



U.S. DEPARTMENT
of ENERGY

ESnet Update

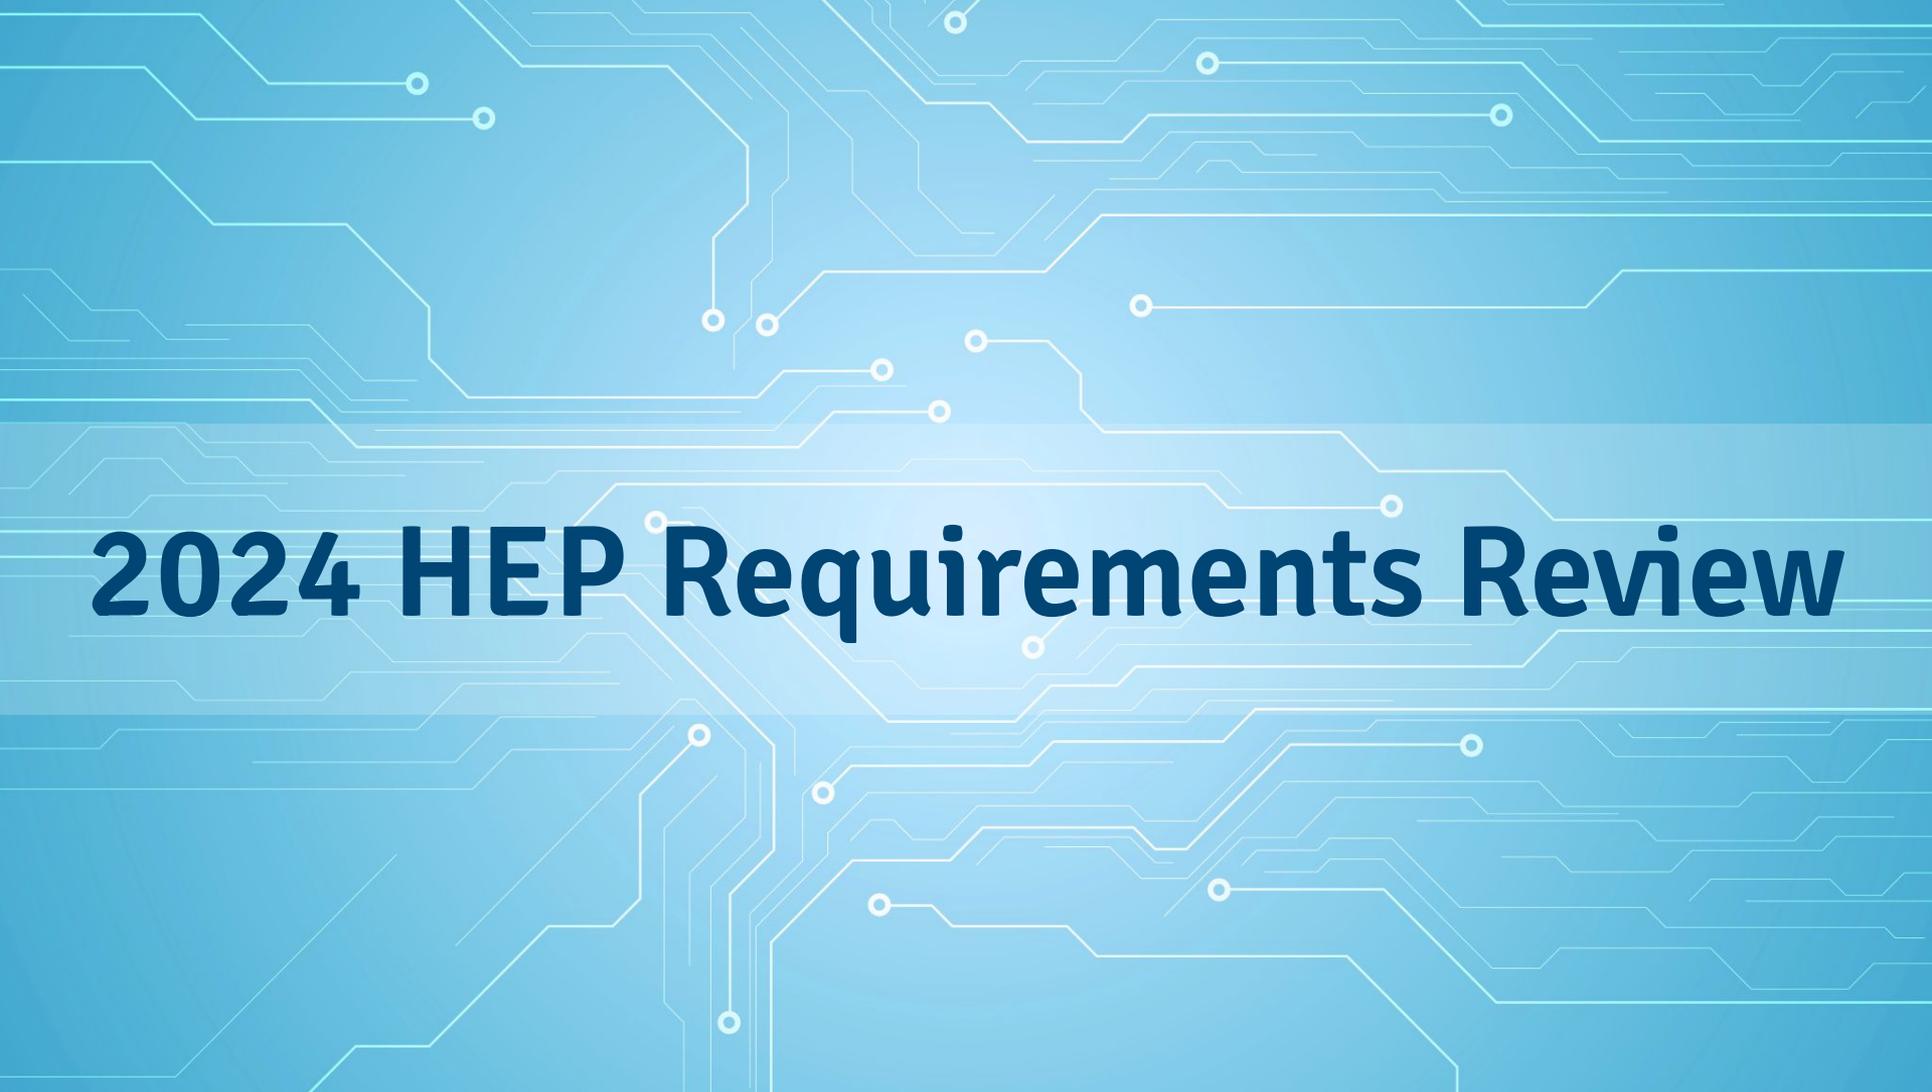


Kate Robinson
ESnet Network Engineering

LHCONE/LHCOPN 55
Karlsruhe, DE
2025-10-07

Outline

- Networking Engineering Updates
- Trans-Atlantic Updates
- Mini Data Challenge 08/2025



2024 HEP Requirements Review

HEP Requirements Review

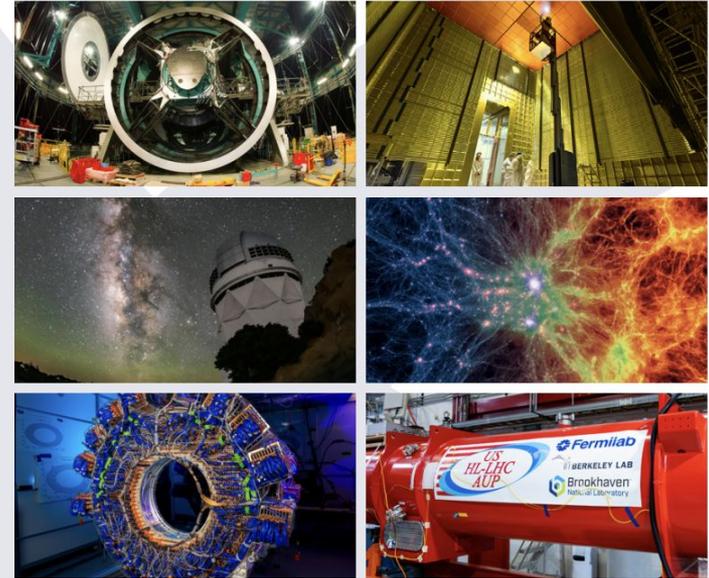
Thank you to everyone who participated in the 2024 HEP Requirements Review. The final report is complete and has been published.

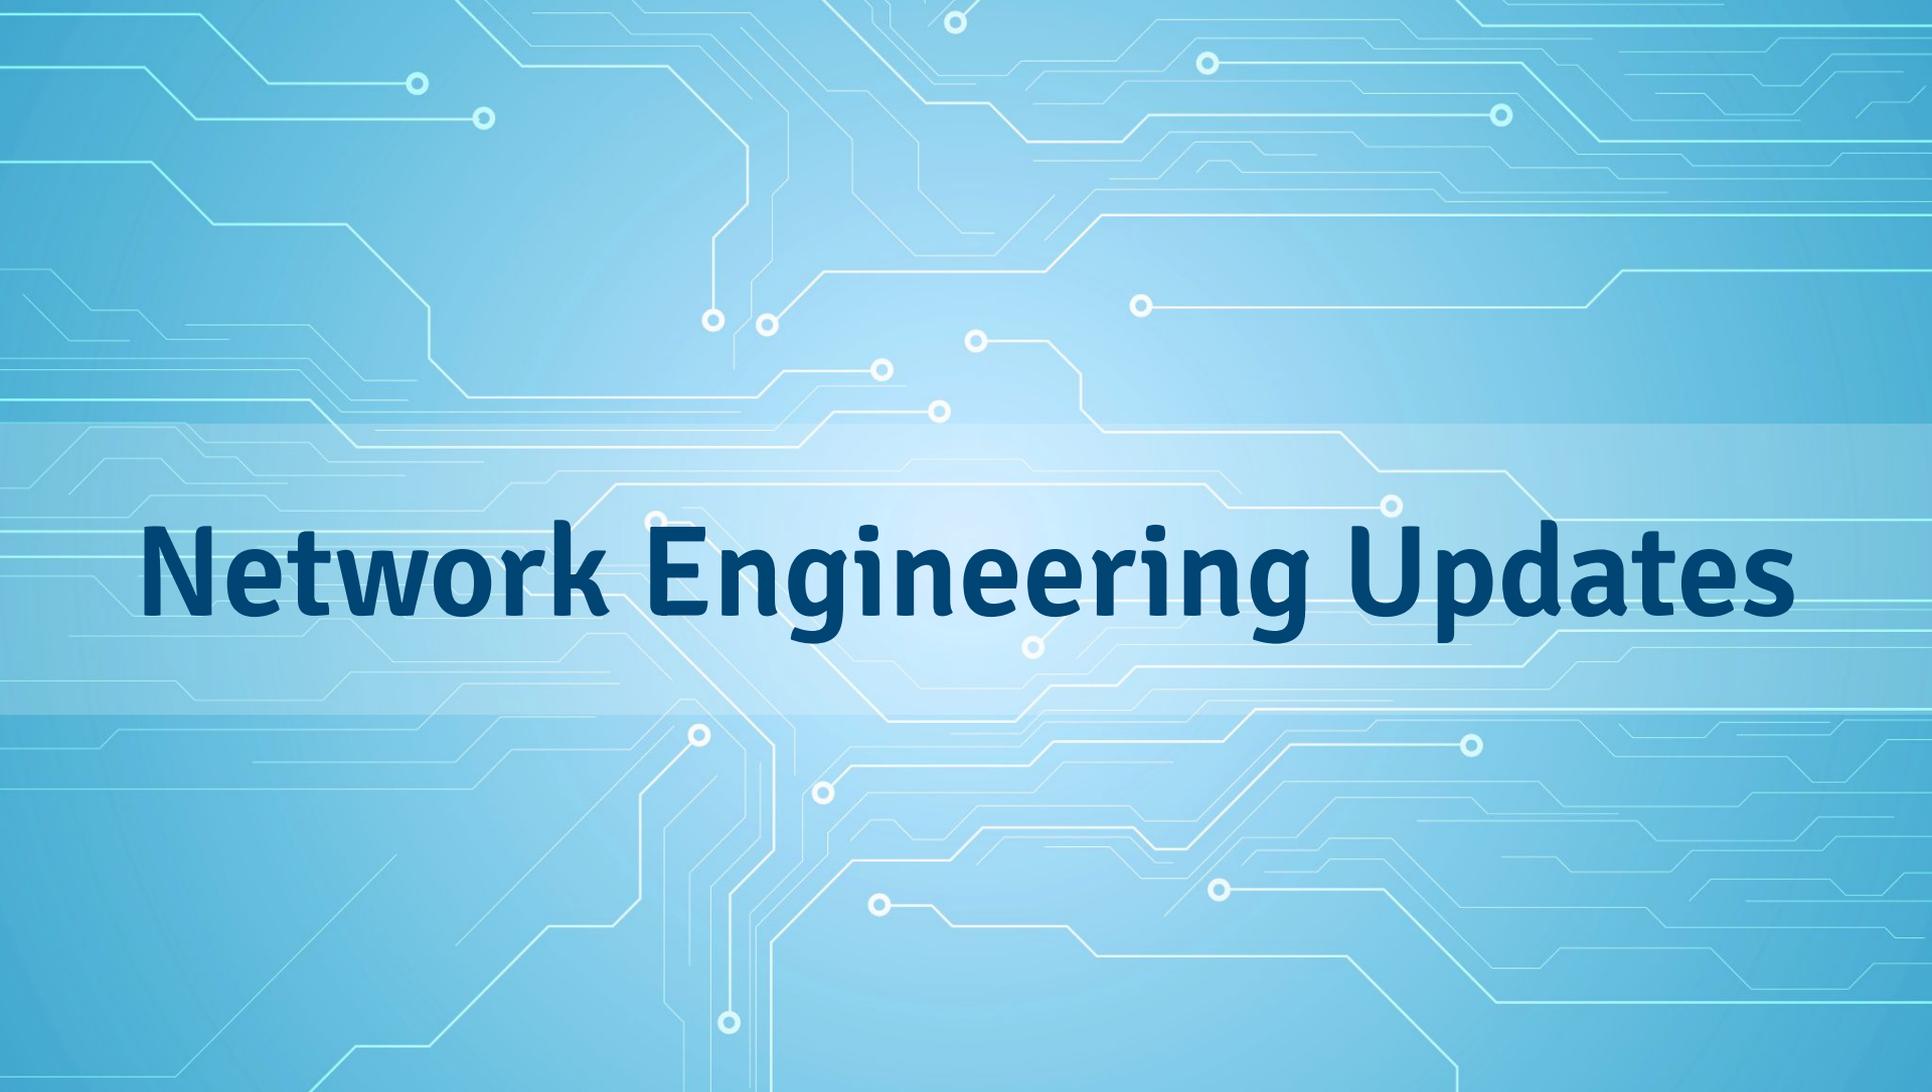
[High Energy Physics Network Requirements Review: Final Report, July 2024–December 2024 \(Technical Report\) | OSTI.GOV](#)



High Energy Physics Network Requirements Review: Final Report

July 2024–December 2024



The image features a blue background with a complex, white circuit-like pattern of lines and nodes, resembling a network diagram or a stylized circuit board. The lines are of varying thickness and form a dense, interconnected web. Small white circles are placed at various points along the lines, representing nodes or data points. The overall aesthetic is clean, modern, and technical.

Network Engineering Updates

OLS Backbone upgrades to 1.2T

Filling in gaps that were less than 1.2T (3x400G)

Completed:

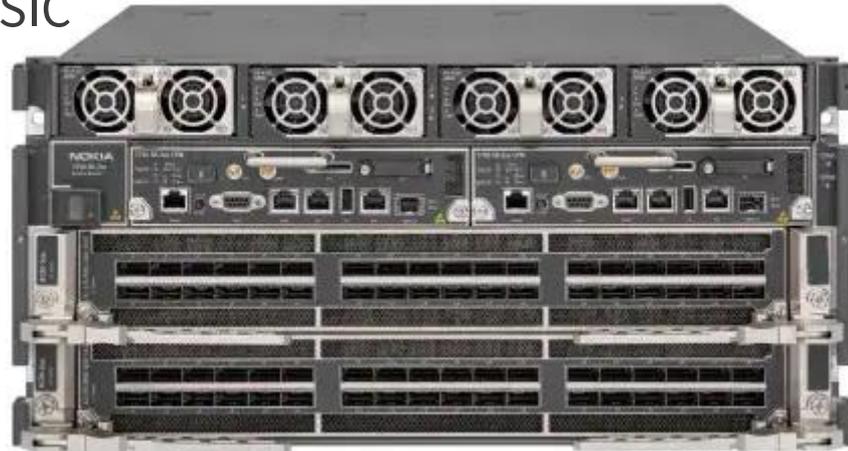
- ATLA-HOUS
- BOST-EQXCH2
- CHIC-KANS
- DENV-ALBQ-ELPA
- KANS-EQXDA3-HOUS
- ATLA-HOUS
- SALT-LASV
- NEWY32AOA-WASH
- CERN-EQXLD8-EQXAM3-CERN*

Still to do:

- BOST-NEWY325HUDSON
- SACR-SALT-DENV
- SACR-SEAT
- SALT-BOIS
- LOSA-SAND

SR-2se Upgrades

- ESnet6 SR-2s (large config) platform
 - Deployed circa 2021
 - Supports max 24 ports 400GE
 - Not enough port density at some PoPs
- As-needed, replacing with SR-2se
 - Next-gen FP5 packet forwarding ASIC
 - 72 port 400GE
- Completed:
 - CHIC, EQXCH2, STAR
 - NEWY32AOA, NEWY32HUDSON
- Planned:
 - WASH, EQXLD8



Nokia OS Upgrades

- Starting process now to upgrade to TiMOS 24.10R5 network-wide
- Motivations:
 - Better support for coherent pluggable optics
 - Bugfix to VRF next-hop rewriting that will allow us to remove leaking between IPT and Base
 - IS-IS multi-topology enhancements will allow us to signal LSPs with IPv6 addresses
- Impacts:
 - In general, ISSU should be transparent
 - Exception: 2se platforms require reboot (30-60 minutes downtime)



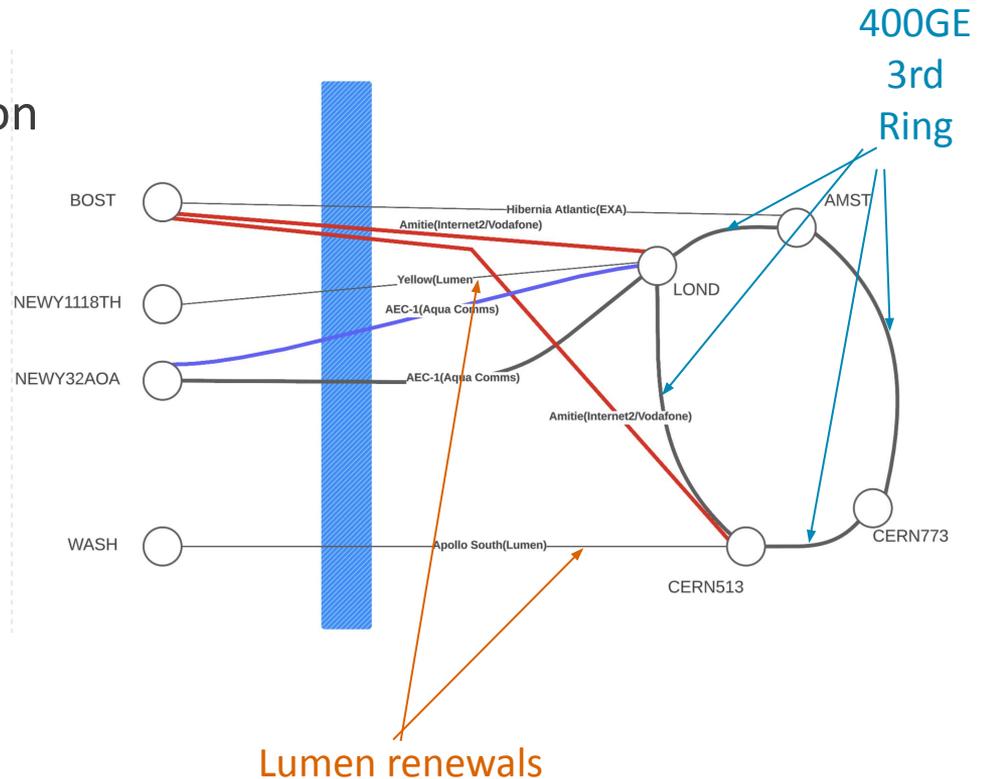
Trans-Atlantic Update

ESnet Trans-Atlantic Strategy

- Analysis from the Experiments' use cases in ESnet's 2020 HEP Requirements Review set multi-year plans into motion
- Components now coming into view
 - Procuring & lighting long-term use of spectrum
 - Diversity of cable routes
 - Partnerships between NRENs
 - Upgrading some lit circuits to 400G
- ESnet is on track to meet and exceed the initial needs of the HL-LHC program by 2027 (3.2 Tbit/sec)
- Currently at 2.6 Tb capacity

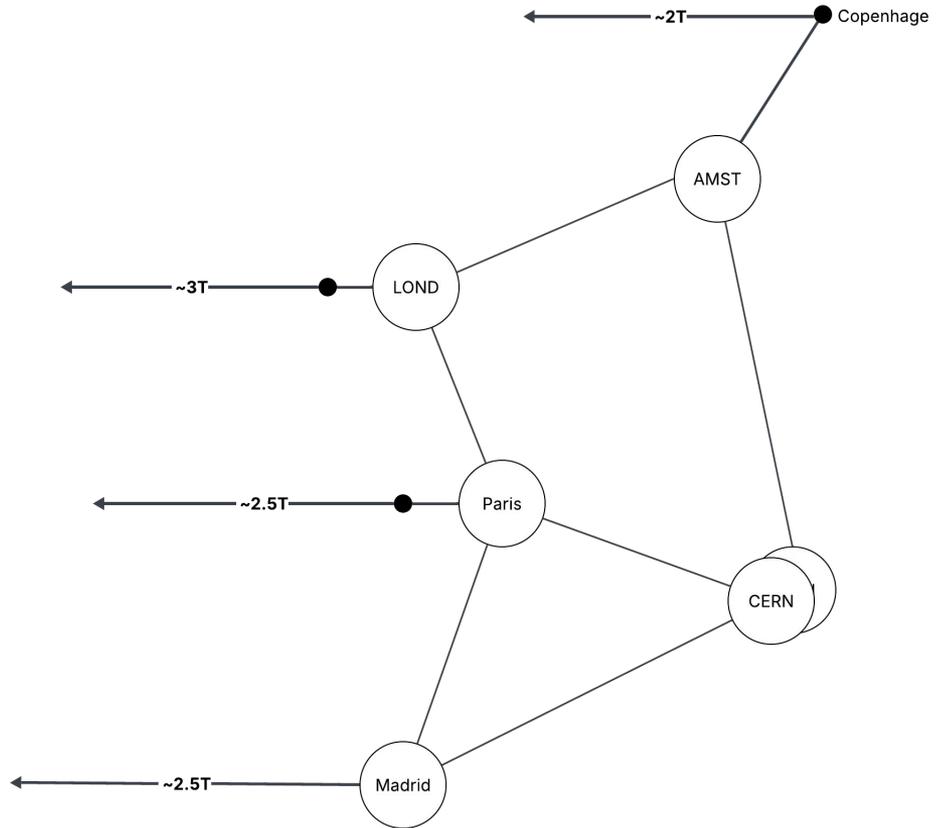
Transatlantic Updates

- AEC-1 Spectrum WL6 Upgrade
- Dunant Spectrum Implementation
- **Lumen 2x100GE Replacement**
 - Term ends 10/2024
 - Replace with GraceHopper 325-EQXLD8
- **EU 3rd 400GE Ring**
 - GÉANT lit services
 - Completed



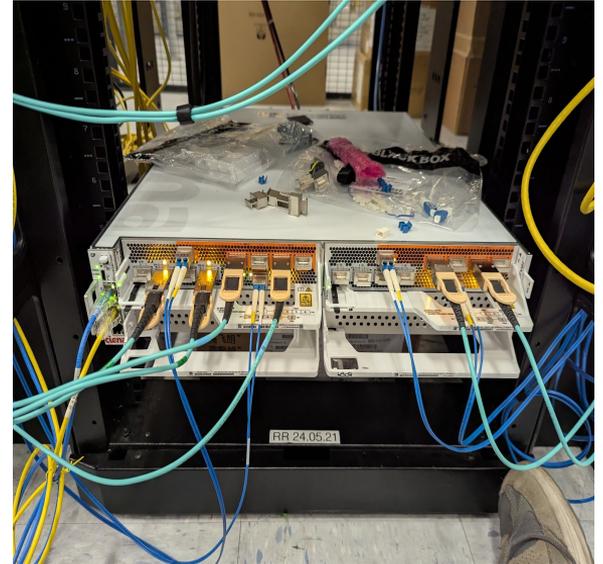
Spectrum on Transatlantic Cables

- NYC-EQXLD8 on AEC-1 (3x400G)
- DC4-Paris on Dunant by Q1 2026
- NYC-EQXAM3 (GEANT) mid 2026
- ATLA-Madrid (GEANT) end 2026?



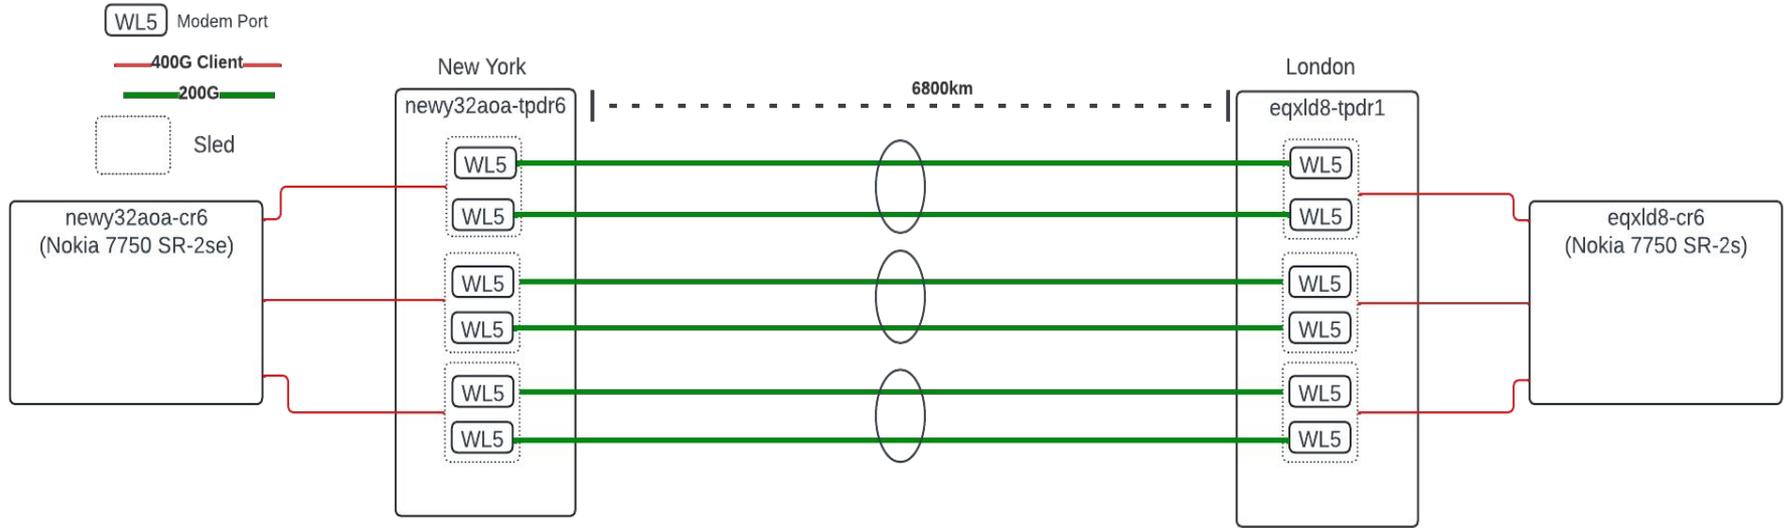
AEC-1 Spectrum WL6 Upgrade

- 1.1Thz Spectrum IRU
NEWY32AOA–Dublin–EQXLD8
- Initially continue to express through Dublin
 - 3x800G line capacity, 200ghz per wave
 - 6x400G client capacity available
 - Wider signal, more capacity, same power
- US side installed, awaiting LD8 delivery
- Planning larger footprint in EQXLD8/Dublin



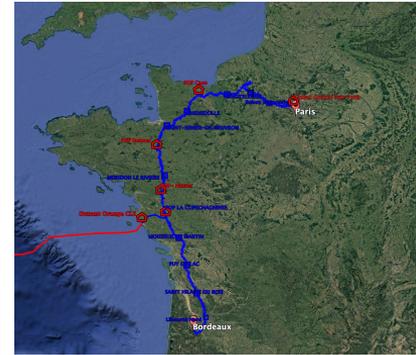
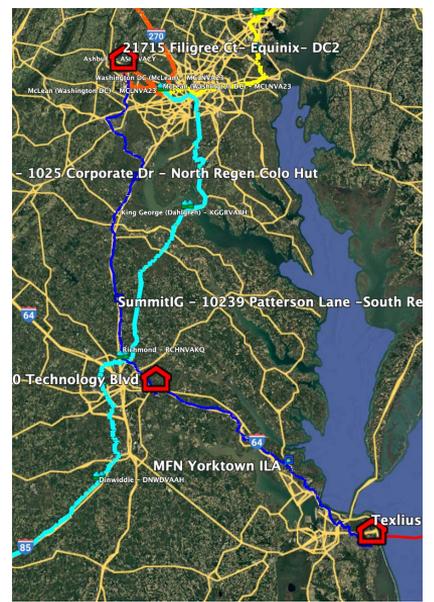
AEC-1 WL5 Implementation

Dec 2024 3x400G



Dunant - Virginia Beach to St. Hilaire de Riez

- 25% FP EQXDC4-PARIS Procurement Complete
- Required spectrum change at Orange CLS
- New POP at Digital Realty PAR1
 - Splice into existing EQXLD8-CERN GEANT path
- Will need future build at VA Beach CLS
- Expecting to bring up circuits by end of Q1 2026

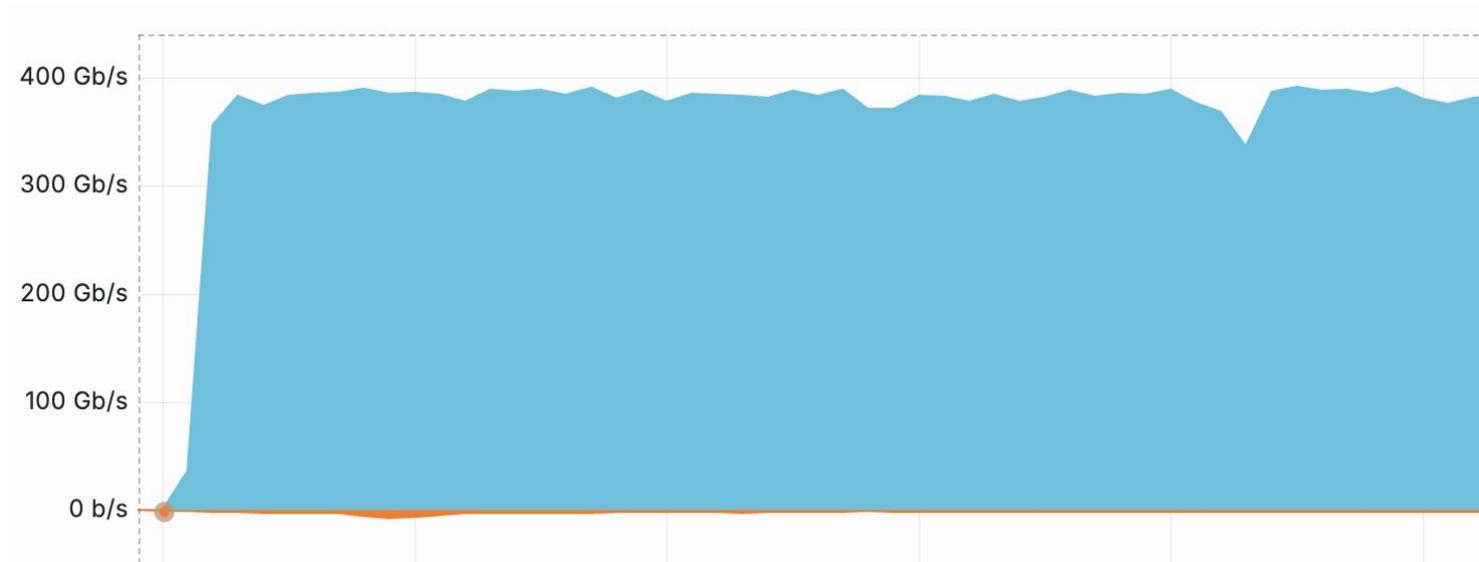


LHCOPN vs LHCONe Comparison on ESnet

- LHCONE
 - Dedicated, Any-to-Any L3VPN with strict AUP
 - Highly parallelized and resilient
 - Multiple peerings, LAGs from sites
 - ECMP throughout ESnet
 - Load balancing across TA paths
- LHCOPN
 - OSCARS as key technology
 - enables shared use of the network by multiple science domains
 - Bandwidth reservation guarantee
 - Preferred QOS profile
 - Point-to-point Layer 2 circuit model
 - Strict path layout
 - Resilient Primary, Secondary, Tertiary circuits
 - Optional reroute capability

Recent LHCOPN testing

- 400G is the current limit of any *one* LHCOPN circuit
 - Now that we have the bandwidth, what is our next step for LHCOPN above 400G?



A background of a network diagram with white nodes and lines on a blue gradient. The nodes are connected by thin white lines, forming a complex web. The nodes themselves are small white circles, some of which are slightly larger and brighter than others. The overall effect is a sense of interconnectedness and data flow.

LHCONE: US Tier-2 Ecosystem

Current US Tier-1 Status

- BNL
 - 1.6 Tbit/s
 - 800G (2x400GE) primary for LHCOPN + R&E connectivity
 - 800G (2x400GE) primary for LHCONE
- FNAL
 - 2.4 Tbit/s
 - 1.6 Tbit/s (4x400GE) primary for LHCOPN & LHCONE
 - 800G (2x400GE) for R&E connectivity

- *In both cases, traffic can failover between links*

ESnet and the US Tier-2 Campuses

- *Typical T2 topology:*
 - ESnet LHCONE --> Regional Network --> Campus --> T2
- 2025-2026: Synchronize DC27 plans with all parties
 - Individual PI's
 - Tier 2 and Departmental Support Staff
 - Campus IT groups
 - Regional Networks
 - R&E Exchange points
- **Data Challenges (plus mini challenges) are driving real outcomes crucial to HL-LHC success**

US Tier-2 Status Updates - Regionals

- In Production:
 - 4x400G to BTAA for MWT2, AGLT2, Wisconsin, Purdue
 - 2x400G to GPN for Nebraska and U-Oklahoma (SWT2)
 - 400G to NEREN for NET2
 - 400G to SOX for Vanderbilt
 - 400G through WIX for U-Florida
 - 2x400G to CENIC for UCSD and Caltech
 - 400G to NOX for MIT
- Pending:
 - 400G to LEARN for UT-Arlington (SWT2) & TACC Tier-3

US Tier-2 Recent Status Updates - Sites

- In Production:
 - UCSD 400G upgrade to CENIC
 - MIT 400G upgrade to ESnet at both Boston and NY
- Pending:
 - Caltech 400G upgrade to CENIC



Mini Data Challenge

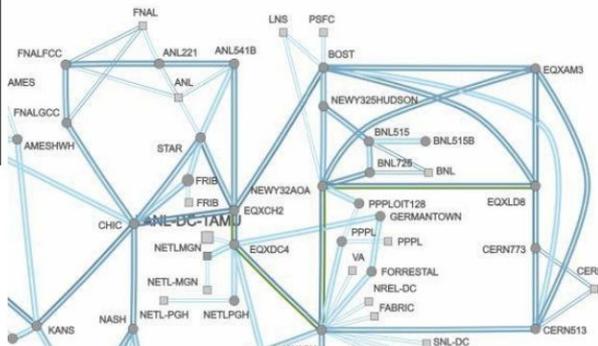
- US-CMS and US-ATLAS conducted a joint mini data challenge in August 2025
- First week was for each experiment to test independently to their US Tier 2 sites with a focus on capacity
- Joint testing on the second week was conducted to target any overlapping contention between Tier 1s and Tier 2s in the United States
- **Result:** No contention detected on ESnet network
- ESnet monitoring services were used including Stardust dashboards with SNMP & flow data as well as high touch data was captured and exported for future analysis

Observation from Joint USATLAS/USCMS Test

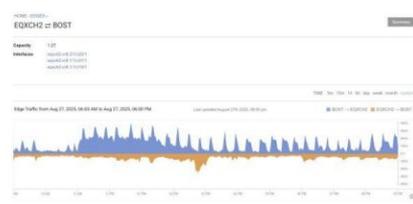
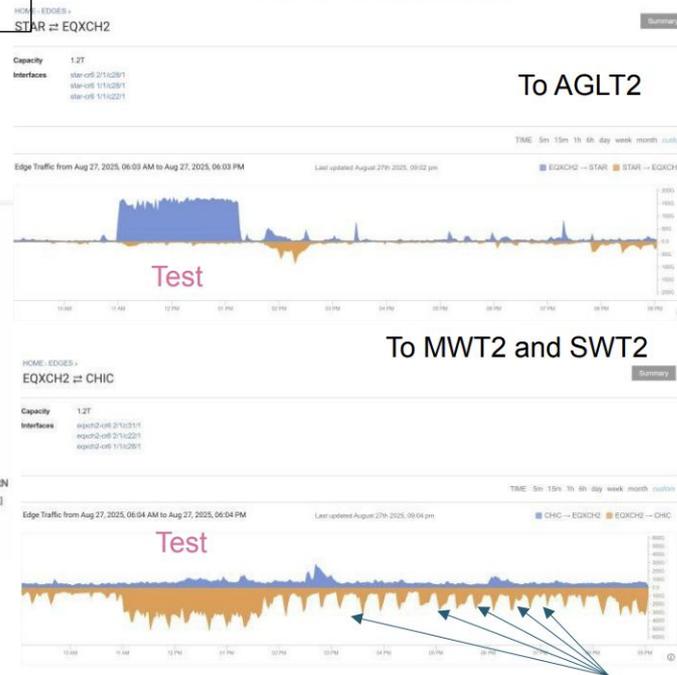
Test: From BNL and NET2 to AGLT2, MWT2 and SWT2 (UTA) and FNAL to all USCMS Tier-2s

No obvious hotspots have been observed. The destination sites were obtaining data at their expected site rates.

Periodic large data from CERN to FNAL is observed during test although it didn't affect the data rate for USATLAS sites.



ESNET monitors



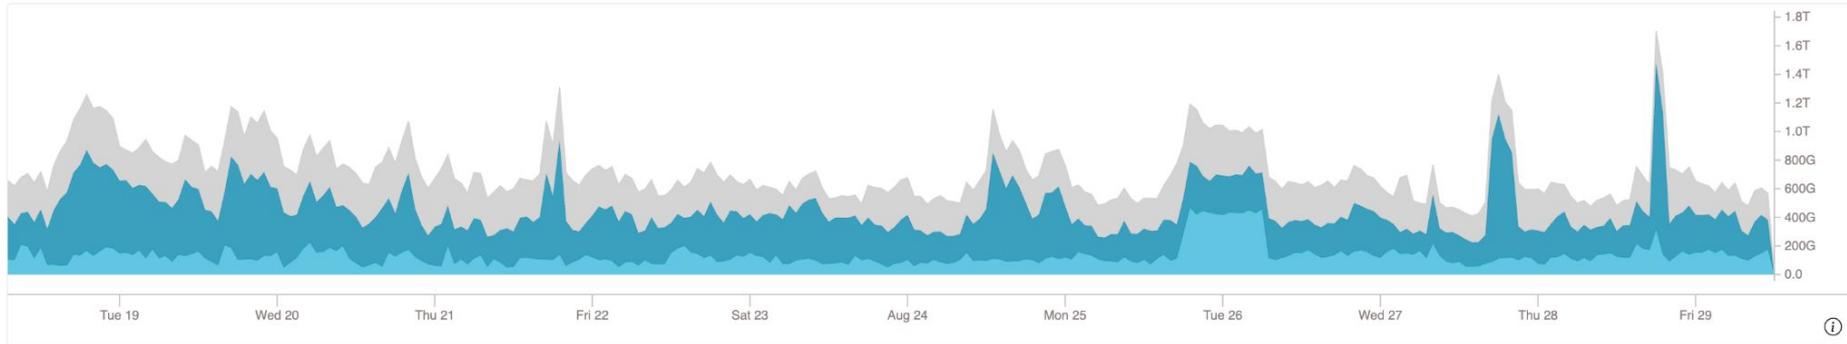
Curious periodic traffic?

Total ESnet Traffic 8/19-8/27

ESnet Traffic from Aug 17, 2025, 10:00 PM to Aug 29, 2025, 03:00 AM

Last updated August 29th 2025, 12:00 pm

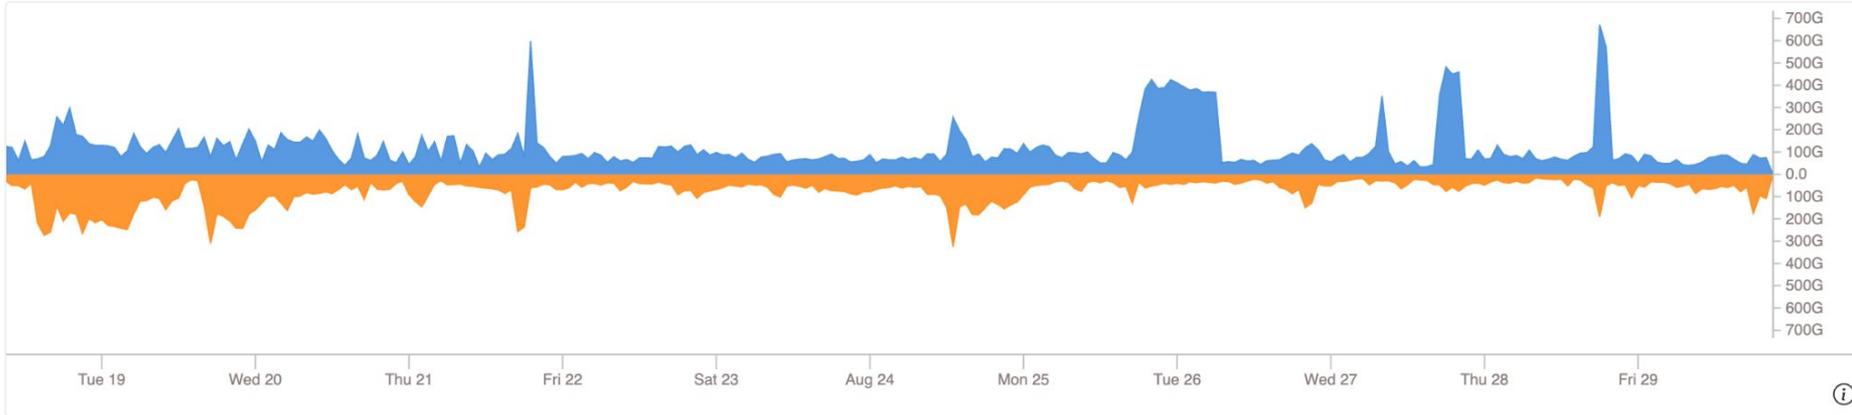
■ OSCARS ■ LHCONE ■ Other



Total Site Traffic

Last updated August 29th 2025, 09:00 pm

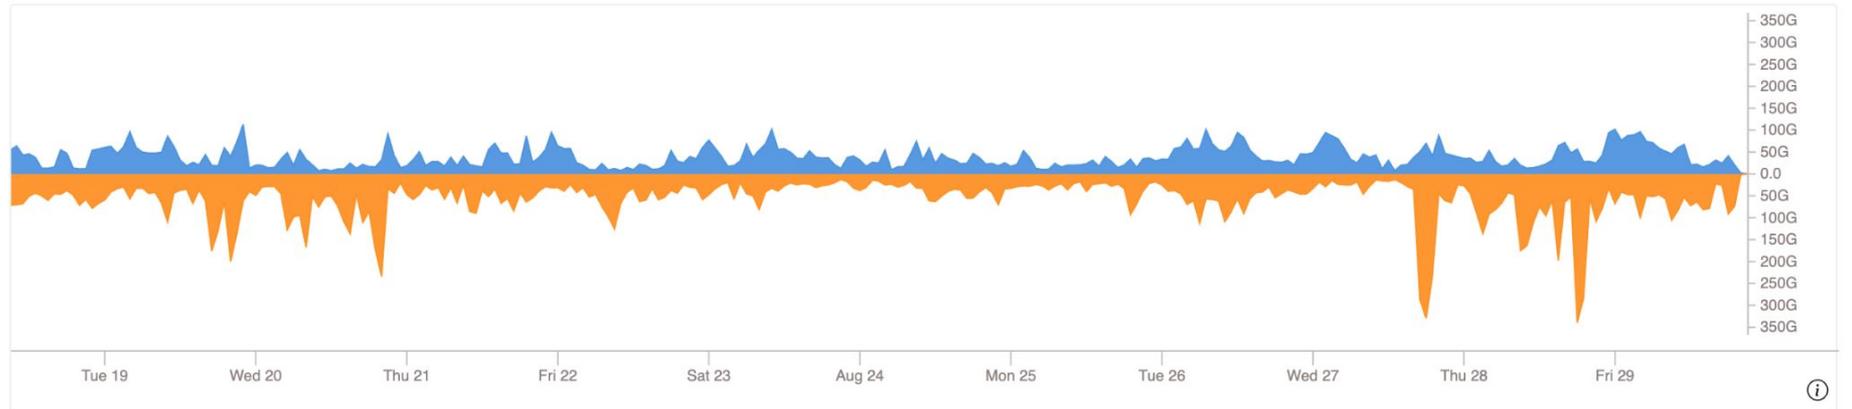
To site From site

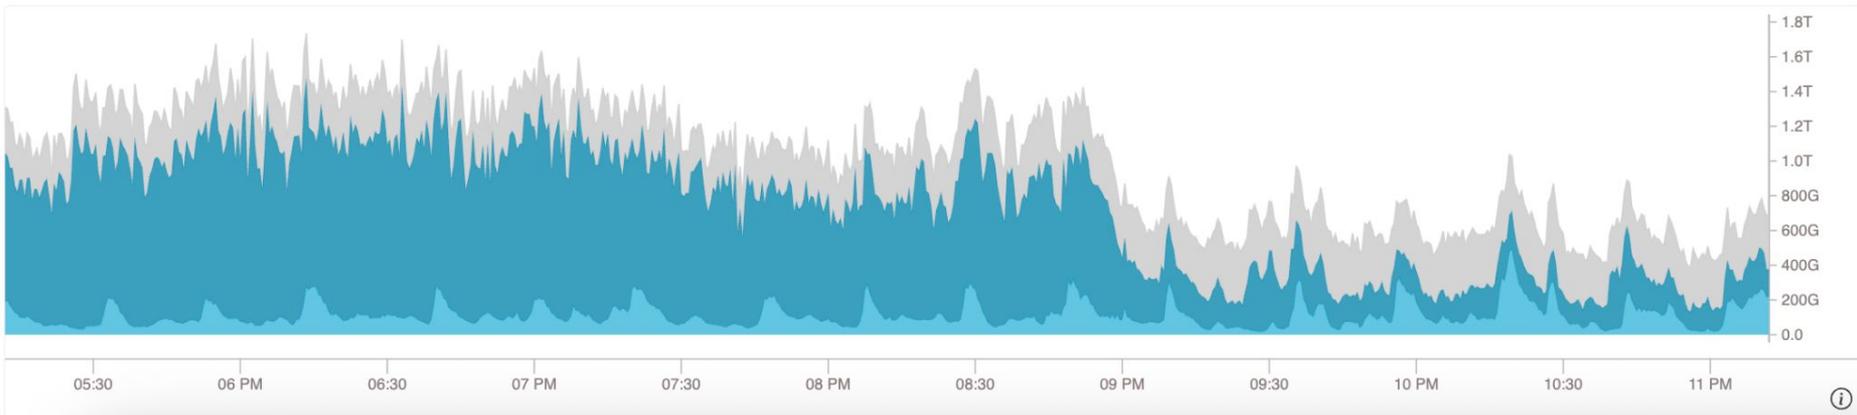


Total Site Traffic

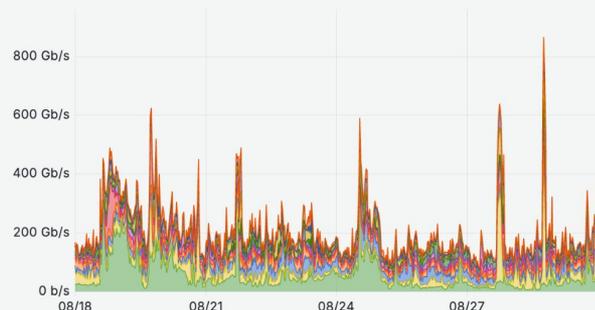
Last updated August 29th 2025, 09:00 pm

To site From site



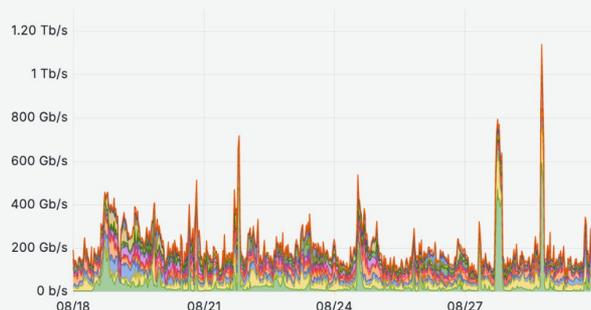


Top 10 Interfaces by Incoming Rate (SNMP)



- fnalfcc-cr6::fna1_se-1600
 ■ bnl725-cr6::bnl_se-101
 ■ cern513-cr6::geant_se-719
- chic-cr6::uchicago_se-204
 ■ chic-cr6::in-gpop_se-6
 ■ eqxam3-cr6::geant_se-707
- seat-cr6::canarie_se-326
 ■ nash-cr6::vanderbilt_se-1605
- chic-cr6::uwmadison_se-203
 ■ cern773-cr6::cern_se-722

Top 10 Interfaces by Outgoing Rate (SNMP)



- fnalfcc-cr6::fna1_se-1600
 ■ chic-cr6::uchicago_se-204
 ■ chic-cr6::in-gpop_se-6
- cern513-cr6::geant_se-719
 ■ bost-cr6::net2_se-1281
 ■ bost-cr6::mit_se-1769
- eqxam3-cr6::geant_se-707
 ■ chic-cr6::uwmadison_se-203
 ■ bnl725-cr6::bnl_se-101
- hous-cr6::uta_se-490

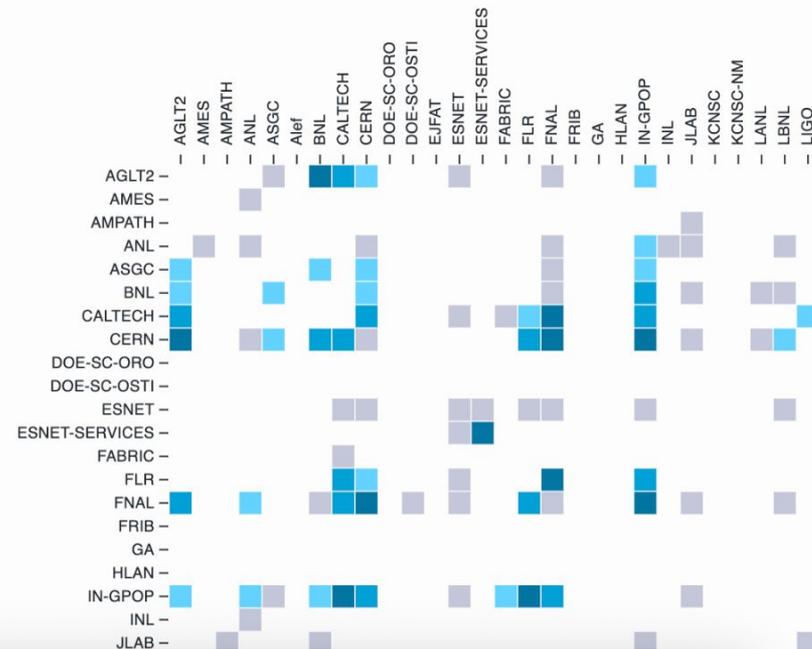
HT Packet Size Explorer

This dashboard explores transfers with large packets (frames > 4096 bytes) between sites using High Touch data. It is NOT recommended to set the time range to longer than 30 minutes.

Top Talkers with large packet transfers

Source	Destination	Total Volume	Num Lar
UCHICAGO	IN-GPOP	779 GB	
FNAL	IN-GPOP	15.9 GB	
IN-GPOP	UIUC	98.2 GB	
IN-GPOP	UCHICAGO	180 GB	
CERN	FNAL	1.05 TB	
VANDERBILT	CALTECH	53.5 GB	
VANDERBILT	UCSD	33.9 GB	
VANDERBILT	FLR	28.9 GB	
CERN	IN-GPOP	2.44 GB	
IN-GPOP	UNL	211 GB	

Large Packets between High Volume Sites ⓘ





U.S. DEPARTMENT
of ENERGY

Thanks!