U.S. DEPARTMENT of ENERGY

BERKELEY LAB

ESnet
ENERGY SCIENCES NETWORK

PoV: Streaming vs. Data Copy

NREN perspective

Eli Dart

Network Engineer, Science Engagement

dart@es.net

TNC 2026

Helsinki, FI

8 June 2026

Our World Is Data-Centric

- Many, many scientific discoveries are impossible without data analysis
- Modern science relies on computing, data, storage, all interconnected using fast networks
- ESnet thinks of itself as a “data circulatory system”
  - This holds for NRENs generally

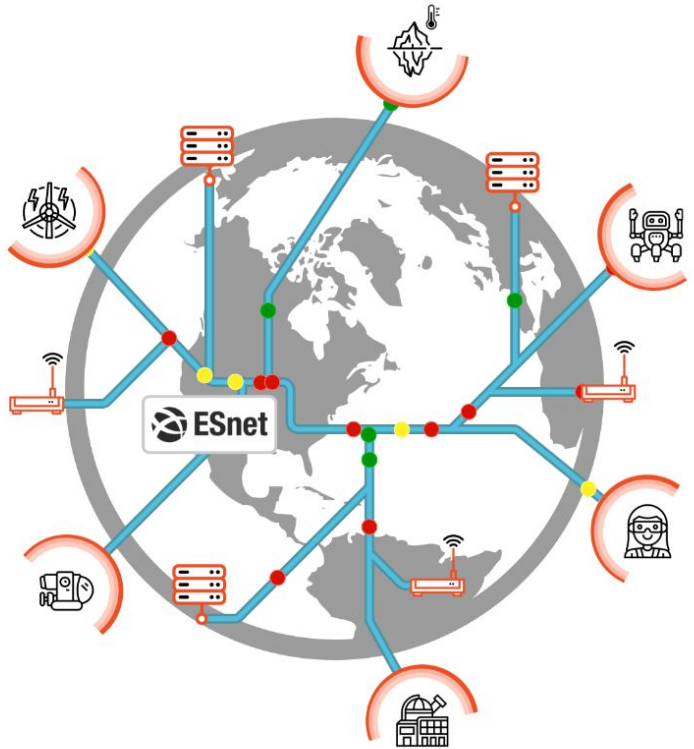

ESnet

Data World Relies On NRENs

- NRENs support the scientific research ecosystem.
- Interconnect national labs user facilities, research institutions, universities, globally
- Provide reliable, high-performance connectivity to global research collaborations, the Cloud, etc.

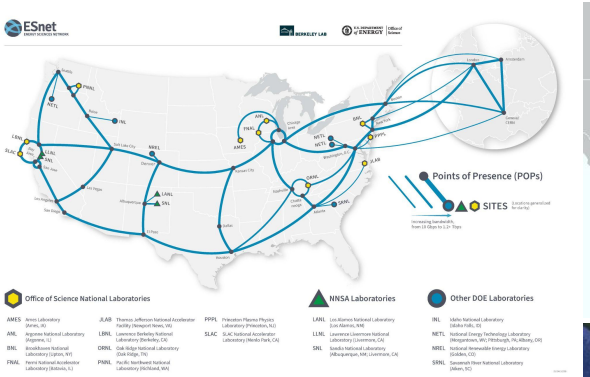

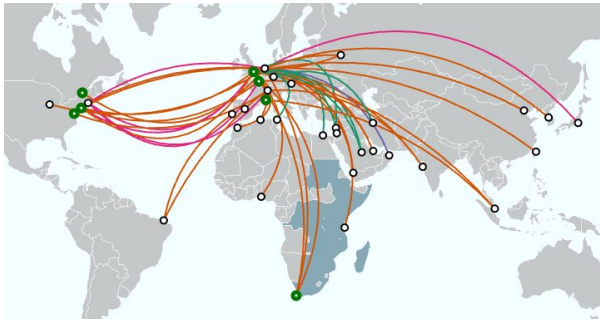

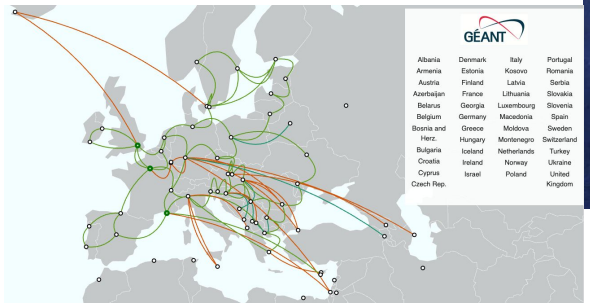

ESnet

A Note About Data Locality

- I hear a lot of people say “don’t move the data”
  - “Just move the compute to the data!”
- Of course! Do that if you can. But there are often difficulties
  - Open access data, closed access computing
    - Code execution is often tightly controlled
    - Scientist often doesn’t have permission to run analysis next to the existing data store → need to move the data
  - Specific hardware required
    - Specialized systems, computing scale, other requirements
    - Move the data to the specific system
  - I don’t have to pay to run on my local system
    - More freedom, creativity if I run there
    - But, I have to move the data to my cluster
- Therefore, we must support data movement in the general case
- **“Data Movement” takes multiple forms**

File Copy Workflow

- Data starts and ends on nonvolatile storage
- Doesn’t matter if it’s HPC, Cloud, DTNs, Laptop, USB stick, etc at each end
- Well understood
- Well hardened tools

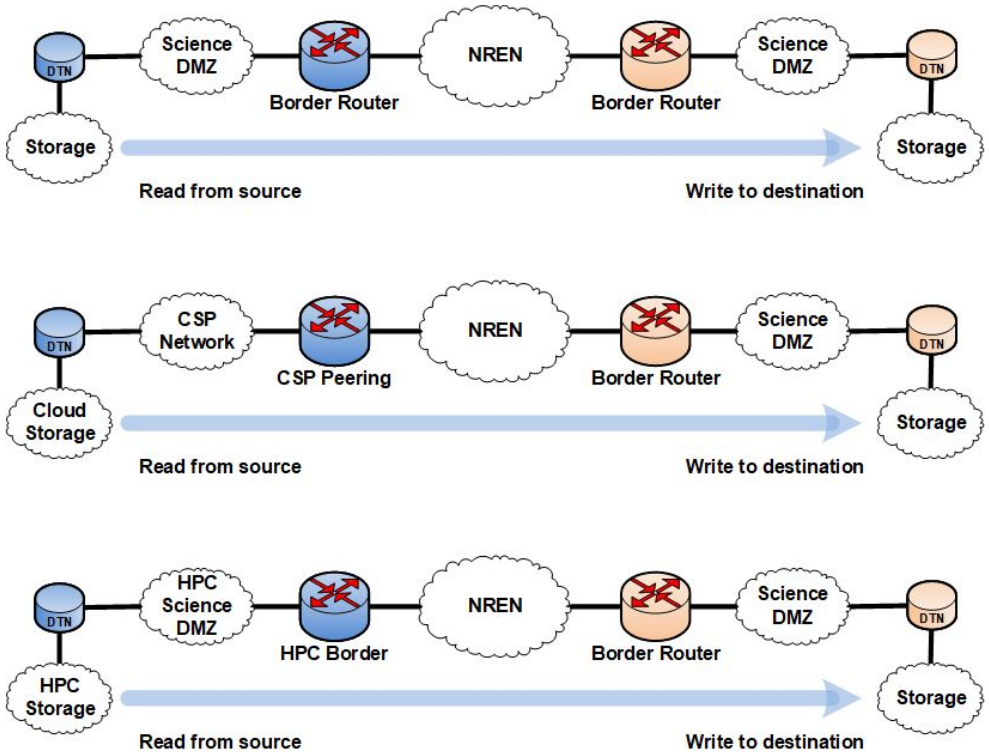

ESnet

Streaming Workflow

- Wide variation among workflows - some examples here
- At least one end is running compute - online
- Key point: **no requirement to copy data to destination storage first**
- Critical dependence on network - performance, reliability, resilience
- Wide variability in toolset

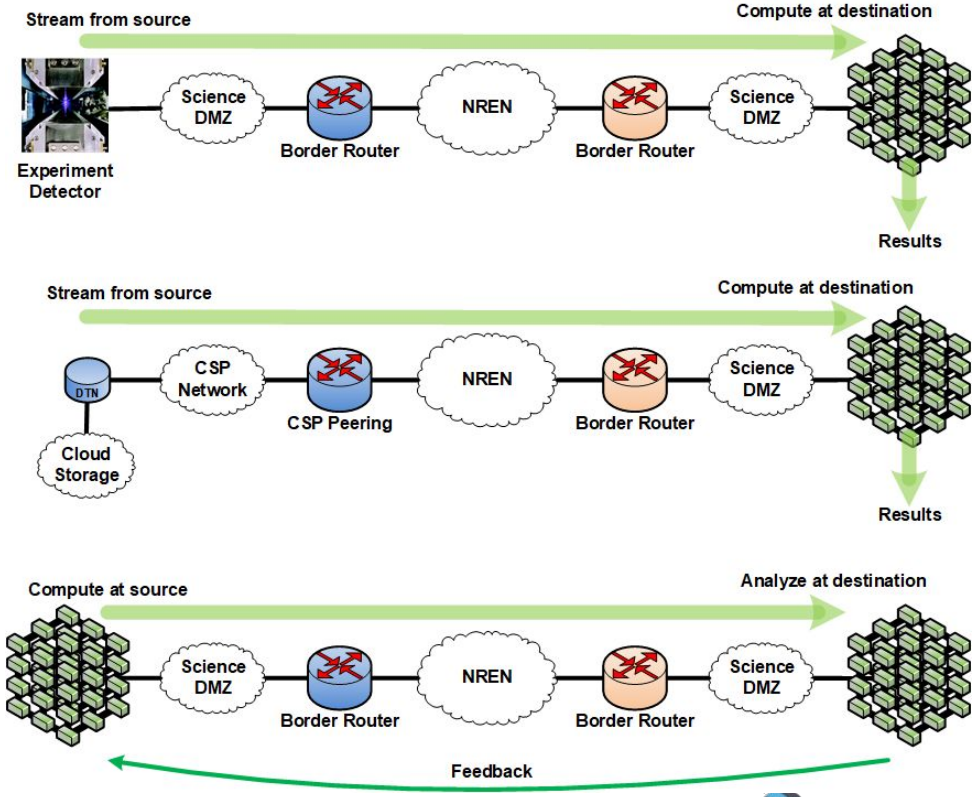

Copy Example: Large Scale Replication From Cloud

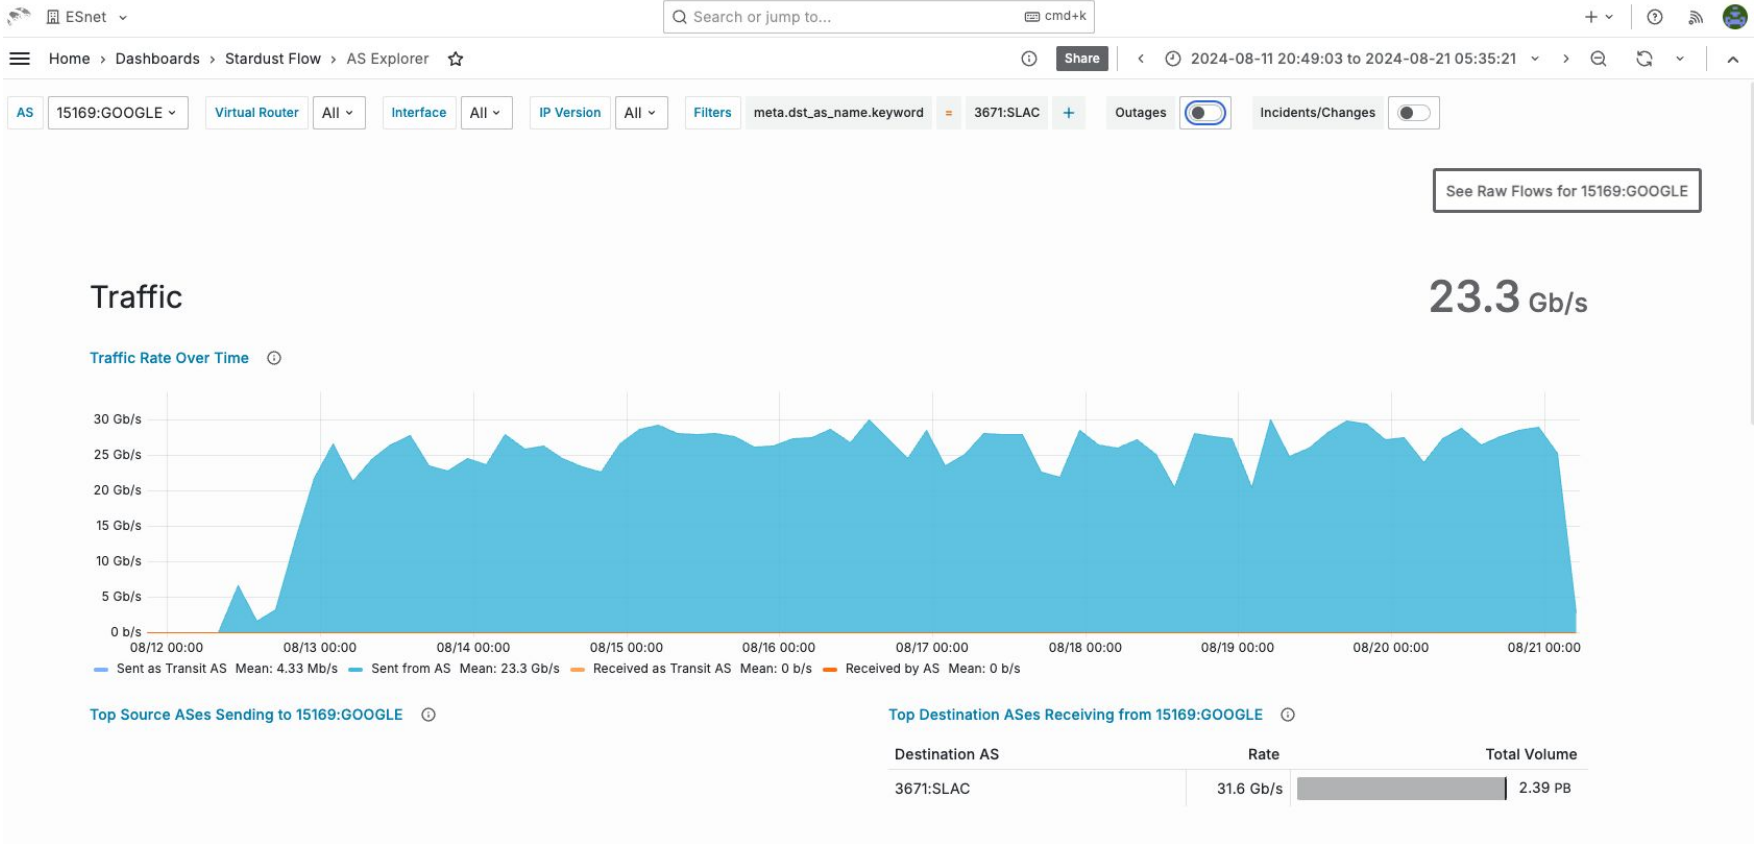

Copy Example: Large Scale Replication From Cloud

- Rubin/LSST project at SLAC
- Large scale data replication
  - 250M files
  - Over 2.3PB
  - Approx. 8.5 days
- Six DTNs at SLAC
  - K8 environment
  - Google Cloud Storage Transfer application
- Traffic traversed three Squid proxies at SLAC

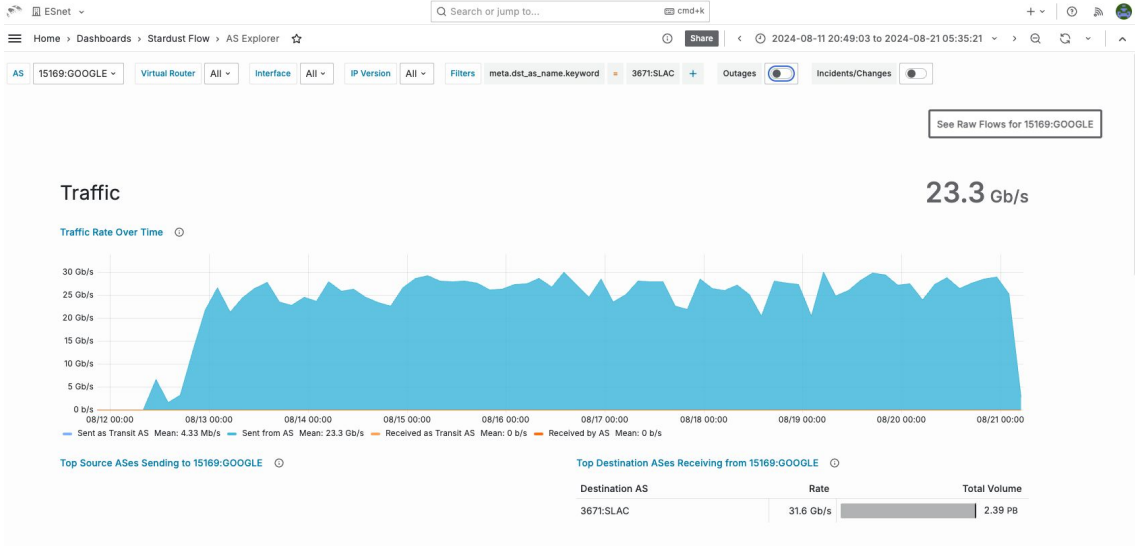

ESnet

Google to SLAC Replication - Key Points

- Petascale data transfer from cloud works!
- Small number of DTNs can achieve good performance
- Containers are viable for high performance applications, including I/O intensive

File Transfer: Discussion

- When the infrastructure is already in place, users just use it
  - Nobody has to ask permission
  - No complex config for users to do
  - No need to debug the end to end path before using it
- Users can concentrate on the workflow
- Simple protocols, high performance
  - Simple policy: Support DTN driven data transfers
  - Simple protocols: Globus, cloud-native object access, etc.
  - Safe to operate when secured with performant technologies, in keeping with Science DMZ
  - Well understood by the science community
- Performance determined by storage on both sides

Storage Becoming A Limitation

- In many cases, storage I/O is now the bottleneck
- Answer: **streaming**
- One or both ends of the data transfer is volatile
  - Instrument
  - Memory
- Removes storage limitations from the workflow

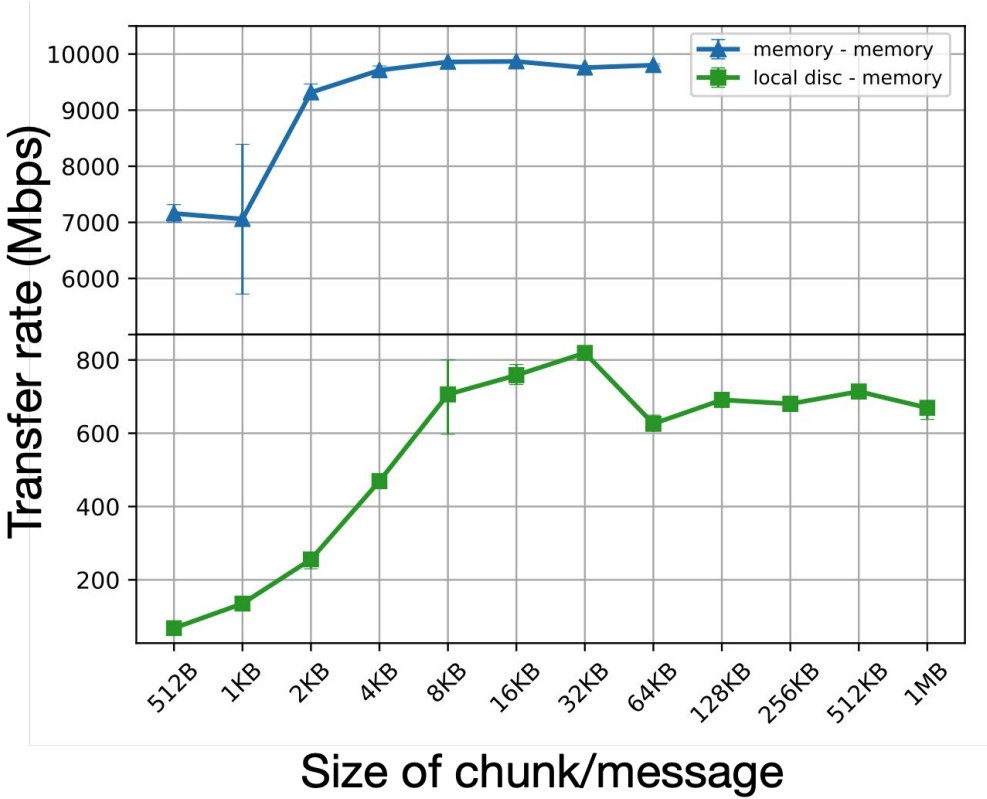

Image: Raj Kettimuthu, ANL

Stream From Site Storage To Cloud

- Workflow reads from site storage into cloud analysis job
- No local storage in the cloud environment
  - Compute job analyzes the data and then drops it
  - Need the data again? Fetch it again!
  - No cloud storage costs
- High performance perimeter required for success
  - Packet loss or other network issues reduce compute performance
  - Easy to do with Science DMZ

Training and Inference in Commercial Cloud

- DUNE collaboration
- ML based data reduction
- Near-real-time inferencing classifiers supporting neutrino physics workflows.
- Hybrid workflow
  - Data and CPU computing at FNAL
  - GPUs in Google
- Impossible without high performance perimeter
- Works great if you do have high performance perimeter

Data From Lab to Cloud For Inference

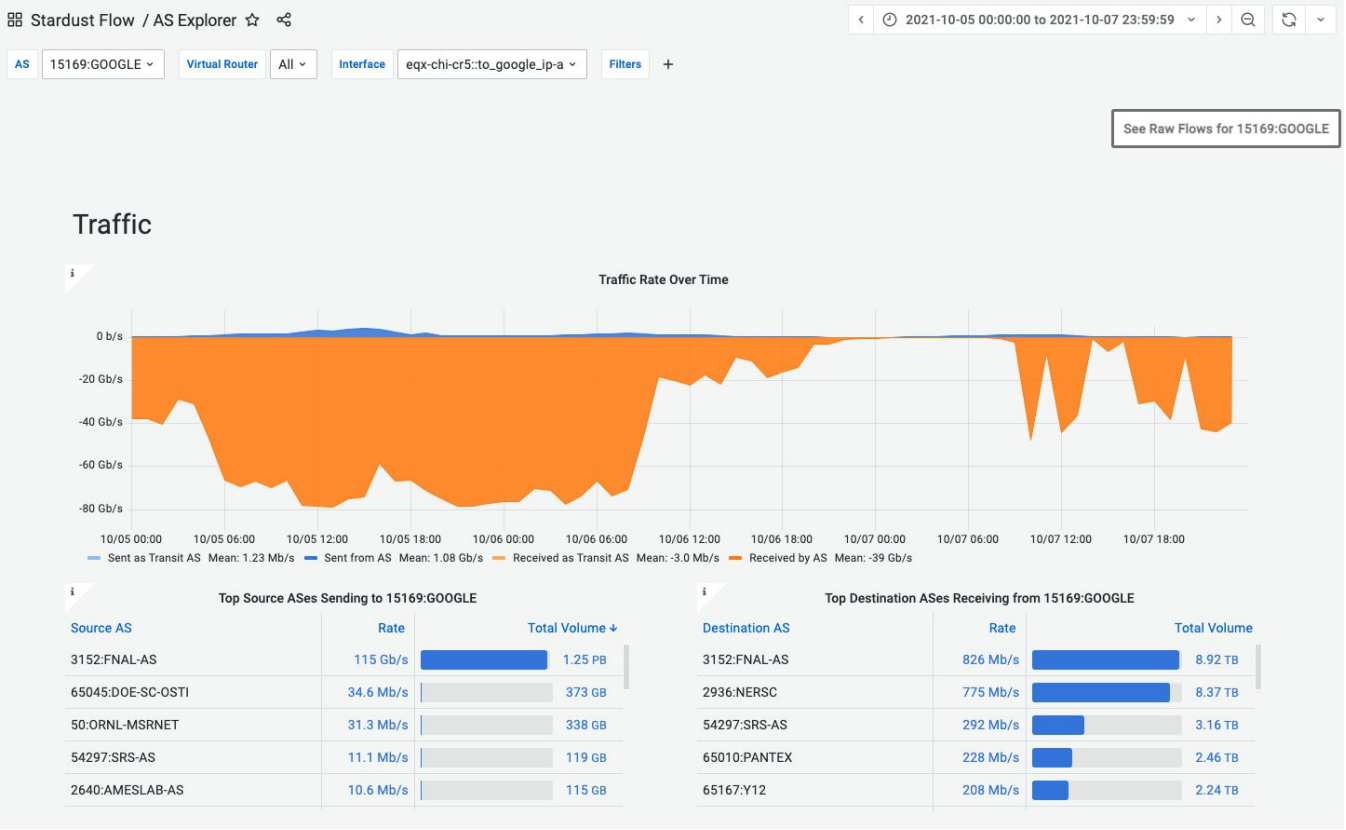

- ~80Gbps for hours at a time
- Over 1PB moved in ~30 hours
- No need to replicate the data first (no file transfer workflow)
- No in-cloud storage costs

ESnet

Site to Cloud Streaming: Discussion

- Data storage and servers are on infrastructure owned by the community/collaboration
  - Storage systems with DTNs
  - POSIX or object store, doesn't matter
- Analysis/compute clients in the cloud fetch the data directly
  - No copying to cloud storage first
  - Just read the data in from remote storage, analyze it
- Site perimeter design: DTNs for outbound data
  - Simple security policy (address and port filters)
  - Safe to operate with performant security technologies
- Science workflow requires high performance
  - Scaling enterprise firewalls to support this is difficult and expensive
  - No added benefit from enterprise firewalls (just address/port)

Stream From Cloud To HPC

- These workflows go in the other direction
- Data comes from the cloud, across the perimeter, and into a running analysis job on a supercomputer
- Can be many many nodes on both sides
- Bandwidth can get quite large
- Not limited by HPC storage I/O
  - No need to wait for data replication
  - Workflow can go faster than the DTN cluster

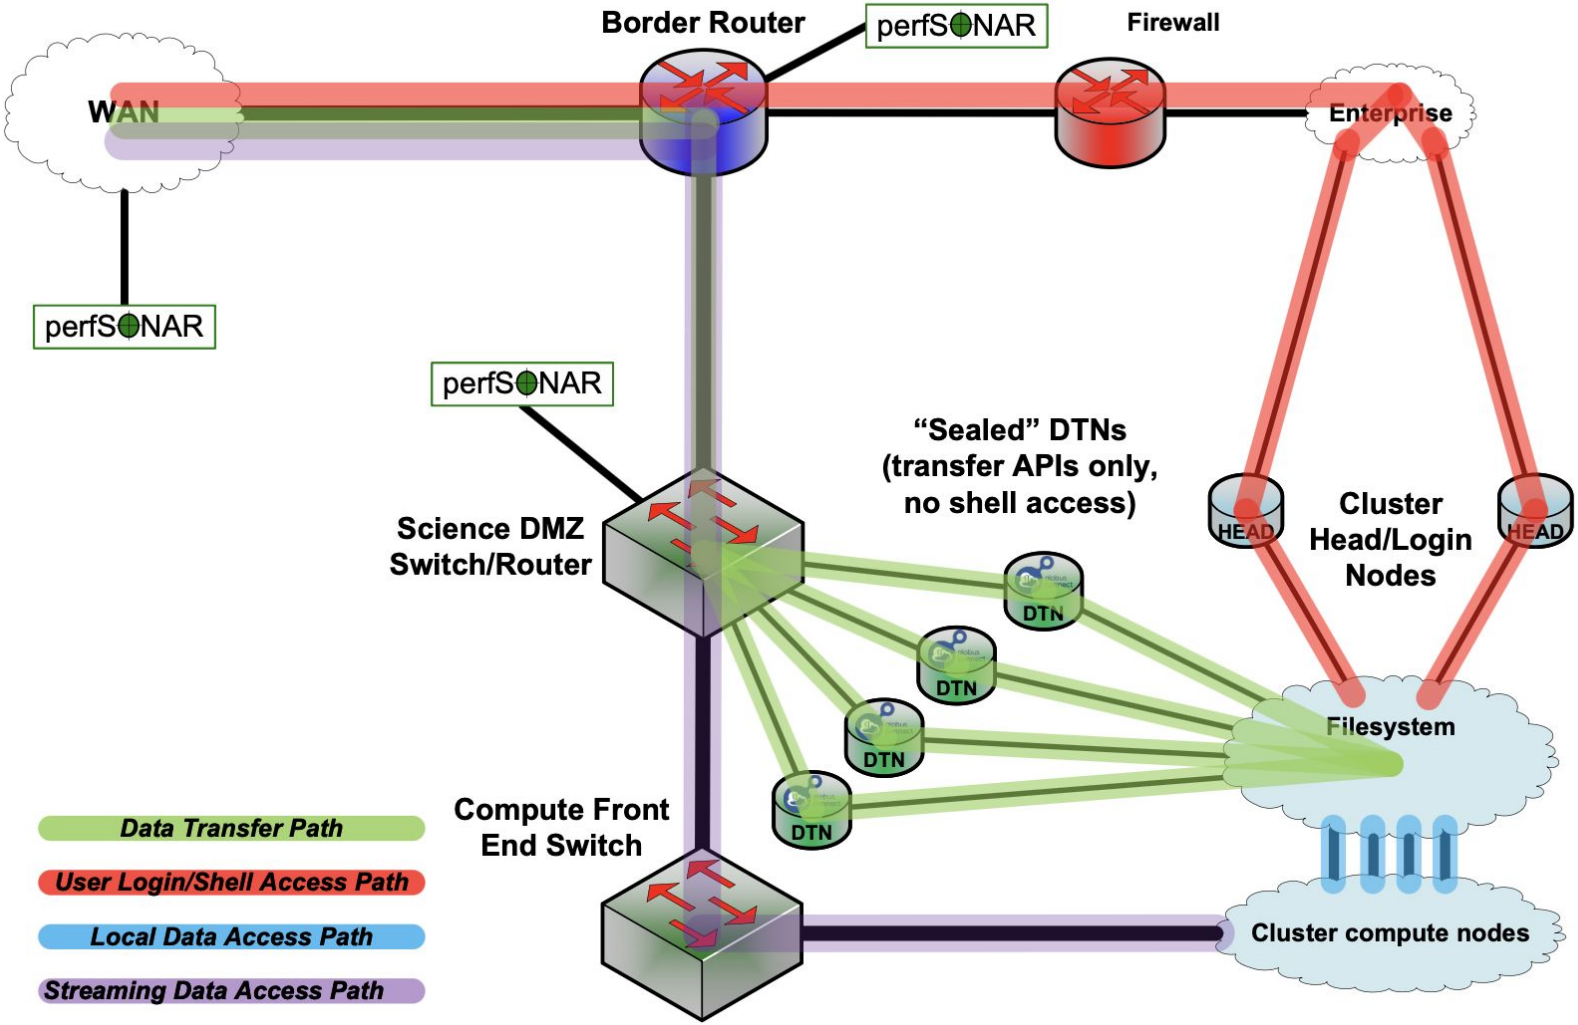
Border Router
perfSONAR
Firewall
WAN
Enterprise
perfSONAR
perfSONAR
"Sealed" DTNs
(transfer APIs only,
no shell access)
Science DMZ
Switch/Router
HEAD
Cluster
Head/Login
Nodes
HEAD
DTN
DTN
DTN
DTN
Filesystem
Compute Front
End Switch
Cluster compute nodes
Data Transfer Path
User Login/Shell Access Path
Local Data Access Path
Streaming Data Access Path

Direct Transfer to Compute: Amazon to NERSC

- Data set in commercial cloud
- Weather forecasting code
- No replication before job launch
- Direct ingest by compute job takes pressure off the HPC facility storage system

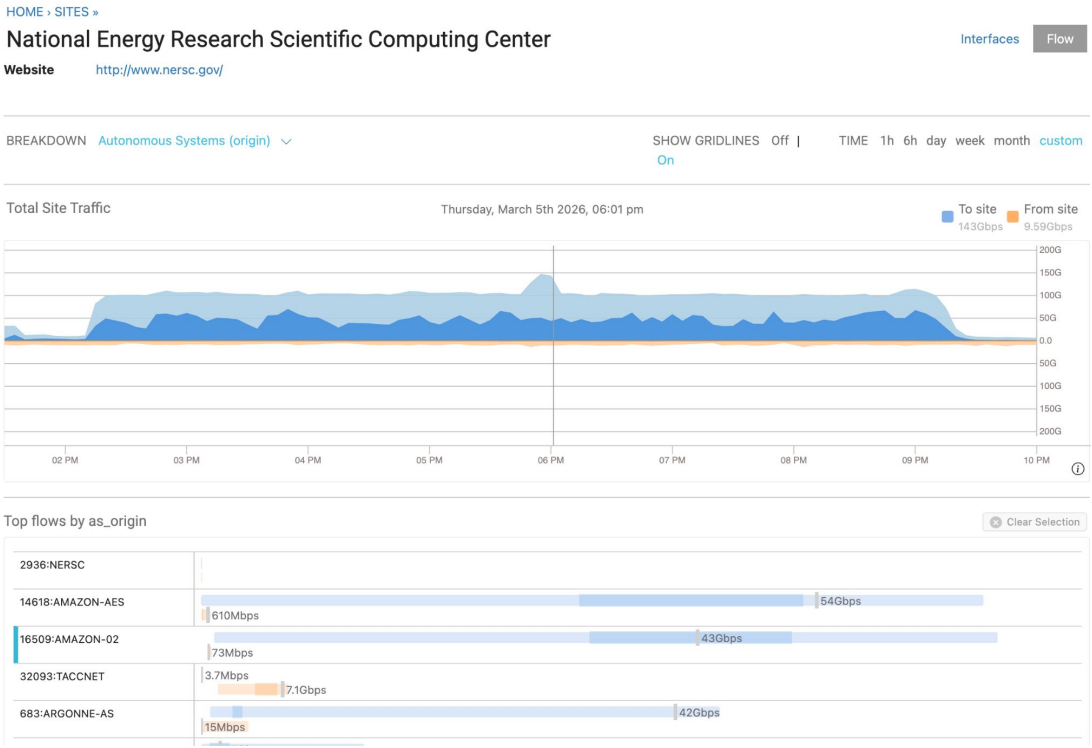

ESnet

Many cloud servers to a few compute nodes

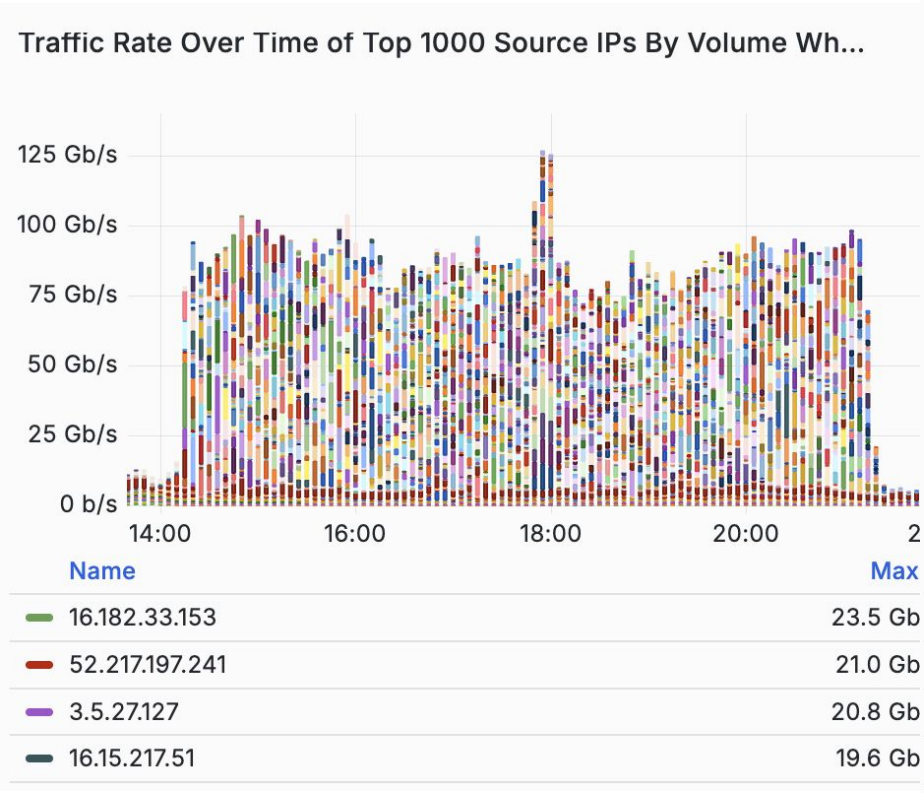

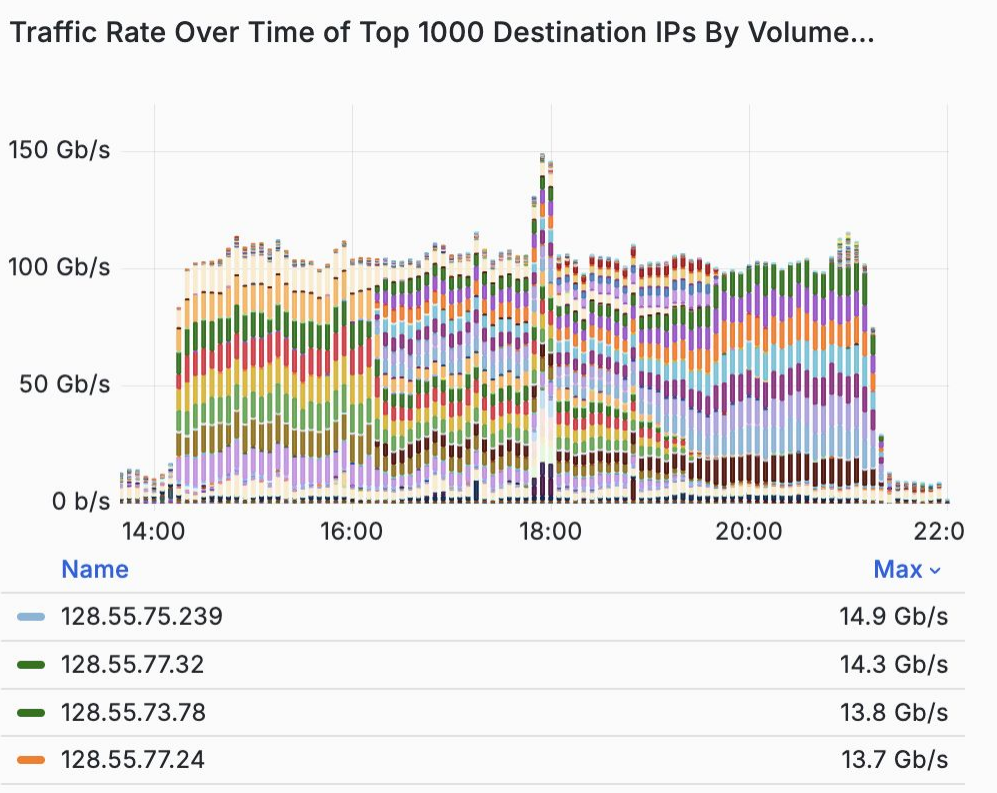

ESnet

Direct Transfer to Compute: Google to NERSC

- Large (over 5PB) public data set in commercial cloud
  - ECMWF ERA5
  - AI-ready
- ML-accelerated earth system model application
- Direct data ingest into running code

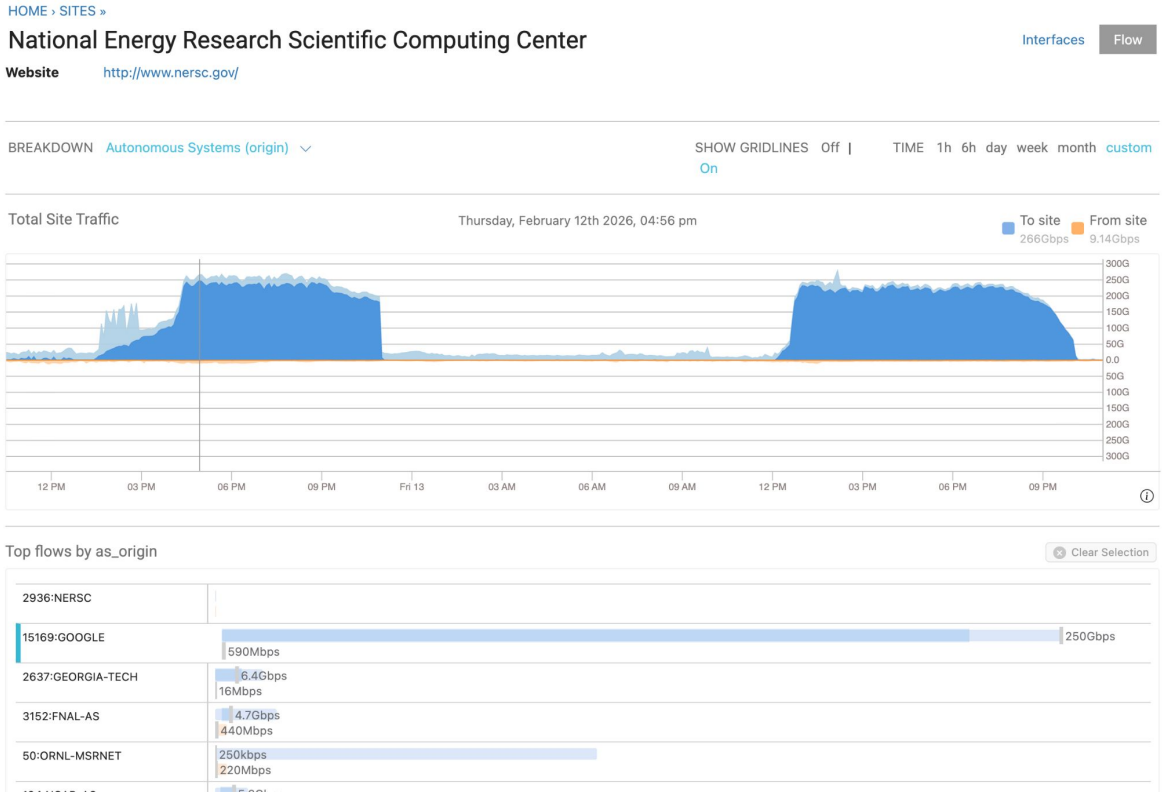

ESnet

Flow statistics: massively parallel

- Host breakdown:
  - Many, many parallel streams in these transfers
- Many cloud hosts delivering over 10Gbps
- User was able to just launch the job and move forward.
  - No special setup
  - No debug before use

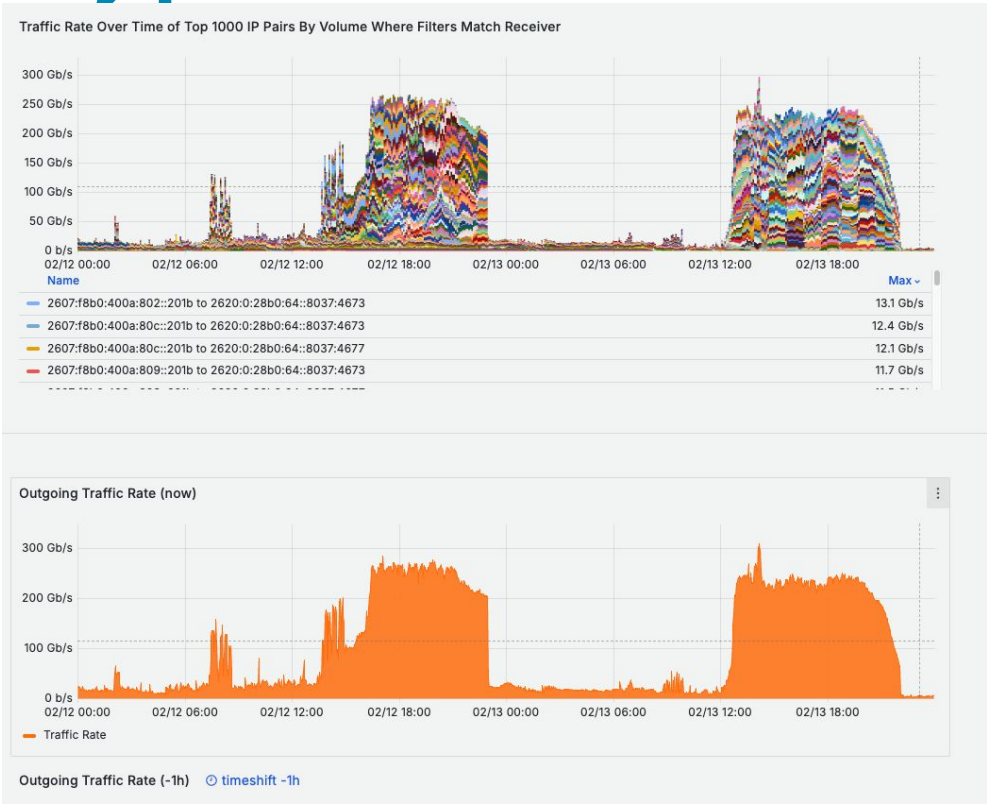

IPv6 Works!

- A lot of this transfer occurred over IPv6
- It was already set up and configured - no need for user to care
- IPv6 is real now, even in commercial cloud environments

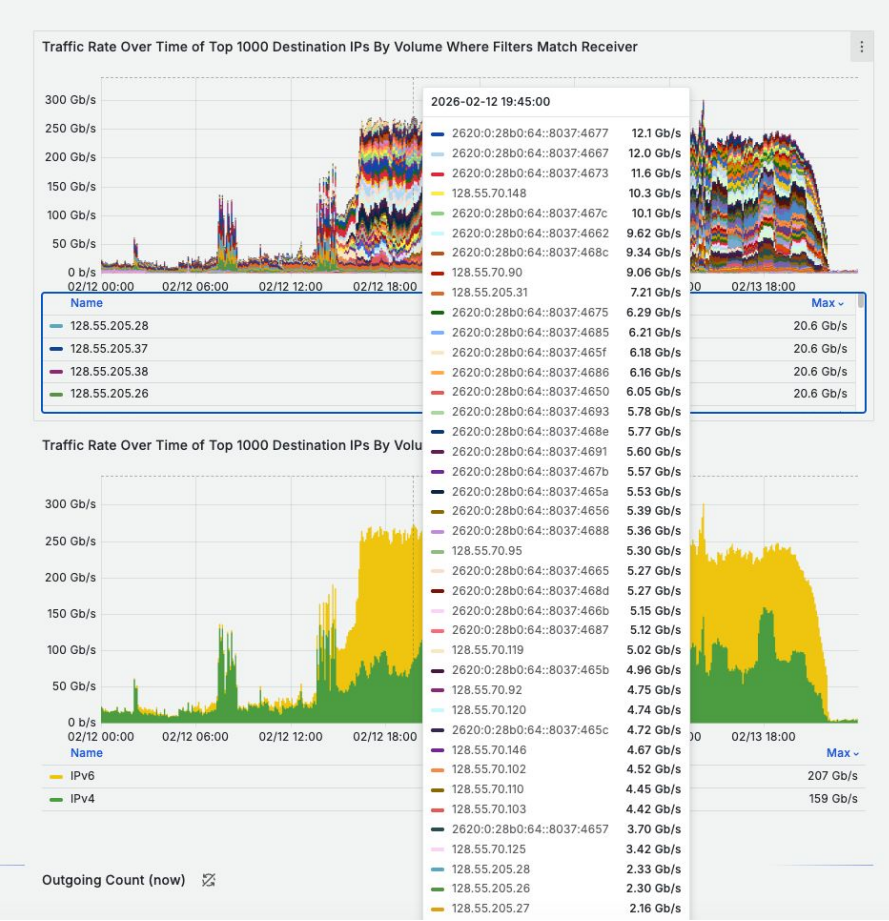

ESnet

Streaming From Detector to HPC: EJFAT

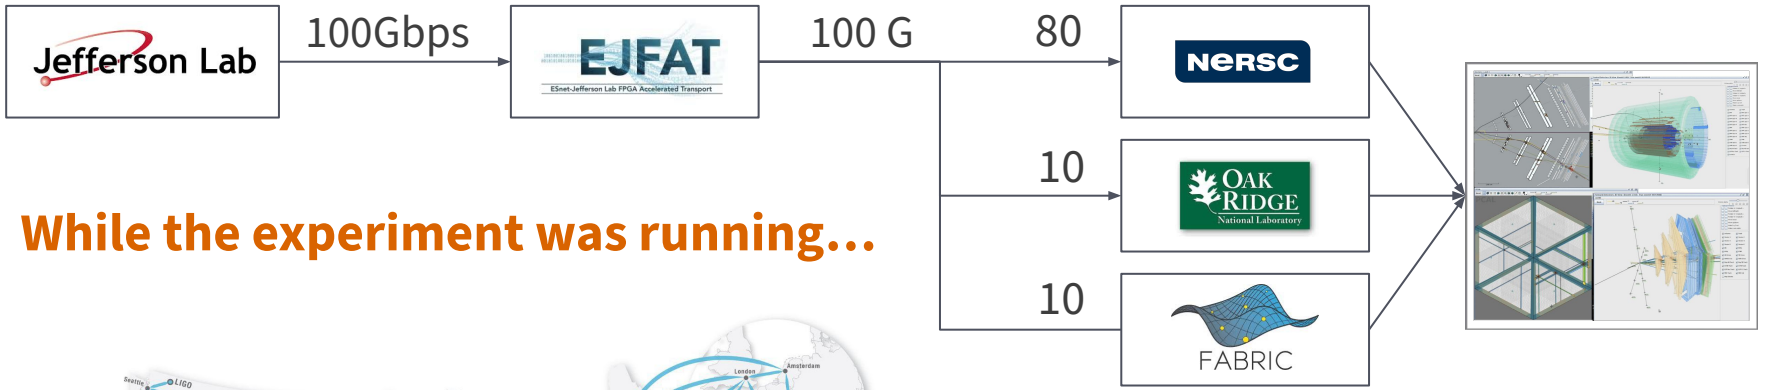

While the experiment was running…

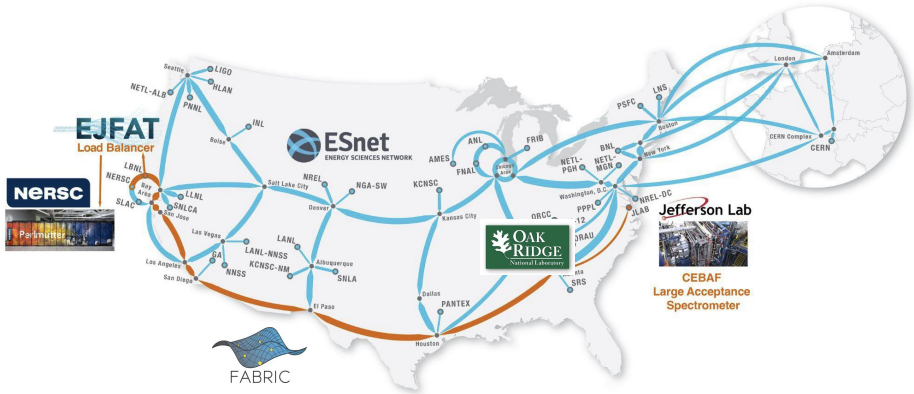

…We added compute nodes at ORNL
…And more compute nodes in FABRIC

With NO event loss.

ESnet

LHC/US-CMS: Strategic Location of Storage

- Future storage resources concentrated at Tier1 site
- Future compute resources concentrated at Tier2 sites
- Rely on network to interconnect the two

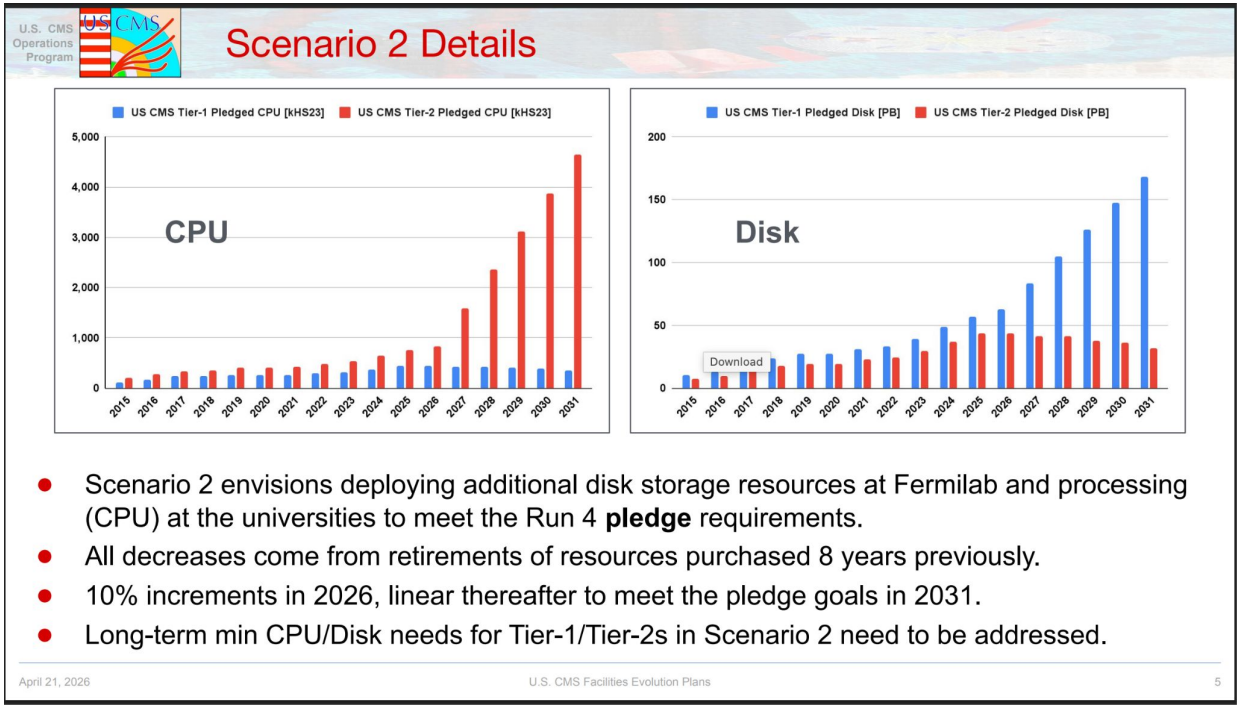

Evolution Of Data-Intensive Science

- Science DMZ enabled high performance file transfer
  - DTNs
  - Performant tools
- Now, we are running into the limitations of storage
  - File transfer before job launch increases friction
  - Storage layer is slow and expensive
- New workflows drive ever higher performance across network boundaries
- But this only works if the whole path is performant
  - Sites, Science DMZs, etc.
  - NRENs in the middle

In conclusion – ESnet's vision:

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

U.S. DEPARTMENT of ENERGY

BERKELEY LAB

ESnet
ENERGY SCIENCES NETWORK

Thanks!

Eli Dart

dart@es.net

https://my.es.net/

https://www.es.net/

https://fasterdata.es.net/