ESnet
ENERGY SCIENCES NETWORK

Accelerating World-Changing Research Collaborations

2022 ANNUAL REPORT

U.S. DEPARTMENT OF ENERGY | Office of Science

Table of Contents

About ESnet

The Energy Sciences Network (ESnet) is the high-performance network user facility of the Department of Energy's Office of Science, delivering highly reliable data transport capabilities and services optimized for the requirements of large-scale integrated scientific research. ESnet is stewarded by the Advanced Scientific Computing Research Program (ASCR) and managed and operated by the Scientific Networking Division at Lawrence Berkeley National Laboratory (Berkeley Lab). Its vision is that scientific progress will be completely unconstrained by the physical location of instruments, people, computational resources, or data. By offering highly specialized, continuously evolving networking systems and innovative services, ESnet provides the foundation for an integrated research infrastructure that enables the nation's scientists to collaborate on some of the world's most important scientific challenges.

Director’s Note

What a year 2022 was for ESnet!

Since 1986, the eponymous network operated by ESnet and headquartered at Lawrence Berkeley National Laboratory has served as the data circulatory system for the Department of Energy, linking all of the DOE's national laboratories, premier scientific instruments, and supercomputing centers, and tens of thousands of DOE-funded researchers and their collaborators. This intricately connected system enables data to move quickly and reliably between sites and scientists, accelerating time to discovery.

Last October, we unveiled ESnet6, the sixth iteration of this vital network. Imagine transplanting an entire cardiovascular system while the patient is awake…and sprinting. I am incredibly proud of the ESnet team for managing to do so smoothly, with minimal disruption for our users, two years ahead of schedule, and under budget.

ESnet6 is the result of significant leaps forward in not only data transfer speeds but also orchestration, automation, measurement, and monitoring. To extend the circulatory system metaphor, neither blood nor data just flow magically to the right places at the right volume and speed by themselves. Sometimes the body needs to orchestrate competing demands for blood from different tissues at different times. ESnet's software, workflows, and services function like the sympathetic and parasympathetic nervous systems and the master controllers of the cardiovascular system in the brain. Our portfolio is designed to meet the ever-increasingly complex needs of modern scientific research, which require dynamic integration of experiment, observation, theory, modeling, simulation, visualization, machine learning (ML), artificial intelligence (AI), and analysis — all of which generate data rapidly approaching the exabyte scale.

Building the Foundation for the DOE's Integrated Research Infrastructure

The vision of the DOE Office of Science's Integrated Research Infrastructure (IRI) initiative, to which ESnet has been a contributor since its inception in 2019, is closely aligned with ESnet's mission: "To empower researchers to meld DOE's world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation."

ESnet6 is the first large-scale infrastructure project designed specifically with IRI in mind. And in 2022, ESnet staff continued to work with the IRI Task Force, participating in IRI Architecture Blueprint Activity discussion groups, requirements gathering, report writing and preparation, and other supporting activities.

These four strategic priorities will guide ESnet's future efforts:

1. **Transforming Operations:** While priority one is to operate a highly performant and robust network, we are actively exploring new architectures, infrastructure enhancements, improvements to business processes, additional orchestration and automation capabilities, and ways to integrate new technologies such as AI and ML — all to improve the resiliency, efficiency, and effectiveness of the user facility.
2. **Expanding Services Portfolio:** Our staff continues to innovate, experiment, prototype, and transition to production new data and network services. In addition, we actively look to expand the modalities through which scientists acquire data, from private 5G to low-Earth-orbiting satellites in remote locations, and potentially through quantum networks.
3. **Increasing Stakeholder Value:** As high-speed, big-data networking experts, we can co-design solutions based on upcoming requirements with our scientific and site user community to ensure that ESnet provides the most value to all of DOE as well as to Research & Education (R&E) community stakeholders worldwide.
4. **Building Accountability and Transparency:** We will continue to foster a culture of accountability and transparency in which our users — and our employees — can perform at their personal best.

Perhaps ESnet's biggest challenge in the coming years will be to increase outreach efforts, to make our vast collection of users aware that this dedicated scientific network — on whose speed and reliability they have long depended unknowingly — also offers tools and services that can optimize their scientific data acquisition, transport, placement, and sharing, thus supercharging their research workflows.

Spurring Collaboration through Cross-Functional Conversations

In 2022, we began tackling that challenge by hosting Confab22, our first-ever gathering for our science users, held in conjunction with the unveiling of ESnet6 in October in Berkeley, California. Confab launched the cross-functional conversations that we hope will lead to many more collaborations in which ESnet can play an active role. We're eager to continue those discussions at Confab23 in October 2023, focusing primarily on the IRI initiative.

This annual report, also ESnet's first, was conceived as fuel for that conversational fire. The following pages highlight some of the tools and services representing progress in operational innovations (page 16) and applied research (page 22) that we are particularly excited about. The projects starting on page 30 showcase how deliberate integration of experimental, networking, and computational facilities can accelerate insight from experiments that rely on the massive datasets generated by large-scale instruments such as genome sequencers, telescope observatories, X-ray light sources, and particle accelerators.

Strengthening Our Human Network

ESnet's other challenge is to continue to manage our growth, with more than 120 employees in 23 states working in many different roles. Our ability to contribute meaningfully to scientific progress and to humanity requires a diversity of backgrounds, skills, and geography. Ultimately, ESnet's success as an organization — and as a network — depends on our reliability and resiliency as well as our speed and ability to innovate. See page 46 for more details on these efforts.

As rewarding as 2022 was for us, I look forward to what our team will achieve over the next few years. It is truly an exciting time to be part of the global science community.

— Inder Monga
Executive Director, Energy Sciences Network (ESnet)

ESnet By The Numbers

99.999%
Site Uptime*

4.86
Overall User Satisfaction Rating (out of 5)

ESNET COMPRISES...

15,000 MILES
of dedicated fiber-optic cable **

75
Router sites**

273
Amplifier sites**

4
100 Gbps trans-Atlantic circuits***

THE TEAM

126
Staff and contractors****

14%
Increase in headcount from CY2021

23
States

* ESnet is measured and expected to provide 99.9% site uptime. The network regularly exceeds that standard, providing near 100% uptime to almost all connected sites.

** In U.S.

*** Additional three 400G circuits coming online in FY23

**** As of Dec. 31, 2022

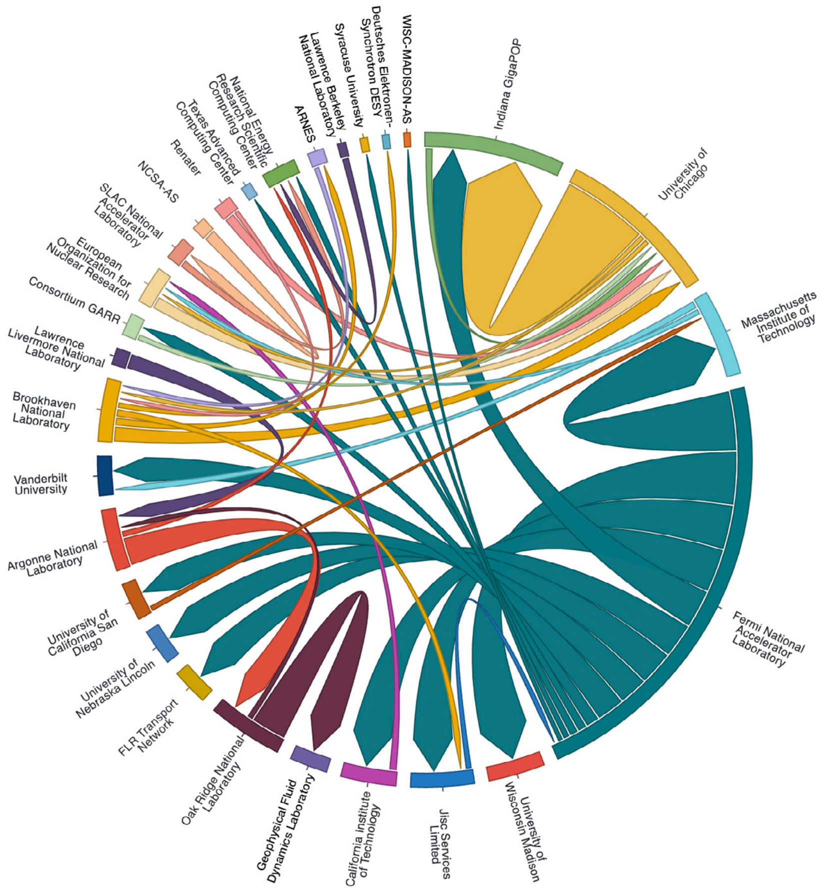

High-Capacity Users This chord diagram depicts the volume of data* moved across ESnet in calendar 2022 between the top 40 pairs of source and destination organizations.

* Of Fermi National Accelerator Laboratory's 199 petabytes sent, 128 PB were associated with the Large Hadron Collider project. Additionally, the 73.5 PB of data moved from University of Chicago to Indiana GigaPOP was entirely attributed to the LHC.

ESNET CONNECTS:

17
National Laboratories

28
DOE User Facilities (including ESnet)

271
R&E, Commercial, and Other Networks (in 5 Countries Including U.S.)

CAPABILITIES:

1.3 EXABYTES

Transited in ESnet in CY2022

46 TBPS

Aggregate Capacity

400 GBPS to 1 TBPS

Network Backbone

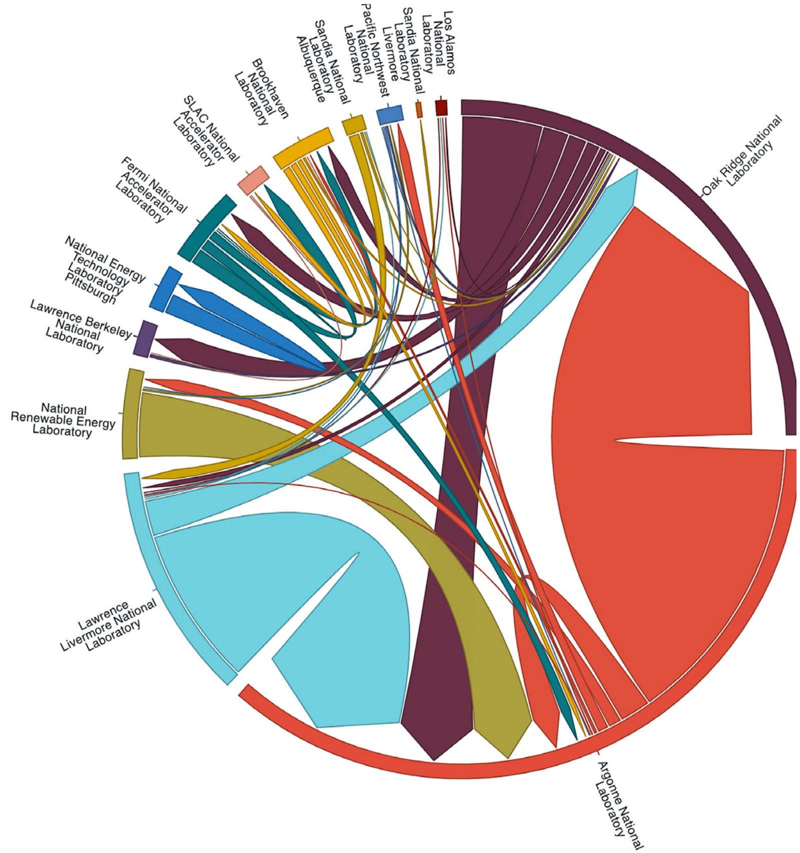

Usage by National Laboratories This diagram depicts the volume of data* moved across ESnet in calendar 2022 between the top 40 pairs of national laboratories.

* Argonne National Laboratory was the largest producer and consumer of inter-lab data, with a combined 34 petabytes of data movement to and from other labs. Of this, 14 PB of data was sent to Oak Ridge National Laboratory, and 7.65 PB was received from Lawrence Livermore National Laboratory as part of an ESGF Data Replication Effort (see page 37).

Evolution of ESnet

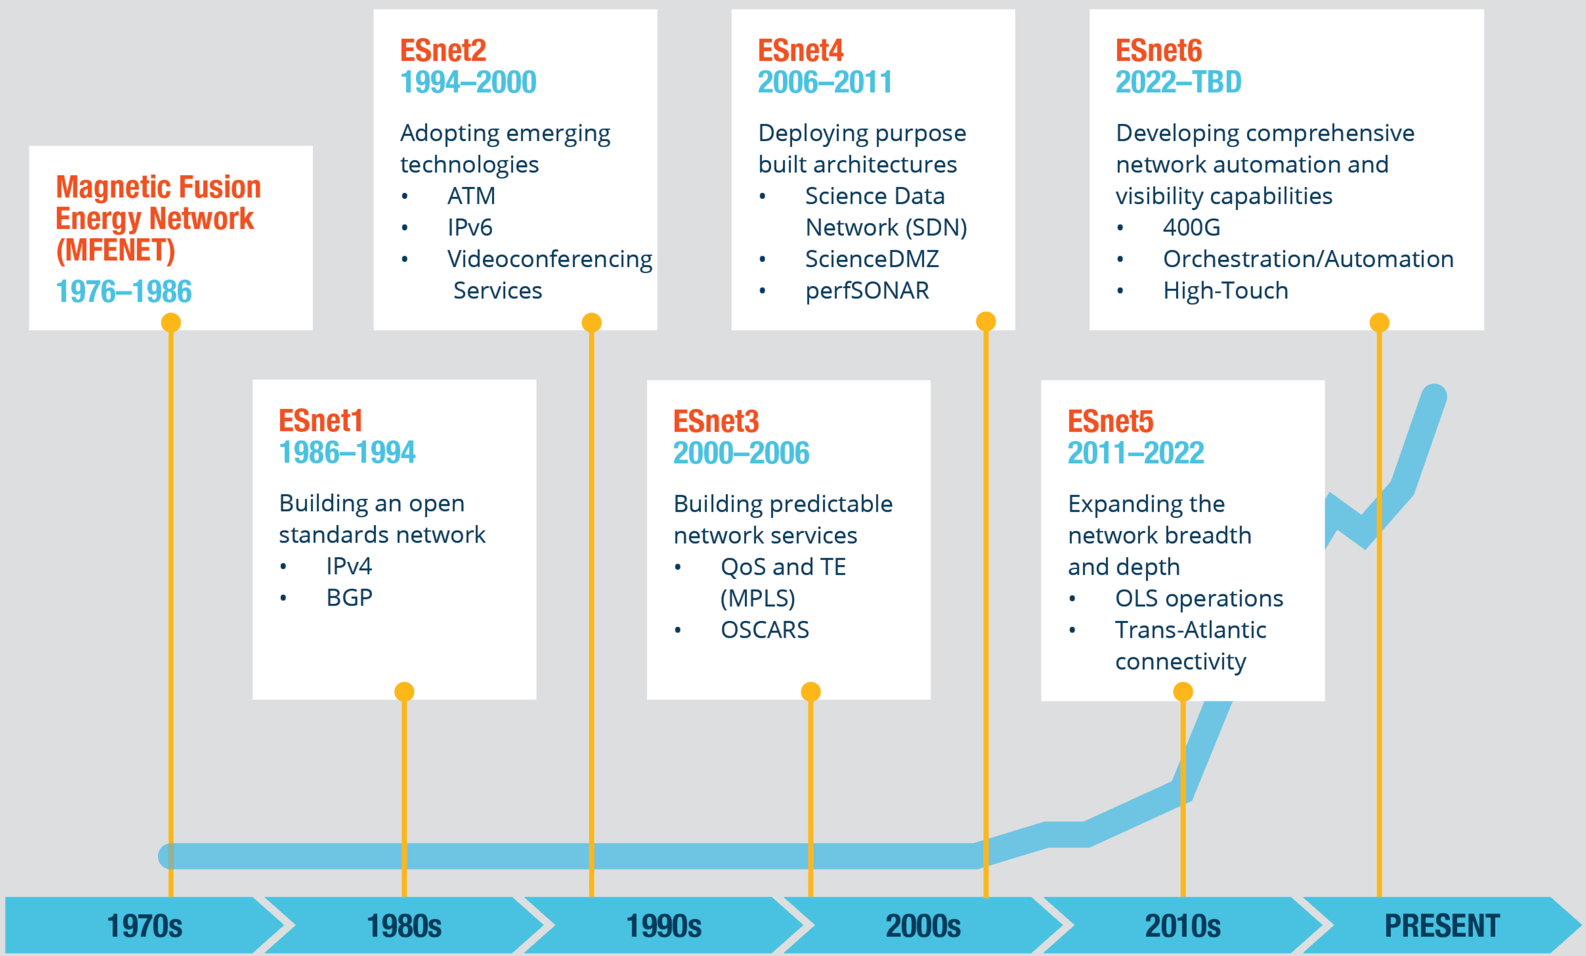

ESnet Traffic: Historic Usage

ESnet's network traffic has grown exponentially in the past two decades. In 2022, ESnet carried more than 1.3 exabytes of data to and from the sites it connects. An exabyte is a nearly incomprehensible amount of data, equal to 1,000 petabytes or one billion gigabytes. This torrent of data is spurring ESnet to innovate in multiple areas to continue to deliver on its mission — and enable DOE scientists to seamlessly move, store, and analyze that data without drowning in it.

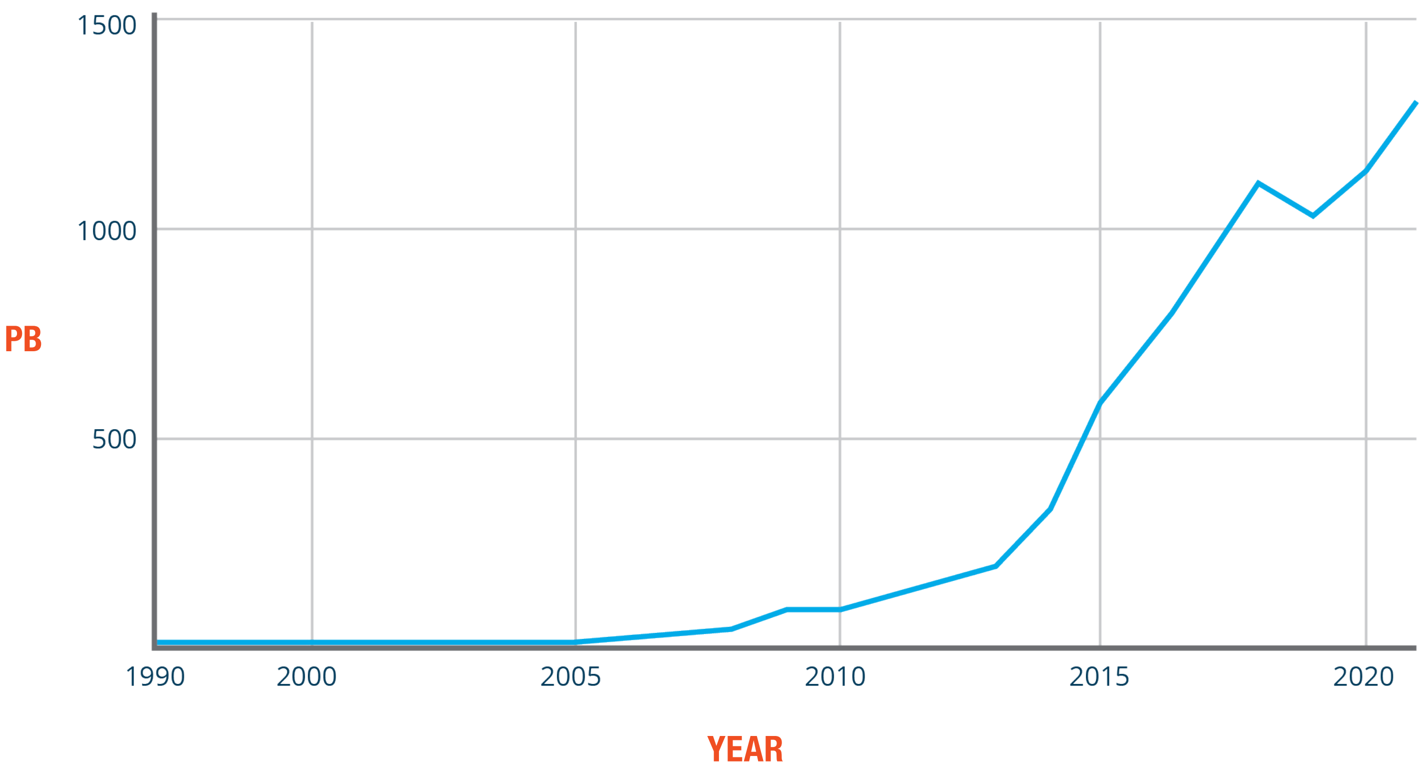

ESnet6: Built for Tomorrow's Scientific Collaborations

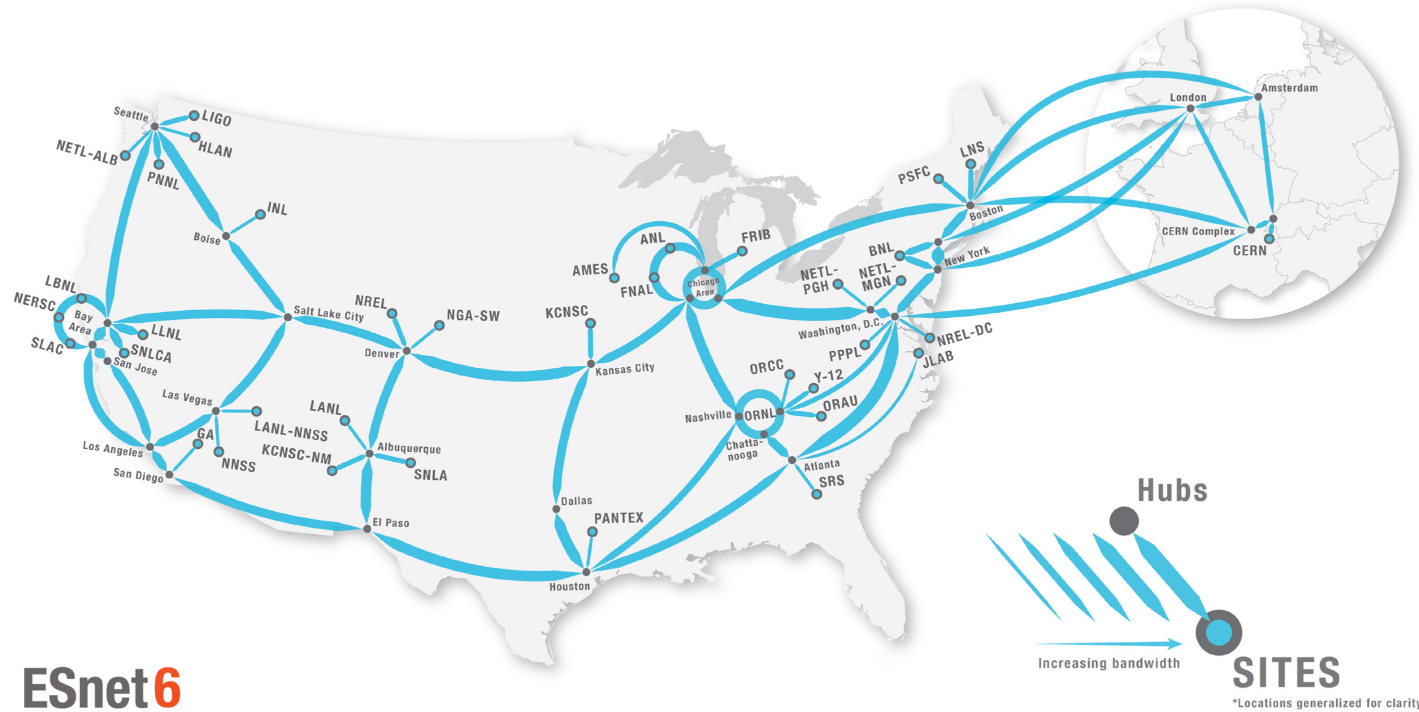

Begun in 2017, ESnet6 was completed in 2022 — two years ahead of schedule and under budget. This was a significant feat, especially considering the hurdles posed by the global pandemic, disruptions in the supply chain, and the task of replacing a network while it was actively in use — not unlike taking apart and re-assembling an airplane full of passengers… in mid-flight.

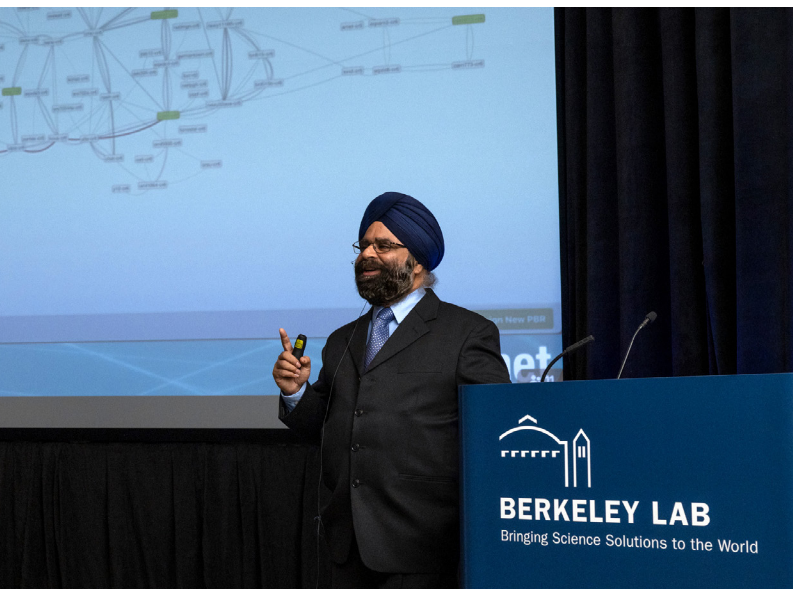

ESnet Executive Director Inder Monga at the ESnet6 unveiling.

ESnet6 was specifically designed to help the DOE research community and its international collaborators harness the deluge of data created by scientific research in the exascale and artificial intelligence era. The project had three simple yet ambitious goals. The first was to build a network designed to handle that exponential growth in data cost-effectively. In 2021, ESnet5 carried more than 1.1 exabytes of science data; by 2022, the total exceeded 1.3 exabytes. Traffic on ESnet has increased nearly by a factor of ten every 5.5 years. Second, to invest in resiliency and reliability that would ensure real-time, seamless performance of data analysis and enable new modalities, such as streaming terabits of data directly from instruments to supercomputers thousands of miles apart. Third, to build flexibility and a foundation strong enough to support new breakthroughs in machine learning, artificial intelligence, quantum information science, and the revolutions yet to come.

ESnet6 represents a transformational change in the way networks are built for research, with improved capacity, resiliency, and flexibility. With more than 46 Terabits per second of aggregate bandwidth deployed, it features a significant increase over prior generations of the network. This capacity boost enables scientists to more quickly move, process, analyze, visualize, share, and store the mountains of research data produced by experiments, modeling, and simulations.

But the new network does more than just increase capacity. For the first time, ESnet built and managed the entire network stack built over a dedicated, nationwide 15,000-mile dark-fiber footprint. With ESnet6, our engineers have developed smart, programmable, and automated services uniquely built to support the multi-petabyte dataflows typical of scientific research today. In addition, they are future-proofed to manage the emerging exabyte data era, streaming data from instruments and high-impact “digital twins” that require predictability and low latency. New capabilities and services include customizable network services via a new automation platform, high-precision telemetry to improve network performance, improved overall network security, and a future programmable API platform for scientists to directly request custom network services. (Many of these innovations are discussed further in this annual report.)

ESnet6 was officially unveiled at a special event attended by DOE and Berkeley Lab leadership, as well as local, state, and federal government representatives, including U.S. Congresswoman Barbara Lee and Asmeret Asefaw Berhe, director of the DOE’s Office of Science. The event also featured keynote talks from Internet pioneer Vint Cerf, vice president and chief internet evangelist for Google and ESnet Policy Board member; and Ian Foster, distinguished fellow at Argonne National Laboratory and professor of computer science at the University of Chicago. Each provided insights into the impact ESnet6 is having on the scientific community and the global Internet.

Additional speakers included Berkeley lab scientists Ann Almgren, senior scientist and the Applied Mathematics Department head, and David McCallen, senior scientist in the Energy Geosciences Division and professor at the University of Nevada, Reno. These researchers shared background on how ESnet impacts several of their exascale computing research projects, including wind energy and earthquake simulations.

Internet pioneer Vint Cerf delivered a keynote.

The event also included a live demonstration of ESnet6's new automation platform, which rapidly configured ESnet6 network paths to support the transfer of large science datasets across the country in under two minutes. In addition, a live, multi-terabyte transfer of earthquake simulation data was also performed between the Oak Ridge Leadership Computing Facility to the National Energy Research Scientific Computing Center in minutes.

Close collaborations with Lumen Technologies, Ciena, Infinera, Nokia, and Xilinx (AMD) provided ESnet with leading-edge equipment and expertise to help bring the new network into service. These partners provided key building blocks, including ESnet's nationwide fiber optic cable footprint, optical switching and routing platforms, and an extreme-scale packet monitoring system.

ESnet's improved capacity, resiliency, and flexibility make it faster, easier, and more efficient for scientists around the world to conduct and collaborate on ground-breaking research. As Vint Cerf said at the unveiling, "The Internet is about information and data flow. With the launch of ESnet6, we are empowering the scientific community with unprecedented new capabilities to help advance the search for life-saving and planet-saving discoveries."

The ESnet6 project was finished more than two years ahead of schedule and under budget. In recognition of this unusual feat, the Department of Energy (DOE) presented ESnet with a special Project Assessment Award. Office of Project Assessment Director Kurt W. Fisher presented it to ESnet Network Services Group Lead Kate Petersen Mace, who as the project's director accepted the award on behalf of the ESnet6 team.

Operational Innovations

ESnet's network is operated as a high-availability network 24 hours per day, 365 days per year, and significant resources are devoted to ensuring its secure, continuous availability. While ESnet is required to deliver "three 9s" — 99.9% availability — ESnet routinely exceeds that standard. In fact, sites with redundant connections to ESnet often see 100% availability (zero downtime) over the course of a year.

Our staff continues to work on enhancing the technical systems that benefit both our internal team and external production services, in relentless pursuit of performance and reliability. These improvements encompass tools for gathering, assessing, and studying network performance data, as well as automating the setup of devices on the production network. In 2022, ESnet's operational innovations included (but were not limited to) the following projects.

Improving Service-Oriented Notifications

In CY2022, ESnet continued to improve its overall user satisfaction ratings, rising to 4.86 in CY2022 from 4.79 (out of 5) in CY2021. Captured via a survey of ESnet Site Coordinator Committee (ESCC) members that was conducted anonymously by the ESCC Chair, this rating has continually improved for several years, attributable to a focus on operational improvements.

This year, the ESnet Network Operations Center (NOC) upgraded the communications workflow to better prioritize outage response and communications based on impact and urgency. Additionally, for most of our IP and OSCARS services (OSCARS is an open-source, ESnet-developed advanced software system for reserving network resources), ESnet began sending targeted notifications to our site users that allow them to determine whether a planned outage or unplanned impairment will affect their services. ESnet also built new capabilities for visibility into site uptime. Next-gen operations will continue to be a key ESnet initiative in 2023, transitioning to service-based notifications, reducing the number of notifications sent to individual sites, and making the information in the notifications easier to understand.

Computer Systems Engineers Charles Shiflett and Rémy Doucet in ESnet’s Berkeley Lab data center.

Open-Source Visualization Tools

Stardust is a time-series data collection, analysis, and visualization platform developed by ESnet and drawing on the expertise gained from the National Science Foundation–funded NetSage project. The platform provides a scalable means to collect and process multiple network measurements, allowing users to create custom dashboards and visuals. It gives ESnet technical staff improved visibility into the performance and behavior of network services while fostering greater collaboration within the R&E community.

In 2022, the Stardust analytics team developed and made public several open-source tools for both data visualization and processing. The most notable were four visualization plug-ins submitted to and accepted by the Grafana platform: a bump chart, slope graph, chord diagram, and matrix. In addition to visualization plug-ins, ESnet also developed an open-source tool for managing Grafana dashboards named Grafana Dash-N-Grab (GDG), available on GitHub.

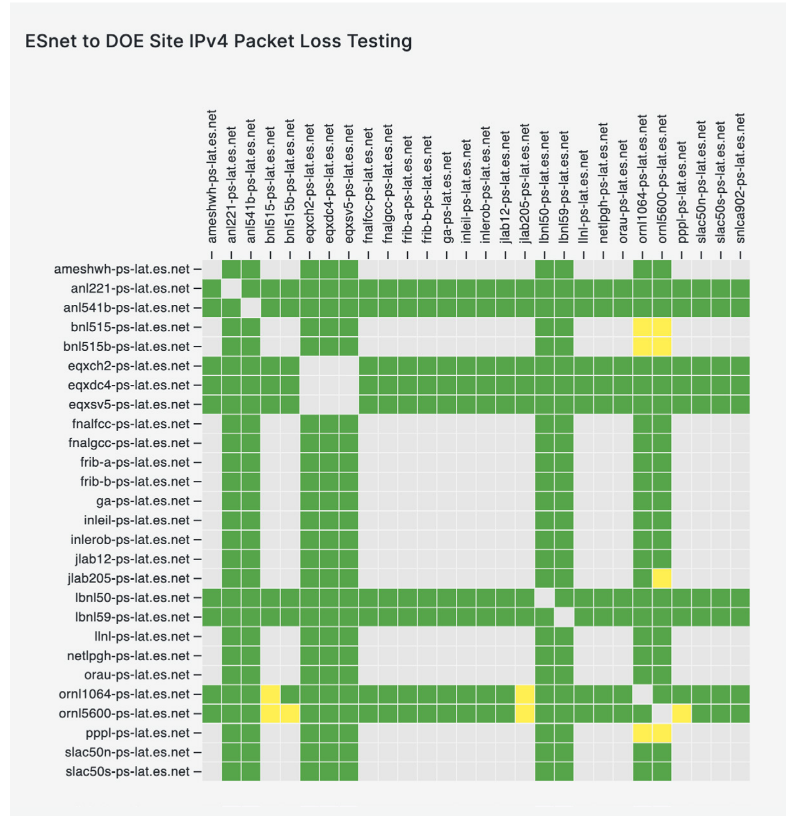

ESnet's new matrix plugin can compare two sets of categorical data, whereas existing Grafana plugins all required timeseries data on both axes. This allows, for example, packet loss between hosts to be visualized as shown.

Upgrading perfSONAR

perfSONAR is a platform for end-to-end network performance measurement and monitoring, used by network performance engineers to help pinpoint and diagnose subtle network performance issues. ESnet is a key member of the highly successful perfSONAR collaboration, along with Internet2, Indiana University, GÉANT, The University of Michigan, and Rede Nacional de Ensino e Pesquisa (RNP). perfSONAR is influential in the Research and Education Network (NREN) community and is the de facto standard platform for network testing and measurement in science networking.

More than 2,000 registered perfSONAR hosts have been deployed on 400 networks in more than 50 countries. For example, the Worldwide Large Hadron Collider (LHC) Computing Grid (WLCG) project has intensive data transmission requirements; a perfSONAR “mesh” has been built between all LHC sites and is monitored with a unified dashboard. In 2022, ESnet performed a hardware upgrade of the 50 hosts it maintains, to enable testing of network speeds up to 100 Gbps, and shared the experience gained in tuning this new hardware with the R&E community.

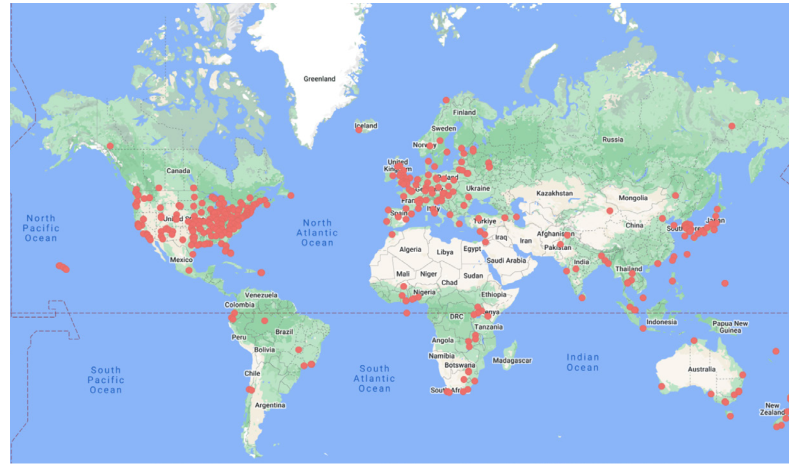

perfSONAR hosts worldwide.

Network Automation Innovations

One of the significant improvements represented by ESnet6 is its service orchestration and automation system. A key driver of this is the open-source Workflow Orchestrator tool, first developed in 2019 by the Netherlands R&E consortium SURF, which helps network administrators both automate (execute repetitive tasks reliably and easily) and orchestrate (add a layer of intelligence to tasks being automated and a complete audit log of changes). ESnet has contributed significantly to this open-source project and partnered with SURF to expand its adoption in the R&E community.

In this architecture, the Orchestrator functions as a quarterback, coordinating calls to various systems such as ESDB, NameSurfer (IPAM), Cisco NSO, and Ansible to implement configuration changes on network elements. In 2022, ESnet added several new workflows that automated the deployment of ESnet6 and migrated existing services to new network devices. In particular, the Internal Host Connectivity (IHC, used for managing direct connections to ESnet from servers) service workflows gained extra functionality. The team implemented the ability to configure BERT, Gateway, perfSONAR, and High Touch hosts. The Infrastructure and Networking teams extensively used the High Touch IHC workflows to push configurations from ESDB to newly deployed hosts. (For more about High Touch, see page 23.)

ESnet also made many operational enhancements to the development process of the Orchestrator application: expansion of the automated test suite improved the quality of the code, and Jira automation increased the team's velocity and ability to monitor workload.

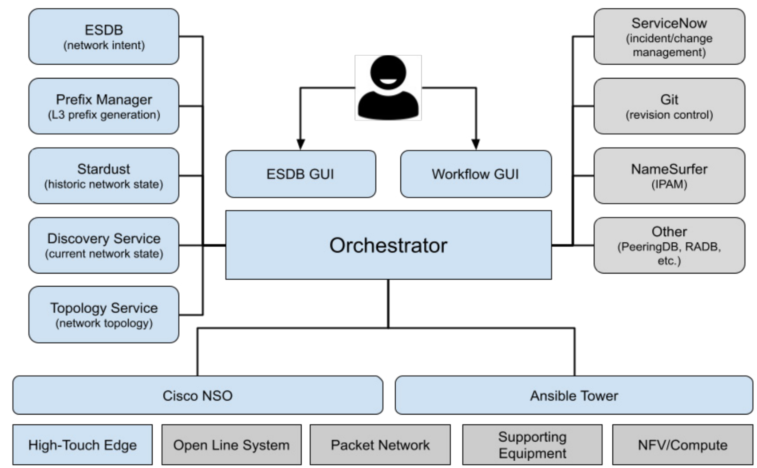

The network service orchestration architecture.

Supporting IPv6 Adoption and Standardization

The new Internet protocol IPv6 was first introduced in 1998 to address the shortfall of unique IP addresses available under IPv4. Despite the improvements in efficiency and security that come with IPv6, adoption in the United States and around the world has been slow.

As the DOE laboratories' scientific network, ESnet has naturally co-led both the DOE IPv6-only Implementation Team and the community of practice groups, providing forward momentum and input on migration off of IPv4 and developing transition mechanisms for that migration path. With the team, ESnet staff organized multiple community practice sessions and worked with many entities, both within the DOE and in the larger federal government space. ESnet team members presented at and participated in the larger USG Federal IPv6 Task Force, providing important input and updates on the recently published NSA IPv6 security guidelines. ESnet staff also authored three active Internet Engineering Task Force drafts, offering information about extensive testing, lab work, and operation deployments relevant to the success of this effort.

Strengthening the Security Landscape

ESnet has built a strong security program with a rich history of expertise in network security monitoring. It is one of the key partners supporting the Zeek open-source project, and in response to the White House's 2022 cybersecurity strategy, ESnet has launched its own Zero Trust program and developed a whitepaper that sets its security strategy for years to come. The security team was also essential in the successful completion of the ESnet6 project: it developed ESnet's first security service as one of the key performance parameters, allowing ESnet to block and isolate traffic quickly on the WAN for customers. The team also released the open-source SCRAM (Security Capture and Release Automation Manager) tool, enabling our partner sites to better protect their networks.

Applied Research Activities

ESnet's research and development activities are integral to our mission of advancing the capabilities of today's networking technologies to better serve the science requirements of the future. We continually investigate, create, develop, and test the services, protocols, routing techniques, and tools necessary to meet the expanding needs of our user community of DOE scientists.

Using Telemetry to Drive High-Touch Services

Being able to identify the kinds of data flowing across the network, and to and from whom, will allow ESnet to optimize future workflows and services. ESnet's High-Touch project provides new data and insights through the use of programmable hardware (Field Programmable Gate Arrays, or "FPGA") and software. The first class of these services offers network telemetry at both the flow and packet levels at multi-100GE speeds, improving visibility into traffic for network engineering, cybersecurity, and networking research. In 2022, ESnet demonstrated a "code complete" system implementation for providing unsampled flow data using a prototype server configuration installed in two ESnet locations. The demonstration showed the improved insights made possible by examining unsampled packet data instead of the sampled or aggregated data that is the norm for technologies commonly available today. Some initial analysis of data gathered from the prototype deployment has led to performance enhancements and improvements in the fidelity of the captured data.

This work, combined with engagement with internal ESnet customers (particularly in network engineering and cybersecurity), will drive future applications built with the base data collection capabilities and data analysis. Deployment of production High-Touch servers with FPGA hardware into the ESnet6 production network began in late CY2022. When completed, two High-Touch servers will be installed at more than 40 ESnet locations, allowing visibility of all traffic entering or exiting the ESnet backbone.

Applying Good SENSE to Workflows

SENSE (SDN for End-to-End Networking @ Exascale) is a multi-network ESnet orchestration and intelligence system providing programmatic-driven networked services to domain science workflows. These services can span multiple domains/sites and be presented to different workflows in a highly customized manner. The SENSE services are referred to as "networked services" because, in addition to the network elements, SENSE can orchestrate compute and storage elements that connect to the network, such as DTNs, Cloud compute, and site-managed network resources.

A key objective of the SENSE system is to allow the workflows to access the network in a manner that best facilitates their objectives. A prototype system was deployed in 2022 at UC San Diego and Caltech, where Large Hadron Collider (LHC) Tier 2 facilities are hosted. In addition, SENSE is being integrated into LHC Compact Muon Solenoid (CMS) workflows that use the Rucio/FTS/XRootD data management and movement system. Although the CMS compute system management provides deterministic resources for individual users, the data transfers and associated network services are unpredictable. By integrating SENSE services, Rucio can identify the dataflow groups that should have a higher priority, which influences network provisioning. ESnet's presentation at the SC22 conference demonstrated this interoperation with SENSE; testing/development is ongoing as part of the CMS experiment support.

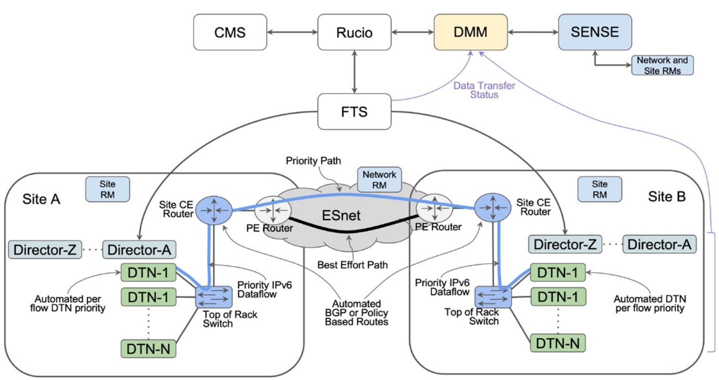

SENSE Integration with Rucio/FTS/XRootD Based Data Transfer Systems.

Sewing FABRIC Together: A National-Scale Programmable Experimental Network Infrastructure

Funded by the National Science Foundation's (NSF's) Mid-Scale Research Infrastructure program, FABRIC is a novel adaptive programmable national research network testbed that allows computer science and networking researchers to develop and test innovative architectures that could yield a faster, more secure Internet. (The acronym FABRIC is Adaptive ProgrammaBle Research Infrastructure for Computer Science and Science Applications.)

FABRIC collaborators include the University of North Carolina Renaissance Computing Institute (RENCI), the University of Illinois Urbana Champaign, the University of Kentucky, Clemson University, and ESnet. An operational national-scale testbed is slated to be unveiled in October of 2023. FABRIC is built on the ESnet6 fiber footprint, and for this effort, ESnet provided engineering and operational support to deploy the FABRIC nodes nationwide (and for FABRIC Across Border internationally). In addition, ESnet supplied consulting to finalize aspects of the experiment infrastructure, including security; formalized operational roles and procedures for production support; co-authored (with RENCI) the FABRIC Network Service Model; and convened the Science Advisory Committee.

ESnet also contributed software development to the project, including a new family of Layer-3 network service to the FABRIC Control Framework, a tool-based federation across multiple testbeds (e.g., FABRIC, Chameleon, and SENSE testbeds), support for Cloud connect services for FABRIC experimenters, and the implementation of an automated process to add external facilities interconnecting the various FABRIC sites.

FABRIC Across Borders (FAB)

Funded by the NSF's International Research and Education Network Connections (IRNC), the FABRIC Across Borders (FAB) project seeks to deploy four FABRIC nodes at the University of Bristol (UK), CERN (Switzerland), University of Amsterdam (The Netherlands), and University of Tokyo (Japan).

As part of this project, ESnet assists with the international node deployments, networking to the core FABRIC network, and working with international teams of researchers to develop advanced domain science workflows in the form of experiments on the combined FABRIC and FAB infrastructure. ESnet staff also continued to lead the connectivity planning for the CERN site deployment and to support experimenter coordination and connectivity requirements for LSST and CMB-S4.

At the Crossroads for the Self-Driving Network with AI Engine Hecate

Imagine a "self-driving network" that combines telemetry, automation, dev-ops, and machine learning to create an infrastructure that is responsive, adaptive, and predictive on the fly. Such a network will require an amalgamation of multiple AI capabilities integrated with the networking infrastructure that can perform various functions to improve network aspects such as performance and traffic engineering.

Scientific research networks pose specific traffic engineering problems:

1. Large science traffic flows are often random, peaking when large facilities run or have multiple spontaneous large data transfers.
2. Science traffic exhibits high variability, encompassing diverse performance requirements such as deadline-driven transfers, low-latency transfers, and long-lived flows that often clog up the network.
3. Network performance is extremely critical in preventing loss. Losing packets can seriously jeopardize the integrity of science results.

The primary goal for network optimization is to ensure the best possible network design and performance while minimizing the total cost. In networks designed to support science, optimizing network performance involves minimizing packet loss and optimizing for both high- and low-latency traffic.

Led by ESnet Planning & Innovation researchers and named for the Greek goddess of the crossroads, Hecate is an AI engine that tackles the traffic engineering component. ESnet's experience in developing graph neural networks (GNNs) to help predict WAN traffic shows that networks, modeled as graphs, produce more accurate predictions than traditional techniques. The team used machine learning to recognize traffic patterns and perform path finding using real-time network telemetry (i.e., network health telemetry data including latency, utilization, and loss), for packet route calculation and real-time dynamic path optimization.

A patent for Hecate was filed in November 2022 — "Autonomous Traffic (Self-Driving) Network with Traffic Classes and Passive/Active Learning" — and is pending. The technology can be licensed through Berkeley Lab's Intellectual Property Office.

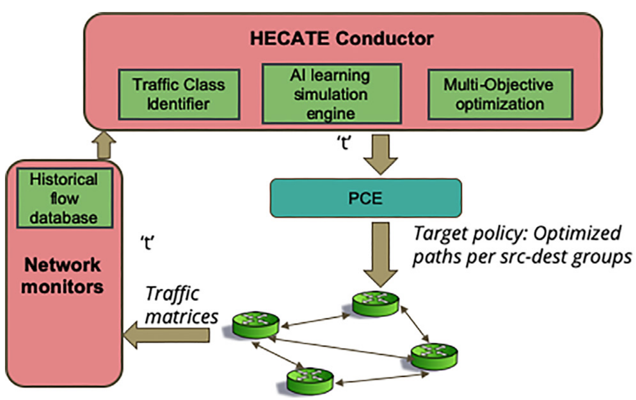

Hecate traffic engineering instantiation.

Using In-Network Data Caching to Reduce Traffic Volume

In large scientific collaborations, multiple researchers frequently need access to the same set of files for various analyses. This results in repeated interactions with large amounts of shared data that are located at considerable distance. These data interactions have high latency due to distance and consume the limited bandwidth available on the WAN. To address this, ESnet installed regional data storage caches as a new networking service to reduce the WAN traffic and data access latency.

In CY2022, ESnet collaborated with LHC's high-energy physics experiment, CMS, to gain insight into networking characteristics and the predictability of network utilization. These insights encompassed network traffic reduction, data throughput performance, and more. ESnet aimed to understand the predictability of network utilization to assist in planning for the deployment of additional in-network caches in the science network infrastructure. By exploring about 3 TB of operational logs, we observed that, on an average day, this cache removed 67.6% of file requests from the wide-area network and reduced the traffic volume on the wide-area network by 12.3TB (or 35.4%). ESnet developed and tested a machine learning model that can accurately predict the cache accesses, cache misses, and network throughput, making the model useful for future studies on resource provisioning and planning. It also deployed two additional caching nodes in Chicago and Boston. Additional deployments and comparative study to better understand in-network caches are planned for CY2023.

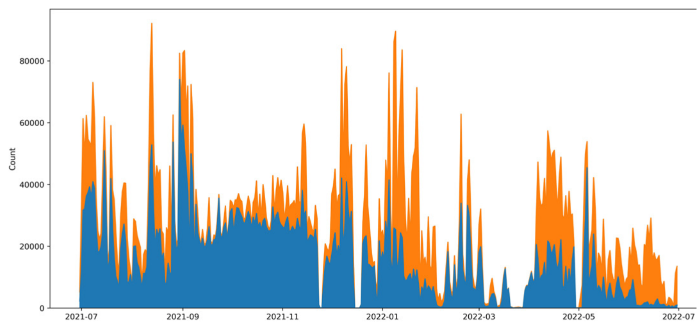

Based on Compact Muon Solenoid operational logs, the graph shows the number of files requested, separated into those that could be satisfied with files in the cache (or hits, in blue) and those that require wide-area network transfers (misses, in orange). In the SoCal Repo, an average day sees about 16,000 file requests as hits, along with 8,000 misses. (Note: Only file requests that miss the cache trigger remote network file transfers.)

Janus Framework Enables DTN-as-a-Service

Data transfer nodes (DTNs) are hosts specifically designed to optimize data transfers. ESnet's Advanced Technologies and Testbed Group has developed the DTN-as-a-Service (DTNaaS) model that explores the concept of a managed data movement service platform that can stage and move data across a wide area within a well-defined and dynamic core network. The model focuses on a prototype controller implementation and automated packaging of DTN transfer software (e.g., Globus) while also seeking to develop a robust framework able to accommodate the many types of software containers that require flexible deployment strategies.

The result is the DTNaaS model Janus, a lightweight orchestration framework built around exposing container configuration and tuning specifically for high-performance data mover applications. By specializing the feature set of the framework, Janus reduces complexity and offers a path to rapid, optimized deployment for common DTN hardware patterns. In particular, Janus addresses a common DTN requirement for multiple, high-speed network attachments in containers with IPv4/IPv6 dual-stack configurations. Janus's deployment for in-network caching, the ESnet testbed, the FABRIC software federation, and ESnet community DTNs necessitated several capability upgrades in the underlying software in 2022. Janus now successfully manages the container images and dual-homed container network attachments on the caching nodes at ESnet's Chicago and Boston sites. Based on these and other successes, the original caching pilot node at Sunnyvale has also begun the process of transitioning to the Janus/DTNaaS deployment model.

EJ-FAT: Real-Time DSP Processing for Edge Compute

The DOE's particle accelerators, x-ray light sources, electron microscopes, and other large science facilities are instrumented with many high-speed data acquisition systems (DAQs) that can produce multiple 100 Gbps data streams for recording and processing. This processing is conducted on large banks of compute nodes in local and remote data facilities. To increase the ability to integrate these experiments with high performance computing (HPC) centers, in CY2022 ESnet collaborated with the Thomas Jefferson National Accelerator Facility (JLab) to develop a prototype real-time load balancer for distributing UDP-encapsulated DAQ payloads into a dynamically allocated set of compute elements.

Named EJ-FAT — for ESnet JLab FPGA Accelerated Transport — the load balancer was implemented using FPGA technology to achieve line rates in excess of 100 Gbps, and has successfully integrated with JLab's ERSAP processing pipeline for end-to-end event processing. The entire workflow has been demonstrated to perform real-time, end-to-end DSP processing, at 100 Gbps.

Designed to support WAN latencies for geographically distributed accelerator facilities and HPC centers, EJ-FAT demonstrates that computing can be integrated with data acquisition without burdening the data acquisition hardware with the details of the computing environment. Through streamlining the data acquisition system, this demonstrates the potential for computing to scale up both in terms of the number of compute nodes and in the sophistication of the computing design. This facilitates the utilization of remote computing for enhanced compute resiliency and enables experiments with high-speed DAQs that may not be able to deploy sufficient local computing resources to fully leverage some experimental modalities. EJ-FAT also makes it easier for a computing facility to host the processing for multiple accelerators or experiments, and facilitates researchers use of the best computing for their experiment configuration.

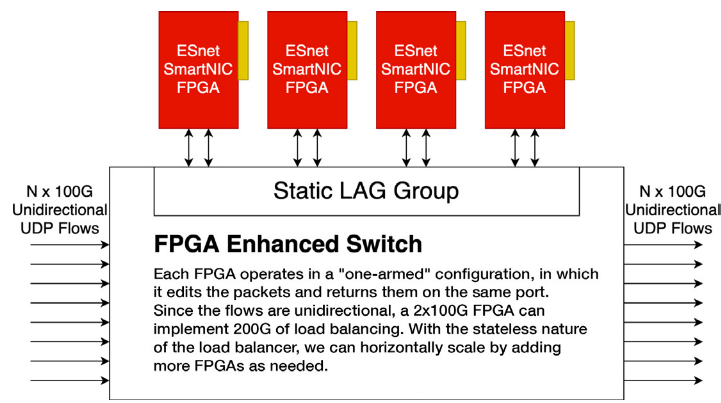

A scalable implementation of the load balancer showing packet processing using ESnet SmartNICs built on AMD-Xilinx Alveo cards.

Designing the QUANT-NET Architecture and Control Plane

Today, quantum networks are in their infancy. Like the Internet, quantum networks are expected to undergo different stages of research and development until they reach practical functionality. Based on laws of quantum mechanics (such as superposition, entanglement, quantum measurement, and the no-cloning theorem), researchers are envisioning quantum networks with novel capabilities that could be transformative to science, the economy, national security, and more.

The Quantum Application Network Testbed for Novel Entanglement Technologies (QUANT-NET), a project funded by DOE/ASCR, brings together expertise and resources at UC Berkeley, the California Institute of Technology, and the University of Innsbruck to build a three-node distributed quantum computing testbed between two sites, Berkeley Lab and UC Berkeley, connected with an entanglement swapping substrate over optical fiber and managed by a quantum network protocol stack. The team plans to demonstrate entanglement between small-scale ion trap quantum processors at both locations. On top of this capability, they will implement the most fundamental building block of distributed quantum computing by teleporting a controlled-NOT gate between two nodes.

In 2022, QUANT-NET researchers made significant progress in the R&D areas of quantum testbed infrastructure, trapped-ion quantum technology, quantum frequency conversion R&D, color center quantum technologies, and QUANT-NET architecture and control plane. Also in 2022, ESnet's quantum network research group completed an initial design of QUANT-NET's architecture, utilizing a centralized control model that is modular, flexible, extensible, and deployable, as well as an initial design of the control plane.

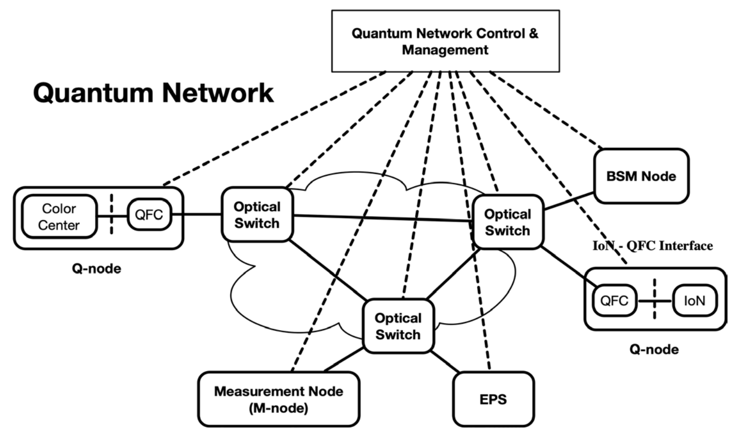

QUANT-NET network architecture (centralized control): modular, flexible, extensible, and deployable.

Scientific Research Collaborations

ESnet's goal is to create the best possible network environment for data intensive, distributed research projects across the DOE ecosystem. Achieving that goal requires strong "co-design" partnerships with various stakeholders. This is a primary focus of our Science Engagement team, but multiple departments within ESnet also actively participate in these collaborative efforts, which range from engaging with individual principal investigators to conducting the formal Requirements Review program.

The following projects provide a sampling of ESnet's efforts to create an environment in which scientific discovery is completely unconstrained by geography — and scientists can remain unburdened and unaware of the sophisticated network infrastructure underlying their work.

The Engagement and Performance Operations Center (EPOC): A Toolbox for R&E Networks

ESnet has been collaborating with the Texas Advanced Computing Center (TACC) as co-PI to establish the Engagement and Performance Operations Center (EPOC), using an NSF grant totaling $3.5 million. (Originally intended for three years, the grant's duration was extended to five years due to pandemic-related delays.) EPOC is designed to provide researchers — and the IT engineers supporting them — with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. The project scope supports five main activities:

- Providing "roadside assistance" through a coordinated Operations Center to actively resolve network performance issues with end-to-end data transfers.
- Conducting application deep dives to establish closer collaboration with application communities, understanding end-to-end research workflows, and assessing bottlenecks and potential capacity problems.
- Enabling network analysis using the NetSage monitoring suite to proactively discover and resolve performance issues. NetSage is currently deployed at eight locations within networks, campuses, and facilities
- Coordinating training activities to ensure effective use of network tools and science support.
- Emphasizing data mobility, through the Data Mobility Exhibition effort.

EPOC's user data provides deeper insight into the needs of scientific research collaborations and thus transforms ESnet's ability to achieve better data transfers. In 2022, the ESnet/TACC EPOC team gave 50 talks at virtual and in-person events; addressed 88 roadside assistance tickets, bringing the total the project has addressed since the start to 277; and completed and published the results of six deep dive events.

Designing the Superfacility Networking Approach

The Superfacility Project at Berkeley Lab, which began in 2019 and concluded in 2022, has been a strategic priority for ESnet. The superfacility model is designed to leverage high-performance computing (HPC) and high-speed networks for experimental science. More than simply a model of connected experiment, network, and HPC facilities, the superfacility program envisioned an integrated ecosystem of compute, networking, and experimental infrastructures, supported by common, shared software, tools, and expertise needed to make those connected facilities easy to use.

Researchers used Stanford's Linac Coherent Light Source to pioneer a new form of X-ray crystallography, in which experimental data was transferred automatically via ESnet to supercomputers at NERSC and back, yielding initial analysis in under 10 minutes — a speed record for this type of experiment. *(Image Credit: Ella Maru Studios and J. Nathan Hohman.)*

A key component of the Superfacility Project has been in-depth engagements with science teams that represent challenging use cases across the DOE Office of Science. The goal was to develop and implement automated pipelines that analyze data from remote facilities at large scale without routine human intervention. In several cases, the project team went beyond demonstrations and now provides production-level services for the experiment teams, which include:

- Dark Energy Spectroscopic Instrument (DESI): Automated nightly data movement from telescope to NERSC and deadline-driven data analysis.
- Linac Coherent Light Source (LCLS): Automated data movement and analysis from several experiments running at high datarate end stations.
- National Center for Electron Microscopy (NCEM): Automated workflow pulling data from the 4D STEM camera to the Cori supercomputer at NERSC for near-real-time data processing.
- LUX-ZEPLIN experiment (LZ): Automated 24/7 data analysis from the dark matter detector (see page 36).

In the next few years, ESnet will be applying what it has learned from the Superfacility Project implementations to operationalize the DOE's Integrated Research Infratructure initiative.

Bolstering Data Acquisition and Transfer for GRETA

GRETA (Gamma Ray Energy Tracking Array) is a new instrument designed to reveal novel details about the structure and inner workings of atomic nuclei, elevating our understanding of matter and the stellar creation of elements. GRETA will be installed at the DOE's Facility for Rare Isotope Beams (FRIB) located at Michigan State University in East Lansing; it will go online in 2024.

GRETA will house an array of 120 detectors that will produce up to 480,000 messages per second — totaling 4 gigabytes of data per second — and send the data through a computing cluster for real-time analysis. Although the data will mostly traverse a network of about 50 meters for on-the-fly analysis, the system has been designed to easily send data to more distant HPC systems as necessary.

Working with the Berkeley Lab nuclear physicists building the GRETA experiment, ESnet is:

- Designing the computing environment to support the experiment's operation and data acquisition; and providing the system architecture, performance engineering, networking, security, and other related expertise.
- Advancing multiple key high-performance software components of the data acquisition system: the Forward Buffer, Global Event Builder, Event Streamer, and Storage Service.
- Developing several simulations of the instrument's various components, allowing for functional and performance testing.

In CY2022, ESnet completed the initial implementation of the Event Streamer; it is now ready for the integration phase.

Strengthening the Energy Grid with the ARIES Project

The National Renewable Energy Laboratory's (NREL) Advanced Research on Integrated Energy Systems (ARIES) "digital twin" project uses ESnet to unite research capabilities at multiple scales and across sectors, creating a platform for understanding the full impact of energy systems integration. ARIES will make it possible to understand the impact and get the most value from the millions of new devices that are connected to the grid daily, such as electric vehicles, renewable generation, hydrogen, energy storage, and grid-interactive efficient buildings.

In early December 2021, ARIES researchers teamed up with the Pacific Northwest National Laboratory (PNNL) and ESnet staff to demonstrate these capabilities during a live demonstration of a microgrid modeling experiment. This multi-laboratory demonstration showed that advanced control systems in the Cordova, Alaska, microgrid could allow it to maintain power to critical resources such as the hospital and the airport during an extreme weather event and loss of hydro-power resources.

During this tightly coupled experiment, NREL simulated the Cordova microgrid while PNNL simulated the advanced control systems. Using its On-demand Secure Circuits and Advance Reservation System (OSCARS), ESnet provided a reliable, low-latency connection so that research equipment at the two laboratories could exchange frequent command and control information.

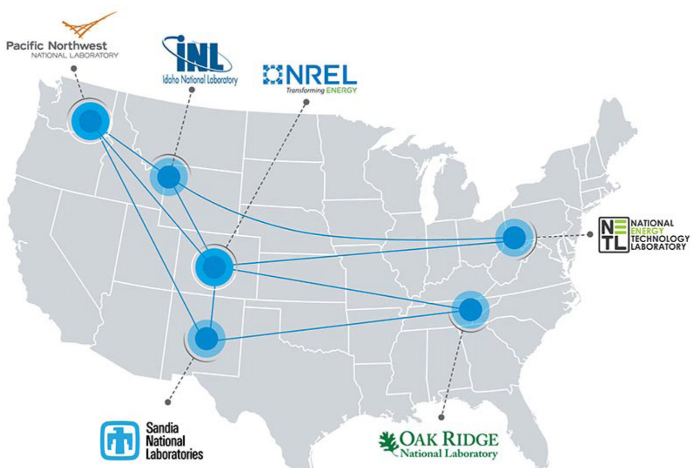

The participants in the ARIES digital twin project.

The round-trip latency between NREL and PNNL was 24 milliseconds, with a variance of approximately 0.02 milliseconds. The average latency in the original SuperLab demonstration was about 27 milliseconds, with a variance of about 11.5 milliseconds. The very low latency variance was vital to the project's success and made exchanging command-and-control information between the two laboratories nearly deterministic.

ARIES experiments continue and are being expanded to explore new use-case scenarios and to involve more national laboratory participants. A quad lab experiment is being planned in 2024 for NREL, Idaho National Laboratory, Sandia National Laboratory, and National Energy Technology Laboratory; ESnet is providing both networking and technical consulting capabilities to support these efforts.

Shaking Up Data Sharing for EQSIM

EQSIM — one of the projects backed by the DOE's Exascale Computing Project — has been pioneering advancements since 2017. Its focus is on assisting researchers in comprehending the impact of seismic activity on the structural soundness of buildings and infrastructure. The EQSIM software was built and tested using the Summit supercomputer at Oak Ridge National Laboratory and the Cori and Perlmutter systems at Berkeley Lab. It uses physics-based supercomputer simulations to predict the ramifications of an earthquake on buildings and infrastructure and create synthetic earthquake records that can provide much larger analytical datasets than have been previously available. Thanks to the new GPU-accelerated HPC systems, the EQSIM team was able to create a detailed, regional-scale model that includes all of the necessary geophysics modeling features, such as 3D geology, earth surface topography, material attenuation, non-reflecting boundaries, and fault rupture.

ESnet was instrumental in enabling the sharing of the massive amounts of data needed to design and operate the EQSIM software across multiple platforms, noted EQSIM principal investigator David McCallen, a senior scientist in Berkeley Lab's Earth and Environmental Sciences Area. "EQSIM workflow is designed around exploitation of multiple platforms and efficient data transfer," McCallen said during the ESnet6 unveiling in October 2022. "ESnet is the glue that stitches it all together as we transfer our simulation data back and forth from one site to the other."

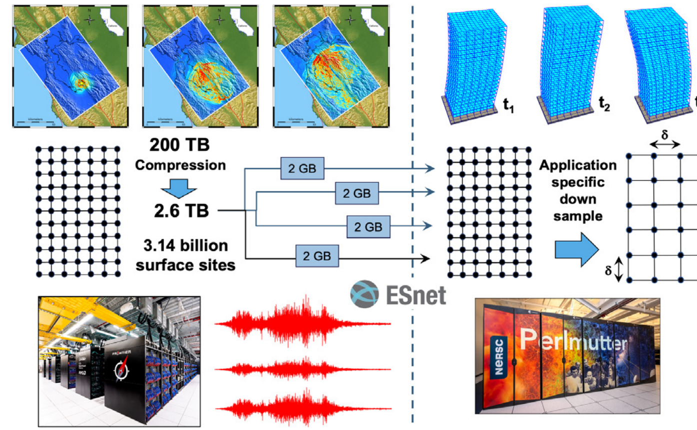

EQSIM data workflow diagram showing the use of two computing facilities connected via ESnet.

Exploring Networking at the Wireless Edge

Wireless capabilities such as private 4G/5G, satellite, and Internet of Things (IoT) mesh network concepts are making it possible for a wide variety of field scientists to deploy mobile and remote sensor systems in ways never before possible. The ability to simplify deployment of ESnet services beyond the optical edge and, using advanced wireless technologies, make it possible for DOE researchers to transfer data seamlessly to DOE HPC centers and user facilities will be a powerful enabler for future science breakthroughs across a wide range of disciplines.

In 2022, ESnet procured and deployed Nokia Digital Automation Cloud private 4G Citizens Band Radio Service (CBRS) hardware — which will be upgraded to the 5G standard in CY2023 — to explore scientific use cases, identify future service needs supporting “wireless edge” field science applications, and conduct wireless networking and edge compute research. The CBRS system is being deployed as a multi-site federation, supporting activities at Berkeley Lab in Berkeley, California, as well as Earth and Environmental Sciences Area (EESA) research in a remote area near Mt. Crested Butte, Colorado. This EESA collaboration has required the deployment of the system using a combination of wireless technologies to support backhaul of sensor data from a sensor field in the Rocky Mountains back to NERSC, including CBRS, Starlink satellite service, Long Range (LoRa) sensor mesh networking, remote off-grid Wifi, and commercial cellular service. ESnet will build upon this capability in future years as part of a focused effort to identify the correct services and technologies to best facilitate the science programs we support.

The EESA-ESnet collaboration in Crested Butte, Colorado, sends sensor data from the field to NERSC via a system utilizing Citizens Band Radio Service hardware, Starlink satellite service, Long Range (LoRa) sensor mesh networking, remote off-grid Wifi, and commercial cellular service. (The white sphere is a Doppler weather radar.)

Accelerating LUX-ZEPLIN’s Search for Dark Matter

ESnet is playing a critical role in facilitating the burgeoning LUX-ZEPLIN (LZ) experiment being used to help researchers find dark matter in the Universe. Located in the Sanford Underground Research Facility (SURF) in South Dakota, the uniquely sensitive LZ detector is designed to detect dark matter in the form of weakly interacting massive particles. Led by Berkeley Lab, the experiment has 34 participating institutions and more than 250 collaborators.

Data collection from the LZ detector began in 2021; at present, about 1 petabyte of data is collected annually. The data pipeline runs 24/7, relying on ESnet and the National Energy Research Scientific Computing Center to consistently manage and transfer that data to multiple facilities. Each day’s batch of raw data is initially sent up to the SURF surface facility, then on to NERSC via ESnet6 for processing, analysis, and storage. The raw and derived data products are also sent to a data center in the United Kingdom for redundancy.

In 2022, the LZ dark matter detector passed a check-out phase of startup operations and delivered its first results. LZ researchers reported that, as a result of that initial run, LZ is currently considered the world's most sensitive dark matter detector.

Replicating a Massive Climate Data Set

In 2022, ESnet enabled the replication of a very large-scale data set of international importance to two HPC facilities. The more than 7 petabytes of climate data belonged to the Earth System Grid Federation (ESGF), a peer-to-peer enterprise system that develops, deploys, and maintains software infrastructure for the management, dissemination, and analysis of model output and observational data. ESGF's primary goal is to facilitate advancements in Earth system science. It is an international, interagency effort led by the DOE and co-funded by the National Aeronautics and Space Administration, National Oceanic and Atmospheric Administration, National Science Foundation, and several international laboratories.

A multi-lab team replicated more than 7 PB of ESGF data from Lawrence Livermore National Laboratory (LLNL) to the Argonne Leadership Computing

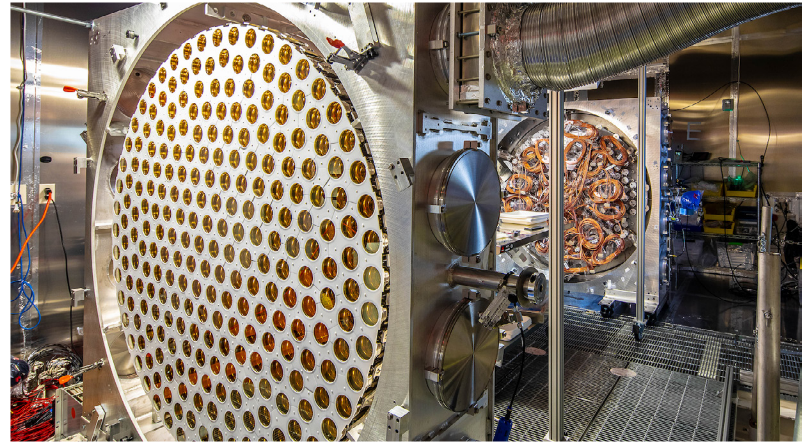

Part of the LUX-ZEPLIN (LZ) dark matter detector, located deep in the Black Hills of South Dakota in the Sanford Underground Research Facility. *(Photo Credit: Matthew Kapust/SURF)*

Facility (ALCF) and Oak Ridge Leadership Computing Facility (OLCF) over ESnet6. The team used the Globus data transfer platform to ensure that the transfer occurred reliably and with high performance. The data sets were transferred from LLNL to either the ALCF or the OLCF and then replicated between them to ensure that both sites successfully transferred a complete copy of the data. The existence of these replicas provides several advantages, including resilience and protection against catastrophic data loss, and improved data proximity to high performance computing centers.

Community Engagement and Knowledge Transfer

ESnet's mission compels us to share knowledge, experience, and findings with both the DOE ecosystem and the global R&E networking community, including commercial networking technology providers. Technical outreach and community building guarantee that we remain at the forefront of understanding emerging scientific needs and technology options while also ensuring visibility for DOE's global science collaborations.

Community engagement and knowledge-transfer activities encompass a variety of approaches, including technical consulting, publishing, and participation in DOE and other U.S. scientific planning endeavors. Additionally, ESnet contributes through keynote speeches, technical presentations, demonstrations, and workshops at numerous high-profile conferences annually.

ESnet Launches First Annual Confab Gathering

In October 2022, ESnet's Science Engagement Group convened and hosted ESnet's first Science User conference to foster cross-interactions between DOE research scientists, users, network engineers, and stakeholders. The focus was different from the longtime ESnet Site Coordinators Committee (ESCC) meeting, a technical forum that brings together IT representatives from the ESnet-connected sites. Held in Berkeley, California, in conjunction with the unveiling of ESnet6, Confab22 sought to showcase how scientists across domains use ESnet6 and its advanced services effectively and to provide a platform for bidirectional conversations, particularly across siloed application and IT communities.

The technical program enabled discussions among the 100-plus attendees about co-design activities including:

- ESnet6 capabilities in development and science user thoughts, with ESnet leadership
- Exascale biology compute needs and network uses, with Dan Jacobson of Oak Ridge National Laboratory
- ARIES project (see page 33) and future "superlab" network uses, with Rob Hovsapian of National Renewable Energy Laboratory
- Field Science and Wireless/5G use cases, with ESnet staff
- DOE Integrated Research Infrastructure initiative with ASCR Facilities Director Ben Brown
- Derecho simulation and data visualization with Paul Ullrich, U.C. Davis, and Brad Carvey, Sandia National Laboratory.

The discussions started at Confab have resulted in several new collaborations between ESnet's Science Engagement group and external stakeholders. The conference will now become a regular ESnet activity driving outreach and user support efforts.

Oak Ridge National Laboratory's Dan Jacobson presents use cases for Exascale Computing Project computational systems biology to the Confab22 audience.

Other Key Conferences

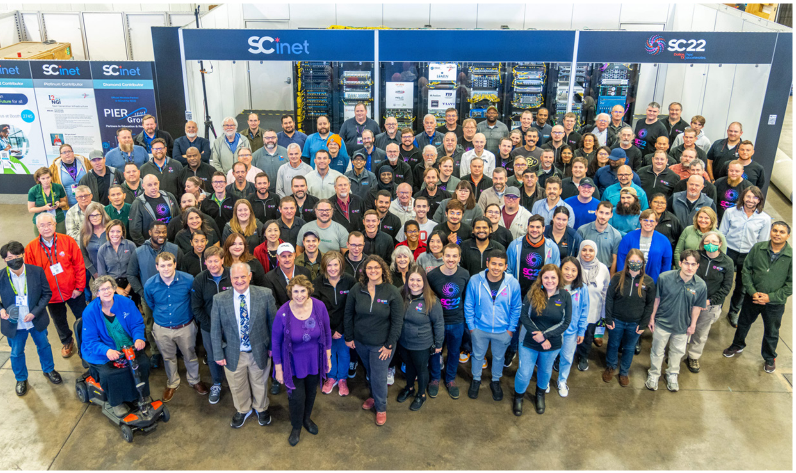

The SCinet22 team of volunteers.

SC

Each year, ESnet participates in the International Conference for High Performance Computing, Networking, Storage, and Analysis, known simply as SC, one of the largest high performance computing conferences in the world, with an emphasis on data-intensive scientific support. The contributions can be divided into two main categories: First, ESnet offers connectivity and services to the conference, facilitating live demonstrations and tests. Second, ESnet staff members actively participate as essential volunteer team members responsible for constructing SCinet — an annual conference network and the “world’s fastest network” for the duration of the conference. This network serves as a testbed for innovative problem-solving strategies tailored to address challenges within the HPC community.

TNC

The largest R&E networking conference, TNC is co-funded by the European Union via the GÉANT (GN4-3) Project. It attracts more than 800 participants from 70+ countries, representing national and regional research and education networks, schools and universities, technology providers, and scientific projects.

Internet2 Technology Exchange

Hosted by Internet2, a community providing a secure high-speed network, Cloud solutions, research support, and services tailored for research and education, the TechX conference brings together the global R&E community.

ESCC

Each year, ESnet convenes the ESnet Site Coordinator Committee to share insights and guide the future of the

ESnet infrastructure. In 2022, the group, made up of technical representatives from ESnet's supported sites, met in person for the first time since 2019 to collaborate on operational, network engineering, and security innovations and challenges.

The Quilt Annual Meeting

The Quilt is a national coalition of non-profit U.S. regional R&E networks representing 43 networks across the country. Members of The Quilt provide advanced network services and applications to 900-plus universities and thousands of other educational and community anchor institutions.

OFC

ESnet participates in talks and demonstrations for OFC, the Optical Fiber Communication Conference and Exhibition sponsored by Optica, IEEE Communication Society, and IEEE Photonics Society, as well as in building OFCnet, an unconventional high-speed network connecting the show floor to a research center in Chicago.

Nick Buraglio was among several ESnet presenters at TNC22.

ACM SIGCOMM

ACM SIGCOMM is the flagship annual conference of the ACM Special Interest Group on Data Communication, an important conference in networking held in Amsterdam in 2022. ESnet staff co-authored the invited paper "Transport Control Networking: Optimizing Efficiency and Control of Data Transport for Data-Intensive Networks" (see page 42) and gave a keynote.

Technical Presentations and Publications

In 2022, ESnet staff delivered 57 technical presentations, published two dozen conference papers and journal articles, and filed one patent. (For a full list, with links where available, visit es.net/publications2022.)

The following examples were selected to illustrate ESnet's work in the applied research areas featured in this report.

Accelerating data movement

J. Dunefsky, M. Soleimani, R. Yang, J. Ros-Giralt, M. Lassnig, I. Monga, F. Wüerthwein, J. Zhang, K. Gao, Y.R. Yang, "Transport Control Networking Optimizing Efficiency and Control of Data Transport for data-intensive Networks." NAI '22: Proceedings of the ACM SIGCOMM Workshop on Network-Application Integration (August 2022): 60-66.

E. Dart, "The Science DMZ: Strategic Future" (technical presentation, Corporation for Education Network Initiatives in California [CENIC] conference, Monterey, CA, September 2022).

AI-driven networking

M. Kiran, N. Buraglio, S. Campbell, "Using Deep Learning to Improve Data Transfers for WAN" (technical presentation, the Asia Pacific Advanced Network [APAN 53] Conference, AI-Driven Networks Working Group, Dhaka, Bangladesh, March 2022).

S. Murugesan, M. Kiran, B. Hamann, G.H. Weber, "Netostat: Analyzing Dynamic Flow Patterns in High-Speed Networks." Cluster Computing 25 (March 2022): 2915–2930.

H. Jin, K. Raghavan, G. Papadimitriou, C. Wang, A. Mandal, P. Krawczuk, L.Pottier, M. Kiran, E. Deelman, P. Balaprakash, "Workflow Anomaly Detection with Graph Neural Networks." Presented at 2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS), Dallas, TX, November 2022.

Automation testing environments

N. Buraglio, C. Cummings, "Building a Realistic Orchestration Validation Environment for netwoRks (ROVER)" (technical presentation, TNC22/GÉANT Conference, Trieste, Italy, June 2022).

Edge compute management

E. Kissel, C. Shiflett, "Janus Container Management and the EScp Data Mover" (Department of Energy booth demonstration, The International Conference for High Performance Computing, Networking, Storage, and Analysis [SC22], Dallas, TX, November 2022).

IETF standards activity

N. Buraglio, C. Cummings, R. White, "Unintended Operational Issues With ULA" Presented at Internet Engineering Task Force (IETF114) conference, V6Ops Network Working Group, Philadelphia, PA, July 2022.

In-network caching

J. Bellavita, A. Sim, K.Wu, I. Monga, C. Guok, F. Wüerthwein, D. Davila, "Studying Scientific Data Lifecycle in On-demand Distributed Storage Caches." SNTA '22: Fifth International Workshop on Systems and Network Telemetry and Analytics (June 2022): 43-50.

Orchestrated and coordinated data movement

F. Wüerthwein, J. Guiang, A. Arora, D. Davila, J, Graham, D. Mishin, T. Hutton, I. Sfiligoi, H. Newman, J. Balcas, T. Lehman, X. Yang, C. Guok, "Managed Network Services for Exascale Data Movement Across Large Global Scientific Collaborations." Presented at 2022 4th Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP), 2022 Supercomputing Conference (SC22), Dallas, TX, November 2022.

Quantum networking

I. Monga, E. Saglamyurek, W. Wu, "Quantum Communication: A Physical Experiment or a Network Paradigm Shift?" (technical presentation, 9th Annual International Workshop on Innovating the Network for Data-Intense Science [INDIS], Dallas, TX, November 2022).

A. Derevianko, E. Figueroa, J. MartÍnez-Rincón, I. Monga, A. Nomerotski, C.H. Peña, N.A. Peters, R. Pooser, N. Rao, A. Slosar, P. Spentzouris, M. Spiropulu, P. Stankus, W. Wu, S. Xie, "Quantum Networks for High Energy Physics." arXiv:2203.16979 [quant-ph] (March 2022).

Wireless integration

M. Kiran, "5G for Science" (technical presentation, 2022 Department of Energy [DOE] Cybersecurity and Technology Innovation Conference, Portland, OR, June 2022).

A. Wiedlea, "Exploring the Wireless Edge for Science: What Does that Mean for Us?" (technical presentation, 2022 Corporation for Education Network Initiatives in California [CENIC] Annual Meeting, Monterey, CA, September 2022).

Understanding real-time streaming

R. Kube, R.M. Churchill, C.S. Chang, J. Choi, R. Wang, S. Klasky, L. Stephey, E. Dart, M.J. Choi, "Near Real-Time Streaming Analysis of Big Fusion Data." Plasma Physics and Controlled Fusion 64, No. 3 (February 2022).

The People of ESnet

ESnet has almost tripled in size since 2017, scaling up to meet the needs of the ESnet6 project and the future demands of the exabyte era. At the end of 2022, the organization numbered 126 employees and contractors, plus an additional 10 students who worked with us during the year.

ESnet's organizational structure is continually being reimagined and reconfigured in response to current and anticipated future needs. The group and department structure following the completion of ESnet6 is shown at right, with staff totals as of December 2022.

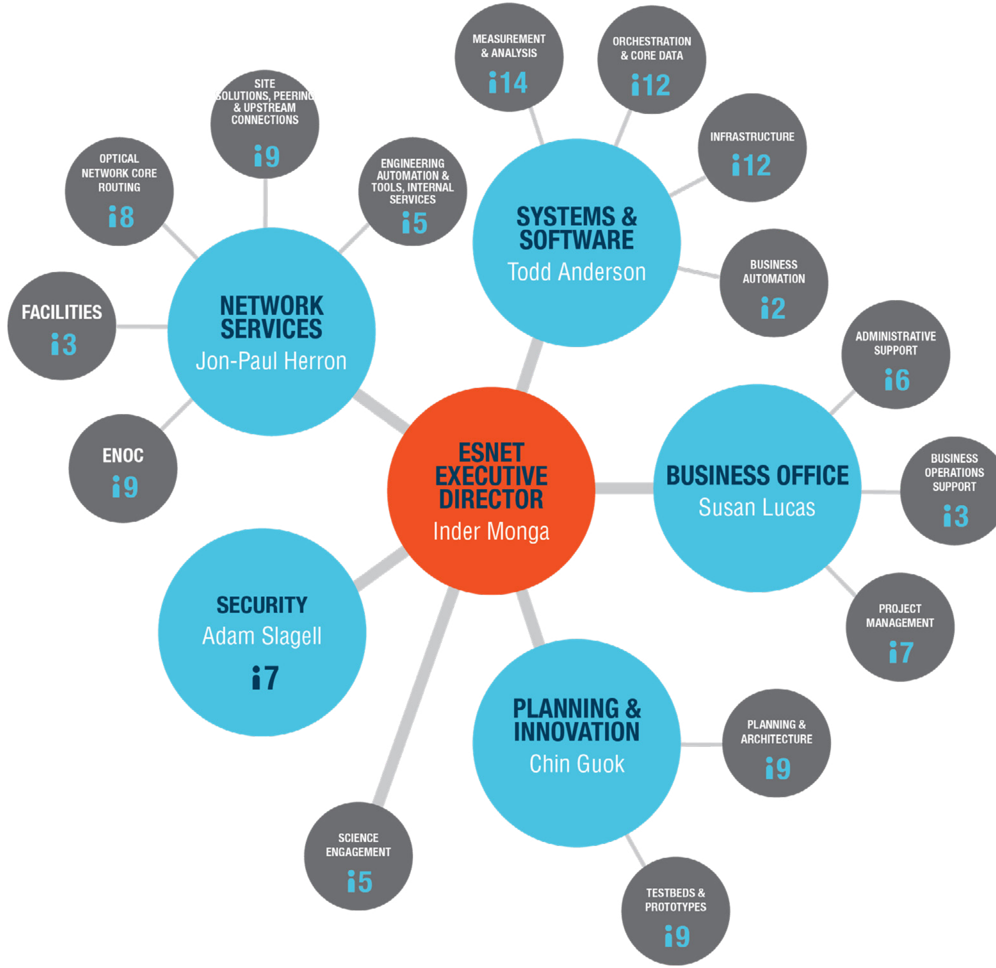

Left: ESnet team snapshot (partial) as of August 2023


ESnet staff at an on-site picnic in 2022.

ESnet's Culture

ESnet has experienced rapid growth while weathering a global pandemic and completing a massive, technically complex project. In 2022, ESnet conducted an organization-wide employee survey. Its results inspired the launch of a new Culture & Engagement Program aimed at further strengthening staff satisfaction and retention rates. Activities included revamping recognition programs, sending a regular staff newsletter, offering training to improve collaboration and strengthen leadership, and planning workshops in which to develop our shared values.

ESnet is proud to be stewarded by Berkeley Lab, which celebrates a rich tradition of multidisciplinary teams working together to bring science solutions to the world that began with founder E.O. Lawrence. We too believe that fostering a diverse workforce — diverse in experiences, perspectives, and backgrounds — and a culture of inclusion are necessary to attract and engage the brightest minds and advance scientific excellence. Berkeley Lab's IDEA (for inclusion, diversity, equity, and accountability) initiative offers robust new data, training, and opportunities that ESnet looks forward to incorporating further into our hiring, mentoring, and Culture & Engagement activities.

Women in IT Networking at SC (WINS)

For the past nine years, ESnet has co-led a workforce development/diversity initiative at the SC conference, called WINS, for Women in IT Networking at SC. The program, a joint effort between ESnet, Indiana University, the Keystone Initiative for Network-Based Education and Research, and the University Corporation for Atmospheric Research (UCAR), was developed to address the gender gap in information technology, particularly in the field of network engineering.

Through a competitive application process, WINS selects and provides travel funding for early- to mid-career women to participate as SCinet volunteers. They are matched with a SCinet team and a mentor based on skills and interest area, which range from wireless networking to IP routing, from fiber-optic cable installation to cybersecurity. In 2022, ESnet served on the WINS management team, participated in developing the application and selection process, and offered full or partial funding to seven attendees to participate in SciNet code sprint, staging and/or setup, and show activities. The participants came from institutions ranging from R1 research schools to community and tribal colleges. In addition to the 2022 WINS cohort, 13 prior WINS participants — many of whom hold leadership roles on the SCinet team — returned with their own institutional funding to participate in SCinet.

Mentoring Students

Every year, ESnet offers 6 to 12 graduate, undergraduate, and high school students paid internships to work with scientists and engineers to address challenges in high-speed scientific networking for 12 to 16 weeks. Some of these students stay on for another semester or return to ESnet later to continue their research.

In 2022, ESnet hosted two Ph.D. students, three master's students, two undergraduates, and two high school students, including students from the DOE Science Undergraduate Laboratory Internship (SULI) program and the Berkeley Lab Experiences in Research (EinR) program. ESnet staff enjoy the opportunity to mentor these potential and early-career professionals, and to widen the career pipeline in keeping with IDEA's principles.

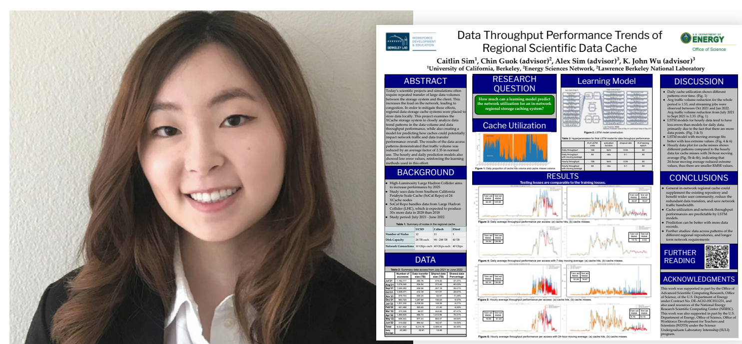

UC Berkeley undergraduate Caitlin Sim had a 2022 summer Science Undergraduate Laboratory Internship with ESnet.


NO PARKING

FOR MORE INFORMATION ABOUT ESNET, CONTACT:

ESnet Communications
Lawrence Berkeley National Laboratory
1 Cyclotron Road
Berkeley, CA 94720-8148

CSCOMMS@LBL.GOV | WWW.ES.NET

ESnet Annual Report Editor: Bonnie Powell

Design: Samantha Trieu, Berkeley Lab Creative Services

Photography: Berkeley Lab and ESnet, except where noted.

DISCLAIMER

23-SN-35047


ESnet
ENERGY SCIENCES NETWORK
BERKELEY LAB