



ESnet

ENERGY SCIENCES NETWORK

ESnet Transatlantic Engineering

Kate Robinson

Eli Dart

Dale Carder

LHCOPN/LHCONE #52

University of Catania

9-11 April 2024



U.S. DEPARTMENT OF
ENERGY

Office of Science




Agenda

- ESnet Transatlantic Current View
- ESnet Transatlantic Future Strategy
- ESnet TA traffic engineering
- GEANT routing policy
- LHCOPN Paths

Trans-Atlantic upgrades

- **Now In Production:**
 - 400G New York - London
 - 400G Boston - CERN
- **Currently underway:**
 - 400G Boston - London (End of April)
- **Trans-Atlantic capacity targets**
 - **3.2T in 2027**, in advance of Run 4



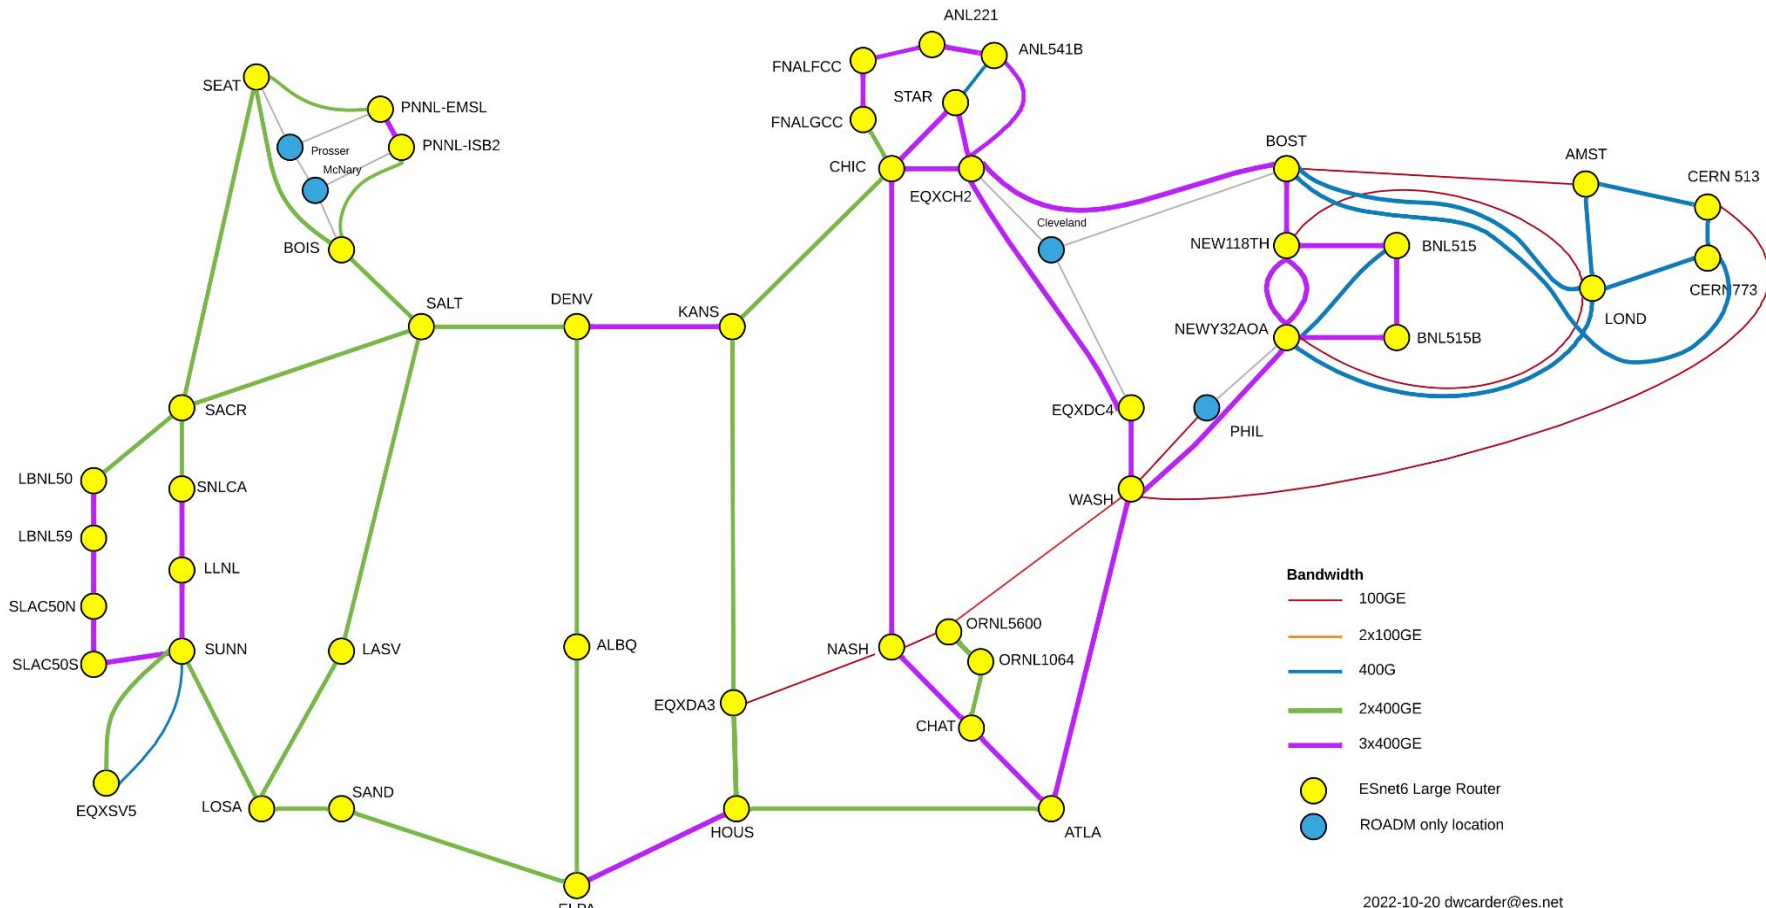
*Assuming funding continues as expected

Amitie / AEC-3

- Procurement complete in coordination with Internet2
- Procured 2x “Managed Spectrum” channels 102.5Ghz each




ESnet Backbone upgrades - Q3 2023



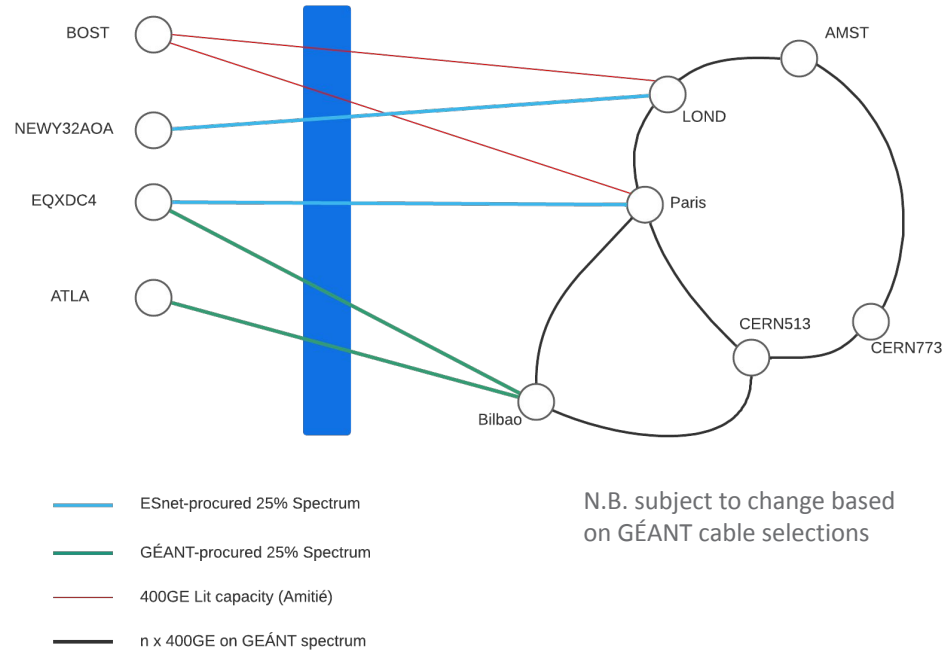
Future US - Europe Connectivity Plans

- Collaborating with GEANT to share spectrum on subsea
- Plan: collectively acquire optical spectrum across 4+ diverse cables
- ESnet-targeted cables are fixed; still some variability in GEANT plans
- Depending on GEANT's spectrum procurement timelines, we may investigate additional 400GE lit services for interim diversity



US - Europe Connectivity: 2027 Plan

- Two additional EU PoPs
 - Paris (firm)
 - Bilbao? (tbd)
- n x 400G EU rings in partnership w/ GÉANT
- Subsea spectrum IRUs are 15+ year contracts
- Subsea spectrum across 4 cables should provide ~10Tbps aggregate capacity to ESnet
 - Meet HL-LHC commitments plus growth



Aquacomms AEC-1 (“Cable-1”) Spectrum

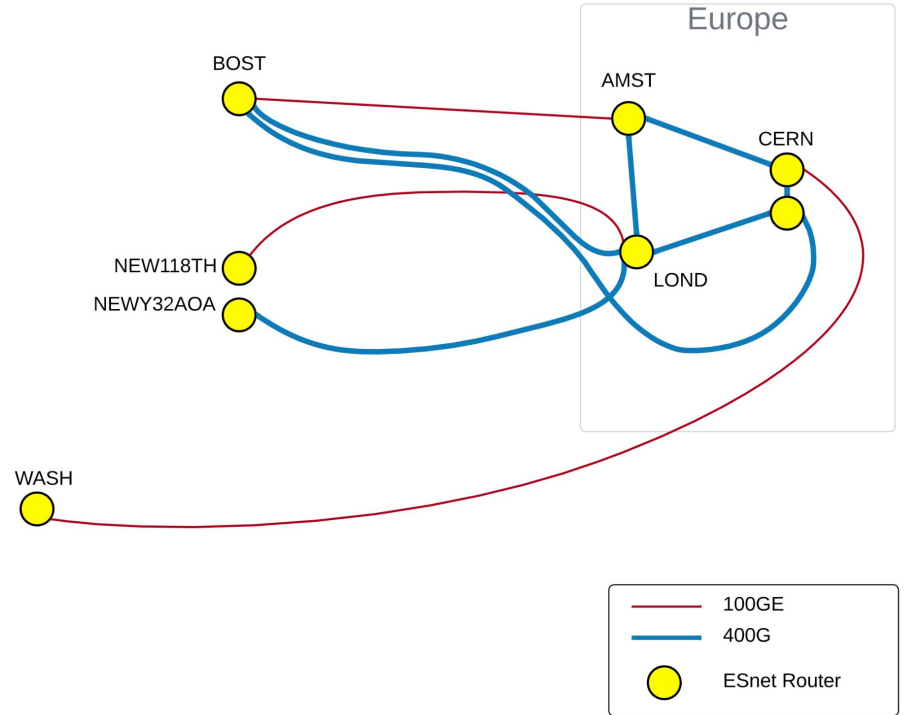
- Procurement complete
- Awaiting spectrum service delivery from Aquacomms
 - Expected Summer 2024
- Follow-on procurement for Dublin colo expected soon
 - Optical-only for regen + GÉANT add/drop
- SLTE (Infinera FlexILS) equipment BOM finalized
 - Requisition to be placed ASAP
 - Installation late summer / early fall
- Target: 3x400G in service NEWY32AOA - EQXLD8 by end of CY2024

Additional Spectrum (“Cable-2”)

- Dunant Cable
- RFP development in progress
- Requesting 25% spectrum from Paris - Ashburn, VA
- Targeting award by end of CY2024

European ring upgrades

- **Now In Production:**
 - 400G European Ring
 - Amsterdam to CERN
 - London to CERN
 - Amsterdam to London
- **Future plans**
 - Additional 1x400G (ETA 9/24)




*Assuming funding continues as expected

Trans-Atlantic IGP Traffic Engineering

- Treat all links possible as "internal to ESnet" (2 x 100g NEAAR/AEC-2)
 - Encapsulate MPLS traffic onto a vlan, pop out the other side.
- Configured policy on ESnet routers that have LHCONE connections to load-balance traffic to/from CERN *across all available TA paths*
 - Overrides shortest-path routing
 - Up to 6 paths to choose from
 - Can weight some paths heavier than others.
- Started with a handful of routers
- Expanded to include all LHCONE Tier-2 connections
 - Some added on week 2 of the data challenge to mitigate traffic on wash-cern link, which is a shortest path for sites in the Southern U.S.

Trans-Atlantic IGP Traffic Engineering Example



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

ESnet Routing policy

- Typical ESnet R&E routing policy is to "cold-potato" all traffic:
 - accept traffic into ESnet as close to the source as possible (ask peers to "hot-potato" to ESnet)
 - deliver traffic as close to where it is going, keeping it on ESnet for as long as we can
 - Coordinated by humans & agreements
 - Implemented via BGP MED attributes indicating relative costs

Path engineering w/ eBGP, DC24 nuances


- During DC24, we wanted to utilize all potential Trans-Atlantic paths (including non-ESnet), particularly in the EU->US direction
 - With late delivery of our Amitie cable, there was some internal concern with (2) 400G LHCOPN circuits potentially impacting LHCONE
- Traffic engineering goal:
 - Allow GEANT to deliver LHCONE traffic from London and Paris to the US on their TA infrastructure, freeing that up from ESnet's links
 - Keep GEANT to ESnet traffic at Amsterdam and Geneva the same
 - (however, we were ready to tweak that if need be)
 - For LHCONE peerings w/ GEANT, ESnet currently not sending BGP MEDs.
 - Worked great, and leveraged the full capability of R&E infrastructure

GEANT Routing Policy

- Prioritise routes from R&E networks over those from Internet Exchanges or commercial peers using local preference.
- Advertise MEDs to our BGP neighbours, these MEDs are derived from our IGP metrics. This ensures that incoming traffic from neighbours enters our network at a point nearest to its destination.
- Neighbors may honor announced MEDs if choose
- Enable NRENs and R&E peers to choose hot potato or cold potato routing approach or a mixture of the two.

** Info provided by GEANT **


LHCOPN - <https://my.es.net/oscars/list>




- Primary LHCOPN paths are each using a diverse 400G TA link.
- There is only 1x400G ring in EU currently but each use a different direction and ESnet entry points into CERN
- Secondary/tertiary paths use diverse 100G paths
- Without protect enabled, failover is manual (FNAL)
- With protect enabled, the primary will fail over to shortest-path protect

BNL LHCOPN DC24

IPv6




IPv4




FNAL LHCOPN v6 DC24

IPv6




IPv4




ESnet Transatlantic Usage peaks during DC24

Total Europe to US Traffic (SNMP) on Transatlantic Links



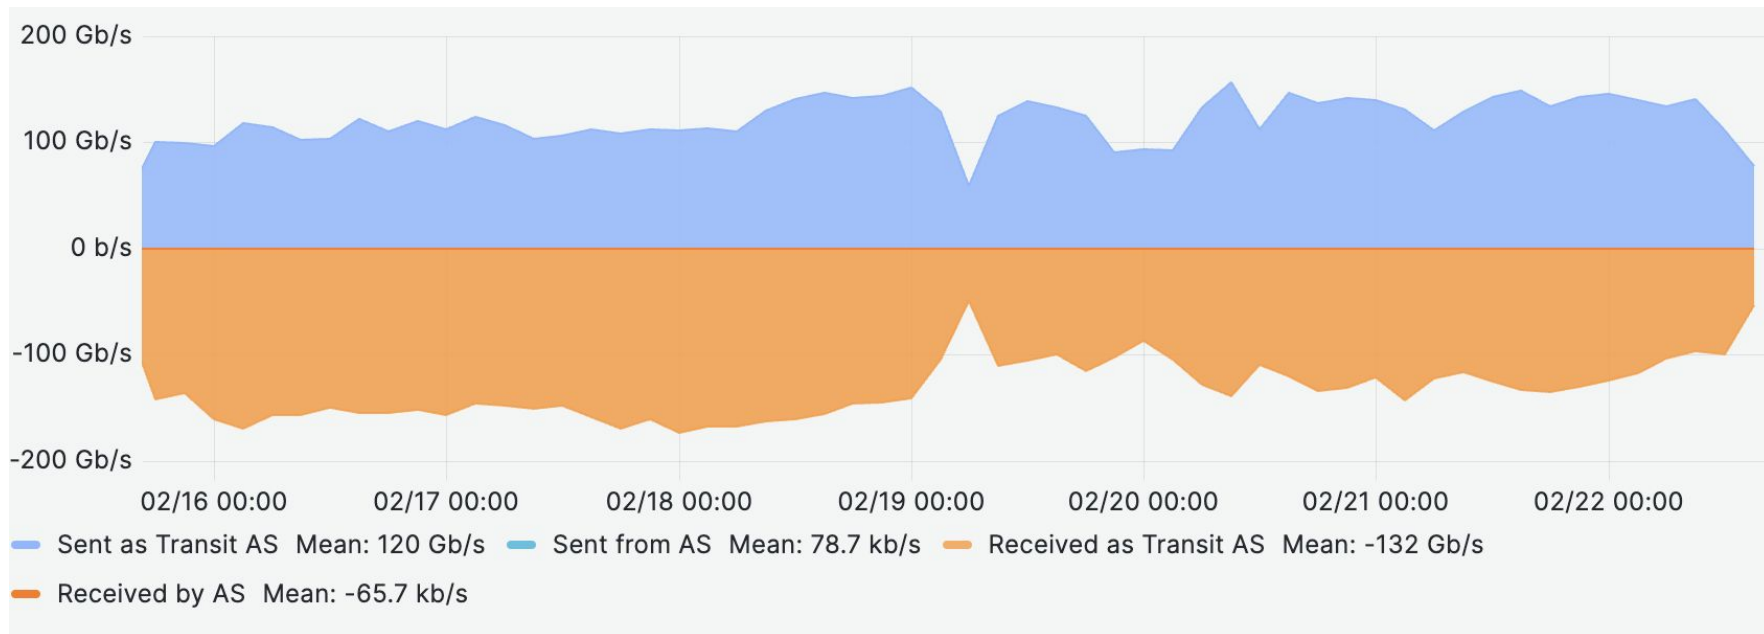
Total US to Europe Traffic (SNMP) on Transatlantic Links



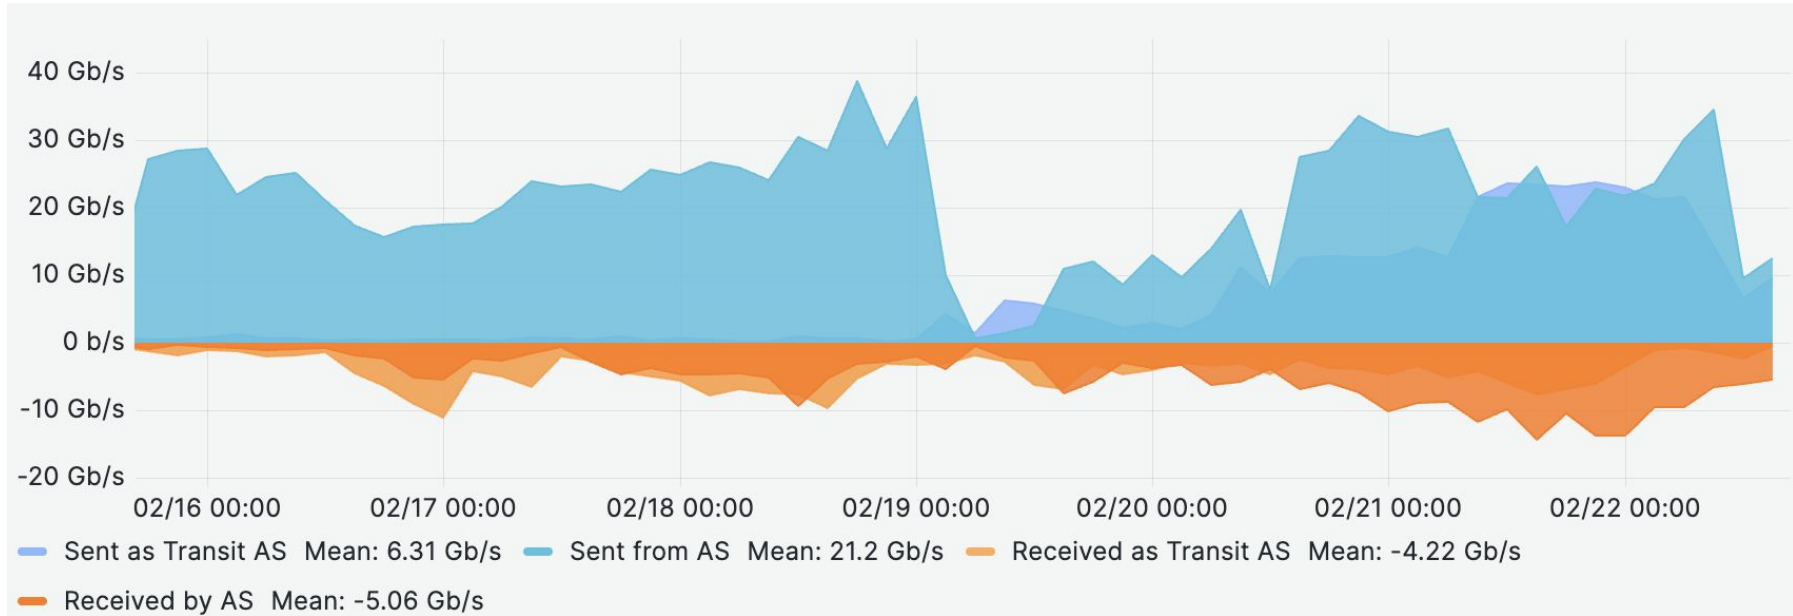
LHCOPN Traffic placement + LHCONE weighted load-balancing

<https://public.stardust.es.net/d/lkFCB5Hnk/lhc-data-challenge-overview?orgId=1&from=1708451852305&to=1708711052305>

ESnet--GEANT Traffic



ESnet--SURFNET Traffic



ESnet--NORDUNET Traffic

