



# ESnet

ENERGY SCIENCES NETWORK

## ESnet Updates

Eli Dart for many others at ESnet

LHCOPN/LHCONE #50  
FZU, Prague CZ  
18-19 April 2023



U.S. DEPARTMENT OF  
**ENERGY**

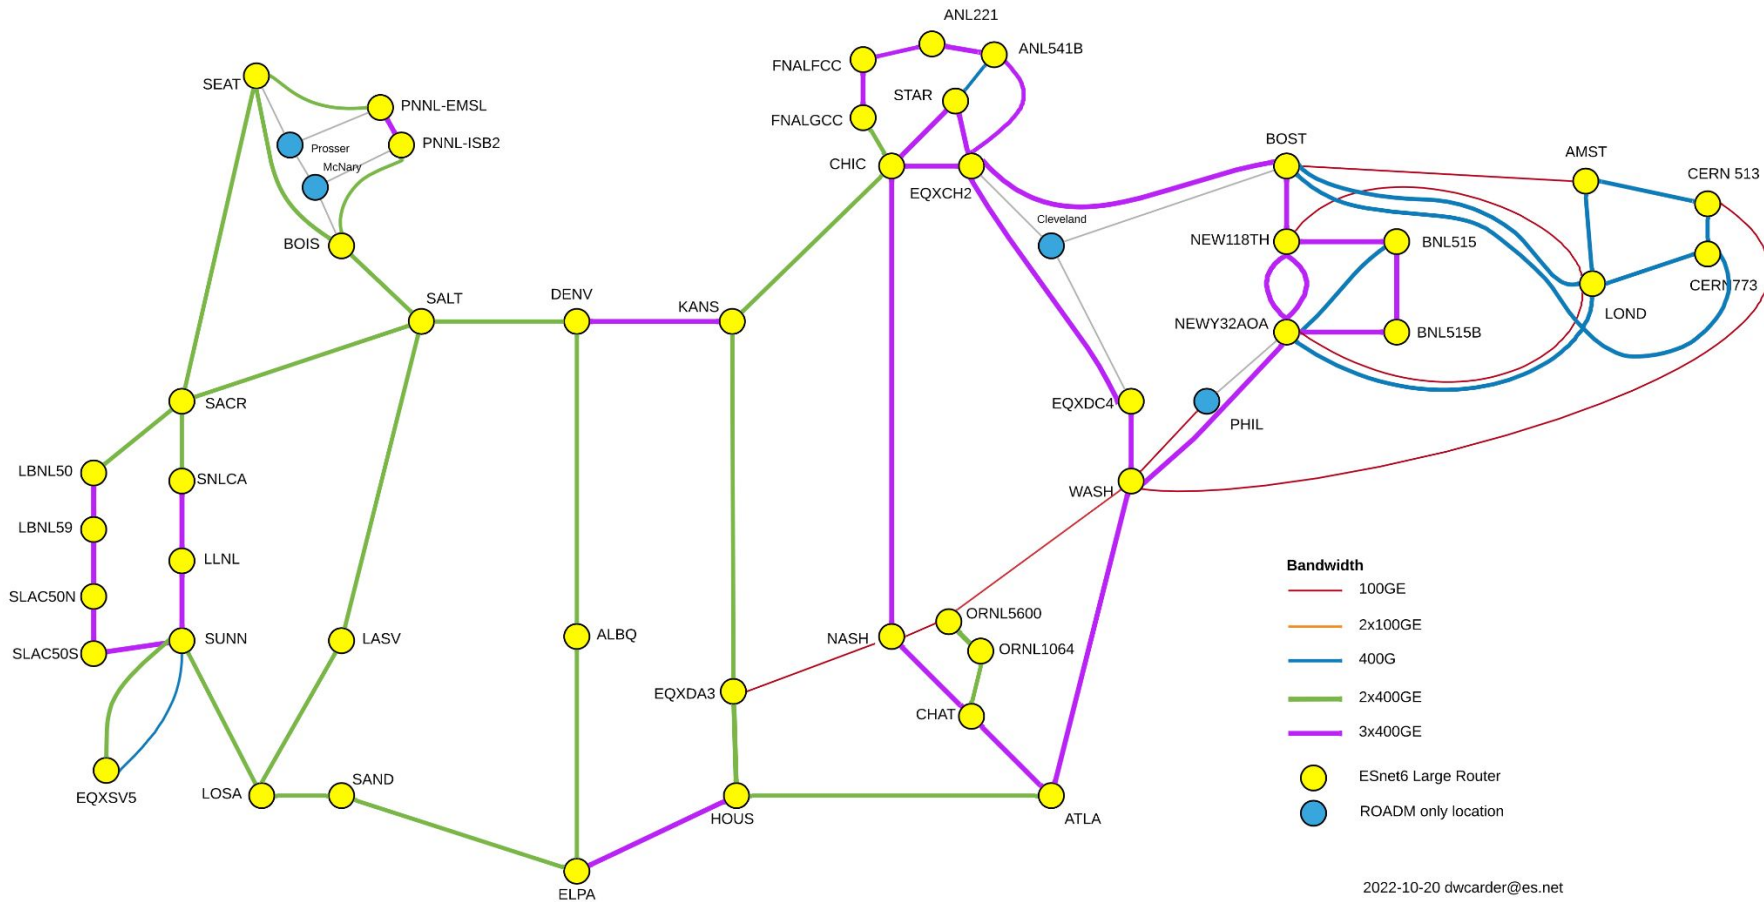
Office of Science



# Agenda

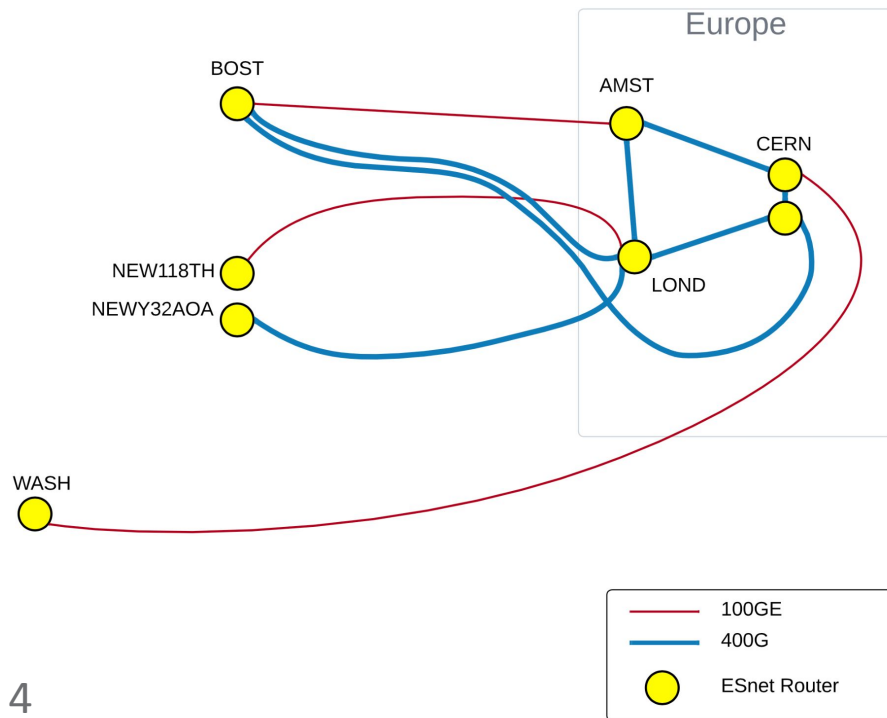
- Note well - this presentation contains the work of many people
- ESnet Backbone Planning
- Trans-Atlantic Networking Plans
- Tier1 and Tier2 connectivity
- perfSONAR plans for LHCONE
- Statistics for DC24

# ESnet Backbone planning - Q3 2023



# Trans-Atlantic & EU ring upgrades

- **Now In Production:**
  - 400G New York - London
- **Currently underway:**
  - 400G Boston - London (late fall)
  - 400G Boston - CERN (late fall)
  - 400G Europe Ring (late fall)
- **Trans-Atlantic capacity targets**
  - 1.5T in advance of DC24
  - ...
  - **3.2T\* in 2027**, well in advance of Run 4



# Transatlantic Cable Landing



# Routing Changes for Load Balancing

- We will have a mix of transatlantic link speeds for a while
- Routing protocols are not capacity-aware or load-aware
- Network Services (Michael Sinatra and others) have been working on a solution
  - ECMP across segment-routed LSPs
  - Put more LSPs on the 400G circuits
  - Initial results are positive
- Consequences:
  - Different flows between host pairs may take different paths with different latencies
    - Load balancing is per-flow, \*not\* per-packet, so we don't anticipate re-ordering or related problems
  - Traceroute may look odd (again)

# ESnet Cloud connectivity today

ESnet6 built physical network into major commercial facilities

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




# DOE Site 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)
  - FNAL - US CMS Tier 1
    - Current: 400G (2 x 100G + 2 x 100G)
    - Near Future: 800G (1 x 400G + 1 x 400G)
  - NERSC
    - Current: **1T** (2 x 400G + 2 x 100)



# 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>



# US Tier 2 Site Connectivity

- ESnet making the rounds talking to *every* US T2 site
- Gathering and helping synchronize plans from
  - Individual PI's
  - Departmental Support Staff
  - Campus IT & CIO
  - Regional Networks
  - R&E Exchange points
- Plans assessed:
  - Vanderbilt, SoX, CalTech, UCSD, CENIC, Nebraska, GPN, Purdue, Wisconsin, OmniPoP, New NET2, UMASSNET, MIT, NoX , LEARN, AGLT2
- Todo:
  - MWT2, UTA

# LHCONE Connectivity Upgrades

- In progress
  - Direct Connection to CENIC in Los Angeles
    - CalTech, UCSD
  - PACWAVE connectivity in Los Angeles
    - International connectivity
  - GPN 400G
    - Nebraska
  - NOX 100G in Boston
    - Supporting New NET2

# perfSONAR

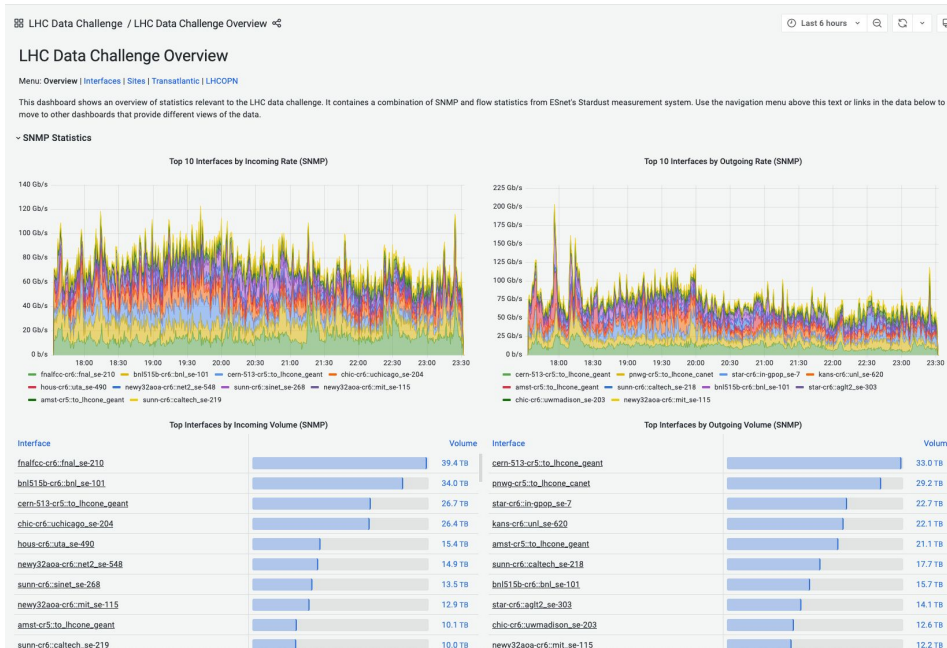
- ESnet perfSONAR fleet is nearing completion
  - Lots of supply chain and related logistics problems
  - Tuning challenges (different hosts than we originally specified)
- Problem: little LHCONE presence
  - Previously there were only a few LHCONE hosts in ESnet
  - Desire to avoid additional hardware deployments for LHCONE
- Solution: LHCONE interfaces on ESnet perfSONAR hosts
  - Nearing completion of IPv6 LHCONE interfaces on many/all hosts
  - Intent is to enable richer test and measurement on LHCONE than was previously available
  - Target availability: May 2023

# perfSONAR version 5

- perfSONAR Version 5 is out!
- Packages for major distros coming soon
- ESnet will upgrade its systems shortly
- We need to have a discussion about community tags in perfSONAR (can we claim 'LHCONE' in the lookup service?)
- We are looking forward to working with the community on testing Tier2 paths

# Statistics for DC24

- Something to think about for tomorrow's discussion
- ESnet's statistics systems can tell us some things about network data
- What do we need to do in preparation for DC24?
- Sites provide statistics - what do we need from the network in the middle?



# Moving Forward Together

- I'm presenting the work of many other people
- Several things are in motion now
  - Transatlantic and backbone upgrades
  - perfSONAR
  - Exchange point and related connectivity for Tier2s
- With the knowledge that resources are finite, what else needs to be done?



# Thanks!

[dart@es.net](mailto:dart@es.net)

Energy Sciences Network (ESnet)

Lawrence Berkeley National Laboratory

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

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

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

