduce the number of climate model runs required for UQ from a prohibitive $O(10^6)$, with standard Markov chain Monte Carlo methods, to a manageable $O(10^3)$

Molecular Biology

Leadership computing enables molecular-level modeling, producing deeper understanding of fundamental biological phenomena at unprecedented scales (time and space), leading to better therapies

OLCF-7

Systematic studies of dynamics leading to better antibiotics

OLCF-6

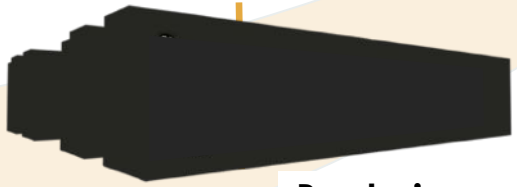

Frontier

Summit

Titan

Bacterium
10^{11} atoms

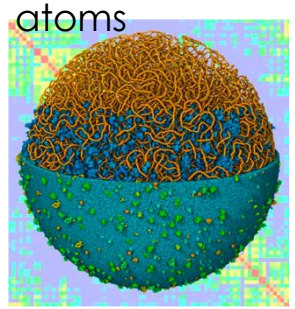

COVID-19
10^{9} atoms

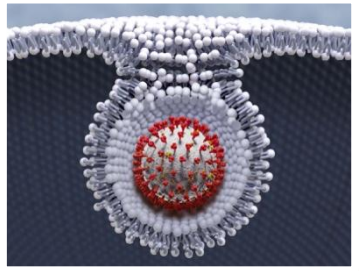

The **nuclear pore complex**
10^{8} Atoms

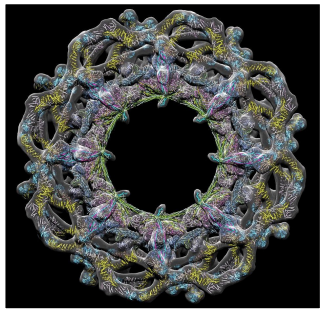

The **ribosome**
10^{7} atoms

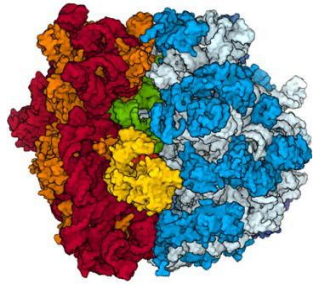

Tobacco mosaic virus
10^{6} atoms

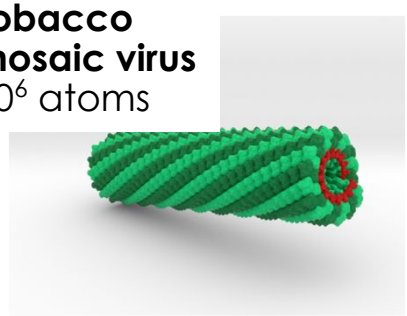

- ML-generated force fields
- AI-enabled correlation analysis, cluster analysis, rare event analysis, etc.
- Structural predictions based on both AI and MD methods

Energy-Efficient Jet Engines

GE Aerospace plans to reduce aviation emissions by 20%

- Larger fans are more efficient
- Nacelle drag offsets gains
  - Lose the nacelle
- Cannot simulate physically
  - No wind tunnel large enough
- Frontier provides enough resolution to understand turbulence
  - Frontier can only model one blade
- GE needs to model whole engine (10-12 blades) and the wing
  - Both level flight and take-off climb
  - The problem size will not fit on Frontier

Below: The novel GE RISE engine featuring open fan blade design.

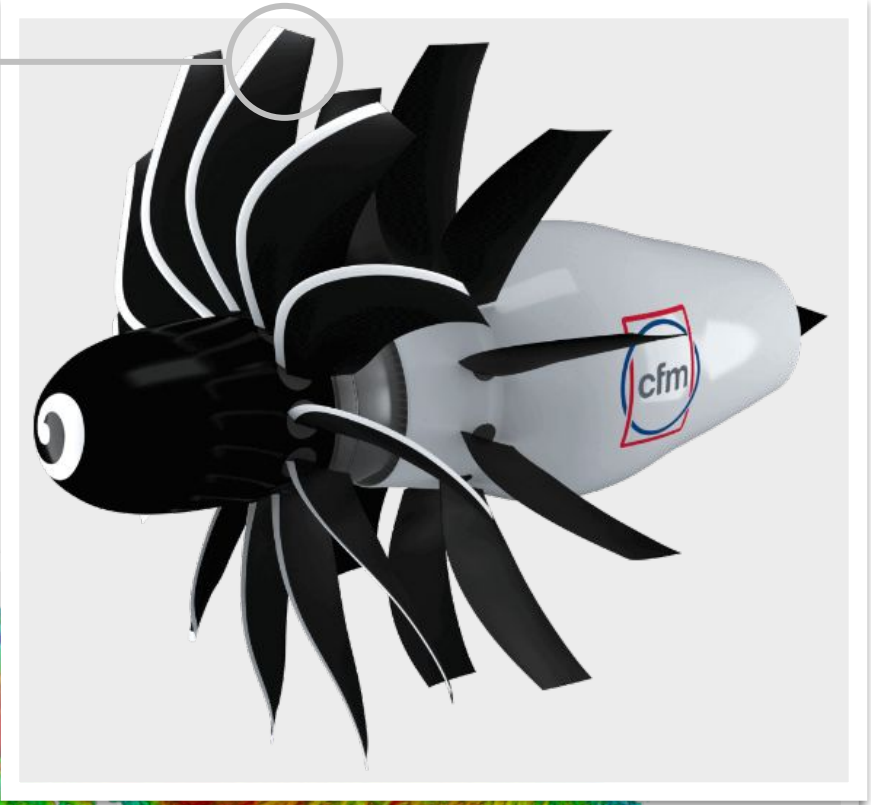

Below: A GE Aerospace visualization of turbulent flow in the tip region of an open fan blade. These models could help reduce CO_2 emissions by more than 20%.

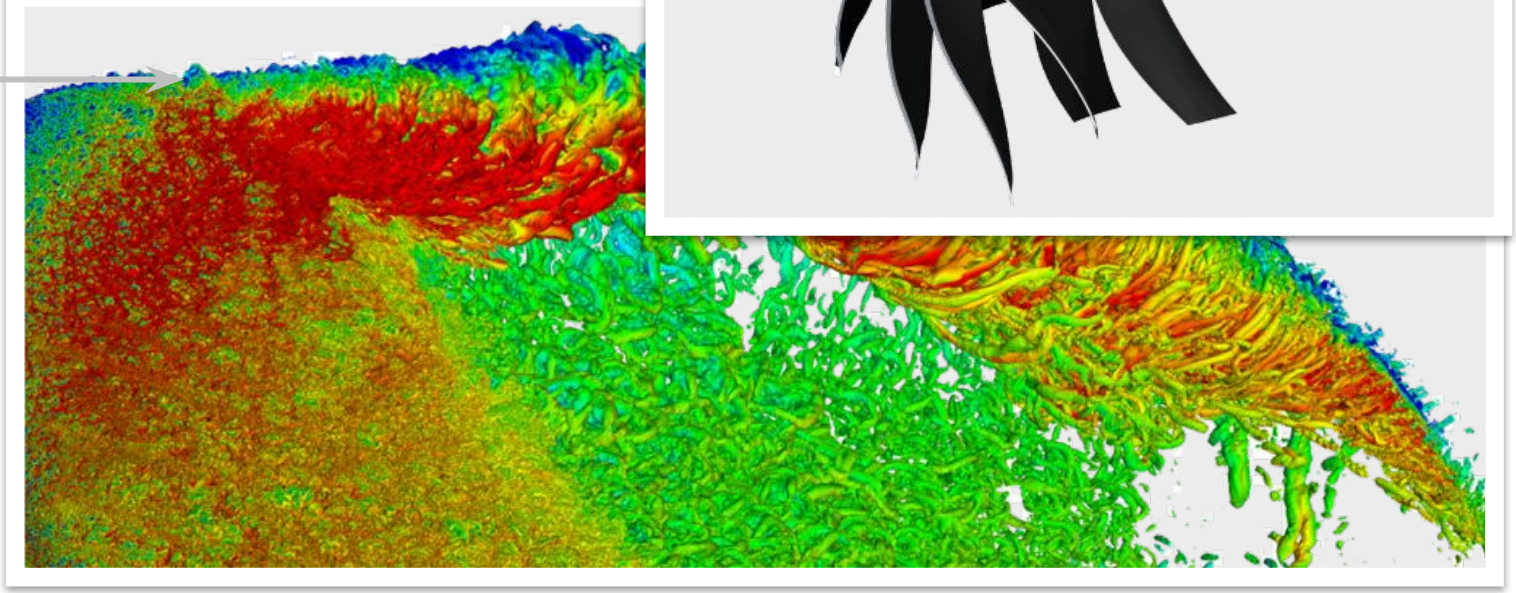

OAK RIDGE National Laboratory

DOE's Integrated Research Infrastructure

Three IRI Blueprint Science Patterns

Time-Sensitive Pattern

- Workflows that have time-critical requirements (i.e., real time) including rapid decision-making, experiment control, coordinating distributed assets, and data capture/reduction

Data Integration-Intensive Pattern

- Analysis of data from multiple sources, e.g., simulations and experiments/observations
- Cross-site data-driven discovery
- AI/ML incorporated into simulations and experiments

Long-Term Campaign Pattern

- Sustained access (several years) to resources at scale, e.g., sustained simulation production and large data (re)processing for collaborative use

https://www.osti.gov/biblio/1984466

LCFs in the Integrated Research Infrastructure Ecosystem

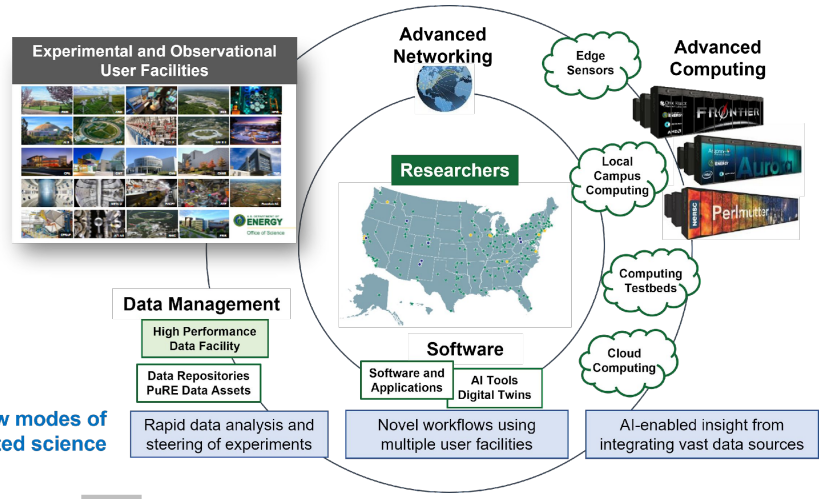

ALCF and OLCF Joint Leadership Computing Mission

- Leadership-scale compute and data storage allocations
- Enable new, novel, and unique data analyses
- Other federal agency and industry users
- Diversity of technology
- Secure computing

NERSC and HPDF High-Performance Production Computing

- SC-driven allocations
- Reliability, active-active failover
- Data wrangling
- Long-term archival storage

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

OAK RIDGE National Laboratory

Changing Advanced Computing Ecosystem

We live in a new world

- We are spoiled; we had >50 years of more for less
- Dennard Scaling ended ~20 years ago
- Slowing (ending?) of Moore's Law
- But processors are still getting X times faster…
  - Packaging (integrating more silicon into a processor)
  - But gains come with a price
    - More silicon area costs more
    - More power
- Adding insult to injury, transistors costs more
  - DOE will struggle to deliver more performance

There is hope

- Harnessing reduced precision to calculate full precision
  - HPL-MxP shows how to use reduced precision plus iterative refinement to solve for full precision
  - Not always guaranteed to converge 😕
- New architectures
  - Data flow, coarse-grained reconfigurable
  - Quantum (for a limited class of algorithms)
- Specialization
  - RISC-V
  - Chiplets + UCIe

OAK RIDGE National Laboratory

OLCF-6's RFP adapts to the new Ecosystem

OLCF-6 System Goals (1/2)

Provide Leadership Modeling and Simulation (ModSim) capabilities

Both strong and weak scaling

Focus on apps using 20-100% of the system

Analysis and data science

Explore and incorporate transformational AI technology to accelerate scientific discovery

Physics-based, validated, high fidelity

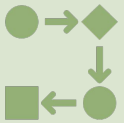

Support DOE's Integrated Research Infrastructure

Extend Leadership capabilities to enable new workflows between facilities

OLCF-6 System Goals (2/2)

Improve Energy-Efficiency

- Co-design with vendors to optimize performance/watt
- Optimize from processor to node to system to whole data center
- Investigate new architectures/accelerators

Expandable to add new capabilities and to support new workloads

- Quantum
- Specialized AI accelerators
- Novel architectures

Tailor storage to meet current and new paradigms

- Parallel file system to support write-optimized ModSim workload
- AI-optimized (i.e., small, random read) storage
- Securely handle protected data

The Elephant in the Room: OLCF-6 Engages the Cloud

- **No cloud vendor wants to sell a stand-alone system**
  - They view a dedicated resource as an extension of their cloud
- **Allow cloud vendors to propose on-prem or off-prem**
- **Known-unknowns risks**
  - How to manage, access telemetry, understand locality, assess energy-efficiency
- **Unknown-unknowns**

- OLCF Cloud Evaluation Report 2022
- New cloud pilot to study integration with OLCF and infrastructure manageability

OAK RIDGE National Laboratory

How the SOS Community can help

Areas of Interest

- Resiliency mitigations
- Power/energy monitoring, control, and optimization
  - Application-level and across the data-center
- Federated ID management
- Multiple resource manager integration
  - E.g., Slurm and K8s
- Data services
  - Streaming (not file based), tiering migration (between flash, HDD, archive)

Areas of Interest

- Hybrid cloud management
- Frontier Digital Twin
- System telemetry analysis tools
- Quantum integration with classical computing
- Many more

Thank you!