ESnet
ENERGY SCIENCES NETWORK

ESnet Update
Fall 2021 LHCOPN/LHCONE Meeting

Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory

Dale W. Carder, ESnet Network Engineering

2021-10-11

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Agenda

- HEP Requirements review revisit
  - Ongoing engagement activities
- ESnet6 deployment status
- Data Challenge 1
- R&D Involvement
  - Packet marking: firefly collector
  - FABRIC

ESnet

HEP Requirement Review Process & Timeline

Program Coordination

Establish collaboration agreement between ESnet, ASCR/ESnet Program Management and HEP Program Management. Identify programs, facilities, PIs to participate in review.

1:1 Case Study Review

ESnet staff meets 1:1 with each case study author(s) 1:1 to review the content, ask questions, clarify information.

Analysis & Review

Analysis of case study materials and output of discussion from 1:1s and Focus Groups leading to a set of findings and recommendations.

Publication

Final publication of report including for distribution among partners/stakeholders.

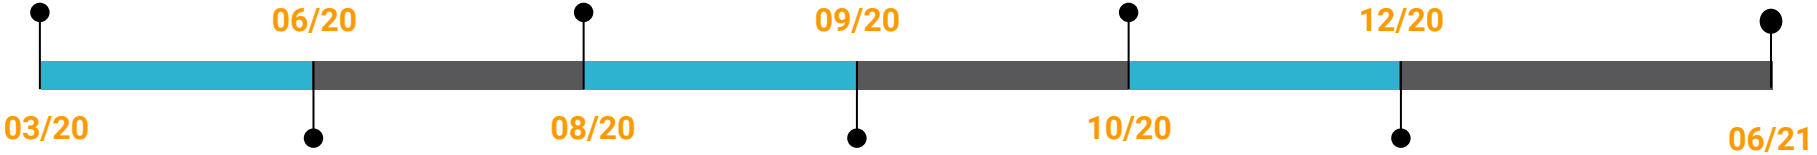

Case Study Development

Each PI/collaboration/facility develops a case study from a provided template.

Focus Groups

Cross-cut moderated discussion bringings together authors from across various case studies to share common challenges, themes.

Assembly & Editing

Creation of report, adding of editorial materials and background.

ESnet

Requirements Review: Outcomes

- Review published:
  - Zurawski, Jason, Brown, Benjamin, Carder, Dale, Colby, Eric, Dart, Eli, Miller, Ken, Patwa, Abid, Robinson, Kate, Rotman, Lauren, and Wiedlea, Andrew. High Energy Physics Network Requirements Review (Final Report, July-October 2020). United States: N. p., 2021. Web. doi:10.2172/1804717.
  - https://doi.org/10.2172/1804717
- Executive summary, Findings, and Actions can be found in the complete report:
  - https://www.osti.gov/servlets/purl/1804717
  - Pgs 1-30
- A 1-Year check-in with all parties in early 2022 will be planned to capture evolving needs and program changes.

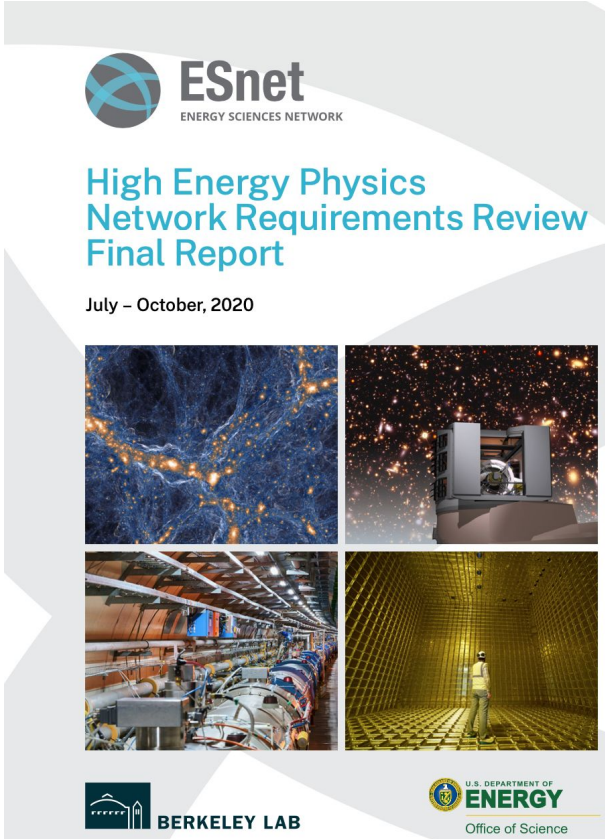

Preparing for Data Deluge (20 years ago …)

- Network throughput testing occurred years before the first beams (e.g. Apr 2006 on the chart)
- It is anticipated that a similar ramp up will be needed for HL-LHC
- Efforts like the DME, along with participating in “data challenges”, are critical

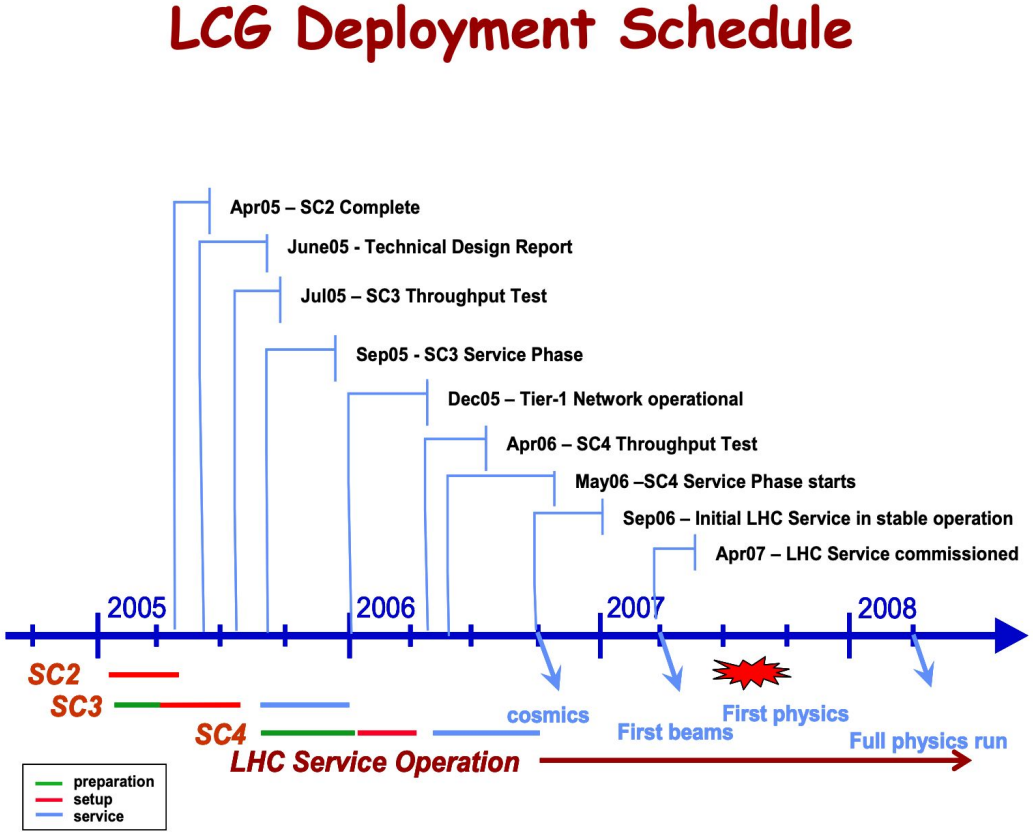

Courtesy Jamie Shiers @ CERN

Outyear planning

- Engagement w/ National, Regional, State, & Campus Networks
  - A holistic view is needed to consider all off the components involved to meet the expectations for large scale science.
- ESnet is trying to get the word out ahead of budgeting & planning, in multiple forums (and funding agencies).

- Give CIOs a heads-up now, ***and continuously*** through 2027.

ESnet

Outyear planning, Tactical examples

- Departmental data transfer nodes
- Campus edge (at least Science DMZ) connectivity upgrades.
- Transponders will need to be replaced.
- Optical systems will need to support very wide channels.
- Regional switching infrastructure port density

ESnet

Preparation for High-Luminosity LHC (2027)

- Current projections:
  - 400G connectivity needed per T2
    - recall there are numerous T2's in the US for ATLAS and CMS, served through regional exchanges
  - Needs likely driven by storage, w/ tradeoffs on placement.
    - ~5-12 PB typical now, possibly 20-40 PB per site in 2027
  - on-network caches?

- ***Data Challenges (full software stack)***
  - ***10% of the target 2021*** ***# should match Run 3***
  - ***30% in 2023*** ***# 2x100 per Tier2?***
  - ***60% in 2025*** ***# probably want 400G per Tier2***
  - ***100% in 2027*** ***# 400G or more per Tier2***

ESnet

ESnet6 the evolution of the ESnet network

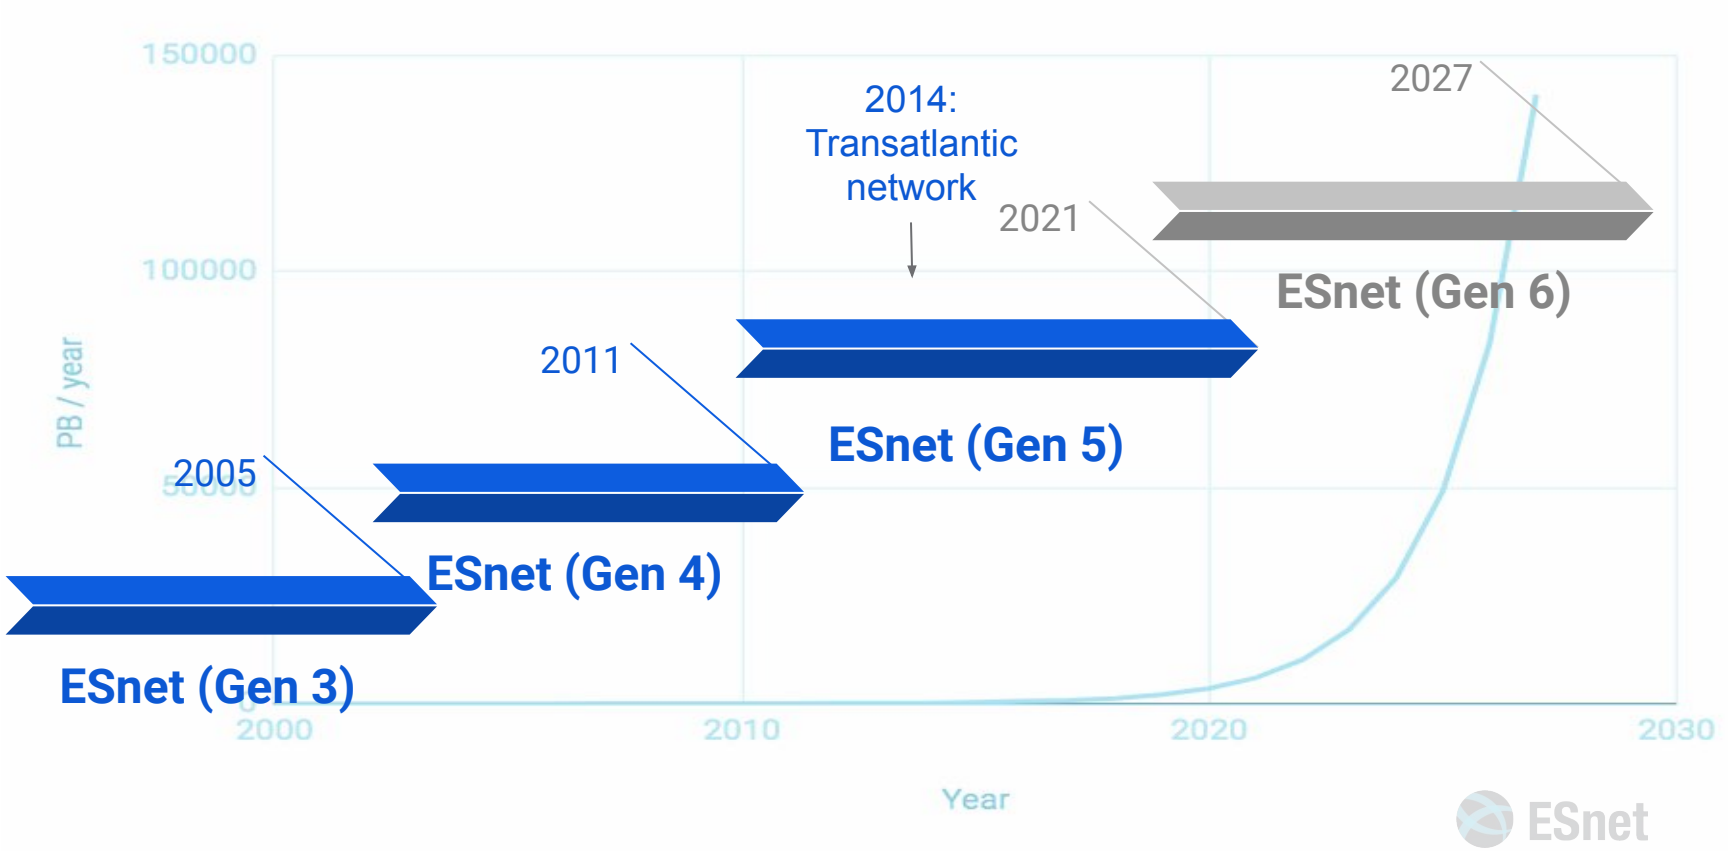

Phase 1 - Done!

15,000+ Miles of Fiber

300+ Colo Facilities

Infinera Flex-ILS Optical System

272 Amplifiers

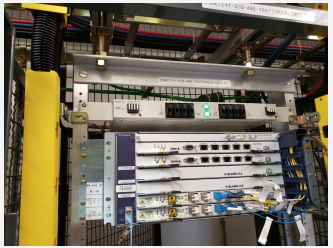

37 ROADMS

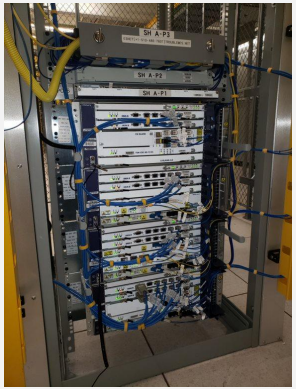

Layer

Assurance | Provisioning

Security | Analytics

Data Model

Snet5 Mgmt Net

Phase 1 Transponders

Open Line System

ESnet Fiber | ANL Fiber

ESnet Colo | ESnet6 Mgr

72 Groove Transponders

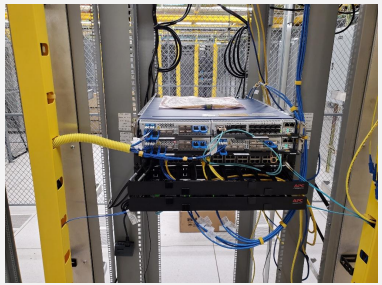

Phase 2 - In Progress

67 Routers
Nokia 7750-SR2

ESnet 6 - Low Touch

41 Transponder Chassis
Ciena Waveserver 5

27 Management Network

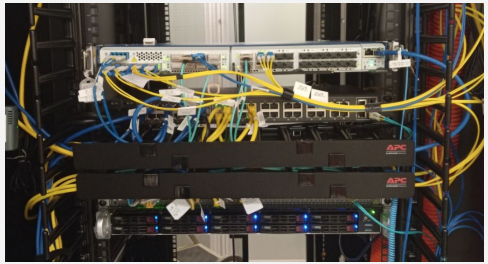

ESnet6 Routers

ESnet6 Mgmt Net

Packet

OLS to Brookhaven

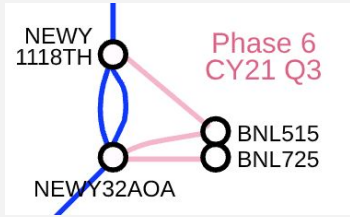

Optical

Phase 2 Transponde

Assurance

Provisioning

ANL Fiber

BNL OLS

Security

Analytics

ESnet6 Mgmt Network

Data Model

ANL Fiber

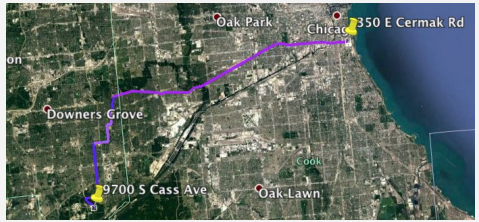

Existing | Phase 1 | Vendor | Vendor + Internal

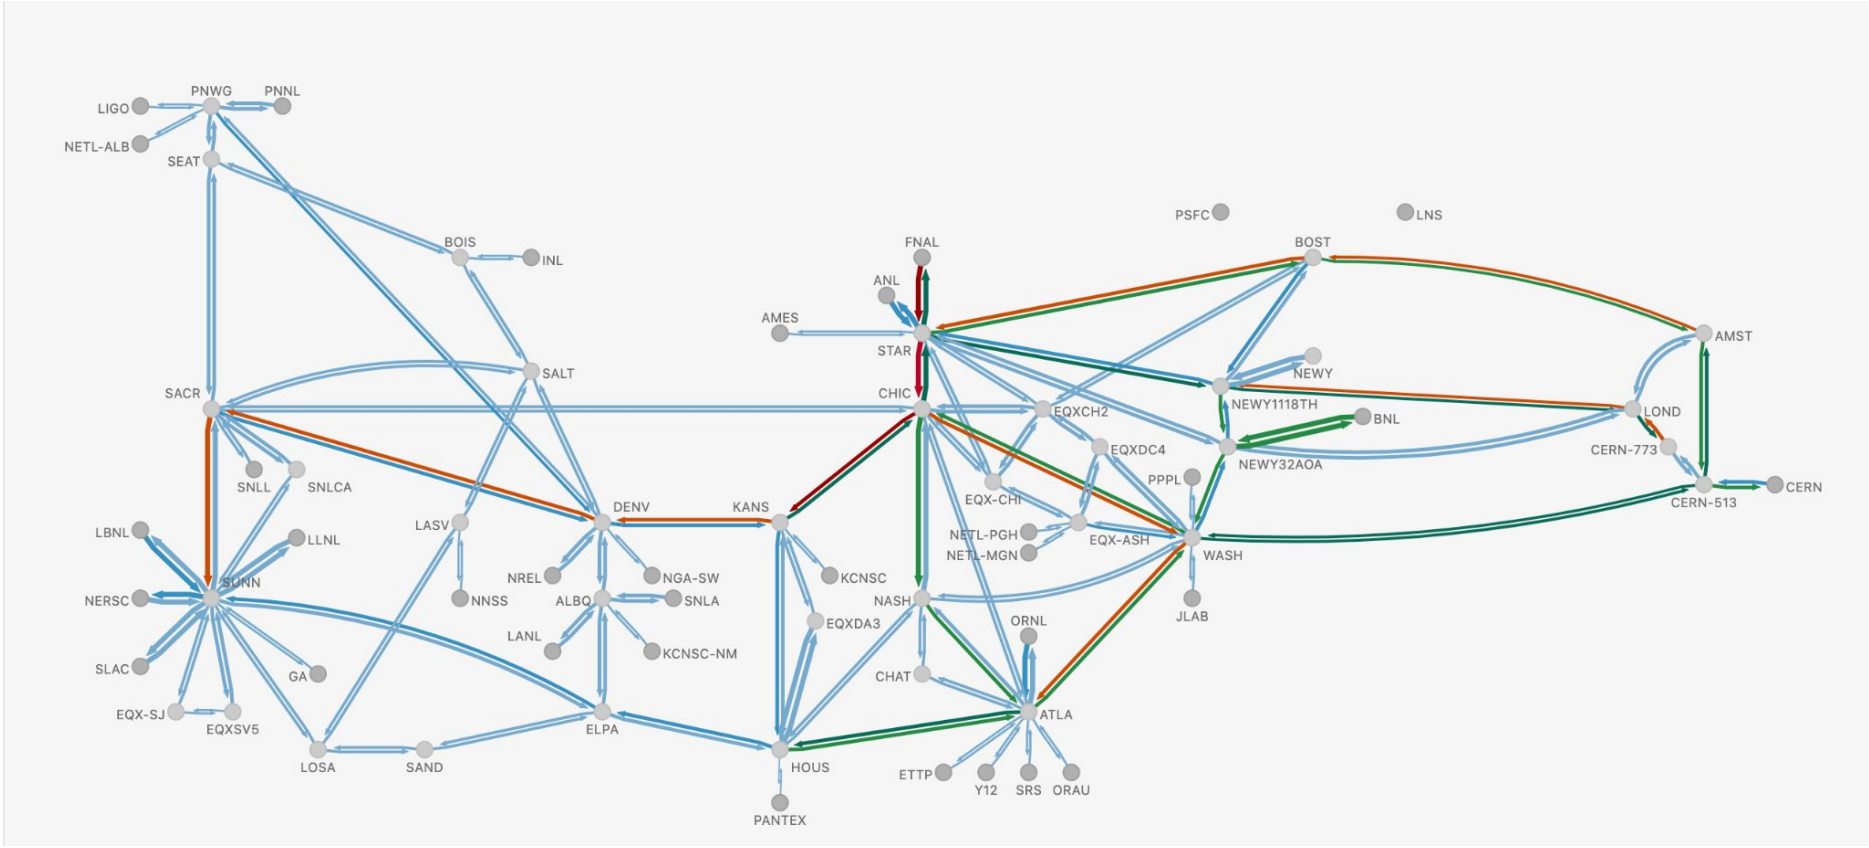

https://my.es.net

ESnet

ESnet6 Deployment Status - Routers

Deployment Forecast					
Deployment Group	Count	Physical Installation Start	Physical Install Complete	Nodes start Transition to Operations	Completion
Large Domestic	40	Jan 2021	Aug 2021	Jun 2021	Nov 2021
Small Domestic Group 1	23	Aug 2021	Apr 2022	Aug 2021	Apr 2022
PNNL & Small Domestic Group 2	5	Nov 2021	Apr 2022	Nov 2021	Apr 2022
Large Europe	4	Jan 2022	Apr 2022	Feb 2022	Jun 2022

ESnet

Delivering ESnet6 Automation

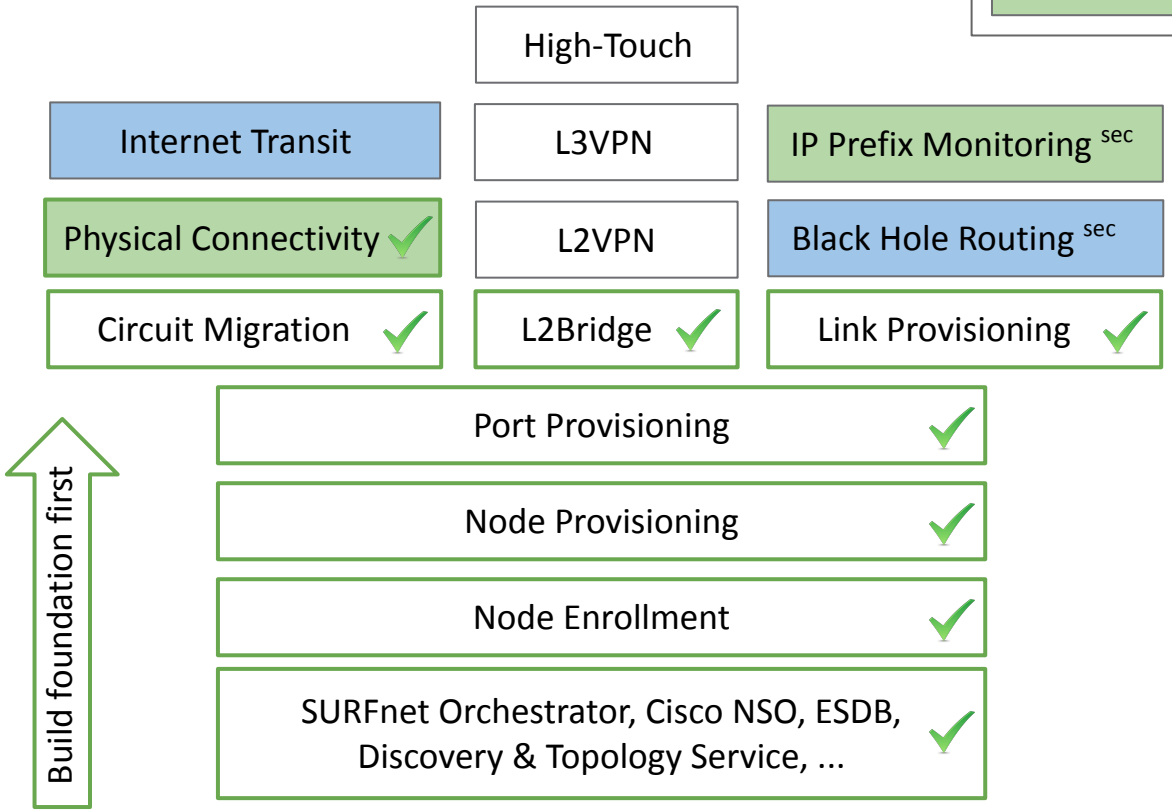

Point & click provisioning

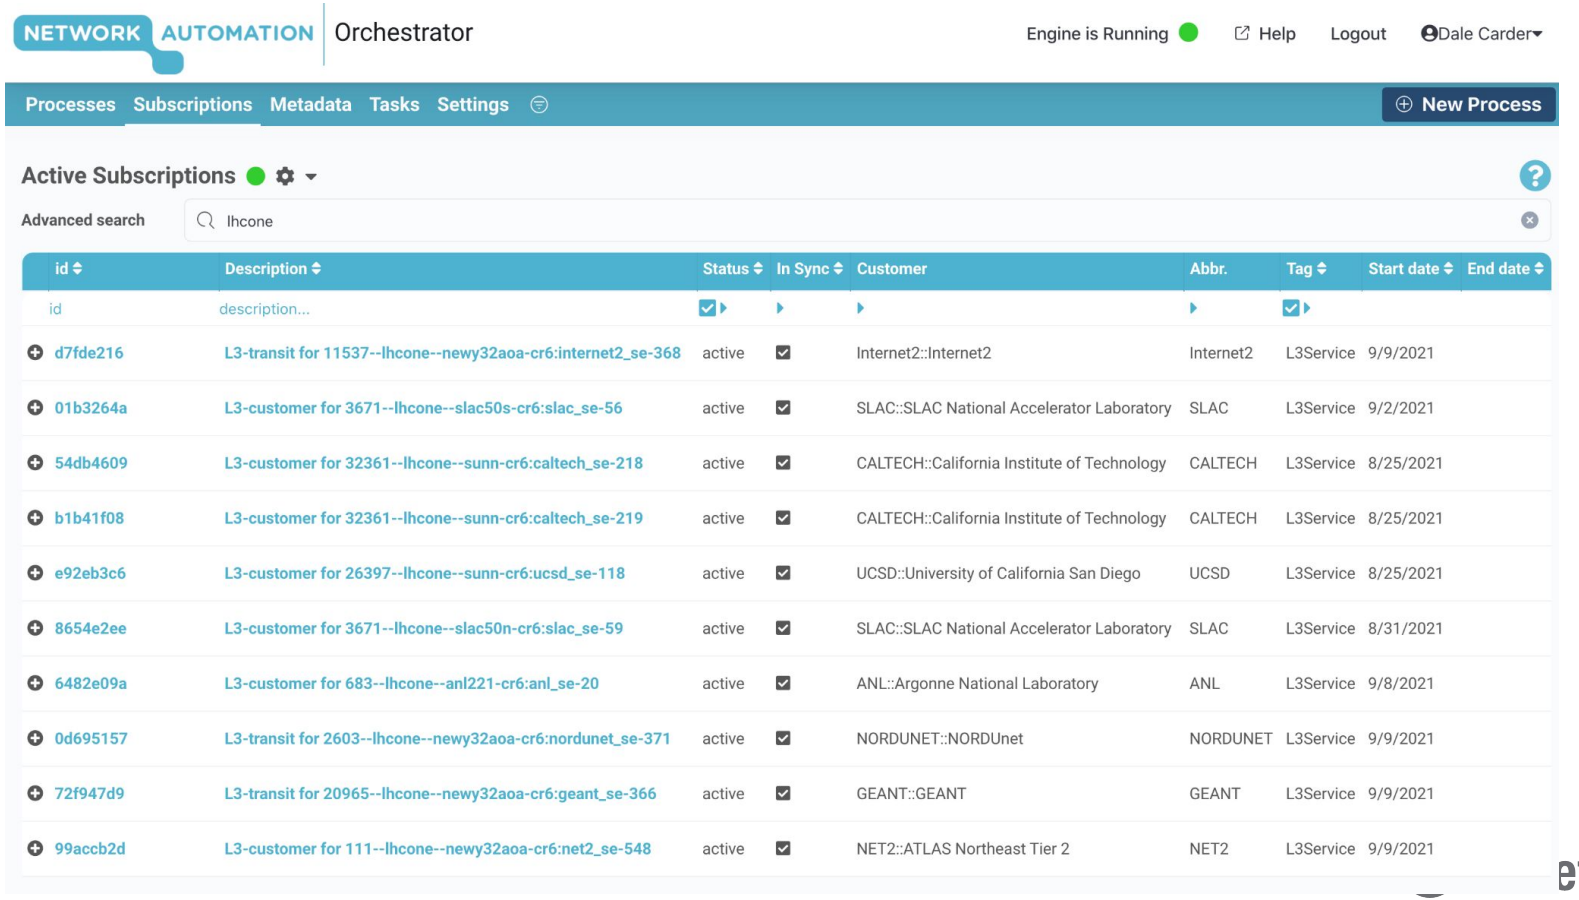

id	Description	Status	In Sync	Customer	Abbr.	Tag	Start date	End date
d7fde216	L3-transit for 11537--lhcone--newy32aoa-cr6:internet2_se-368	active	☑	Internet2::Internet2	Internet2	L3Service	9/9/2021	
01b3264a	L3-customer for 3671--lhcone--slac50s-cr6:slac_se-56	active	☑	SLAC::SLAC National Accelerator Laboratory	SLAC	L3Service	9/2/2021	
54db4609	L3-customer for 32361--lhcone--sunn-cr6:caltech_se-218	active	☑	CALTECH::California Institute of Technology	CALTECH	L3Service	8/25/2021	
b1b41f08	L3-customer for 32361--lhcone--sunn-cr6:caltech_se-219	active	☑	CALTECH::California Institute of Technology	CALTECH	L3Service	8/25/2021	
e92eb3c6	L3-customer for 26397--lhcone--sunn-cr6:ucsd_se-118	active	☑	UCSD::University of California San Diego	UCSD	L3Service	8/25/2021	
8654e2ee	L3-customer for 3671--lhcone--slac50n-cr6:slac_se-59	active	☑	SLAC::SLAC National Accelerator Laboratory	SLAC	L3Service	8/31/2021	
6482e09a	L3-customer for 683--lhcone--anl221-cr6:anl_se-20	active	☑	ANL::Argonne National Laboratory	ANL	L3Service	9/8/2021	
0d695157	L3-transit for 2603--lhcone--newy32aoa-cr6:nordunet_se-371	active	☑	NORDUNET::NORDUnet	NORDUNET	L3Service	9/9/2021	
72f947d9	L3-transit for 20965--lhcone--newy32aoa-cr6:geant_se-366	active	☑	GEANT::GEANT	GEANT	L3Service	9/9/2021	
99accb2d	L3-customer for 111--lhcone--newy32aoa-cr6:net2_se-548	active	☑	NET2::ATLAS Northeast Tier 2	NET2	L3Service	9/9/2021	

Data Challenge 1 - Lots Of Packets

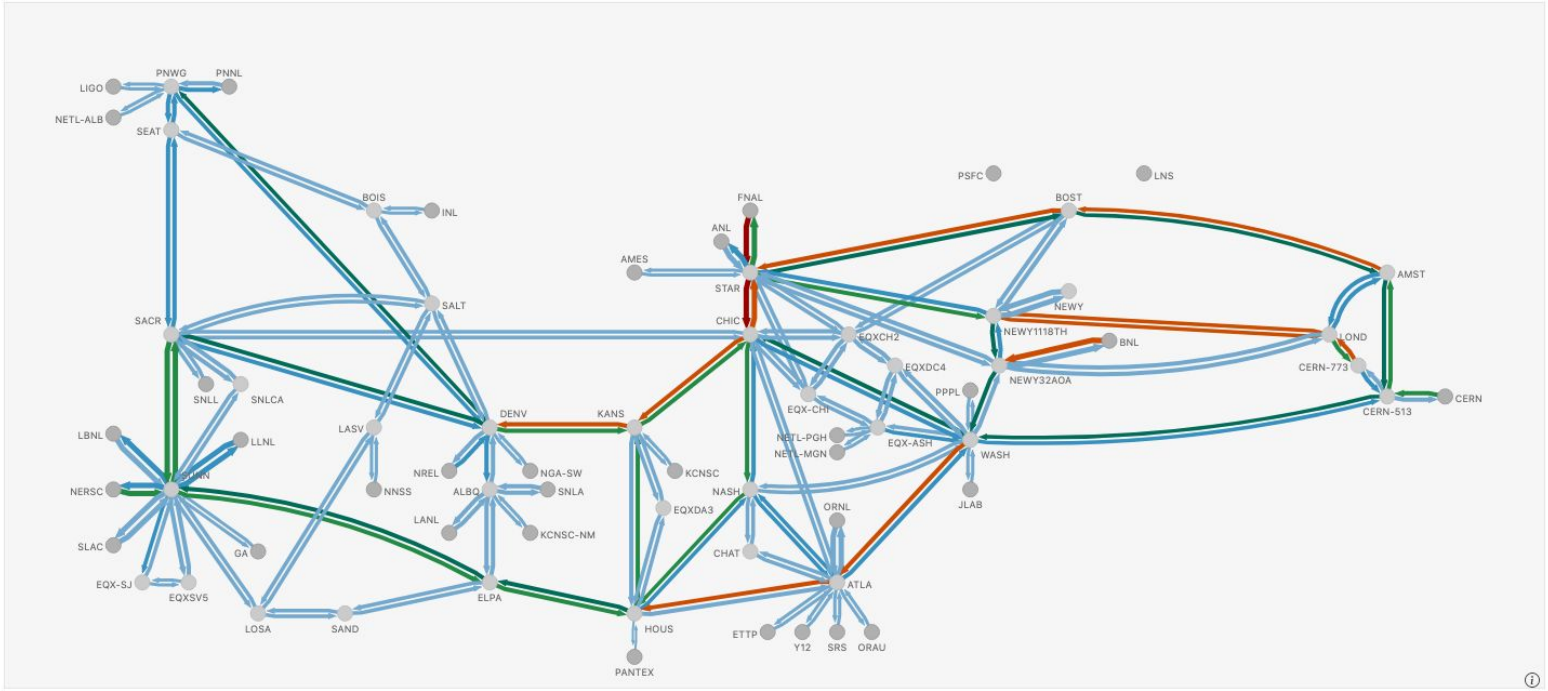

Why transatlantic outages take a while

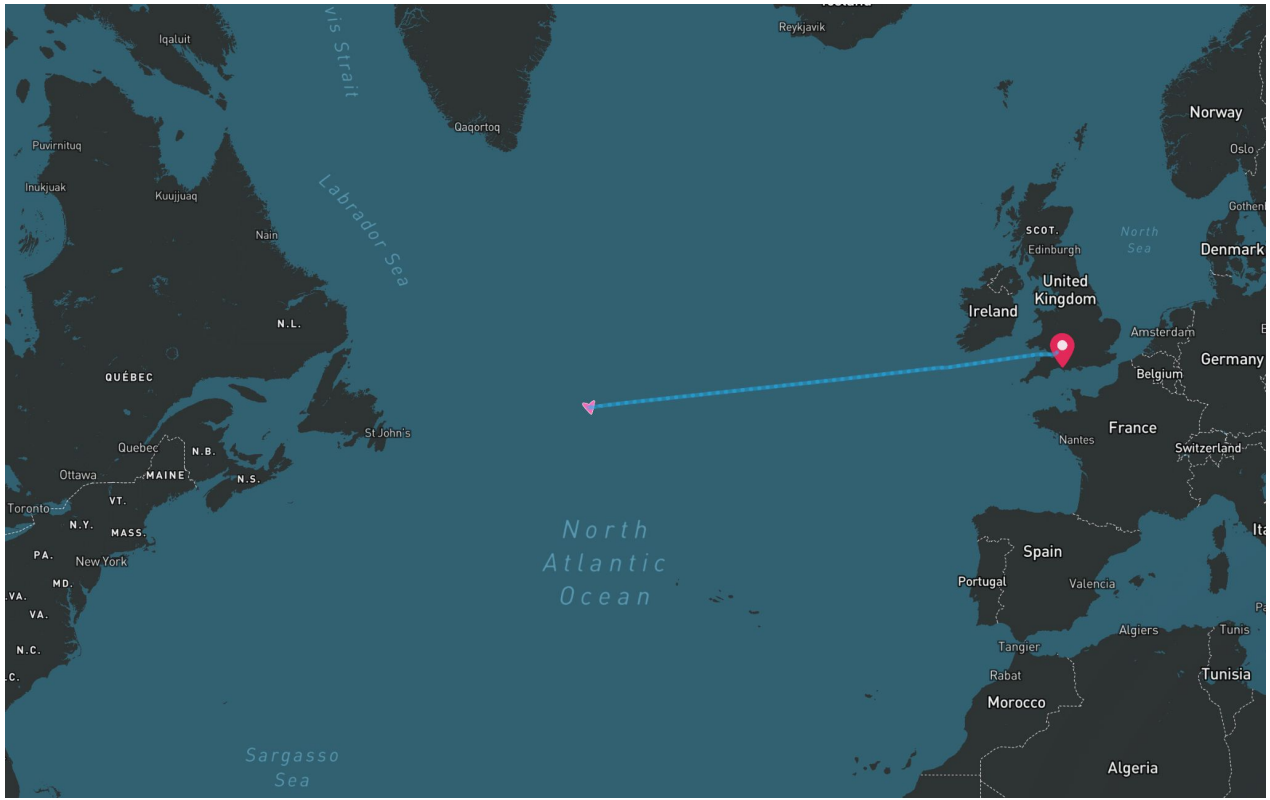

ESnet

Data Challenge 1: Traffic Rebalancing

- ESnet6 transition is partially complete
- CHIC-KANS path was saturated by LHC plus other traffic (Climate)
- Brought ESnet6 capacity into production opportunistically
- ESnet6 capacity will be much greater soon

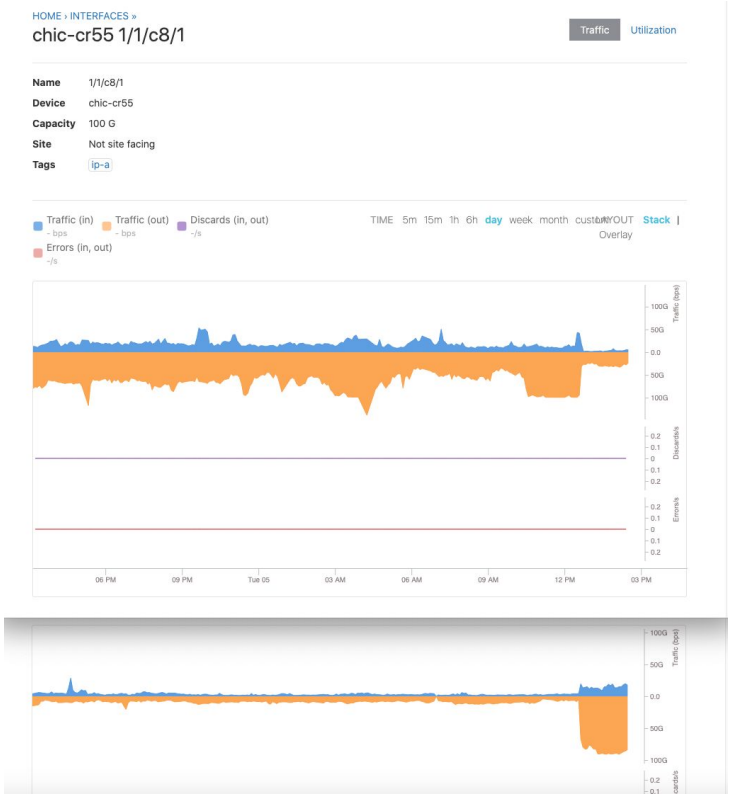

Data Analysis Using Cloud

- LHC experiments can use Cloud effectively (as demonstrated in the past)
- Other collaborations are also exploring Cloud
- Example: DUNE
  - CPU compute and data are at FNAL
  - GPUs in Google
- Ran concurrently with Data Challenge

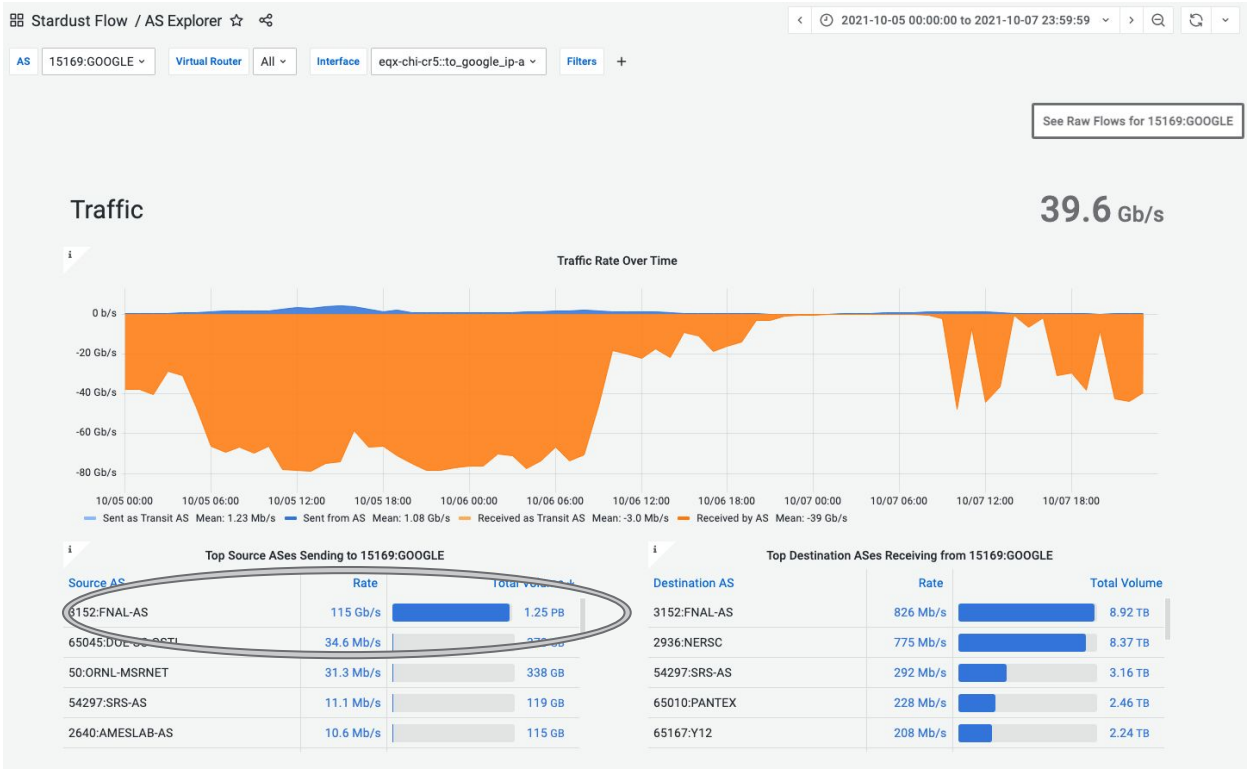

ESnet Stardust: Network Stats

https://public.stardust.es.net/d/IkFCB5Hnk/lhc-data-challenge-overview?orgId=1

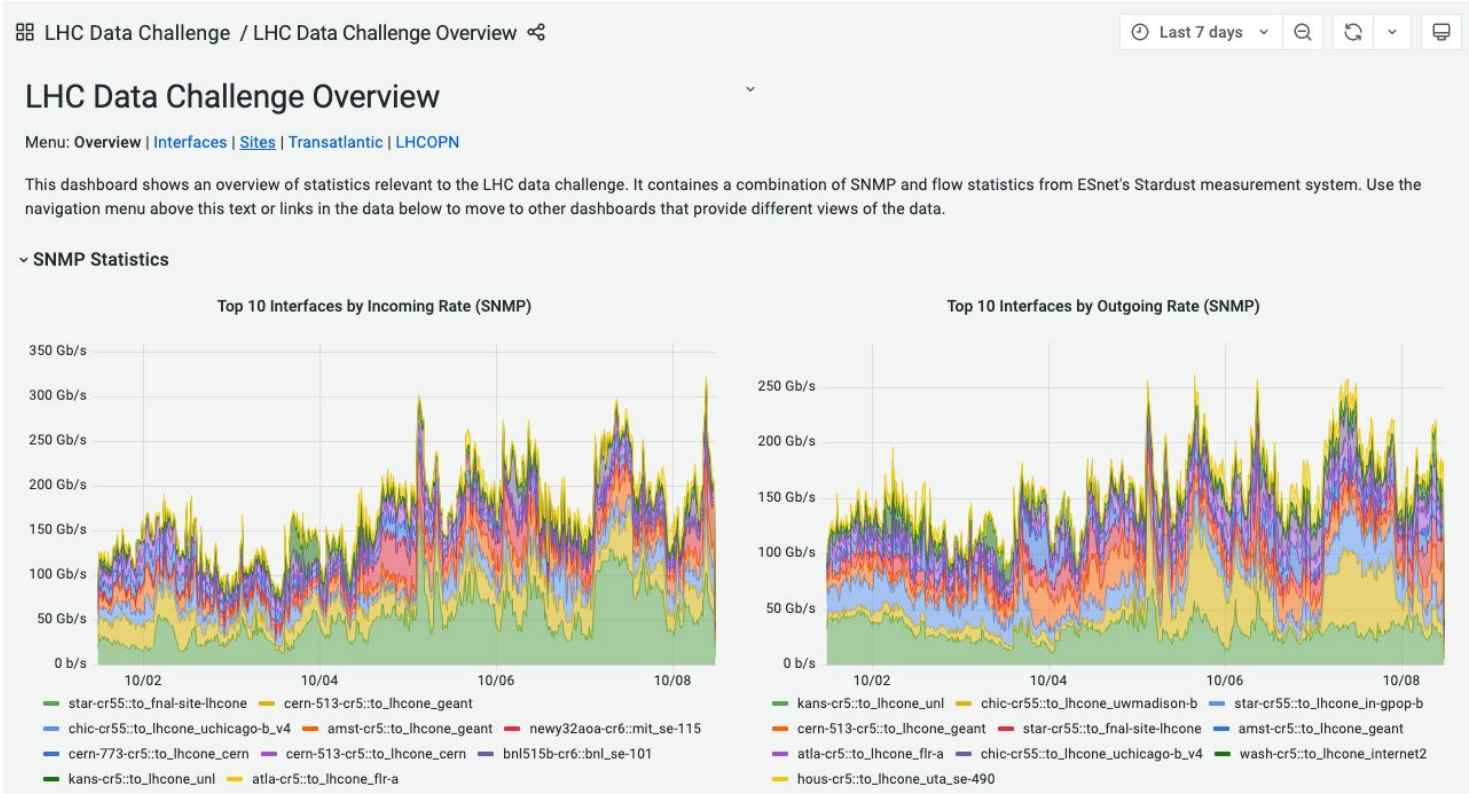

ESnet

ESnet Stardust: Network Stats

https://public.stardust.es.net/d/IkFCB5Hnk/lhc-data-challenge-overview?orgId=1

Currently experiencing likely software bug w/ ipfix export

- "only" affecting MPLS traffic encapsulated through ESnet6 routers
- looks like v6 fields truncated and stuffed into v4 ipfix template, and/or other bit alignment issue

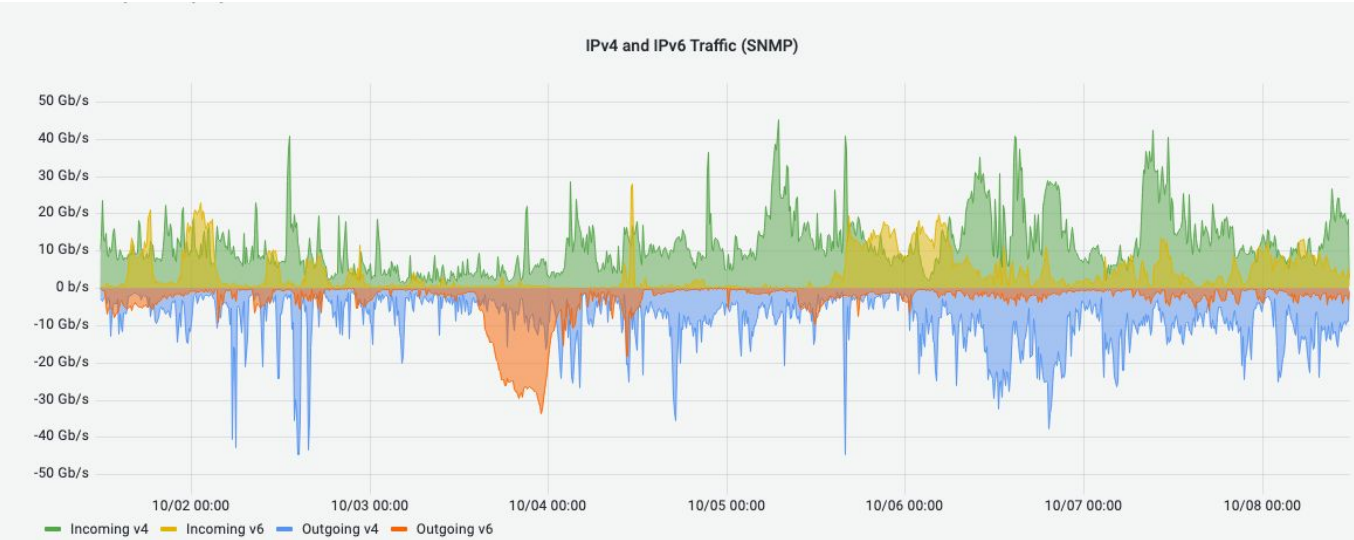

R&D Efforts

Firefly messages

- Part of the Network Technical Working Group organized by Marian Babik and Shawn McKee
- Complementary to packet marking
- Easier/faster to implement than IPv6 flow labels in the short term
- Running in a limited/trial capacity for Data Challenge 1
  - Goal is to get some experience with data export and correlation
  - Richard Cziva (ESnet) working on data analysis
- More coming from Shawn and Marian
- Collaborators: Stacey Sheldon, Richard Cziva, Dale Carder, Shawn McKee, Marian Babik, Garhan Attebury, Eric Lancon + BNL networking, Andy Hanushevsky, others - thanks!

ESnet

ESnet is receiving firefly messages!

- We are receiving firefly messages from multiple Data Challenge sites
  - CERN
  - ATLAS: BNL, AGLT2
  - CMS: Caltech, UNL
- Firefly json over syslog/udp transport


```
Oct  6 04:09:01 dcdoor02 1 2021-10-06T04:09:01.044638+00:00 dcdoor02.usatlas.bnl.gov flowd - firefly-json - {"flow-id": {"protocol": "tcp", "afi": "ipv6", "dst-ip": "2620:0:210:1::41", "src-port": 443, "src-ip": "2620:0:210:1::de", "dst-port": 38780}, "version": 1, "flow-lifecycle": {"state": "end", "end-time": "2021-10-06T04:09:01.044259+00:00", "start-time": "2021-10-06T04:08:49.838699+00:00", "current-time": "2021-10-06T04:09:01.044656+00:00"}, "context": {"experiment-id": 16, "application": "flowd v0.1.2", "activity-id": 1}}
```


ESnet

FABRIC progress - Deployment Status
Phase 1 Deployment (September 2021)

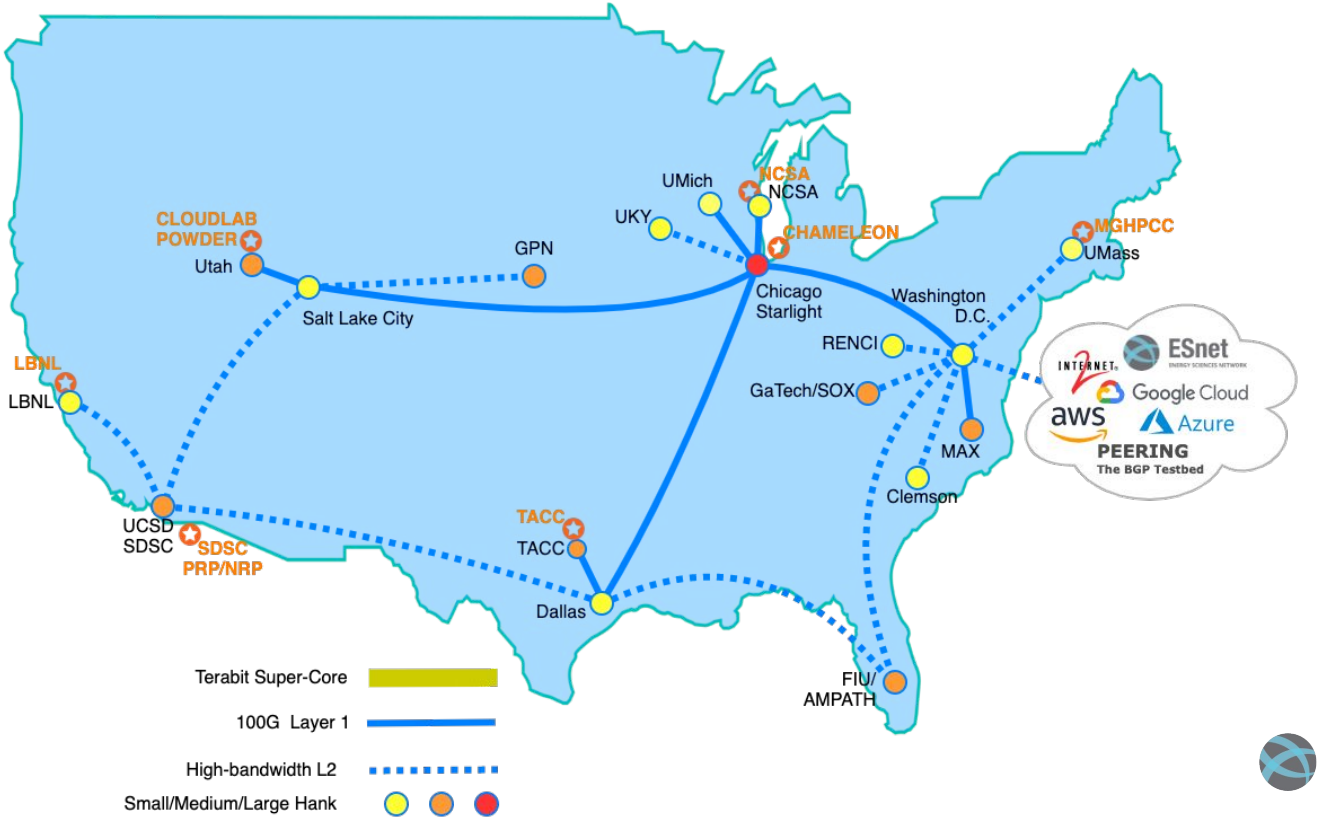

ESnet

FAB (FABRIC Across Borders)

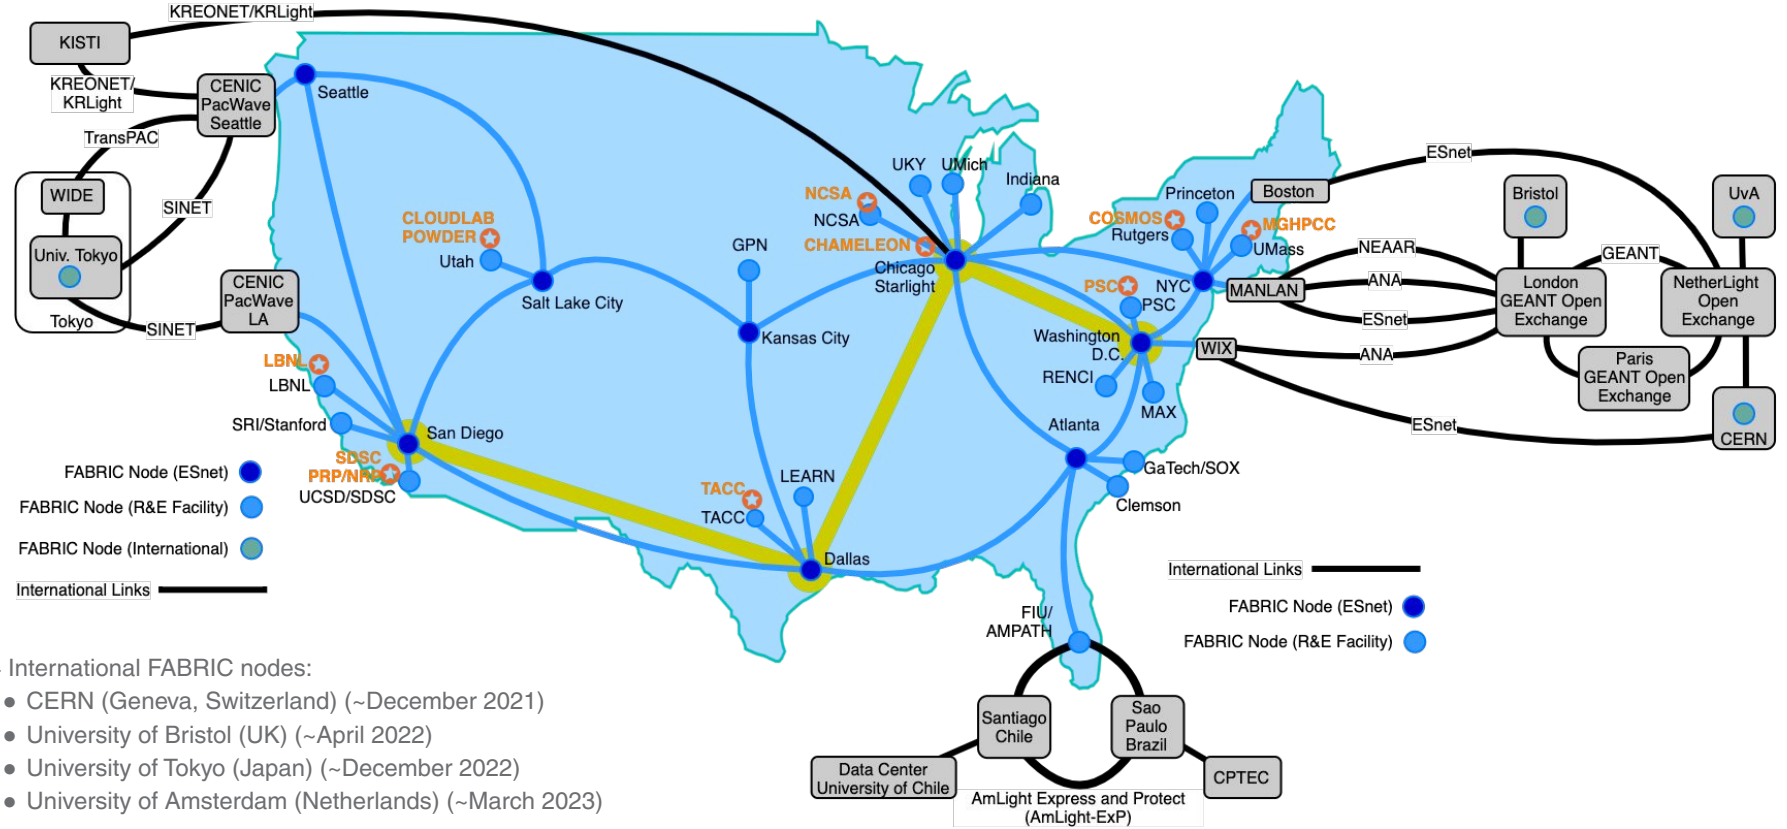

4 International FABRIC nodes:

- CERN (Geneva, Switzerland) (~December 2021)
- University of Bristol (UK) (~April 2022)
- University of Tokyo (Japan) (~December 2022)
- University of Amsterdam (Netherlands) (~March 2023)

ESnet