

ASCR Facilities Division FY 2021 Highlights

Advanced Scientific Computing Advisory Committee meeting
September 29, 2021



U.S. DEPARTMENT OF
ENERGY

Office of
Science

ASCR Facilities Division highlights, FY 2021

- We achieved significant ALCC program improvements, highlighted by the ALCC Working Group spanning stakeholder DOE program offices.
- The ASCR Facilities continue to innovate in our Requirements Reviews
 - Leadership Computing for clean energy and advanced manufacturing
 - Leadership Computing for biodefense applications
 - ESnet High Energy Physics program and Fusion Energy Sciences program
- The ASCR Integrated Research Infrastructure Task Force published a white paper that has positioned us for deeper engagement with non-ASCR user facilities and program offices.

ALCF operations and project highlights, FY 2021



ALCF-3 Project

- Modification to Intel Build Contract completed in February
 - 10 new racks were added
 - PCR approved for additional power for the new racks
 - GPU changed to PVC-XT
- Authorization was provided to Intel for ordering material for phase 1 (scheduled to complete in Q1 of CY2022)



Polaris Testbed


- Installation of cabinets and the cooling system is running on schedule.
- Polaris is equipped with 560 2nd and 3rd Gen AMD EPYC processors and 2,240 NVIDIA A100 Tensor Core GPUs
- Polaris will deliver approximately 44 PF of peak double precision performance and nearly 1.4 EF of theoretical artificial intelligence (AI) performance
- Will be available for testing and porting applications in 2022



Core Operations

- ALCF exceeded target metrics for system availability, INCITE hours delivered, and capability hours delivered.
- Theta GPU user availability and production started in January 2021.
- Storage system growth investments continue.


ALCF Notable Science Highlights, FY2021



[Energy efficient buildings](#)


Using Theta, scientists built a digital twin of over 178,000 buildings, showing that 99% of buildings will benefit from existing energy saving measures, potentially offsetting 500-3000 lb of CO₂.

[Energies. 2021; 14\(1\):132](#)



[Cancer research at Exascale](#)


Using Aurora's unique capabilities, this Early science project will provide new insights into spread of cancer cells through blood and will guide new treatment paths and drug delivery.



[Machine learning for solar energy](#)

Scientists used ML enhanced, data-driven simulations of light absorption at water-photoelectrode interface, creating a new pathway for next-generation technologies, from solar fuel cells to water treatment systems.

[Chem. Sci., 2021,12, 4970-4980](#)



[COVID-19 drug discovery](#)

Using ML techniques on Theta and Theta GPU, this project analyzed data from billions of molecules, screening 4 billion potential drug candidates.

ESnet operations and project highlights, FY 2021

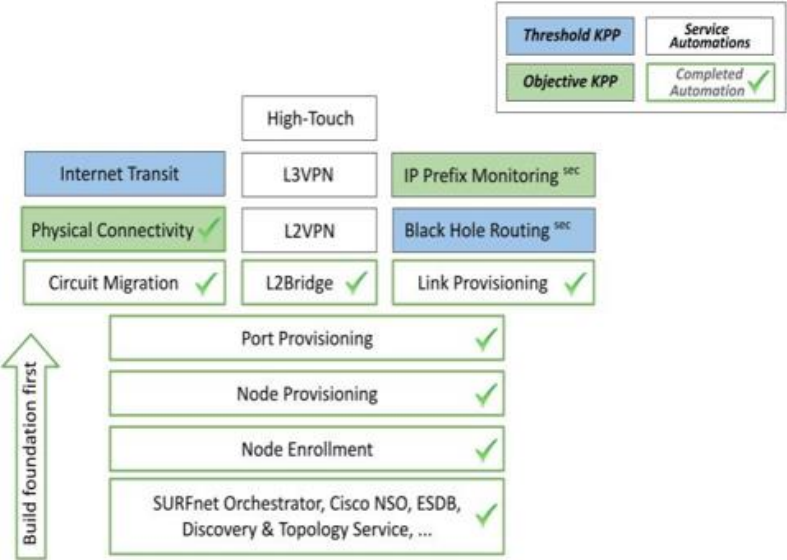
Optical: Installation of ESnet6 substrate across 300 locations; optical service has fully transitioned from ESnet5 to ESnet6; decommissioning of ESnet5 optical completed Q1, FY21.

Programmability & Automarion: Software automation of network orchestration workflows are fundamentally transforming deployment of ESnet6; Automation threshold KPP complete by Q2, FY22.

Network: Significant progress in deployment of capacity and new routers; Network threshold KPP expected to be complete by Q2, FY22.



— Initial state: no waves migrated
— Mid state: some waves migrated
— End state: all waves migrated



15.5 Tbps



40




Core operations: ESnet has delivered above its specified reliability metric of 99.9% throughout the pandemic.

ESnet notable science highlights, FY 2021

Climate
Modelling

Key Climate Change projections in the IPCC Assessment Report 6.


ESnet enabled massive transfer of CMIP6 global model data from the Earth System Grid Federation (ESGF) to DOE's NERSC supercomputing system. Analysis produced IPCC's projections of future climate change.



Health &
BioScience

Near-real time imaging of the SARS-CoV-2 virus proteases.


ESnet's high-speed link from DOE's Linac Coherent Light Source (LCLS) to NERSC allowed remote collaborators to rapidly examine SARS-CoV-2 viral structure, contributing key data for next-gen COVID-19 treatments¹.



Astronomy &
Astrophysics

Analysis of Dark Energy's effect on expansion of the universe.


ESnet's robust on-demand services for petascale data transfer between DOE supercomputing centers enabled development of mock sky catalogs to understand Dark Energy Spectroscopic Instrument (DESI) data².



Networking
R&D

In-Network Caching and Data Transfer Node as a Service (DTNaaS)

ESnet successfully demonstrated a 3-fold reduction in network usage and significant lower data-fetch latencies via *in-network caching*, and a new design architecture for *data transfer service virtualization*³.



1. Durdagi et al. 2020. Cold Spring Harbor Laboratory. <https://doi.org/10.1101/2020.09.09.287987>.
2. Ruiz-Macias et al. 2020. arXiv [astro-ph.GA], <http://arxiv.org/abs/2010.11283>.
3. Zettar zx Evaluation for ESnet DTNs, <https://www.es.net/assets/Uploads/zettar-zx-dtn-report.pdf>.

NERSC operations and project highlights, FY 2021



Safe, successful facility upgrade and system Phase 1 installation, during a pandemic year, are significant accomplishments that required many changes to work schedules and protocols, with many restrictions to on-site access and team-work.

- Facility upgrade reached substantial completion in January.
 - The upgrade provides an additional 12.5 MW of power, along with associated cooling and management infrastructure.
- Physical placement of Phase I system completed in April.
 - The NERSC and HPE teams continue to integrate, configure, and test the system while observing COVID-19 related protocols.
- Perlmutter dedication ceremony held in May.

Perlmutter listed as:


- **#5 in Top500:** 64.6Pflop/s performance
- **#3 in Top500 HPCG:** 1.91 HPCG-Pflops
- **#6 in Green500:** 25.55 gigaflops/watt power efficiency

NERSC data center awarded **DOE Sustainability Award for Innovative Approach to Sustainability**

NERSC notable science highlights, FY 2021

Building Innovative Nanotubes for Energy Applications

Scientists have learned how to control “excitons” in nanotubes to reduce energy losses and tune the colors emitted by LEDs. *Chemistry of Materials*



Land Ice Contributions to 21st Century Sea Level Rise

Scientists quantify the effect of melting land ice on the global sea level.


Nature



Machine Learning Technique Adds Depth, Breadth & Speed for Sky Survey Analysis

Self-supervised representation learning overcomes shortcomings of existing methods for enabling discovery.



Astrophysical Journal Letters



Collisions of Light Produce Matter/Antimatter by Colliding Packets of Light

Analysis of data from the STAR experiment has produced definitive evidence for this phenomena predicted more than 80 years ago.

Physical Review Letters




Scientists propose source of unexplained solar jets

Coronal jets, which can affect satellites and power grids, could be caused by globs of plasma emerging from the sun in ball shapes that resemble magnetic shapes known as spheromaks.

Physics of Plasmas

Breakthrough in Efficiency of Membranes for Batteries, Fuel Cells and Water Purification

Researchers have designed a membrane with “solvation cages” that can greatly speed up the flow of the desired molecules through an interface.



Nature

OLCF operations and project highlights, FY 2021



Summit

- Operated Summit, the Nation's most capable HPC - sustaining US leadership
- #1 on Top 500 list, June 2018-June 2020, #2 on the current list.
- 200 petaflop IBM AC922 system, links more than 27,000 NVIDIA Volta GPUs with more than 9,000 IBM Power9 CPUs to provide unprecedented opportunities for the integration of artificial intelligence (AI) and scientific discovery
- New high memory racks added for COVID applications.
- Summit will continue to support INCITE, ALCC, DD and COVID applications in FY 2022



Frontier

- A focus on reducing technical risks put the Frontier project on path to deliver exascale capabilities in calendar year 2021 – ensuring DOE delivers on Agency Priority Goal.
- Site Prep complete ahead of schedule.
- Installation of Frontier hardware started in August 2021 and will be complete in October 2021.
- Frontier will deliver approximately 1.5 ExaFlops of peak double precision performance and nearly xx EF of theoretical artificial intelligence (AI) performance
- Will be available for testing and porting applications in 2022.
- Competitive allocations begin July 2022 (ALCC, priority access; General access through INCITE – January 2023)



Core Operations


- OLCF exceeded target metrics for Agency Priority Goal, system availability, INCITE hours delivered, and capability hours delivered.
- Effectively managed CAREs Act funding to support numerous COVID research teams.
- Maintained uninterrupted Summit operations and Frontier site preparation during the pandemic without compromising safety.
- Developed secure container, Citadel, for sensitive personal health data.
- Maintained user access to resources for data processing and visualization, ARM testbed, and Quantum computing.




OLCF notable science highlights, FY 2021

Novel understanding of SARS-CoV-2 and a new method for studying disease

(ACM Gordon Bell Special Prize) Imaging techniques can provide snapshots of viruses such as SARS-CoV-2, but these fall short of capturing the dynamic movements of viral proteins. Researchers have built a first-of-its-kind workflow based on AI and run on the OLCF to simulate the virus's spike protein in numerous environments—the most comprehensive simulation of the virus performed to date.



A snapshot of a visualization of the SARS-CoV-2 viral envelope comprising 305 million atoms. Image Credit: R. Amaro, UC San Diego; A. Ramanathan, ANL



An organic solvent (yellow) and water (blue) separate & form nanoclusters on the hydrophobic and hydrophilic sections of cellulose (green), driving the efficient deconstruction of biomass. Credit: M. Lehman/ORNL.

Production of Biofuels and Bioproducts

The mechanism by which solvents drive efficient biomass deconstruction is not understood. Simulations suggest the success of pretreatment lies in the formation of tetrahydrofuran-rich nanoclusters on the nonpolar cellulose surfaces and hydrophobic lignin and equivalent water-rich nanoclusters on polar cellulose surfaces. This finding can be employed to maximize deconstruction of biomass and facilitate the separation of its components for upgrading to energy and materials.




Weather Forecasting's Exascale Future

Researchers from ECMWF and ORNL achieved the first global simulation of the Earth's atmosphere at a 1-square-kilometer average grid-spacing for a full 4-month season. The milestone improved the resolution of the "European Model," which currently operates at 9-kilometer grid-spacing for routine weather forecast operations. It also serves as the first step in an effort to create multi-season atmospheric simulations at high resolution, pointing toward the future of weather forecasting—one powered by exascale supercomputers.

Meteosat Second Generation satellite image verifying simulation accuracy. Image courtesy ECMWF

Witnessing Our Cosmic Dawn


The cosmic Dawn simulation project used the OLCF to model the Universe in its first billion years during the Epoch of Reionization (EoR) as it changed from a globally neutral to ionized state—a transition driven by the first stars to form—in proto-galaxies. As such, Cosmic Dawn (CoDa) III, follows the structuration of matter under the influence of gravity, pulling gas into clumps and compressing them until they become star-forming galaxies.




Enlarged image of a subregion of CoDa III 8h-1cMpc across, containing two large galaxy progenitors (similar to the Milky Way's Andromeda) and a galaxy group progenitor (similar to Centaurus A).

Advanced Turbine Engine Simulation

General Electric (GE) designs gas turbine jet engines that are world-leading in efficiency, emissions, and durability. A GE team used the OLCF to complete first-of-their-kind 3D flow simulations that are providing breakthrough insights with unprecedented speed and exceptional detail and accuracy and the potential to influence the engine design process.



higher-resolution model of turbulent flow (bottom) made possible using the Summit supercomputer. Image Credit: GE Research



New Pathway for COVID-19 Inflammatory Response

In the lung fluid cells of COVID-19 patients, researchers found an increased expression of enzymes that can trigger the production of bradykinin and a decreased expression of enzymes that would break it down. They also found that some genes in the lung cells increased the production of hyaluronic acid, a substance that can trap around 1,000 times its own weight in water to form a hydrogel. This may point to new drug targets worth exploration.

A normal blood vessel, top, is compared with a blood vessel affected by excess bradykinin, which permits fluid, shown in yellow, to leak out and allows immune cells, shown in purple, to squeeze their way out of blood vessels. Credit: J Smith/ORNL