National Energy Research Scientific Computing Center

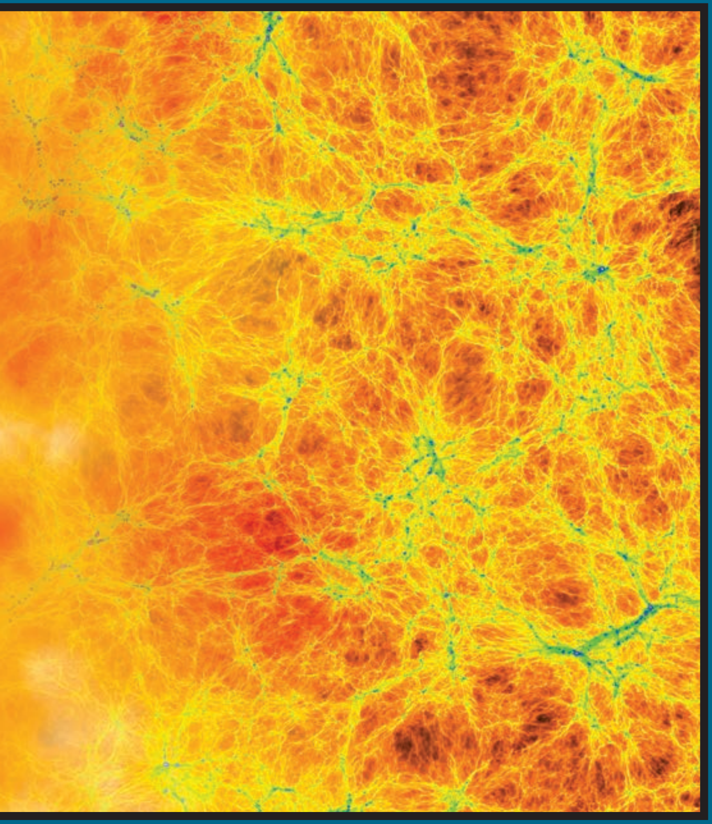

2018 ANNUAL REPORT

U.S. DEPARTMENT OF ENERGY | Office of Science

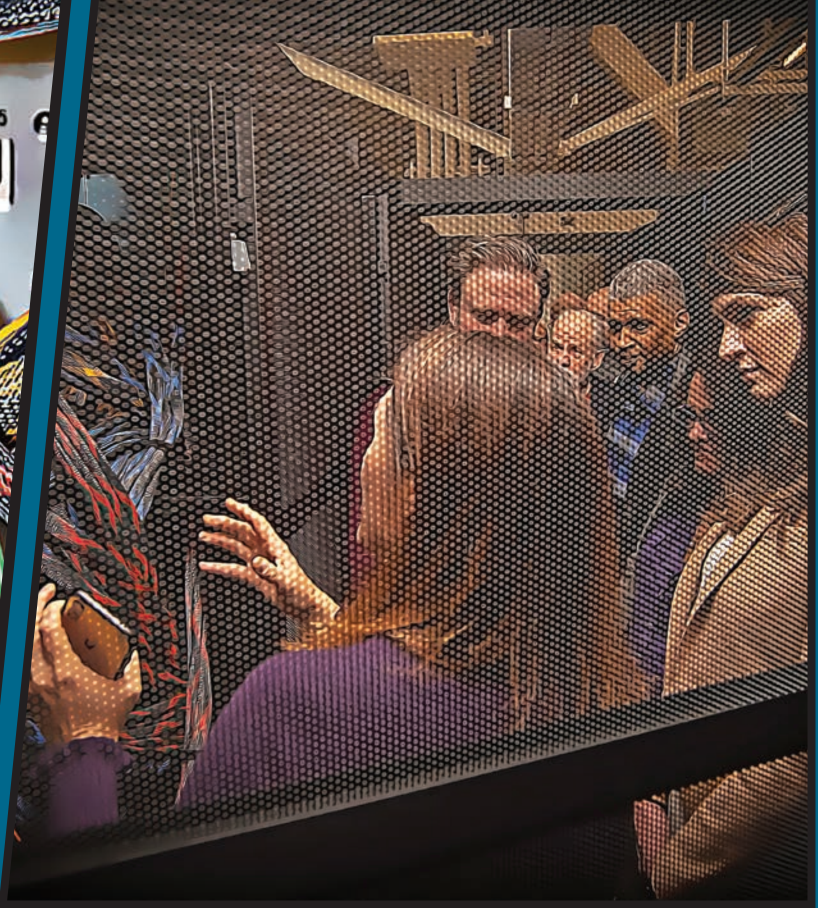

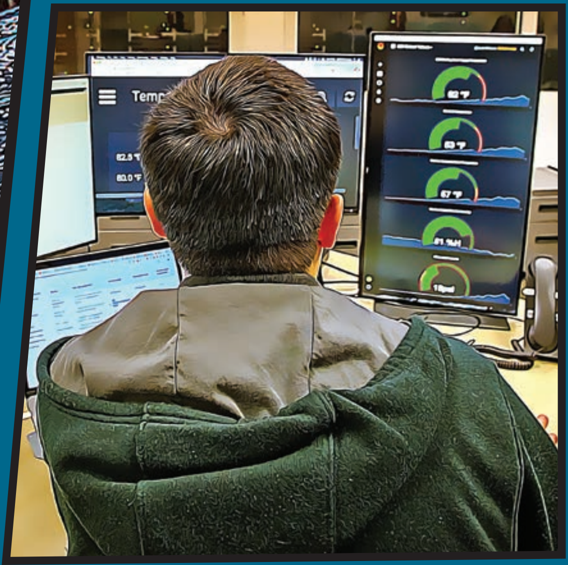

Ernest Orlando Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley, CA 94720-8148

This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231

Image Credits: Thor Swift, Berkeley Lab

Cover Simulations: Andrew Myers, Zarija Lukic

National Energy Research Scientific Computing Center

2018 ANNUAL REPORT

Table of Contents

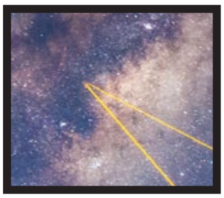

Page **12**

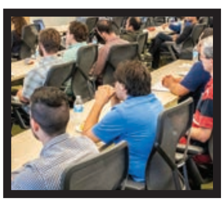

Page **39**

Page **48**

Page **49**

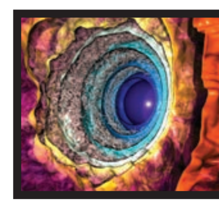

Page **57**

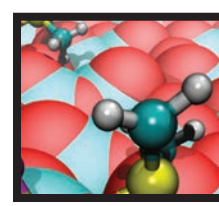

Page **61**

Director’s Note

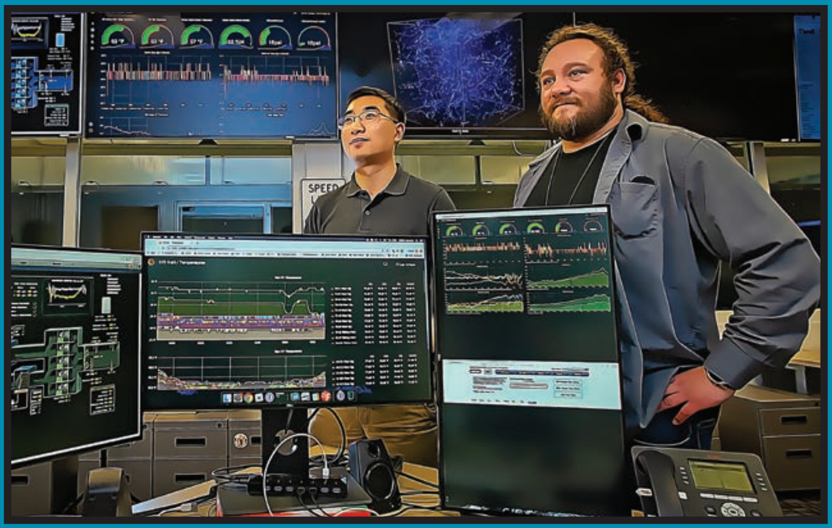

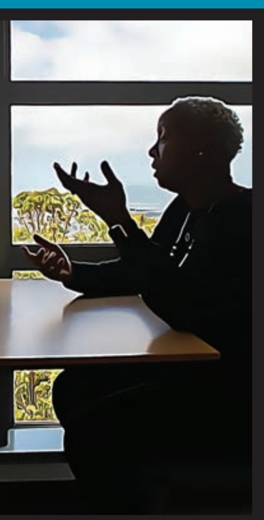

2018 was another year of scientific impact at the National Energy Research Scientific Computing Center, the mission high performance computing facility for the DOE Office of Science. NERSC is managed by Lawrence Berkeley National Laboratory and funded by DOE's Advanced Scientific Computing Research Office. Its mission is to accelerate scientific discovery at the DOE Office of Science through HPC and extreme data analysis.

One of the highlights of the year was the announcement of our next supercomputer, Perlmutter, a heterogeneous Cray Shasta system featuring NVIDIA GPUs, AMD CPUs, the Cray Slingshot interconnect, and an all-flash scratch file system. Named in honor of Nobel Prize-winning Berkeley Lab scientist Saul Perlmutter, it is the first NERSC supercomputer designed with both large-scale simulations and extreme data analysis in mind. Slated to be delivered in late 2020, Perlmutter is optimized for science, and we are collaborating with Cray, NVIDIA, and AMD to ensure that it meets our users' computational and data needs, providing a significant increase in capability and a platform to continue transitioning our broad workload to energy-efficient architectures.

NERSC supports thousands of scientists from all 50 states working on ~700 research projects spanning the range of SC scientific disciplines. Preparing our users for Perlmutter's arrival is of paramount importance. We are also focused on supporting the growing data analysis and machine learning workloads at NERSC, particularly those originating from DOE's experimental and observational facilities. In 2018, NERSC staff continued to collaborate with application teams through the NERSC Exascale Science Application Program to prepare codes for pre-exascale and exascale architectures, with an emphasis on increasing parallelism, data locality, and vectorization. Some NESAP collaborations focusing on data analysis workloads have seen application or kernel performance increases of up to 200x, making it feasible to run them at large scales on HPC platforms. And NERSC staff continue to lead trainings and tutorials to increase user capabilities on our systems.

The impacts of these efforts are showing as the large NERSC user base continues to transition to advanced architectures. Of the ~350 projects that use 99% of the total time at NERSC, in 2018 63% used more than 100,000 hours per month on the energy-efficient manycore Cori KNL nodes, compared to 51% in 2017.

Finding additional parallelism and improving vectorization in applications is key to moving codes to exascale architectures; the next step for users is to identify components of applications that can benefit from GPU acceleration. Toward this end, NERSC has deployed a GPU cabinet on Cori that will be used by NESAP teams to port and prepare applications for Perlmutter, whose GPUs will offer accelerated

compute horsepower, large memory, and high-bandwidth data transfer in the desired power envelope to run simulations, machine learning, and artificial intelligence on data streaming from telescopes, particle accelerators, sensors, microscopes, and more.

This is important because, at present, approximately 25% of NERSC users are part of a DOE experimental facility analyzing data at NERSC, and more than 35% of NERSC Principal Investigators note that the primary role of their project is to analyze data from DOE's experimental facilities, create tools to analyze data from these facilities, or combine large-scale simulation with data analysis from facilities. Engaging with these facilities is key to understanding their requirements and developing capabilities that will support many experiments. To coordinate our efforts, we have built an internal "Superfacility Project" that brings together staff from across the organization to partner with scientists at experimental and observational facilities.

Supporting the advanced analytics requirements of these new workloads is also critical, so NERSC has invested heavily in deploying analytics and machine learning tools and infrastructure at scale. The number of users of these packages has grown dramatically in the past year. NERSC has also built up staff expertise in analytics and machine learning, launching popular tutorials at SC18 and hosting a NERSC "Data Day" training session and Machine Learning for Science summit at Berkeley Lab. In addition, NERSC staff were authors on the International Conference on Machine Learning Best Paper "Graph Neural Networks for IceCube Signal Classification." And 2018 closed out with NERSC staff winning a Gordon Bell prize at the annual SC conference for "Exascale Deep Learning for Climate Analytics," in collaboration with NVIDIA and Oak Ridge Leadership Computing Facility staff.

In 2018, NERSC initiated and completed a number of activities that improved our operational security and efficiency. To reduce the risk of compromised user accounts, NERSC decided to make multi-factor authentication mandatory for the allocation year 2019. This meant converting the bulk of our users from using passwords and certificates to using MFA at every authentication. To reduce the impact on user workflows and patterns of usage, NERSC developed an innovative approach to MFA called SSH Proxy, which allows users to generate an SSH key that can be used to log in for a limited amount of time; implemented an extensive communications campaign; held virtual office hours to help NERSC users get started with MFA; and created thorough documentation.

2018 also saw the launch of a major power and cooling upgrade in Shyh Wang Hall to support the Perlmutter system. When NERSC was planning the additional construction needed to support Perlmutter, one consideration was whether a new mechanical substation would be required for the additional power load of the new system. Using data from the OMNI monitoring data collection system, it was determined that an additional mechanical substation was unnecessary, resulting in a cost savings of $2 million.

And while most NERSC hardware relocated to Wang Hall in 2015, the tape libraries remained in operation at the Oakland Scientific Facility through 2018. Allowing temperature and humidity to vary somewhat in the Wang Hall data center may be energy efficient, but it is harmful to tape. As a result, the existing libraries at OSF could not be relocated to Wang Hall without a solution for environmental containment. NERSC investigated several alternatives and ultimately worked with IBM to deploy a new set of library systems that feature self-contained humidity and temperature controls, providing a suitable environment for tape.

As you can see, 2018 was a busy and exciting year for NERSC. Our staff continues to expand, as does our user base and workload, and we are working hard to meet these needs and challenges as we look ahead to the next generation of supercomputing architectures and applications.

Sudip Dosanjh
NERSC Division Director

NERSC SYSTEMS 2018

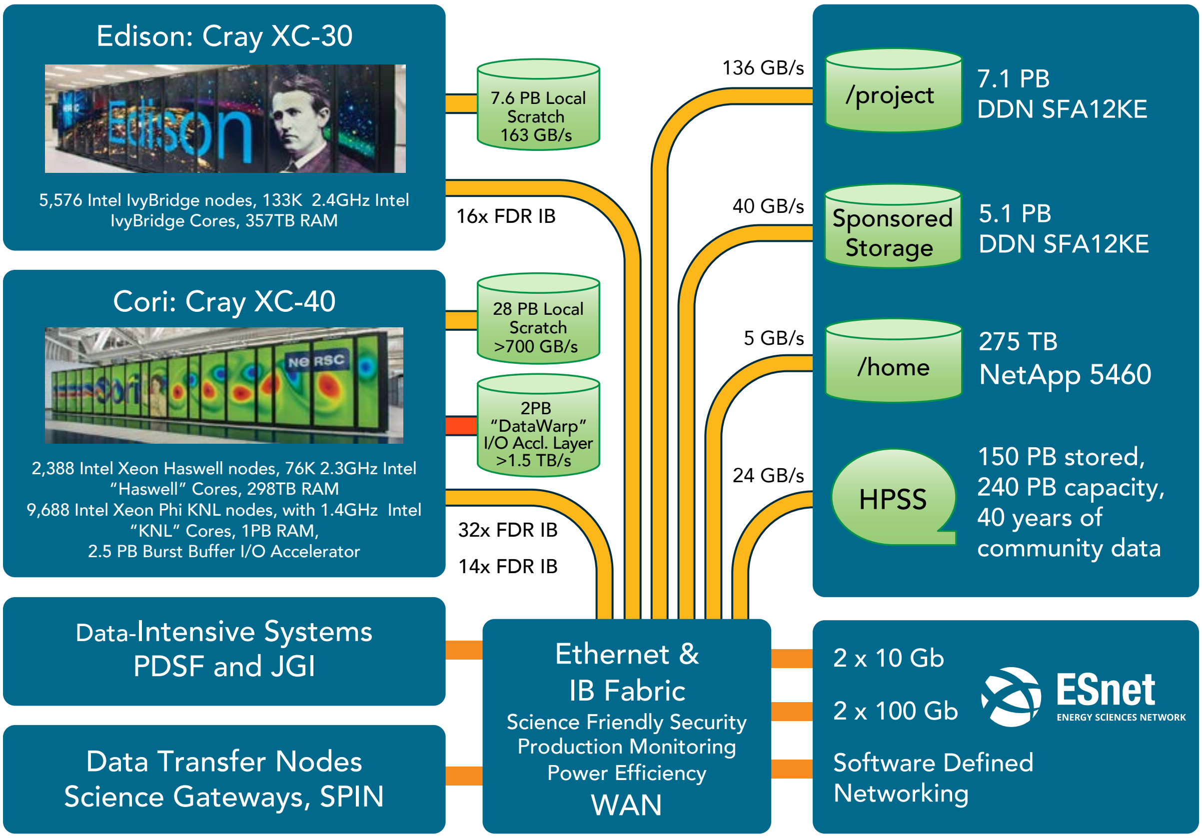

SPECIFICATIONS FOR NERSC 2018 COMPUTATIONAL SYSTEMS

System Name	System Type	Processor Type	Processor Speed (GHz)	Peak Performance	Number of Compute Nodes	Aggregate Memory	Memory Per Node	Inter-connect	Scratch Disk Space
Edison	Cray XC30	Intel Ivy Bridge	2.4	2.57 PF/sec	5,586	357 TB	64 GB	Aries	7.6 PB
Cori Haswell	Cray XC40	Intel Haswell	2.3	1.92 PF/sec	2,388	298 TB	128 GB	Aries	28 PB
Cori KNL	Cray XC40	Intel Knights Landing	1.4	28 PF/sec	9,688	1.09 PB	96 GB	Aries	28 PB

2018 Usage Demographics

BY THE NUMBERS

10,051 Million

Computing Hours Provided in 2018

>7,000 Annual Users FROM **649** Institutions + National Labs

PUBLICATIONS INVOLVING NERSC RESOURCES: **2,584**

2018 NERSC USAGE BY DOE OFFICE OF SCIENCE PROGRAM

(NERSC Hours Charged in Millions)

- **2,147** Basic Energy Sciences
- **1,154** Biological and Environmental Research
- **1,076** High Energy Physics
- **910** Fusion Energy Sciences
- **763** Nuclear Physics
- **166** Advanced Scientific Computing Research

2018 NERSC USAGE BY INSTITUTION TYPE

(NERSC Hours Charged in Millions)

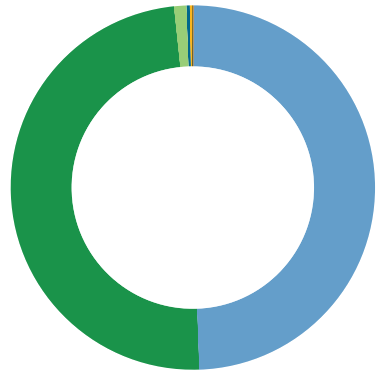

- **4,806** Universities
- **4,748** DOE Labs
- **111** Other Government Labs
- **21** Small Business
- **15** Private Labs
- **6** Industry
- 5 Nonprofits

2018 NERSC USAGE BY DISCIPLINE

(NERSC Hours Charged in Millions)

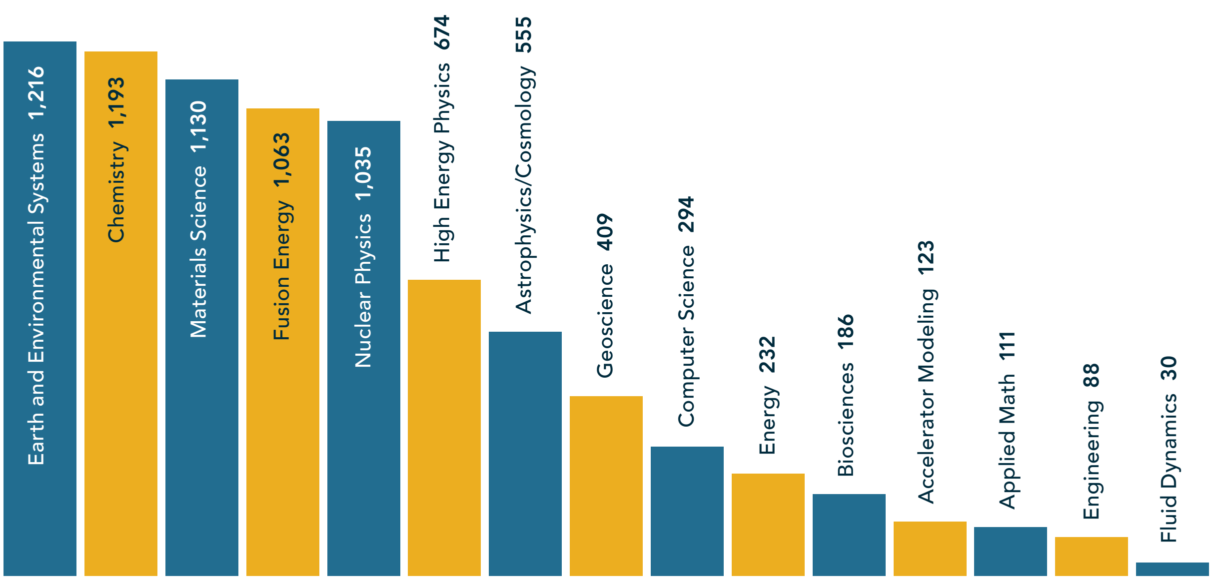

2018 DOE & OTHER LAB USAGE AT NERSC

(NERSC Hours Charged in Millions)

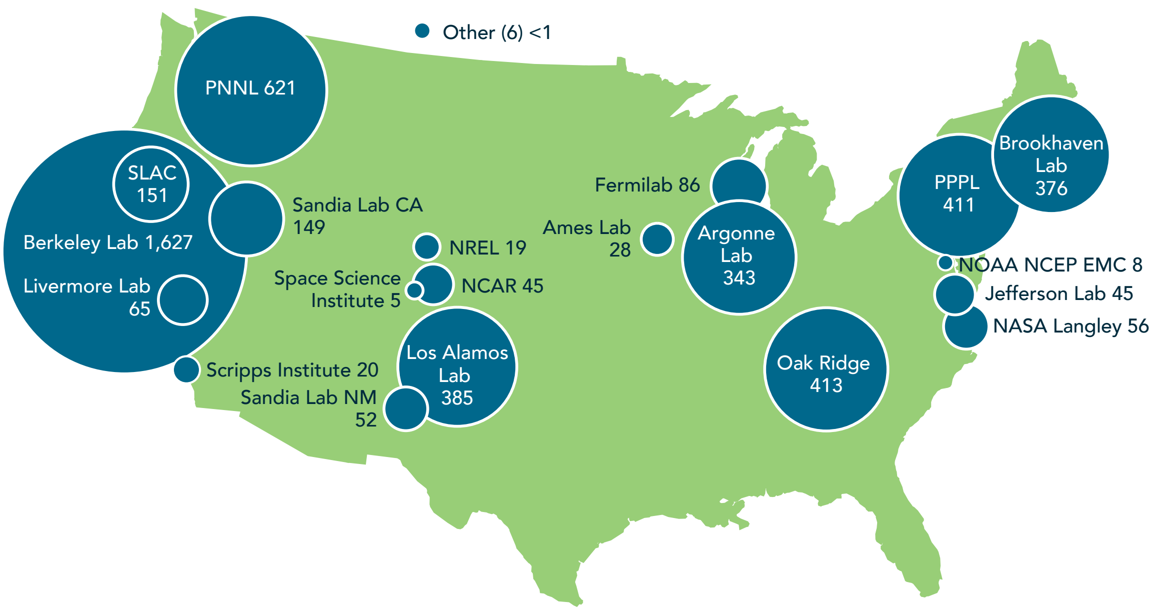

2018 NERSC USERS BY STATE

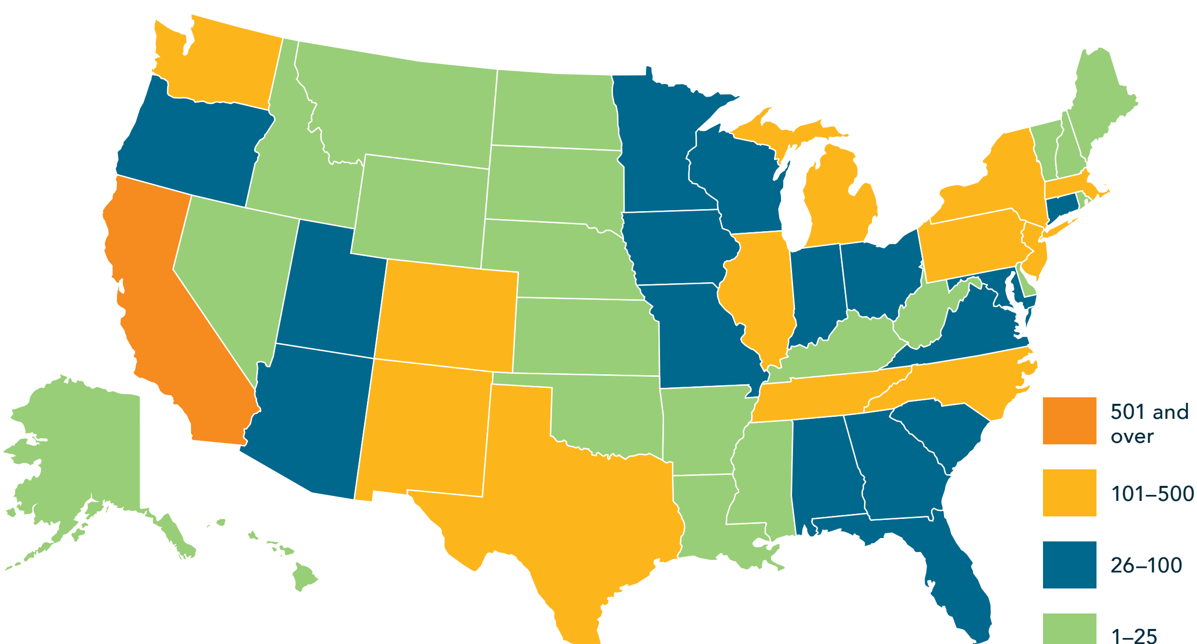

2018 ACADEMIC USAGE AT NERSC

(NERSC Hours Charged in Millions)

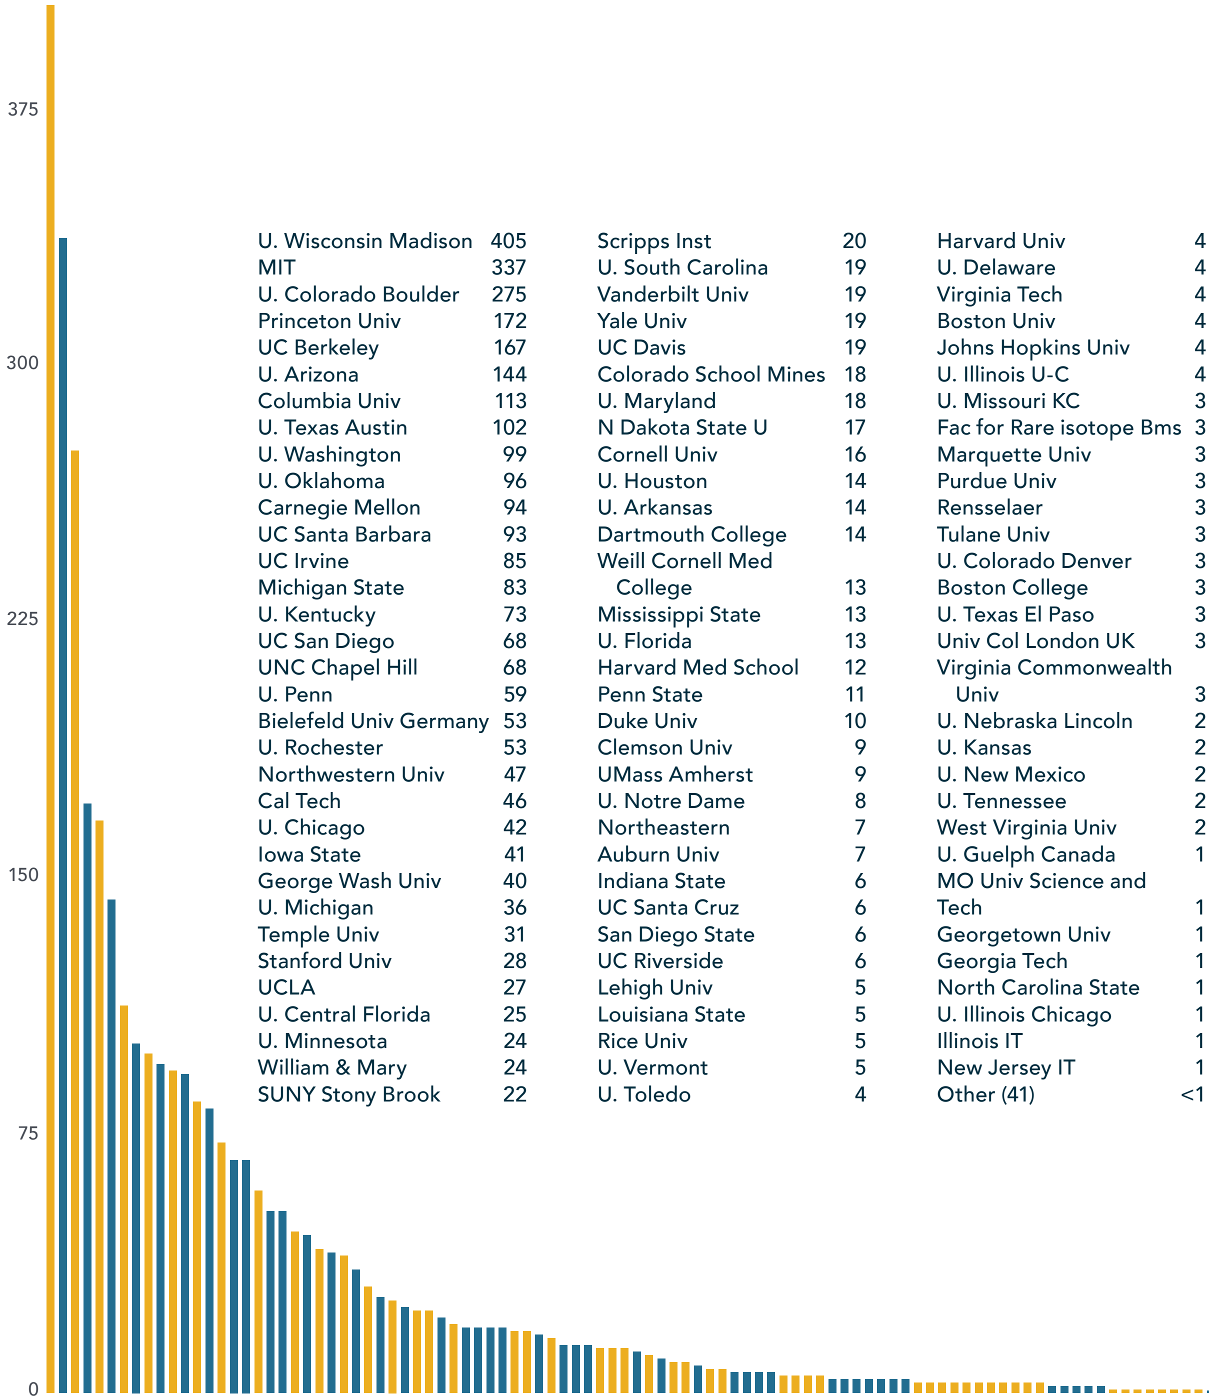

U. Wisconsin Madison	405
MIT	337
U. Colorado Boulder	275
Princeton Univ	172
UC Berkeley	167
U. Arizona	144
Columbia Univ	113
U. Texas Austin	102
U. Washington	99
U. Oklahoma	96
Carnegie Mellon	94
UC Santa Barbara	93
UC Irvine	85
Michigan State	83
U. Kentucky	73
UC San Diego	68
UNC Chapel Hill	68
U. Penn	59
Bielefeld Univ Germany	53
U. Rochester	53
Northwestern Univ	47
Cal Tech	46
U. Chicago	42
Iowa State	41
George Wash Univ	40
U. Michigan	36
Temple Univ	31
Stanford Univ	28
UCLA	27
U. Central Florida	25
U. Minnesota	24
William & Mary	24
SUNY Stony Brook	22
Scripps Inst	20
U. South Carolina	19
Vanderbilt Univ	19
Yale Univ	19
UC Davis	19
Colorado School Mines	18
U. Maryland	18
N Dakota State U	17
Cornell Univ	16
U. Houston	14
U. Arkansas	14
Dartmouth College	14
Weill Cornell Med College	13
Mississippi State	13
U. Florida	13
Harvard Med School	12
Penn State	11
Duke Univ	10
Clemson Univ	9
UMass Amherst	9
U. Notre Dame	8
Northeastern	7
Auburn Univ	7
Indiana State	6
UC Santa Cruz	6
San Diego State	6
UC Riverside	6
Lehigh Univ	5
Louisiana State	5
Rice Univ	5
U. Vermont	5
U. Toledo	4
Harvard Univ	4
U. Delaware	4
Virginia Tech	4
Boston Univ	4
Johns Hopkins Univ	4
U. Illinois U-C	4
U. Missouri KC	3
Fac for Rare isotope Bms	3
Marquette Univ	3
Purdue Univ	3
Rensselaer	3
Tulane Univ	3
U. Colorado Denver	3
Boston College	3
U. Texas El Paso	3
Univ Col London UK	3
Virginia Commonwealth Univ	3
U. Nebraska Lincoln	2
U. Kansas	2
U. New Mexico	2
U. Tennessee	2
West Virginia Univ	2
U. Guelph Canada	1
MO Univ Science and Tech	1
Georgetown Univ	1
Georgia Tech	1
North Carolina State	1
U. Illinois Chicago	1
Illinois IT	1
New Jersey IT	1
Other (41)	<1

Innovations

Preparing for NERSC's Next Supercomputer: Perlmutter

In 2018, one of the biggest developments at NERSC was the announcement of its next supercomputer. In October the center revealed that it had signed a contract with Cray for a pre-exascale "Shasta" machine, which is slated to be delivered in the latter part of 2020. It is the first NERSC system specifically designed to meet the needs of large-scale simulations and to facilitate analyzing massive datasets from scientific experimental facilities, a growing challenge for scientists across multiple disciplines.

The new supercomputer, a heterogeneous system comprising both CPU-only and GPU-accelerated cabinets, will more than triple the computational power currently available at NERSC. It will feature a number of innovations designed to meet the diverse computational and data analysis needs of NERSC's user base, including a new Cray system interconnect, "Slingshot," that is designed for data-centric computing; NVIDIA GPUs with new Tensor Core technology; direct liquid cooling; and an all-flash scratch filesystem that will move data at a rate of more than 4 terabytes/sec.

The system has been named "Perlmutter" in honor of Dr. Saul Perlmutter, an astrophysicist at Berkeley Lab and a professor of physics at the University of California, Berkeley. He was awarded the 2011 Nobel Prize in Physics for his contributions to research that revealed that the expansion of the universe is accelerating. He currently leads the Berkeley Institute for Data Science and the international Supernova Cosmology Project and co-directs the Berkeley Center for Cosmological Physics.

Perlmutter and his team were among the first scientists to welcome NERSC when it arrived at Berkeley Lab in 1996, and they quickly took advantage of NERSC's computing horsepower. They sifted through enormous sets of data to find rare events like supernovae to build a pipeline of measurements that would upend expectations at the time about the history of the expansion of the universe. The results that emerged after years of observation and computation not only played a vital role in their Nobel-winning research, they were also an example of the impact that large datasets have on scientific discovery today.

Throughout 2018, preparations for the Perlmutter system were a key driver of the center's day-to-day strategy, which includes transitioning the broad Office of Science workload to advanced exascale architectures, evaluating new energy-efficient architectures, and supporting the growing data analysis and machine learning workloads, particularly those originating from DOE's experimental and observational facilities. This section discusses some of the specific efforts undertaken by NERSC in 2018 in preparation for Perlmutter's deployment.

TRANSITIONING USERS TO ADVANCED ARCHITECTURES

At present, transitioning our broad user base to energy-efficient advanced architectures as represented by pre-exascale and future exascale systems is a prime NERSC initiative. To make good use of the Knight's Landing (KNL) architecture on Cori, users must increase both application parallelism and data locality to take advantage of the larger core count and longer vector units on KNL, both key features that will be found on exascale systems. NERSC tracks the percentage of projects that are spending more than 50% of their allocation on KNL. Excluding the set that uses less than 10% of total NERSC hours, in 2017 36% of projects used 50% of their hours on KNL. In 2018, nearly 50% of projects used the majority of their hours on KNL.

The number of users and projects using Cori KNL increased in 2018 relative to 2017 as users continued to transition their workload to KNL. Some 282 users charged at the rate of 2 million NERSC Hours per year in 2018, compared to 198 in 2017. In 2018, of the projects that accounted for 99% of NERSC usage, 62.5% used more than 100,000 hours per month on KNL, compared to 51.5% in 2017. In addition, Cori KNL is supporting a mix of jobs, including large-scale calculations. In the final quarter of CY18, 32% of cycles on Cori KNL were consumed by jobs using 1,024 or more nodes.

Optimizing applications to run well on Cori KNL is a critical stepping stone to running effectively on Perlmutter's CPU and GPU architectures. NERSC has continued to communicate this point to our users as part of its exascale readiness strategy, emphasizing performance portability and the importance of on-node parallelism and data locality. Toward this end, we continue to actively engage the user community through user training, workshops, and hackathons (more on this in Engagement and Outreach).

DEPLOYING AN EXTERNAL GPU PARTITION

Finding additional parallelism and improving vectorization in applications are two fundamental steps in moving codes to exascale architectures. For NERSC users, the goal is to identify components of applications that can benefit from GPU acceleration. To help users prepare for the Perlmutter system, in 2018 NERSC deployed a GPU partition on Cori that supports application performance tuning for GPUs. This cabinet will be used by the next generation of NESAP teams to port and prepare applications for Perlmutter.

The partition comprises 18 Cray CS-Storm 500NX nodes, each with 8 NVIDIA Tesla V100 accelerators, one Intel Skylake processor, and 384 GiB of memory. The onboard GPUs are connected via NVLink, and four dual-port EDR Infiniband NICs provide network connectivity to the nodes. Four of these ports connect to an internal Infiniband network to enable fast MPI communications. The remaining four ports connect the GPU nodes directly to NERSC's internal 100 GbE network, the NGF network, and Cori's Lustre network.

The software configuration of this system is notable from both a user perspective and a systems management perspective. Although the GPU systems were procured as an entirely distinct cluster, NERSC presents these nodes to users as an extension of Cori. NERSC runs the same operating system (CLE from Cray) and configuration used on Cori's login nodes for the GPU systems. Thus, to users, nodes in the GPU partition look and feel like a Cori node. The partition uses the same file systems, applications, and development tools for the GPU nodes as Cori, providing a familiar environment for users. The GPU partition's management environment also parallels Cori's environment by using the SMW software from Cray. This has enabled NERSC to use SMWflow — our git-based software configuration process — to manage multiple systems, such as the GPU cabinet.

In addition to easing the development, testing, and deployment of new software across NERSC's systems, by adopting this process we have been able to unify the systems software between the large-scale Cray systems and any add-on nodes or racks, such as the GPU rack on Cori.

TELEMETRY AND MONITORING SYSTEMS FOR PERLMUTTER

The challenge of monitoring a computational center grows as the center deploys larger and more diverse systems. As system size increases, it becomes harder to discern the problem from the noise. Staff often experience alert fatigue, an occurrence when so many alerts come in that the actual problem is obscured by false alarms or alarms for issues that are symptoms of the core problem.

NERSC has partially resolved this issue by ensuring that most alerts are actionable and that multiple alerts for common problems, such as node outages, do not arise. For the future, NERSC is investigating solutions such as Kubernetes, Prometheus, Elastic, and machine learning techniques to help pinpoint system issues. In addition, we are looking at possibly integrating these multiple tools into a single view.

As systems become larger, more heterogeneous, and more complex, existing monitoring methods will become unmanageable. Traditional methods of monitoring node-based services no longer fit the way systems are configured and managed. New computing paradigms that use containers such as Kubernetes allow for the system to monitor itself and fix some failures by automatically restarting common services or moving the service into a fresh container. A more innovative approach for monitoring these complex systems — in which nodes do not always have the same function when placed back into service — is to focus on the purpose of the node rather than the actual node it runs on. Instead of monitoring the node's function, we can monitor its configuration and let that guide how we manage its alerts.

NERSC is investigating tools such as Prometheus that are specifically designed to monitor a dynamically changing environment built upon containers and services. The tool is open source, is customizable in the way we get data from clients, supports the Simple Network Management Protocol, has better alerting service compared to Nagios, and can provide metrics for what is being monitored

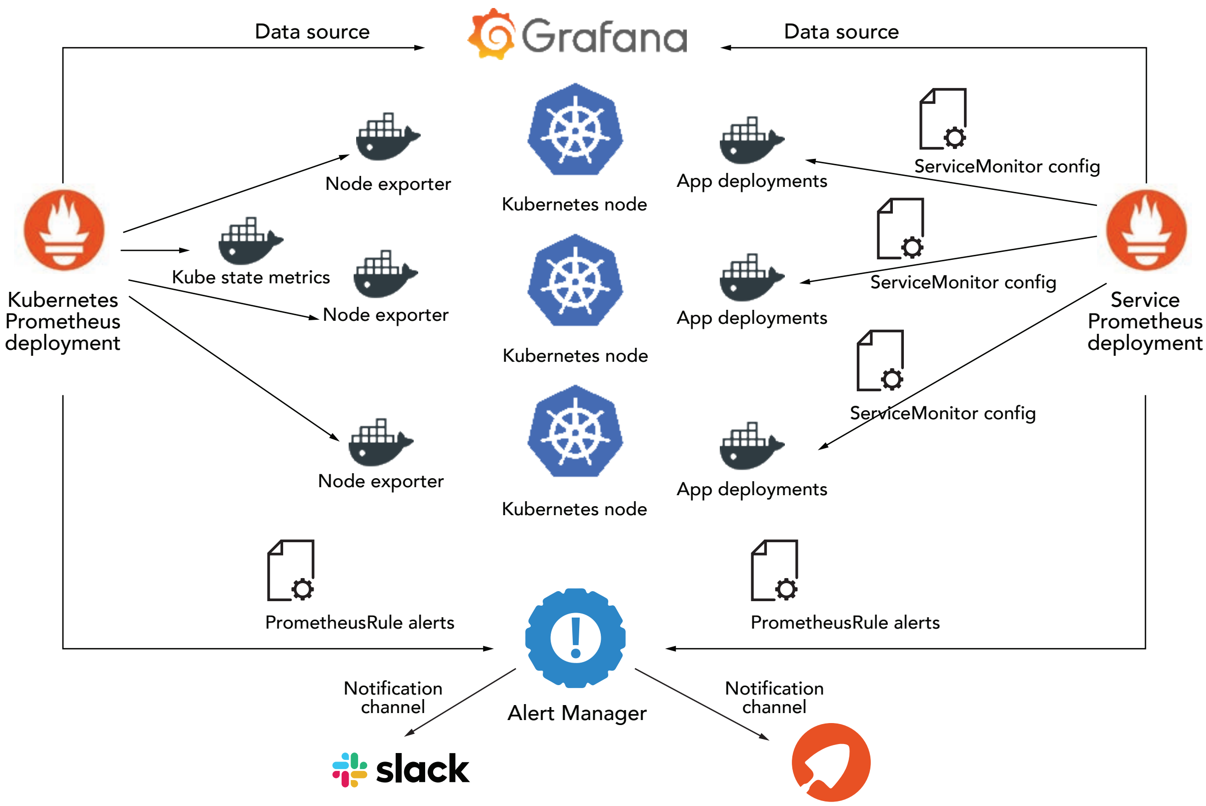

The future of monitoring for the Perlmutter supercomputer.

without using any computational resources. Prometheus can also group alert types, reducing the number of alerts sent or received and thus addressing the alert fatigue issue.

Further, as problems are recorded, the tool will be able to do predictive analysis. For example, it could monitor the fill rate of a file system and send the alert hours or days prior to the actual problem occurring so it can be fixed before it causes a wider issue. This function will help staff transition from a reactive to a proactive standard of operation.

Machine learning implemented within the monitoring infrastructure offers additional potential capabilities. Elastic contains a machine learning module that also provides predictive analysis as long as the same type of problem has been previously recorded. The predictive module can potentially provide a hierarchical alerting structure. Similar to Prometheus, these added functionalities can help staff identify issues more quickly or possibly even before they occur.

With the sources of monitoring data continuously growing, NERSC recognizes the need to integrate these new data sources with the data collection system to provide much-needed information about its systems — data that is currently unavailable. Data must be collected in a single location, with new data added as it becomes available. Creating tools to provide an interface to view the integrated data gives NERSC the ability to leverage the information to adapt to the scale of new systems. By considering and investigating several options now, NERSC will be prepared when Perlmutter comes online in 2020.

Research for Next-Generation Systems

HETEROGENEOUS SSI: A NEW PERFORMANCE METRIC

Perlmutter is slated to be NERSC's first system with large numbers of both CPU-only and GPU-accelerated nodes. Evaluating the performance of Perlmutter's complex heterogeneous platform is challenging because these node types have remarkably different performance characteristics that are exacerbated by the algorithmic differences between applications.

Performance metrics historically used by NERSC, such as Sustained System Performance (SSP), have no notion of heterogeneity. Simple modifications like adding SSP contributions from each partition are inadequate because they fail to describe GPU performance in the frequently encountered case where one or more applications cannot run on the GPU partition; the GPUSSP would collapse to zero, even if the performance of all other applications is outstanding.

Our solution was to develop the heterogeneous Sustained System Improvement (hetero-SSI) metric, which we presented during the Performance Modeling, Benchmarking, and Simulation (PMBS) workshop at SC18. A critical novelty within hetero-SSI is the introduction of affinity parameters, which optimally distribute applications to the most performant partitions. Hetero-SSI was used to optimize the relative sizes of the CPU and GPU partitions on Perlmutter and will remain

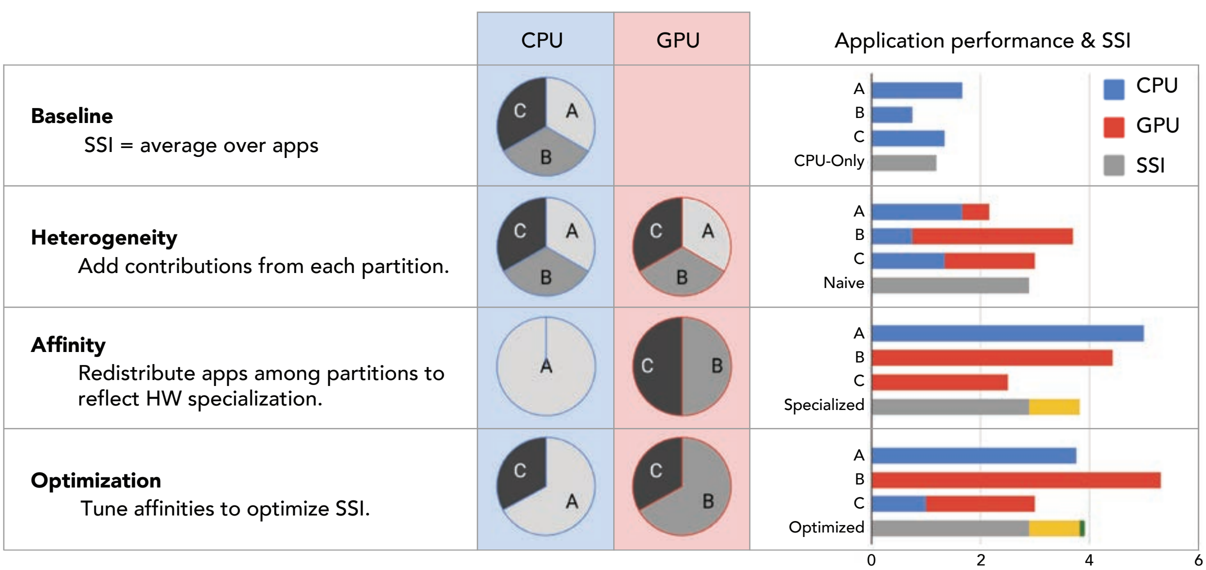

Summary of the heterogeneous Sustained System Improvement metric.

relevant for future systems as they incorporate more diverse mixtures of specialized accelerators that have been optimized for specific workloads.

As shown in the figure below, the new hetero-SSI performance metric adopts three incremental changes relative to the "baseline" (non-heterogeneous) SSI. The baseline SSI simply averages the performance of applications when each is allocated an equal fraction of the CPU partition; "naive" incorporation of heterogeneity adds the GPU contribution to the performance of each application before averaging. The second change recognizes that applications that run better on GPUs should run only on GPUs (and vice versa). This strict affinity between applications and partitions reflects users' likely initial response to the availability of GPU hardware and leads to the "specialized" increment to SSI, shown by the yellow bar. The final stage determines the optimal distribution of applications across partitions. This decreases the aggregate (CPU+GPU) performance of some applications for the benefit of others, resulting in a net increase in SSI (shown by the green bar). Although this example uses only two partitions, hetero- SSI can be easily applied to any number of accelerators.

NERSC AND CRD RESEARCH COLLABORATIONS

The Computing Sciences program at Berkeley Lab is anchored by two major user facilities — NERSC and ESnet — and the Computational Research Division (CRD), which has major activities in applied math, computer science, data science, and domain-specific partnerships in computational modeling and data analysis. The facilities serve a large cross-section of the DOE science community, with NERSC supporting several thousand researchers from universities, national laboratories, and industry, and ESnet connecting the DOE community to one another and the rest of the Internet for distributed team science.

In 2018 NERSC established a number of collaborations with researchers in CRD. These collaborations are designed to create opportunities to advance Berkeley Lab Computing Sciences Area strategic goals, strengthen CRD-NERSC-ESnet ties, prepare for future DOE calls for proposals, and work together on projects that bridge current CRD projects and NERSC/ESnet goals. The research topics identified for collaboration were based on three key strategic priorities for NERSC over the next five years:

- Application readiness for NERSC-9
- Superfacility model
- NERSC-9 workflow accelerators and NERSC-10 pathfinding

In October, researchers from CRD presented specific ideas for collaborative work with NERSC. These were discussed and reviewed by NERSC staff who ultimately selected five collaborative areas:

- Application performance modeling and enablement
- Exploring computational accelerator technologies for NERSC-10
- Jupyter tools for HPC
- Mixed-precision computing
- Complex workflow support and optimization

IMPROVING NERSC CONNECTIVITY OVER THE WIDE AREA NETWORK

NERSC data transfer nodes (DTNs) are a set of systems used for transferring data into and out of the facility and its systems. The DTNs provide high-speed access to the center's file systems and HPSS tape archive through several popular data transfer tools.

In 2018 the network infrastructure connecting the DTNs was upgraded to 100Gb. Individual DTNs now have 100Gb network connections providing an aggregate bandwidth of 1.2TB for the DTN pool. To improve performance over the wide area network, systems are

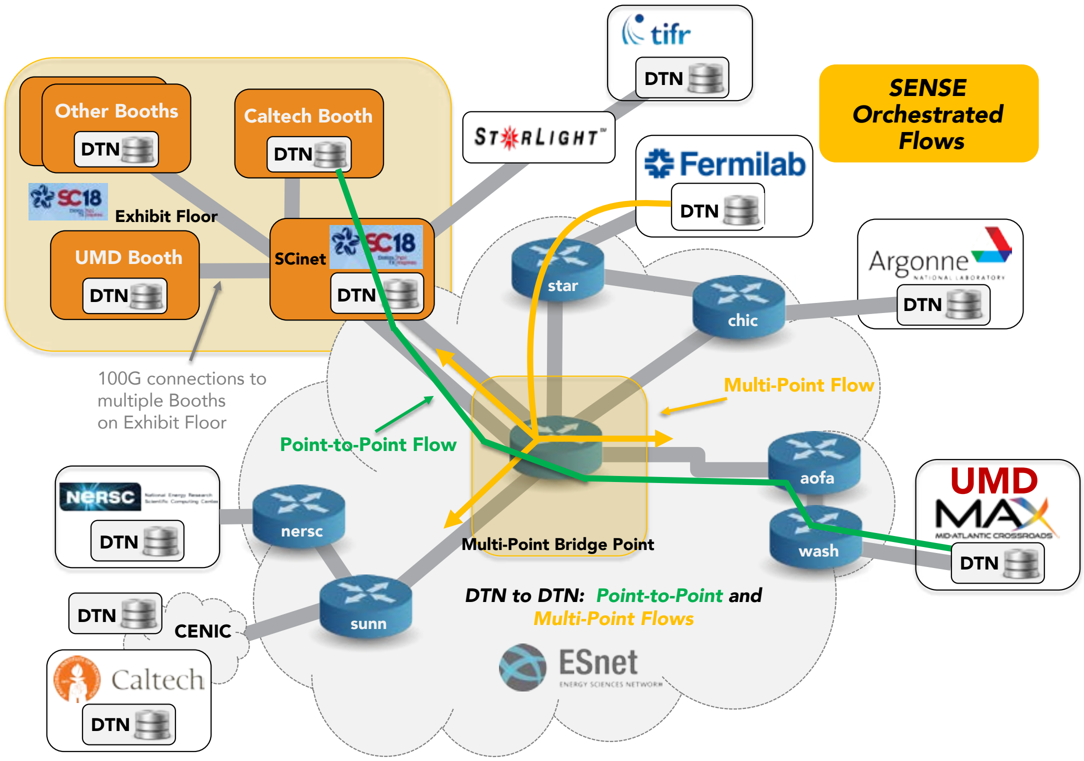

Overview of the SENSE research project.

configured for traffic pacing, which helps stabilize transfer rates for individual streams. This network update improves connectivity between NERSC and other computational, experimental, and observational facilities.

The DTN pool also provides an endpoint for exploring software-defined networking (SDN) over the wide area network. The Software-defined network for End-to-end Network Science at Exascale (SENSE) research project aims to build smart network services whose features can be programmed through a simple intuitive interface. A joint effort between ESnet, Berkeley Lab's Computational Research Division, Caltech, Fermilab, Argonne National Lab, and the University of Maryland, the goal of SENSE is to grant scientific applications the ability to manage the network as a schedulable resource, similar to how computational and facility instruments are managed, thereby including network scheduling in complex workflows spread across multiple facilities. It also complements future superfacility infrastructure efforts that aim to build tighter integration between NERSC and remote facilities.

The network team has also deployed 100Gb infrastructure to support the newly installed tape libraries at Shyh Wang Hall. The new connectivity provides 400Gb aggregated bandwidth, preparing the archive for increased workloads from the Perlmutter system when it is installed, and from remote facilities.

HIERARCHICAL ROOFLINE DEVELOPMENT

The Roofline performance model — developed by Berkeley Lab's Samuel Williams and first published in 2009 — is widely used for characterizing and optimizing HPC codes. New Roofline achievements in 2018 include the integration of Roofline into Intel Advisor, the extension of DRAM Roofline to Hierarchical Roofline, and the proposal of Roofline data collection approaches using open-source tools such as LIKWID.

In 2018, there were two major advances to the Roofline model. First, a data-collection methodology was proposed for hierarchical Roofline on NVIDIA GPUs. This methodology uses the Empirical Roofline Toolkit to collect compute and bandwidth peaks, which provide a more realistic set of Roofline ceilings than vendor specifications themselves. It also uses NVIDIA's Nvprof help understand the various nuances of HPC performance on GPUs, such as memory coalescing and thread predication.

In addition, an empirical Roofline methodology was proposed to help quantitatively assess performance portability across multiple architectures. The methodology empirically captures a more realistic set of performance bounds for both CPUs and GPUs using ERT; factors in the true cost of different floating-point instructions, such as divides, when counting FLOPs; incorporates the effects of different memory access patterns, such as strided memory accesses; and is proven to be effective in producing an accurate "architectural efficiency" for performance portability analysis. The highly parameterizable kernel GPP from BerkeleyGW was deployed to validate the

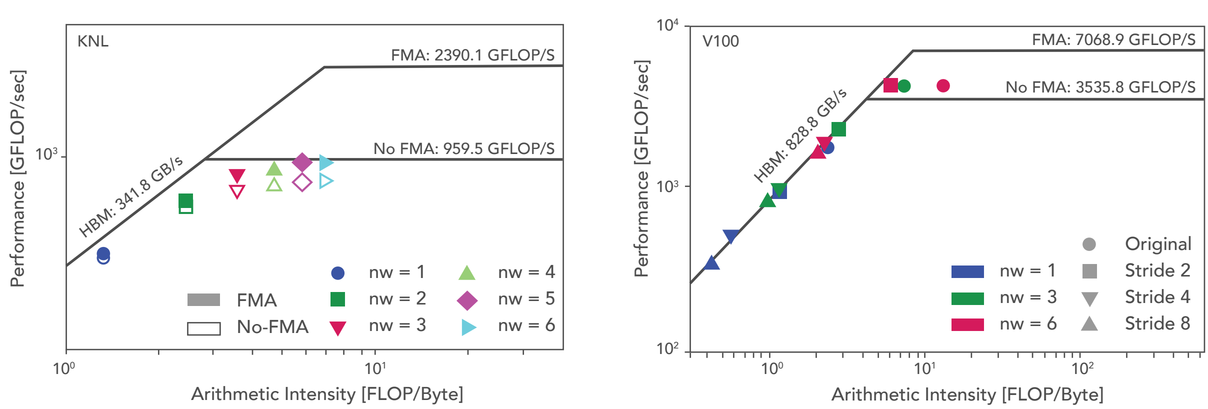

Roofline analysis for the impact of FMA/no-FMA instruction mix on performance on KNL (left), and strided memory access pattern on performance on NVIDIA V100 (right).

tool to collect application performance data, such as runtime, FLOPs, and data movement. The Hierarchical Roofline incorporates all four memory levels — L1 (unified cache), L2, device memory, and system memory — into a single figure and provides more insights into performance analysis. Proxy applications such as GPP from BerkeleyGW and HPGMG from AMReX were used to validate the effectiveness of this methodology and to effectiveness of this methodology on two modern architectures: Intel Xeon Phi and NVIDIA V100 GPUs. It is also used to facilitate the discussion around various nuances in performance portability analysis, such as appropriate FLOP counting for divides and appropriate pairing of application performance and Roofline ceiling. Additional work is ongoing in collaboration with NVIDIA.

NEW TOTAL KNOWLEDGE OF I/O FRAMEWORK

In 2017 NERSC published an initial version of the Total Knowledge of I/O (TOKIO) framework, which provides a foundation for connecting multiple sources of monitoring data to establish a holistic understanding of I/O performance. In 2018 a new release of the TOKIO framework was published, with a new set of companion tools designed to work specifically with Cray ClusterStor storage appliances such as those deployed with NERSC's Edison and Cori systems. The architecture and functionality of this expanded release of TOKIO was presented at the 2018 Cray User Group in Stockholm, Sweden; this version of the TOKIO framework is now in use on Argonne's Theta system.

Also, in collaboration with Argonne Mathematics and Computer Science and the Argonne Leadership Computing Facility (ALCF), TOKIO was applied to study how I/O performance varies on the production file systems at NERSC and ALCF over the course of an entire year. This work, presented at SC18 in a paper titled "A Year in the Life of a Parallel File System," applied a variety of statistical techniques to quantify the degree to which different forms of resource contention and storage system degradations cause abnormally poor I/O performance over three different time scales. These results identified specific sources of contention that, if mitigated using new quality-of-service (QOS) techniques, would dramatically reduce I/O performance variation. Such advanced QOS controls will be available in the Perlmutter system.

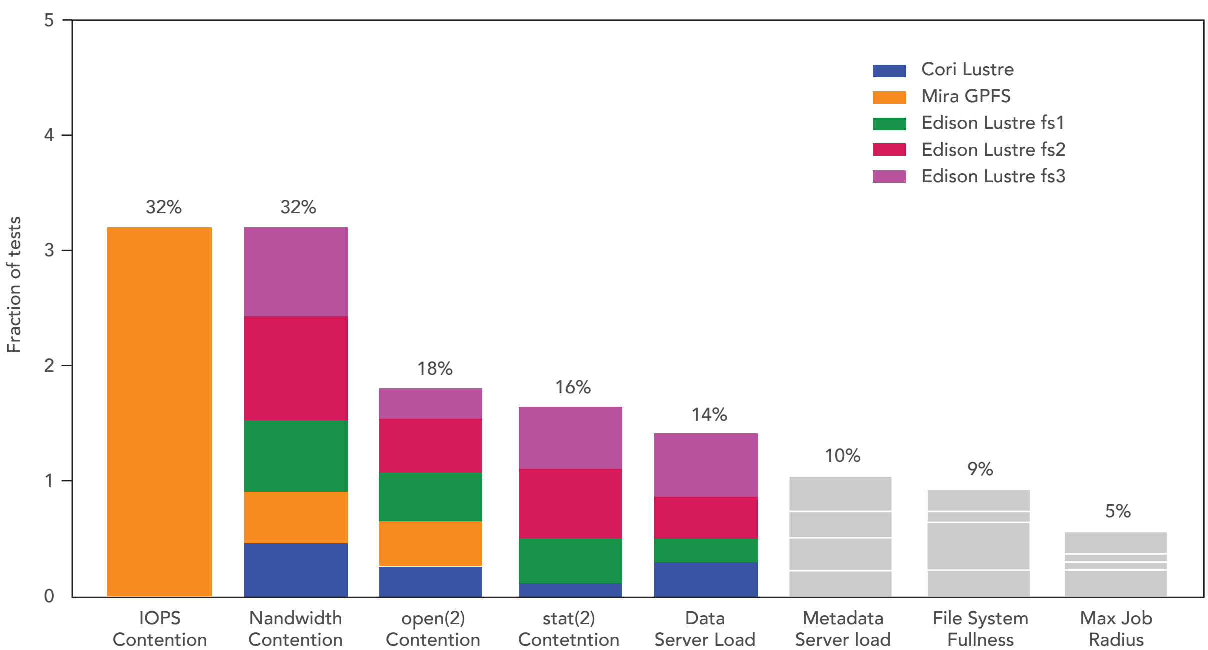

Causes of transient I/O performance losses that occurred over a one-year period at NERSC and ALCF. Multiple factors may contribute to a single incident. "IOPS Contention" was observed only on Mira because it relied on data not collected at NERSC. The grey bars indicate causes that occurred too infrequently to establish strong statistical significance.

UPDATE ON THE STORAGE 2020 ROADMAP

In 2017 NERSC published a technical report presenting a detailed storage roadmap and vision to address the overwhelming data-storage challenges the science community is expected to face over the coming decade. The report, "Storage 2020: A Vision for the Future of HPC Storage," used feedback from the Exascale Requirement Reviews, surveys of current and emerging storage technologies, and extensive discussions with peers and collaborators within NERSC and other DOE facilities to define a reference storage architecture for HPC centers. Using the requirements reviews and a detailed workload analysis, NERSC identified four data storage categories required by the user community: temporary, campaign, community, and "forever" (long-term).

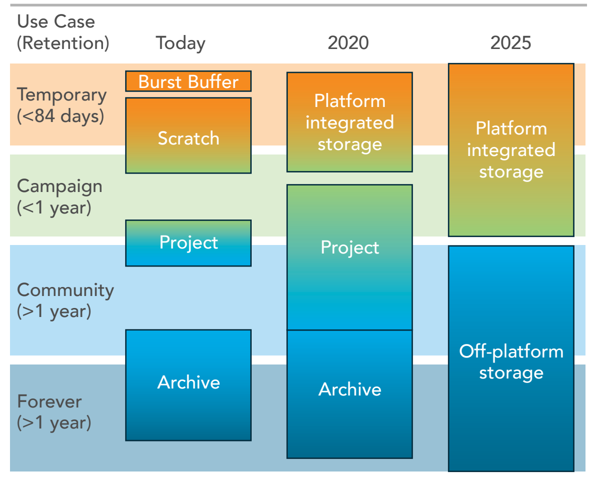

Evolution of the NERSC storage hierarchy between today and 2025.

Throughout 2018, NERSC used the findings presented in the report to drive areas of innovation and research and to guide operational decisions for the center's storage deployments at all levels of the storage hierarchy. At the highest performing Temporary tier, NERSC committed to deploying an all-flash Lustre file system with the Perlmutter system and began working with Cray to ensure that the deployment will meet functional and performance expectations. NERSC also embarked on a major refresh of the Campaign/Community tier. As part of that effort, NERSC issued an RFP for a 200- to 300-petabyte capacity-focused storage system. In the Forever tier, NERSC reinvested in its tape archive, deploying two new tape libraries with integrated cooling that are uniquely suited to protect the tape media from the temperature and humidity variations in the energy-efficient computer room in Shyh Wang Hall (for more on this, see Center News).

The Storage 2020 report also guided areas of work focused on exploring longer term research in storage hardware and software. The report recognized that continuing to increase the performance and size of parallel file systems that were first developed for hard disk drives and much smaller computational systems will be technically challenging. It also identified object store technologies as a potential alternative high-capacity, high-performance storage system. Object stores are designed with a focus on scalability and have found widespread adoption in commercial cloud computing environments. Among the challenges to using object stores for HPC applications is that few applications can perform I/O directly to an object store.

Using middleware between the object store and an HPC application is one possibility for addressing this. To explore this approach, we developed an HDF5 virtual object layer plugin for Intel DAOS, Ceph RADOS, and Openstack Swift object stores. Three HPC applications — VPIC, H5Boss, and BDCATS — were used to evaluate the object stores. The results — presented at the 2018 IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems in a paper titled "Evaluation of HPC Application I/O on Object Storage Systems" — found that for some workloads object stores perform better than a traditional parallel file system, laying the groundwork for NERSC to further explore the use of object stores.

UPDATE ON THE EXAHDF5 COLLABORATION

The ExaHDF5 project, composed of developers and researchers at NERSC, ALCF, and the non-profit The HDF Group, researches improvements to, implements capabilities in, makes production releases of, and supports users of the HDF5 I/O middleware package in HPC environments. The ExaHDF5 team has been funded to work with ECP application developers to improve their knowledge and effective use of HDF5 and to enhance the HDF5 library for use in exascale environments.

As part of the 2018 year of ECP-ExaHDF5 funding, NERSC delivered the following feature milestones:

- The virtual object layer, which allows interception of HDF5 public API at runtime to access data in alternate ways, such as in an object storage system, other file formats, or using system-specific caching techniques.
- Data Elevator, a write-caching virtual object layer connector that intercepts HDF5 writes and redirects them to use burst buffer storage transparently.
- Topology-aware I/O in HDF5, to take advantage of the underlying storage network topology for enhancing the I/O performance of HDF5 applications.
- Full single writer/multiple readers application access, which supports advanced workflows by allowing simultaneous, lock-free access to HDF5 data files from one writing application and many reading applications.

In addition to delivering functionality for HPC applications funded by ECP, the ExaHDF5 team is continuing its project to expand and enhance HDF5 to store experimental and observational data throughout the DOE complex. As part of this effort, the team is researching enhancements to HDF5 to record provenance information, store sparse and stream-oriented data, perform multi-writer concurrent I/O, and update the schema of HDF5 files over time.

Machine Learning and Deep Learning for Science

Another key component of enabling large-scale data analysis at DOE's experimental facilities is supporting the advanced analytics requirements of these increasingly complex workloads. Toward this end, NERSC is investing heavily in deploying analytics and machine learning tools and infrastructure at scale.

DEEP LEARNING, EXTREME WEATHER, AND THE 2018 GORDON BELL PRIZE

Building expertise in the intersection of HPC and artificial intelligence is crucial for NERSC; our deep learning user community is steadily growing and is interested in exploring increasingly sophisticated methods on increasingly complex datasets.

The common perception in the HPC community is that high-productivity tools such as PyTorch or TensorFlow favored by many NERSC deep learning practitioners offer rapid prototyping capabilities but fall short on performance and scalability. NERSC researchers have

NERSC and NVIDIA collaborators shared the ACM SC18 Gordon Bell award for scalability and time to solution with another team led by ORNL.

proven that these two important properties need not be mutually exclusive; over the course of eight months in 2018, a NERSC team working closely with collaborators from NVIDIA prototyped a network from scratch and optimized its performance and scaling on ORNL's Summit supercomputer, the world's largest HPC system.

These efforts led to a 2018 Gordon Bell Prize for applying an exascale-class deep learning application to extreme climate data and breaking the exaop (1 billion billion calculations) computing barrier for the first time with a deep learning application.

Climate scientists rely on some of the fastest computers in the world to run high-fidelity simulations, but those simulations produce massive amounts of data that create their own computational challenge. This prompted the collaborators to apply deep learning methods to extract detailed information from climate data produced at NERSC. Rather than computing simple quantities such as average global temperature, these methods discover and locate features like hurricanes and atmospheric rivers, which are important in characterizing extreme weather patterns and their impact.

For this project, the Berkeley Lab-led team used TensorFlow, a high-productivity framework, to scale deep learning code on up to 4,560 nodes on ORNL's Summit system. The performance enhancements are not hardwired to the climate application, but can be generally applied to other scientific deep learning applications that require extreme scaling. The team hopes that his project has provided a blueprint for AI+HPC projects for the future.

NEW METHODS FOR DETECTING AND CHARACTERIZING EXTREME WEATHER EVENTS

Two PhD students who first came to Berkeley Lab as summer interns in 2016 are spending six months a year at the Lab through 2020 developing new data analytics tools that could dramatically impact climate research and other large-scale science data projects.

Grzegorz Muszynski is a PhD student at the University of Liverpool, U.K. studying with Vitaliy Kurlin, an expert in topology and computational geometry. Adam Rupe is pursuing his PhD at the University of California at Davis under the supervision of Jim Crutchfield, an expert in dynamical systems, chaos, information theory and statistical mechanics. Both are also working in NERSC's Data & Analytics Services (DAS) group, and their PhDs are being funded by the Big Data Center (BDC), a collaboration between NERSC, Intel, and five Intel Parallel Computing Centers.

From the get-go Muszynski and Rupe's projects have focused on addressing a grand challenge in climate science: finding more effective ways to detect and characterize extreme weather events in the global climate system across multiple geographical regions and developing more efficient methods for analyzing the ever-increasing amount of simulated and observational data. Automated pattern recognition is at the heart of both efforts, yet the two researchers are approaching the problem in distinctly different ways: Muszynski is using various combinations of topology, applied math, and machine learning to detect, classify, and characterize weather and climate patterns, while Rupe has developed a physics-based mathematical model that enables unsupervised discovery of coherent structures characteristic of the spatiotemporal patterns found in the climate system.

"When you are investigating extreme weather and climate events and how they are changing in a warming world, one of the challenges is being able to detect, identify, and characterize these events in large data sets," said Karthik Kashinath, a computer scientist and engineer in the DAS group who leads multiple BDC climate science projects. "The volume at which climate data is being produced today is just insane, so there is a strong need to automate the process of discovering structures in data." There is also a desire to find climate data analysis methods that are reliable across different models, climates, and variables.

Muszynski and Rupe are making steady progress toward meeting these challenges. Muszynski has developed a framework of tools from applied topology and machine learning that complement existing tools and methods used by climate scientists and can be mixed and matched depending on the problem to be solved. In addition, the method is "threshold-free," a key advantage over

> *"As far as we are aware, this is the only completely unsupervised method that does not require training data. In addition, it covers every potential structure and pattern you might be looking for in climate data, and you don't need preconceived notions of what you are looking for."*
>
> **— Adam Rupe, PhD student at NERSC**

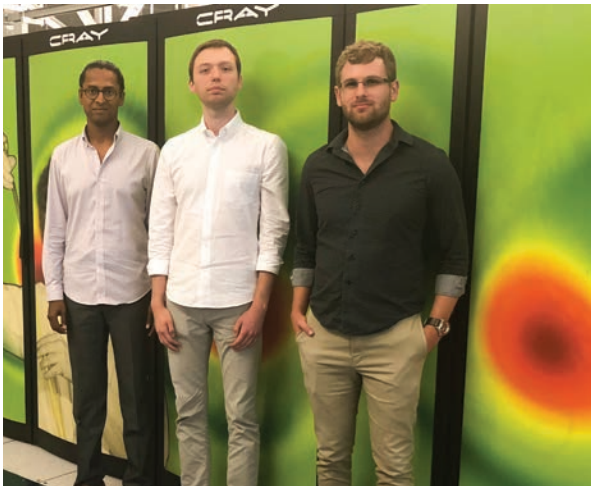

Left to right: Karthik Kashinath, Grzegorz Muszynski, and Adam Rupe of NERSC.

existing data analysis methods used in climate research. As part of this work, Muszynski parallelized his codebase on several nodes on NERSC's Cori system to accelerate the machine learning training process, which often requires hundreds to thousands of examples to train a model that can classify events accurately.

While topology has been applied to simpler, smaller scientific problems, this is one of the first attempts to apply topological data analysis to large climate data sets. The results so far have been impressive, with notable reductions in computational costs and data extraction times. In addition, Muszynski's framework "doesn't really care where you are on the globe," Kashinath said. "You can apply it to atmospheric rivers in North America, South America, Europe — it is universal."

Rupe is also thinking outside the box, applying physics rather than machine or deep learning to analyze data from complex nonlinear dynamical systems and using physical principles associated with organized coherent structures to find these structures in the data. In particular, his model relies on computational mechanics to look for local causal states that deviate from a symmetrical background state; these local causal states provide a principled mathematical description of coherent structures and a constructive method for identifying them directly from data.

Rupe's approach requires novel and unprecedented scaling and optimization on Cori for multiple steps in the unsupervised discovery pipeline, including clustering in very high-dimensional spaces and clever ways of data reuse and feature extraction, Kashinath noted. Among other things, they have discovered that this approach offers a powerful alternative to machine and deep learning by enabling unsupervised segmentation and pixel-level identification of coherent structures without the need for labeled training data.

"As far as we are aware, this is the only completely unsupervised method that does not require training data," Rupe said. "In addition, it covers every potential structure and pattern you might be looking for in climate data, and you don't need preconceived notions of what you are looking for. The physics helps you discover all of that automatically."

COSMOFLOW: HARNESSING THE POWER OF DEEP LEARNING TO BETTER UNDERSTAND THE UNIVERSE

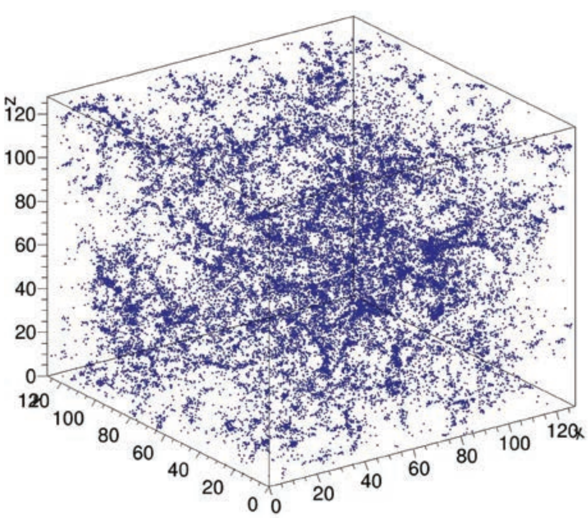

Example simulation of dark matter in the universe used as input to the CosmoFlow network.

A Big Data Center collaboration between computational scientists at NERSC and engineers at Intel and Cray has yielded another first in the quest to apply deep learning to data-intensive science: CosmoFlow, the first large-scale science application to use the TensorFlow framework on a CPU-based HPC platform with synchronous training. It is also the first to process 3D spatial data volumes at this scale, giving scientists an entirely new platform for gaining a deeper understanding of the universe.

Cosmological "big data" problems go beyond the simple volume of data stored on disk. Observations of the universe are necessarily finite, and the challenge that researchers face is how to extract the most information from the observations and simulations available. Compounding the issue is that cosmologists typically characterize the distribution of matter in the universe using statistical measures of the structure of matter in the form of two- or three-point functions or other reduced statistics. Methods such as deep learning that can capture all features in the distribution of matter would provide greater insight into the nature of dark energy. First to realize that deep learning could be applied to this problem were Siamak Ravanbakhsh and his colleagues; however, computational bottlenecks when scaling up the network and dataset limited the scope of the problem that could be tackled.

Motivated to address these challenges, the CosmoFlow team laid out three primary goals for this project: science, single-node optimization, and scaling. The science goal was to demonstrate that deep learning can be used on 3D volumes to learn the physics of the universe. The team also wanted to ensure that TensorFlow ran efficiently and effectively on a single Intel Xeon Phi processor node with 3D volumes, which are common in science but not so much in industry, where most deep learning applications deal with 2D image data sets. And finally, they wanted to ensure high efficiency and performance when scaled across thousands of nodes on the Cori supercomputer.

Thus CosmoFlow was designed to be highly scalable; to process large, 3D cosmology datasets; and to improve deep learning training performance on modern HPC supercomputers such as the Intel processor-based Cray XC40 Cori supercomputer at NERSC. CosmoFlow is built on top of the popular TensorFlow machine learning framework and uses Python as the front end. The application leverages the Cray PE Machine Learning Plugin to scale the TensorFlow Deep Learning framework to more than 8,000 nodes. It also benefits from Cray's DataWarp I/O accelerator technology, which provides the I/O throughput required to reach this level of scalability.

During the initial phase of the CosmoFlow project, the researchers identified framework, kernel, and communication optimization that led to more than 750x performance increase for a single node. They also solved problems that limited scaling of deep learning techniques to 128 to 256 nodes. In a technical paper presented at SC18, the CosmoFlow team described the application and initial experiments using dark matter N-body simulations produced using the MUSIC and pycola packages on the Cori supercomputer at NERSC. In a series of single-node and multi-node scaling experiments, the team was able to demonstrate fully synchronous data-parallel training on 8,192 of Cori with 77 percent parallel efficiency and 3.5 Pflop/s sustained performance.

MACHINE LEARNING SPEEDS SCIENCE DATA SEARCHES

As scientific datasets increase in size and complexity, the ability to label, filter, and search this expanding body of information has become laborious, time-consuming, and sometimes impossible. With this in mind, a team of researchers from Berkeley Lab and UC Berkeley are developing innovative machine learning tools to pull contextual information from scientific datasets and automatically generate metadata tags for each file. Scientists can then search these files via a new web-based search engine for scientific data called Science Search.

Today, search engines are ubiquitously used to find information on the Internet, but searching science data presents unique challenges. For example, Google's algorithm relies on more than 200 clues to achieve an effective search. These clues can come in the form of key words on a webpage, metadata in images, or audience feedback from billions of people when they click on the information they are looking for. In contrast, scientific data comes in many forms that are radically different than an average web page, require context that is specific to the science, and often lack the metadata to provide the context required for effective searches.

As a proof-of-concept, in 2018 the Berkeley team began working with staff at Berkeley Lab's Molecular Foundry to demonstrate the concepts of Science Search on images captured by the facility's instruments, making a beta version of the platform available to Foundry researchers. At the Foundry, researchers from all over the world apply for time and then travel to Berkeley Lab to use extremely specialized instruments free of charge, sometimes collecting up to a terabyte of data in under 10 minutes. Users then need to manually sift through this data to find quality images with "good resolution" and save that information on a secure shared file system or on an

“One of the reasons it is possible for us to build a tool like Science Search is our access to resources at NERSC.”

— Gonzalo Rodrigo, a Berkeley Lab postdoctoral researcher

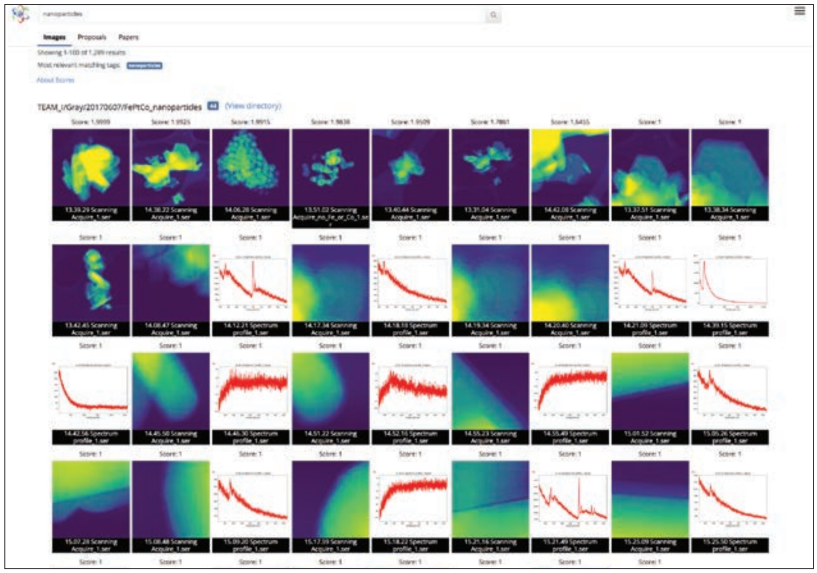

This screen capture of the Science Search interface shows how users can easily validate metadata tags that have been generated via machine learning, or add information that hasn’t already been captured.

external hard drive that they eventually take home with them to analyze. Because it is tedious and time consuming to manually add notes to terabytes of scientific data and there is no standard for doing it, most researchers just type shorthand descriptions in the filename. This might make sense to the person saving the file, but often doesn’t make much sense to anyone else. The lack of real metadata labels eventually causes problems when the scientist tries to find the data later or attempts to share it with others.

To address this challenge, Science Search uses convolutional neural networks to mine the “science ecosystem” — including instrument timestamps, facility user logs, scientific proposals, publications, and file system structures — for contextual information. The collective information from these sources — including timestamp of the experiment, notes about the resolution and filter used, the user’s request for time — provides critical contextual information. The team has also put together an innovative software stack that uses natural language processing and other machine learning techniques to extract contextual keywords about the scientific experiment and automatically create metadata tags for the data.

Because scientific instruments are generating an ever-growing body of data, all aspects of the Science Search engine needed to be scalable to keep pace with the rate and scale of the data volumes being produced. The team achieved this by setting up their system using NERSC’s Spin tool. Spin is a Docker-based edge-services technology developed at NERSC that can access the facility’s high performance computing systems and storage on the back end (for more about Spin, see Engagement and Outreach).

“One of the reasons it is possible for us to build a tool like Science Search is our access to resources at NERSC,” said Gonzalo Rodrigo, a Berkeley Lab postdoctoral researcher who is working on the natural language processing and infrastructure challenges in Science Search. “We have to store, analyze, and retrieve really large datasets, and it is useful to have access to a supercomputing facility to do the heavy lifting for these tasks. NERSC’s Spin is a great platform to run our search engine; it is a user-facing application that requires access to large datasets and analytical data that can only be stored on large supercomputing storage systems.”

USING NEURAL NETWORKS TO ENHANCE HIGH ENERGY PHYSICS SIMULATIONS

In a plenary talk at the 2018 Computing in High Energy and Nuclear Physics conference, Steve Farrell — a machine learning engineer at NERSC — presented findings from a study in which a team of Berkeley Lab Computing Sciences researchers used generative adversarial networks (GANs) to speed simulations in high energy physics (HEP) studies.

At the 2018 CHEP conference, NERSC's Steve Farrell presented findings from a study in which Berkeley Lab Computing Sciences researchers used generative adversarial networks to speed simulations in high energy physics research.

Orchestrating particle collisions and observations at facilities like CERN, where groups of protons collide with one another 40 million times per second, is a massive scientific undertaking. The goal is to understand the interactions and decays of fundamental particles and look for anomalies from the expected behavior. But analyzing the enormous amounts of data produced from the experiments is becoming an overwhelming challenge. Large event rates mean physicists must sift through tens of petabytes of data per year, including a large volume of simulated data essential for modeling those collisions and performing analysis. Traditionally, HEP simulations of collision events in experimental facilities take a particle-by-particle approach that is time-consuming and computationally expensive.

One newer approach involves applying deep learning generative models to the data-simulation process. Some research groups have focused on single particles as they enter a certain section of the detector, using a deep learning generative model to simulate the particle's "showering" — how it sprays out as it collides with material. But Farrell and colleagues at Berkeley Lab are taking a different approach: using deep learning to train a GAN to learn the distribution of full detector images. With this modeling approach — which has also been successful at generating large, weak lensing cosmology convergence maps — the team was able to generate full particle physics events and large, weak lensing cosmology convergence maps.

The approach to training a GAN model is like a two-player game in which one neural network is trying to generate realistic looking samples and the other is trying to distinguish which ones are generated and which ones are real, Farrell noted. However, while it is a powerful technique, it can be difficult and unstable. "Much like how two people may try to learn to play chess by only playing against each other; they may get great against each other but their skills may not generalize," he said. "They may not actually learn how to play chess."

Despite the challenges, Farrell and colleagues have shown that their model can produce whole event images with reasonably good fidelity that exhibit the desired physical characteristics. They do this by applying standard physics reconstruction techniques to these images, clustering the energy deposits, and creating particle object representations. "Then we look at the distributions of those and compare to real events to decide how well it is doing. And basically we get a model that can learn to capture the distributions pretty well."

This sort of work is mission critical for the future of the Large Hadron Collider and related experiments because of the enormous data rate and stringent requirements on the precision of our analyses, where ultrafast, high-fidelity simulations will be required for maximizing the scientific output of these unique datasets.

Engagement and Outreach

NESAP Breaks New Ground in 2018

In 2018, NERSC staff continued to collaborate with application teams as part of the NERSC Exascale Science Application Program (NESAP) to prepare codes for next-generation pre-exascale and exascale architectures, focusing on increasing parallelism, data locality, and vectorization.

Some collaborations focusing on data analysis workloads have seen application or kernel performance increases of up to 200x, making it feasible to run them at large scales on HPC platforms.

Highlights from 2018 from the NESAP for Data program include progress on accelerating the Dark Energy Spectroscopic Instrument (DESI) spectroscopic extraction pipeline on Cori. While the goal was to accelerate the codes' performance on the Cori KNL partition, performance on Cori Haswell and Edison increased as well. These efforts make it feasible for DESI to move into commissioning in 2019 and run their processing pipeline on KNLs.

Efforts to port the x-ray light source tomographic reconstruction code TomoPy to GPUs began in the fall of 2018. The focus is on iterative reconstruction methods that are computationally intensive and a good candidate for porting to GPUs. With the arrival of GPU nodes on Cori, we are able to demonstrate >200x speedup on some reconstruction tasks compared to the Edison baseline. Other NESAP projects in 2018 included a Roofline analysis of the Geant Monte Carlo code and improving scalability of the ATLAS experiment's pipeline.

NESAP staff also developed and evaluated several performance portability case studies. As an example, a NESAP postdoc documented a materials science case study of tensor contractions using CUDA, OpenMP, and OpenACC to target the Cori KNL system as well as Summit at the Oak Ridge Leadership Facility at Oak Ridge National Laboratory. This work was presented and published at the 5th annual WACCPD workshop at SC18.

ENGAGING WITH THE EXASCALE COMPUTING PROJECT

Berkeley Lab is one of six major national laboratories playing a key role in the Department of Energy's Exascale Computing Project (ECP), which aims to deliver breakthough modeling and simulation solutions that analyze more data in less time. NERSC engaged with the ECP in a variety of ways in 2018:

Access to NERSC systems for application development and software technology projects: NERSC provided 750M NERSC Hours to support ECP application development and software technology projects. Usage by ECP teams totaled just over 600M NERSC Hours.

PathForward: NERSC staff have been active in the PathForward projects with Intel, AMD, NVIDIA, IBM, HPE, and Cray, acting as technical representatives and subject matter experts. NERSC and Berkeley Lab Computing Sciences staff have also been involved in deep dive meetings held at Intel, AMD, and Cray.

Software technologies: NERSC identified packages of interest and ranked the relevance of the software development projects in the ECP Software Technologies area.

Continuous integration: NERSC contributed to the design and implementation of a CI/CD system that can run and submit compute jobs on HPC resources. Staff deployed a testing GitLab instance, integrated with the local test and development systems, that is being used to run CI/CD workflows by ECP project members and local staff.

Training: NERSC collaborated with ECP and the other ASCR facilities to present training classes. For example, "Jupyter in HPC" was presented in February with an emphasis on using Jupyter to expose HPC resources at NERSC, ALCF, and OLCF; and NERSC staff contributed to the ATPESC summer school covering high-performance I/O topics.

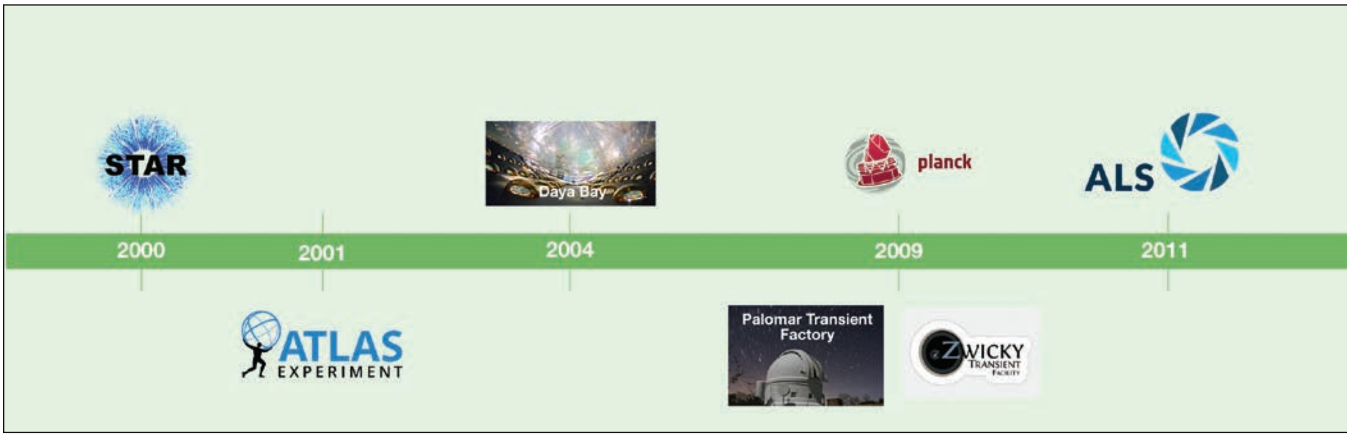

Timeline of prototype pipelines integrating multiple facilities.

Applications: A significant number of ECP application teams are collaborating with NERSC staff to ready their codes for exascale through the NESAP program. These include Lattice QCD, NWChemEx, WDMApp, WarpX, ExaStar, ExaSky, Subsurface, Exabiome, E3SM, and ExaFEL (Exascale for Free Electron Lasers). For example, as part of the ExaFEL project, ESnet has developed a network application client for SLAC's Linac Coherent Light Source (LCLS) that allows their workflow to meet bandwidth guarantees by shaping the network path. In 2018, NERSC, ESnet, and LCLS successfully demonstrated the client in an ECP milestone. NERSC staff are also involved in projects such as the AMReX codesign center and ExaHDF5 leadership.

NERSC LAUNCHES INTERNAL SUPERFACILITY PROJECT

Over the years, NERSC's workload has expanded to include more projects that are analyzing data from DOE's experimental and observational facilities. In 2018, nearly 25% of NERSC users were part of a DOE experiment or facility project analyzing data at NERSC, and more than 35% of NERSC Principal Investigators noted that the primary role of their project is to analyze data from DOE's experimental facilities, create tools to analyze data from these facilities, or combine large-scale simulation with data analysis from facilities. Each SC Office is well represented by these data projects, with BES having the largest absolute number of projects and HEP having the largest percentage of projects (56%).

Engaging with these facilities is key to understanding their requirements and developing capabilities that will support many experiments. NERSC has recognized the increased importance of this new workload to current and future operations and, in response, in 2018 initiated an internal Superfacility Project to:

- Coordinate these user facility engagements
- Anticipate their cross-cutting challenges in moving, managing, analyzing, and sharing data
- Develop an action plan to address the needs of these users.

The Superfacility Project draws upon NERSC talent across all departments, with contributions from ESnet and the Computational Research Division, to produce the best possible outcome for experimental and observational science users. Initial superfacility engagements include the ALS, LCLS, and National Center for Electron Microscopy (BES); DESI, LZ, and Large Synoptic Survey Telescope Dark Energy Science Collaboration (HEP); and the Joint Genome Institute (BER). In 2018, NESAP for Data efforts addressed code performance issues of ALS users and the DESI experiment.

ENABLING EDGE SERVICES USING CONTAINERS ON SPIN

Spin, NERSC's Scalable Platform Infrastructure, was developed to enable NERSC users and staff to quickly deploy edge services on NERSC systems using Docker-based containers. Container-based approaches are becoming the preferred method to deploy and maintain services, and Spin is used to host science gateways, web applications, databases, workflow managers, and other edge services. In 2018, the Spin team offered four training sessions to nearly 80 users and used NERSC's new documentation system at **https://docs.nersc.gov/** to provide extensive documentation and hands-on examples to cover container utilization beyond basic Docker maintenance. Users hail from institutions such as Berkeley Lab, the Joint Genome Institute, UCLA, and the DESI partner Laborat.rio Interinstitucional de e-Astronomia in Brazil. By December 2018, Spin had grown to host more than 90 services in production and 140 services in development. Notable services include:

The Materials Project, an interactive portal for materials science used by researchers to deep dive into a large information set on known and predicted materials. The Materials Project has grown to serve a community of more than 65,000 users and has more than 1,000 active users per day. Originally hosted on a single virtual machine, the Materials Project migrated to a microservices infrastructure on Spin, with mongo databases also hosted at NERSC and the cloud, allowing easy horizontal scalability to meet increasing traffic and usage demands.

The Open Chemistry Data Platform, a prototype data portal that uses JupyterLab, NERSC's NEWT API, and custom software libraries to enable computational chemistry calculations. The platform was designed by Kitware in collaboration with Berkeley Lab's Computational Chemistry, Materials, and Climate Group through an SBIR.

NERSC's JupyterHub service, which has 300 users and manages an average of 100 notebook servers on any given day.

NERSC's Rstudio service, with more than 55 users.

UPDATE ON INDUSTRY PARTNERSHIPS

NERSC's industry partnerships are critical to maintaining and communicating relationships with key stakeholders. These partnerships bring together researchers from government labs, academia, and the private sector in a wide range of domains that connect scientific discovery with technological innovation. In parallel with these projects, which are designed to advance HPC technologies, NERSC collaborates with industry through applied science programs that bring HPC speed to challenging engineering and manufacturing problems.

In 2018 NERSC had 13 industry-led computing repos allocated through the DOE Mission Science, ASCR Leadership Computing Challenge, and NERSC Director's Reserve programs. One of the ALCC awards includes an umbrella allocation at NERSC in support of the DOE Advanced Manufacturing Office's HPC4Manufacturing program. A 2018 HPC4Manufacturing project was a continuation of 2017 work between Berkeley Lab's Computational Research Division and PPG Industries to decrease the energy required to paint automobiles, which consumes 70% of the total energy used in automobile assembly. This project involves developing key mathematical computational fluid mechanics algorithms, and deploying them using DOE supercomputing facilities, of which NERSC is a primary source, to model rotary bell

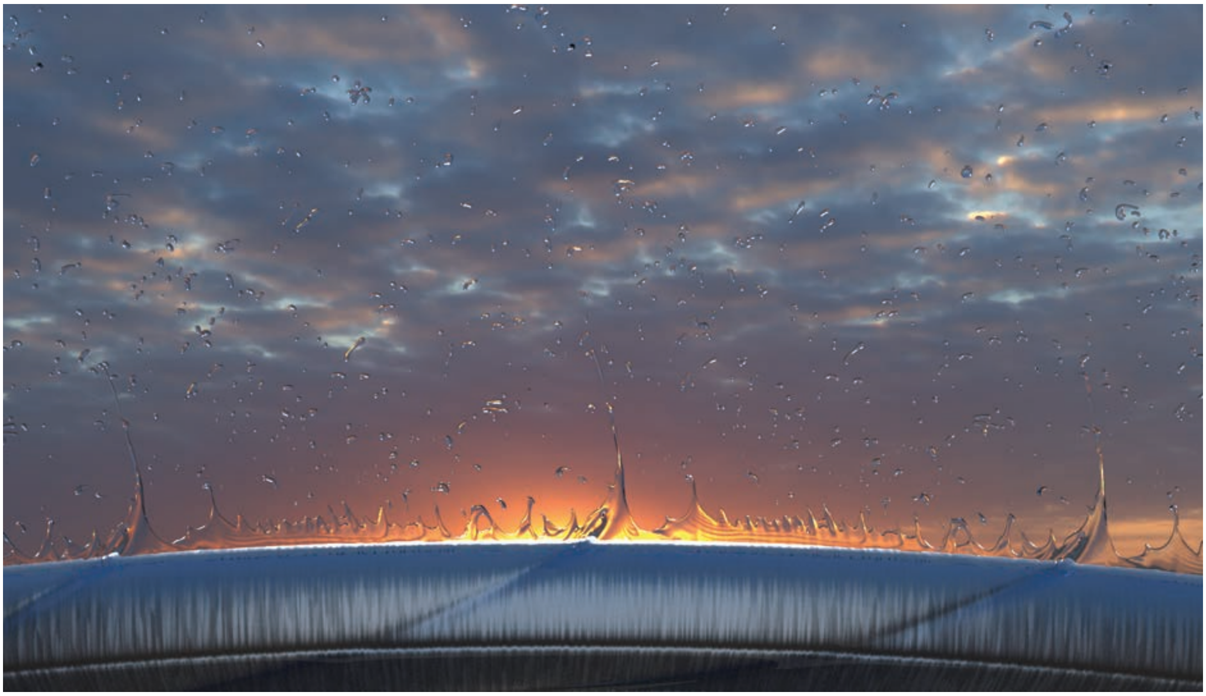

A simulation of high-resolution petascale modeling of liquid atomization, computed at NERSC using a state-of-the-art HPC numerical framework as part of a DOE HPC4Manufacturing project.

atomization devices that are used to paint automobiles. With this information, new coatings that atomize well at higher flow rates can be developed to increase productivity and reduce booth size, thus delivering significant energy savings and enhancing manufacturing competitiveness.

The Principal Investigator on this project is Reza Rock, Sr., of PPG Industries; the National Lab Partners are Robert Saye and James Sethian of Berkeley Lab's Mathematics Group.

NERSC HOSTS FIRST BIG DATA SUMMIT AT BERKELEY LAB

Some 75 researchers and data scientists attended the Big Data Summit 2018, hosted in July at Berkeley Lab by NERSC. Co-organized by NERSC, Intel, and Cray, the inaugural event featured a series of presentations highlighting results from advanced data analytics and management projects on NERSC's Cori supercomputer (see table, next page). The projects are being led by NERSC, Intel, Cray, and five Intel Parallel Computing Centers (IPCCs).

"Over the past year, the Big Data Center collaboration has made strong strides in pushing the frontier of capability applications on the Cori platform," said Prabhat, who leads the Data and Analytics Services team at NERSC and helped organize the event. "The Big Data Summit was held to share our results with the broader community, solicit feedback, and foster new collaborations. Active participation from Intel, Cray and the IPCCs — all members of the Big Data Center collaboration — made this event a success, and we are looking forward to new results in the future."

NERSC plans to hold a second Big Data Summit, with an emphasis on results from advanced data analytics and machine learning projects on the NERSC Cori system, in July 2019.

PRESENTATIONS AT THE BIG DATA SUMMIT

"Big Data Center"	*Prabhat, NERSC*
"HPC and ML — The New Frontier of Compute"	*Victor Lee, Intel*
"Data Science and Machine Learning at Supercomputer Scale"	*Mike Ringenburg, Cray*
"Distributed Training of Generative Adversarial Networks for Fast Detector Simulation on Intel Xeon HPC Cluster"	*Vikram Saletore and Hans Pabst, Intel; Damian Podareanu and Valeriu Codreanu, SURFsara B.V.; Sofia Vallecorsa and Federico Carminati, CERN*
"Deep Probabilistic Models for Astronomy"	*Jeffrey Regier, University of California, Berkeley*
"Project DisCo: Physics-based Discovery of Coherent Structures in Spatiotemporal Systems"	*Adam Rupe, James P. Crutchfield and Ryan G. James, University of California, Davis*
"Working Towards Distributed Inference Compilation at Scale"	*Frank Wood, University of British Columbia (formerly Oxford University)*
"Neutrino Detection with Graph Neural Networks"	*Nicholas Choma and Joan Bruna, New York University*
"Machine Learning and Topological Data Analysis: Applications to Pattern Classification in Fluid and Climate Simulations"	*Vitaliy Kurlin and Grzegorz Muszynski, University of Liverpool*
"Deep Dive into TensorFlow: Single and Multi-Node Optimizations on Intel Xeon"	*Ashraf Bhuiyan, Intel; Steven Farrell, NERSC*
"High Productivity Languages — Python"	*Heidi Pan, Intel; Rollin Thomas, NERSC*
"CosmoFlow: Using Deep Learning to Learn the Universe at Scale"	*Debbie Bard, NERSC*
"Data Management and I/O Deep Dive"	*Quincey Koziol, NERSC*

TRAINING WORKSHOPS HIGHLIGHT MACHINE LEARNING FOR SCIENCE

A new training emphasis at NERSC in 2018 was on machine learning and deep learning for science. The annual NERSC Data Day, which emphasized NERSC's users' machine learning capabilities, included introductory machine learning talks, an overview of the NERSC machine learning stack, practical guides to machine learning performance and scaling, and a hands-on machine learning scaling session. The event gave NERSC the opportunity to interact with many users and lab researchers to understand the challenges they face in applying machine learning to their domains, discussing with them NERSC's existing capabilities and how to extend them to better serve their needs. In a follow-up survey, 80% of attendees gave the Data Program a high rating, with requests for similar events in the future.

NERSC staff also led the inaugural Machine Learning for Science workshop, which was attended by more than 100 members of the Berkeley Lab science community. Also, at SC18, NERSC staff teamed up with Cray to conduct an all-day machine-learning tutorial that was attended by over 125 people. The tutorial featured a live Jupyter-based hackathon that enabled participants to run deep learning jobs on the Cori system. Both events were well received by the community; scientists attending were able to better appreciate the fundamentals of machine learning, deep learning, NERSC's software capabilities, real-world use cases, and best practices for scaling deep learning on HPC systems.

STEM COMMUNITY SUPPORT

NERSC engages with the broader community, including local colleges and universities and graduate students, through a variety of activities. NERSC has been a supporter of student cluster competitions and the DOE Computational Science Graduate Fellowship review meeting for many years, and in 2018 branched out to participate in the DOE's Cyberforce Competition for the first time. Since its inception, NERSC has hosted and mentored students — particularly those interested in STEM (science, technology, engineering, and math) — through school outreach, internships, the NERSC Student Program, and more. NERSC also conducts tours of the NERSC facility and computer room for high school, community college, and university groups. These tours sometimes include presentations and discussions about HPC and careers in computer science and scientific computation.

"With this event, we're hoping to build a pipeline of students who care about cybersecurity in a national lab environment and get them inspired to work at Berkeley Lab."

— Brent Draney, networking and security group lead at NERSC

CYBERFORCE COMPETITION

In 2018, NERSC joined with six other DOE labs in hosting 66 college student teams across the U.S. for the DOE CyberForce Competition, a computing event at

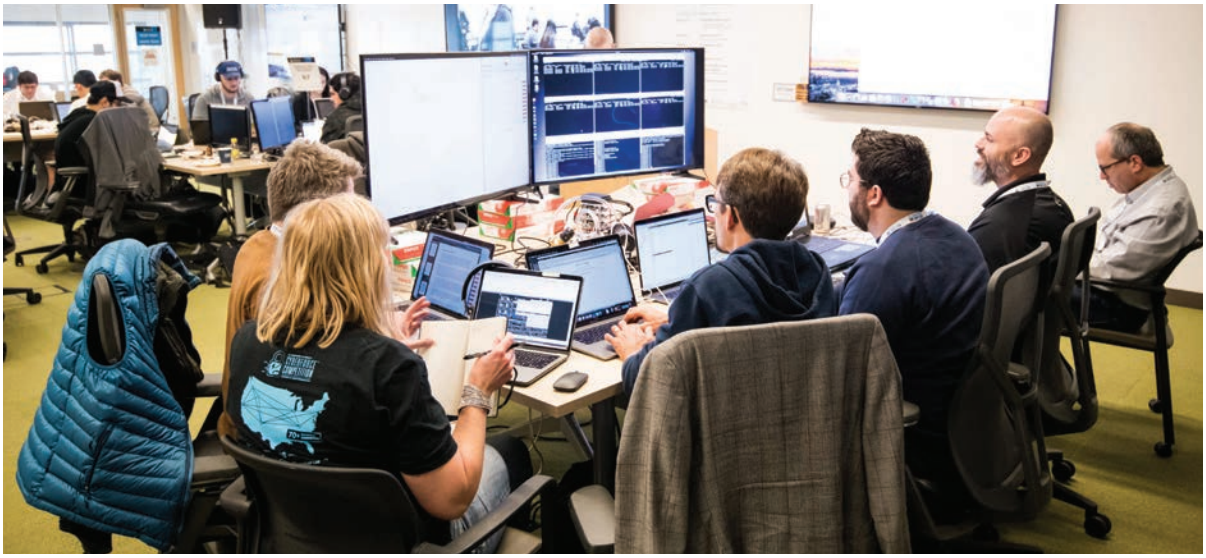

Four teams of students from California colleges competed in a regional CyberForce Competition at Berkeley Lab, organized by NERSC staff and volunteers.

which student "blue teams" compete against each other, acting as system administrators defending a simulated cyber-physical infrastructure from "red-team" attackers. This competition is one of many ways the DOE promotes the development of a workforce of cyber professionals with competencies relevant to the energy sector. It also raises awareness of cutting-edge cybersecurity and critical infrastructure innovation happening in the Department and across the national laboratories.

Since Argonne first hosted the event in 2016, enthusiasm among academia, industry, and government alike has grown rapidly and the competition has continually expanded, attracting more students and sponsors each year. The day-long event is held simultaneously on-site at multiple DOE labs, each hosting a number of teams from universities in their respective regions. Each lab provides logistics and the simulated environments, a volunteer "green team" acting as end users putting support demands on the blue teams, plus judges and technical support staff. The competition challenges students to respond to a scenario based on a real-world challenge of vital importance: protecting the nation's energy-related critical infrastructure from cyber threat.

In the 2018 competition, NERSC hosted blue teams from UC Berkeley, UC Davis, Cal State University San Bernardino, and Embry-Riddle Aeronautical University; red team volunteers from Cal State University San Bernardino, UC Berkeley, private sector professionals, and NERSC staff; and white and green team volunteers from across Berkeley Lab. The blue teams defended a simulated oil transportation network, power delivery system, and HPC system, while the green team stressed the students with demands to run IT services, fix problems, attend sudden meetings, and perform unrelated tasks such as making paper airplanes. Approximately 50 Berkeley Lab volunteers contributed to the success of the NERSC event.

Each national lab's regional winner and DOE's national winner took home a medal; the team from UC Davis came in first in the regional competition at Berkeley Lab and fifth overall in the national competition. All teams had the option to take home the high-performance cluster and industrial control system they configured for the competition.

The Laney College SC18 Student Cluster competition team and advisors: Jameelah Mercer, Steve Leak (advisor), Dan Gens (advisor), Nahuel Pregot, Khuong Dinh, Simon Tekleab Dair, Dave Liang, Jonnie Williams (advisor), and Nigel Brown.

According to Brent Draney, who leads the networking and security group at NERSC and served as the Lab's site lead in organizing the regional competition, all participants — whether they won or lost — also took home something else: real-world experience in a national lab setting. "With this event, we're hoping to build a pipeline of students who care about cybersecurity in a national lab environment and get them inspired to work at Berkeley Lab," he said.

SC18 STUDENT CLUSTER COMPETITION

First held in 2007, the Student Cluster Competition at the annual Supercomputing Conference was developed to immerse undergraduate and high school students in HPC. In this real-time, non-stop, 48-hour challenge, teams of students assemble small computer clusters on the exhibit floor, racing to complete real-world workloads across a series of applications.

In 2018, NERSC collaborated with Laney College — a community college in Oakland, California — and Cray to train and send a team of students to the Student Cluster Competition at SC18. The students at Laney come to California from across the world and from all stages of life, with many returning to college to study for a second

The 12-week CS summer student program offers undergraduate and graduate students in science and engineering the chance to gain research experience with NERSC, ESnet, and the Computational Research Division.

career. All of those who participated in the Student Cluster Competition were new to HPC and spent the year diving headlong into all aspects of HPC, from system architecture and administration to benchmarking, building and running scientific software, analyzing and reporting on results, and recovering from an (artificially induced) disaster.

Since the competition, most have continued their involvement in HPC, some locally in the ongoing Laney Supercomputing Club (formed around preparation for the competition), some as interns at NERSC and Berkeley Lab, and some moving on to their next educational institution and joining the local HPC community there.

SUMMER STUDENT PROGRAM

In June 2018, the Computing Sciences Summer Student Program welcomed dozens of college students from around the world to Berkeley Lab. During their two months at the Lab, the students had the opportunity to tour a state-of-the art supercomputer facility, peer into the heart of massive supernovae using math and computing, and participate in hands-on programming workshops.

The 12-week program, initially launched in 2010, offers undergraduate and graduate students in science and engineering the chance to gain research experience with ESnet, NERSC, and the Computational Research Division. In addition to completing a research project, the students are given the opportunity to attend weekly talks, tour Berkeley Lab facilities (including NERSC's computer room and the ALS), and present a poster outlining their summer project. As a member of this select group of students, their participation in the program will contribute important and valuable insights to their future professional development.

"Every year we get more than 90 students during the summer coming to the lab through the summer student program," said Osni Marques of the Computational Research Division, who chairs the program. "It is an opportunity for them to get to know each other, network, exchange ideas on their research assignments and be exposed to career opportunities. "

UC MERCED GRAD STUDENTS VISIT NERSC

In April 2018, a group of PhD students from UC Merced, all participants in the National Science Foundation's Interdisciplinary Computational Graduate Education (ICGE) program, visited Berkeley Lab to learn about the many ways computing can be applied to research problems across multiple science domains. Computational Research Division (CRD) Director David Brown hosted the group, and there were talks from NERSC's Prabhat and CRD's

UC Merced graduate students participating in the National Science Foundation's Interdisciplinary Computational Graduate Education program pose in front of the Cori supercomputer at NERSC.

Sean Peisert, Julian Borrill, Deb Agarwal, Mike MacNeil, Andy Nonaka, and Silvia Crivelli. The group also took a tour of the NERSC computer room.

Brown and Deputy Lab Director Horst Simon serve on the advisory board for the ICGE program. The program is designed to enhance student success and reduce graduate student attrition rates in the computational sciences, particularly for underrepresented minorities and first-generation students, while fostering computational and data analytic skills within an interdisciplinary framework.

"Fifty percent of PhD students nationwide drop out, and that number is even higher among minority students," said Prof. Marjorie S. Zatz, principal investigator for the UC Merced ICGE program and vice provost and graduate dean at the university. "By providing these opportunities and multiple mentorships, we provide a very safe environment for these students to learn and explore and manage interactions. It is a mix of learning coding, working on projects, successfully navigating through graduate school, and moving into the business world."

NERSC STAFF PARTICIPATE IN RESEARCH CONFERENCES, WORKSHOPS, AND MORE

In 2018, NERSC continued to be very active in the HPC community, with an emphasis on staying engaged in the research and development of next-generation computing systems. This includes participating in numerous vendor meetings associated with the APEX (NERSC, LANL, and SNL) and CORAL (ORNL, ANL, and LLNL) collaborations as they plan for their respective systems to be deployed in the 2019 to 2022 timeframes.

NERSC continues to be engaged in the Exascale Computing Project (ECP), in particular the PathForward and Hardware Integration areas. For PathForward, NERSC staff are the points of contact for milestone reviews. NERSC also participates in numerous status calls, ECP annual meetings, and the PathForward biannual face-to-face meetings. In addition, NERSC staff continue to work on ASCR research projects, either as PIs or collaborators.

NERSC staff have historically participated in the OpenMP and Fortran committees and continued to do so in 2018. In addition, in 2018 NERSC joined the OpenACC committee and the Standard Performance Evaluation Corporation (SPEC) to participate in the creation of the SPEC HPC 2020 benchmark suite. NERSC also participates in the MPI Forum.

NERSC is highly involved in HPC-relevant conferences, workshops, and standards committees each year. NERSC staff participate at all levels in these activities, organizing, chairing, and presenting numerous publications on its diverse set of research topics. These activities provide an opportunity to present research done by NERSC and track developments within the broader HPC and research communities. In 2018, NERSC staff contributed to hundreds of events, including:

International supercomputing conferences: International Parallel and Distributed Processing Symposium (IPDPS), International Supercomputing Conference (ISC), The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC).

ECP activities: ECP Hardware Evaluation (HE) review, Pathforward Spring and Fall reviews, Intel Pathforward deep dives.

Cray and Intel Xeon-Phi events: Cray User Group (CUG), Cray Compass System Management. Workshop, Intel eXtreme Performance Users Group (IXPUG), Intel Knights Landing hackathon.

GPU events: NVIDIA's GPU Technology Conference (GTC), OLCF GPU hackathons at UofC Boulder, NCSA, Brookhaven and OLCF.

Programmatic meetings: Salishan Conference on High Speed Computing, CORAL-2 Technical Evaluation Team meeting, Cray Quarterly meetings.

Python conferences for data analytics: PyCon, PyBay, JupyterCon, AnacondaCon.

Machine Learning: NeurIPS, ICMLA.

High Energy Physics events: HEPiX, International Conference on Computing in High Energy and Nuclear Physics (CHEP), Geant4 collaboration meeting, ATLAS collaboration meeting.

Climate events: AGU, AMS, Climate Informatics.

Security: USENIX Security Symposium, BroCon, CyberForce Defense Competition Meeting, New Security Paradigms Workshop.

Standards Committee meetings: MPI, OpenMP and Fortran. In June 2018 NERSC hosted the annual joint meeting of the ISO/IEC JTC1/SC22/WG5 international working group and the INCITS PL22.3 task group that together are responsible for maintaining and updating the Fortran programming language standard. Activities during this meeting achieved most of the remaining tasks leading to the publication of the "Fortran 2018" standard at the end of 2018.

Storage: Storage System and Input/Output Workshop (SSIO), USENIX Conference on File and Storage Technologies (FAST), Large Tape User Group (LTUG), Lustre User Group (LUG), OpenZFS Developer Summit, High Performance Storage Systems (HPSS) User Forum.

Center News

Setting the Stage for Perlmutter's Arrival

In 2018, NERSC initiated and completed a number of activities that improved operational security and efficiency and helped NERSC prepare to support the Perlmutter system and the center's increasingly broad workload.

MULTI-FACTOR AUTHENTICATION

To reduce the risk of compromised user accounts, in 2018 NERSC decided to make multi-factor authentication (MFA) mandatory for the allocation year 2019. This meant that in 2018, NERSC had to convert the bulk of its users from using passwords and certificates to using MFA at every authentication.

NERSC's strategy was to make it as easy as possible for users to sign up for and use MFA. To reduce the impact on user workflows and patterns of usage, NERSC developed an innovative approach to MFA called SSH Proxy that allows users to generate an SSH key that can be used to log in for a limited amount of time (default 24 hours). NERSC also developed thorough documentation, implemented an extensive communications campaign, and held virtual office hours to help NERSC users get started with MFA.

When designing the MFA implementation, NERSC carefully considered measures to increase the ease of use and minimize impact on users. Of particular interest were means that could also minimize additional burden to NERSC staff. Among the design goals of the MFA infrastructure was to make the enrollment process simple for users and to minimize the load on support staff. To this end, NERSC chose free, software-based, one-time password authenticators that users can download and install themselves and that did not require a hard token. NERSC added features to NIM (NERSC's identity management system) to allow users to provision their authenticator. Thus, users were able to install and configure MFA for their account with minimal impact on support staff.

Most NERSC users have access to a smartphone, so the primary free authenticator NERSC selected for users is Google Authenticator, which can be installed on most smartphones. For those users who do not own a smartphone, are unable to use a smartphone in their environment, or have a physical disability that prevents them from using a smartphone, NERSC added support for the Authy free authenticator, which is available for Windows and MacOS and also as a Chrome browser plug-in. The majority of users were able to avail themselves of at least one of these options, but some chose other options, such as OnePass or Firefox plugins, without NERSC support.

In addition, NERSC wanted to make it possible for users whose tokens were not immediately available to still be able to work. For this purpose, NERSC developed the ability in NIM for users to provision back-up one-time passwords that a user could print and store in a safe place to be used in case of emergency.

The majority of users set up an authentication token on their smartphone. A typical consumer changes to a new device at least every couple of years, potentially resulting in the user no longer being able to access NERSC resources. Without an easy way for users to manage their authentication tokens, including a way to recover after losing access to a phone or accidentally removing an active token, users might be unable to work and would require extensive staff assistance. To address circumstances in which a user loses access to all of their active tokens, for whatever reason, NERSC developed a self-service mechanism for users to set up a new token.

Because an MFA system is inherently more complex and potentially more fragile than conventional single password authentication, another design goal of the MFA infrastructure was to ensure that users would not be affected by authentication failures due to maintenance or outages on the infrastructure itself. Operations Technology Group and Infrastructure Services Group staff built in redundancy, monitoring, and emergency controls to mitigate this risk. First, the MFA infrastructure was built with a redundant “shared-nothing” design, enabling components to fail or be taken offline for maintenance non-disruptively. Then, a complete suite of real-time checks was added to NERSC’s center-wide monitoring framework that continuously exercise the MFA backend, performing authentication on frequent intervals and producing alerts for Operations to dispatch to staff. As a last resort, the system was designed with a “kill switch” that can be quickly activated to temporarily deactivate the token authentication subsystem during a loss of service. The redundancy, monitoring, and “kill switch” features of the MFA infrastructure have all been used to prevent, detect, and respond to events that would have otherwise impacted users.

A third infrastructure design goal was to minimize the impact of MFA on users’ ability to conduct their research, particularly for users with complex workflows that may require data transfer into or out of NERSC. MFA by its nature impedes automated workflows, as it requires a user to enter a new password with each connection. While the MFA policy allows for exemptions to the requirement in cases where MFA would have a significant impact on a user’s scientific work, NERSC wanted to minimize the number of users who had exemptions. A key innovation, SSH Proxy, was instrumental to achieving this goal.

SSH PROXY

As NERSC prepared to require MFA for users starting in the 2019 allocation cycle, one major concern was providing the best user experience possible while also increasing the security of the center. This was of particular interest for remote login to NERSC systems using the common secure shell (SSH) protocol.

Most NERSC users are accustomed to using SSH with “key pairs” to securely avoid repeated password prompts for remote logins during a work session. Previously, users registered an SSH key pair in the account management system, which allowed remote login to NERSC systems after proving ownership of the key pair (typically with a password). However, this method is not compatible with MFA, and it gave NERSC only coarse control over the lifetime and limitations of the use of that key pair. NERSC considered relying solely on the existing MyProxy grid certificate service, which could use MFA to supply certificates to be used with SSH, but unfortunately this requires a special SSH client, and the grid certificate technology is already deprecated by the broader grid community. What was required was an approach that would allow a user to continue to use standard SSH clients, avoid repeated password and MFA prompting, and, ideally, limit the lifetime and usable scope of the key to enhance compliance with stricter NERSC policies.

MFA TRANSITION PLAN

With a point-in-time deadline for users to switch to using MFA, NERSC was concerned about the impact on support staff should the thousands of NERSC users wait until the deadline and then enroll in MFA all at once. To minimize this impact, NERSC set a goal of having 1,000 existing users enrolled by December 2018. This approach not only served to reduce the number of users enrolling at once, but helped NERSC identify and correct any issues with the enrollment process ahead of the deadline. To achieve this goal, NERSC conducted an extensive communications campaign along several fronts and the User Engagement Group (UEG) developed thorough, clear documentation for users on how to enroll and use MFA.

The campaign to inform users of the plan kicked off in August 2018 with an announcement by the NERSC Director that MFA would become mandatory in the new allocation year. Over the following months, NERSC management sent memos to users encouraging enrollment and testifying to the ease of use, and NERSC published profiles on the NERSC website of users who had successfully transitioned to using MFA, which included a discussion of their research and the minimal impact of MFA on their workflows. In November and December 2018, UEG reached out to projects with high usage and low MFA enrollment, encouraging them to enroll before the deadline and offering help if there were any questions. UEG included reminders and announcements of new features in the NERSC weekly newsletter, and MFA was the topic for three of the weekly podcasts. Finally, UEG held several rounds of virtual "office hours" where support staff spent the day online in a Zoom video conference room, and users could dial into and ask questions.

The campaign was an unqualified success. Not only did NERSC achieve the goal of 1,000 users in early December 2018, by the allocation year rollover more than 2,500 users had enrolled. As many as 300 users per day had enrolled, with only a small percentage requiring help from support staff.

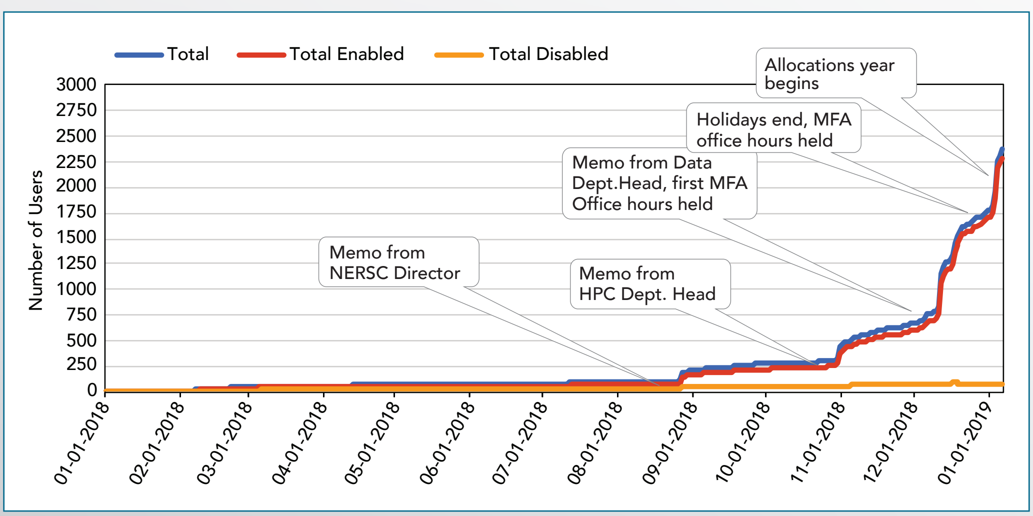

User enrollment in MFA in 2018, from initial rollout to mandatory requirement.

To meet this need, NERSC developed SSH Proxy, a service loosely modeled after MyProxy but with greater ease of use and flexibility. The SSH Proxy service is a RESTful-like service that allows a user to authenticate and obtain a limited-use SSH key in return. The SSH Proxy uses MFA, accepting the user's NERSC username, password, and one-time password. NERSC provided users a shell script that performs the authentication and places the resulting key pair in the appropriate location, ready to use. By default, this SSH key pair can be used for 24 hours, so a typical user would use the SSH Proxy once at the beginning of the day, and then use the limited-use key pair throughout the course of a day, very similar to users' previous usage patterns.

NERSC is also using the SSH Proxy to accommodate automated workflows and tasks. These use cases often need to run without human intervention for months at a time. Requiring the user to re-authenticate every 24 hours for workflows that operate continuously for days, weeks, or longer has a significant impact on usability. Therefore, we designed the SSH Proxy to support "scopes" that can have different constraints. A scope might allow a longer lifetime (say, a week) but narrowly limit where the key can be used, easing workflow operation while offering a positive risk/ benefit trade-off.

The SSH Proxy leverages existing functionality available in common SSH client software, making the approach durable. NERSC also chose to deploy the SSH Proxy service using Docker, which allowed us to easily make it redundant, scalable, and secure; to finely control software version releases; and to perform upgrades with minimal risk of user impact.

NERSC plans to release the SSH Proxy service as open source on GitHub. The required approvals for release have already been obtained. NERSC is planning to do a final code review prior to releasing it publicly. NERSC plans to publish details on this in the future and is interested in sharing details of this innovation with other facilities that may be interested in offering a similar capability for similar use cases such as automated workflows.

OMNI ANALYSIS IMPACTS FACILITY UPGRADE COSTS

NERSC was able to save about $2 million on an additional substation for its next HPC system, Perlmutter, by leveraging its innovative data-collection system, OMNI (Operations Monitoring and Notification Infrastructure).

Perlmutter, arriving in late 2020, will be the largest system to date at NERSC and requires a facility upgrade to support it. At the time of initial analysis for the upgrade, the recommended scope of work included a $2 million mechanical substation. But by using OMNI, NERSC conducted an analysis of operational data and determined that the new substation was not required.

The Berkeley Lab Facilities Master Specification document dictates that new substation requirements are to be determined by using the total peak power usage of each device (as specified by the manufacturer) that will be powered by the substation or by an analysis of the facility using at least one year's worth of operational data showing why an upgrade was not needed.

Initially, using the total peak power usage method indicated that NERSC would need to install the additional mechanical substation. But analyzing the OMNI data yielded a different result. OMNI has been collecting operational data of the facility, including computational and mechanical power usage, over the past two years. By analyzing multiple sources of OMNI data — such as job load on the current system, climate, the wet bulb temperature, and air temperature in relation to power — the Perlmutter facilities team was able to demonstrate that the actual demand on the mechanical substation was much lower than the total peak usage rating of each device.

Our analysis found that, from January 1, 2018 to March 31, 2019, the total power load of all of the compute and HPC non-compute substations was stable relative to the mechanical substation. As expected, in the warmer months there is a higher-than-average overall demand on the mechanical substation, but this demand never exceeded 1 megawatt.

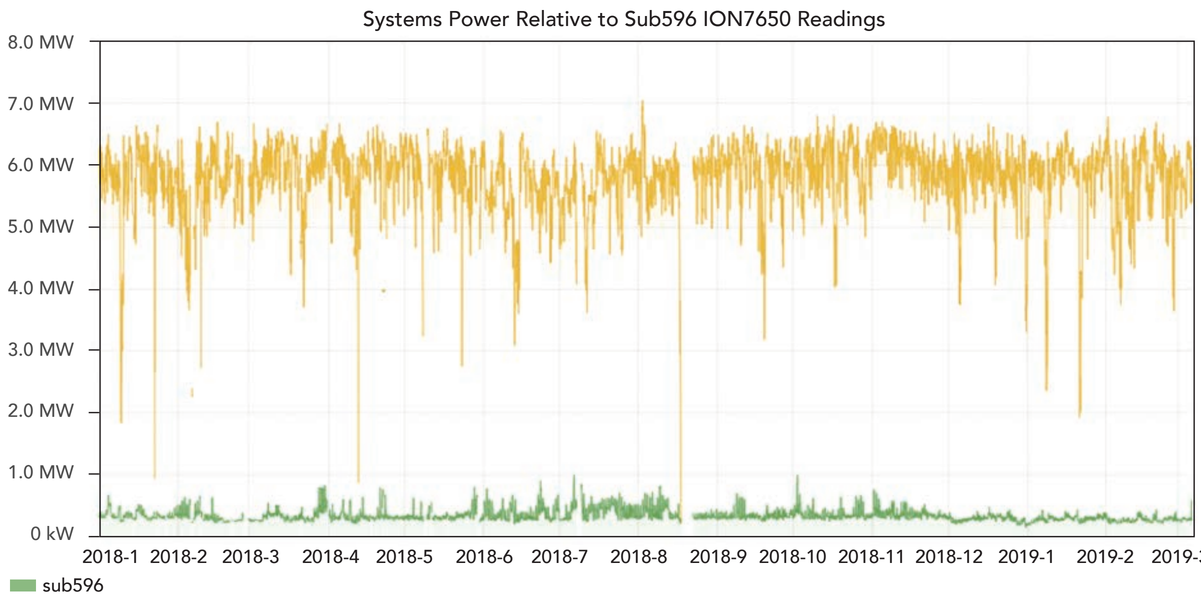

Total power load from the compute substations are stable relative to the mechanical substation.

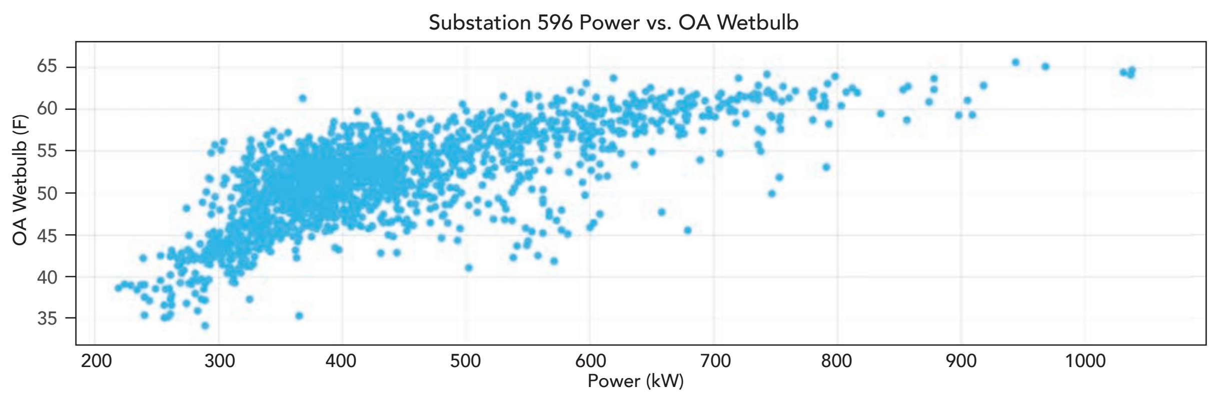

Scatterplot of outside air wet bulb temperature relative to mechanical substation 596 power.

We also looked at the load on the mechanical substation relative to the outside air wet bulb. The wet bulb temperature measures the evaporative cooling and accounts for both the temperature and the moisture in the air. Previous analysis of data in OMNI has shown that the outside air wet bulb temperature has the largest influence on the overall power usage effectiveness of the facility. By analyzing a scatterplot of outside air wet bulb temperature relative to mechanical substation power, NERSC staff verified that the higher readings (above an 800 kW power draw) occur on days when the wet bulb temperature is around 58° F and above (i.e., days when it is both hot and humid outside). Analysis of the plot also found that the majority of hours over the 15-month timespan are concentrated in that 300–500 kW power consumption range where the wet bulb reading fluctuates between 45–55° F.

This demonstrates that the Berkeley climate and evaporative cooling have a significant impact in maximizing energy efficiency and reducing unnecessary costs.

Ultimately, having the data available in a central location enabled the NERSC facilities team to analyze this data and conclude that Perlmutter's load would not require a new substation.

NERSC'S RESPONSE TO NORTHERN CALIFORNIA WILDFIRES

Wildfires in California during 2018 affected NERSC operations when smoke drifted toward the Bay Area. In November, smoke from the Camp Fire in Butte County, California, caused unhealthy air for two weeks in Berkeley and Oakland. NERSC's Shyh Wang Hall includes monitoring equipment to measure the particulate counts both inside and outside the facility. NERSC's response to the low air quality was to reduce the level of outside air into the Shyh Wang Hall facility to keep the particulate count low and protect both equipment and air quality for staff. Staff were also authorized to work from home if their home environment provided better air quality than was at the Lab.

Particulate matter can be harmful to hardware, where it can accumulate in delicate areas such as between disk platters, in tape drives, and on the tapes themselves, potentially voiding the warranty of the equipment. NERSC has particulate sensors that can detect different sizes. The Camp Fire started on November 8 and burned for more than two weeks; by analyzing data collected by these sensors during those two weeks, NERSC staff found that the increase in particulate count at 1 micron measured both outside and inside the building. The particulate counts inside the building remained lower after adjustments were made.

The Bay Area Air Quality Management District has neighborhood sensors that monitor the air quality in the surrounding areas. The calculation of air quality includes particulate matter, but also four other categories of air pollution. Decreasing outside air intake within the NERSC computer room and office areas allowed for better air quality inside than outside.

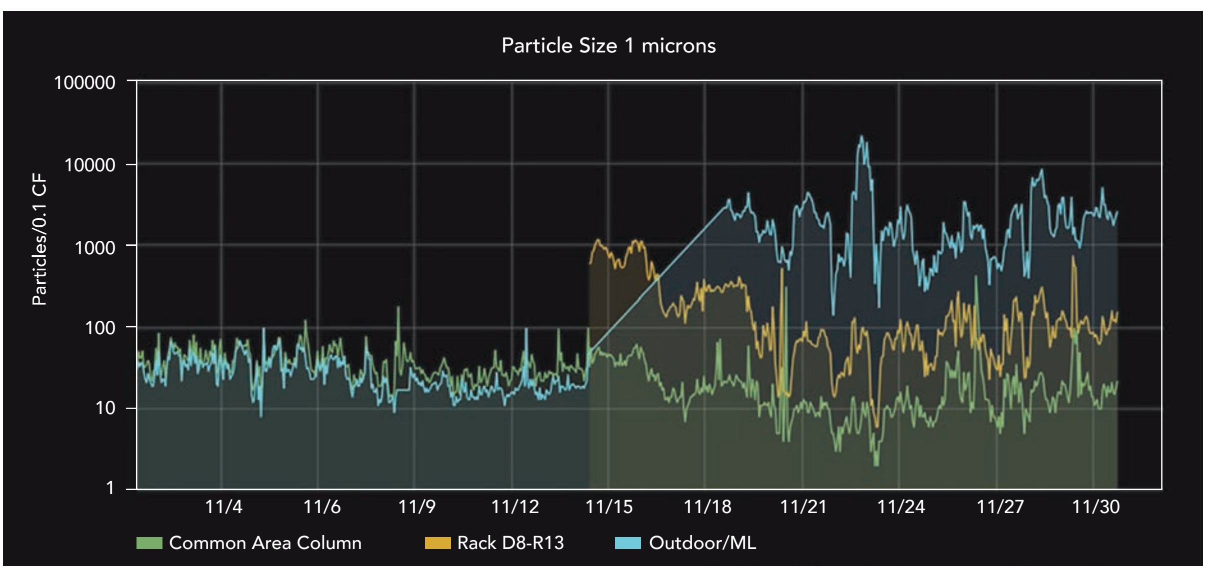

The various sizes of particles being monitored throughout the NERSC computer room floor during the Camp Fire in November 2018.

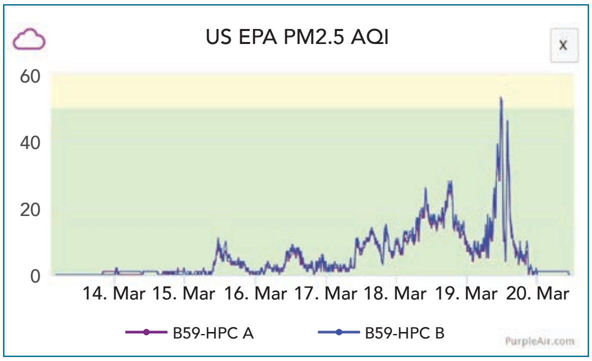

Air quality reading from the Shyh Wang Hall facility in March 2019.

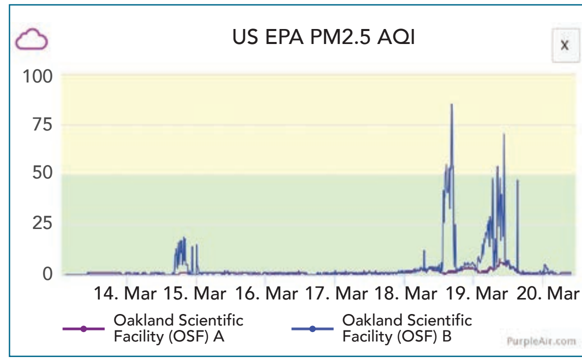

Air quality reading from the Oakland Scientific Facility in March 2019.

In comparison to the Shyh Wang Hall facility, the Oakland Scientific Facility, NERSC's previous home — which still hosted the HPSS tape libraries at the time of the Camp Fire — did not have a system to efficiently move air. The particulate count and air quality within the Oakland facility was poor enough that NERSC made the business decision to shut down the HPSS tape libraries for three days to ensure the hardware remained safe.

Following our experience during the Camp Fire, NERSC installed air quality sensors for monitoring the pollution within the facility, supplementing the particulate sensors and the neighborhood air quality sensors. We continue to monitor data from the air quality sensors in the Shyh Wang Hall facility and the Oakland facility

NEW TAPE LIBRARIES INSTALLED IN SHYH WANG HALL

When NERSC relocated from its former home in Oakland, California to Berkeley Lab's Shyh Wang Hall in 2015, not everything came with. The last remaining task in the heroic effort of moving the supercomputing facility and all of its resources is 43 years of archival data that's stored on thousands of high performance tapes in Oakland.

Those 130 petabytes of experimental and simulation data have to be electronically transferred to new tape libraries at Berkeley Lab, a process that will take up to two years — even with an ESnet 400 Gigabit "superchannel" now in place between the two sites.

Tape for archival storage makes sense for NERSC and is a key component of NERSC's storage hierarchy. Tape provides long-term, high capacity stable storage and only consumes power when reading or writing, making it a cost-effective and environmentally friendly storage solution. And the storage capacity of a tape cartridge exceeds that of more commonly known hard disk drives.

Allowing temperature and humidity to vary somewhat in the Shyh Wang Hall data center may be energy efficient, but it is harmful to tape. As a result, the existing libraries at OSF could not simply be relocated to Shyh Wang Hall without a solution for environmental containment. Options included off-site third-party data centers, cloud providers, and building out a climate-controlled subroom within the Shyh Wang Hall data center. Some of these were not ideal from a data stewardship perspective, and none of them were cost effective.

Ultimately, NERSC staff worked with IBM to deploy a new, cost-effective technology that allows the tapes to remain on-site at NERSC. The IBM tape libraries feature self-contained humidity and temperature control to provide the ideal operating environment. The new technology also delivers the capacity needed to keep pace with archival storage demand, providing a petabyte of storage per square foot. NERSC deployed two of these libraries in the Fall 2018 — currently the largest deployment of this technology in the world.

NERSC's new tape libraries have self-contained humidity and temperature control to provide the ideal operating environment.

Once the technology was selected, NERSC needed to figure out how to transition to the new libraries while minimizing impact to users. An electronic transfer of 130PB of data to new cartridges over an ESnet 400 Gb connection was initiated in 2018 and will take up to 18 months to complete. In addition, 3,000 cartridges were physically moved between the two locations. A total of 30 PB were moved this way across a 10-day timeframe through a process that required ejecting the tapes, having a courier transport them, and having staff import tapes in their new location (a process we called "sneakernet") — all with minimal disruption to NERSC users. In comparison, without sneakernet, moving this much data across the network would have taken 100 days.

NERSC RECOGNIZED BY NASA FOR CONTRIBUTIONS TO PLANCK MISSION

Launched in 2009, the Planck satellite was a joint project between the European Space Agency and NASA. It spent four years capturing the oldest light in the universe, the cosmic microwave background (CMB), which was emitted approximately 13.8 billion years ago.

These observations contributed to the most detailed map of the early universe ever created and supplied further evidence for an early phase of accelerated cosmic expansion, called inflation, during which the seeds of all structures, from galaxies to planets, were sown.

In addition to pushing the limits of cosmological research, Planck also pushed the boundaries of high performance computing, and one of the major U.S. contributions to this international collaboration was in data analysis. A pioneering agreement between the DOE and NASA guaranteed that NERSC would provide the high performance computing support essential to the mission's success. And for that support, in 2018 NERSC was recognized with a NASA Group Achievement Award.

"We could not have done all of the science data analysis that was done without high performance computing, and the largest Planck computing jobs were all run at NERSC. Planck, the greatest cosmology mission yet flown, depended critically on NERSC resources and the excellent technical support that we received throughout the project," said Charles Lawrence of NASA's Jet Propulsion Laboratory and project scientist for the U.S. Planck Project, who nominated NERSC for the award.

The Planck project has computed on six generations of NERSC flagship systems, with each system providing about a 10x increase in computing power. The project also spanned four NERSC Directors, including Horst Simon, Bill Kramer (interim), Kathy Yelick, and Sudip Dosanjh.

"George Smoot's Nobel Prize winning analysis of the Cosmic Microwave Background Explorer data using supercomputers of their time showed us the potential impact that these resources could have on CMB research. In the mid 1990s NERSC started supporting CMB research with projects like BOOMERANG, and the collaboration demonstrated significant scientific impact. The Planck mission continues this tradition of CMB computing at Berkeley Lab," said Horst Simon, Deputy Director of Berkeley Lab and former NERSC Director.

ICECUBE RESEARCH GARNERS BEST PAPER AWARD AT IEEE MACHINE LEARNING CONFERENCE

Nicholas Choma, a student at New York University whose work is supported in part by the Big Data Center collaboration (which includes NERSC), received the Best Paper Award at the 2018 International Conference on Machine Learning and Applications in Orlando, Florida.

The paper, "Graph Neural Networks for IceCube Signal Classification," involves some of the first research to apply a new form of deep learning on graphs and has yielded some very promising results for the IceCube project, noted Lisa Gerhardt, a big data architect at NERSC and a co-author on the paper. The paper was presented by Choma on behalf of the IceCube Collaboration, which provided the simulation and much of the analysis paradigm.

"Tasks involving the analysis of geometric data have recently gained prominence in the machine learning community, giving birth to a rapidly developing field of geometric deep learning," the team wrote. "In this paper we study the application of graph neural networks (GNNs) to the challenging problem of neutrino detection in the IceCube observatory." Through this work they were able to demonstrate the effectiveness of their GNN architecture on a task classifying simulated IceCube events, showing that it could outperform a traditional physics-based method and classical 3D convolution neural networks.

This work is the result of a collaboration between NERSC and members of Berkeley Lab's Nuclear Science Division (NSD) who are active in the IceCube Collaboration. The NSD contingent included Spencer Klein and Tomasz Palczewski, who provided the baseline (non-machine-learning) comparison point. Other co-authors were Wahid Bhimji and Prabhat of NERSC; Zahra Ronaghi, a former Berkeley Lab post-doc now with NVIDIA; Federico Monti of Universita dell Svizzera Italizana; Michael Bronstein of Imperial College; and Joan Bruna of New York University.

IceCube is a neutrino observatory located at the South Pole whose primary purpose is to look for high-energy neutrinos that are produced by the same cosmic particle accelerators that produce ultra-high energy cosmic rays.

IMPROVING COMMUNICATION WITH NERSC USERS

NERSC communicates with users through its website, MyNERSC user portal, and a comprehensive weekly email that is sent to all active users. NERSC also holds monthly NERSC User Group (NUG) webinars. The NUG Executive Committee is frequently queried for feedback about potential policy changes. NERSC staff are involved at various levels with many science communities, with some of the strongest engagements in materials science, nuclear physics, high energy physics, science at light sources, and genomics. In 2018, NERSC introduced some innovative new methods of communicating with users.

NERSC WEEKLY PODCAST

During 2018 NERSC created a weekly podcast entitled "NERSC User News," in which topics relevant to NERSC users are presented in an audio format. The podcasts are freely available on iTunes, Google Play, Spotify, and other platforms, where users can subscribe to the series. The podcast complements the weekly email newsletter — which emphasizes current news and opportunities — with magazine-like featured content.

The format of the podcast is a one-on-one interview with a NERSC expert on a particular topic. For example, the first podcast, released in May 2018 just before a lengthy maintenance on Cori, addressed the common user question of why a maintenance takes such a significant amount of time. NERSC strives for at least 10 minutes of content for each podcast, but some interviews have gone as long as 30 minutes. Other topics in 2018 included the HPSS system, the August medium-voltage switchgear maintenance that required a complete power outage to the whole building, machine learning, running jobs at scale, the Slurm scheduler, and a series introducing the new Perlmutter machine. Three podcast episodes covered MFA as it progressed from optional to required for the new allocation year and gave updates as functionality improved.

REAL-TIME SUPPORT WITH VIRTUAL OFFICE HOURS

Written user support tickets are an effective medium for solving complex problems: exact error messages can be cut and pasted, and users and consultants can experiment with a question or suggestion and record detailed answers. But some questions are better answered via a conversation; for example, when filling out an ERCAP application, helping a user to identify needs, and requesting appropriate resources are more easily achieved without the latency of a ticket.

The university tradition of "office hours" is an effective support model for such questions, but it is difficult when your "students" (or in NERSC's case, "users") are distributed across the U.S. and around the world. Fortunately, videoconferencing has reached a level of quality that enables a solution. For several days during the ERCAP process and again in the lead up to MFA becoming mandatory, NERSC held virtual office hours via the combination of a physical room, a large AV system, and Zoom conferencing. This allowed staff to help users in real time and, with the aid of screen sharing, provide fast and effective support while reducing what would otherwise be a large ticket load on NERSC staff.

During the transition to mandatory MFA, NERSC assisted more than 100 users via virtual office hours. Zoom breakout rooms enabled multiple consultants to help multiple users in parallel when volume was high.

IMPROVEMENTS TO DOCUMENTATION FRAMEWORK

In 2018, NERSC developed **https://docs.nersc.gov** to simplify updating and maintaining the user-facing web documentation. Good technical documentation is important for the effective use of any system and is a key part of the user experience. NERSC is widely recognized as a leader in providing in-depth examples on a broad range of topics, as evidenced by a high ranking in Google search results for many search terms. However, maintaining a comprehensive set of web content requires a significant investment of staff time, including tasks such as adding new content, ensuring continuing correctness as software and policy details change, and coordination with external sources.

https://docs.nersc.gov is designed to meet these challenges and more. NERSC evaluated how it was maintaining documentation and wanted to make it easier to add new content, encourage a more open review process, and enable others to easily contribute to the documentation. To achieve this, NERSC adopted the same processes used for source-code updates. The source content is hosted and maintained in an open, version-controlled repository. NERSC staff and users can contribute updates by issuing a pull request. This enables our users and non-NERSC staff stakeholders to propose contributions that can be reviewed by NERSC staff. For example, an OpenMPI developer at Los Alamos National Laboratory has contributed details and examples for use of this software on Cori.

The source of the site is based on Markdown and plain text files and rendered into a responsive static website, but it can also be easily viewed in plain text form. This format also enables continuous testing of the content: correctness of examples, valid links, etc. Finally, the site is fully instrumented to enable data-driven refinement and enhancement based on user interaction with the content.

NERSC STAFF SHOWCASE ANALYTICS EXPERTISE AT AGU FALL MEETING

NERSC's presence at the annual fall meeting of the American Geophysical Union — held December 10–14, 2018, in Washington, D.C. — was larger than ever before, thanks in part to the center's strides in machine- and deep-learning-based data analytics for two fundamental challenges in climate science: pattern recognition for tracking and discovery in large climate datasets, and emulation of complex dynamical processes.

With a total of 15 talks and posters, the NERSC team covered a wide range of analytics techniques, including machine learning, deep learning, applied topology, computational mechanics, and other novel statistical methods applied to problems in extreme weather and climate science. A significant portion of the work that was presented at the meeting involved research output from the Big Data Center and NERSC-supported internships. At the AGU meeting, all the NERSC graduate students funded by these programs presented either a talk or a poster or both.

These presentations highlight how NERSC is pushing the boundaries of advanced data analytics and bringing cutting-edge climate informatics solutions to the most pressing problems in climate science today.

"Towards Exascale Deep Learning for Climate Science" (Invited)	*Prabhat, et al*
"Lessons Learnt from Applying Deep Learning to Scientific Problems at NERSC" (Invited)	*Prabhat, et al*
"Using Deep Learning for Probabilistic Detection of Extreme Weather"	*Ankur Mahesh, et al*
"Probabilistic Detection of Extreme Weather Using Deep Learning Methods" (Invited)	*Ankur Mahesh, et al*
"NEX-AI: A Cloud and HPC Agnostic Framework for Scaling Deep Learning and Machine Learning Applications for Earth Science"	*Sangram Ganguly, et al*
"Tracking and Predicting Extreme Climate Events"	*Soo Kyung Kim, et al*
"Deep Learning Recognizes Weather and Climate Patterns" (Invited)	*Karthik Kashinath, et al*
"Physics-informed Generative Learning to Emulate Unresolved Physics in Climate Models"	*Jinlong Wu, et al*
"Towards Generative Deep Learning Emulators for Fast Hydroclimate Simulations"	*Adrian Albert, et al*
"Climate Science, Deep Learning, and Pattern Discovery: The Madden-Julian Oscillation as a Test Case"	*Benjamin Toms, et al*
"Project DisCo: Physics-based Discovery of Coherent Structures in Spatiotemporal Systems"	*Adam Rupe, et al*
"Towards Topological Pattern Detection Methods in Climate Science: Application to Atmospheric Blocking Events"	*Grzegorz Muszynski, et al*
"Deep Learning on the Sphere: Convolutional Neural Networks on Unstructured Meshes"	*Chiyu Jiang, et al*
"Using a Statistical Tropical Cyclone Genesis Model for Assessing Differences in Climate Scenarios and Geographic Basins"	*Arturo Fernandez, et al*
"Climatenet: Bringing the Power of Deep Learning to the Climate Community via Open Datasets and Architectures"	*Prabhat, et al*

Science Highlights

NERSC runs highly available, stable, complex, world-class HPC resources with extremely usable development and software environments, allowing its thousands of users to be very scientifically productive. With more than 2,500 publications in 2018, we can only share a sample of NERSC users' science here. The following were chosen to represent the breadth of scientific research and data-focused projects supported by NERSC.

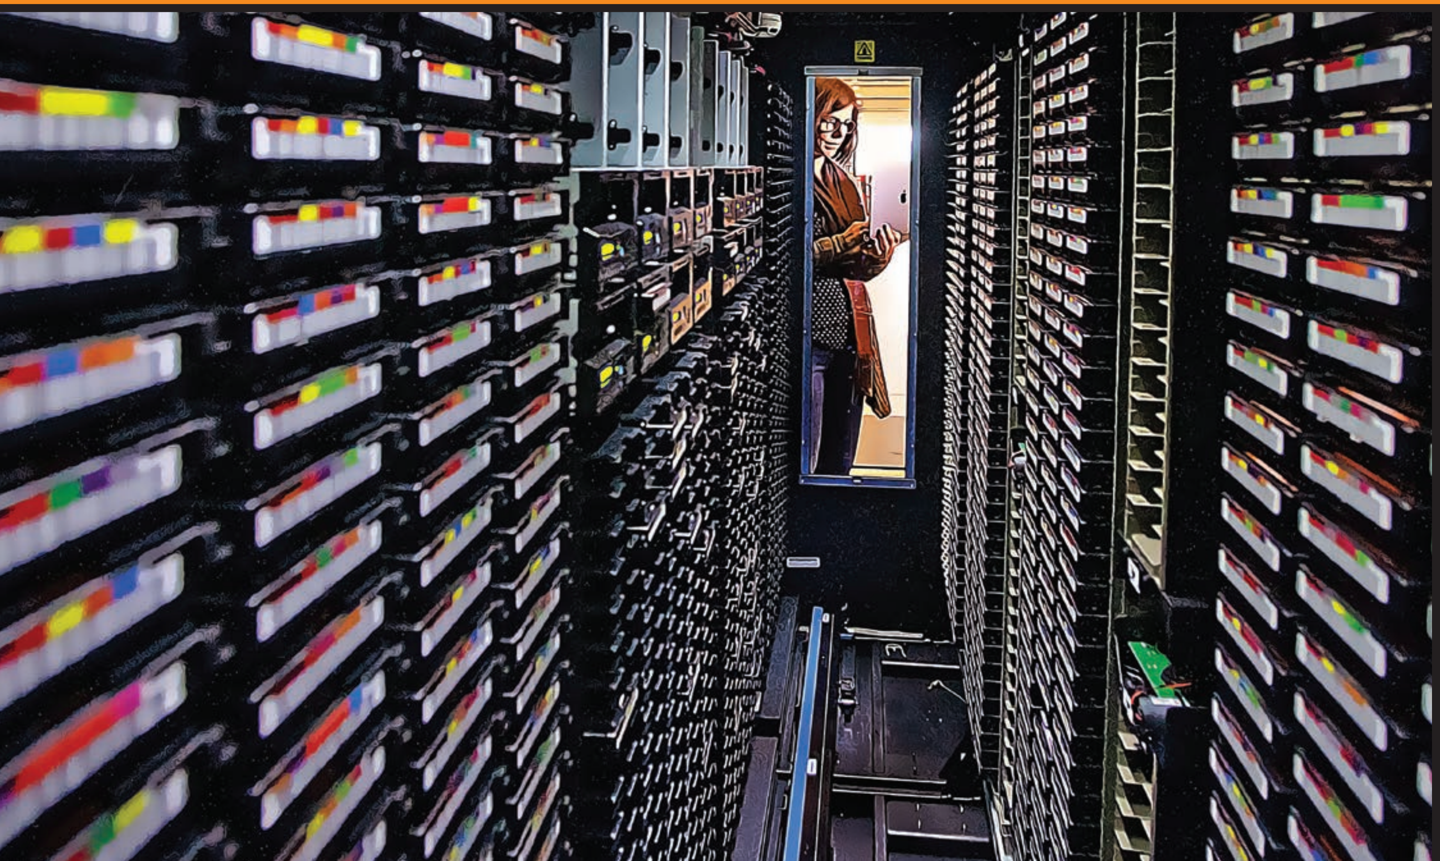

Self-healing Cement Aids Oil & Gas Industry

Basic Energy Sciences

SCIENTIFIC ACHIEVEMENT

Pacific Northwest National Laboratory (PNNL) researchers used a month's worth of simulations at NERSC to help them understand the characteristics of a unique "self-healing" cement — a composite of cement and a polymer — that can repair itself in as little as a few hours. Such a product would be a boon to the oil and gas industry; wellbore cement for geothermal applications has a lifespan of only 30 to 40 years, and when the cement cracks, which happens often, repairs can top $1.5 million per well.

SIGNIFICANCE AND IMPACT

With thousands of subsurface energy development wells running annually, this technology could have a dramatic impact on the cost of energy production. This research was featured on the cover of the journal *Applied Materials & Interfaces* in January 2018.

RESEARCH DETAILS

The researchers constructed a first-of-its-kind model based on density functional theory that simulates what occurs inside the cement/polymer system and used it in a series of simulations run on NERSC supercomputers. They used large-scale *ab initio* molecular dynamics to model about 900 atoms over a month's worth of computing, involving ~500,000 NERSC Hours. NERSC developed and provided optimized builds and build instructions for the code CP2K that was used for these calculations.

Principal Investigator: Vassiliki-Alexandra Glezakou, Pacific Northwest National Laboratory

Journal Citation: Manh-Thuong Nguyen, et al, *ACS Applied Materials and Interfaces*, 10:3011-3019; January 24, 2018, doi: 10.1021/acsami.7b13309

Full Story: http://bit.ly/NERSCcement

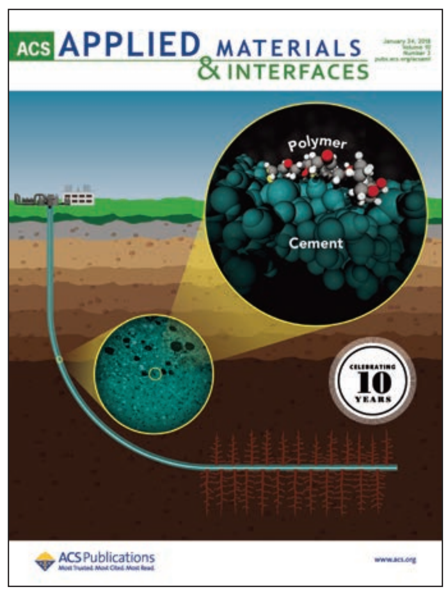

Large-scale *ab initio* molecular dynamics were used to determine the fundamental interactions between polymer and cement in self-healing cement composites.

Improving Data Reconstruction for the STAR Experiment

Nuclear Physics

SCIENTIFIC ACHIEVEMENT

Physicists from Brookhaven National Laboratory (BNL) and Berkeley Lab used Cori to reconstruct data collected from the STAR (Solenoidal Tracker at RHIC) nuclear physics research facility at BNL, an advance that dramatically reduces the time it takes to make detailed data available for scientific discoveries. STAR consists of a large, complex set of detector systems that measure the thousands of particles produced in each collision event. Detailed analyses of billions of such collisions have enabled STAR scientists to make fundamental discoveries and measure the properties of the quark-gluon plasma.

SIGNIFICANCE AND IMPACT

The researchers reconstructed multiple datasets collected by the STAR detector during particle collisions at the Relativistic Heavy Ion Collider (RHIC). By running multiple computing jobs simultaneously on the allotted supercomputing cores, the team transformed raw data into "physics-ready" data at the petabyte scale in a fraction of the time it would have taken using in-house high-throughput computing resources — even with a two-way transcontinental journey via ESnet. Preparing raw data for analysis typically takes many months, making it nearly impossible to provide such short-term responsiveness

RESEARCH DETAILS

Several technologies developed at NERSC allowed STAR to build a highly fault-tolerant, multi-step data-processing pipeline that could scale to a practically unlimited number of nodes with the potential to dramatically reduce the time it takes to process data for many experiments. For example, Shifter, a NERSC-developed container for HPC systems, has a perNodeCache feature that was used to solve the problem of slow startup times due to the need to build a database cache on Lustre. After fine-tuning their methods based on the initial tests, the researchers started scaling up, initially using 6,400 computing cores on Cori; in later tests they utilized 25,600 cores. The end-to-end efficiency of the entire process — the time the program was running (not sitting idle, waiting for computing resources) multiplied by the efficiency of using the allotted supercomputing slots and getting useful output all the way back to BNL — was 98 percent.

Principal Investigator: Jeff Porter, Lawrence Berkeley National Laboratory

Journal Citation: M. Mustafa et al, *J. Phys.*: Conf. Ser., 898, 082023

Full Story: http://bit.ly/NERSCstarexperiment

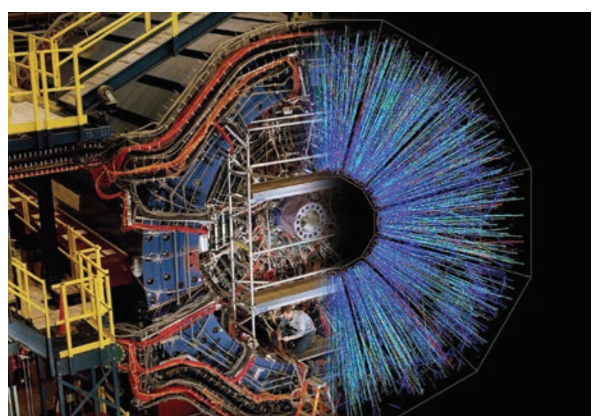

The STAR detector at the Relativistic Heavy Ion Collider nuclear physics research facility at Brookhaven National Laboratory.

Illuminating the Behavior of Luminous Blue Variables

Nuclear Physics

SCIENTIFIC ACHIEVEMENT

Three-dimensional simulations run at NERSC, the Argonne Leadership Computing Facility (ALCF), and NASA provided new insights into the behavior of a unique class of celestial bodies known as luminous blue variables (LBVs) — rare, massive stars that can shine up to a million times brighter than the Sun. This work was featured on the cover of *Nature* in September 2018.

SIGNIFICANCE AND IMPACT

Astrophysicists are intrigued by LBVs because their luminosity and size dramatically fluctuate on a timescale of months. They also periodically undergo giant eruptions, violently ejecting gaseous material into space. Although scientists have long observed the variability of LBVs, the physical processes causing their behavior are still largely unknown, and traditional one-dimensional models of star structure are inadequate for studying LBVs.

RESEARCH DETAILS

Researchers from UC Santa Barbara, UC Berkeley, and Princeton University ran 3D simulations to study three different LBV configurations. All the simulations included convection, the action when a warm gas or liquid rises while its cold counterpart sinks. During the outburst phase, the 3D simulations predict that convection causes a massive star's radius to irregularly oscillate and its brightness to vary by 10 to 30 percent on a timescale of just a few days — in agreement with current observations.

The team used millions of computing hours at NERSC on the Cori system, at ALCF, and at NASA. In addition to spending about 5 million CPU hours at NERSC on the early phase of the project, Jiang's team used another 10 million CPU hours running part of the 3D simulations. NERSC also helped the researchers develop and optimize methods and explore parameter space prior to making big runs at all three centers. NERSC consultants provided assistance, including helping debug and optimize I/O issues on Cori.

Principal Investigator: Yan-Fei Jiang, UC Santa Barbara

Jiang, Yan-Fei, et al, Nature, 561:498, September 27, doi: 10.1038/s41586-018-0525-0

Full Story: http://bit.ly/NERSClbvstars

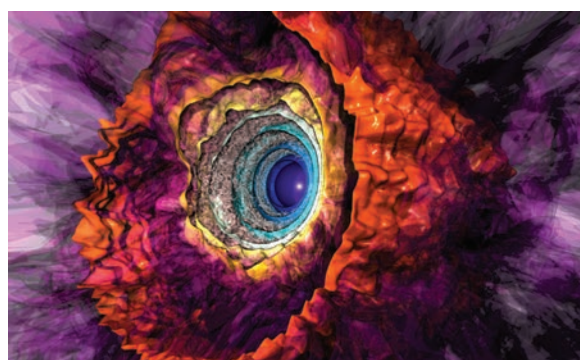

Intense light from the star's core pushes against helium-rich pockets in the star's exterior, launching material outward in spectacular geyser-like eruptions. The solid colors denote radiation intensity, with bluer colors representing regions of larger intensity. The translucent purplish colors represent the gas density, with lighter colors denoting denser regions.

Magnetic Islands Confine Fusion Plasmas

Fusion Energy Sciences

SCIENTIFIC ACHIEVEMENT

Magnetic islands, bubble-like structures that form in fusion plasmas, can grow and disrupt the plasmas and damage the doughnut-shaped tokamak facilities that house fusion reactions. Researchers from Princeton Plasma Physics Laboratory (PPPL) ran large-scale computer simulations at NERSC to produce a new model that could be key to understanding how the islands interact with the surrounding plasma as they grow and lead to disruptions. The multiscale kinetic interaction of the magnetic islands with a high-temperature background plasma was simulated in the realistic geometry of a tokamak fusion energy device for the first time.

SIGNIFICANCE AND IMPACT

The PPL research team found that, unlike in previous simplified simulations and analytic studies that assumed flattening of plasma pressure, the islands can allow particles to orbit across them, drive 3D plasma flows that suppress turbulence, and keep the plasma pressure from being flat, which is required to achieve sustained energy generation in the presence of unavoidable islands. Details of the simulated fluctuation, flow, and confinement are consistent with experimental observations from the KSTAR tokamak in South Korea. These findings could significantly influence our understanding of plasma disruption precursion events and how to suppress them.

RESEARCH DETAILS

Working with the XGC code developed at PPPL, the team modeled magnetic islands using plasma conditions from KSTAR. The plasma structure around the islands proved markedly different from standard assumptions, as did their impact on plasma flow, turbulence and plasma confinement, agreeing with experiments. This study used 6.2 million hours on Cori. XGC receives optimization support from the NERSC Exascale Scientific Applications Program. Going forward, a larger scale computer could allow the XGC code to start from the spontaneous formation of the magnetic islands and show how they grow, in self-consistent interaction, with the sheared plasma flow and plasma turbulence. The results could lead to a way to prevent disastrous disruptions in fusion reactors.

Principal Investigator: Amitava Bhattacharjee, Princeton University

Full Story: http://bit.ly/NERSCmagneticislands

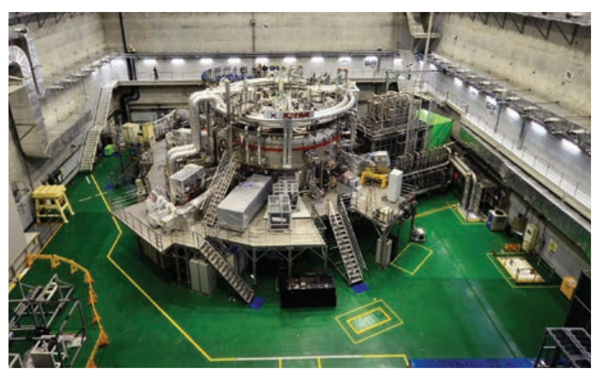

Korean Superconducting Tokamak Advanced Research facility (KSTAR)

NOvA Experiment Finds Evidence of Antineutrino Oscillation

High Energy Physics

SCIENTIFIC ACHIEVEMENT

The NOvA neutrino experiment, in collaboration with the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC-4) program and the HEPCloud program at DOE's Fermi National Accelerator Laboratory, was able to perform the largest-scale analysis ever to support the recent evidence of muon antineutrinos oscillating into electron antineutrinos over long distances, a phenomenon that has never been unambiguously observed.

SIGNIFICANCE AND IMPACT

Precisely measuring how neutrinos and antineutrinos change from one type into another, and then comparing them, will help scientists unlock the secrets that these particles hold about how the universe operates.

RESEARCH DETAILS

Running on NERSC's Cori and Edison supercomputers, NOvA used nearly 35 million core hours in a 54-hour period. This unprecedented amount of computing enabled scientists to do analysis of real data coming off the detector at a rate 50 times faster than achieved in the past and carry out some of the most complicated techniques used in neutrino physics, allowing them to dig deeper into the seldom seen interactions of neutrinos. The first round of analysis was done within 16 hours. Researchers were able to view what data from the experiment, and in less than six hours the entire collaboration could examine it. Executing the same analysis on a single desktop computer would take 4,000 years.

Members of the User Engagement, Data Science Engagement, Data and Analytics Services, and Computational Systems Groups coordinated large-scale reservations on the Cori and Edison supercomputers on an accelerated schedule to meet the needs of the NOvA experiment.

Principal Investigatosr: Oliver Gutsche, Jim Kowalkowski, Fermilab

Full Story: http://bit.ly/NERSCnova

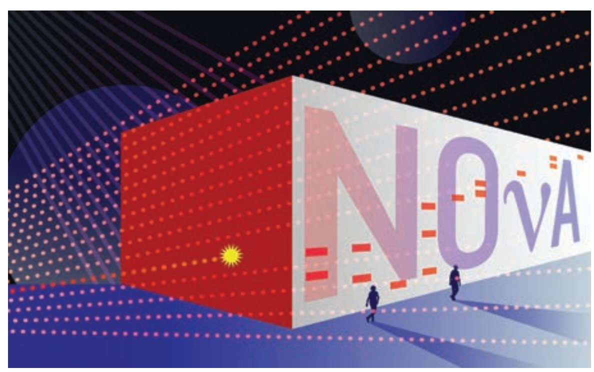

The Fermilab NOvA neutrino experiment has seen strong evidence of muon antineutrinos oscillating into electron antineutrinos over long distances, a phenomenon that has never been unambiguously observed.

Ice Sheet Modeling Exposes Vulnerability of Antarctic Ice Sheet

Biological and Environmental Research

SCIENTIFIC ACHIEVEMENT

Researchers from Berkeley Lab, Swansea University, and the University of Bristol used a highly resolved model of the Antarctic Ice Sheet to systematically examine the vulnerability of its floating ice shelves to regional collapse and the potential for large resulting contributions to sea-level rise.

SIGNIFICANCE AND IMPACT

The biggest uncertainty in near-future sea-level rise comes from the Antarctic Ice Sheet. Antarctic ice flows in relatively fast-moving ice streams. At the ocean, ice flows into enormous floating ice shelves that push back on their feeder ice streams, buttressing them and slowing their flow. Melting and loss of ice shelves due to climate changes can result in faster-flowing, thinning, and retreating ice, leading to accelerated rates of global sea-level rise.

RESEARCH DETAILS

To learn where Antarctica is vulnerable to ice shelf loss, the research team divided it into 14 sectors, applied extreme melting to each sector's floating ice shelves in turn, then ran their ice-flow model 1,000 years into the future for each case. They used the DOE SciDAC-supported BISICLES adaptive mesh refinement ice sheet model, which resolves flow down to 1km resolution, and found that the greatest vulnerability came from extreme thinning applied to any of the three ice shelves connected to West Antarctica, where much of the ice sits on bedrock lying below sea level.

This project comprised 35,000 years of high-resolution, full-continent Antarctic simulations and used more than 1 million CPU hours on NERSC's Cori system. BISICLES uses the Chombo block-structured adaptively refined rectangular grid solver framework, developed at Berkeley Lab. Chombo was the focus of optimization efforts undertaken through the NESAP program by former postdoc Andrey Ovsyannikov.

Principal Investigator: Dan Martin, Berkeley Lab

Journal Citation: Martin, D.F., et al, *Geophysical Research Letters*, 46, DOI 10.1029/2018GL081229.

Full Story: http://bit.ly/NERSCicesheet

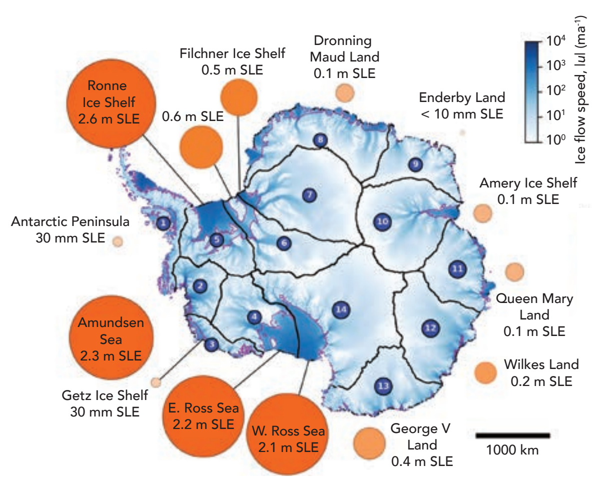

The BISICLES ice sheet model gives researchers insight into potential ice-shelf loss in each of its 14 sectors.

Batteries Get a Boost from 'Pickled' Electrolytes

Basic Energy Sciences

SCIENTIFIC ACHIEVEMENT

Researchers at Argonne National Laboratory used computer simulations at NERSC to help reveal the mechanism behind a common additive known to extend the life of lithium-ion batteries. At its heart is a chemical reaction similar to pickling. By better understanding the mechanism for the cathode-protective action by the phosphite, battery developers can find new ways to achieve and improve this process.

SIGNIFICANCE AND IMPACT

The energy storage system of choice for electric vehicles is the lithium-ion battery, but the cathode in these batteries limits their maximum energy storage capacity. Researchers have identified high-capacity cathode materials that operate at high voltage, but their long-term use remains problematic. As the cell charges and discharges, the materials react with the liquid electrolyte degrading it and stalling performance. Mixing a performance-boosting additive into the liquid electrolyte — in this case, trimethylsilyl phosphite (TMSPi) — can stall decomposition by adding a protective layer to the cathode surface. Until now, however, scientists didn't understand how TMSPi worked.

RESEARCH DETAILS

Using a series of simulations carried out at NERSC, ANL researchers found that TMSPi itself isn't directly involved in protecting the cathode. Instead a chemical derivative (PF_2OSiMe_3) does the job, slowly formin as the lithium salt in the electrolyte reacts with TMSPi. Uncovering this mechanism was a computationally intensive process; the researchers did two-step, high-level DFT calculations to achieve the necessary electronic convergence using several configurations of molecules and surfaces. A typical set of calculations on NERSC's Cori supercomputer used 288 processing cores and ran for over 24 hours. With so many different configurations to test, the researchers did around 60 simulations requiring about 41,500 computing hours.

Principal Investigatosr: Hakim Iddir, Argonne National Laboratory

Journal Citation: C. Peebles, et al, *J. Phys. Chem. C*, 2018122189811-9824, April 20, 2018.

Full Story: http://bit.ly/NERSCbatteriesboost

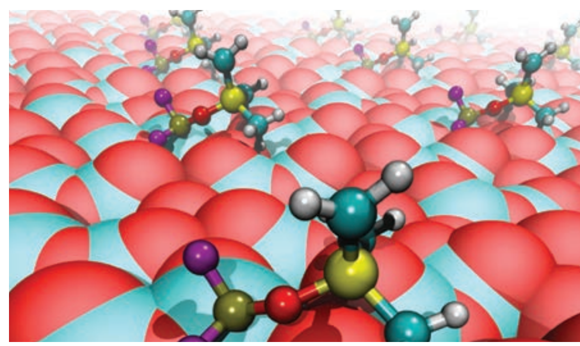

This illustration shows "pickled" electrolyte molecules (PF_2OSiMe_3) binding to reaction centers on a Li-On battery's cathode surface. For the ball-and-stick molecules attached to cathode surface, olive green indicates phosphorus (P); purple, fluorine (F); red, oxygen (O); and structure above oxygen, $SiMe_3$.

2018 Publications

In 2018, NERSC users reported more than 2,500 peer-reviews published papers that involved NERSC resources. The reported number of publications is based on self-reporting by users in a free-form field on applications for 2019 project renewals (this process is known as "ERCAP"). While this method has served NERSC well for a number of years, it also has shortcomings, including being time consuming for NERSC staff to process, increasing the possibility of duplicate or invalid entries, and missing publications produced by projects that did not apply for renewal.

Throughout 2018, NERSC worked to develop a new system that will be used for 2019 reporting and tracking. The new system does a keyword search of the Web of Science (WOS) database to gather publications that acknowledge NERSC and will have a web interface for users to enter publications that have not been identified by the WOS search. Users will be incentivized to enter data in two ways:

- Publications and summaries entered there will be eligible to be chosen as a NERSC highlight for the web or DOE and will also be eligible for one of the annual NERSC Achievement Awards.
- Data will be pulled directly from this new publications database and automatically incorporated into allocation requests for project renewals for 2020.

The system for collecting data from the WOS is already in place and is able to identify about 60+ percent of the total number of publications reported via ERCAP (**http://bit.ly/2018publications**). The new system will use DOIs as identifiers, thus ensuring uniqueness of reported publications and verification of publication dates.

To see a comprehensive list of publications and presentations by NERSC staff in 2018, go to **http://bit.ly/2018staffpublications**.

Allocation of NERSC Director's Discretionary Reserve (DDR)

In 2018, NERSC used its DDR pool of 850 million NERSC hours to enable strategic projects and to award time to projects that were selected through the competitive NERSC High-Impact Science at Scale program. DDR resources were allocated in three categories:

High-impact science at scale | Strategic projects | Industrial partners

The High-Impact Science at Scale program allocated time from the NERSC Director's Discretionary Reserve based on responses to a call for proposals. The call solicited applications from projects that could produce high-impact science using the unique capabilities of Cori at scale but did not have adequate time to accomplish their goals through the usual ERCAP allocations process. Each project had a NERSC staff member assigned to help the teams prepare their code and workflows for extreme scale computing. Staff arranged reservations of resources and debugged problems as they arose during the runs.

NERSC used some of its DDR allocation to enable Strategic Projects to support research in key areas of interest to NERSC, Berkeley Lab, and the Office of Science. For example, the "Simulation and Analysis for the GlueX Detector" project allowed the GlueX experiment to test their workflows at NERSC as a feasibility study.

In addition to these projects, NERSC supported allocations for 44 exploratory and 11 education projects. Exploratory projects are aimed at supporting new PIs to get started on their research or existing PIs developing new lines of research, and they must be approved by a DOE allocation manager before they are accepted. Education projects are requested by instructors at U.S. colleges and universities to train their students in HPC as part of a course of study. Altogether these 55 projects used less than 15 million NERSC Hours in 2018.

NERSC also engages with industry when the Center identifies opportunities to apply HPC capabilities to industry-sponsored open research.

NERSC HIGH-IMPACT SCIENCE AT SCALE PROJECTS FOR 2018

PI Name	Institution	Project Title	NERSC Hours Used	Max Nodes Used
Baron, Edward	U. Oklahoma	Scaling and I/O in the Largest 3D Radiative Transfer Calculation of the Solar Photosphere and Chromosphere	22,041,634	9,440
Feng, Yu	Berkeley Lab	Simulating the DESI, LSST and CMB-S4 Universe	21,465,997	8,300
Guo, Fan	Los Alamos Lab	3D Kinetic Simulations of Nonthermal Particle Acceleration in Flaring Magnetic Reconnection	41,456,737	2,048
Jiang, Yanfei	UC Santa Barbara	Global 3D Radiation Hydrodynamic Simulations of Mass Loss from Wolf-Rayet Stars	96,355,710	1,877
Karasev, Valentin	U. Rochester	Density functional theory calculations of the transport properties in the high energy density regime	13,219,233	512
Leung, Lai-Yung Ruby	PNNL	Ultra-high resolution atmospheric simulations of water cycle processes and extremes	42,542,819	1,349
Li, Jun	UC San Diego	Hot electron scaling and energy coupling in nonlinear laser plasma interactions	32,527,861	2,048
Marom, Noa	Carnegie Mellon	Data driven discovery of singlet fission materials	59,178,705	640
Marom, Noa	Carnegie Mellon	Computational discovery of singlet fission materials	10,325,507	1,024
Stanier, Adam	Los Alamos Lab	Probing the physics of magnetic reconnection: from fusion energy to space plasmas	12,923,529	2,048
Steinbrecher, Patrick	Bielefeld Univ Germany	Baryon-Strangeness Correlations in strongly interacting matter	65,534,437	2,600
Timmermans, Ben	Berkeley Lab	Multi-resolution surrogate modeling for simulated climate extremes and event attribution studies	38,282,744	1,056
Trebotich, David	Berkeley Lab	Chombo-crunch: extreme scale simulation of flow and transport in heterogeneous media	188,147,459	9,200
Wang, Xiao	Harvard Med School	Consensus Equilibrium Method for Extreme-Scale CT Image Reconstruction	14,134,177	9,464
Xantheas, Sotiris	PNNL	Benchmarks for NWChem/SPEC libraries	59,258,259	6,144

NERSC STRATEGIC PROJECTS FOR 2018

PI Name	Institution	Project Title	NERSC Hours Used
Antypas, Katerina	Berkeley Lab	NERSC Application Readiness for Future Architectures	5,211,188
Bhatele, Abhinav	Livermore Lab	Performance Analysis, Modeling and Scaling of HPC Applications and Tools	1,038,738
Borrill, Julian D.	Berkeley Lab	Modeling Polarized Galactic Foregrounds For Cosmic Microwave Background Missions	14,791,844
Bouchard, Kristofer E.	Berkeley Lab	Collaborative Research in Computational Neuroscience - CRCNS	25,642,527
Gerber, Richard A.	Berkeley Lab	NERSC fapesp (Brazil) collaboration	202,244
Jenn, Alan	UC Davis	Investigating bioenergy with carbon capture and sequestration using the U.S. TIMES	46,772
Nogales, Eva	Berkeley Lab	NERSC Supported Structuromics: solving structures to bridge the gap between sequence and function	1,357,150
Prabhat, Mr	Berkeley Lab	Data Intensive Science R&D	356,003
Romps, David M.	Berkeley Lab	Atmospheric dynamics	1,511,404
Solar-Lezama, Armando	MIT	Software Synthesis for High Productivity ExaScale Computing	47,273
Song, Yun S.	UC Berkeley	Population genomic analysis of whole genome sequencing data	148,268
Watson, Chip	Jefferson Lab	Simulation and Analysis for the GlueX Detector	14,215,240

NERSC INDUSTRIAL PARTNERS GIVEN A DDR ALLOCATION IN 2018

PI Name	Institution	Project Title	NERSC Hours Used
Fryman, Joshua	Intel Inc.	Intel Exascale R&D Pathfinding Architecture Studies (FFWD-2 and XStack)	0 (4M allocated)
Lau, Calvin	Tri Alpha Energy	Stability and Turbulent Transport of Field-Reversed Configuration (FRC) Plasmas	1,523,594
Nugent, Peter E. (Coordinator)	Berkeley Lab	High Performance Computing for Manufacturing	23,586,321

APPENDIX A:

NERSC Users Group Executive Committee

OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH

Jeff Hammond, *Intel*

Brian Van Straalen, *Lawrence Berkeley National Laboratory*

OFFICE OF BASIC ENERGY SCIENCES

Alexander Dunn, *Lawrence Berkeley National Laboratory*

Donny Winston, *Lawrence Berkeley National Laboratory*

Paul Kent, *Oak Ridge National Laboratory*

OFFICE OF BIOLOGICAL AND ENVIRONMENTAL RESEARCH

Samson Hagos, *Pacific Northwest National Laboratory*

Alicia Clum, *Joint Genome Institute*

David Wong, *U.S. Environmental Protection Agency*

OFFICE OF FUSION ENERGY PHYSICS

Christopher Holland, *University of California, San Diego*

Nathan Howard, *MIT*

Orso Meneghini, *General Atomics*

OFFICE OF HIGH ENERGY PHYSICS

James Amundson, *Fermilab*

Weiming An, *UCLA*

Zarija Lukic, *Lawrence Berkeley National Laboratory*

OFFICE OF NUCLEAR PHYSICS

Balint Joo, *Jefferson Lab*

David Lawrence, *Jefferson Lab*

David Tedeschi, *Lawrence Berkeley National Laboratory*

MEMBERS AT LARGE

Rob Egan, *Joint Genome Institute*

Jerry Jenkins, *Hudson/Alpha Institute for Biotechnology*

Pieter Maris, *Iowa State University*

Angelo Rossi, *University of Connecticut*

APPENDIX B:

Office of Advanced Scientific Computing Research

The mission of the Advanced Scientific Computing Research (ASCR) program is to discover, develop and deploy computational and networking capabilities to analyze model, simulate and predict complex phenomena important to the Department of Energy (DOE). A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

To accomplish its mission and address those challenges, the ASCR program is organized into two subprograms: Mathematical, Computational and Computer Sciences Research; and High Performance Computing and Network Facilities

The Mathematical, Computational and Computer Sciences Research subprogram develops mathematical descriptions, models, methods and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales. The subprogram also develops the software to make effective use of advanced networks and computers, many of which contain thousands of multi-core processors with complicated interconnections, and to transform enormous data sets from experiments and simulations into scientific insight.

The High Performance Computing and Network Facilities subprogram delivers forefront computational and networking capabilities and contributes to the development of next-generation capabilities through support of prototypes and testbeds.

Berkeley Lab thanks the program managers with direct responsibility for the NERSC program and the research projects described in this report:

ASCR PROGRAM

Barbara Helland, *Associate Director*

Julie Stambaugh, *Financial Management Specialist*

Lori Jernigan, *Program Support Specialist*

Tameka Morgan, *Administrative Specialist*

Christopher Miller, *AAAS Fellow*

FACILITIES DIVISION

Ben Brown, *Director (Acting); Physical Scientist, ESnet Program Manager*

Betsy Riley, *Computer Scientist, ALCC Program Manager*

Carolyn Lauzon, *Physical Scientist, ALCC Program Manager*

Claire Cramer, *Physical Scientist, REP Program Manager*

Sonia Sachs, *Computer Scientist, ALCF*

Robinson Pino, *Computer Scientist, REP Program Manager*

Christine Chalk, *Physical Scientist, ORLC Program Manager, CSGF Program Manager*

Sally McPherson, *Program Assistant*

RESEARCH DIVISION

Robinson Pino, *Director (Acting)*

Teresa Beachley, *Program Assistant*

Laura Biven, *Mathematician, Data & Visualization*

Randall Laviolette, *Physical Scientist, SciDAC Application Partnerships*

Thomas Ndousse-Fetter, *Computer Scientist, Network Research*

Ceren Susut, *Physical Scientist, SC Program SAPs*

Rich Carlson, *Computer Scientist, Collaboratories/Middleware*

Steven Lee, *Physical Scientist, Base, Math: Algorithms, Models, Data*

Lucy Nowell, *Computer Scientist, Computer Science*

Angie Thevenot, *Program Support Specialist*

APPENDIX C:

Acronyms and Abbreviations

ACM
Association for Computing Machines

ACS
American Chemical Society

ALCC
ASCR Leadership Computing Challenge

ALS
Advanced Light Source, Lawrence Berkeley National Laboratory

ANL
Argonne National Laboratory

API
Application Programming Interface

APS
American Physical Society

ASCII
American Standard Code for Information Interchange

ASCR
Office of Advanced Scientific Computing Research

BDC
Big Data Center

BER
Office of Biological and Environmental Research

BES
Office of Basic Energy Sciences

BMS
Building Management System

BNL
Brookhaven National Laboratory

CCM
Cluster Compatibility Mode

CERN
European Organization for Nuclear Research

CESM
Community Earth Systems Model

CFD
Computational Fluid Dynamics

CI/CD
Continuous Integration and Continuous Deployment

CMB
Cosmic Microwave Background

CO2
Carbon dioxide

CPU
Central Processing Unit

CRD
Computational Research Division, Lawrence Berkeley National Laboratory

CSE
Computational Science and Engineering

DARPA
Defense Advanced Research Projects Agency

DAQ
Data Acquisition System

DDR
Double Data Rate

DESI
Dark Energy Spectroscopic Instrument

DFT
Density Functional Theory

DNS
Direct Numerical Simulation

DOE
U.S. Department of Energy

DOI
Digital Object Identifier

DSL
Dynamic Shared Library

DTN
Data Transfer Node

DVS
Data Virtualization Service

ECP
Exascale Computing Project

EFRC
DOE Energy Frontier Research Center

EMSL
Environmental Molecular Science Laboratory, Pacific Northwest National Laboratory

EPSI
SciDAC Center for Edge Physics Simulations

ERD
Earth Sciences Division, Lawrence Berkeley National Laboratory

ERT
Empirical Roofline Toolkit

ESnet
Energy Sciences Network

eV
Electron Volts

FDM
Finite Difference Method

FEC
Forward Error Correction

FES
Office of Fusion Energy Sciences

FICUS
Facilities Integrating Collaborations for User Science

FLOPS
Floating Point Operations

FTP
File Transfer Protocol

GB
Gigabytes

Gbps
Gigabits Per Second

GPU
Graphics Processing Unit

GUI
Graphical User Interface

HDF5
Hierarchical Data Format 5

HEP
Office of High Energy Physics

HPC
High Performance Computing

HPC4Mfg
High Performance Computing for Manufacturing

HPSS
High Performance Storage System

HTML
Hypertext Markup Language

HTTP
Hypertext Transfer Protocol

I/O
Input/Output

IEEE
Institute of Electrical and Electronics Engineers

InN
Indium Nitride

IPCC
Intel Parallel Computing Center; Intergovernmental Panel on Climate Change

IPM
Integrated Performance Monitoring

iPTF
intermediate Palomar Transient Factory

ITER
An international fusion energy experiment in southern France

ITG
Ion Temperature Gradient

IXPUG
Intel Xeon Phi Users Group

JCESR
Joint Center for Energy Research Storage

JET
Joint European Torus

JGI
Joint Genome Institute

KNL
Knights Landing Processors

LED
Light-emitting Diode

LANL
Los Alamos National Laboratory

LCLS
Linac Coherent Light Source

LLNL
Lawrence Livermore National Laboratory

MIT
Massachusetts Institute of Technology

MODS
Monitoring of Data Services

MOF
Metal Oxide Framework

MPI
Message Passing Interface

MPP
Massively Parallel Processing

MSI
Mass Spectrometry Imaging

NCAR
National Center for Atmospheric Research

NESAP
NERSC Exascale Scientific Application Program

NEXAFS
Near Edge X-ray Absorption Fine Structure

NGF
NERSC Global Filesystem

NIH
National Institutes of Health

NIM
NERSC Information Management

NOAA
National Oceanic and Atmospheric Administration

NP
Office of Nuclear Physics

NPLQCD
Nuclear Physics with Lattice QCD

NSF
National Science Foundation

NUG
NERSC Users Group

NVRAM
Non-volatile Random Access Memory

OLCF
Oak Ridge Leadership Computing Facility

OMNI
Operations Monitoring and Notification Infrastructure

OpenMP
Open Multi-Processing

OpenMSI
Open Mass Spectrometry Imaging

PB
Petabytes

PDACS
Portal for Data Analysis services for Cosmological Simulations

PDSF
Parallel Distributed Systems Facility, NERSC

PF
Petaflop

PI
Principal Investigator

PiB
Pebibyte

PIC
Particle-In-Cell Simulations

POSIX
Portable Operating System Interface

PSII
Photosystem II

PNNL
Pacific Northwest National Laboratory

PPPL
Princeton Plasma Physics Laboratory

PUE
Power Usage Effectiveness

QCD
Quantum Chromodynamics

QUBITS
Quantum Bits

SC
DOE Office of Science

SciDAC
Scientific Discovery Through Advanced Computing

SDN
Software-defined Networking

SIAM
Society for Industrial and Applied Mathematics

SLURM
Simple Linux Utility for Resource Management

SPACK
Super Computing Package Manager

TACC
Texas Advanced Computing Center

TB
Terabytes

TOAST
Time Ordered Astrophysics Scalable Tools

TOKIO
Total Knowledge of I/O Framework

URL
Universal Resource Locator

VASP
Vienna Ab initio Simulation Package

VM
Virtual Machine

WAN
Wide Area Network

FOR MORE INFORMATION ABOUT NERSC, CONTACT:

Kathy Kincade

NERSC Communications
1 Cyclotron Road
Berkeley Lab, MS593024B
Berkeley, CA 94720-8148
Email: kkincade@lbl.gov
Phone: 510-495-2124
Fax: 510-486-4300

NERSC's Web Site

www.nersc.gov

NERSC Annual Report Editor

Kathy Kincade

Contributing Writers and Editors

Katie Antypas, Brian Austin, Elizabeth Bautista, Deborah Bard, Shane Canon, Brandon Cook, Chris Daley, Jack Deslippe, Doug Doerfler, Richard Gerber, Rebecca Hartman-Baker, Damian Hazen, Doug Jacobsen, Thorsten Kurth, Glenn Lockwood, Betsy MacGowan, Prabhat, David Skinner, Cory Snavely, Rollin Thomas, Linda Vue, Margie Wylie

Design

Design, layout, illustration, photography, and printing coordination:
Berkeley Lab Creative Services


OPPER

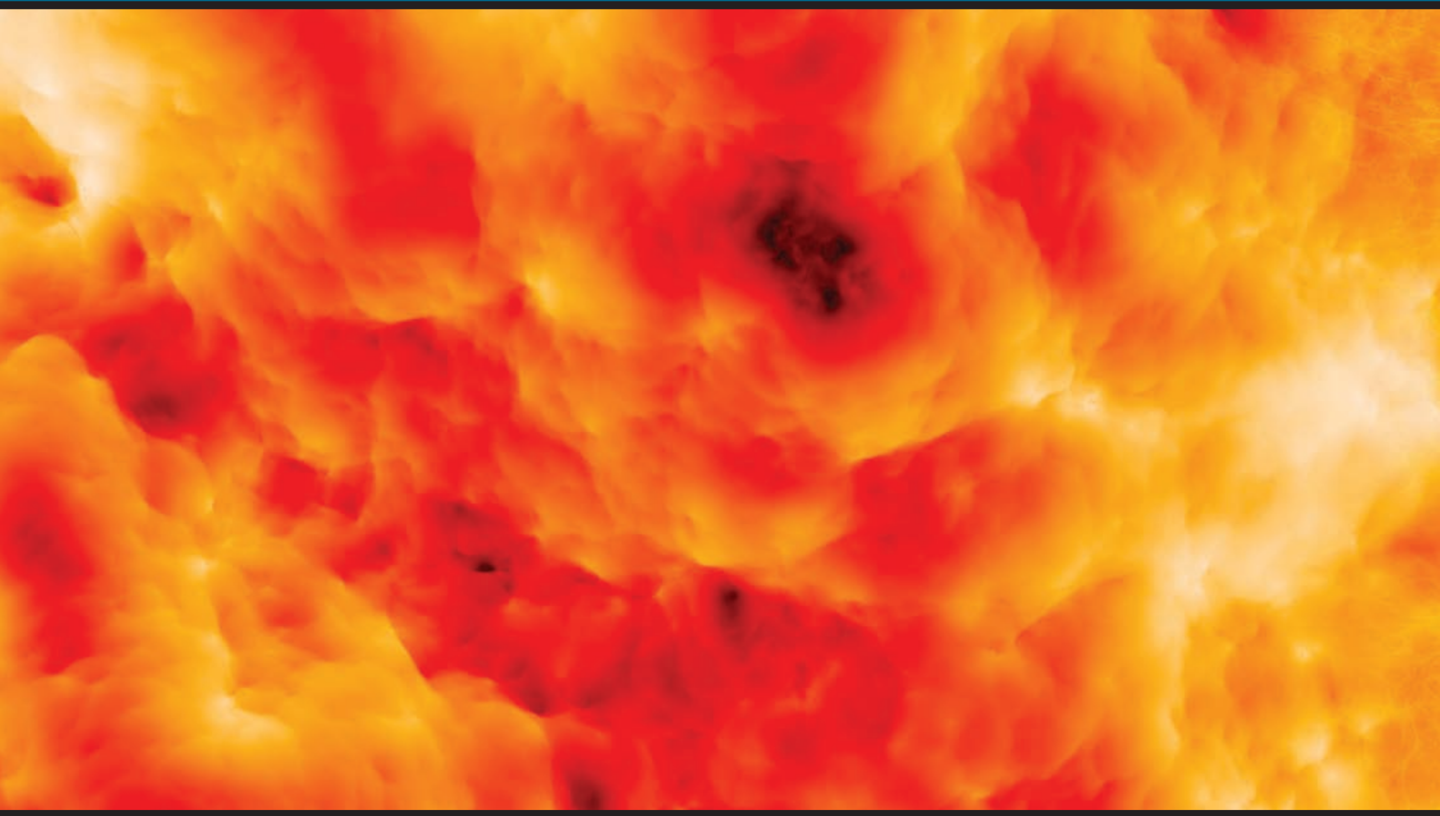

BERKELEY LAB

NERSC

19-NERSC-60017