




Programming for Perlmutter


Jack Deslippe

NESAP Lead
Application Performance Group Lead


NERSC is the mission High Performance Computing facility for the DOE SC



7,000 Users
800 Projects
700 Codes
2000 NERSC citations per year




Simulations at scale



Data analysis support for
DOE's experimental and
observational facilities

Photo Credit: CAMERA



2 -



BERKELEY LAB
Bringing Science Solutions to the World



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Lawrence Hall of Science

NERSC Virtual Tours - Monthly in 2021 <https://visits.lbl.gov/>



Berkeley Lab

**Advanced Light Source
(ALS)**


**Wang Hall
NERSC
ESnet**

**University of
California,
Berkeley**



NERSC has a dual mission to advance science and the state-of-the-art in supercomputing

- We collaborate with computer companies years before a system's delivery to deploy advanced systems with new capabilities at large scale
- We provide a highly customized software and programming environment for science applications
- We are tightly coupled with the workflows of DOE's experimental and observational facilities – ingesting tens of terabytes of data each day
- Our staff provide advanced application and system performance expertise to users

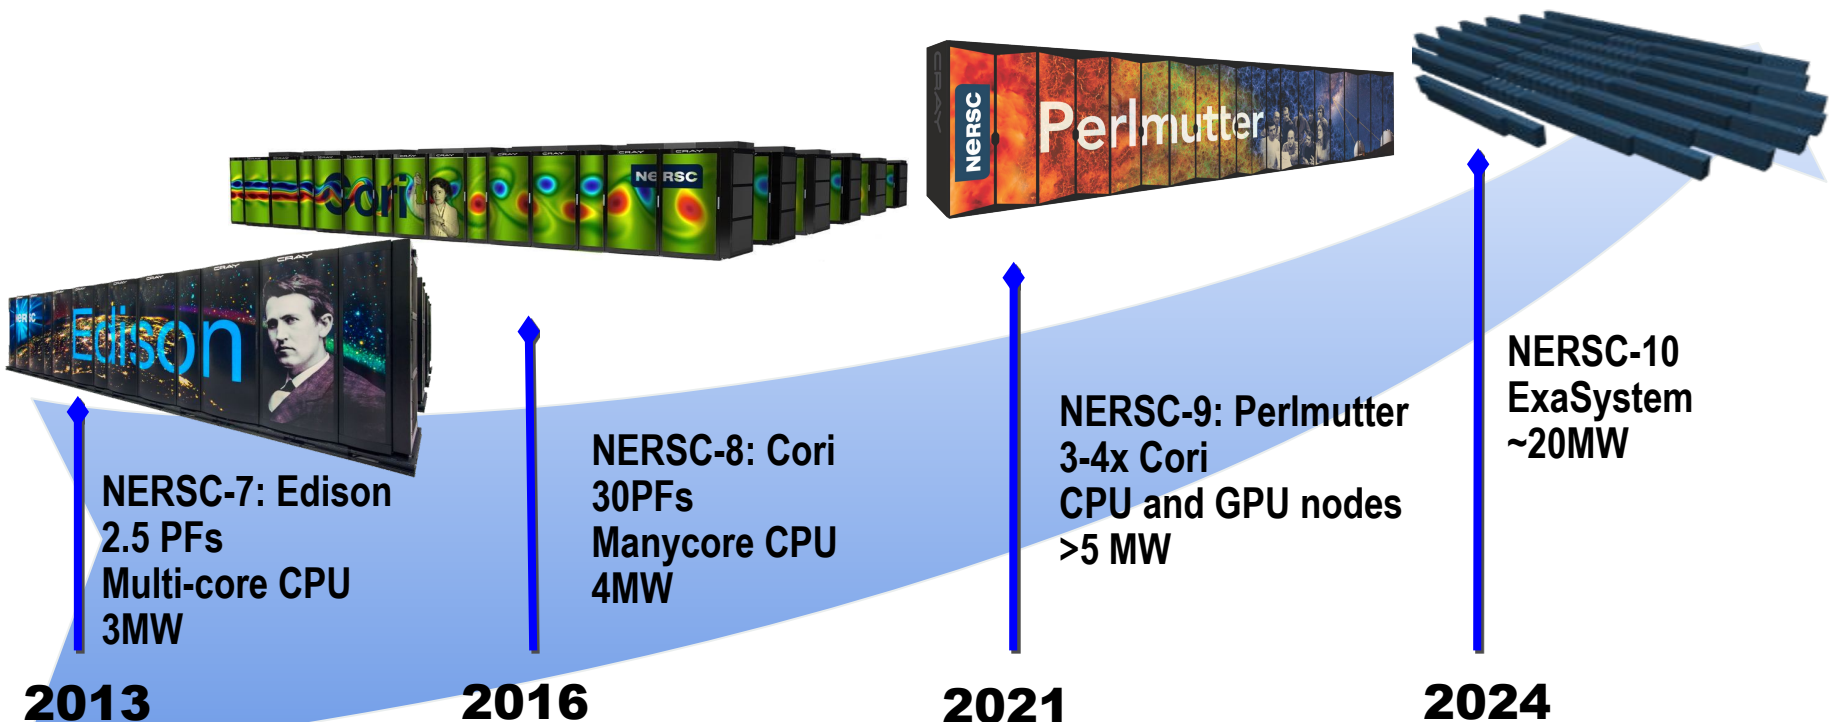



NERSC-9 will be named after Saul Perlmutter

- Winner of 2011 Nobel Prize in Physics for discovery of the accelerating expansion of the universe.
- Supernova Cosmology Project, lead by Perlmutter, was a pioneer in using NERSC supercomputers to combine large scale simulations with experimental data analysis
- Login “saul.nersc.gov”



NERSC Systems Roadmap





- HPE Cray System with 3-4x capability of Cori
- GPU-accelerated (GPU/CPU) and CPU-only nodes
- HPE Cray Slingshot high-performance network
- All-Flash filesystem
- Application readiness program (NESAP)


Phase I: Arrived on February 22, 2021

- 1,536 GPU-accelerated nodes
- 1 AMD “Milan” CPU + 4 NVIDIA A100 GPUs per node
- 256 GB CPU memory and 40 GB GPU high BW memory
- 35 PB FLASH scratch file system
- User access and system management nodes

Phase II Addition: Arrives This Fall

- 3,072 CPU only nodes
- 2 AMD “Milan” CPUs per node
- 512 GB memory per node
- Upgraded high speed network
- CPU partition will match or exceed performance of entire Cori system

Perlmutter Phase 1 Arrival



February 22, 2021

Perlmutter Programming Environments

	GPU Support	FORTRAN/C /C++	OpenACC 2.x	OpenMP 5.x	CUDA	Kokkos / Raja	Cray MPI
NVIDIA							
CCE							
GNU				(Community Effort)			
LLVM				(Community Effort)			

Vendor Supported

NERSC Supported

OpenMP NRE partnership with NVIDIA

- Agreed upon subset of OpenMP features to be included in the NVIDIA (was PGI) compiler
- OpenMP test suite created with micro-benchmarks, mini-apps, and the ECP SOLLVE V&V suite
- 5 NESAP application teams partnering with NVIDIA to add OpenMP target offload directives
- The production OpenMP offload compiler was released in April 2021.
 - Beta compiler was delivered in November 2020.
 - Offload capability was integrated into the NVIDIA HPC SDK compiler suite in December 2020.

A screenshot of the Berkeley Lab Computing Sciences website. The header includes the Berkeley Lab logo, the text "BERKELEY LAB COMPUTING SCIENCES" and "LAWRENCE BERKELEY NATIONAL LABORATORY", and the U.S. Department of Energy logo. Navigation links include Home, About, News & Media, Seminars, Careers, Awards, Safety, and For Staff. A search bar is also present. The main content area features a news article titled "NERSC, NVIDIA to Partner on Compiler Development for Perlmutter System" dated March 21, 2019. The article text states: "The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (Berkeley Lab) has signed a contract with NVIDIA to enhance GPU compiler capabilities for Berkeley Lab's next-generation Perlmutter supercomputer." It further details the partnership, mentioning the DOE's announcement in October 2018 and the specifications of the Perlmutter supercomputer, including its use of NVIDIA GPUs with Tensor Core technology and AMD EPYC CPUs.

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

BERKELEY LAB COMPUTING SCIENCES

A-Z INDEX | PHONE BOOK | CAREERS | SHARE | FOLLOW

Home About News & Media Seminars Careers Awards Safety For Staff

Home » News & Media » News » NERSC, NVIDIA to Partner on Compiler Development for Perlmutter System

NEWS & MEDIA

News

CS In the News
InTheLoop

NERSC, NVIDIA to Partner on Compiler Development for Perlmutter System

MARCH 21, 2019

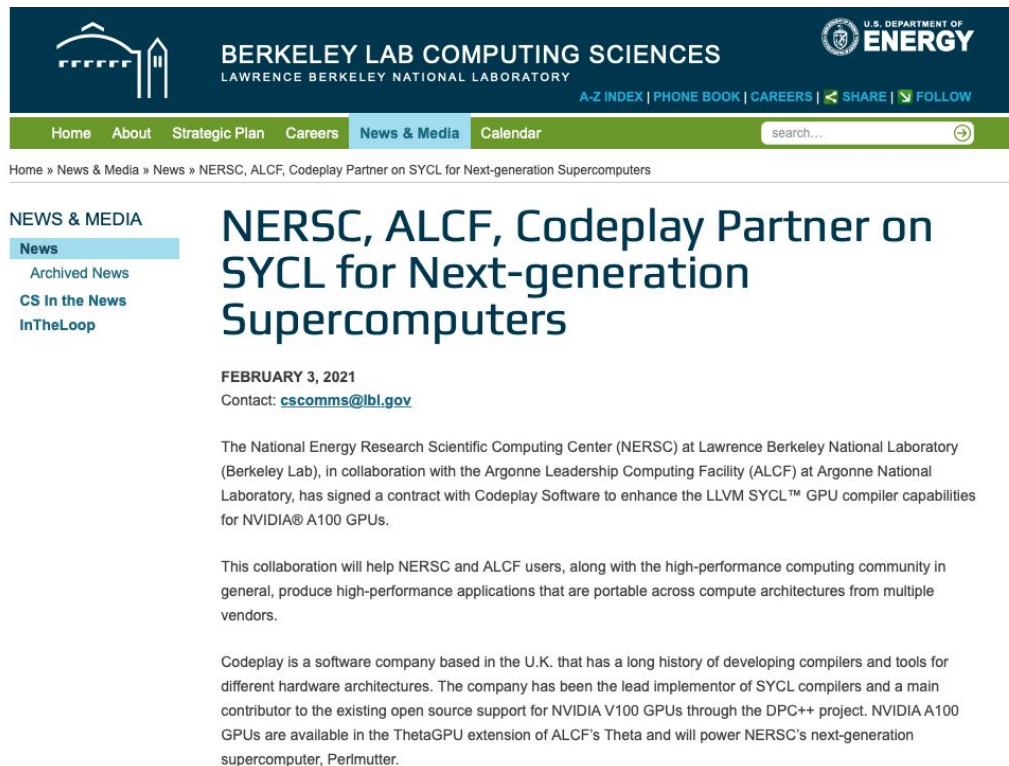
The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (Berkeley Lab) has signed a contract with NVIDIA to enhance GPU compiler capabilities for Berkeley Lab's next-generation Perlmutter supercomputer.

In October 2018, the U.S. Department of Energy (DOE) announced that NERSC had signed a contract with Cray for a pre-exascale supercomputer named "Perlmutter," in honor of Berkeley Lab's Nobel Prize-winning astrophysicist Saul Perlmutter. The Cray Shasta machine, slated to be delivered in 2020, will be a heterogeneous system comprising both CPU-only and GPU-accelerated cabinets. It will include a new Cray system interconnect designed for data-centric computing; NVIDIA GPUs with new Tensor Core technology; CPU-only nodes based on next-generation AMD EPYC CPUs; direct liquid cooling; and an all-flash scratch filesystem that will move data at a rate of more than 4 terabytes/sec.



NERSC, ALCF and Codeplay partnership on SYCL

- Target SYCL 2020 (latest specification) support on Ampere A100 GPUs
- Open LLVM based compiler
- Provides Portability for Apps Developed for Aurora
- Extensions for A100
 - Asynchronous Copy
 - Asynchronous Barrier
 - Tensor core types/ APIs



The screenshot shows the Berkeley Lab Computing Sciences website. The header includes the Berkeley Lab logo, the text "BERKELEY LAB COMPUTING SCIENCES" and "LAWRENCE BERKELEY NATIONAL LABORATORY", and the U.S. Department of Energy logo. Navigation links include "A-Z INDEX", "PHONE BOOK", "CAREERS", "SHARE", and "FOLLOW". A green navigation bar contains "Home", "About", "Strategic Plan", "Careers", "News & Media", and "Calendar", along with a search bar. Below the navigation bar, a breadcrumb trail reads "Home » News & Media » News » NERSC, ALCF, Codeplay Partner on SYCL for Next-generation Supercomputers". The "NEWS & MEDIA" section is active, with sub-links for "News", "Archived News", "CS In the News", and "InTheLoop". The main article is titled "NERSC, ALCF, Codeplay Partner on SYCL for Next-generation Supercomputers" and is dated "FEBRUARY 3, 2021". The contact information is "Contact: cscmms@lbl.gov". The article text states: "The National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (Berkeley Lab), in collaboration with the Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory, has signed a contract with Codeplay Software to enhance the LLVM SYCL™ GPU compiler capabilities for NVIDIA® A100 GPUs." It further explains: "This collaboration will help NERSC and ALCF users, along with the high-performance computing community in general, produce high-performance applications that are portable across compute architectures from multiple vendors." The final paragraph states: "Codeplay is a software company based in the U.K. that has a long history of developing compilers and tools for different hardware architectures. The company has been the lead implementor of SYCL compilers and a main contributor to the existing open source support for NVIDIA V100 GPUs through the DPC++ project. NVIDIA A100 GPUs are available in the ThetaGPU extension of ALCF's Theta and will power NERSC's next-generation supercomputer, Perlmutter."

Many NERSC Applications Actively Working w/ OpenMP Offload

- BerkeleyGW
- Quantum ESPRESSO
- WEST
- GAMESS
- QMCPACK
- LAMMPS
- US QCD Apps
- Kokkos Apps w/ OpenMP Back-end