



ESnet

ENERGY SCIENCES NETWORK

ESnet Updates

Kate Robinson

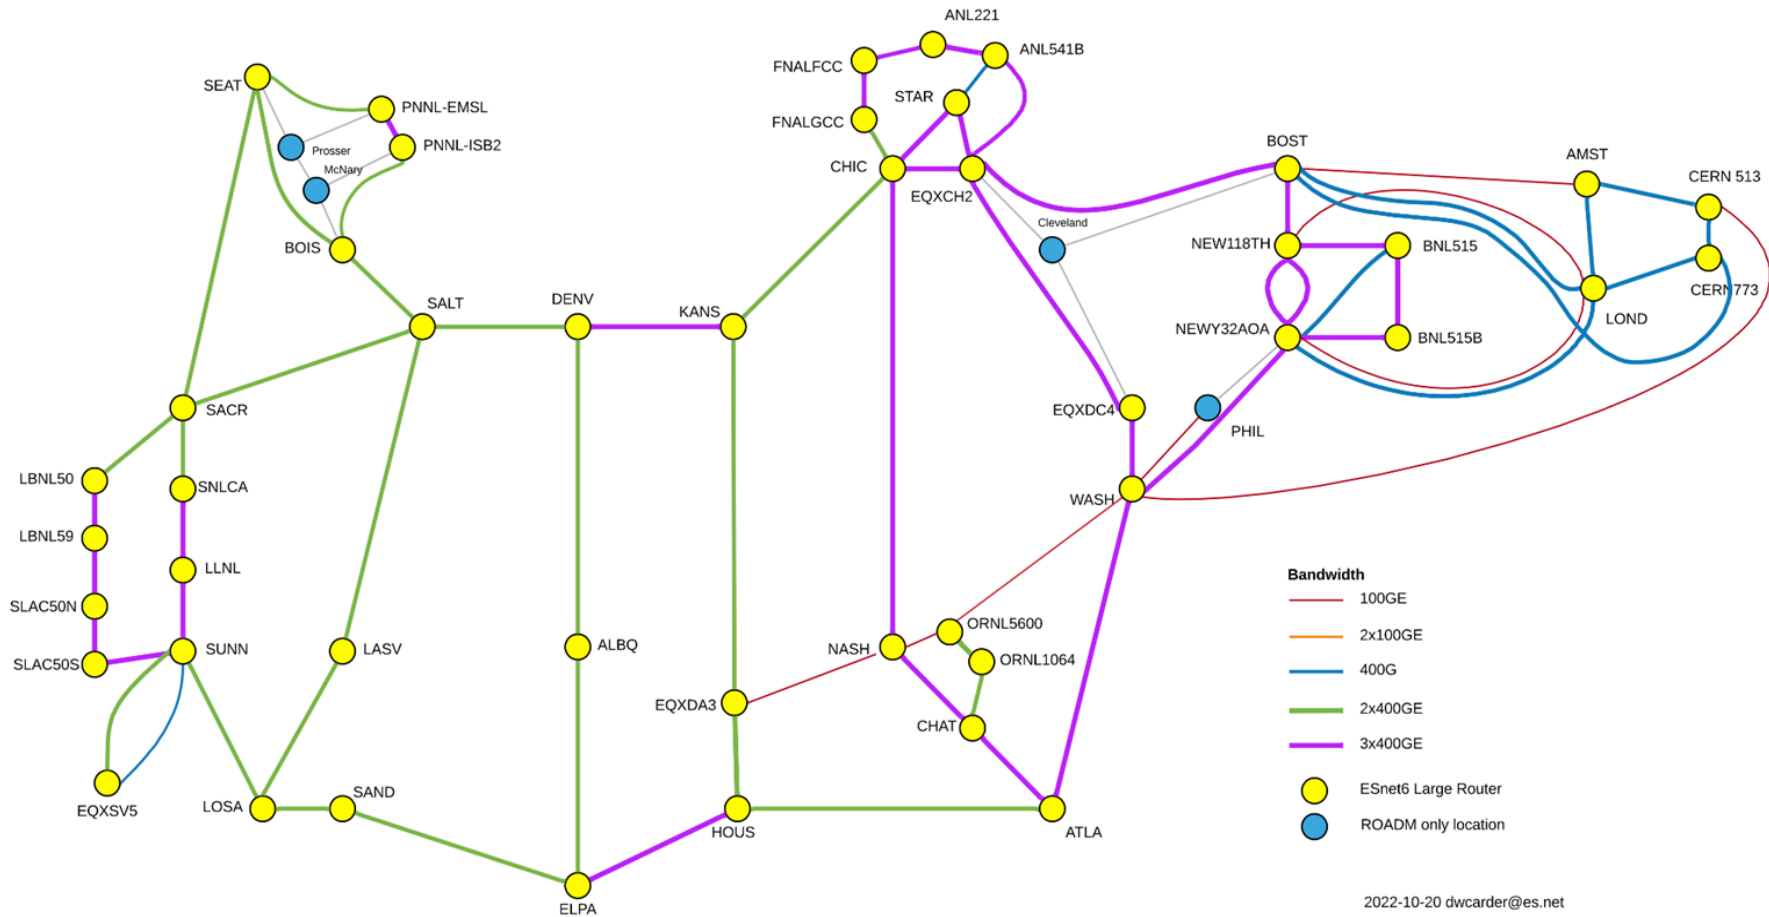
LHCOPN/LHCONE #51
University of Victoria
18-19 October 2023



Agenda

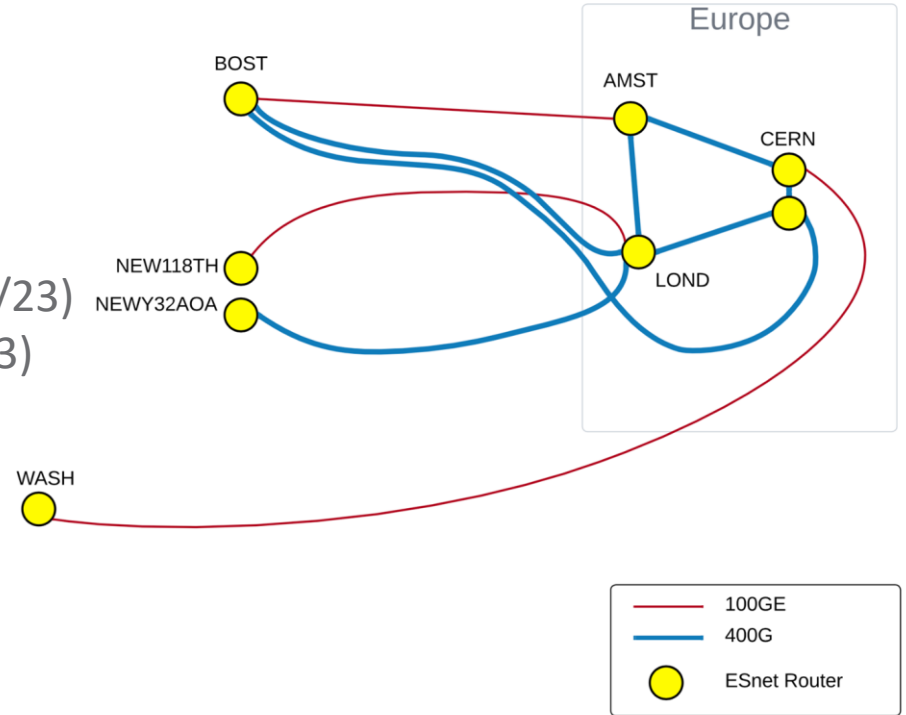
- ESnet Backbone Planning
- Trans-Atlantic Upgrades
- European Ring Upgrades
- Tier1 and Tier2 connectivity
- perfSONAR plans for LHCONE
- Statistics for DC24
- In-network data caching pilot updates

ESnet Backbone upgrades - Q3 2023



Trans-Atlantic upgrades

- **Now In Production:**
 - 400G New York - London
- **Currently underway:**
 - 400G Boston - London (Estimated 12/23)
 - 400G Boston - CERN (Estimated 12/23)
- **Trans-Atlantic capacity targets**
 - 1.5T in advance of DC24
 - **3.2T in 2027**, in advance of Run 4



*Assuming funding continues as expected

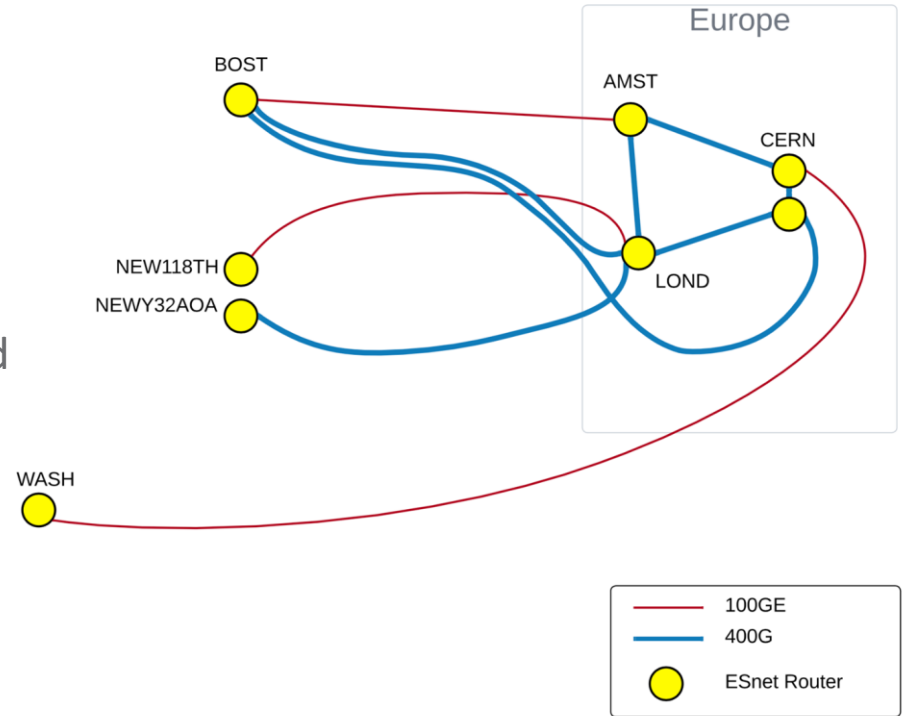
Amitie / AEC-3

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



European ring upgrades

- **Now In Production:**
 - 400G European Ring
 - Amsterdam to CERN upgraded
 - London to CERN upgraded
 - Amsterdam to London upgraded
- **Future plans**
 - Additional 1x400G



*Assuming funding continues as expected

ESnet Cloud connectivity today

- via private fiber interconnects
 - 5x100G to Google (one more pending)
 - 3x100G to Oracle
- via fabrics
 - 6x100G to Microsoft
 - 6x100G to Amazon


Private Cloud Interconnects to nearly any provider

- 5 locations (each 2x100G) to PacketFabric
 - OSCARS connectivity across ESnet
 - possibility for API-based provisioning end-to-end

ESnet participated in a trial of Google Cloud for use as a Tier 2

US Tier 2 Site Connectivity

- Report identified connectivity needs for T2 Sites
- Projections codified into data challenges
- Data Challenges (full software stack)
 - 1: 10% of the target 2021 # Successful
 - 2: 25% in 2024 # On Track
 - Interested in understanding what happens after DC24
- Continuing engagement with US Tier2 sites and the network paths to get to them



<https://doi.org/10.2172/1804717>



Regional Connectivity Upgrades

- ESnet has been focusing on upgrading connections to regional exchange points that connect us to US Tier 2 sites in preparation for DC24
 - CENIC in Los Angeles (in progress)
 - CalTech, UCSD
 - GPN 400G (completed)
 - Nebraska
 - NOX 100G in Boston (completed)
 - Supporting New NET2
 - SOX (in progress)
 - Vanderbilt
 - OmniPoP (in progress)
 - Michigan
 - Michigan State
 - U-Chicago
 - IU-Bloomington
 - UIUC
 - Purdue
 - UW-Madison

US Tier 1 Connectivity

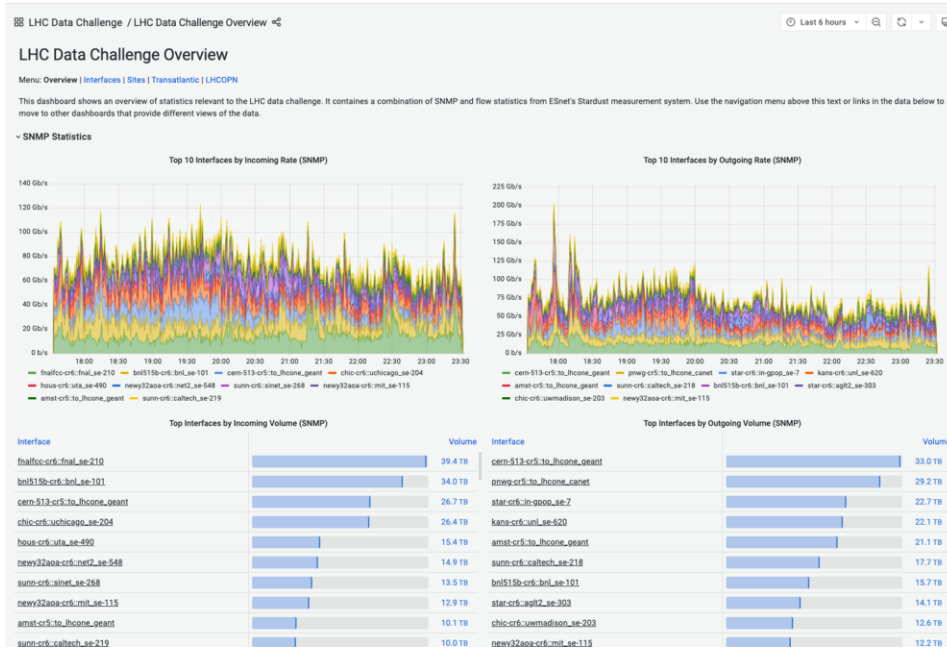
- ESnet6 installed routers collocated at our sites
- Most are connected to our optical system at 1.2Tbit + redundancy
- We are now ready to accommodate upgrades as sites are able
 - BNL - US ATLAS Tier 1
 - Current: 300G (2 x 100G + 1 x 100G)
 - Near Future: 800G (1 x 400G + 1 x 400G) Estimated 12/2023
 - FNAL - US CMS Tier 1
 - Current: 400G (2 x 100G + 2 x 100G)
 - Near Future: 800G (1 x 400G + 1 x 400G)

perfSONAR

- ESnet has a robust perfSonar deployment, but previously there were only a few LHCONE hosts in ESnet (maybe just 1)
- We recently deployed interfaces on many current ESnet perfSonar servers in LHCONE as we wanted to avoid additional hardware deployments for LHCONE
- IPv6 only perfSonar interfaces in LHCONE
- Selected perfSonar hosts nearest to US Tier 1 and 2 sites
- [LHCONE](#) dashboard
- [ATLAS](#) mesh dashboard
- [CMS](#) mesh dashboard
- What other hosts (especially non-ESnet hosts) should be added to this?

Statistics for DC24

- ESnet has been working to improve our statistics collection and display in advance of DC24
- Please take a look and let us know if there is anything additional you would like to see from ESnet
- [Public dashboard](#)



In-network caching pilot

- Scientific experiments and simulations generate increasingly large volumes of data
- Growing data volumes result in demanding data movement requirements
- Data stored at a small number of locations is challenging for geographically distributed collaborations (e.g. LHC)
- Significant portions of popular datasets are used by many researchers
- Storage caching allow data sharing among users in the same region
 - Reduces repeated data transfers over the wide-area network
 - Decreases data access latency
 - Increases data access throughput
 - Improves overall application performance

Data caching current use cases

- US CMS: Southern California Petabyte Scale Cache
 - Regional storage cache for US CMS user analysis at Caltech and UCSD
 - 24 federated XCache nodes: Approximately 2PB of total storage capacity
 - 12 nodes at UCSD: each with 24 TB, 10 Gbps
 - 10 nodes at Caltech: each with storage sizes ranging from 96TB to 388TB, 40 Gbps
 - 1 node at LBNL59 (ESnet): 44 TB storage, 40 Gbps
- US CMS: Chicago
 - Regional cache for U. Wisc Madison, Purdue and Notre Dame
 - 6 federated XCache nodes: Approximately 345 TB in total
 - 5 nodes at U. Wisc Madison: each with 35TB, 10Gbps
 - 1 node at CHIC (ESnet): 184TB on 100Gbps (LHCONE)
- US CMS: Boston
 - Regional cache for MIT
 - 1 node at BOST (ESnet): 184TB on 100Gbps (LHCONE)
- OSG/OSDF: London and Amsterdam

Data caching future plans

- Immediate (2023-2024)
 - Pilot project with US ATLAS
 - Testing multi-service platform on BOST node with DTNasS containers, targeting a BNL's VP
 - Pilot project with OSDF at LOND and AMST (mainly for LIGO and DUNE)
 - Continue discussion and possible pilot work with DUNE
 - Comparing characteristics from different regional caching nodes
 - Additional node deployment
 - CMS and OSDF, multi-use case, at Atlanta
 - Upgrade storage capacity at CHIC and BOST
- Near-term (2024-2025)
 - Possible caching node expansion for LHC experiments
 - Longer term traffic projection
 - Prediction on data distribution to connect to the traffic engineering (possible connection to the high touch data)
- Beyond (2025-2026)
 - Monitoring testbed
 - Making conclusions on the caching pilot
 - Storage-integrated networking testbed
 - Edge storage, longer term storage for pre-staged data with some (simple) computing power
 - Metadata management testbed (Rucio?)

Thanks!

katerobinson@es.net

Energy Sciences Network (ESnet)

Lawrence Berkeley National Laboratory

<http://my.es.net/>

<http://www.es.net/>

<http://fasterdata.es.net/>



LHCOPN

- We have *one* 400GE link across the Atlantic, plus 3x100GE dedicated, plus some hand-wavy fraction of another link (shared with NEEEAR)
- We now have 1x400 between all of our Europe nodes
- Therefore, we are able to provide (and are providing) a 400G clean path for the primary circuit only
- NOTE: If any of the 400G links fail, the primary will fail over to shortest-path protect (all three BNL circuits have protection), which means that you'll share the 100G BOST-AM3 link with FNAL's LHCOPN protect circuit. This is one of the trade-offs with using protection on the primary--if we lose the primary, we can't shift the traffic onto a link that we know isn't being used by the other tier1.
- By the end of CY23 (or maybe a bit thereafter):
- We will have *three* 400GE links across the Atlantic (BOST-CERN513 and BOST-LD8 will be added to 32AOA--LD8) and 2x100G (WASH--CERN513 and 111 8th--LD8)
- Therefore, we will have the ability to have all three OPN circuits on a 400GE clean path
- At some point, we will have 2x400 (instead of 1x400) between all of our Europe nodes
- We should still be able to manage path diversity and component diversity across the 3 LHCOPN circuits
- Note: The estimate for the new 400G TA circuits is end of November 2023, but I am also accounting for acceptance testing (and any kinks getting worked out), cross-connect delays, LOA delays, and other random slippage; hence "end of CY23."