

# Pushing the Boundaries of Biomedical Science with Frontier

**Georgia Tourassi, PhD**

Director, National Center for Computational Sciences

Director, Oak Ridge Leadership Computing Facility

Oak Ridge National Laboratory

June 22, 2022

Presented at the CompBioMed Center of Excellence

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

THERE'S FAST,  
and there is  
**FRONTIER**  
fast.



At 1.1 exaflops of performance,  
Frontier is faster than

the next **7** systems combined  
on the Top500;

and the next **68** systems combined  
after that;

and the next **393** systems combined  
after that.

Source: May 30, 2022 Top500 release

# OAK RIDGE NATIONAL LABORATORY'S FRONTIER SUPERCOMPUTER



- 74 HPE Cray EX cabinets
- 9,408 AMD EPYC CPUs, 37,632 AMD GPUs
- 700 petabytes of storage capacity, peak write speeds of 5 terabytes per second using Cray Clusterstor Storage System
- 90 miles of HPE Slingshot networking cables

TOP500

#1

1.1 exaflops of performance on the May 2022 Top500.



GREEN500

#1, #2

62.04 gigaflops/watt power efficiency on a single cabinet.

52.23 gigaflops/watt power efficiency on the full system.



HPL-AI

#1


6.88 exaflops on the HPL-AI benchmark.



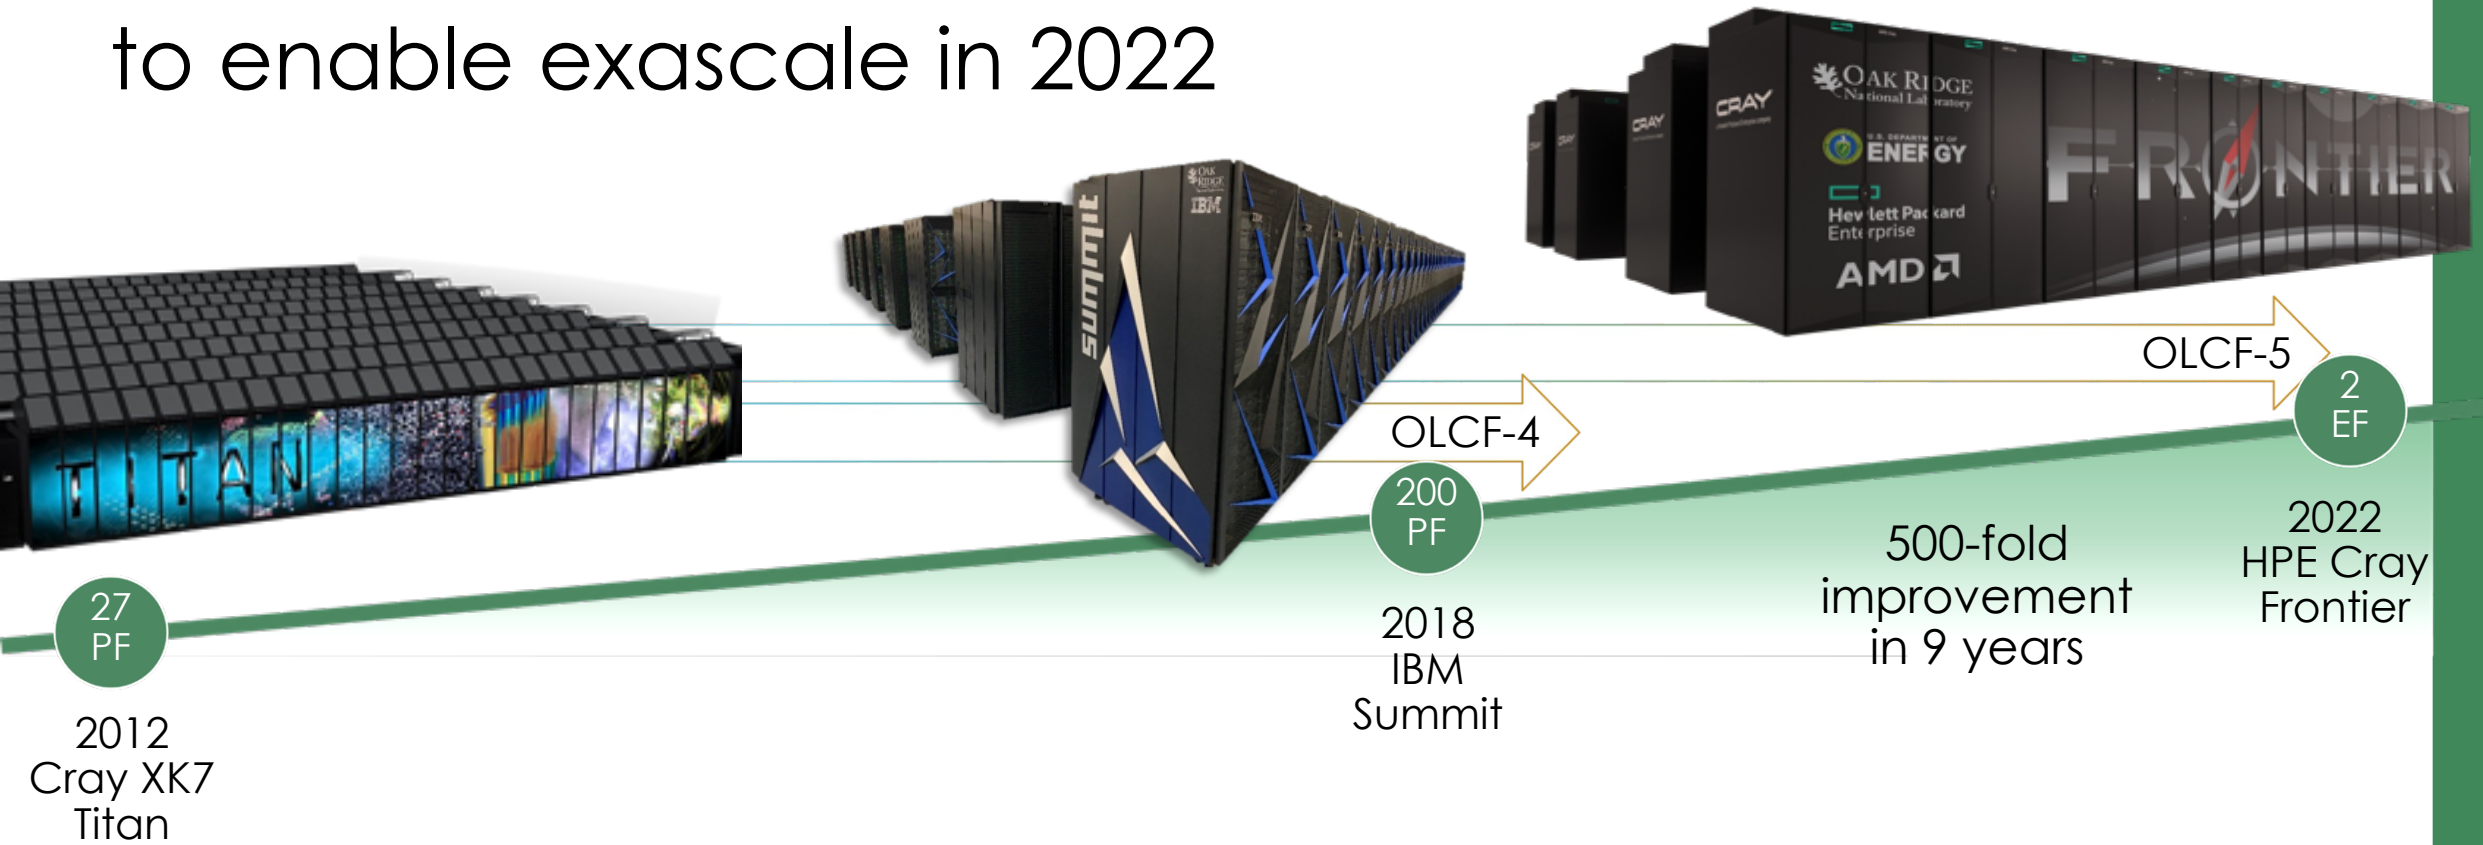
Sources: May 30, 2022 Top500 release

# ORNL has systematically delivered a series of leadership-class systems

On scope • On budget • Within schedule



# We have built on this record of success to enable exascale in 2022



## Steady progress of 10x per generation



# Updated power and cooling to support Frontier

2.5 miles x8 lines of power



New energy plant



29MW of power











Computing  
and data

Advancing  
artificial  
intelligence

Supporting  
traditional  
modeling &  
simulation

Scale  
computing  
& data analytics  
to exascale  
and beyond

Image courtesy of NASA





# On our journey to Exascale, we found an architecture that could excel at simulation, data analytics, and artificial intelligence

As supercomputers got larger and larger, we expected them to be more specialized and limited to just a small number of applications that can exploit their growing scale

We found that “Summit” architecture with few, large-memory, multi-GPU nodes excels at:



- Data analytics – CoMet bioinformatics application for comparative genomics.  
**Has achieved 2.36 ExaOps** mixed precision (FP16-FP32) on Summit (2018 Gordon Bell Winner)
- Deep Learning – Climate: neural network learns to detecting extreme global weather patterns  
**Has achieved a sustained throughput of 1.0 ExaOps** (FP16) on Summit


**Frontier Exascale computer uses and improves on Summit’s successful architecture**

# Oak Ridge has been working on Energy Efficient Computing for a Decade

Since 2009 the biggest concern with reaching Exascale has been energy consumption

In 2009 a target was set at 20 MW for 1 EF

- **ORNL pioneered GPU use in supercomputing** beginning in 2012 with Titan thru today with Frontier. Significant part of energy efficiency improvements.
- **200x reduction in energy per (peak) FLOPS** from Jaguar to Frontier at ORNL
- ORNL achieves additional energy savings from using warm water cooling in Frontier (32 C).  
**ORNL Data Center PUE= 1.03**
- **Frontier and its TDS are #1 and #2 on Green500**
  - Fronter TDS Ranked #1 62 Gflops/Watt
  - Frontier Ranked #2 52 Gflops/Watt





## One more word on power efficiency

- One cabinet of Frontier has a 10% higher HPL than all of Titan
  - While only using 309 kW compared to the Titan's 7 MW




One Cabinet  
2.5 m<sup>2</sup>




200 Cabinets  
404 m<sup>2</sup>

# Frontier science and application readiness



# OLCF enabling science across domains




## Identification of **drug compounds** for treatment of **COVID-19**

Simulation of different compounds docking to the SARS-CoV-2 coronavirus' spike protein

ORNL and Univ. of Tenn. – **Summit**

Image: Compound (gray) calculated to bind to the SARS-CoV-2 Spike Protein (cyan) to prevent it from docking to the Human Angiotensin-Converting Enzyme 2 (ACE2) receptor (purple).




## Demonstration of **superconductivity**

Discovery that hydrogen atoms in a metal hydride material are much more tightly spaced than had been predicted for decades

ORNL – **Titan**

Image: Illustration of a zirconium vanadium hydride atomic structure at near ambient conditions as determined using neutron vibrational spectroscopy and Titan.




## Detection of **cancer**

Multi-objective neural network to research cancer pathology on tumors in biopsy images

ORNL – **Summit**

Image: A portion of a whole-slide cancer biopsy image where tumors appear in the orange overlay. It was labeled using the MENNDL neural network.




## Discovery of **quantum materials**

Use of AI to find patterns in neutron scattering data that can lead to an understanding of the physics inside quantum or complex magnetic materials.

ORNL – **Titan, Eos, Rhea, and Summit**

Image: Atomic structure of Dy<sub>2</sub>Ti<sub>2</sub>O<sub>7</sub>, a spin ice that is glass-like at low temperatures.



## Simulation of **quantum transport**

**Gordon Bell Prize winner**


Simulation of the transport of electric charge carriers through nanoscale materials

ETH Zurich – **Summit**

Image: As electrons flow through transistors, such as the ones on circuit boards, they generate heat that dissipates into the environment around them. This project could inform production of new semiconductors with optimal heat-evacuating properties.



## Simulating the Spread




*When the ventilation jet and exit door are on opposite sides of the room, a dead zone is formed in which aerosol particles can linger. Image Credit: Ramesh Balakrishnan, Argonne National Laboratory*

A team at Argonne National Laboratory used Summit to simulate the spread of aerosols through a modern elementary school classroom layout.

Balakrishnan, Ramesh, Rao Kotamarthi, and Paul Fischer. "Large Eddy Simulation of Isothermal and Non-Isothermal Turbulent Flows in Ventilated Classrooms." Paper presented at the 74th Annual Meeting of the APS Division of Fluid Dynamics, Phoenix, AZ, November 21–23, 2021.

## Pathways for Inhibiting COVID-19




*An ORNL-led team studied the SARS-CoV-2 spike protein in the trimer state, shown here, to pinpoint structural transitions that could be disrupted to destabilize the protein and negate its harmful effects. Image Credit: Debsindhu Bhowmik, ORNL*

An ORNL team ran molecular dynamics simulations on Summit, then analyzed the output to produce a complete molecular picture of the "spike" protein on the surface of the SARS-CoV-2 virus.

Serena H. Chen, M. Todd Young, John Gounley, Christopher Stanley, and Debsindhu Bhowmik, "How Distinct Structural Flexibility within SARS-CoV-2 Spike Protein Reveals Potential Therapeutic Targets," IEEE International Conference on Big Data (2021). doi: 10.1109/BigData52589.2021.9671323

## Gordon Bell Nomination for COVID-19 Reproduction




*The team built an integrated workflow that models the experimental Cryo-EM volumetric data as a finite element mesh, which is then simulated using fluctuating finite element analysis (FFEA). To annotate the mesh with the strengths of the protein-protein interactions, the researchers used all-atom molecular dynamics (AAMD) simulations and employed AI-techniques to automatically bridge the FFEA and AAMD simulations. Credit: Gorgun, Trifan, Ramanathan*

A team led by Arvind Ramanathan at Argonne National Laboratory used Summit to peer inside the intricacies of how the SARS-CoV-2 virus reproduces itself.

Anda Trifan, et al. "Intelligent Resolution: Integrating Cryo-EM with AI-driven Multi-resolution Simulations to Observe the SARS-CoV-2 Replication-Transcription Machinery in Action." Paper Presented at the International Journal of High Performance Computing Applications, 2021

## Aerosolized SARS-CoV-2 viral particle



*Visualization of an aerosolized SARS-CoV-2 Particle. Image Credit: #COVIDisAirborne Team*

A team led by Rommie Amaro of the University of California, San Diego has used the Summit supercomputer to model an aerosolized SARS-CoV-2 viral particle for the first time.

Dommer, Abigail, et al. "#COVIDisAirborne: AI-Enabled Multiscale Computational Microscopy of Delta SARS-CoV-2 in a Respiratory Aerosol." Paper to be presented at SC21: The International Conference for High Performance Computing, Networking, Storage and Analysis, St. Louis, MO, November 2021.


## Stepping up CANDLE



John Gounley, an ORNL computational scientist in the Advanced Computing for Health Sciences Section of the Computational Sciences and Engineering Division at ORNL, poses with visualizations from his red blood cell modeling work. Image Credit: Genevieve Martin, ORNL

John Gounley is working on codes for the CANCER Distributed Learning Environment (CANDLE) project, which will provide deep learning methodologies to accelerate cancer research on machines like Frontier, the upcoming exascale system at ORNL.

## CSB Protein in DNA Repair



This structural model of Rad26-RNA polymerase II complex from cryo-electron microscopy sheds light on the mechanism of transcription-coupled DNA repair. The Rad26 molecular motor (in light and dark green) binds to the upstream DNA duplex (light blue) of a stalled RNA polymerase II and redirects the DNA path.

Using the Summit supercomputer, a Georgia State University (GSU) research team has revealed the structural mechanism of the Cockayne Syndrome B (CSB) protein in transcription-coupled DNA repair.

"Mechanism of Rad26-assisted rescue of stalled RNA polymerase II in transcription-coupled repair." *Nature Communications*, 2021

## Swift Protein Modeling




Ensemble refinement of molecular motors with CryoFold. Image Credit: Abhishek Singharoy, ASU

A team led by Abhishek Singharoy at Arizona State University is using Summit to model molecular structures in fine detail. The team recently simulated millions of structures and matched the results with experimental data to gain new insights into how proteins transition to different shapes, or conformations.

Mrinal Shekhar, et al., "CryoFold: Determining Protein Structures and Data-Guided Ensembles from Cryo-EM Density Maps," *Matter* 4, no. 10 (2021): 3195–216, <https://doi.org/10.1016/j.matt.2021.09.004>.

## Protein Functions at Genome Scale




This protein drives key processes for sulfide use in many microorganisms that produce methane, including *Thermosiphon melanesiensis*. Researchers used supercomputing and deep learning tools to predict its structure, which has eluded experimental methods such as crystallography. Image Credit: Ada Sedova, ORNL

A team of scientists led by ORNL and the Georgia Institute of Technology is using supercomputing and revolutionary deep learning tools to predict the structures and roles of thousands of proteins with unknown functions.

Mu Gao et al., "High-Performance Deep Learning Toolbox for Genome-Scale Prediction of Protein Structure and Function." Paper presented at the IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC) *Physics of Plasmas* 28, no. 6 (2021).

## Simplified Protein Binding



Dimers, protein structures consisting of two bound proteins, from the Dockground database. The interface at which the proteins meet is shown in dark gray. Image credit: ORNL


A team led by Sharon Glotzer at the University of Michigan used the Summit supercomputer at ORNL to model lock-and-key interactions between proteins to study their binding behaviors.

Fengyi Gao, Jens Glaser, and Sharon Glotzer, "The Role of Complementary Shape in Protein Dimerization," *Soft Matter* (2021), doi:10.1039/D1SM00468A.

# Growing Biological/Biomedical Applications



# Synthetic Biology



# Bioimaging



## Precision Medicine



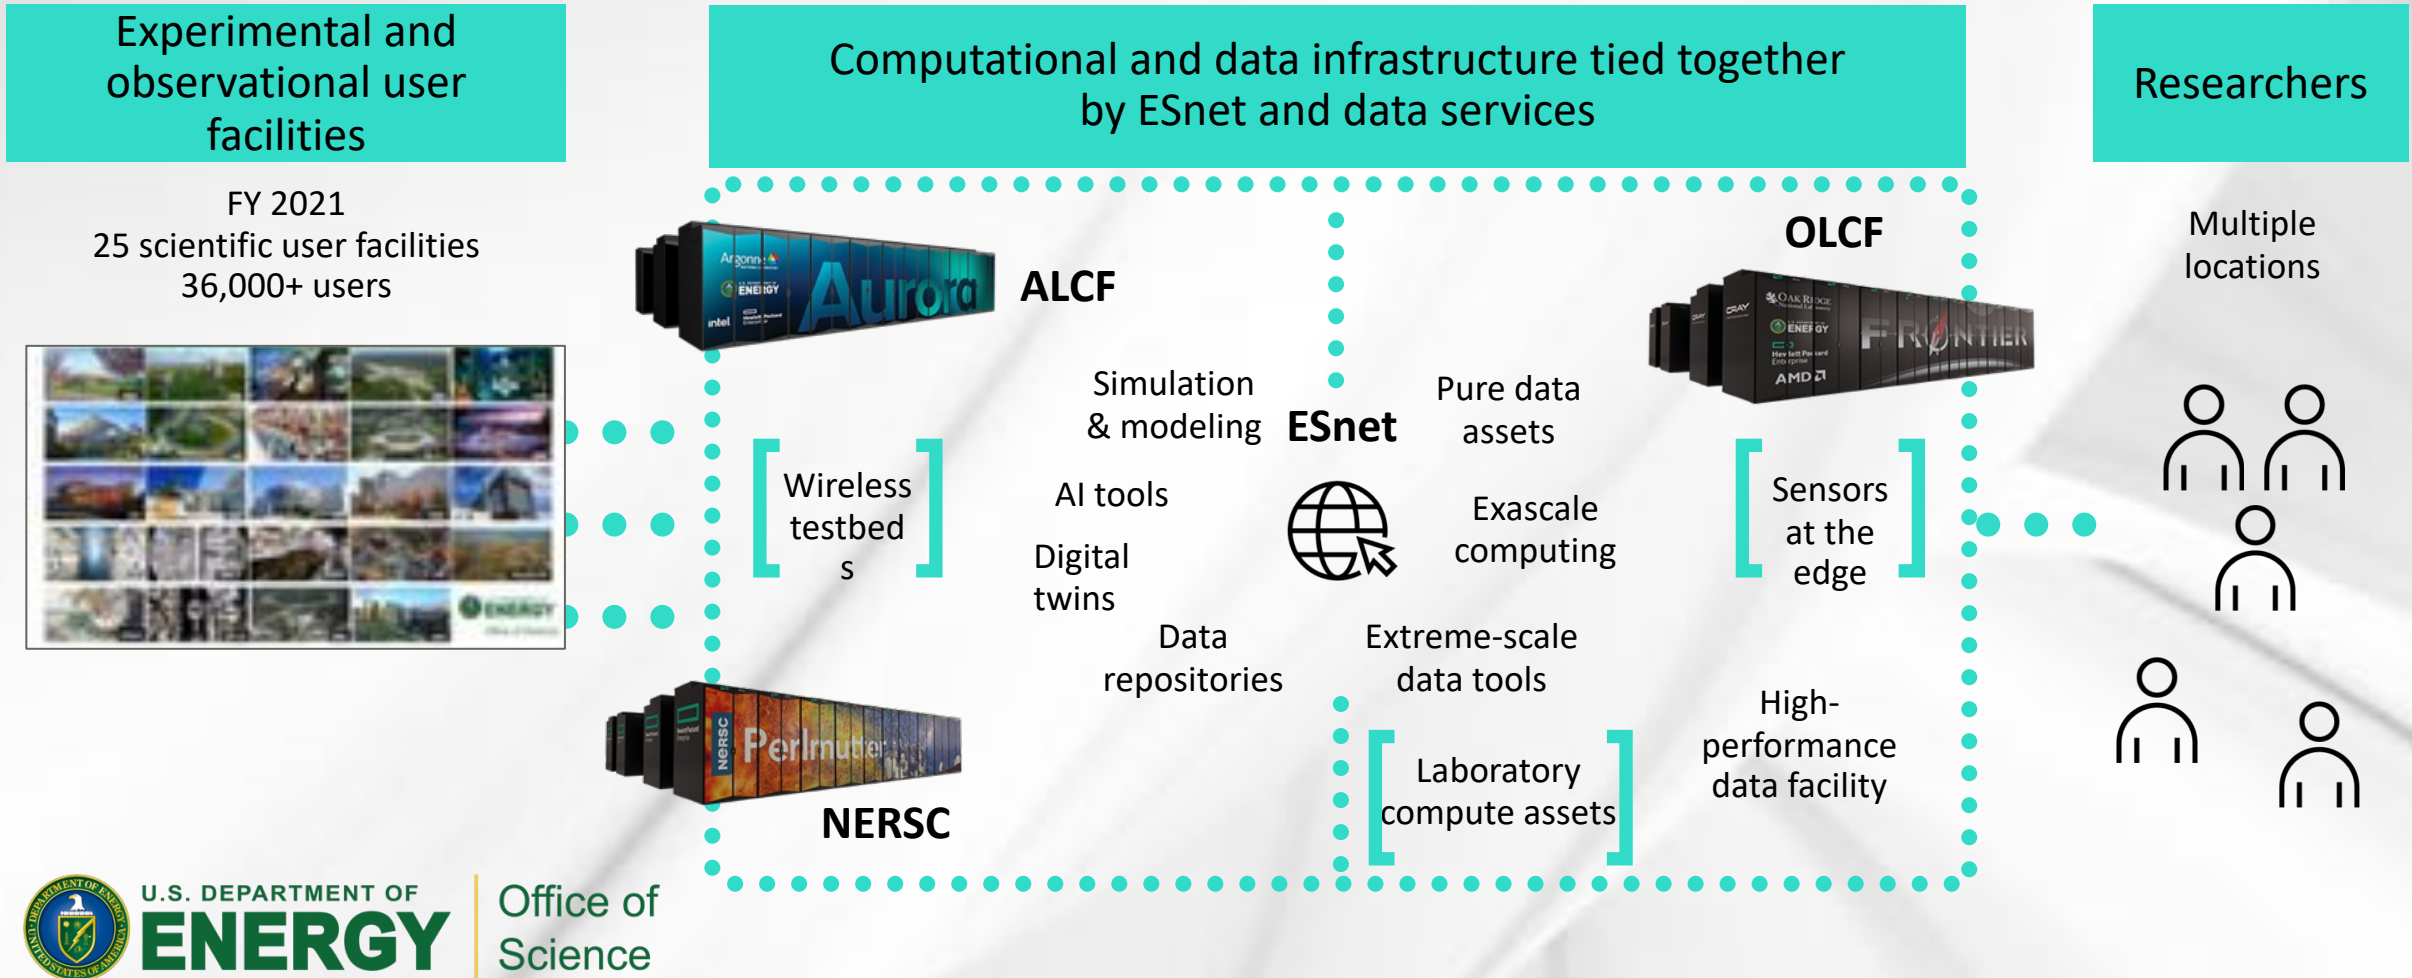
## Digital Twin



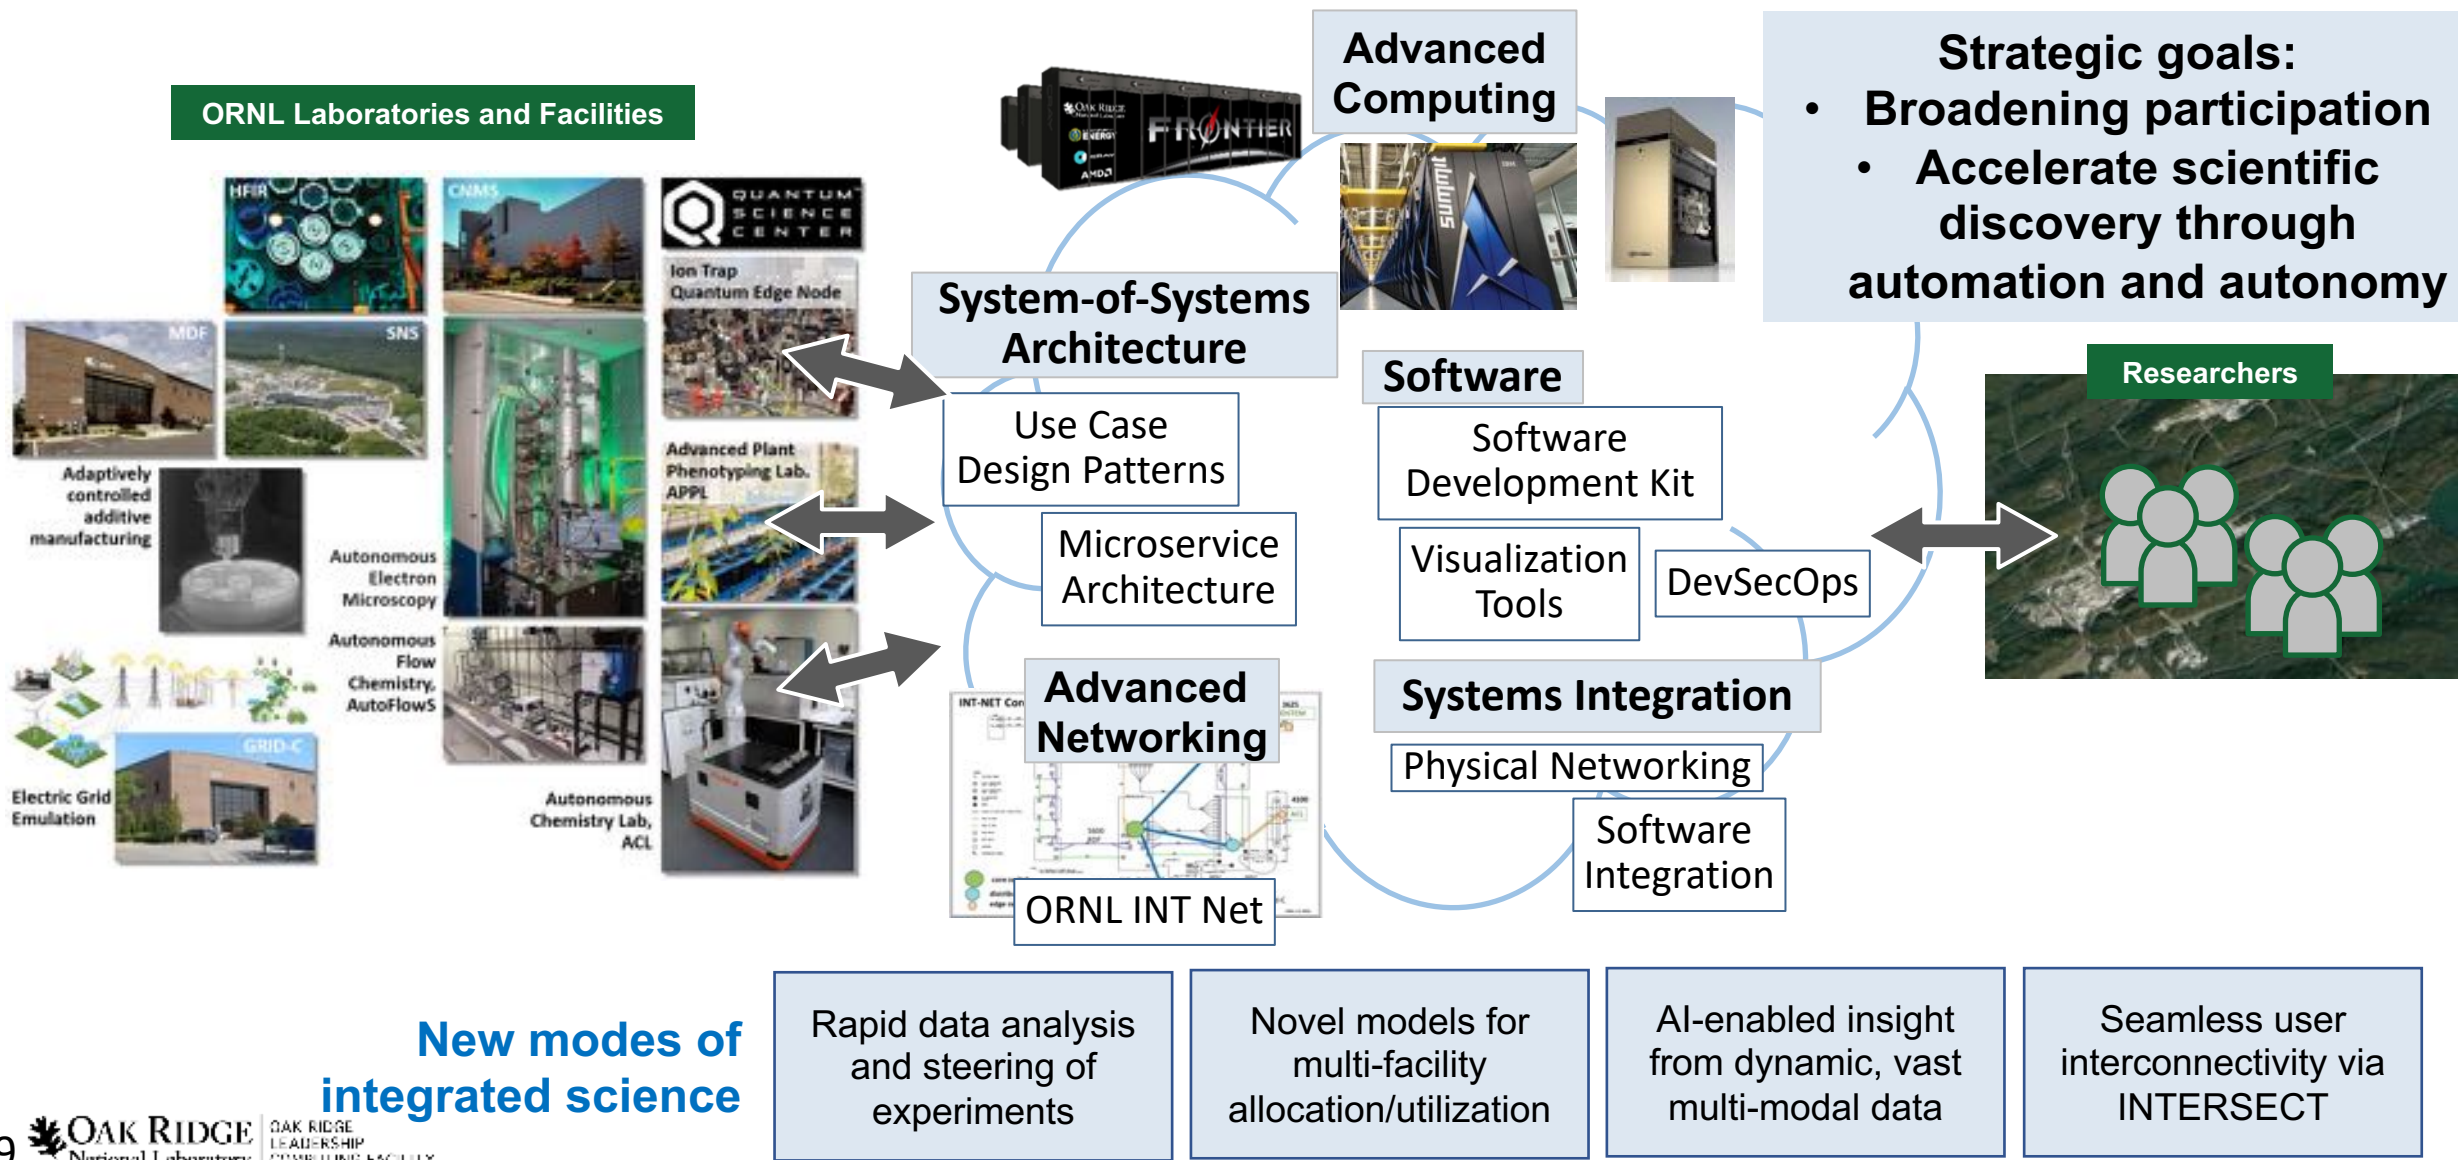
## Drug Discovery



# Incipient ecosystem: Office of Science User Facilities



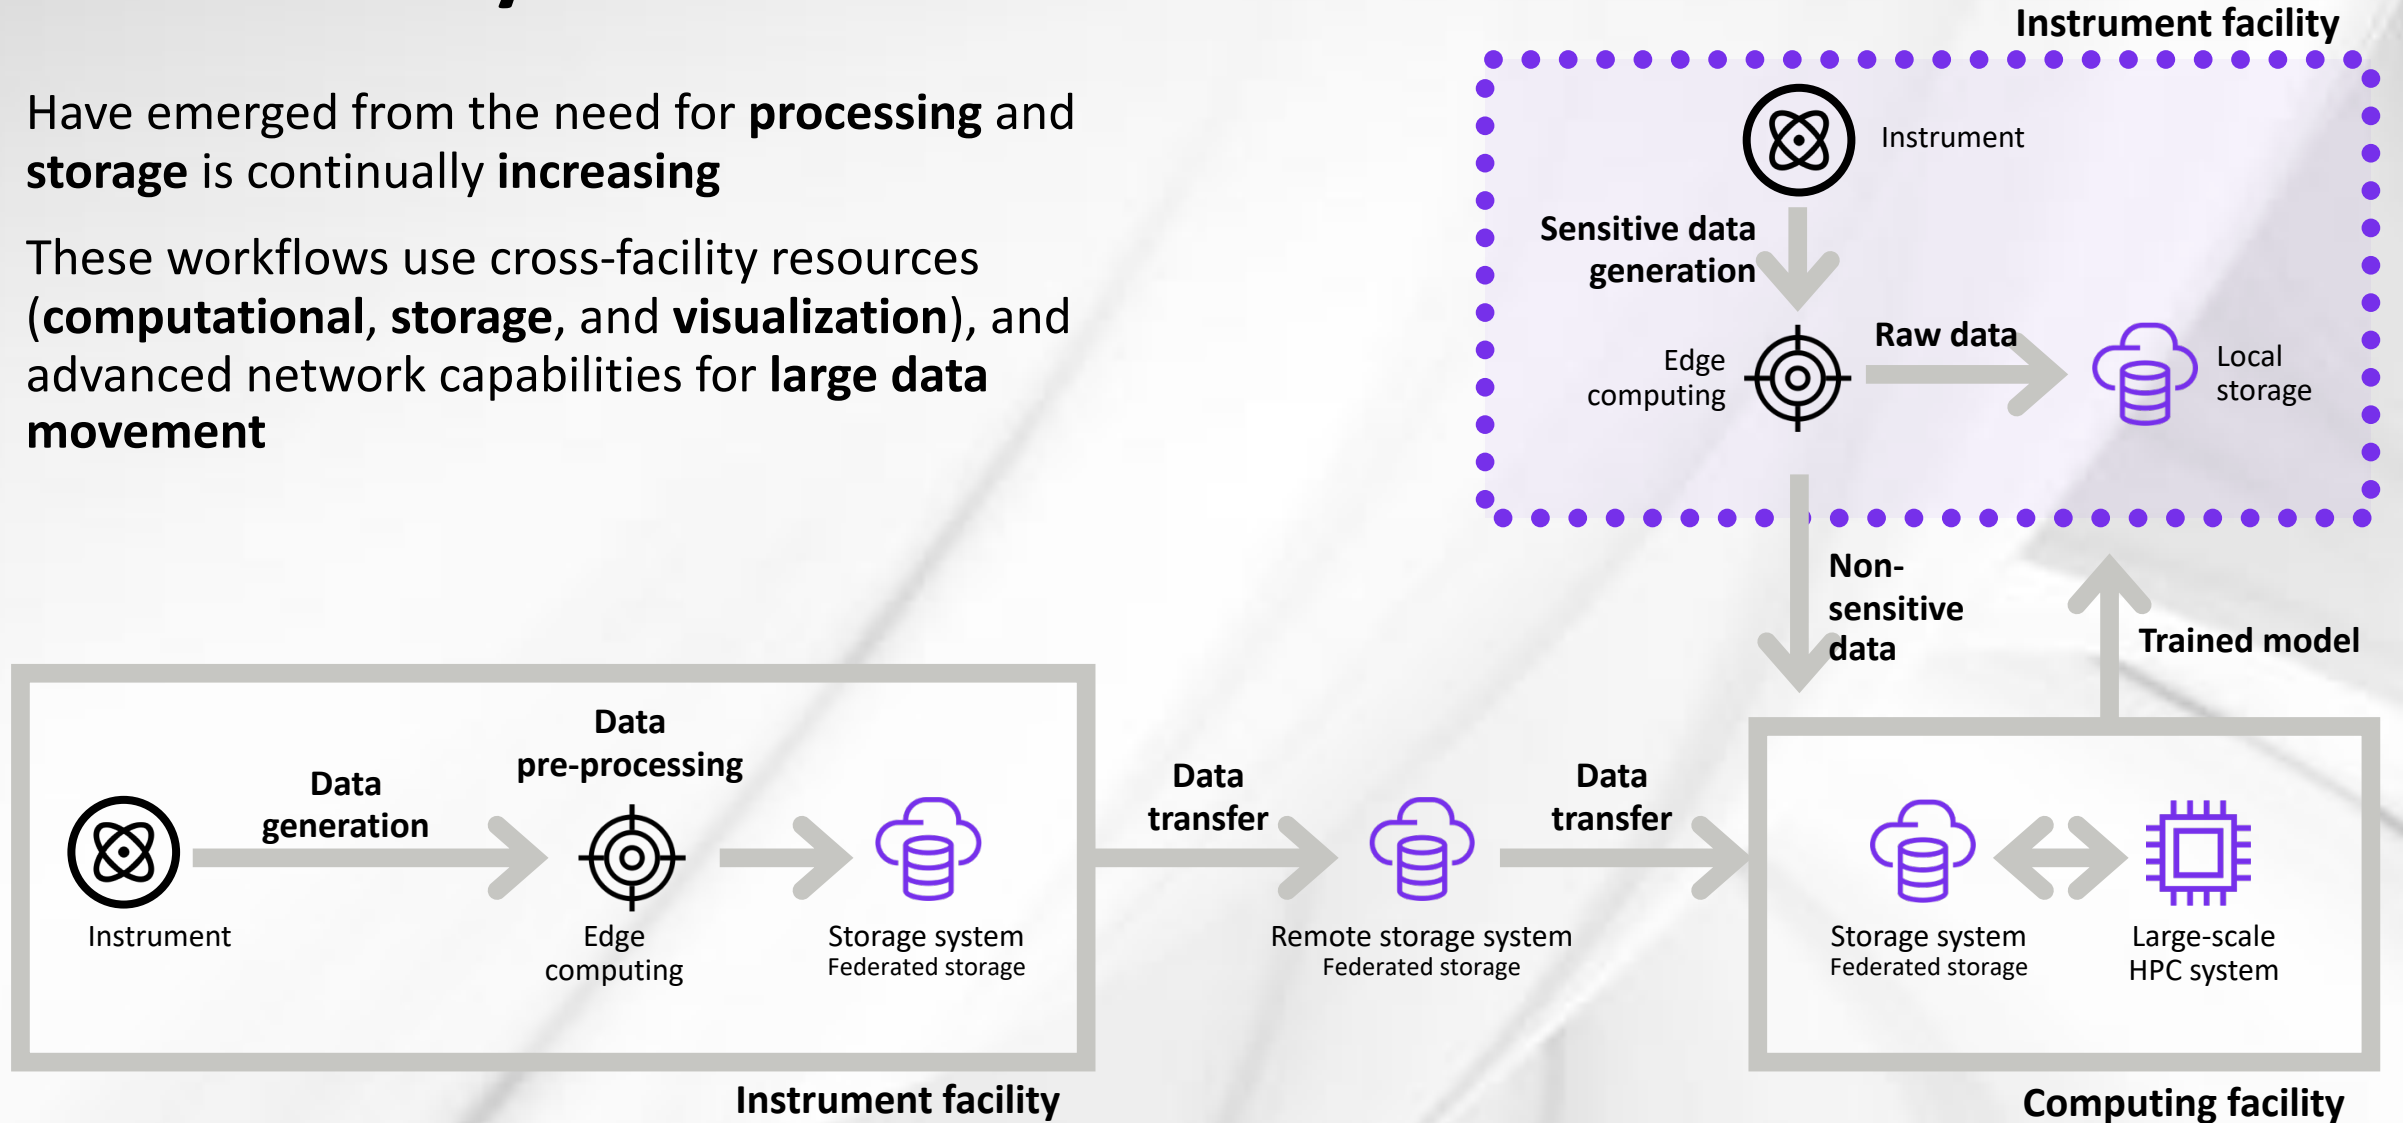
# The Exascale Ecosystem Builds Upon our Experience with INTERSECT and our Initiatives



# Cross-facility workflows

Have emerged from the need for **processing** and **storage** is continually **increasing**


These workflows use cross-facility resources (**computational, storage, and visualization**), and advanced network capabilities for **large data movement**



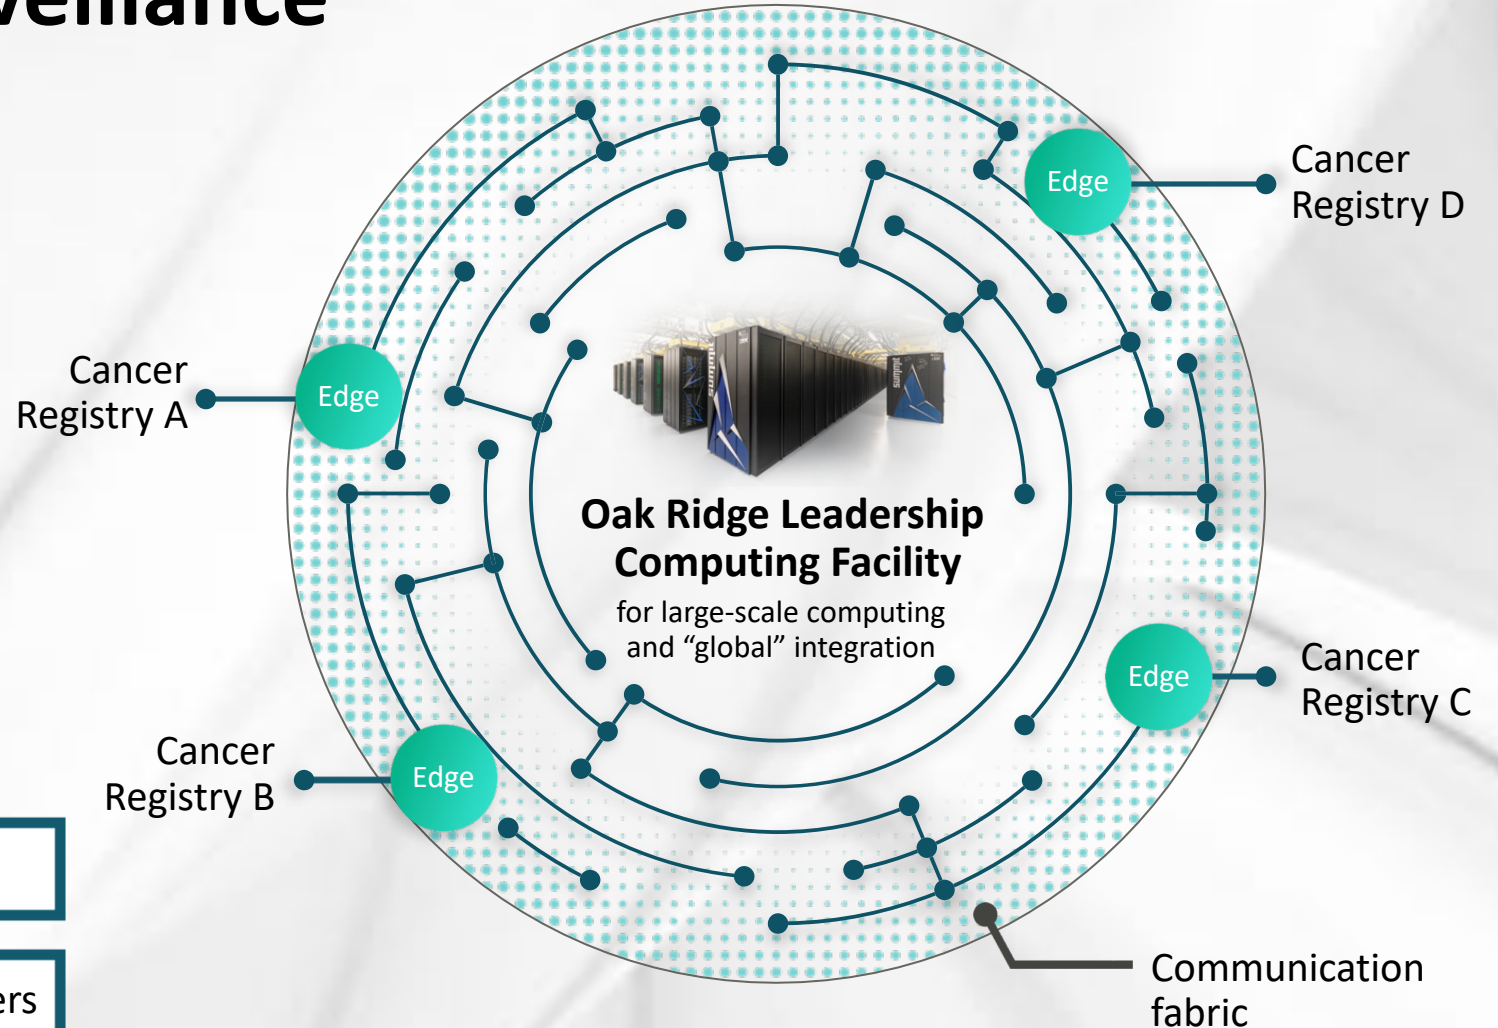

EMERGING NEED: Federated learning with privacy-preservation



# Cross-agency workflow: AI-enabled cancer surveillance



\*Subcontract under New Mexico  
\*\*Three regions represent the state of California: Greater Bay, Los Angeles, and Greater California



# OLCF Innovations in Privacy Preserving Technologies

## CITADEL

*A workflow for scalable analysis of protected data*

Enables computing of **protected data assets**

*(e.g., Protected Health Information and International Traffic in Arms Regulations)*

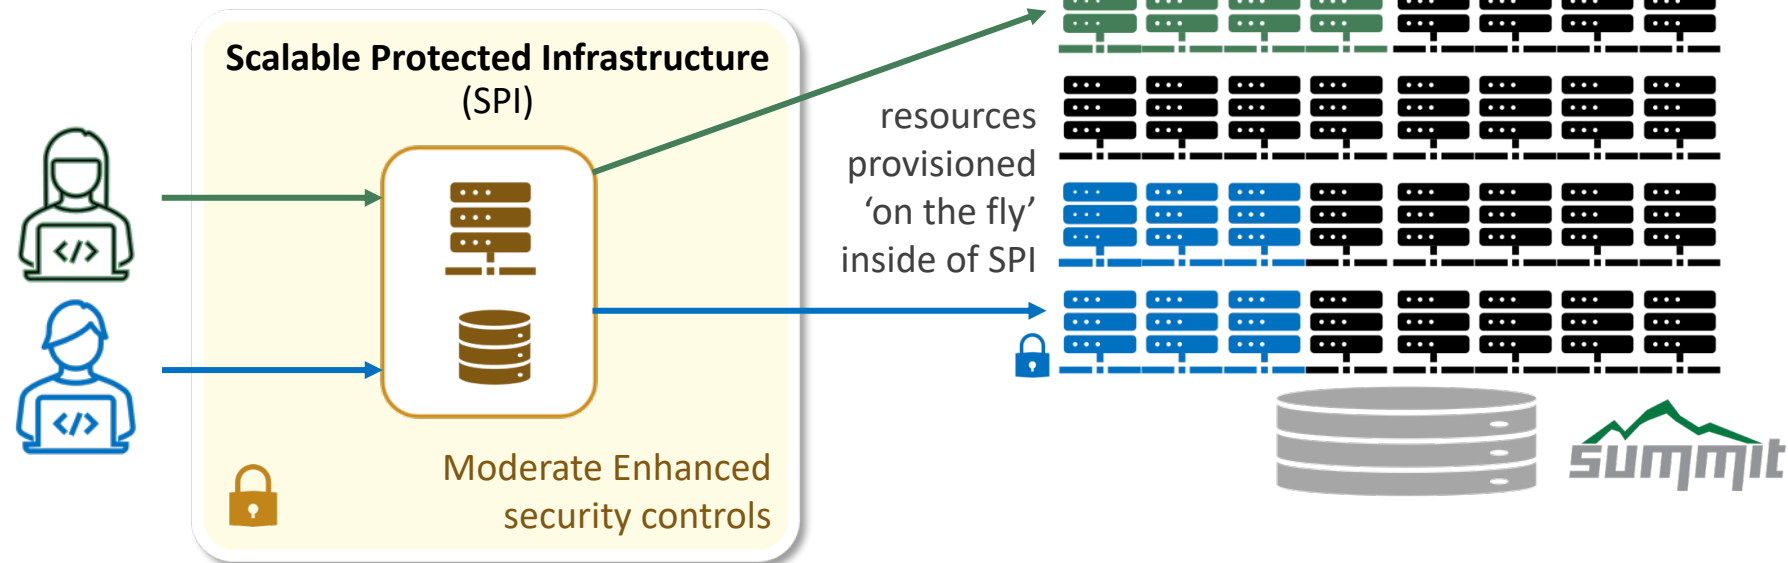
Comply with **national data protection standards**

*(i.e., FISMA, HITECH and NIST SP 800-66)*

Authority to Operate (ATO) received in 2021



Allows **sensitive data** to be processed on an isolated **protected environment**



# Allocation Programs

## INCITE




The INCITE program receives 60% of the available LCF hours annually

### INCITE 2020

- Summit : 22,455,974 node hours consumed by 39 awarded projects.
- Theta : 20,876,433 node hours consumed by 14 awarded projects.

## ALCC



The ALCC program receives 20% of LCF hours annually

### ALCC 2020

- Summit : 6,524,243 node hours consumed by 52 awarded projects (ALCC 2020 + ALCC 2021)
  - In 2020, the ALCC program provided the hours for the HPC COVID-19 Consortium which allocated 21 projects and 2,328,000 node hours
- Theta : 10,700,00 node hours consumed by 36 awarded projects (ALCC 2020 + ALCC 2021)
  - 14 HPC COVID-19 Projects, consumed 2,380,000 node hours

## ECP



The LCF provides resources for the ECP AD and ST teams for their work in preparing for Frontier and Aurora Access

- In CY 2020 the OLCF provided 1,912,122 Summit Node hours in support of the ECP goals
- In CY 2020 700,000 Theta node hours supported ECP
- Up to 10% of the allocable hours

**The remaining 10% of the allocatable hours go to the LCF Director's Discretionary Program**



Resources mentioned in this talk:

<https://www.olcf.ornl.gov/>

<https://www.doeleadershipcomputing.org/>

<https://science.osti.gov/ascr/Facilities/Accessing-ASCR-Facilities/ALCC>



**Thank you!**

# Questions/Discussion?

