U.S. DEPARTMENT OF
ENERGY
Office of Science

National Energy Research Scientific Computing Center

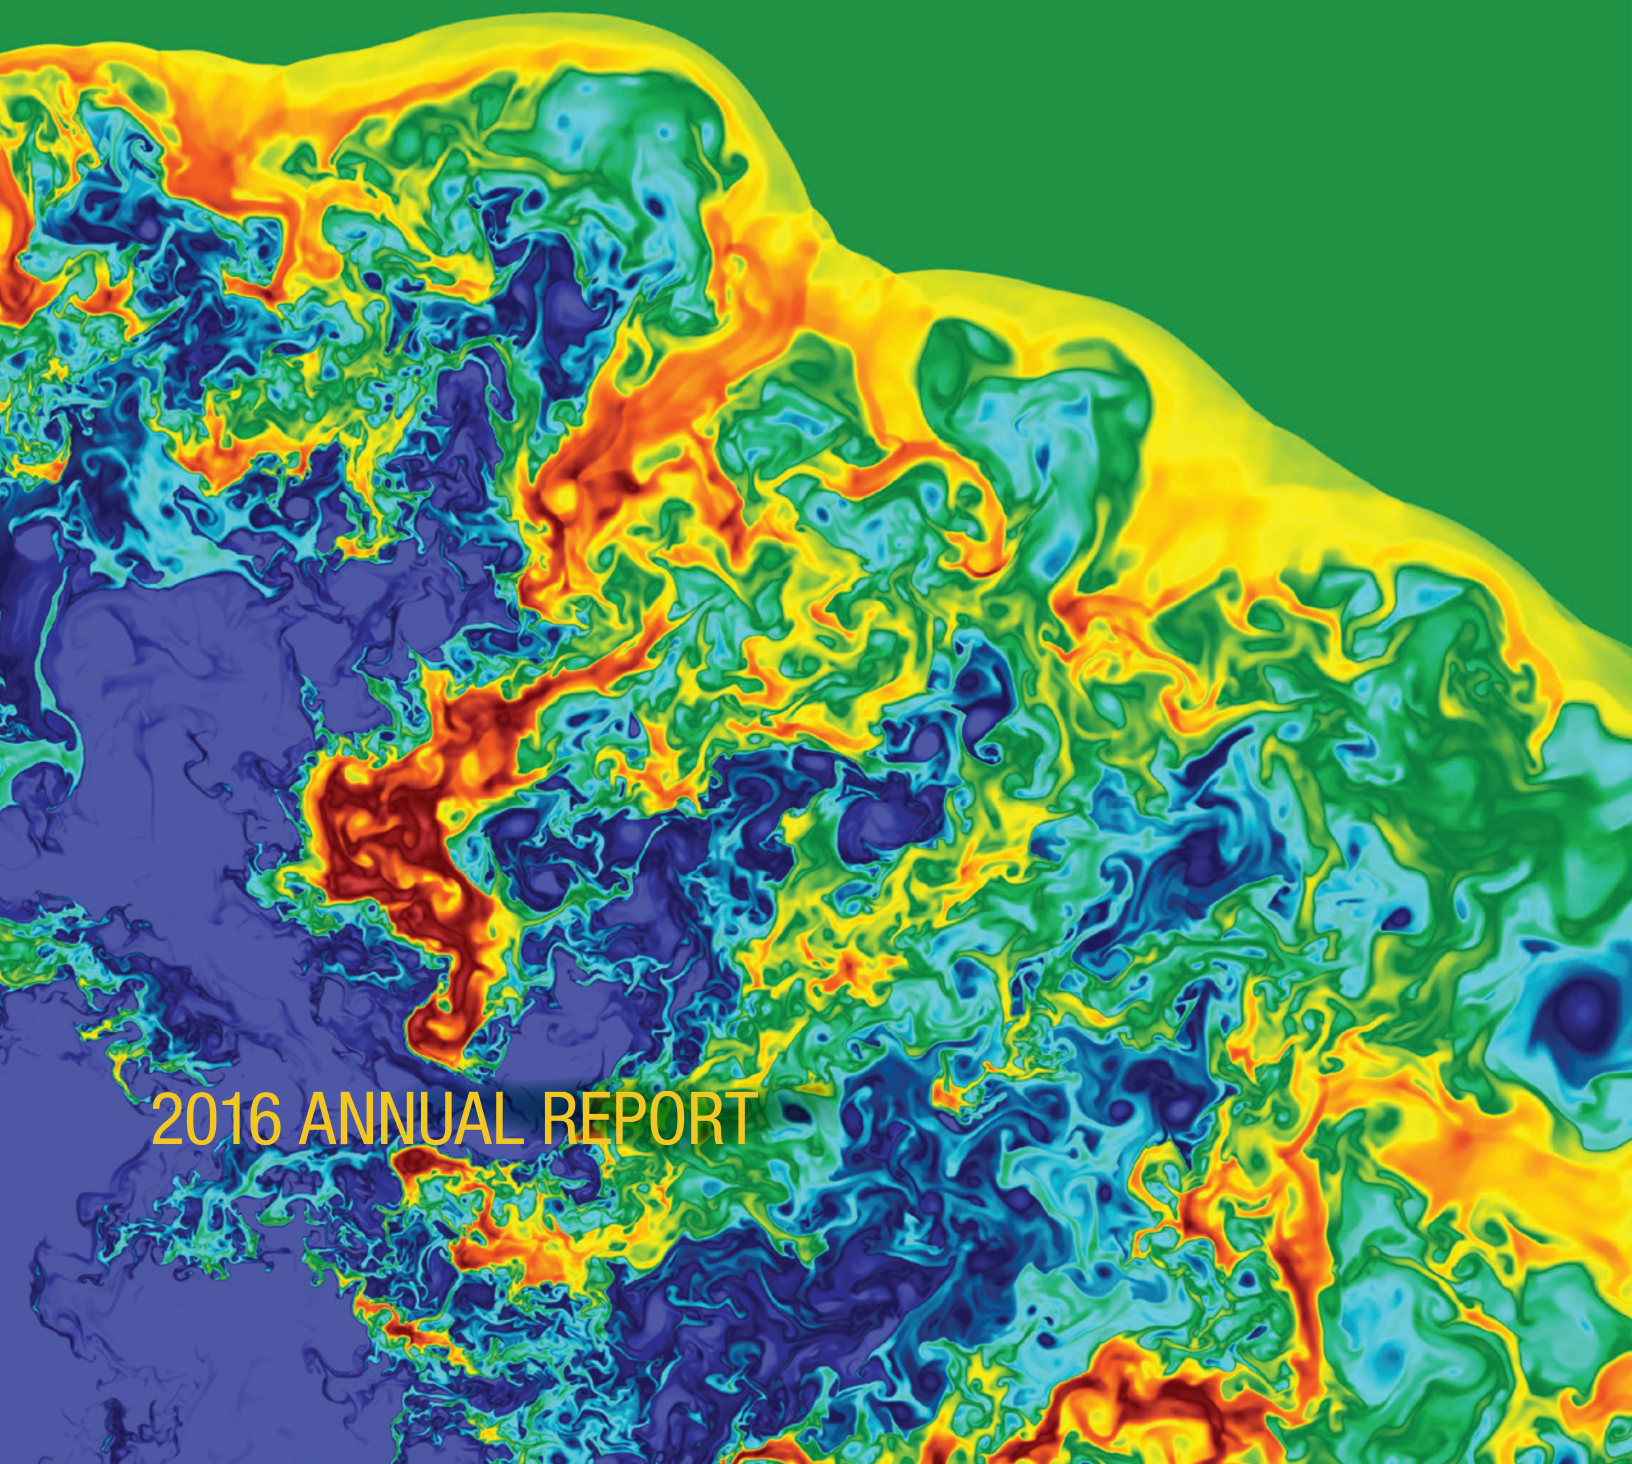

2016 ANNUAL REPORT

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:

FRONT COVER: Ken Chen, National Astronomical Observatory of Japan (story p. 24)

BACK COVER (LEFT TO RIGHT): H. Seo, M. Govoni and G. Galli, University of Chicago (story p. 26); David Trebotich, Lawrence Berkeley National Laboratory (story p. 23); Jean-Luc Vay, Lawrence Berkeley National Laboratory (story p. 25)

THIS PAGE, PAGE 44, PAGE 60-61: Marilyn Chung, Lawrence Berkeley National Laboratory

National Energy Research Scientific Computing Center

2016 ANNUAL REPORT

Table of Contents

Director's Note: Looking to the Future

While 2015 was a year of transition for the National Energy Research Scientific Computing Center (NERSC), 2016 was a year in which we looked to the future—from infrastructure, staffing and computing resources to helping users make a smooth transition to the next generation of supercomputing architectures and beyond.

For more than 40 years, NERSC has been providing cutting-edge HPC resources and expertise to the Department of Energy's Office of Science research community. The center specializes in advanced architectures capable of both world-class simulation and massive data analytics. NERSC has long been known for the productivity of its users; in 2016, our nearly 7,000 users reported more than 2,200 peer-reviewed published papers that involved NERSC resources.

One of our most notable achievements in 2016 was the full deployment of our newest supercomputer, Cori, a Cray XC40 with a peak performance of 30 petaflop/s that is the fifth fastest system in the world, according to the November 2016 TOP500. Preparing for and deploying Cori in the new Shyh Wang Hall facility at Berkeley Lab was a major undertaking that required dedicated support across all of NERSC throughout the year. The system was delivered in two phases:

- The Cori Data Partition (also known as Cori Haswell), installed in late 2015, comprises 12 cabinets and more than 2,300 Haswell compute nodes. It was customized to support data-intensive science and the analysis of large datasets through a combination of hardware and software innovations.
- Cori KNL, installed in mid-2016, added another 52 cabinets and more than 9,600 Intel Xeon Phi Knights Landing (KNL) compute nodes, making Cori the largest supercomputing system for open science based on these processors.

An increasing part of NERSC's workload involves data-intensive computing. Dramatically growing scientific datasets from experimental and observational sources require petascale-plus computing capabilities for analysis, and NERSC users are increasingly coupling large-scale simulations and data analysis. Fortunately, with new energy-efficient architectures such as the KNL processors that power Cori, we have the potential to address much larger problems than we could in the past. Toward this end, in 2016 a number of new capabilities were added to Cori to support users' increasingly complex workflows:

- A "burst buffer" for improved I/O
- A real-time queue that enables time-sensitive analyses of data
- Software defined networking to allow scientists to co-schedule networking bandwidth, compute resources and burst buffer bandwidth
- Container-based tools to help users run a wider range of software more easily and securely.

Another key area for NERSC in 2016 involved continuing to transition the user community to Cori's manycore architecture through NESAP, the NERSC Exascale Science Application Program. NESAP was launched in 2014 as a collaboration of NERSC staff and post-docs, application developers and third-party library and tool developers (Intel and Cray). Throughout 2016 NERSC held more than 10 NESAP "dungeon sessions" to assist users in optimizing their codes for a number of scientific applications running on Cori. NESAP applications are now reaping the benefits of these activities, with optimized NESAP runs averaging 3x better performance on Cori KNL than they would have achieved without the optimization efforts.

Following these successes, in 2016 NERSC launched NESAP for Data, which, like NESAP, joins application teams with resources at NERSC, Cray and Intel. While the initial NESAP projects involve mostly simulation codes, NESAP for Data targets science applications that process and analyze massive datasets acquired from DOE-supported experimental and observational sources, such as telescopes, microscopes, genome sequencers, light sources and particle physics detectors.

Part of NERSC's strategy around application performance on Cori also includes leveraging community expertise to help NESAP teams and contributing knowledge and techniques developed through NESAP back to the HPC community. Much of this has been accomplished through NERSC's involvement with the Intel Xeon Phi Users Group and the NERSC User Group. In addition, in 2016 NERSC held two dozen user training events and workshops designed to help users take full advantage of NERSC's evolving HPC resources. And more than 50 NERSC staff members gave presentations and tutorials at conferences and workshops in the U.S. and internationally, further helping us build relationships and facilitate communication with the HPC community and our growing user base.

As we look ahead to the next 12 months, we are excited about our involvement in several key DOE initiatives that benefit the HPC community, such as the Exascale Computing Project and HPC4Mfg. Also, in partnership with the ALCF and OLCF, we have completed the remaining exascale requirements reviews workshops, targeting the 2020 and 2025 time frames. And we are moving forward with NERSC-9, a system due in 2020 that will meet the needs of extreme computing and data users by accelerating workflow performance, plus provide a vehicle for the demonstration and development of exascale-era technologies.

All of this bodes well for NERSC's future and the future of scientific computing and the DOE Office of Science.

Sudip Dosanjh
NERSC Division Director

NERSC User/Usage Demographics

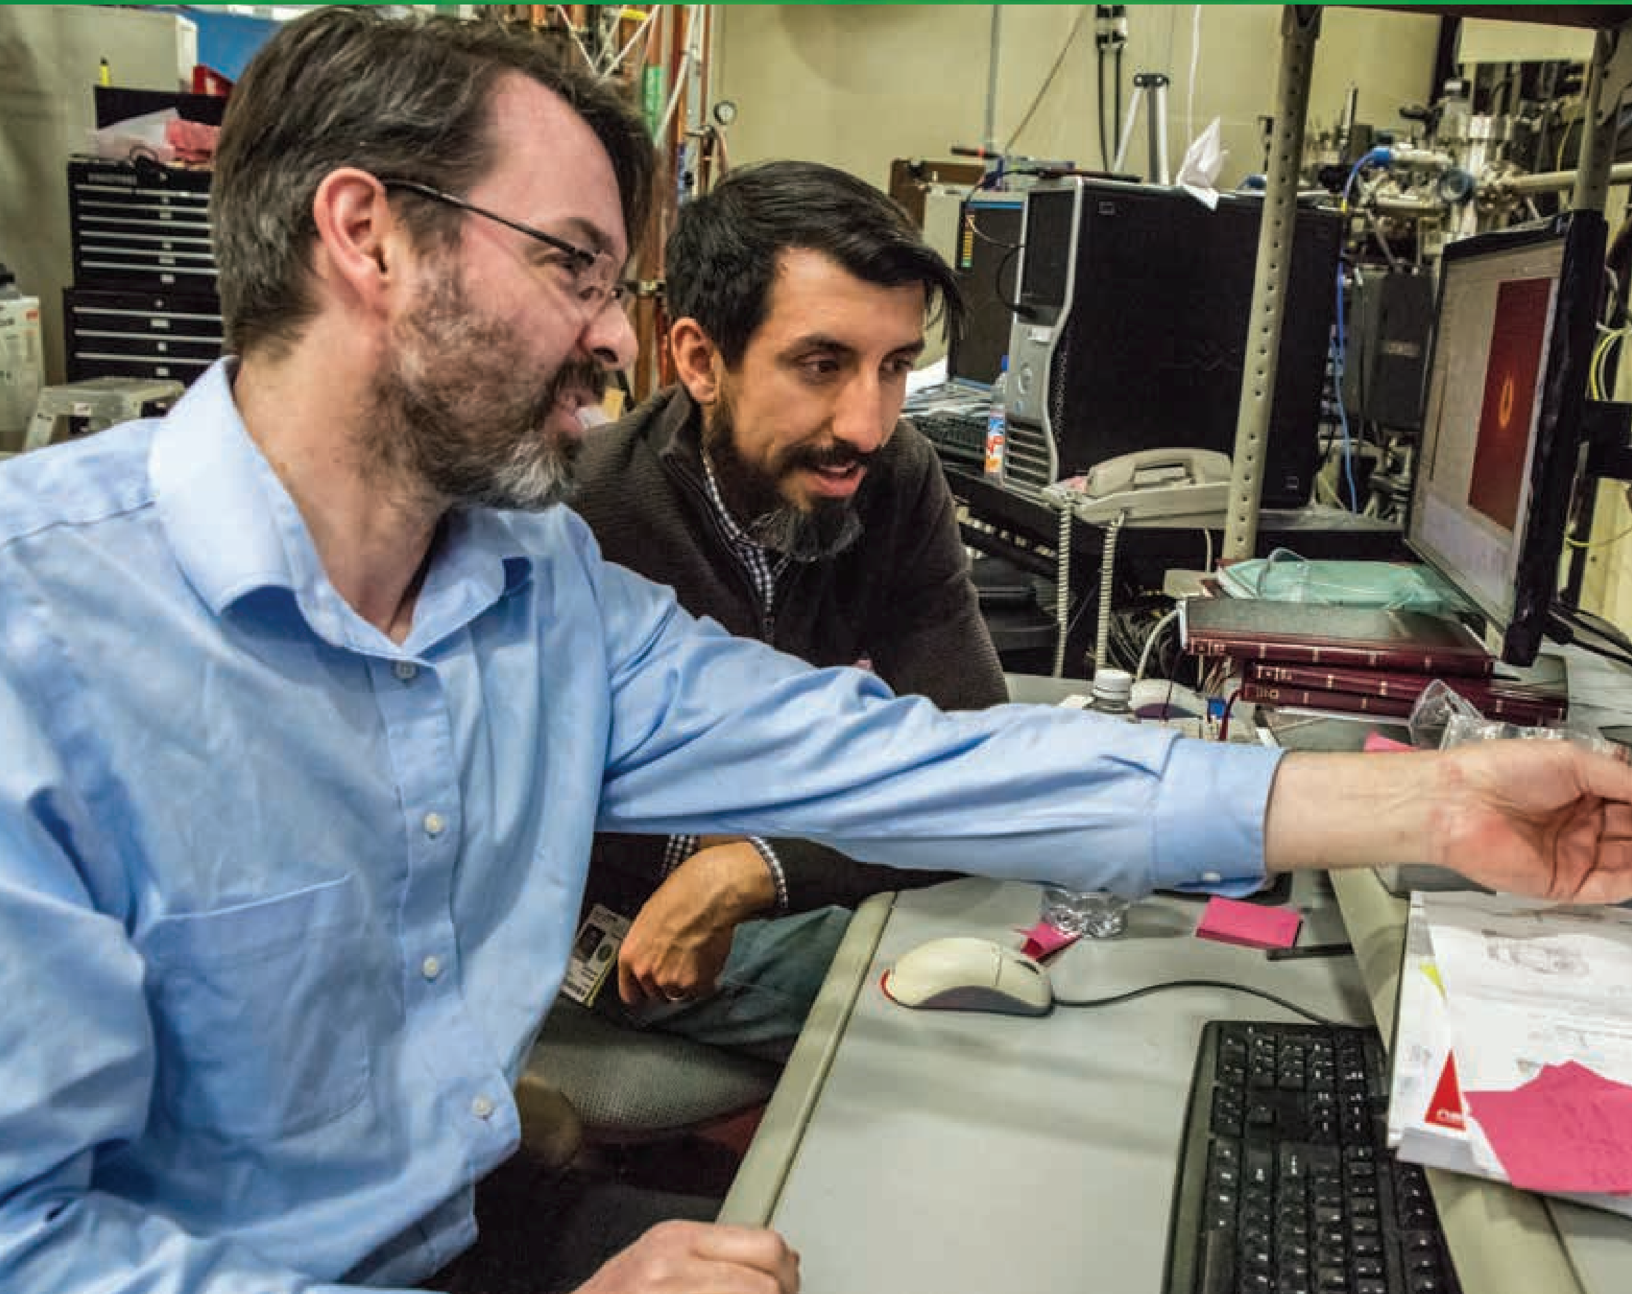

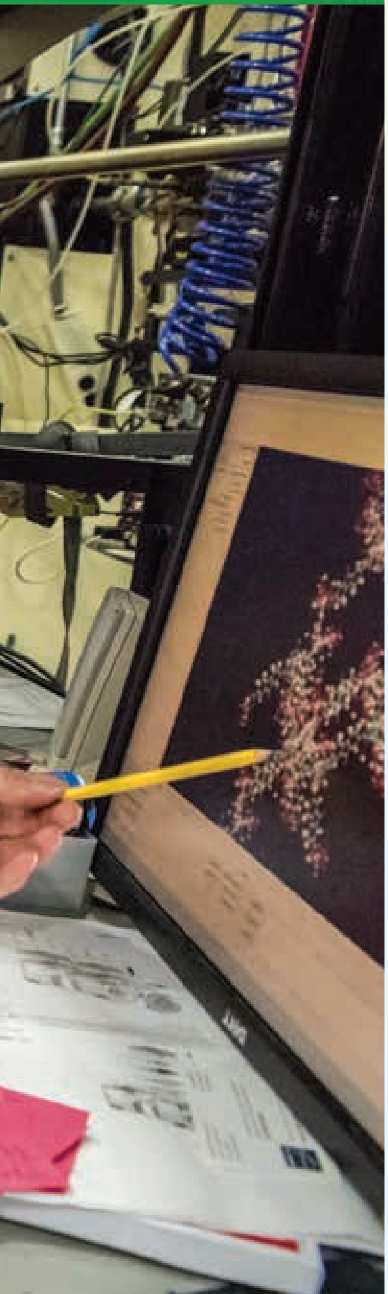

2016 NERSC Usage by DOE Program Office

(NERSC HOURS IN MILLIONS)

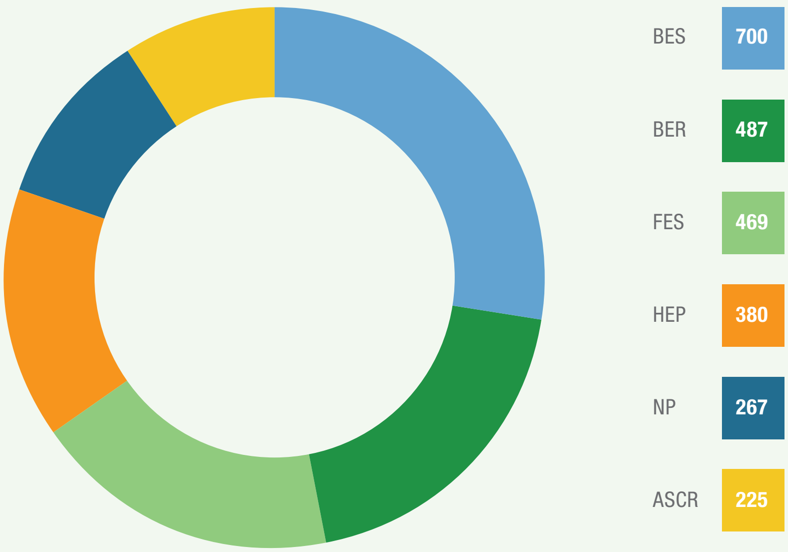

2016 NERSC Usage by Institution Type

(NERSC HOURS IN MILLIONS)

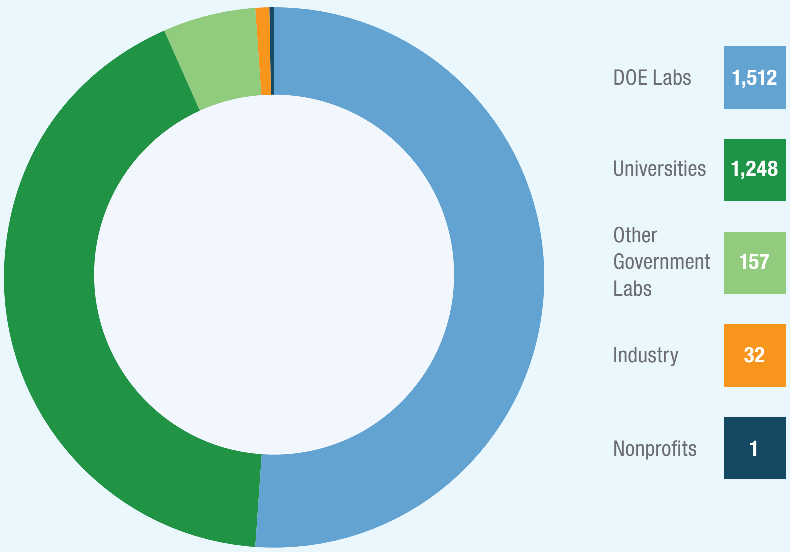

2016 DOE & Other Lab Usage at NERSC

(NERSC HOURS IN MILLIONS)

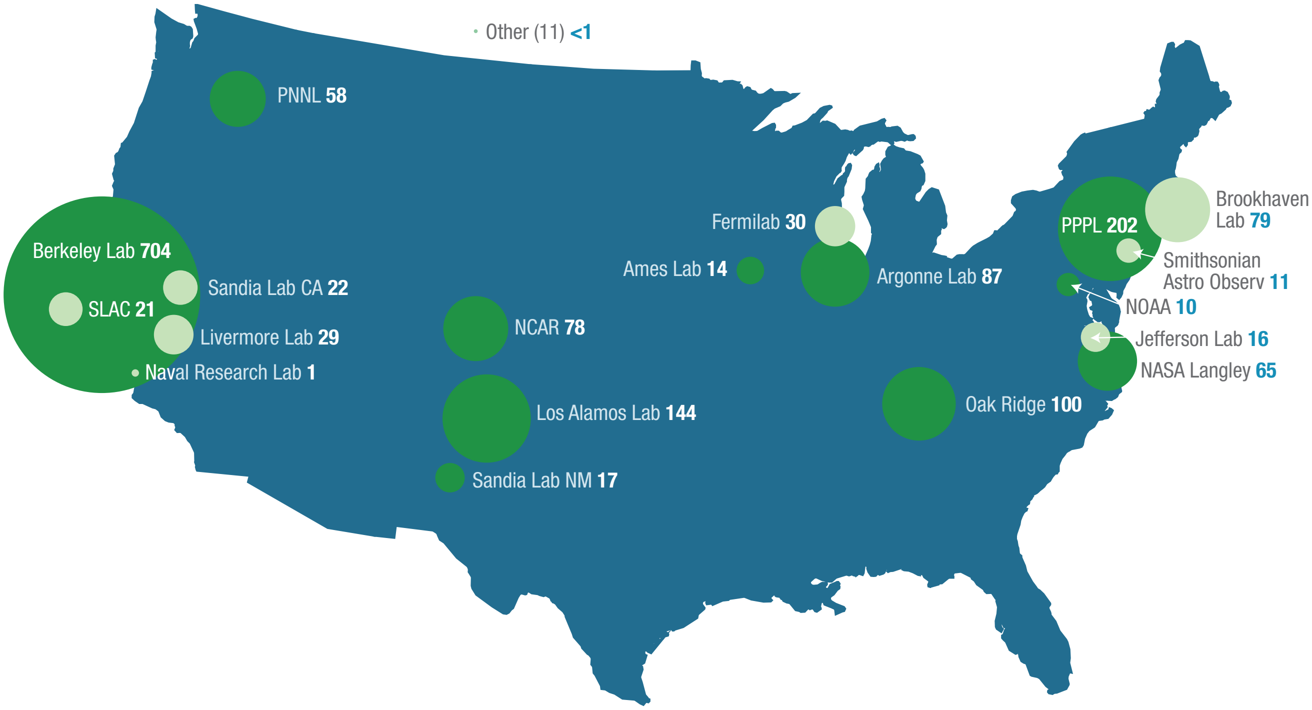

2016 Academic Usage at NERSC

(NERSC HOURS IN MILLIONS)

UC San Diego	141	U. Colorado Boulder	18	Tulane Univ	8
U. Arizona	121	U. Michigan	18	UC Santa Barbara	8
MIT	110	Johns Hopkins Univ	17	Iowa State	7
UC Berkeley	92	Temple Univ	17	U. Texas Austin	7
Princeton Univ	91	U. Rochester	17	Vanderbilt Univ	7
UCLA	60	Northwestern Univ	16	Columbia Univ	6
U. Kentucky	58	U. Maryland	15	Cornell Univ	6
UC Irvine	39	U. Oklahoma	14	Harvard Univ	6
George Wash Univ	35	Colorado School Mines	13	LA Tech Univ	6
U. Chicago	34	Florida State	11	Stanford Univ	6
U. Wisconsin Madison	32	U. Central Florida	10	U. Florida	6
U. Missouri KC	26	Dartmouth College	9	Clemson Univ	5
U. Illinois U-C	24	U. Washington	9	Oregon State	5
U. Penn	20	Auburn Univ	8	Penn State	5
Cal Tech	19	North Carolina State	8	U. South Carolina	5

2016 NERSC Users by State

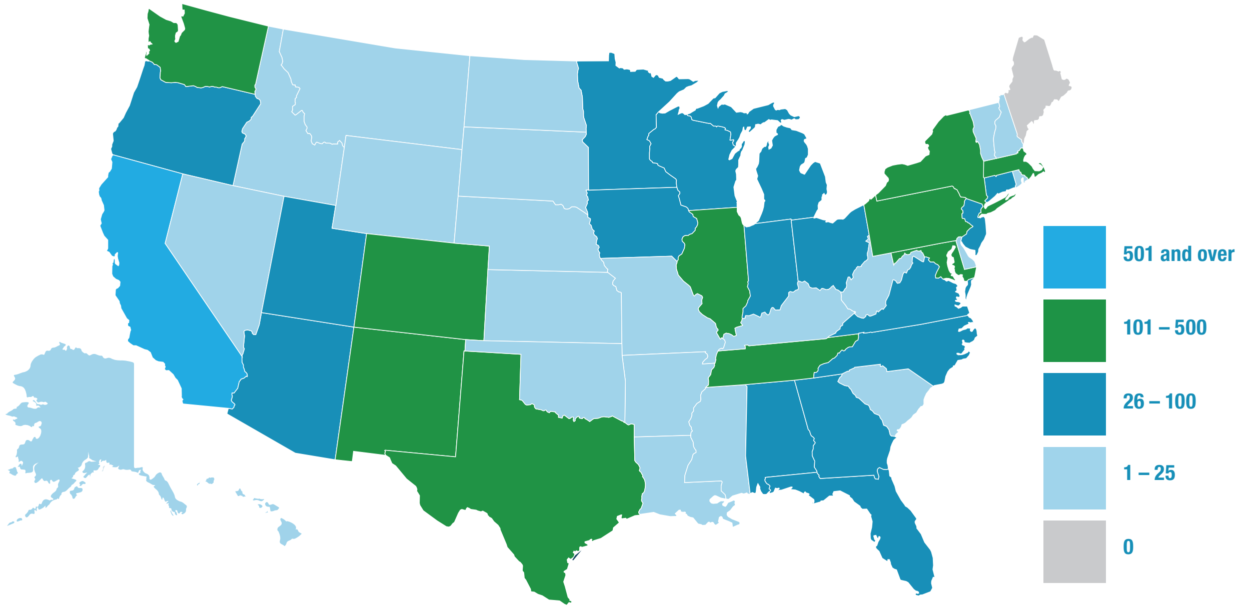

California	2,419	New Jersey	128	Georgia	58	Missouri	20	Montana	7
Illinois	397	North Carolina	104	Utah	42	South Dakota	18	Hawaii	6
New York	268	Michigan	98	Ohio	41	Delaware	17	Mississippi	5
Tennessee	267	Maryland	94	Arizona	38	Kentucky	16	West Virginia	5
Washington	226	Florida	84	Minnesota	37	Oklahoma	16	Wyoming	5
Massachusetts	191	Connecticut	76	Alabama	36	Kansas	15	Nebraska	4
Pennsylvania	182	Wisconsin	73	Oregon	33	North Dakota	12	Vermont	4
Texas	165	Indiana	67	District of Columbia	26	Arkansas	10	Nevada	3
Colorado	162	Virginia	65	Louisiana	24	New Hampshire	10	Alaska	2
New Mexico	134	Iowa	62	South Carolina	22	Rhode Island	9	Idaho	1

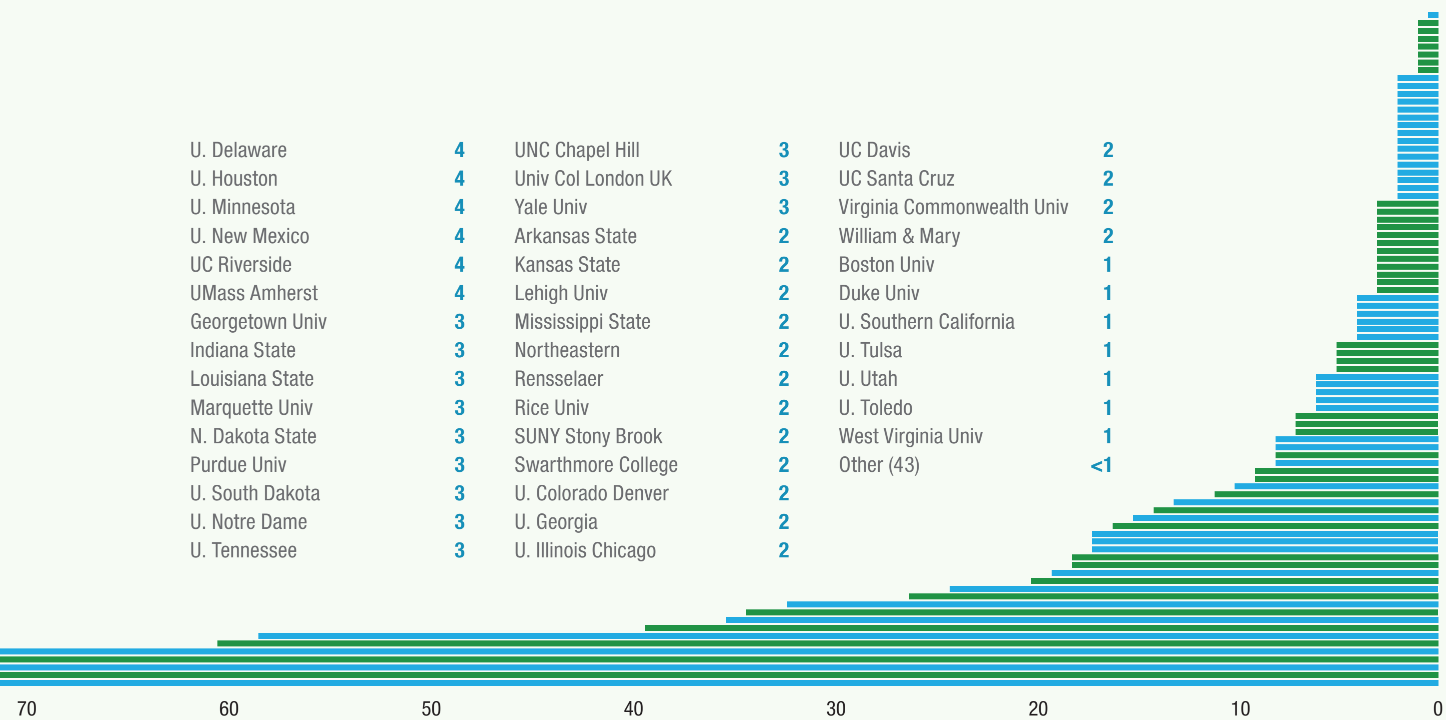

U. Delaware	4	UNC Chapel Hill	3	UC Davis	2
U. Houston	4	Univ Col London UK	3	UC Santa Cruz	2
U. Minnesota	4	Yale Univ	3	Virginia Commonwealth Univ	2
U. New Mexico	4	Arkansas State	2	William & Mary	2
UC Riverside	4	Kansas State	2	Boston Univ	1
UMass Amherst	4	Lehigh Univ	2	Duke Univ	1
Georgetown Univ	3	Mississippi State	2	U. Southern California	1
Indiana State	3	Northeastern	2	U. Tulsa	1
Louisiana State	3	Rensselaer	2	U. Utah	1
Marquette Univ	3	Rice Univ	2	U. Toledo	1
N. Dakota State	3	SUNY Stony Brook	2	West Virginia Univ	1
Purdue Univ	3	Swarthmore College	2	Other (43)	<1
U. South Dakota	3	U. Colorado Denver	2		
U. Notre Dame	3	U. Georgia	2		
U. Tennessee	3	U. Illinois Chicago	2		

2016 NERSC Users by Country

North America:	5,853
United States of America	5804
Canada	41
Puerto Rico	7
Mexico	1

South America:	30
Chile	17
Brazil	10
Argentina	2
Colombia	1

Africa:	7
South Africa	6
Nigeria	1

Middle East/ Asia Pacific:	335
People's Republic of China	165
Japan	38
Republic of Korea (South)	35
Australia	27
India	25
Taiwan, Province of China	21
Israel	17
Singapore	2
Bangladesh	1
Malaysia	1
New Zealand	1
Pakistan	1
Saudi Arabia	1

2016 NERSC Usage By Discipline

(NERSC HOURS IN MILLIONS, INCLUDES DOE PRODUCTION, ALCC AND DIRECTOR'S RESERVE PROJECTS)

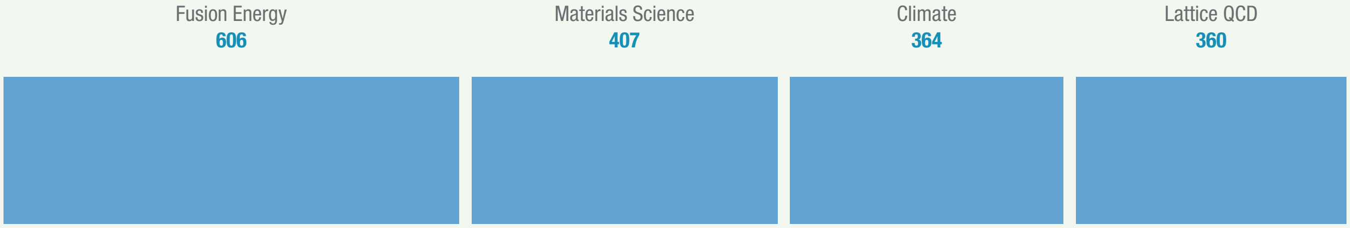

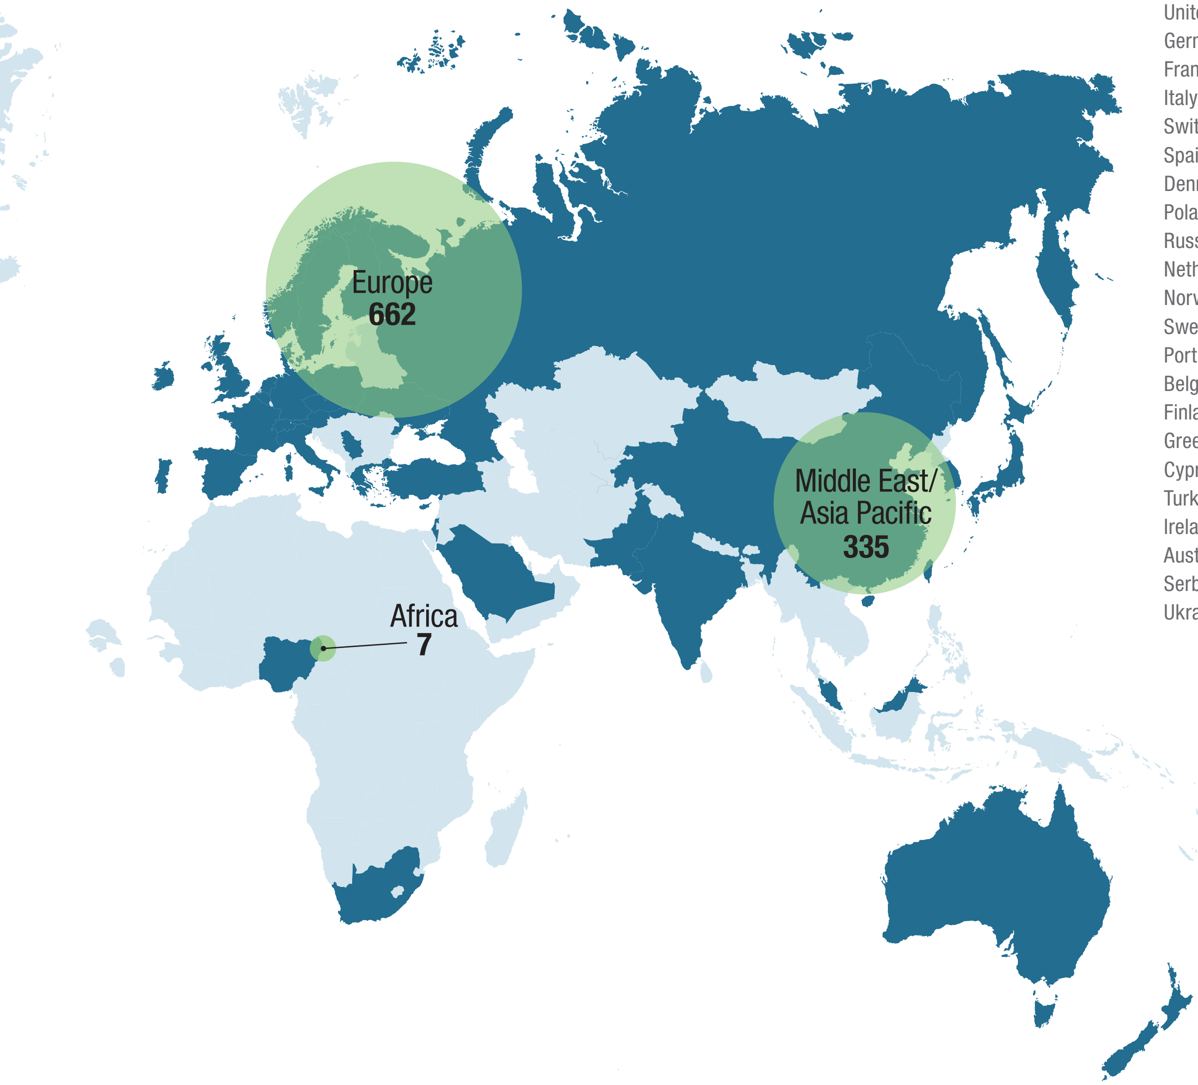

Europe:	662
United Kingdom	172
Germany	96
France	87
Italy	61
Switzerland	51
Spain	37
Denmark	32
Poland	32
Russian Federation	18
Netherlands	11
Norway	11
Sweden	11
Portugal	10
Belgium	8
Finland	6
Greece	5
Cyprus	4
Turkey	4
Ireland	3
Austria	1
Serbia	1
Ukraine	1

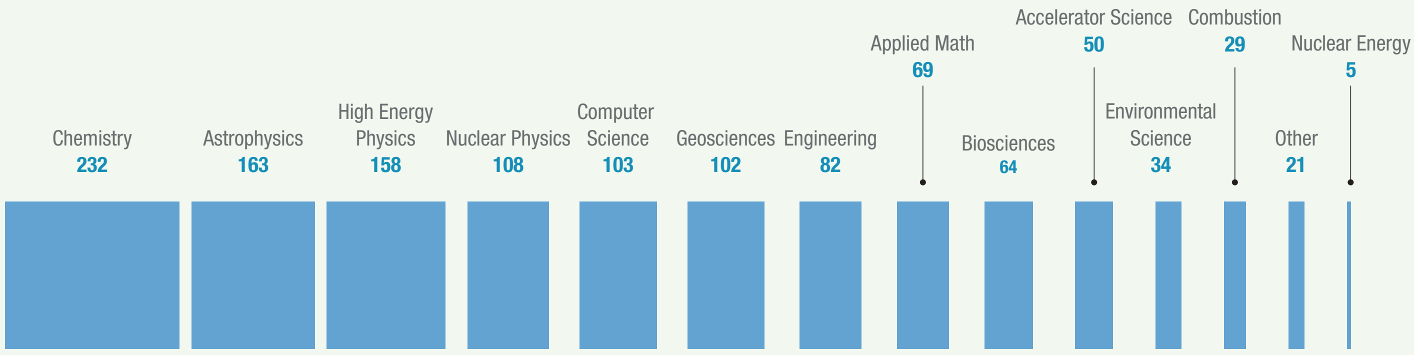

Innovations

For more than 40 years, NERSC has been providing cutting-edge HPC resources and expertise to the Department of Energy's Office of Science research community. The center specializes in reliable and advanced architectures capable of both world-class simulation and massive data analytics. In 2016, NERSC continued its tradition of developing innovative tools and processes designed to provide its nearly 7,000 users with additional services, higher performance and operational efficiency. These innovations align with our three strategic focus areas: data science and analytics, exascale and operational excellence.

Data Science and Analytics

Nearly 7,000 users employ NERSC's supercomputing platforms to tackle problems across the sciences, ranging in scale from astronomical to organismal, from molecular all the way down to subatomic physics. Typical data set sizes range from 100 gigabytes to petabytes. While NERSC has state-of-the-art computational and storage resources to handle the logistics, the real challenge is in determining scalable analytics methods and software frameworks.

For example, extreme weather events pose great potential risk to ecosystems, infrastructure and human health. Analyzing extreme weather in the observed record (satellite and weather station products) and characterizing changes in extremes in simulations of future climate regimes is an important task. Similarly, upcoming astronomical sky surveys will obtain measurements of tens of billions of galaxies, enabling precision measurements of the parameters that describe the nature of dark energy. Here again, developing the appropriate tools to help astronomers more efficiently sift through all of this data is critical.

Here are some of the data science and analytics innovations NERSC rolled out in 2016 to help users meet these needs.

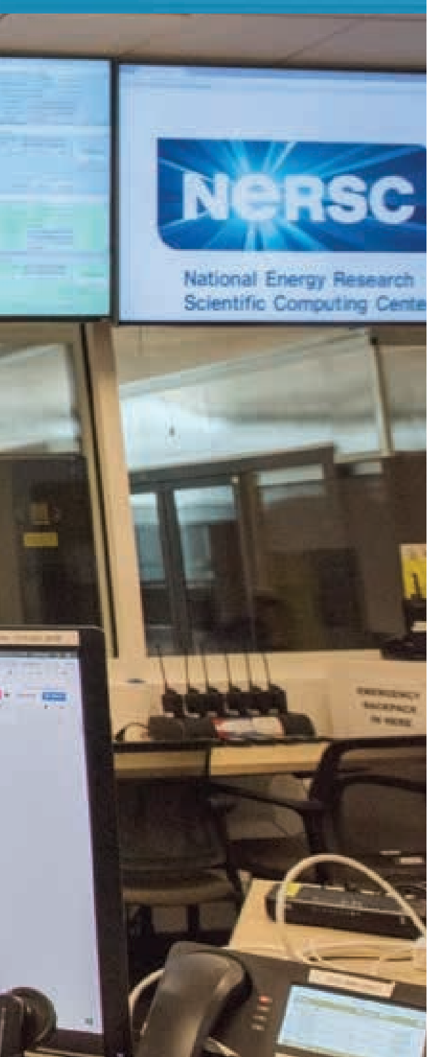

Accelerating Data-Intensive Workflows with SDN

Detectors and imaging facilities coming online in the next three to five years are expected to produce data in excess of 1 terabyte per second. These data-intensive workloads require systems that have the ability to ingest and process data from scientific instruments and sensor networks. However, a major challenge is enabling HPC systems to effectively ingest data from these instruments.

This dilemma prompted NERSC, in partnership with Cray, to explore new ways to more efficiently move data in and out of NERSC's Cori supercomputer, a Cray XC40. One promising solution is to integrate software defined networking (SDN) capabilities to allow scientists at experimental facilities to co-schedule networking bandwidth, compute resources and Burst Buffer bandwidth.

Software defined networking encompasses several kinds of network technologies to make the network as agile and flexible as the virtualized server and storage infrastructure found in the modern data center. Incorporating SDN into the NERSC workflow requires being able to provide dynamic scheduling and provisioning of networking end points that map to compute resources in Cori.

NERSC engineers initially focused on basic capabilities to enable high-bandwidth outbound connectivity to external systems, which is critical for many common workflows in high energy physics space and real-time data analysis. Toward this end, NERSC

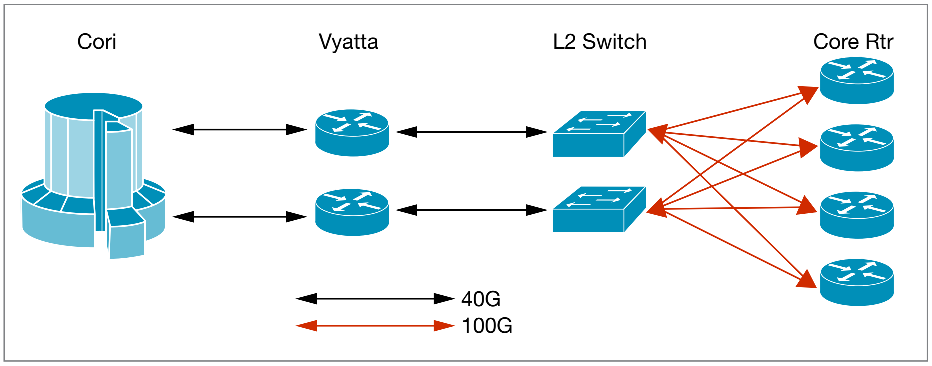

High-level network architecture diagram for the software-defined network connected to Cori.

deployed software-based routers running Brocade's vRouter coupled with special "bridge" nodes connected to Cori's high-speed Aries network, enabling traffic to efficiently route from Cori compute nodes to external networks. These routers are also associated with an SDN controller based on the OpenDaylight Controller. Eventually this controller will be integrated with SLURM (Simple Linux Utility for Resource Management)—a powerful open source, fault-tolerant and highly scalable resource manager and job scheduling system developed by SchedMD—to enable jobs to specify their network requirements and have the network respond accordingly.

NERSC is partnering with ESnet to explore how this capability can be integrated and extended across the wide area network (WAN). The ultimate goal is to enable end users from experimental facilities to provision end-to-end connectivity and bandwidth that extend into systems like Cori to enable real-time analysis.

Optimizing Cori Queues to Increase User Efficiency

When the new Cori Haswell partition came online in 2016, it was immensely popular with users, resulting in a queue backlog that sometimes extended for weeks. The intermixing of data-intensive and traditional HPC workloads on the system created a job mix that was challenging for SLURM to schedule in a single iteration.

To address this problem and help Cori serve these workloads more efficiently, NERSC worked to analyze job patterns and user needs. This analysis determined that a larger debug queue and increased capability for high throughput jobs would help the data community, many of whom were trying out a system like Cori for the first time. So NERSC engineers worked to optimize the scheduling algorithm. By changing the way job priority was handled, the scheduling algorithm was able to scan the entire backlog of jobs in one pass. This meant a larger pool of jobs could be considered for backfill, and the system could be used more effectively.

NERSC made two major adjustments to the priority and scheduling algorithms in SLURM to allow the HPC and data workloads to happily coexist. The typical prioritization scheme on most data-intensive and high-throughput computing systems uses fair-share or related algorithms to prioritize jobs. This scheme is appropriate on systems where there is little variation in job size, but on the Cori Haswell nodes there is frequently a 200-fold variation in resource requests between jobs, and on Edison and Cori KNL there can be a 2,000-fold variation. This variation in job size is not amenable to fair-share based algorithms, since fair-share and related algorithms reorder jobs constantly in order to provide the desired mix of user jobs.

The first major adjustment to SLURM, in partnership with SchedMD, was to devise a new backfill scheduling algorithm that can more effectively exploit opportunities for backfill. This

NERSC made two major adjustments to the priority and scheduling algorithms in SLURM to allow the HPC and data workloads on Cori to happily coexist.

is done by ensuring that all jobs above a particular priority level (called a "bright line") are able to reserve resources. We use a priority scheme where all jobs are submitted at or below that resource reservation bright line. We then use job aging to progressively move jobs from below the bright line to above it. Since all jobs age at the same rate, and all jobs are submitted at or below the priority bright line, this ensures that jobs being considered for resource reservations never reorder. This limits the computational time spent on resource reservations, which is expensive, to only those jobs that require resource reservations to start. It also allows the quantity of jobs considered for reservations to grow and shrink based on queue conditions. Jobs below the bright line are now only considered for backfill. These kinds of backfill calculations are very quick and efficient, evaluating thousands of jobs per scheduling cycle.

The second major change to our SLURM was a modification we made internally and plan to submit to SchedMD for general inclusion in SLURM. With this modification, SLURM only allows a configurable number of jobs (or resources requested by jobs) per user to age into the high priority segment. This has the effect of allowing jobs from more users to start over a period of time.

Overall, these changes to the SLURM scheduler reduced the average wait time for short debug jobs on Cori from hours to less than 10 minutes and greatly improved the efficiency of scientific computing at NERSC.

slurm
workload manager

Data Transfer Nodes Get an Upgrade

Data transfer nodes (DTNs) are specially configured systems that provide the best performance for transferring data into and out of the center. NERSC has been collaborating with ESnet and other facilities on the Petascale DTN project to achieve transfer rates of one petabyte per week between major facilities—initially NERSC, Argonne, Oak Ridge and NCSA. This is translated to regularly achievable rates over 15 Gbs on real example science datasets. Because transfers over the WAN have high latency between sender and receiver, moving data over the WAN at fast rates requires careful configuration of all the devices along the data path. To meet this goal, the NERSC storage and networking teams, along with ESnet, worked to implement several major enhancements.

On the DTN servers, the storage team upgraded the operating system to take advantage of new features in the TCP stack. This included a feature that provides traffic pacing. Pacing smoothed out transfer speeds, especially when moving data in parallel streams, resulting in less variability and higher rates overall. The storage team also mounted the Cori scratch filesystem directly on the DTNs, so NERSC users are now able to write directly into the scratch filesystem from outside of NERSC. For some workflows, this saved copying data between filesystems.

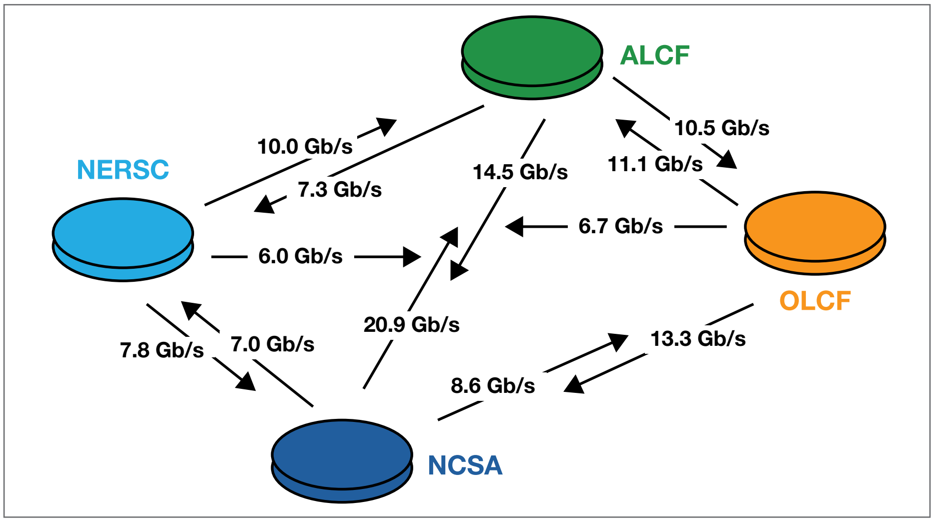

The Petascale DTN project is designed to achieve and maintain the ability to move data at the rate of roughly one petabyte per week between major ASCR and HEP facilities.

Working with ESnet researchers, NERSC also deployed a passive network monitoring tool called tstat that allows analysis of all network flows on the busy production systems. Historically, this has been difficult to achieve because of the volume of data that needs to be captured. The tstat tool uses targeted sampling of network flows, which allowed us to identify flows that were underperforming. The NERSC networking team used the output to fine-tune network devices, increasing the bandwidth between NERSC internal and border routers to 2 x 100 Gpbs and configuring switches to better handle WAN traffic, further improving transfer performance.

Further enhancements are planned for the DTN cluster in 2017. These include increasing the size of the DTN pool; targeting DTN nodes for specific purposes such as GridFTP transfers, interactive work or transfers to HPSS; and increasing the network speed the systems are capable of to 100 Gpbs.

The Path to Exascale

Years of scientific domain knowledge and code development have been invested in scientific applications that are now considerably more sophisticated than they were when the research community made the transition from vector to parallel supercomputing two decades ago. In terms of sheer size, some of the community codes are an order of magnitude larger than they were in the early 1990s. Also, many of the significant architectural changes expected in the exascale era will be manifest in the next generations of computers, such as NERSC's Cori system, and in future systems at other DOE national labs.

To help move the DOE research community along the path to exascale computing, NERSC has spearheaded or participated in several projects designed to help our users make the transition to next-generation systems and architectures. For example, in 2016 NERSC—in partnership with the ALCF and OLCF—completed the remaining exascale requirements reviews workshops with FES, BER, NP and ASCR and began working on a requirements review cross-cut meeting and report. NERSC is also participating in DOE's FastForward and

DesignForward programs, which are focused on the development of hardware and software for exascale systems. NERSC staff have been involved in the writing, release and evaluation of the follow-on Pathforward program, which is part of the Exascale Computing Project.

Here are some detailed examples of exascale-related innovations implemented at NERSC in 2016.

Roofline Model Enhances Manycore Code Optimization Efforts

The Roofline Model, a software toolkit developed at Berkeley Lab over the past decade to better understand supercomputer performance, is now being used to boost application performance for researchers running codes on manycore architectures at NERSC and other supercomputing facilities.

This work, spearheaded by computer scientists in Berkeley Lab's Computational Research Division, is helping to address the growing gap between processor performance and memory performance—the "memory wall"—that can limit how well many scientific applications perform. Even with the fastest processors, if data gets bogged down moving in and out of memory, the processors end up waiting for data to process, which slows overall performance.

With the popularity of manycore processors—such as the KNL processors in Cori—and the growing diversity of computing architectures, supercomputers are increasingly complex, making it difficult for users to achieve sustained application performance across different architectures. But understanding memory bandwidth limits can help explain the gap between theoretical and observed performance. Thus performance models and tools that aid users' understanding of bandwidth limits for a particular application are crucial.

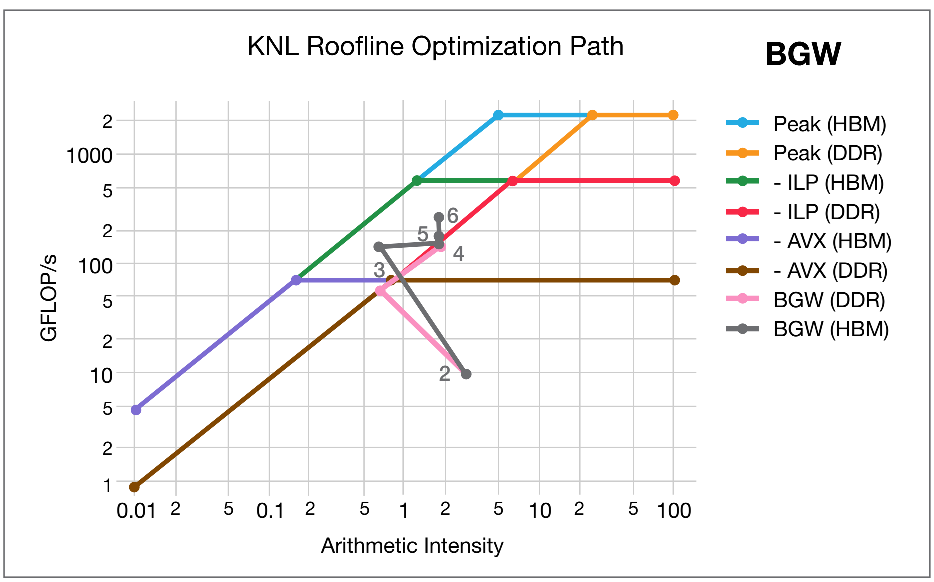

The Roofline Model is designed to enable users to assess the quality of attained performance of an application by combining data locality (how much data needs to move in and out of the processor), bandwidth (the speed at which that data can be moved) and different parallelization paradigms into a single performance figure, yielding a two-dimensional graph that clearly plots the bandwidth bottlenecks that must be eliminated to speed up the application.

The Roofline Model has been used for a number of years to characterize supercomputing systems and architectures, and it is now helping users who write scientific applications for manycore systems see how changing parameters of their code can improve performance. In 2016, it was expanded to both visualize and guide application optimization, and new tools have been developed to support this.

For example, through a collaboration with Intel, the Roofline team—which includes NERSC staff—has been working to incorporate Roofline into Intel's Vector Advisor Tool, a vectorization optimization and shared memory threading assistance tool. The resulting Intel Advisor Roofline Tool allows users to collect Roofline performance data, including accurate FLOPs counts that correctly consider different vector registers and masks using Intel's PIN tool, as well as accurate data traffic from various levels of the cache memory hierarchy. NERSC use cases have driven feature development in the Intel tool, including Intel's new DRAM Roofline and hierarchical Roofline features.

NERSC Launches NESAP for Data

Following in the footsteps of the popular NERSC Exascale Science Applications Program (NESAP) launched in 2014, in 2016 NERSC issued a call for proposals for NESAP for Data. Through NESAP, NERSC partners with code teams and library and tool developers to prepare and optimize their codes for the Cori manycore architecture. Like NESAP, the NESAP for Data program joins application teams with resources at NERSC, Cray, and Intel. It is jointly managed by NERSC's Data Analytics and Services, Data Science Engagement and Application Performance groups.

While the initial NESAP projects involved mostly simulation codes, NESAP for Data is targeting data-intensive science applications that rely on processing and analysis of massive datasets acquired from experimental and observational sources such as telescopes, microscopes, genome sequencers and light sources. The first round of selected NESAP for Data projects, announced in January 2017, are:

- Dark Energy Spectroscopic Instrument Codes; Stephen Bailey, Berkeley Lab (HEP)
- Union of Intersections Framework; Kris Bouchard, Berkeley Lab (BER)
- Cosmic Microwave Background Codes (TOAST); Julian Borrill, Berkeley Lab (HEP)
- ATLAS Simulation/Analysis Code; Steve Farrell, Berkeley Lab (HEP)
- Tomographic Reconstruction; Doga Gursoy, Argonne (BES)
- CMS Offline Reconstruction Code; Dirk Hufnagel, Fermilab (HEP)

NERSC, Argonne Address Emerging Storage Needs

As the storage hierarchy deepens to accommodate the bandwidth and capacity requirements of extreme-scale systems, the number of individual storage devices and servers that contribute to application performance is increasing. To understand I/O performance in these increasingly complex architectures, NERSC has been collaborating with Argonne National Laboratory to develop tools and techniques that combine performance data from multiple sources through the I/O stack.

One result of this effort has been the creation of a Lustre-specific module for the Darshan I/O profiling library developed at Argonne. This module, created by in NERSC's Advanced Technologies Group, allows Darshan to track the specific Lustre object storage targets (OSTs) accessed by an application's I/O and identify specific devices or servers that are degrading overall performance.

In a paper presented at the 5th Workshop on Extreme-Scale Programming Tools, held in conjunction with SC16, this module was used to attribute poor checkpointing performance in an extreme-scale cosmology application to the temporary failover of two Lustre object storage servers. The Darshan Lustre module has been integrated into Darshan and was released publicly in Darshan version 3.1 in September 2016.

Operational Excellence

OTG Implements New Facility Monitoring Tools

NERSC's Operations Technology Group (OTG) continues to improve its facility monitoring tools. In 2015 OTG implemented an aggregation interface that provided an easier way to visualize, diagnose and resolve alerts, and in 2016 OTG scaled down the process overhead on the systems themselves. NERSC uses Nagios—which provides enterprise-class open source IT monitoring, network monitoring and server and applications monitoring—as its main monitoring tool and has enhanced it with innovations in generating the alerts. Instead of a compute node running the Nagios client and signaling an alert, an external function checks each node for the service to determine if it is still running.

In 2016 this monitoring process was implemented in three new areas: the Datawarp and Burst Buffer on Cori, the SLURM scheduler and KNL node status. These innovations in monitoring offer a significant advantage: they are single-path communication, SMW to Nagios server. This means Nagios does not initiate the plugin, so compute nodes are not interrupted; all the data is gathered from an external source, and the only additional code that has to reside on the system is a small BASH script on the SMW. As a result, this method minimizes security considerations.

In addition, the process is simple to manage administratively, with little overhead on the system. It is also scalable; passive check results are lighter on the Nagios process because only changes get reported, and no plugin processing is performed on the system. A Nagios freshness check mechanism is used to ensure that the plugin is functioning; if not, it is automatically restarted.

This method of collecting facilities data provides robust searching capabilities. Staff can quickly isolate problems by searching against "SLURM downtime + reason strings." Nodes in use by a particular job or user can also be easily identified, allowing staff to correlate multiple separate node events or identify nodes that may be affected later. With 11,000+ nodes total in NERSC systems and multiple concurrent events that can affect hundreds of nodes each, this searching/filtering ability is indispensable for enabling staff to make sense of the alarms, diagnose them and eventually start a process of restoring systems back to their functional state.

Centralized Data Collection Provides New Analysis, Correlation

Working with data collected throughout 2016, NERSC OTG implemented three major interfaces for staff to use to analyze the data that OTG collects: Kibana, Grafana and a RESTful API. Kibana comes with Elasticsearch—a distributed, RESTful search and analytics engine—to help explore the raw data, providing an almost interactive way to view data. Grafana provides a way to view multiple data sources such as time series data, making it easier to build graphs for visualization. The RESTful API allows users to write their own programs to access data and create their own outputs or sample data. These interfaces empower staff to create queries, analyze data and look for patterns that allow them to solve issues system administrators had not considered.

This approach to data collection for Cori has led to multiple innovations:

- **Cray community availability of data:** In collaboration with NERSC, Cray created a plugin architecture called xtpmd that allows sites to independently write their own plugins to gather the system environmental data from their systems. Not only is this environment available to the community, it helps NERSC specifically collect data about power, temperature, fan speeds, water temperature and jobs running for all the nodes, slots and cabinets. Prior to this innovation, collecting this data was error prone because the data needed to be parsed from text files. The data can now be streamed real-time into our collection infrastructure for analysis and visualization.

- **Burst Buffer data validation:** Consolidating Cori system data helped support our understanding of the Burst Buffer, whose storage devices are solid-state drives. Collecting read and write statistics and visualizing them allows us to see user access patterns for this dynamic storage method. We are able to verify and correlate that, as expected, users are able to use Burst Buffer for faster storage access and are able to stage their data to this device to obtain better performance for their jobs.

- **KNL processor environmental anomaly and patterns detection:** The data collection infrastructure now has several months of data and we are able to analyze environmental data and events. In one example, the data collected allowed NERSC staff to run an analysis that enabled Cray to adjust a thermal alarm setting for the system.

NERSC plans to migrate these data collection methods to Edison.

Storage Library Automation Goes Beyond the GUI

Maintaining stability within the storage libraries can be achieved through preventive and proactive actions. In the past, these efforts would take hours and involved a series of manual processes as staff worked with the graphical user interface (GUI) on the console. A challenge has been to find a way to use more efficient ways of interacting with the system. Toward this end, in 2016 OTG staff implemented automation processes that have achieved some level of efficiency in the following areas:

- **Oracle tape libraries:** Staff investigated a tool called SikuliX, an open-source research project at the User Interface Design Group at MIT. The tool uses image recognition to identify and control GUI components. This has great potential to ease management of the Oracle robotic tape libraries, which use a GUI for managing many aspects of the system. Early successes include the ability to programmatically collect and submit logs to Oracle Support, monitor the operational state of the libraries and detect impending failures before they occur.

- **Automating the tape reclaim process:** Tapes in NERSC's High Performance Storage System (HPSS) use both a digital and physical barcode. As users delete files, tapes become sparsely used and an administrative process is needed to periodically survey the tape population and copy data from sparse tapes to new volumes. The process frees up the sparse tapes for reuse; however, there is an additional step of reallocating the now empty tape to HPSS. This reallocation activity, called "reclaim," involves multiple steps and is a time-consuming process. OTG needed an efficient way to complete the work without extensive staff intervention. To address this need, staff implemented a series of scripts to automate this multi-step, detail-oriented process for the IBM tape library. The process included safety protocols to prevent harming tape cartridges that contain useful data. The scripts identify tapes for reclaim, interact with the HPSS software to remove the cartridges from the system and then eject the volume to be physically relabeled by OTG staff.

Leveraging Network Infrastructure Consolidation

As part of the transition to Wang Hall, the NERSC networking group designed a new "mesh" topology for the switches infrastructure in the new facility. The mesh topology encapsulates the full configuration of every switch in the fabric into a single configuration, and then places that configuration on every node. The end result is that any switch in the mesh can quickly become the "master" node, or fabric orchestrator, and control the entire mesh in case of a failure.

In addition, the switches themselves are now configured using "bonded" connections into the network. This configuration and the mesh design cause minimal disruption to the network, even when a single link fails or when software needs to be upgraded. Through the bonded connections, one switch can be upgraded while the second switch maintains the connectivity for all the hosts connected to that switch in bonded mode. This innovative design removes any single point of failure and provides redundancy, fault tolerance and configuration efficiency.

Science Highlights

This section presents a selection of research highlights from 2016 based on computations and simulations run at NERSC, illustrating the breadth of science supported by the center for all six DOE program offices.

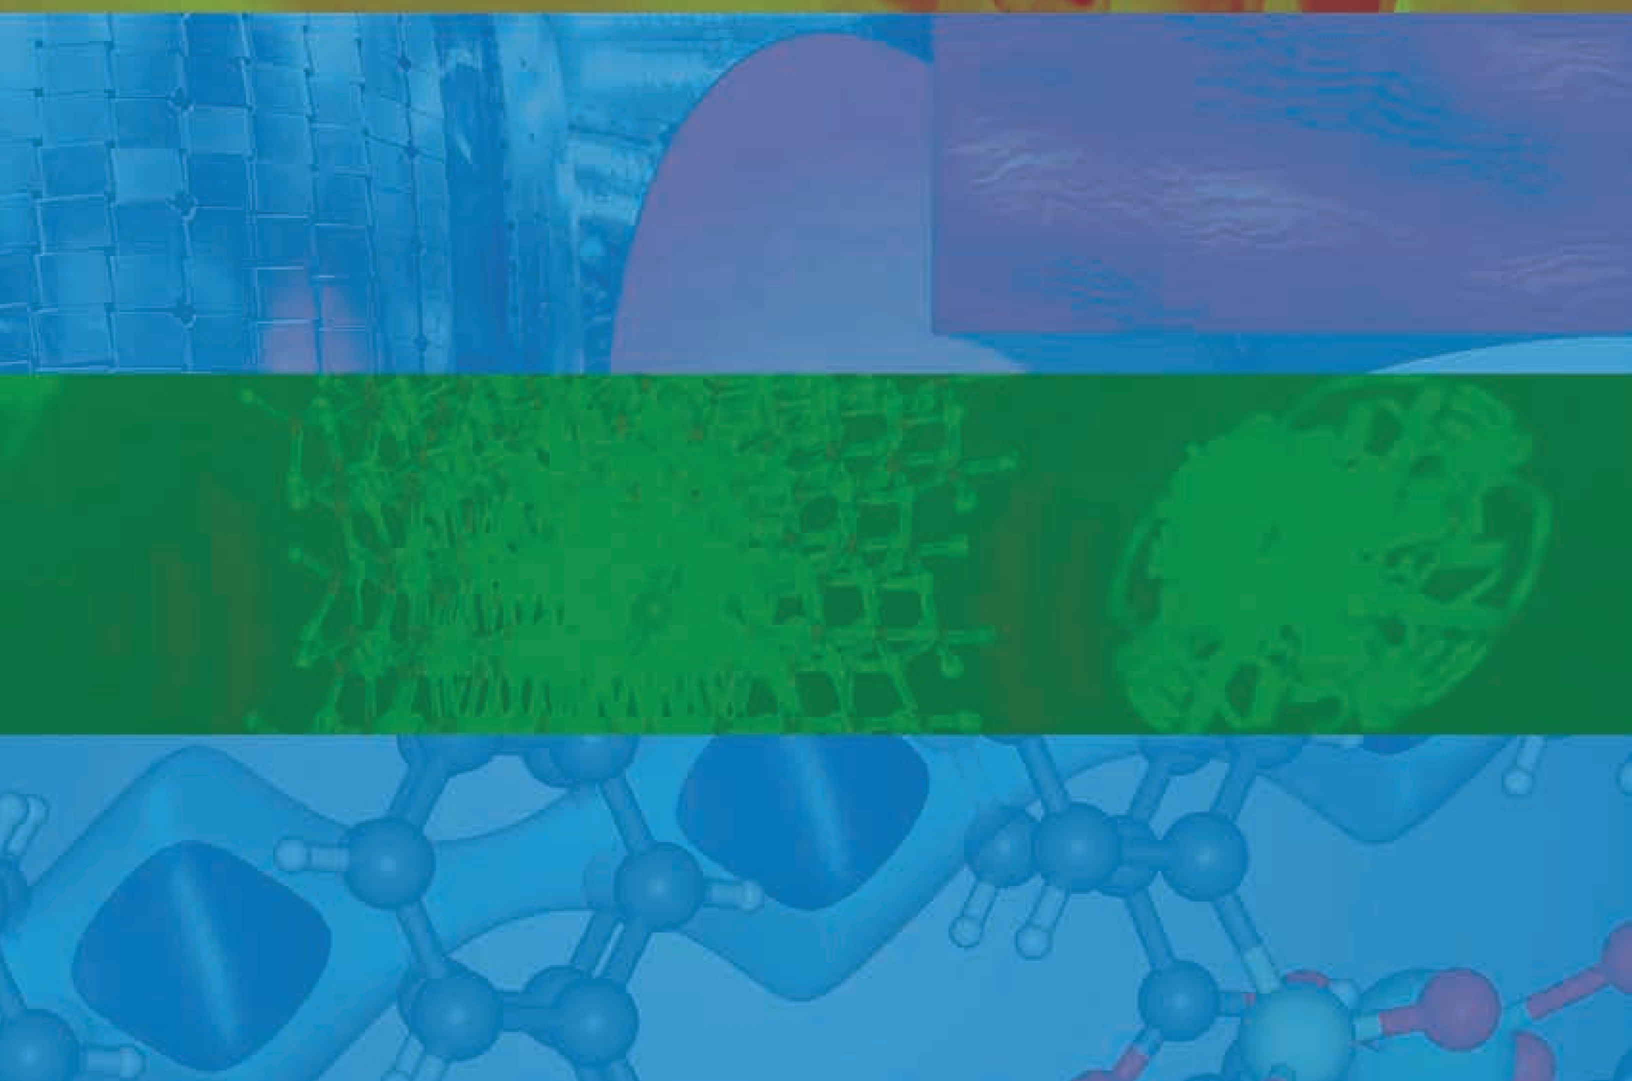

HPC4Mfg: Reducing the Cost of Paper Manufacturing

Advanced Scientific Computing Research

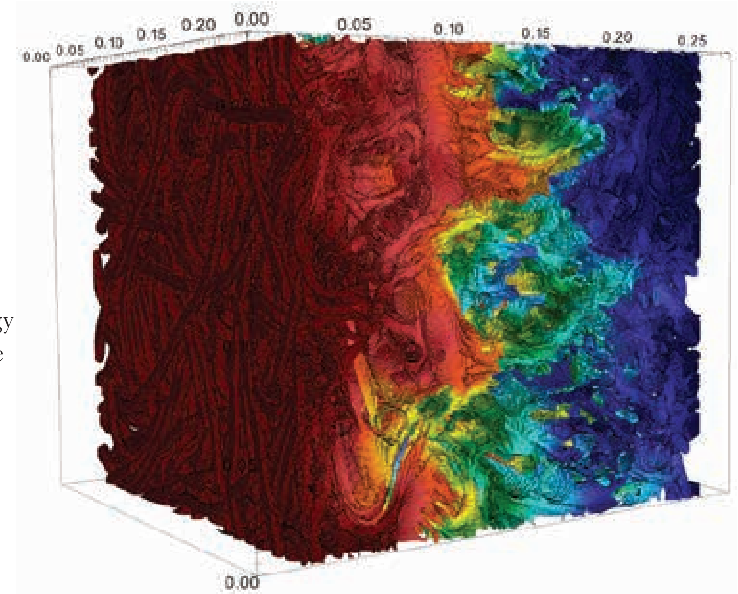

Scientific Achievement

Simulations run at NERSC through a unique collaboration between Berkeley Lab, Lawrence Livermore National Laboratory (LLNL) and an industry consortium (The Agenda 2020 Technology Alliance) could help U.S. paper manufacturers significantly reduce production costs and increase energy efficiencies. The project is part of the DOE's HPC for Manufacturing (HPC4Mfg) initiative, a multi-lab effort to use high performance computing to address complex challenges in U.S. manufacturing.

SIGNIFICANCE AND IMPACT

The papermaking industry ranks third among the country's largest energy users, behind only petroleum refining and chemical production, according to the U.S. Energy Information Administration. To address this issue, the researchers are using advanced supercomputer modeling techniques to identify ways that paper manufacturers could reduce energy and water consumption during the papermaking process. The first phase of the project targeted "wet pressing"—a process in which water is removed by mechanical pressure from the wood pulp into press felts that help absorb water from the system before it is sent through a drying process. If manufacturers could increase the paper's dryness by 10-15 percent during the wet pressing, it would save paper manufacturers up to 20 percent of the energy used in the drying stage and as much as $400 million for the industry annually.

As part of one of the first HPC4Mfg projects, a full-scale microscale flow model developed at Berkeley Lab was used to model the complex pore structures in press felts used in paper manufacturing.
Image: David Trebotich, Lawrence Berkeley National Laboratory

RESEARCH DETAILS

The researchers used a computer simulation framework, developed at LLNL, that integrates mechanical deformation and two-phase flow models, and a full-scale microscale flow model, developed at Berkeley Lab, to model the complex pore structures in the press felts. The model's flow and transport solver in complex geometries, developed by Berkeley Lab's David Trebotich, is based on the Chombo software libraries developed at the lab and is the basis for other application codes, including Chombo-Crunch, a subsurface flow and reactive transport code that has been used to study reactive transport processes associated with carbon sequestration and fracture evolution. Trebotich ran a series of production runs on NERSC's Edison system and was successful in providing his LLNL colleagues with numbers from these microscale simulations at compressed and uncompressed pressures, which improved their model.

Principal Investigator:
David Trebotich, Lawrence Berkeley National Laboratory

Full Story:
http://bit.ly/NERSCHPC4Mfg

The Invisible Chaos of Superluminous Supernovae

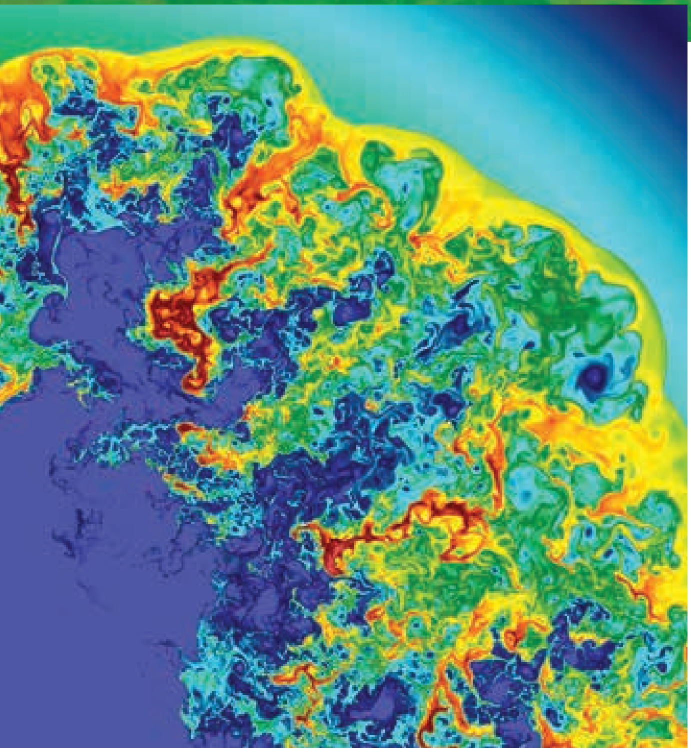

A 2D superluminous supernova simulation run on NERSC's Edison system.
Image: Ken Chen, National Astronomical Observatory of Japan

High Energy Physics

Scientific Achievement

To better understand the physical conditions that create superluminious supernovae, astrophysicists ran 2D simulations of these events using supercomputers at NERSC and a compressible astrophysics code, CASTRO, developed in Berkeley Lab's Computational Research Division.

SIGNIFICANCE AND IMPACT

Sightings of a rare breed of superluminous supernovae—stellar explosions that shine 10 to 100 times brighter than normal—have been perplexing astronomers for years. In particular, they are confounded by the extraordinary brightness of these events and their explosion mechanisms. But the first 2D simulations of these supernovae, run on NERSC's Edison system, are giving researchers new insights into their properties.

RESEARCH DETAILS

This is the first time anyone has simulated superluminous supernovae in 2D; previous studies have only modeled these events in 1D. By modeling the star in 2D the researchers captured detailed information about fluid instability and mixing that isn't possible to obtain from 1D simulations. These details are important to accurately depict the mechanisms that cause the event to be superluminous and explain their corresponding observational signatures such as light curves and spectra.

In addition to NERSC supercomputing resources, the research team used the CASTRO code to create multi-dimensional simulations of the superluminous supernovae. CASTRO is an adaptive mesh refinement hydrodynamics code developed at Berkeley Lab by scientists Ann Almgren and John Bell that is commonly used in astrophysics research.

Principal Investigator:
Stan Woosley, University of California, Santa Cruz

Journal Citation:
K. Chen, S. Woosley, T. Sukhbold, "Magetar-powered Supernovae in Two Dimensions. I. Superluminous Supernovae," *The Astrophysical Journal*, Vol. 832, No. 1, December 2016

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

The Incredible Shrinking Particle Accelerator

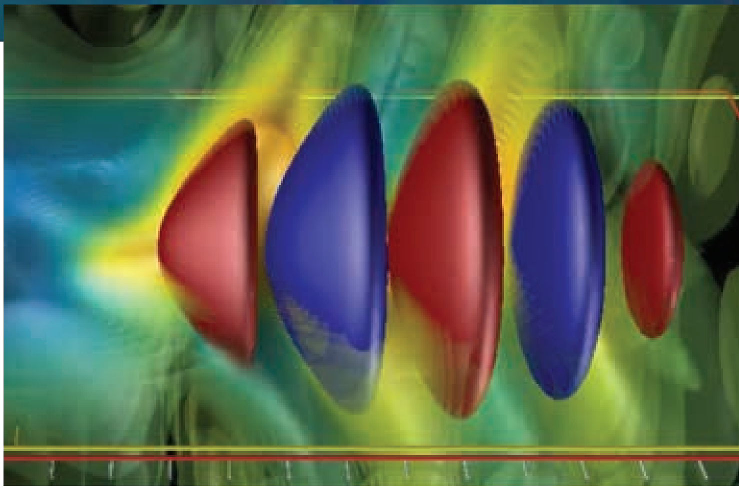

High Energy Physics
Advanced Scientific Computing Research

Scientific Achievement

A team of Berkeley Lab researchers developed WarpIV, a new data analysis/visualization toolkit designed to help speed particle accelerator research and design by enabling *in situ* visualization and analysis of accelerator simulations at scale.

Laser wakefield particle accelerators offer the prospect of less costly and much smaller accelerators. Using the waves created by a laser shot through plasma they "surf" particles to higher speeds.
Image: Jean-Luc Vay, Lawrence Berkeley National Laboratory

SIGNIFICANCE AND IMPACT

Long valued for their role in scientific discovery and in medical and industrial applications such as cancer treatment, food sterilization and drug development, particle accelerators still occupy a lot of space and carry hefty price tags. To take full advantage of the societal benefits of particle accelerators, game-changing improvements are needed in the size and cost of accelerators.

Thus efforts are under way to make this technology more affordable and accessible by shrinking size and cost without losing capability. *In situ* visualization tools such as WarpIV will help physicists design and develop significantly smaller accelerators and reduce the amount of supercomputing time needed to do so.

RESEARCH DETAILS

To create WarpIV, the team combined two software tools already widely used in high energy physics: Warp, an advanced particle-in-cell simulation framework, and VisIt, a 3D scientific visualization application. Together, these tools give users the ability to perform *in situ* visualization and analysis of their particle accelerator simulations at scale—while the simulations are still running and using the same HPC resources—reducing memory usage and saving computer time.

The researchers ran a series of ion accelerator simulations in 2D and 3D to analyze Warp IV's performance and scalability and found significant quantitative differences in the models. This highlights the need for both 3D simulations and *in situ* visualization for accurate modeling of new types of particle accelerators.

Principal Investigators:
Jean-Luc Vay, Oliver Rübel; Lawrence Berkeley National Laboratory

Journal Citation:
O. Rübel, et al, "*In situ* Visualization and Analysis of Ion Accelerator Simulations using Warp and VisIt," *IEEE Computer Graphics and Applications*, 36, May-June 2016, doi: 10.1109/MCG.2016.62

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

Toward Better Materials for Quantum Computing

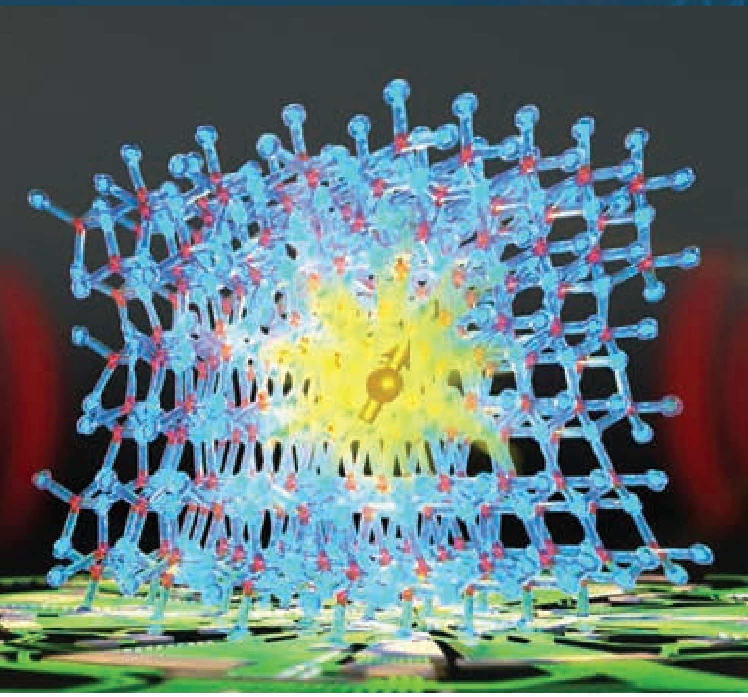

Artist's rendition of an engineered nitrogen vacancy in aluminum nitride. *Image: University of Chicago*

Basic Energy Sciences

Scientific Achievement

One of the leading methods for creating qubits (quantum bits) involves exploiting the structural atomic defects in diamonds. However, using diamond is both technically challenging and expensive. Using NERSC supercomputing resources, researchers from the University of Chicago and Argonne National Laboratory found a way to engineer an analog defect in an alternative material, aluminum nitride, which could lead to reductions in the cost of manufacturing quantum computers.

SIGNIFICANCE AND IMPACT

Quantum computers have the potential to break common cryptography techniques, search huge datasets and simulate quantum systems in a fraction of the time it takes today's computers. But engineers first need to be able to control the properties of qubits, a unit akin to a "bit" in classical computing that is used to store quantum information. At present, one of the most promising solid-state qubits is created when a nitrogen atom occupies a place near a vacant site in a diamond's carbon lattice; this defect is called a nitrogen-vacancy center in diamond.

RESEARCH DETAILS

Using NERSC's Edison supercomputer, the researchers found that by applying strain to aluminum nitride, structural defects can be created in the material that may be harnessed as qubits similar to those seen in diamond. Their calculations were performed using different levels of theory and the WEST code, developed at the University of Chicago, which allowed them to accurately predict the position of the defect levels in the band-gap of semiconductors. NERSC resources enabled the research team to simulate these quantum defects and accurately simulate the surrounding environment.

Principal Investigator:
Giulia Galli, Marco Govoni, University of Chicago

Journal Citation:
H. Seo, M. Govoni, G. Galli, "Design of Defect Spins in Piezoelectric Aluminum Nitride for Solid-State Hybrid Quantum Technologies," *Nature Scientific Reports*, Vol. 6, Article No. 20803, Feb. 16, 2016, doi: 10.1038/srep20803

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

Solving a Plasma Turbulence Mystery

Fusion Energy Sciences

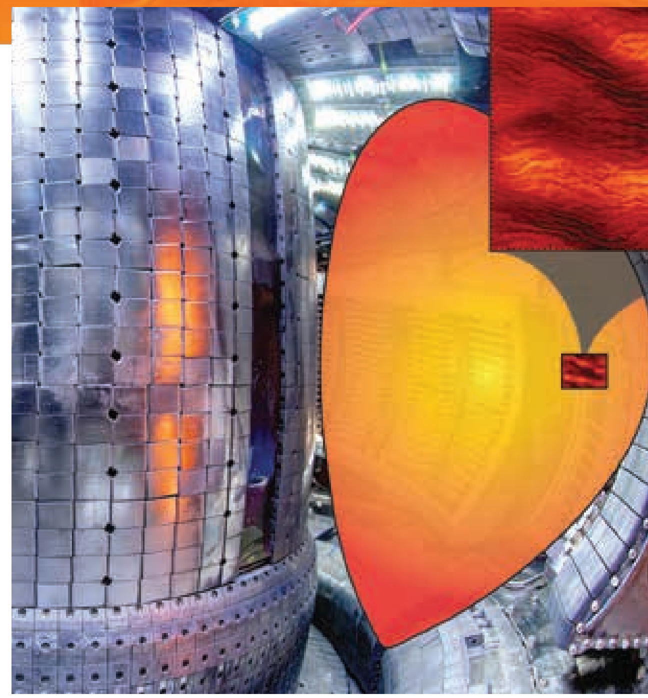

Scientific Achievement

A series of simulations run on Edison found that interactions between turbulence at the tiniest scale (electrons) and turbulence at a scale 60 times larger (ions) accounts for the mismatch between theoretical predictions and experimental observations of heat loss in tokamak fusion reactors.

SIGNIFICANCE AND IMPACT

This research is helping physicists better understand what influences the behavior of the plasma turbulence that is driven by the intense heating necessary to create fusion energy. It also yielded answers to long-standing questions about plasma heat loss that have previously stymied efforts to predict the performance of fusion reactors and could help pave the way for this alternative energy source.

For many years, predictions from leading theories have been unable to explain how much heat loss is coming from electrons in fusion plasma; in this study, the researchers show that by using a coupled model that captures both large- and small-scale turbulence simultaneously, they can reproduce the experimental electron heat losses.

RESEARCH DETAILS

2D and 3D simulations were run on NERSC's Edison system over a two-year period. The multi-scale simulations used the gyrokinetic model implemented by the GYRO code developed by Jeff Candy at General Atomics. By performing high-resolution multi-scale simulations, the team simultaneously modeled multiple turbulence instabilities that have previously been treated in separate simulations. The study took between 100 million and 120 million core hours on Edison. Each simulation required about 15 million hours of computation running on 17,000-30,000 processors.

High-res image of the inside of the Alcator C-Mod tokamak, with a representative cross-section of a plasma. The inset shows the approximate domain for one of the multi-scale simulations and a graphic of the plasma turbulence in the multi-scale simulation.
Image: Nathan Howard, Massachusetts Institute of Technology

Principal Investigator:
Chris Holland, University of California, San Diego

Journal Citation:
N.T. Howard, C. Holland, et al, "Multi-scale gyrokinetic simulation of tokamak plasmas: enhanced heat loss due to cross-scale coupling of plasma turbulence," *Nuclear Fusion*, Vol. 56, Number 1, 2016, doi: 10.1088/0029-5515/56/1/014004

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

Computer Models Work to Improve Nuclear Fuel Recycling

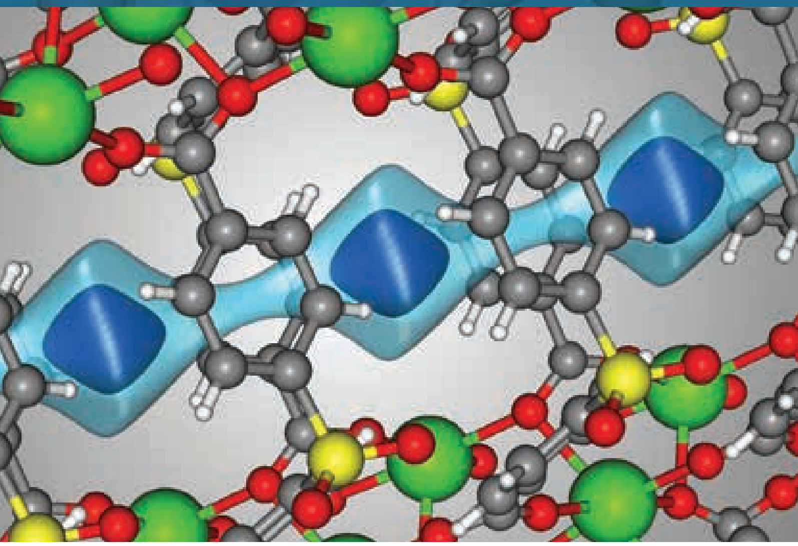

Basic Energy Sciences

Scientific Achievement

A team of computational scientists collaborated with experimentalists to explore the characteristics of materials known as metal-organic frameworks (MOFs), identifying a new material that could aid in nuclear fuel recycling and waste reduction by capturing certain gases released during reprocessing.

A new material could be used to separate xenon and krypton gases from the waste produced in recycling spent nuclear fuels using less energy than conventional methods. In this illustration, the blue surfaces represent the contours of the potential energy of a xenon atom inside the material. *Image: Cory Simon, Lawrence Berkeley National Laboratory*

SIGNIFICANCE AND IMPACT

Recycling nuclear fuel can reuse uranium and plutonium—the majority of the used fuel—that would otherwise be destined for waste. One important step is collecting the radioactive gases xenon and krypton, which arise during reprocessing. Conventional technologies to remove these radioactive gases operate at extremely low, energy-intensive temperatures. By working at ambient temperature, the new material—dubbed SBMOF-1—could be used to separate xenon and krypton gases from the waste produced in recycling spent nuclear fuels using less energy. This would make reprocessing cleaner and less expensive; in addition, the reclaimed materials could be reused commercially.

RESEARCH DETAILS

Using NERSC resources, the researchers sifted through 125,000 molecules to pinpoint one that would be ideal for separating some radioactive gases from spent nuclear fuel without having to use cryogenics. They developed an approach to assess the performance of materials based on their easily computable characteristics. In this case, seven different characteristics were necessary for predicting how the materials behaved; the use of machine learning techniques greatly sped up the material discovery process by eliminating those that didn't meet the criteria.

Armed with these new computational and experimental insights, the researchers can explore SBMOF-1 and other MOFs further for nuclear fuel recycling. These MOFs might also be able to capture other noble gases such as radon, a gas known to pool in some basements.

Principal Investigator:
Maciej Haranczyk, Lawrence Berkeley National Laboratory

Journal Citation:
D. Banerjee, C. M. Simon, et al, "Metal-Organic Framework with Optimal Adsorption, Separation, and Selectivity towards Xenon, *Nature Communications*, June 13, 2016, doi: 10.1038/ncomms11831

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

A Peek Inside the Earliest Moments of the Universe

Nuclear Physics

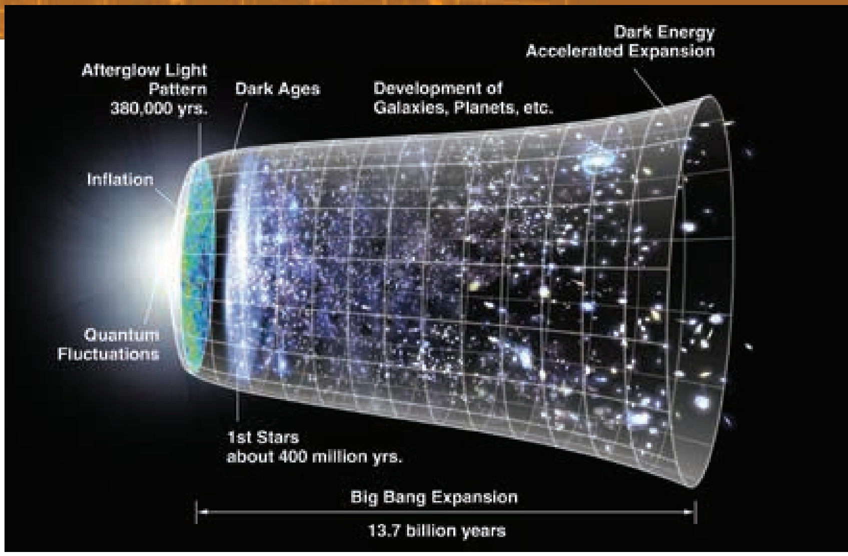

Scientific Achievement

Researchers from the NPLQCD (Nuclear Physics With Lattice QCD) Collaboration used lattice quantum chromodynamics (LQCD) calculations to better understand big bang nucleosynthesis, a process that occurred in the first few minutes following the Big Bang, and precisely measure the nuclear reaction rate that occurs when a neutron and proton form a deuteron. U.S. manufacturing.

Lattice quantum chromodynamics calculations run at NERSC shed new light on big bang nucleosynthesis, a process that occurred in the first few minutes following the Big Bang.
Image: NASA

SIGNIFICANCE AND IMPACT

One of the most critical aspects of big bang nucleosynthesis is the radiative capture process, in which a proton captures a neutron and fuses to produce a deuteron and a photon. Understanding this process and the precise determination of the rate at which it occurs is critical to understanding big bang nucleosynthesis, fusion in the sun and solar neutrino oscillations. These results are the first to be derived directly from the fundamental theory of the strong nuclear force, QCD.

RESEARCH DETAILS

Using NERSC's Edison system and the Chroma LQCD code developed at Jefferson Lab, the researchers ran a series of calculations with quark masses that were 10-20 times the physical value of those masses. These were the first LQCD calculations of an inelastic nuclear reaction, the process whereby a neutron and proton combine to form a deuteron, the nucleus of "heavy" hydrogen.

The calculations also determined the short-distance, two-nucleon interactions with the electro-magnetic field that significantly contribute to the low-energy cross-sections for the radiative capture process. This work reinforces the utility of combining LQCD calculations with low-energy effective field theories describing multinucleon systems.

Principal Investigator:
Martin Savage, University of Washington

Journal Citation:
S. Beane, E. Chang, W. Detmold, K. Orginos, A. Parreno, M. Savage, B. Tiburzi (NPLQCD Collaboration), "*Ab initio* Calculation of the np→dγ Radiative Capture Process," *Physical Review Letters,* 115, 132001, September 24, 2015, doi: 10.1103/PhysRevLett.115.132001

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

Finding the Roots of MJO Modeling Mismatches

Climate researchers are using field data and NERSC supercomputers to better model how the Madden-Julian Oscillation operates and understand its interaction with regional weather systems around the world. *Image: Pacific Northwest National Laboratory*

Biological and Environmental Research

Scientific Achievement

Pacific Northwest National Laboratory researchers used field data and NERSC supercomputers to better model how the Madden-Julian Oscillation (MJO)—a moving disturbance of clouds, rainfall, winds and pressure—operates and gain new understanding of its interaction with regional weather systems around the world. Their work shows the importance of accurately representing the interaction of clouds with the environmental air for modeling the MJO.

SIGNIFICANCE AND IMPACT

Better understanding of the MJO is vital for improving weather forecasting. The unpredictability of the MJO, which occurs every 30 to 60 days and has worldwide impact, makes weather forecasting challenging in many regions of the world, including the Western U.S., Australia and South Asia. Despite decades of work, however, simulating the MJO in climate models and understanding the instabilities that drive it remain difficult. In particular, there has been a lack of fine enough resolution to capture the cloud processes that lie at the core of MJO dynamics.

RESEARCH DETAILS

Using NERSC's Edison supercomputer and data gathered during a field campaign over the Indian Ocean, the researchers identified the processes that are responsible for too much precipitation in climate models, especially during the low-rainfall period of the MJO signal. The simulations, which used 2,400 cores on Edison and about 250,000 computer hours, found that the mismatches are related to the fact that most models get the relationship between environmental moisture and precipitation wrong, producing more precipitation in the models than is observed for the same moisture content in the environment.

Principal Investigator:
Samson Hagos, Pacific Northwest National Laboratory

Journal Citation:
S. Hagos, Z. Feng, C. Burleyson, et al, "Moist Process Biases in Simulations of the Madden-Julian Oscillation Episodes Observed during the AMIE/DYNAMO Field Campaign," *Journal of Climate,* February 2016, doi: 10.1175/JCLI-D-15-0349.1

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

Toward Cost-Effective Electrolyte Polymer Fuel Cells

Basic Energy Sciences

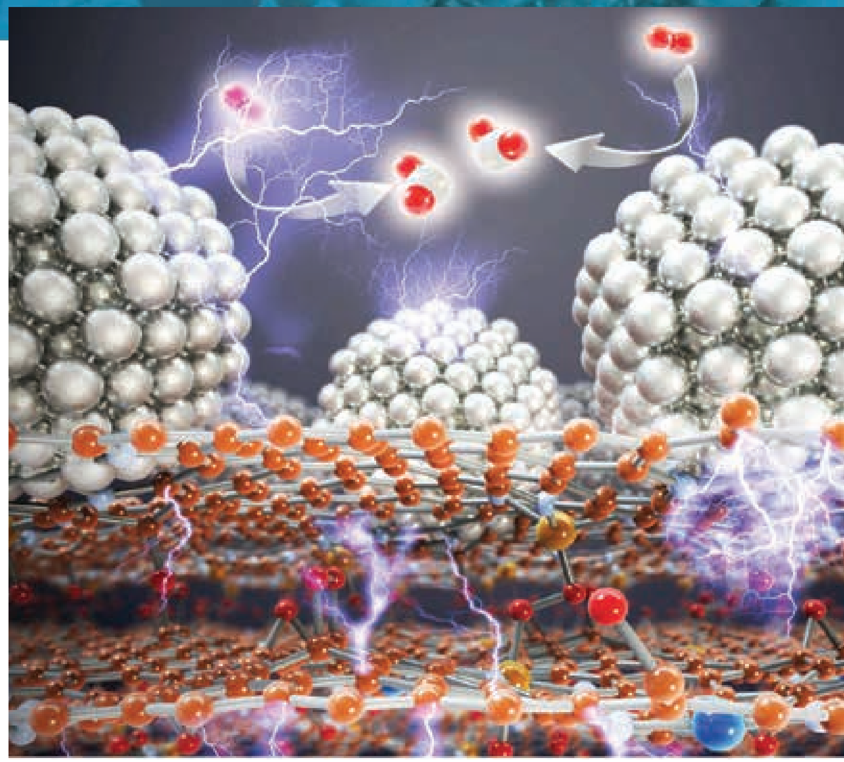

Calculations run at NERSC helped demonstrate how polymer electrolyte fuel cells can be made to run more efficiently and produced more cost-effectively by reducing the amount of a single key ingredient: platinum. *Image: Journal of Materials Chemistry A*

Scientific Achievement

An international team of researchers running density functional theory (DFT) calculations at NERSC demonstrated how polymer electrolyte fuel cells (PEFCs) can be made to run more efficiently and produced more cost-effectively by reducing the amount of platinum needed.

SIGNIFICANCE AND IMPACT

As we enter the age of hybrid, electric and self-driving cars, interest remains high in finding the next generation of fuel cell technology that is low cost, long lasting and mass producible. In recent years, fuel cell research and development efforts have been focused on improving performance and reducing costs through the development of new materials and better water and heat management processes.

Polymer electrolyte fuel cells (PEFCs) are considered the leading candidate for use in transportation applications and for megawatt-size power generation systems. But figuring out how to produce these fuel cells efficiently and cost-effectively at mass scale has proven to be a challenge. The culprit? Platinum—one of the key ingredients in PEFCs—is expensive.

RESEARCH DETAILS

The researchers used about 300,000 core hours on NERSC's Edison system to run their calculations, and test a non-precious-metal-based cathode material containing only a small amount of platinum. Their analysis confirmed that platinum nanoclusters on metal-nitrogen doped ordered mesoporous porphyrinic carbon significantly enhance the material's oxygen reduction reaction activity. Looking ahead, they believe that DFT-based simulations can guide the design of new catalytic materials that have high activity and stability.

Principal Investigator:
YongMan Choi, SABIC Technology Center

Journal Citation:
S. Hwang, Y. Choi, et al, "Enhancement of oxygen reduction reaction activities by Pt nanoclusters decorated on ordered mesoporous porphyrinic carbons," *Journal of Materials Chemistry A*, April 2016, 4, 5869-5876, doi: 10.1039/c5ta09915c

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

User Support and Outreach

In 2016, the NERSC Program supported 6,915 active users from universities, national laboratories and industry. These users came from 49 U.S. states and the District of Columbia and a total of 45 countries.

To meet the computing and data analytics needs of our users, NERSC employs a diverse set of consultants from the User Engagement, Application Performance and Data Science Engagement groups to provide the first line of support linking NERSC and its user community. These consultants, all of whom have master's degrees and/or P.hD.s in science or computational science, are experts in HPC as well as various science domains. They are responsible for problem management and consulting, helping with user code optimization and debugging, strategic project support, web documentation and training, third-party applications and library support, running Office of Science-wide requirements reviews and coordinating the NERSC User Group (NUG). NERSC's consultants and account support staff are available to users via email and an online web interface. Basic account support (password resets, resetting login failures) is available online or via the NERSC operations staff, 24 x 7, 365 days a year.

NESAP: A Collaborative Success

To help NERSC users make the transition from systems with traditional CPUs to the Cori system's energy-efficient manycore processors, in 2014 the center launched the NERSC Exascale Science Application Program (NESAP). The first phase of NESAP has been a collaboration between NERSC, the users and the Cori system vendors (Cray and Intel). NERSC chose to partner with approximately 40 different code teams in preparing their applications for Cori. Twenty of the teams received a commitment of staff time from NERSC (in the form of a NERSC staff liaison and/or a NESAP postdoc), and all teams received access to test systems, training and vendor resources.

The program is still going strong. NESAP highlights in 2016 included:

- Completion of 12 of 16 Intel "dungeon sessions"
- Hiring of eight postdocs: Bill Arndt, MPAS; Taylor Barnes, Quantum ESPRESSO; Kevin Gott, PARSEC; Tuomas Koskela, XGC1; Mathieu Lobet, WARP (he has since assumed a permanent position at NERSC); Tareq Malas (he has since taken a permanent position at La Maison de la Simulation); Andrey Ovsyannikoc, Chombo-Crunch; and Zahra Ronaghi, Tomopy
- Participation of multiple Cray engineers on NESAP projects
- Many user training sessions provided by Cray, Intel and NERSC staff on KNL, tools and application performance
- Multiple "hackathons" hosted at NERSC, including a NESAP all-hands Cori hackathon in November.

NERSC hosted a number of training events and workshops in 2016 to help users improve their skills across a variety of applications and codes.

The NERSC facility plays a central role in the NESAP program, acting as a bridge between the user community and the architects and engineers at the processor vendor (Intel) and the HPC integrator (Cray). For the scientists who make up the majority of NERSC's user base, code optimization can be a daunting task, given that the KNL processor has multiple possible novel architecture features to explore just at the node level:

- Manycores: Each KNL node on Cori has 68 cores (compared to 24 on Edison across two sockets)
- More vector parallelism per core that, combined with a more powerful instruction set, is capable of performing 32 floating point operations (FLOPs) each cycle
- The KNL processor lacks a shared L3 cache amongst the cores but contains 16GB of high-bandwidth MCDRAM right on the package that supports well over 400 GB/s of stream bandwidth

The NESAP effort has sped up our partner codes significantly—by about 3x on Cori KNL—and has improved the applications on Cori Haswell and Edison as well by an average of 2x for codes currently reporting. We are working more closely with the 20 additional tier 3 applications and communicating the lessons learned to the general NERSC user community via an active training schedule and comprehensive web documentation.

Following these successes, in 2016 NERSC expanded the NESAP program to include NESAP for Data, selecting six science application teams to participate in the new program. Like NESAP, NESAP for Data joins application teams with resources at NERSC, Cray and Intel; however, while the initial NESAP projects involve mostly simulation codes, NESAP for Data targets science applications that process and analyze massive data sets acquired from DOE-supported experimental and observational sources (For more information about the NESAP for Data program, see p. 18.)

NERSC staff work closely with users during workshops and hackathons, providing hands-on assistance and feedback.

NERSC Expands User Trainings and Workshops

NERSC regularly holds training events and workshops aimed at improving the skills of users of NERSC resources. In 2016 we held new user training in conjunction with the annual NERSC Users Group meeting and offered a training for users new to the KNL architecture just before Cori KNL came online. We also offered training for NERSC users in particular application areas. For example, users of the materials science applications VASP and Quantum Espresso had a training session focused on those two applications in June, and in November NERSC hosted Martijn Marsman from the University of Vienna, the primary developer of the VASP code, to present a three-day workshop on VASP, attracting more than 150 attendees.

NERSC also secured training from computer science experts. At the beginning of the year we hosted Michael Klemm and Bronis R. de Supinski, both prominent members of the OpenMP standards committee, who gave a tutorial on advanced OpenMP topics for NERSC users. In March, NERSC hosted a hands-on training for advanced users on the Cray tools, featuring experts from Cray. And in June a training session on programming and optimization for Cori Haswell with Cray experts was held.

Additionally, we held two trainings on open-source performance tools developed at other institutions. In July, we hosted presenters from Juelich and LLNL for a Score-P and Scalasca training, and in August we hosted a presenter from the Barcelona Supercomputing Center for a Paraver and BSC Tools training.

NERSC experts were involved in training future computational scientists at two student-oriented workshops:

- At the Computational Science Graduate Fellowship Annual Review Workshop, NERSC presented a session on parallel algorithm design and co-presented another session on network troubleshooting

- At the Richard Tapia Celebration of Diversity in Computing, NERSC experts again presented a session on parallel algorithm design and gave a workshop on cybersecurity as part of the National Laboratory Day workshops.

EVENT	DATE(S)	BRIEF SUMMARY
Advanced OpenMP Training	February 4	One-day advanced OpenMP tutorial from two members of OpenMP committee
NUG New User Training & Hackathon	February 21	Parallel new user training for beginners and hackathon for advanced users to optimize their codes
Materials Science Application Training	June 10	How to use pre-installed materials science applications VASP and Quantum Espresso at NERSC
Cori Phase I Programming and Optimization Training	June 13–16	Training presented by Cray and NERSC about the new Cori XC40 and how to use it and optimize software for it
Score-P and Scalasca Performance Tools Training	July 26	Overview of how to use two popular performance analysis tools, presented by their developers
Paraver & BSC Tools Training: Understanding Application Performance with Paraver	August 5	Overview of how to use performance tools, presented by their developer
Data Day	August 22-23	A series of talks and tutorials on the latest data-focused tools for scientific computing
2016 C++ Summit	October 17–18	Tutorials and presentations from leaders in the C++ community
Cori KNL Training	November 3	Introduction to using KNL architecture and various tools on Cori
VASP Workshop	November 9–11	Learn to use VASP more efficiently via lectures and hands-on tutorials, presented by the developer of VASP
NESAP Workshop and Hackathon	November 29–December 1	In-person workshop for NESAP teams to learn to use performance tools

The table above lists key training events held at NERSC in 2016.

In addition, NERSC staff gave more than 50 presentations and tutorials at conferences and workshops in the U.S. and internationally. These events included SC16, IXPUG, Cray Users Group Meeting, JGI User Meeting Workshop, HEPiX (Spring and Fall), ISC 2016, Intel HPC Developers Conference, APS March Meeting, NCAR Multicore 6 Workshop, American Geophysical Union Fall Meeting, SIAM Conference on Parallel Processing for Scientific Computing and the 251st American Chemical Society National Meeting & Exposition.

NERSC, in collaboration with ALCF, OLCF and the IDEAS project, also hosted a series of webinars on Best Practices for HPC Developers. Presentations were given by HPC software development experts on a range of topics, including distributed version control, testing and documenting, parallel I/O and performance analysis and optimization. The webinars used NERSC's Zoom webinar platform and have been posted on the NERSC Training YouTube channel.

STEM Outreach to the Community

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. In 2016, NERSC supported a number of STEM-related efforts for students:

- A group of 50 students from Dougherty Valley High School in San Ramon, CA visited NERSC in May, where they toured the computer room and participated in lively discussions about the facility and how supercomputers work. They asked great questions, such as "In the future, will there be supercomputers the size of a microchip?" and "What would happen to the supercomputers if there was a huge earthquake or tsunami?" All were AP Computer Science students; most will be majoring in computer science or engineering in college.

Top: Fifty computer science students from Dougherty High School in San Ramon, Calif. toured NERSC in May.

Bottom left: NERSC's first-ever student cluster competition team attended ISC16.

Bottom right: NERSC HPC consultant Brian Friesen (second from left) helped Cray Inc. surprise an inquisitive Southern California high school student.

- In June, five former NERSC interns and one then-current intern joined forces to participate in the Student Cluster Competition at ISC16, marking the first time ever that NERSC fielded a student cluster competition team. The all-female team, whose members came from across the U.S. as well as Canada and Puerto Rico, comprised two high schoolers and four college undergrads. The team was guided by Rebecca Hartman-Baker, acting group lead for NERSC's User Engagement Group. Hartman-Baker and the team worked closely with team sponsors Intel and Cray to select and procure the parts they used to build their cluster.
- When Cray received a letter in February from a high school student in Southern California asking if one of their computers could help him solve a complex math problem —the probability of surviving a zombie apocalypse—Cray reached out to NERSC for some support. While Cray engineers determined that the equation was too large to compute the answer exactly or store it in computer memory (and sent the student, Hugo Villanueva of John F. Kennedy Middle College High School, a two-page letter explaining in detail why), they decided this was a great opportunity to pay a surprise visit to the young man's school to share more. Cray enlisted the help of NERSC HPC consultant Brian Friesen, who talked to the students about his supernovae research, the infrastructure involved in keeping NERSC supercomputers up and running and the types of classes the students should take to work with supercomputers for a living.

Application Portability Activities

Performance portability is a significant issue for NERSC users, who typically have accounts and allocations at multiple HPC centers that include GPUs, CPUs and manycore processors. To help address this challenge, NERSC participated in a number of portability activities in 2016:

- Through NESAP, NERSC is actively promoting a code optimization strategy that improves applications across multiple architectures. We highlight the need to identify and exploit more on-node and vector parallelism, which is relevant to both GPU and manycore x86 architectures. We discourage vendor-specific code, such as the use of explicit AVX intrinsics and TBB. While explicitly accessing/allocating data in MCDRAM on the KNL processors typically involves Intel-specific directives, the identification of "hot" arrays or data structures additionally helps identify data for placement in GPU memory.

- NERSC is also working with OLCF and ALCF on an Office of Science portability effort to identify the performance trade-off of portability on multiple applications and will jointly publish user guidelines and best practices in 2017. Members of NERSC's Application Performance Group are participating in this effort.

- NERSC played a prominent role in organizing the performance portability workshop held in Phoenix in April 2016. Multiple staff members attended and presented and are organizing the 2017 event. NERSC staff participated in a CORAL OpenMP for GPUs hack-a-thon at IBM Yorktown in September. They ported multiple BoxLib routines (a NESAP and ECP code) to portable OpenMP 4.x with target directives and compared performance of non-portable optimized codes on CPUs and GPUs.

- NERSC has two staff participating directly with the OpenMP and MPI standards bodies. In the past year, NERSC actively surveyed users of MPI, OpenMP and emerging programming models to find deficiencies in standards and promote changes that benefit HPC users.

- NERSC hosted a two-day C++ workshop (October 17-18) focused on portability and standardization of C++ parallel frameworks. C++ standards committee members, code teams, Kokkos, Raja, UPC++, Charm++ and HPX were all represented. Half-day tutorials on HPX and Kokkos were given to the user community.

NERSC Hosts its First 'Data Day'

In August, the Data and Analytics Services (DAS) team hosted an entirely new NERSC event: Data Day. Designed for researchers who use, or are interested in using, NERSC systems for data-intensive work, the event drew more than 70 in-person and 30 remote attendees. Interest from industry, academia and a wide range of disciplines within DOE labs made for a diverse and enthusiastic audience.

Participants enjoyed a full day of talks, demos and tutorials on the latest data-focused tools for scientific computing, including forays into machine learning, Python, Spark, visualization tools, the Cori Burst Buffer and data management. A lively poster session rounded out the day's activities, showcasing work and work-in-progress utilizing NERSC resources for experimental and observational research. The event was followed by a half-day hackathon for those who wanted to put their new knowledge to work creating a custom workflow with Spark. Buoyed by overwhelmingly positive reviews, NERSC plans to make Data Day an annual event.

Using Containers to Optimize Science Gateways

NERSC offers users the option to share data with collaborators and others via web-enabled science gateways. For many research teams, this approach is the most practical way to reach scientists outside the immediate NERSC user community.

One excellent example of the future of such gateways is the new Sloan Digital Sky Survey (SDSS) Mirror website. The science gateways team within DAS worked with SDSS

NERSC's first "Data Day," held in August, was so popular the center plans to hold the event annually.

researchers to develop a new service that allows users to create sandboxed virtual containers to manage their services. The SDSS portal takes advantage of Docker container technology, which allows users to create a lightweight virtual machine, called a container, that presents a private view of the operating system to the user. Users can install and manage custom software dependencies and services in this container, independent of the base host. Containers can also be easily ported from one Docker host to another without any changes and then moved to host sites outside of NERSC if users wish. The SDSS team has enabled data mirroring and Subversion version control services under this platform and allows the project to manage these services directly without NERSC staff intervention. This effort was designed to pilot a more general-purpose service through which we expect a significant number of our gateways to be deployed as Docker containers.

Another Docker-based science gateway, Tomostore, created for Berkeley Lab's Center for Advanced Mathematics for Energy Research Applications (CAMERA), introduced the use of CKAN, a dataset sharing platform built on open web standards. Combining CKAN and Docker enabled rapid deployment of a feature-rich tomography data-sharing site in record time. Tomostore was the product of a collaboration between DAS engineers and scientists at Berkeley Lab's Advanced Light Source.

HPC Community Embraces Burst Buffer

Following the deployment of the Burst Buffer on Cori and the Burst Buffer Early User Program in 2015, in 2016 NERSC made a major push to promote the effectiveness of Burst Buffer architectures in the HPC community through conferences, workshops and working groups.

The Burst Buffer Early User Program has been instrumental to these efforts. During 2016, several user applications were optimized for the Burst Buffer across dozens of applications from different projects and science domains. The primary use case for the Burst Buffer architecture was to accelerate checkpoint/restart operations; however, the use cases at NERSC have turned out to be much more diverse, ranging from coupling complex workflows and staging intermediate files to database applications.

This work was showcased in several papers published by NERSC staff in 2016 (see "Publications," p. 55). It also led to two NERSC presentations at SC16. The first, "Scientific Workflows at DataWarp-Speed: Accelerating Data-Intensive Science using NERSC's Burst Buffer" was published as part of the 2016 PDSW-DISCS workshop, held in conjunction with

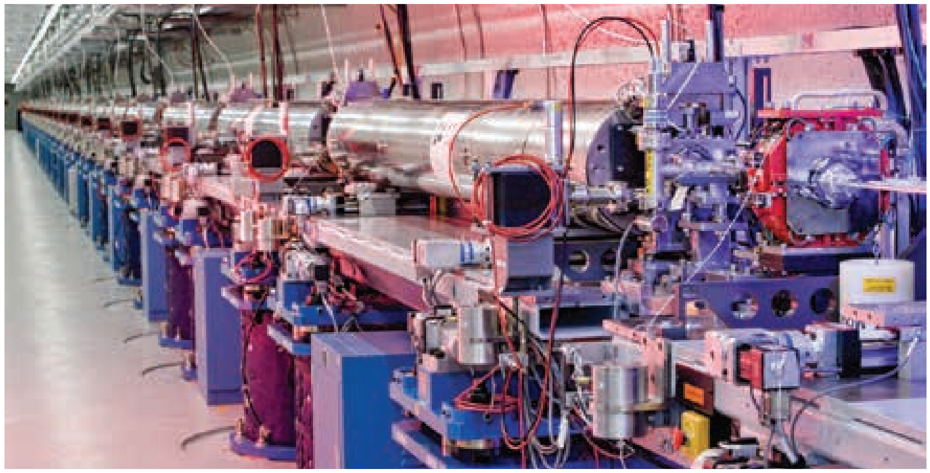

The Linac Coherent Light Source (LCLS) at SLAC uses NERSC's Shifter tool for beamline processing.
Image: SLAC

SC16. That paper looked at using the Burst Buffer to accelerate the workflow of the Berkeley Lab Chombo-Crunch application, a high performance subsurface simulator used for modeling pore scale reactive transport processes associated with carbon sequestration. The second, "Performance Characterization of Scientific Workflows for the Optimal Use of Burst Buffers," was a lightning talk at the Workflows in Support of Large-Scale Science 2016 workshop.

It is worth noting that the close collaboration between NERSC and Cray has resulted in many changes to the DataWarp software (Cray's I/O accelerator on which the Burst Buffer is based). These changes have significantly improved Burst Buffer performance (based on NERSC use cases). Following these improvements, the Burst Buffer was opened to all users at the beginning of 2017, and users are typically seeing a 5x-10x improvement in I/O compared to Lustre.

While the Cray XC40 includes the first commercial, production deployment of a Burst Buffer architecture, the use of non-volatile storage in supercomputers is still a nascent feature and will evolve to more complex and higher level abstractions in exascale systems. Relatedly, NERSC is leading (in collaboration with other DOE sites) the Tiered Storage Working Group. This is a group of DOE, industry and academic representatives convened to move forward ideas for unifying APIs to access nonvolatile storage in HPC systems. These discussions focus on creating opportunities to address traditional I/O challenges in new ways as the storage hierarchy deepens and changes. The working group had a kick-off meeting in August 2016 at Lawrence Livermore National Laboratory and a face-to-face in February 2017 that was hosted by NERSC.

Shifter Moves into Production

Shifter, a modified, lightweight container deployment system developed at NERSC in conjunction with Cray, was deployed on all major NERSC systems in 2016. Shifter has greatly improved usability of the HPC systems by allowing users to choose and build their own environments, making it easier for them to install difficult software stacks and dependencies. Shifter can use common frameworks like Docker, so images can be used at other sites to facilitate scientific reproducibility. Shifter also provides several performance benefits, with the fastest shared library loading time and the ability to mount a per-node XFS file for very fast local-node I/O.

A number of groups have incorporated Shifter into their scientific computing at NERSC:

- The Linac Coherent Light Source (LCLS) at SLAC uses Shifter for beamline processing. Their software installation process is very complicated and depends heavily on Python and shared libraries. Shifter offers ease of installation and good performance of the LCLS software stack.

- STAR (Solenoid Tracker at RHIC [Relativistic Heavy Ion Collider]) at Brookhaven National Lab developed a deployment scheme where their database with detector conditions is copied into the XFS file, bringing database query times from >30 minutes on the Lustre file system to less than 1 second. In 2017, this scheme will be used for a 25M CPU hour simulation campaign.
- Engineers at NERSC and Fermilab worked together to develop a framework to bring HEPCloud jobs to NERSC. The software for most of the HEPCloud groups relies on CVMFS to deliver their terabyte-sized software stacks to the compute resources. Since CVMFS is not installed on the Cray systems at NERSC, NERSC engineers used Shifter to build mammoth images (>300 GB after compression) to deliver the CMS software environment for a pilot campaign. Using this Shifter image, NERSC contributed 2 million CPU hours to three separate CMS production campaigns for data and Monte Carlo reconstruction, simulation and generation. This work laid the foundation for future HEPCloud runs at NERSC.
- Shifter containers for the DayaBay neutrino detector and LZ dark matter experiment have also been developed. DayaBay and LZ traditionally run data analysis and simulation jobs on NERSC's PDSF cluster (a ~3,000-core Linux cluster). These codes are embarrassingly parallel, with each job using a single core. By utilizing the task farmer workflow engine and using a scalable Shifter image in which one copy of their conditions database was deployed for every 1,216 jobs, we demonstrated that both workflows could run on Cori at much higher concurrency than is available on PDSF, leading to faster scientific insights.

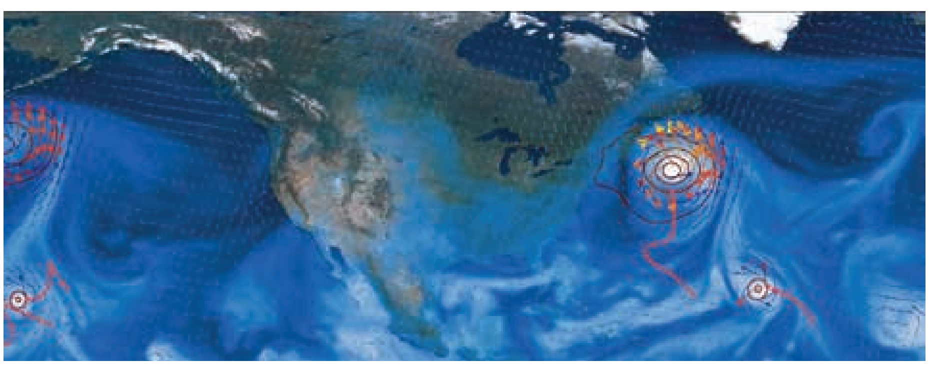

Scientists are increasingly looking to HPC systems and deep learning tools to address large data analytics problems, such as extreme weather events. *Image: Prabhat, Lawrence Berkeley National Laboratory*

NERSC Offers Users Deep Learning Support

Deep learning is enjoying unprecedented success in a variety of commercial applications, but it is also finding its footing in science. Just a decade ago, very few practitioners could have predicted that deep learning-powered systems would surpass human-level performance in computer vision and speech recognition tasks.

At Berkeley Lab, we are confronted with some of the most challenging data analytics problems in science. For example, extreme weather events pose great potential risk on ecosystem, infrastructure and human health. Analyzing extreme weather data from satellites and weather stations and characterizing changes in extremes in simulations of future climate regimes is thus an important task. Similarly, upcoming astronomical sky surveys will obtain measurements of tens of billions of galaxies, enabling precision measurements of the parameters that describe the nature of dark energy. But in each case, analyzing the mountains of resulting data poses a daunting challenge for the scientific community.

Users are thus increasingly looking to HPC systems to solve their large data analytics problems and to apply sophisticated analysis tools like deep learning and machine learning on their large data sets. Toward this end, in 2016 NERSC expanded its support for deep learning, creating an all-encompassing environment that includes libraries for Theano, TensorFlow, Keras, Lasagne, XGBoost, and Caffe. Our goal is to eliminate the hassle of building and linking software and changing paths required to set up a machine learning environment, replacing it instead with a single module load so that users can worry about science instead of tedious software build details. A notebook kernel is included so users can use the environment in a Jupyter notebook. Scientists from many fields, including geosciences, high energy physics, climate, fusion and astrophysics, have already made use of this package.

Jupyter Gets an Upgrade

Jupyter is a powerful, literate-computing web application for creating and sharing “notebooks” that contain code, equations, graphs and text. Since 2015, NERSC has provided users a Jupyter service running on a science gateway node. This configuration gives users access to the NERSC Global Filesystem (NGF) for interactive data exploration and analysis.

While the Jupyter service has proved popular, NERSC users have also requested additional features and suggested improvements. Two issues stood out across a number of consulting tickets submitted in 2016. First, Lustre scratch filesystems were inaccessible from the science gateway. This meant that users had to orchestrate staging of data from scratch to NGF, and then only if the data sets were small enough. Second, the Python stack on the science gateway was different from that on the Cray systems. As a consequence, code run in a notebook might not be portable to the Cray systems due to dependency mismatches. These shortcomings affected user workflows in science disciplines as diverse as climate, metabolomics and cosmology.

Over the course of a few weeks in the summer of 2016, NERSC was able to address these and other shortcomings by leveraging a reserved Cray login node, GSI-enabled SSH and LDAP authentication to launch a Jupyter service on Cori. Top users of the Jupyter service have been invited to try out the new service, and the feedback has been positive. More important, by solving this problem for users we have taken steps toward enabling interactive supercomputing through Jupyter at NERSC.

NERSC User Support: Two 2016 Success Stories

Celeste astronomy code: A team of researchers from Berkeley Lab, Intel, Julia Computing and MIT were working to meet the IPDPS 2017 paper submission deadline in September 2016. Unfortunately, their Celeste code, which finds and characterizes astronomical objects from telescope data, encountered some MPI RMA errors. The developer, Kiran Pamnany at Intel, who implemented the global array tool kits (called Garbo) with MPI one-sided messaging for the Celeste code, came to NERSC looking for someone with experience with MPI-3 RMAs on Cray systems to help address the MPI RMA errors.

A User Support consultant read the user code (written in C and Julia) and was able to isolate the problem as a Cray MPI RMA implementation issue and create a workaround by modifying

The Celeste code has been used to analyze datasets from the Sloan Digital Sky Survey, which has created the most detailed 3D maps of the Universe ever made. *Image: SDSS*

the MPI RMA calls to use point-to-point synchronization instead of the collectives. While that helped, the code still failed later with other MPI errors, so NERSC prepared a test case for Cray so that they could reproduce the error on their internal systems without building the whole dependent software package. Cray identified and fixed a bug in their MPI RMA code. NERSC requested early access to the patched library before its official release, confirmed the bug fix and provided it to the Celeste team.

NERSC's support made it possible for the team to submit their IPDPS paper in time and achieve a number of first results, as mentioned in a press release issued by Julia Computing in November 2016 (http://juliacomputing.com/press/2016/11/28/celeste.html).

Vienna *Ab initio* Simulation Package (VASP): VASP is a computational materials science code that has been one of the top two codes at NERSC for many years. It alone consumes more than 10 percent of computing cycles each year at NERSC, and there are more than 850 registered VASP users at NERSC. To ensure that the Cori KNL system is used efficiently, it is critical to get VASP code and users ready for this new architecture.

Toward this end, NERSC has been involved with the VASP code optimization effort through NESAP. So far the officially released VASP is still an MPI-only code; however, in collaboration with Intel, the Zuse Institute Berlin and NERSC, the VASP developers have completed the transition from the MPI-only to an MPI+OpenMP hybrid code, which has been highly optimized for KNL. NERSC's User Engagement Group (UEG) has been focusing on getting VASP users ready for KNL and exploring the KNL specifics relevant for the build and runtime setup to help users achieve optimal performance. Here are some highlights from 2016:

- In November 2016, UEG consultants hosted a three-day VASP workshop at NERSC, inviting Martijn Marsman at University of Vienna, the MPI+OpenMP hybrid VASP code developer. There were more than 150 registrations each day (1/3 were in-person attendees from all over the U.S.). The workshop featured lectures and hands-on interactions. The lectures were recorded and made available on the YouTube NERSC training site. For the first time, the workshop attendees ran the pre-release MPI+OpenMP hybrid VASP code.

- UEG consultants explored the KNL specifics relevant for the build and runtime setup—such as the NUMA/MCDRAM modes, hyper-threading, hugepages, task/thread/memory affinity, MPI tuning options and core specialization—and carried out extensive performance tests with the MPI+OpenMP hybrid VASP on Cori KNL. Through this effort they were able to resolve an initial VASP performance issue on Cori KNL and have now achieved optimal performance.

- In collaboration with the VASP developers, NERSC is hosting a beta-testing program for the hybrid VASP on Cori KNL. In addition, a Cori KNL early user access program for VASP and QE was launched that allows users to access the privileged queue, with 24-hour wall limit and fast queue turnaround.

Center News

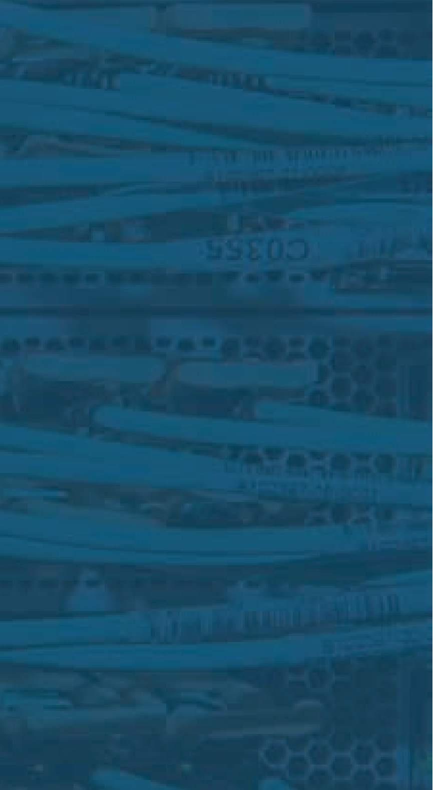

Cori Makes Official Debut

While 2015 was a year of transition for NERSC, with a new building, systems retirements and a large influx of new staff, 2016 was a year where we looked to the future. Most notably, the center's newest supercomputer, Cori, a Cray XC40 with a peak performance of 30 petaflop/s, was fully deployed and is now helping NERSC's nearly 7,000 users broaden the scope of their science research.

Preparing for and deploying Cori in a brand new facility was a top priority in 2016. The system was delivered in two phases. The Cori Data Partition—also known as Cori Haswell—was installed in late 2015 and comprises 12 cabinets and more than 2,300 Haswell compute nodes. It was customized to support data-intensive science and the analysis of large datasets through a combination of hardware and software configurations and queue policies. Cori KNL, installed in mid-2016, added another 52 cabinets and more than 9,600 Intel Xeon Phi KNL compute nodes, making Cori the largest supercomputing system for open science based on these processors. The KNL processors have three distinct features that are new for users and offer the potential for increased application performance: longer vector units, high-bandwidth memory and an energy-efficient manycore architecture.

The combined Cori system is the first to be specifically designed to handle the full spectrum of computational needs of DOE researchers, as well as emerging needs in which data- and computer-intensive work are part of a single workflow—a model that will be important in the coming exascale era.

Thus in 2016, NERSC staff incorporated a number of innovative capabilities into Cori to support these new workflows:

- A Burst Buffer for improved I/O. Based on the Cray DataWarp I/O accelerator, the Burst Buffer on Cori is a 1.5PB layer of NVRAM storage designed to move data in and out of the processor cores more quickly, which improves the overall performance of the system. This helps researchers make more effective use of the system.

- A real-time queue for time-sensitive analyses of data. Users can request a small number of on-demand nodes if their jobs have special needs that cannot be accommodated through the regular batch system.

- Software defined networking (SDN). SDN capabilities have been integrated into Cori to allow scientists with scientific data-intensive workloads at experimental facilities to co-schedule networking bandwidth, compute resources and Burst Buffer bandwidth.

- Container-based tools. To help users better access, manage and analyze their increasingly large data sets, NERSC has enabled Docker-like container technology on its systems through a customized, scalable software toolkit known as Shifter. Shifter leverages container-based computing to help supercomputer users run a wider range of software more easily and securely. A new version of Shifter was released in August 2016, improving its functionality and scalability.

The Cori system was delivered in two phases. The Cori Data Partition—also known as Cori Haswell—comprises 12 cabinets and more than 2,300 Haswell compute nodes. Cori KNL, added another 52 cabinets and more than 9,600 Intel Xeon Phi KNL compute nodes. *Image: Marilyn Chung, Lawrence Berkeley National Laboratory*

Jupyter Notebooks Open up New Possibilities on Cori

NERSC is using Jupyter notebooks to help users more easily access and use the center's Cori supercomputer. Although a number of university and NSF-supported computing centers are using Jupyter, NERSC is the first DOE supercomputing center to do so.

Jupyter notebooks are an open-source web application that allows users to create and share documents that contain live code, equations, visualizations and explanatory text. Uses include: data cleaning and transformation, numerical simulation, statistical modeling and machine learning.

NERSC originally offered the Jupyterlab notebook for multi-user environments on a "Science Gateway" node; when users log in, they receive a notebook for their projects that provides access to NERSC services, files and jobs.

In 2016, NERSC began running a Jupyterlab prototype directly on Cori, providing the center's users with three key benefits:

- Direct access to large datasets, including terabytes written to the scratch filesystem. Since Jupyterlab will run directly on Cori, users will have direct access to the scratch and project filesystems, allowing them to write and view the data.
- Access to the job queueing system, allowing users to submit jobs, query the batch systems and look at results.
- Easy login with username and password.

The web-server version of Jupyter has been integrated into six different user projects at NERSC, including OpenMSI, a cloud-based platform hosted at NERSC and built on Jupyter that allows mass spectrometry imaging data to be viewed, analyzed and manipulated in a standard web browser. Users can log in from their own computers using the notebook, see all the OpenMSI files, add ions of particular interest (which the notebook then grabs from data stored at NERSC and inserts into the notebook) and then center the samples on a grid for analysis.

NERSC/ESnet/Ciena Collaborate on Chip Design

In 2016, computer scientists and mathematicians from NERSC and ESnet worked with engineers at Ciena, a leading networking company, to speed up the process by which Ciena validates the design of its ASIC (application-specific integrated circuit) chips. The collaboration grew out of the existing relationship between Ciena, a pioneer in high-bandwidth optical transport technology, and ESnet, which uses Ciena products to support its high-speed network.

To help Ciena better utilize computational methods in designing and producing their next-generation optical networking products, staff in Berkeley Lab's Computational Research Division demonstrated the feasibility of accelerating computational verification of forward error correction (FEC) codes, which are commonly used in optical transmission equipment for controlling errors in data transmission. To do this they modified a random number generator library called MRG8 (multiple recursive generator with 8th-order recursion), then used 8 million supercomputing hours at NERSC to test and validate the efficacy of the enhanced FEC codes in Ciena modems. The parameter study made use of a task farmer developed at NERSC to support high-throughput parameter research such as the Ciena FEC code study.

Using the parallel processing resources at NERSC had a dramatic impact on the time it took for Ciena to validate the design of its ASIC chips. *Image: Ciena*

To show the benefits of this technology, the team simulated sending nine quadrillion bits of data in a noisy environment, where the channel impairments caused about 500 trillion of these bits to be received in error. The FEC mechanism corrected all 500 trillion errors and ensured us that the bit error rate is below 10^{-16}.

Using the parallel processing resources at NERSC had a dramatic impact on the time it took to run the experiments and validate the design of the FEC algorithm. With NERSC resources, Ciena was able to significantly expand the scope of its study on a compressed timescale.

HPC4Mfg Program Expands

In August 2016, the DOE's High Performance Computing for Manufacturing (HPC4Mfg) Program announced $3.8 million in funding for 13 new industry projects in the second round of the HPC4Mfg program. Berkeley Lab was selected to partner with five of the new projects, several of which will use NERSC resources.

HPC4Mfg, established in March 2015, is designed to create an ecosystem that allows experts at DOE national labs to work directly with U.S. manufacturers to teach them how to adopt or advance their use of high performance computing to address challenging problems in manufacturing. The program gives participants access to world-class supercomputers and scientific expertise from Lawrence Livermore National Laboratory, which leads the program, as well as Lawrence Berkeley and Oak Ridge national laboratories.

Here are three of the HPC4Mfg projects announced in August that are using NERSC resources:

- **PPG Industries:** "Modeling Paint Behavior During Rotary Bell Atomization." PPG, one of the largest suppliers of automotive paint, will use supercomputers at NERSC to develop new coatings that can speed up production. The U.S. auto industry spray-applied over 60 million gallons of paint in 2014, much of it applied by an electrostatic rotary bell atomizer. Most bells can apply paint at 20 percent higher throughput than other methods, but atomization, which controls how fast the paint can be applied, suffers when the rate of fluid delivery is increased. PPG will model paint behavior during rotary bell atomization to come up with new kinds of paint that can be applied more quickly.

- **The American Chemical Society Green Chemistry Institute:** "Accelerating Industrial Application of Energy-Efficient Alternative Separations." Distillation in the chemical industry accounts for roughly 10 percent of energy use in the U.S. While methods such as porous mass separating agents (MSAs) could achieve the same process using much less energy, the manufacturing community needs to gain a better understanding of how MSAs influence the flow process before this method can be applied to industrial applications. Porous membranes appear capable of replacing distillation in many cases at a fraction of the energy use, but the changes in technology infrastructure required for widespread adoption are so profound that they are beyond the reach of any one company acting alone. Through the HPC4mfg collaboration with NERSC, this project will kick-start the fundamental understanding of membrane separations. In addition to providing the HPC resources and support, NERSC is designing and implementing the workflow tools to run the large volumes of simulations required to scan the parameter space of porous membranes.

- **Sepion Technologies:** "Improving the Manufacturability, Performance and Durability of Microporous Polymer Membrane Separators for Li–S Batteries using First Principles Computer Simulations." Energy systems on board aircraft are rapidly being electrified by the aviation industry. To meet industry targets for these systems, next-generation batteries with high energy density are essential. Efforts to commercialize lightweight, energy-dense lithium-sulfur (LiS) batteries have been stalled by problems with the battery's membrane, which limits battery lifetimes. Sepion will use NERSC resources to address key challenges in membrane manufacturing that could lead to longer lifetimes for LiS batteries.

NERSC's SBIR Partner Program Still Going Strong

NERSC has been working with industry partners through the DOE's Small Business Innovation Research (SBIR) grant program for many years. Since 2011, industry researchers have been awarded more than 90 million hours at NERSC via the SBIR program. In 2016 10 SBIR projects were active at NERSC, ranging from device design for free electron lasers to renewable energy modeling and the development of new HPC tools.

Among the SBIR users working with NERSC is Matrix Sensors, which seeks to develop gas sensors that are selective to CO_2 and methane that are far cheaper than existing solutions. Lowering the cost of such sensors is an important factor in the wider deployment in manufacturing processes and energy-efficient buildings. After conducting a scaling study, this company has grown its HPC ambitions and submitted new proposals under the HPC4Mfg program.

SBIRS COMPUTING AT NERSC IN 2016	COMPANY
Full Wave 3-D Modeling of RF Fields in Hot Magnetized Nonuniform Plasmas	FAR-TECH Inc.
Parallel Simulation of Electron Cooling Physics and Beam Transport	Tech-X Corp
Gyrotron Design and Evaluation using New Particle-in-Cell Capability	Tech-X Corp
Integrated Predictive Systems for Solar Energy with Modeling, Post Processing and Machine Learning	Vertum Partners
HPC Tool for Massively-Parallel Thermal Simulation of 3D and Dense Electronic Circuits	Capesym
Combinatorial Prediction of Gas Uptake in Metal-Organic Frameworks for Gas Sensor Applications	Matrix Sensors
Development of Vorcat for HPC Cloud-Based Complex Energy Applications	Vorcat, Inc.
Robust Optimization of Optics	NanoPrecision Products
A Comprehensive High Performance Predictive Tool for Fusion Liquid Metal Hydromagnetics	HyPerComp Inc
Charm++: An Object Oriented Parallel Programming System with Adaptive Runtime	Charmworks Inc

Director's Achievement Award Recognizes NERSC Staff

Members of NERSC's Wang Hall Energization Team were among several Berkeley Lab staff honored in November 2016 with a Director's Award for Exceptional Achievement. The award was given "For the safe and successful energization of Wang Hall and the Cori computer made possible by the close teamwork and integration of contributors from across NERSC, Facilities, EH&S, Protective Services and Procurement."

Here are the NERSC personnel who were on the Energization Team:

- Jeff Grounds
- Ernie Jew
- Brent Draney
- Thomas Davis
- Tina Declerck
- Elizabeth Bautista
- Jeff Broughton

The Director's Awards program at Berkeley Lab rewards significant achievements of Lab employees. Each year, these awards are given for accomplishments, leadership, collaboration, multi-disciplinary science, cross-divisional projects and commitment to excellence in support of the Lab's mission and strategic goals. Berkeley Lab employees submit nominations of individuals and teams, which are then reviewed by a Lab-wide Director's Awards committee.

NERSC 'Burst Buffer' Paper Wins Best Paper Honor

Wahid Bhimji (right), shown with CUG Program Chair Andrew Winfer, accepted the CUG Best Paper award on behalf of the NERSC team.

A paper outlining NERSC's Burst Buffer Early User Program and the center's pioneering efforts to test drive the technology using real science applications on the Cori supercomputer won the Best Paper award at the 2016 Cray User Group (CUG) meeting. Based on the Cray DataWarp I/O accelerator, the burst buffer on Cori is designed to move data in and out of the processor cores more quickly, which improves the overall performance of the system. This helps researchers make more effective use of the computing resource.

In August 2015, NERSC put out a Burst Buffer Early User Program call for proposals, asking NERSC's users to describe use cases and workflows that could benefit from accelerated I/O. NERSC received over 30 responses from the user community and ultimately chose to support 13 applications teams, plus give an additional 16 teams early access to the burst buffer hardware without dedicated support from a NERSC staff member.

The CUG paper describes the experiences of five use cases in terms of performance measurements and lessons learned during their first few months of working with the burst buffer on Cori. The use cases in the paper represent a broad range of science and applications:

- **Nyx/BoxLib:** cosmology simulation code
- **Chombo-Crunch + VisIt:** simulation and visualization of carbon sequestration processes
- **VPIC-IO:** simulation and analysis of plasma physics simulations
- **_TomoPy and SPOT:** real-time image reconstruction of Advanced Light Source and Advanced Photon Source data
- **ATLAS/Yoda:** simulation and data analysis for the LHC ATLAS detector

Personnel Transitions

Promotions

Richard Gerber: Richard Gerber, who joined NERSC in 1996 as a consultant, is now head of NERSC's High-Performance Computing (HPC) Department, formed in early 2016 to help the center's 7,000 users take full advantage of new supercomputing architectures and guide and support them during the ongoing transition to exascale. For the past year, Gerber served as acting head of the department, which comprises four groups: Advanced Technologies, Application Performance, Computational Systems and User Engagement.

Cory Snavely: Cory Snavely, a senior computer systems engineer who joined NERSC in March 2015, was named group lead for the Infrastructure Services Group (ISG). He had been acting lead of ISG since the group was created as part of a reorganization at NERSC in early 2016.

Damian Hazen: Damian Hazen, who has been with NERSC since 2001, was named group lead for the Storage Systems Group. Hazen had been acting lead since October 2015. Prior to that he worked primarily in the Storage Systems Group as an administrator and programmer for HPSS, but was also part of the Networking Group for two years.

New Hires

Brandon Cook: As a member of NERSC's Application Performance Group, Cook is providing HPC support for the facility's users, helping the NERSC Exascale Applications Program teams prepare for computing on next-generation architectures and contributing to the development of user web interfaces. Before coming to Berkeley Lab, Cook was a postdoc at Oak Ridge National Laboratory where he studied and applied scalable methods for electronic structure calculations.

Quincey Koziol: As a principal data architect in NERSC's Data & Analytics Services Group, Koziol is helping to lead the facility's data management efforts. Some of his day-to-day tasks include investigating object storage technologies, helping to define the storage sub-system for the NERSC-9 supercomputer and providing technical leadership on the HDF5 project. Koziol comes to NERSC with more than 25 years' experience in tackling data management challenges, most recently as the principal software architect for the HDF5 project.

Thorsten Kurth: As an HPC consultant who joined NERSC in February 2016, Kurth is working with the application readiness team to deliver optimized codes for Cori. He also acts as liaison for defining and demonstrating application portability between NERSC, ALCF and OLCF. Prior to NERSC he was a postdoc in the Nuclear Science Division at Berkeley Lab. He received his Ph.D. from the University of Wuppertal, Germany, in 2011.

Stephen Leak: As a consultant in NERSC's User Engagement Group, Leak is helping the facility's users effectively harness HPC resources to achieve their research goals. He has also been involved in getting key applications ready to run on Cori KNL. Before joining NERSC, Leak was part of the HPC team at New York University, where he provided user and application support.

Kirill Lozinskiy: As a senior HPC storage systems analyst, Lozinskiy is working on solutions for identifying and classifying archival data in the High Performance Storage Systems (HPSS). He is also working to integrate NERSC's Global Filesystems with HPSS, which will allow for greater flexibility in moving data between the systems, and providing administrative support for the archival storage infrastructure, NERSC Global Filesystems and data transfer nodes. Before coming to NERSC, Lozinskiy was a senior HPC storage systems administrator at the Broad Institute, a biomedical and genomic research center in Cambridge, Mass.

Mario Melara: As a computer systems engineer at NERSC, Melara is continuing his work on a new package manager called Spack, a package management tool designed to support multiple versions and configurations of software on a wide variety of platforms and environments. Before joining NERSC, Melara spent most of his career as a laboratory researcher studying dementia at UC Davis and UC San Francisco.

Abe Singer: As a security analyst in NERSC's Security and Networking Group, Singer is involved with security planning, implementation, monitoring and incident response. Currently his main project is overseeing NERSC's implementation of the DOE's Multi-factor Authentication mandate. Before coming to NERSC, Singer was the chief information security officer for LIGO, the Laser Interferometer Gravitational Wave Observatory, which was based at Caltech and partly funded by NSF.

Becci Totzke: As project coordinator for NERSC, Totzke is providing logistical support for the NERSC-9 and NERSC-8 projects, essentially ensuring that they meet the mission requirements of the DOE Office of Science. Before joining NERSC, she spent 13 years as the association manager for the American Peptide Society, a non-profit membership society of 1,200 peptide chemists, scientists and enthusiasts.

Tony Wildish: As an HPC consultant at NERSC, Wildish supports the Joint Genome Institute. Originally from England, Wildish has spent most of his career at CERN in Geneva, Switzerland. Shortly after obtaining his Ph.D. in high energy physics from Imperial College, London, Wildish went to CERN as a fellow in the ALEPH Online Group, subsequently moving on to the LHC's ATLAS and CMS experiments.

Post-docs

In 2016, NERSC hired eight post-docs to work on specific coding projects:

- Bill Arndt, MPAS
- Taylor Barnes, Quantum ESPRESSO
- Kevin Gott, PARSEC
- Tuomas Koskela, XGC1
- Mathieu Lobet, WARP (he has since assumed a permanent position at NERSC)
- Tareq Malas (he has since taken a permanent position at La Maison de la Simulation)
- Andrey Ovsyannikoc, Chombo-Crunch
- Zahra Ronaghi, Tomopy

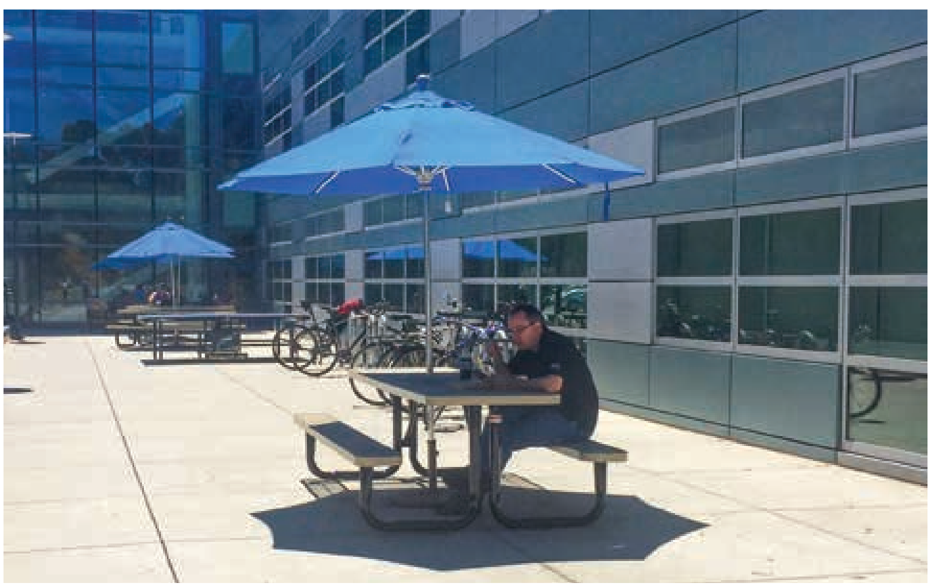

Allocation of NERSC Director's Reserve of Computer Time

Each year, 10% of NERSC time is available for the Director's Reserve. This time is used to support projects of strategic importance to the mission of NERSC and Berkeley Lab; examples in 2016 included ESnet/industrial partner research on high-speed network prototypes, a study of the migration of radioactive contaminants through subsurface media, an ARPA-E project on using solar-heated air to drive a turbine to generate electricity and X-Stack PI meeting demos. In 2016, NERSC allocated 144 million of its 300 million hours to Directors Reserve time and kept the remainder to improve throughput for general NERSC users when the Cori system was down for system integration.

Here is a list of the top 10 Director's Reserve projects for 2016 in terms of hours charged.

PI NAME	ORGANIZATION	PROJECT TITLE	SCIENCE CATEGORY	NERSC HOURS CHARGED
Monga, Inder	Berkeley Lab	Parallel Forward Error Correction Verification and Scaling Prototype for High Speed Optical Networks	Engineering	9,743,190
Prendergast, David	Berkeley Lab	Exploring Dynamics at Interfaces Relevant to Mg-ion Electrochemistry from First-principles (Molecular Foundry)	Materials Science	9,500,724
Steefel, Carl	Berkeley Lab	FP-RadRes: High Performance Flux Prediction Towards Radiological Resilience	Environmental Science	5,724,212
Pearlstein, Arne	U. Illinois	Anchored Solar-Driven Vortex for Power Generation	Engineering	5,181,259
Prabhat, Mr	Berkeley Lab	High Performance Data Analytics	Computer Science	4,844,788
Cruz Silva, Eduardo	Berkeley Lab	Next Generation Semiconductors for Low Power Electronic Devices	Materials Science	4,558,570
Bhatele, Abhinav	Livermore Lab	Performance Analysis, Modeling and Scaling of HPC Applications and Tools	Computer Science	3,266,485
Fryman, Joshua	Intel Inc.	Intel Exascale R&D Pathfinding Architecture Studies (FFWD-2 and XStack)	Computer Science	2,221,804
Kent, Paul R.	Oak Ridge	NESAP: Extending the Capabilities of Quantum Espresso for Cori	Materials Science	2,000,960
Michalakes, John G.	NOAA	Next Generation Global Prediction System (NGGPS) Benchmarking	Climate Research	1,998,665

Publications

In 2016, NERSC's nearly 7,000 users reported more than 2,200 peer-reviewed published papers that involved NERSC resources. In addition, NERSC staff contributed more than 150 papers to scientific journals and conferences, showcasing the center's continued involvement in HPC hardware and software development and expanding expertise in data-intensive computing, data storage and analysis, scientific workflows, code optimization, machine learning and quantum computing.

To see a comprehensive list of publications and presentations by NERSC staff in 2016, go to bit.ly/NERSCpublications.

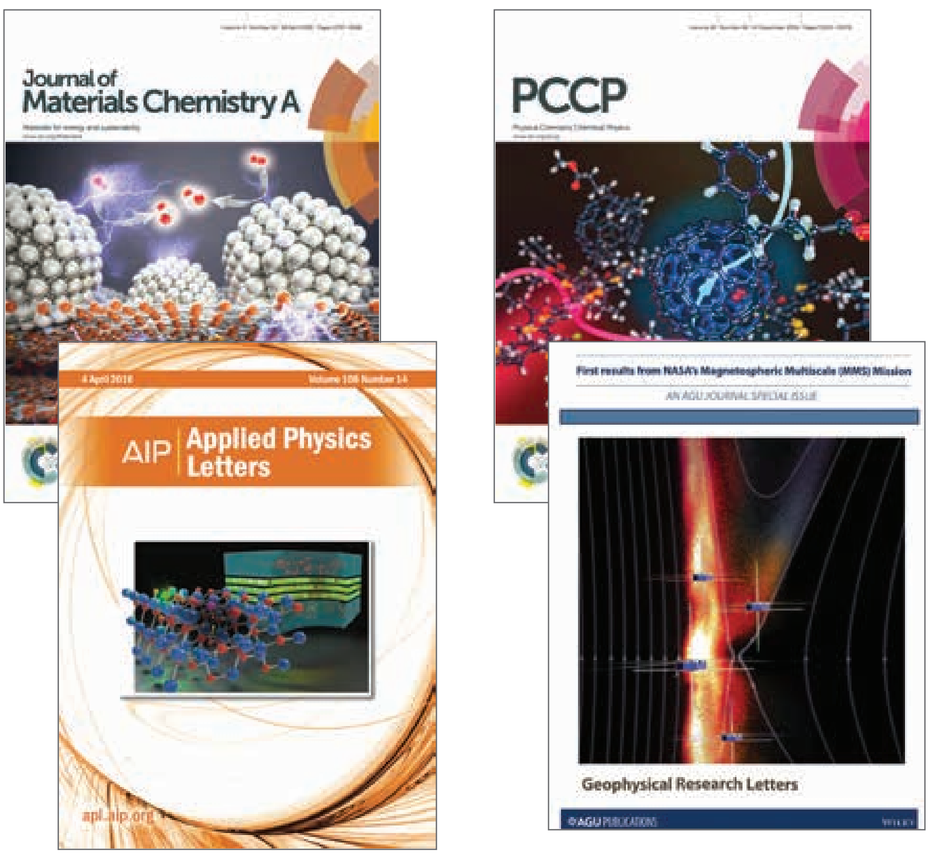

Appendix A

NERSC Users Group Executive Committee

OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH

Milan Curcic, *University of Miami*
Jeff Hammond, *Intel*
Brian Van Straalen, *Lawrence Berkeley National Laboratory*

OFFICE OF BASIC ENERGY SCIENCES

Eric Bylaska, *Pacific Northwest National Laboratory*
Monojoy Goswami, *Oak Ridge National Laboratory*
Paul Kent, *Oak Ridge National Laboratory*

OFFICE OF BIOLOGICAL AND ENVIRONMENTAL RESEARCH

Samson Hagos, *Pacific Northwest National Laboratory*
Gary Strand, *National Center for Atmospheric Research*
David Wong, *U.S. Environmental Protection Agency*

OFFICE OF FUSION ENERGY PHYSICS

Christopher Holland, *University of California, San Diego*
Nathan Howard, *Massachusetts Institute of Technology*
David Green, *Oak Ridge National Laboratory*

OFFICE OF HIGH ENERGY PHYSICS

Ted Kisner, *Lawrence Berkeley National Laboratory*
Zarija Lukic, *Lawrence Berkeley National Laboratory*
Frank Tsung, *University of California, Los Angeles (Chair)*

OFFICE OF NUCLEAR PHYSICS

Balint Joo, *Jefferson Lab*
Nicolas Schunck, *Lawrence Livermore National Laboratory*
Michael Zingale, *Stony Brook University*

MEMBERS AT LARGE

James Amundson, *Fermilab*
Carlo Benedetti, *Lawrence Berkeley National Laboratory*
David Hatch, *University of Texas, Austin*

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, *Contractor*

FACILITIES DIVISION

Christine Chalk, *Director (Acting)*

Ben Brown, *Physical Scientist, ESnet Program Manager*

Betsy Riley, *Computer Scientist, ALCC Program Manager*

Carolyn Lauzon, *Physical Scientist, ALCC Program Manager*

Sonia Sachs, *Computer Scientist, ALCF*

Robinson Pino, *Computer Scientist, REP Program Manager*

Dave Goodwin, *Physical Scientist, NERSC Program Manager*

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

Sally McPherson, *Program Assistant*

RESEARCH DIVISION

Steven Lee, *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*

Abani Patra, *Mathematician, Multiscale Mathematics*

Steven Lee, *Physical Scientist, Base, Math: Algorithms, Models, Data; SCiDAC – Math Centers & Institutes (Acting)*

Lucy Nowell, *Computer Scientist, Computer Science*

Angie Thevenot, *Program Assistant*

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

BER
Office of Biological and Environmental Research

BES
Office of Basic Energy Sciences

BNL
Brookhaven National Laboratory

CCM
Cluster Compatibility Mode

CERN
European Organization for Nuclear Research

CESM
Community Earth Systems Model

CFD
Computational Fluid Dynamics

CMB
Cosmic Microwave Background

CO_2
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

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

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

FLOPS
Floating Point Operations

FTP
File Transfer Protocol

GB
Gigabytes

Gbps
Gigabits Per Second

GPU
Graphics Processing Unit

GUI
Graphical User Interface

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

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

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

OpenMP
Open Multi-Processing

OpenMSI
Open Mass Spectrometry Imaging

OSF
Oakland Scientific Facility

PDACS
Portal for Data Analysis services for Cosmological Simulations

PDSF
Parallel Distributed Systems Facility, NERSC

PI
Principal Investigator

PIC
Particle-In-Cell Simulations

PB
Petabytes

PSII
Photosystem II

PNNL
Pacific Northwest National Laboratory

PPPL
Princeton Plasma Physics Laboratory

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

TACC
Texas Advanced Computing Center

TB
Terabytes

URL
Universal Resource Locator

VASP
Vienna Ab initio Simulation Package

WAN
Wide Area Network

FOR MORE INFORMATION ABOUT NERSC, CONTACT:

Jon Bashor
NERSC Communications
1 Cyclotron Road
Berkeley Lab, MS 59R3022B
Berkeley, CA 94720-8148
Email jbashor@lbl.gov
Phone (510) 486-5849
Fax (510) 486-4300

NERSC's Web Site
nersc.gov

NERSC Annual Report Editor
Kathy Kincade

Contributing Writers and Editors
Jon Bashor, Kathy Kincade, Linda Vu and Margie Wylie
Berkeley Lab Computing Sciences Communications Group

Katie Antypas, Elizabeth Bautista, Deborah Bard, James Botts, Shane Canon, Jim Craw, Jack Deslippe, Richard Gerber, Lisa Gerhardt, Rebecca Hartman-Baker, Damian Hazen, Doug Jacobsen, Craig Lant, Betzy MacGowan, Jackie Scoggins, David Skinner, Cory Snavely, Becci Totzke
NERSC

Design
Design, layout, illustraton, photography and printing coordination
Berkeley Lab Creative Services

17-NERSC-3962