

Highlights of High Performance Computing to be Presented by the U.S. Department of Energy at SC23

Researchers from across the U.S. Department of Energy (DOE) National Laboratory system will showcase their latest findings in computing at [SC23, the International Conference for High-Performance Computing, Networking, Storage and Analysis](#) taking place on Nov. 12–17 in Denver.

The DOE has been a global leader in high-performance computing for decades, overseeing the development of many of the world's most powerful supercomputers. The significant computing resources within the DOE have enabled novel discoveries in energy, materials sciences, artificial intelligence, and many other scientific disciplines. These resources resulted in multiple Gordon Bell Awards over the decades.

Attendees at SC23 can learn more about these awards and more at the DOE booth (#243). Don't miss the featured talks, technical demonstrations, and roundtable discussions happening at the booth on Nov. 14-16. Topics include exascale systems, quantum computing, digital twins and more.

Scheduled speakers and demos are noted below. Visit the DOE booth website (scdoe.info) for the latest information.

Featured Speaker Schedule:

Tuesday, Nov. 14

- 10: 45 a.m. Stephen Lin, Sandia National Laboratories (SNL)**
“Performant simulations of laser-metal manufacturing processes using adaptive, interface-conformal meshing techniques”
Parts manufactured using laser-metal manufacturing techniques, such as laser welding and laser powder bed fusion additive manufacturing, are highly sensitive to the process details, which govern the formation of defects and distortions.
- 11:30 a.m. Moderator: Meifeng Lin, Brookhaven National Laboratory (BNL)**
Panelists: Kerstin Kleese van Dam, BNL; Eric Stahlberg, Frederick National Laboratory for Cancer Research; and Peter Coveney, University College London
“Digital Twins and the Rise of the Virtual Human”
Biomedical Digital Twins, an innovative fusion of medical science and information technologies, are poised to transform medicine and improve patient care. These new approaches to medicine are just beginning to impact the fundamental understanding of biology and influence technologies available for preventing, detecting, diagnosing, and treating disease.
- 1 p.m. Dominic Manno, Los Alamos National Laboratory (LANL)**
“Leveraging Computational Storage for Simulation Science Storage System Design”
LANL has been engaged in how to exploit computation near data storage devices for simulation analysis. The activities in this area will be outlined with some early results and directions for the future will be covered.

- 1:45 p.m.** **Lori Diachin, Lawrence Livermore National Laboratory (LLNL); Andrew Siegel, Argonne National Laboratory (ANL); Michael Heroux, SNL; Richard Gerber, National Energy Research Scientific Computing Center (NERSC)**
“Delivering a Capable Exascale Computing Ecosystem - for Exascale and Beyond”
A close look at the massive co-design, collaborative effort to build the world's first capable computing ecosystem for exascale and beyond. With the delivery of the U.S. Department of Energy’s (DOE’s) first exascale system, Frontier, in 2022, and the upcoming deployment of Aurora and El Capitan systems by next year, researchers will have the most sophisticated computational tools at their disposal to conduct groundbreaking research.
- 2:30 p.m.** **Sudip Dosanjh, Lawrence Berkeley National Laboratory (LBL)**
“The Next 50 Years: How NERSC is Evolving to Support the Changing Mission Space and Technology Landscape”
As NERSC heads into 2024 and its 50th anniversary, the center already has its sights set on the next 50 years. This is second nature for NERSC, which is always looking ahead to what is coming next in science, research, and technology as well as planning, installing, and refining its systems and services in response - and anticipation.
- 3:15 p.m.** **Mahantesh Halappanavar, Pacific Northwest National Laboratory (PNNL)**
“ExaGraph: Graph and combinatorial methods for enabling exascale applications”
Combinatorial algorithms in general and graph algorithms in particular play a critical enabling role in numerous scientific applications. However, the irregular memory access nature of these algorithms makes them one of the hardest algorithmic kernels to implement on parallel systems.
- 4 p.m.** **Marc Day, National Renewable Energy Laboratory (NREL)**
“Adaptive computing and multi-fidelity strategies for control, design and scale-up of renewable energy applications”
We describe our ongoing research in adaptive computing and multi-fidelity modeling strategies. Our goal is to use a combination of low- and high-fidelity simulation models to enable computationally efficient optimization and uncertainty quantification.

Wednesday, Nov. 16

- 10:45 a.m.** **Rachana Ananthkrishnan and Tom Uram, ANL**
“Nexus: Pioneering new approaches to integrating scientific facilities, supercomputing and data”
Nexus (anl.gov/nexus-connect), an Argonne initiative, connects experimental, computing, and storage facilities at Argonne and beyond into an interconnected scientific infrastructure. In so doing, it allows researchers to engage unique DOE resources at Argonne and beyond to tackle challenging questions, making the difficult easy, and the previously impossible conceivable.
- 11:30 a.m.** **Ammar Hakim, Princeton Plasma Physics Laboratory (PPPL)**
“Algorithms for Relativistic Kinetic and Fluid Simulations of Plasmas”
In this talk I will present recent innovations in the design of novel algorithms for relativistic simulations of plasmas. These span the regimes from fluids to kinetic continuum models,

with applications ranging from laser-plasma interactions to the plasma environment around black holes.

1 p.m. Trevor Wood, Oak Ridge National Laboratory (ORNL)

“Exascale use and impact on the future of flight”

To evaluate the potential efficiency delivered by an open fan architecture while simultaneously reducing noise levels, designers need to understand airflow behavior around the blades of the engine, including the complex physics of turbulence, mixing, and potential for flow for separation.

1:45 Sarom S. Leang, Ames National Laboratory (Ames)

p.m. “Enabling QM Study of Actual Length Scale Chemistry and Material Science”

The GAMESS (General Atomic and Molecular Electronic Structure System) software is a versatile computational chemistry package that facilitates understanding and predicting molecular structures and reactions from the quantum mechanical viewpoint.

2:30 Jamie Bramwell and Aaron Skinner, LLNL

p.m. “High-Performance Multiphysics Applications at LLNL”

High-performance multiphysics codes at LLNL support a variety of applications essential for predictive science at the NNSA. Our codes are large, complex, tailored to our applications, and represent decades of investment.

3:15 Graham Heyes, Jefferson Lab (JLab)

p.m. “Research projects at JLab evaluating technologies for the High Performance Data Facility”

Several research and development projects are underway at JLab in partnership with other labs and ASCR facilities, particularly ESnet, to evaluate technologies and inform the conceptual design of the High Performance Data Facility (HPDF). This will be a new scientific user facility specializing in advanced infrastructure for data-intensive science.

4 p.m. Charlie Catlett, ANL

Prasanna Balaprakash, ORNL

“Community Collaboration Around Large AI Models for Science”

A new international effort is bringing together researchers interested in creating large-scale generative AI models for science and engineering problems with those who are building and operating large-scale computing systems. The goal is to create a global network of resources and expertise to facilitate the next generation of AI and maximize the impact of projects by avoiding duplication of effort.

Thursday, Nov. 16

10:45 Damien Lebrun-Grandie, ORNL

a.m. Christian Trott, SNL

“Kokkos ecosystem - Sustaining performance portability at the exascale era”

We have entered the exascale era of high-performance computing (HPC), and with it comes the challenge of writing software that can achieve high performance on a wide variety of heterogeneous architectures. The Kokkos Ecosystem is a performance portability solution

which addresses that challenge through a single source C++ programming model that follows modern software engineering practices.

Technical Demonstrations

Monday, Nov. 13

- 7:00 p.m.** **Demo Station 1:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
“AppSysFusion: “Always-on Monitoring” on Sandia HPC Systems”
- Demo Station 2:** Prasana Balaprakash, Feiyi Wang, Sajal Dash, Junqi Yin, Dan Lu, Ashwin Aji, Leon Song (ORNL)
“Accelerating scientific discoveries with DeepSpeed for Science and AMD-powered Frontier exascale supercomputer”
- 8:00 p.m.** **Demo Station 1:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
“AppSysFusion: “Always-on Monitoring” on Sandia HPC Systems”
- Demo Station 2:** Prasana Balaprakash, Feiyi Wang, Sajal Dash, Junqi Yin, Dan Lu, Ashwin Aji, Leon Song (ORNL)
“Accelerating scientific discoveries with DeepSpeed for Science and AMD-powered Frontier exascale supercomputer”

Tuesday, Nov. 14

- 10:00 a.m.** **Demo Station 1:** Christian Trott; Bruno Turcksin; Daniel Arndt; Nevin Liber; Rahul Kumar Gayatri; Nevin Liber; Sivasankaran Rajamanickam; Luc Berger-Vergiat (ANL, LBL, ORNL, SNL)
“Achieving performance portability with Kokkos”
- Demo Station 2:** Yao Xu; Gene Cooperman; Rebecca Hartman-Baker (LBL)
“Transparent Checkpointing on Perlmutter for Long-Running Jobs”
- 11:00 a.m.** **Demo Station 1:** Hannah Parraga, Michael Prince (ANL)
“Empowering Scientific Discovery at the APS with Integrated Computing”
- Demo Station 2:** Thomas Applencourt; Abhishek Bagusetty (ANL)
“oneAPI and SYCL”
- 12:00 p.m.** **Demo Station 1:** Mariam Kiran; Anastasiia Butko; Ren Cooper; Imtiaz Mahmud, Nirmalendu Patra; Matthew Verlie (ORNL, LBL)
“5G on the Showfloor”
- Demo Station 2:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
“AppSysFusion: “Always-on Monitoring” on Sandia HPC Systems”

- 1:00 p.m.** **Demo Station 2:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"
- 2:00 p.m.** **Demo Station 1:** Brad Richardson; Magne Haveraaen (LBL)
"Fortran generics for 202y"
- Demo Station 2:** Jean Luca Bez; Hammad Ather; Suren Byna; John Wu (LBL)
"Drishti: Where is the I/O bottleneck?"
- 3:00 p.m.** **Demo Station 1:** Christine Simpson, Tom Uram, Rachana Ananthakrishnan, David Schissel,
Hannah Parraga, Michael Prince
"Flexible cross-facility experimental data analysis at ALCF"
- Demo Station 2:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"
- 4:00 p.m.** **Demo Station 2:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"
- 5:00 p.m.** **Demo Station 2:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"

Wednesday, Nov. 15

- 10:00 a.m.** **Demo Station 1:** Caetano Melone (LLNL)
"Dynamically Allocating Resources for Spack CI Builds"
- Demo Station 2:** Marco Minutoli (PNNL)
"Maximizing the Influence at the ExaScale with Ripples"
- 11:00 a.m.** **Demo Station 1:** Caetano Melone (LLNL)
"Dynamically Allocating Resources for Spack CI Builds"
- Demo Station 2:** Imran Latif (BNL)
"CPU performance and power Optimization"
- 12:00 p.m.** **Demo Station 1:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"
- Demo Station 2:** Mariam Kiran; Anastasiia Butko; Ren Cooper; Imtiaz Mahmud, Nirmalendu Patra; Matthew Verlie (ORNL, LBL)
"5G on the Showfloor"

- 1:00 p.m.** **Demo Station 1:** Mariam Kiran; Muneer Alshowkan; Brian Williams; Joseph Chapma (ORNL)
"Quantum Networks a Reality"
- Demo Station 2:** Christine Simpson, Tom Uram, Rachana Ananthakrishnan, David Schissel, Hannah Parraga, Michael Prince (ANL)
"Flexible cross-facility experimental data analysis at ALCF"
- 2:00 p.m.** **Demo Station 1:** Yatish Kumar (LBL)
"Open Source ESnet P4 FPGA smartNIC"
- Demo Station 2:** Imran Latif (BNL)
"CPU performance and power Optimization"
- 3:00 p.m.** **Demo Station 1:** Flavio Castro; Joaquin Chung; Se-young Yu (ANL)
"SciStream: Architecture and Toolkit for Data Streaming between Federated Science Instruments"
- Demo Station 2:** Sam Wellborn; Bjoern Enders; Peter Ercius; Chris Harris; Deborah Bard (LBL)
"Live Streaming of Large Electron Microscope Data to NERSC"
- 4:00 p.m.** **Demo Station 1:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"
- Demo Station 2:** Charles Shiflett (LBL)
"Long-distance high-speed data transfer with EScp"
- 5:00 p.m.** **Demo Station 1:** Ann Gentile, Jim Brandt, Benjamin Schwaller, Tom Tucker (SNL)
"AppSysFusion: "Always-on Monitoring" on Sandia HPC Systems"

Thursday, Nov. 16

- 10:00 a.m.** **Demo Station 1:** Christian Trott; Bruno Turcksin; Daniel Arndt; Nevin Liber; Rahulkumar Gayatri; Nevin Liber; Sivasankaran Rajamanickam; Luc Berger-Vergiat (ANL, LBL)
"Kokkos ecosystem beyond performance portable code"
- Demo Station 2:** Christian Mayr (SNL)
"The SpiNNaker2 Neuromorphic Computing Architecture"
- 11:00 a.m.** **Demo Station 1:** Flavio Castro; Joaquin Chung; Se-young Yu (ANL)
"SciStream: Architecture and Toolkit for Data Streaming between Federated Science Instruments"
- Demo Station 2:** Christian Mayr (SNL)
"The SpiNNaker2 Neuromorphic Computing Architecture"

1:00 p.m. **Demo Station 1:** Mariam Kiran; Muneer Alshowkan; Brian Williams; Joseph Chapma (ORNL)
“Quantum Networks a Reality”

2:00 p.m. **Demo Station 1:** Mariam Kiran; Muneer Alshowkan; Brian Williams; Joseph Chapma (ORNL)
“Quantum Networks a Reality”

Learn more about the national lab activities at the [DOE booth website](#). Find details about the SC23 conference — including a [full schedule of events](#) — at the [SC23 website](#).

Follow the national labs conference participation on social media at @natlabshpc and #DOEatSC, #SC23, #HPCAccelerates.