



U.S. DEPARTMENT
of ENERGY

ESnet Update



Dale W. Carder
Kate Robinson
ESnet Network Engineering

LHCONE/LHCOPN 54
Manchester, UK
2025-03-18

Outline

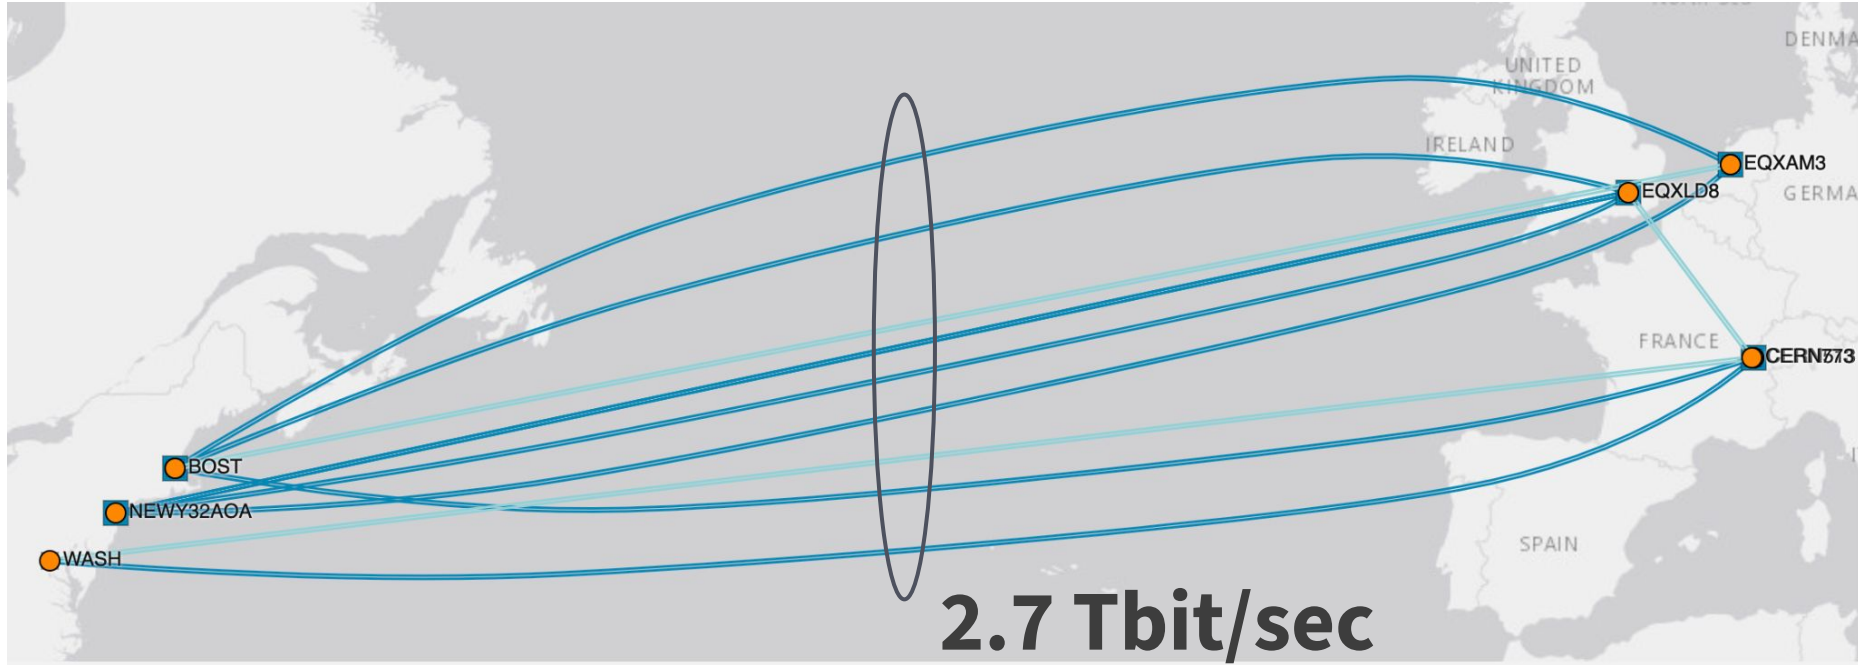
- Trans-Atlantic Updates
- LHCOPN Futures
- LHCONE & the U.S. Tier-2 Ecosystem

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)

* assuming funding continues as expected

Current ESnet Transatlantic Capacity



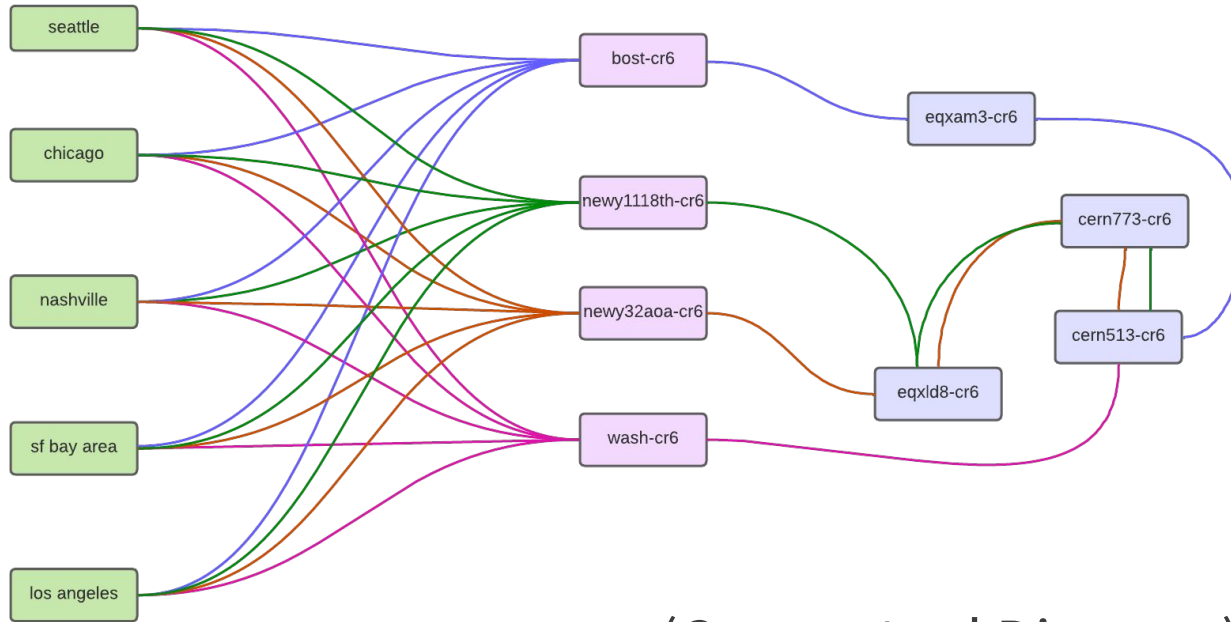
Trans-Atlantic and Europe Connectivity

- ESnet operating 2.7Tbit/sec for our backbone currently
- Additional cable & spectrum to be lit by the end of the year.
 - ESnet will also be expanding our Europe ring.
- Other networks are also bringing in substantial additional resources.
 - Community discussions are underway on mutual backup arrangements.
- The efforts now shift to the Experiments to make effective use of these massive investments.

Load Balancing TA Links for LHCONE

- By treating all links possible as "internal to ESnet", end to end traffic engineering is tractable vs a piecemeal approach
 - See Michael Sinatra's talk(s) on Segment Routing
 - 1: https://www.youtube.com/watch?v=QW7_vq1MWOY
 - 2: <https://www.youtube.com/watch?v=mwVlPlu9OEs>
- ESnet routers /w LHCONE connections load-balance traffic to/from Europe *across all available TA paths*
 - Overrides shortest-path routing
 - Up to 6 paths to choose from
 - Can weight some paths heavier than others.


LHCONE Load Balancing




... And many more routers, where
LHCONE sites connect

(Conceptual Diagram)

Trans-Atlantic Utilization: Europe --> U.S.



Trans-Atlantic Utilization: U.S. --> Europe




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 - Success!

- 400G is the current limit of any *one* LHCOPN circuit
 - Contrast to LHCONE, which can use parallel paths



LHCOPN Future - How do we want to grow?

- Think about service composition
 - Complexity tradeoffs between sites and networks
 - No free lunch
- ex. a) Multiple LHCOPN circuits in parallel?
 - Each circuit would have a specifically enumerated topology
 - ECMP support and extra BGP sessions between endpoints
- ex. b): Multiple LSPs for one circuit?
 - Sites keep single bgp session, traffic spread across LSPs
 - Do you tear down the path if some or a majority of LSPs fail?

A background of a network diagram with white nodes and lines on a blue gradient. The nodes are scattered across the frame, with some lines connecting them, creating a web-like structure. The overall aesthetic is clean and modern, typical of a corporate or technical presentation.

LHCONE: US Tier-2 Ecosystem

ESnet and the US Tier-2 ecosystem

- *Typical T2 topology:*
 - ESnet LHCONE --> Regional Network --> Campus
- 400G Now viable at regional networks
 - Favorable economics in most newer form factors
 - Supply-chain & backlogs resolved
 - Optic prices increasingly better
- 2025 Goal:
 - Work with each regional network on readiness
 - Plan for 400G per T2

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
- Pending:
 - 2x400G to CENIC for UCSD and Caltech (later today?!!!)
- Planning:
 - 400G to NOX for MIT
 - 400G to LEARN for UT-Arlington (SWT2) & TACC Tier-3

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**



U.S. DEPARTMENT
of ENERGY

Thanks!

Dale W. Carder
Kate Robinson
ESnet Network Engineering

LHCONE/LHCOPN 54
Manchester, UK
2025-03-18

