ESnet
ENERGY SCIENCES NETWORK

ESnet: BES Network Requirements for Multifacility Workflows

Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory

Eli Dart, Science Engagement Engineer
Jason Zurawski, Science Engagement Engineer

December 7, 2022

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Topics

- ESnet Overview
- ESnet Requirements Program
  - BES 2022 Specifics
  - Key Findings/Opportunities/Actions

- Credits:
  - Inder Monga, ESnet director
  - Jason Zurawski, ESnet requirements program lead

Mission network

Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.

Vision:

Accelerate Scientific Discovery

The ESnet user facility:
Data Circulatory system for all 28 SC facilities

>30,000 science users

ASCR High End Computing (HEC)

Argonne Leadership Computing Facility (ALCF)

Oak Ridge Leadership Computing Facility (OLCF)

National Energy Research Scientific Computing Center (NERSC)

ASCR High Performance Scientific Network

Energy Sciences Network (ESnet)

BES X-Ray Light Sources


Advanced Photon Source (APS)


Linac Coherent Light Source (LCLS)


Stanford Synchrotron Radiation Light Source (SSRL)


Advanced Light Source (ALS)


National Synchrotron Light Source II (NSLS-II)

BES Nanoscale Science Research Centers (NSRCs)


Center for Functional Nanomaterials (CFN)


Center for Integrated Nanotechnologies (CINT)


The Molecular Foundry (TMF)


Center for Nanophase Materials Sciences (CNMS)


Center for Nanoscale Materials (CNM)

BES Neutron Scattering Facilities


Spallation Neutron Source (SNS)


High Flux Isotope Reactor (HFIR)

BER


Joint Genome Institute (JGI)

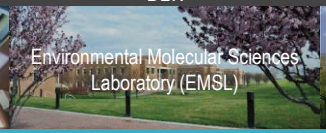
Environmental Molecular Sciences Laboratory (EMSL)


Atmospheric Radiation Measurement (ARM) user facility

FES


National Spherical Torus Experiment - Upgrade (NSTX-U)


DIII-D National Fusion Facility (DIII-D)

HEP


Facility for Advanced Accelerator Experimental Tests (FACET)


Fermilab Accelerator Complex


Accelerator Test Facility (ATF)

NP

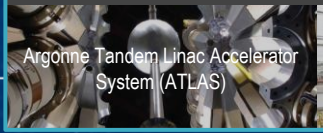
Argonne Tandem Linac Accelerator System (ATLAS)


Continuous Electron Bean Accelerator Facility (CEBAF)

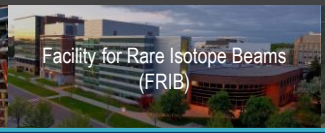
Facility for Rare Isotope Beams (FRIB)

Relativistic Heavy Ion Collider (RHIC)

U.S. DEPARTMENT OF ENERGY
Office of Science

ESnet

ESnet6 lays the foundation for future of data-intensive DOE science

Enough base capacity and ability to cost-effectively add more provides unconstrained access to data, no matter how big or distributed

- **15,000 miles** of fiber across the continental US
- **300 leased colocation spaces** installed with ESnet optical equipment
- **46.1 Tbps** aggregate capacity deployed
- **400Gbps - 1 Tbps** services available
- New fiber spans acquired to **increase reliability and reduce latency**

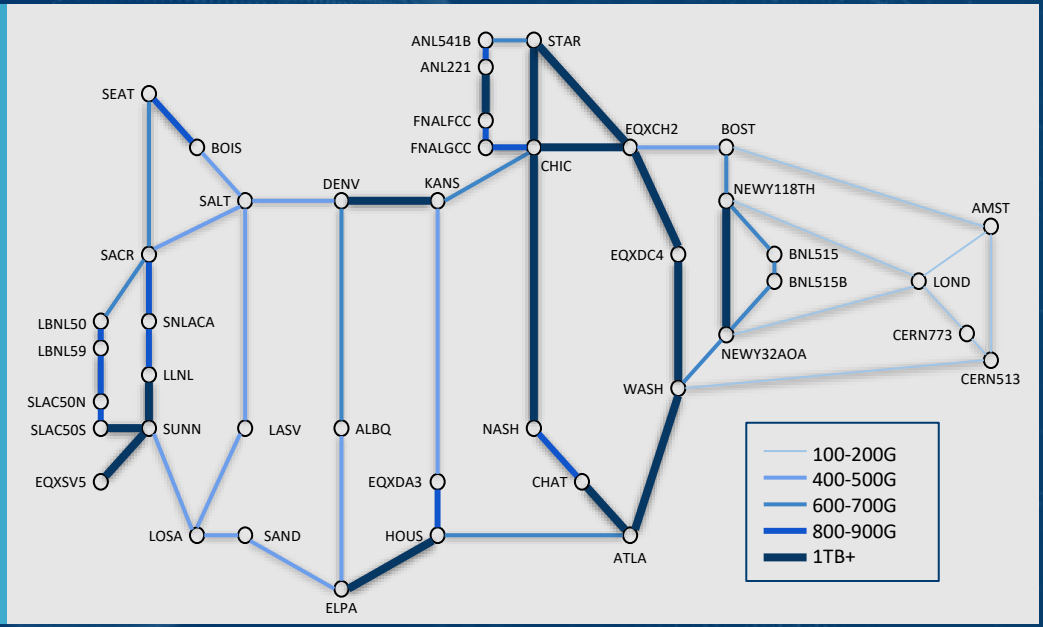

LUMEN Infinera ciena NOKIA ESnet

ESnet6 Serves All DOE National Laboratories

- More than just facilities
  - National Lab staff
  - Institutional computing
- Broad global connectivity
  - Science networks
  - Universities
  - Commodity Internet
  - Cloud providers
  - International

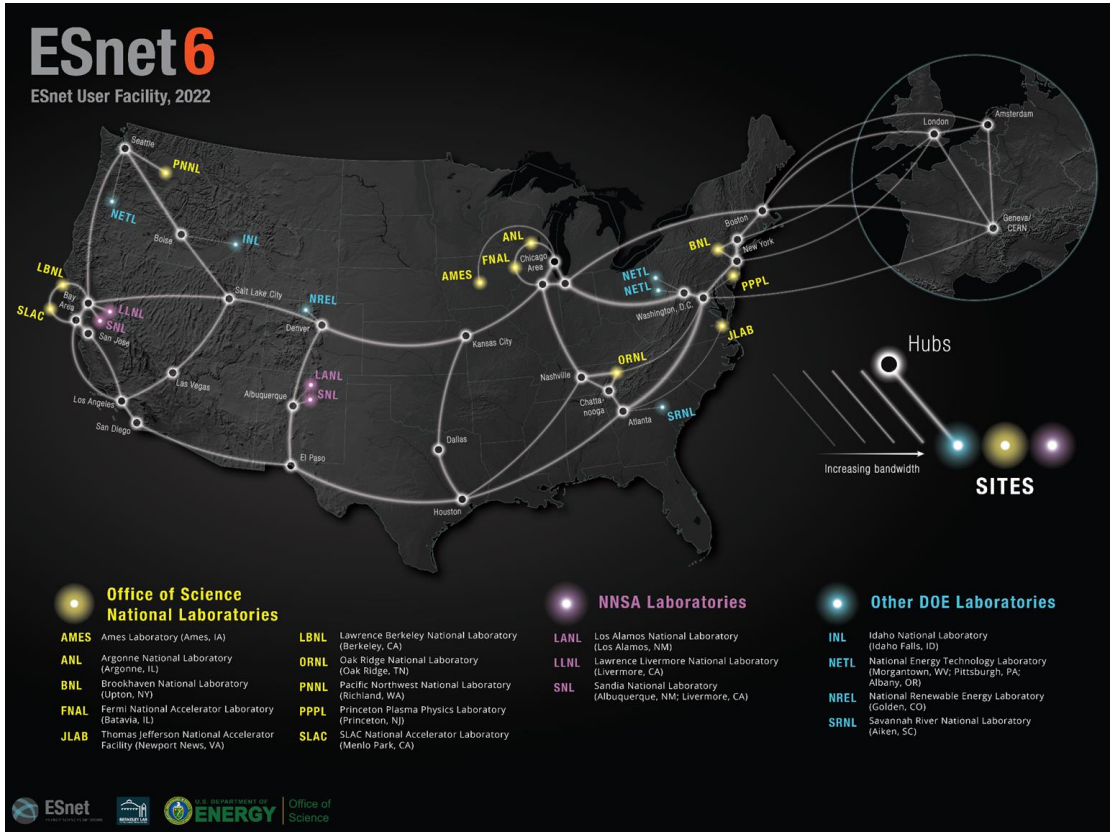

Requirements Program: Strategic Context

- ESnet's Requirements program facilitates conversations between ESnet, our sites, our program management, and the other DOE program offices.
- ***We learn from each other***
- Not only do we gain insight into the data strategy of programs, we can teach programs and scientists about the value of ESnet
  - Best practice, design patterns (Science DMZ, DTNs, perfSONAR)
  - Benefits of using high-performance networking in effective ways

ESnet

Requirements Program Overview

ESnet's core partnership program created to comprehensively evaluate:

- ***Major science experiments and facilities***, both in operation and planned.
- ***The process of science*** used for knowledge discovery, and including scientists' interactions with the instruments and facilities.
- ***The volume of data*** produced now, and anticipated in the future, with an emphasis on geographical location of where the data must be shared, computed and/or stored.
- ***The current technology capabilities*** (network, computation, storage, and software stack) used by each science collaboration/facility as well as any planned upgrades, additions or improvements.

ESnet

Review Outcomes

- Identify technical gaps, bottlenecks, and opportunities that can be addressed by ESnet ***with new or existing services*.**
- ***Forecast network capacity needs***, particularly sites anticipating data increases for ***informed investments*** in bandwidth, and services by ESnet and DOE.
- Ensure the program ***continues to be forward-looking*** even in very challenging future environments.
- Build a robust relationship between ESnet and the BES community ***to ensure continued alignment*** in the short- and long-term.

ESnet

Selected Requirements Review Findings

- A full discussion of all findings would take more time than we have here today
- What follows is a brief subset of the material
  - Organized around one aspect: multi-facility workflows
  - This was identified as an important area of emphasis
- There are many other important findings - see full report:
  - https://www.osti.gov/biblio/1899590
  - Pages 1-36

ESnet

Summarized Findings:

Facility Data Scale

- Over the next decade the combined data generation rates for some BES user facilities will reach the exabyte per year range (e.g., multiple orders of magnitude more data than is generated today). To be successful:
  - Hundreds of PFLOPs of processing power must be available to fully analyze the data
  - Strong networking capabilities must exist to link facility instruments to edge, local, campus, and centralized computing facilities reliably, with low latency.
  - Terabit per second networking and beyond will be required to handle the large amounts of data expected in the coming years.

ESnet

Summarized Findings:

Scientific Data Management

- Upgraded BES light source data generation rates may be too large for traditional file-based workflows, which will facilitate a move to streaming-based workflows directly to computer system memory
- This will require robust facility and Laboratory networking, and increases in data processing capabilities throughout the DOE scientific complex.
- Key elements of a future data management strategy for the light sources include a common API for accessing network and computing resources, parallel data-transfer tools, high-performance data transfer, network performance monitoring, reservations, and dynamic network provisioning.

ESnet

Summarized Findings:
Scientific Workflow

- Future adaptations to the generalized workflow used to support BES science will involve:
  - Integrated software to manage data movement, processing, and analysis.
  - More automated methods to handle adaptive/autonomous experiment steering (e.g., integration with simulation capabilities).
  - Categorizing and sharing of data with other national resources.
  - Increases in data volume due to the upgraded capabilities

ESnet

Summarized Findings:

Computational and Storage Requirements

- BES user facilities face three primary challenges in the coming years related to computation, storage, and networking:
  - The current storage capacities at BES light source facilities are often insufficient to handle the expected data volumes for future upgrades. The most data-intensive facilities are expected to have steady-state data rates approaching 10 PB/year, and peak data rates during burst acquisition of >1 TB/hour.
  - Access to computation is particularly limited for experimental and theory groups. When it cannot be used locally, it must be found either within the DOE ecosystem or at commercial providers.
  - Transfer speeds are currently sufficient, but will need to keep pace with data volumes and a growing need to access external computation.

ESnet

Summarized Findings:

Multifacility Computational Workflows

- Users at BES facilities often leverage the complete ecosystem of available DOE capabilities.
- The process of science often involves combining and analyzing data from multiple sources.
- As data rates and complexity continue to increase, sufficient networking connectivity, bandwidth, and reliability is required to connect measurement facilities, a computing facility, and user home institutions to enable effective data management and analysis.

ESnet

Summarized Findings:

Domestic Networking for Data Mobility

- BES user facilities that leverage DOE HPC resources will continue to expand in the coming years, placing heavy emphasis on ESnet connectivity to manage traffic across the country.
- DOE HPC facilities must keep pace with upgrades at BES user facilities, particularly the light sources, to ensure ample capacity to support the scientific mission.
- By the end of the decade, data aggregated at rates of multiple terabits per second (100–300 Gbps within three years and exceeding 1 Tbps in five years) may flow via ESnet from any of the DOE light sources to any of the DOE HPC facilities.

ESnet

Summarized Findings:

Emerging Needs

- Autonomous experiment steering at BES light source facilities will be crucial to future experiment success. To accomplish this usage pattern sufficient and reliable bandwidth is required, along with scalable and sustainable data storage resources, on-demand access to large-scale computing systems, transparent access to facilities and systems through federated identity and shared data protocols and software tools and infrastructure to facilitate the development of scientific data workflows.
- The rapid advancement in AI/ML algorithms, improved shared workflows, and the advent of exascale computational resources make it possible to create a physically informed virtual platform to perform experimentation: digital twins.

ESnet

Summarized Findings:
Emerging Needs (cont)

- A number of high-performance computation and quantum simulation users and collaborations currently use multiple facilities and locations to accomplish science goals that cross the DOE, commercial, and university complex.
- This work uses a variety of platforms for hybrid computing, and requires a strategy that tightly integrates classical and quantum resources for an overall multi-tiered approach to computing.

ESnet

Multi-Facility Workflows Are Key

- Many different case studies point to multi-facility workflows as important going forward
  - Light sources and HPC
  - Digital Twins
  - Experiment steering using computing
  - Quantum + Classical computing
- ESnet looks forward to collaborating with our partners in BES
  - Combining ESnet6 capabilities with BES facilities
  - Working together to improve scientific productivity using networking
  - Evolving the state of the art as the technology landscape changes

2022 BES Requirements Review: By the Numbers

- ***18*** Case Studies
- ***116*** Authors, Editors, and Participants
- ***17*** Organizations
  - 1 Federal Agency
  - 10 DOE Laboratories / Facilities
  - 6 Universities
- ***~1 Year*** Process (Dec 2021 through November 2022)
- ***~300*** Pages in Final Report
- ***~23*** 'Actions' that may require follow up

ESnet

Case Studies

- Advanced Light Source (ALS)
- Advanced Photon Source (APS)
- National Synchrotron Light Source II (NSLS-II)
- Linac Coherent Light Source (LCLS)
- Stanford Synchrotron Radiation Light Source (SSRL)
- High Flux Isotope Reactor (HFIR) and Spallation Neutron Source (SNS)
- Center for Functional Nanomaterials (CFN)
- Center for Integrated Nanotechnologies (CINT)
- Center for Nanoscale Materials (CNM)
- Center for Nanophase Materials Sciences (CNMS)
- The Molecular Foundry
- Autonomous Experiment Steering for BES Facilities
- BES Design and Development of Digital Twin Strategies
- Multifacility Experimentation and Analysis Workflows:
  - X-ray Light Source Perspective
  - Neutron Scattering Perspective
  - Nanoscale Science Research Center (NSRC) Perspective
- Use of the ESnet for Quantum Simulations of Materials and Molecules
- The Materials Project: Status and Future Directions

ESnet

Participating Orgs

- Federal
  - Department of Energy Office of Science

- DOE Labs & Facilities
  - Ames National Laboratory (Ameslab)
  - Argonne National Laboratory (ANL)
  - Brookhaven National Laboratory (BNL)
  - Energy Sciences Network (ESnet)
  - Lawrence Berkeley National Laboratory (LBNL)
  - Los Alamos National Lab (LANL)
  - The National Energy Research Scientific Computing Center (NERSC)
  - Oak Ridge National Laboratory (ORNL)
  - Sandia National Laboratories (Sandia)
  - SLAC National Accelerator Laboratory (SLAC)

- Universities
  - Dartmouth College
  - Howard University
  - Iowa State University
  - University of Chicago
  - University of Tennessee
  - University of Texas

ESnet

ESnet
ENERGY SCIENCES NETWORK

Thanks!

Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory

Eli Dart, Science Engagement Engineer
Jason Zurawski, Science Engagement Engineer

December 7, 2022

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB