ESnet
ENERGY SCIENCES NETWORK

Stardust

ESnet Network Measurement Environment

Andrew Lake <andy@es.net>

March 8, 2022

Measuring Networks

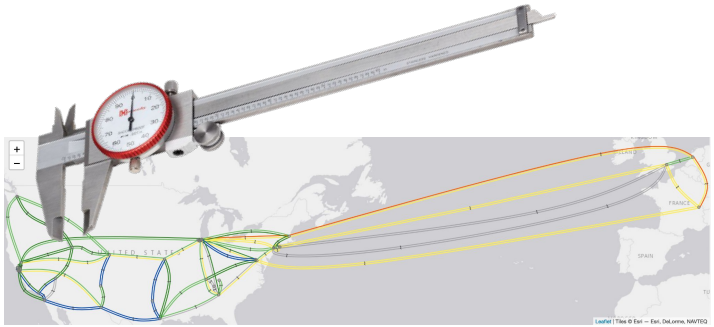

- To effectively design and operate networks, you need to measure their use and performance.
  - How much unused capacity do we have?
  - What data is traversing the network?
  - What level of performance are users seeing?

- Each of these questions requires a different measurement approach
  - Interface counters let us see how much capacity is used on our networks
  - Network Flow accounting data acts as phone call records telling us who is talking to who
  - Active testing is used to measure ability to send data through the network quickly

ESnet

Telling stories with empirical data

- We have lots of different types of stories to tell using the data we collect
  - **Operational Stories:**
    - ***What*** *just happened to the network in the northeast?*
    - ***Who*** *is the source of that giant surge in new traffic?*
  - **Planning/Engineering Stories:**
    - **When** will we run out of capacity to Europe?
    - **Where** are we seeing performance degradation?
  - **Community Engagement Stories:**
    - **How** has a project's use of the network changed in last year, and is there anything we can do facilitate better data transfer?
  - **Value Proposition Stories:**
    - ***Why*** *is my organization awesome? full technicolor.*
- A good story is more than points on a graph, it needs to relate back to the real world things people care about: *who, what, where, when, why, how*
- At ESnet, we augment our measurements with this context to more effectively tell stories
  - This contextual metadata comes from numerous sources itself

ESnet

Stardust: A new approach to help us tell stories

- Stardust was conceived to provide a flexible user empowering multi-data source network measurement environment with strong metadata support.

- Interdisciplinary team from within ESnet
  - Software, Systems and Network engineers
  - 20+ years of industry and NREN experience
  - Multiple service “generations”
  - Extensive open source collaboration experience

- Project Est. Fall 2020, service in production by April 2021
  - Initiated to address measurement gaps identified by internal review.
  - Build PoC and evaluate against commercial alternatives.
  - First objective was to replace legacy traffic analysis system.
  - Fully featured measurement environment is current focus

- Using a pragmatic approach tailored to our strengths
  - Build on past collaborations, integrating where we can, and innovate where it makes sense.
  - Generate effective measurements and combine with relevant metadata.
  - Develop effective visuals suited to our problem domain.

ESnet

Stardust

Network Measurement and Analysis for ESnet

Extensible / Open Architecture

NSF NetSAGE project derived

Leverage Open Source components where we can, and innovate where it makes a difference.

Multiple access methods

Dashboards, Indexed APIs and “Raw”

Today, we are focused on users creating and sharing visual dashboards.

In the future, we expect direct programmatic access will become increasingly common for ML work and external collaboration.

Multi Datasource

Extensible and Open

Traffic Accounting, Link and Resource Use, Performance Testing Results, Others not yet invented.

Each has its own set of measurements to which we add a common core set of metadata.

Flexible aggregation

Time Frame and Relationships

The metadata we add to measurements is used to summarize data to tell stories, and having common metadata allows us to use multiple data sets in a story.

- How are researchers moving science data and how has that changed over the last 3 years?
- What just caused that huge spike in traffic on the links to europe in the last 15 minutes and is that likely impacting data transfers?

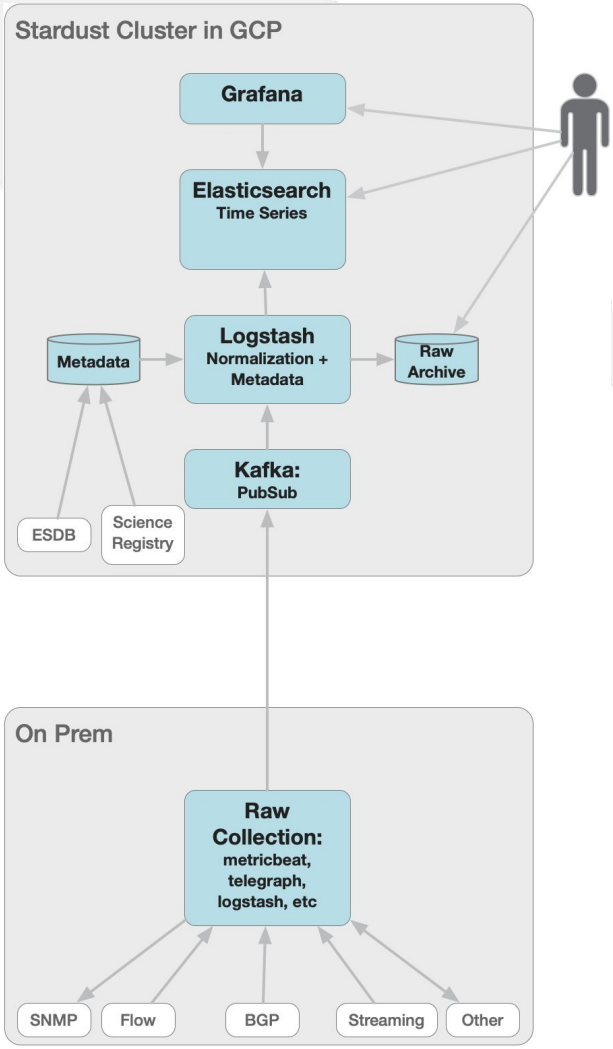

ESnet

Stardust Dashboards

- Combination of curated and custom dashboards used internally by ESnet staff members to visualize and contextualize measurements

Not This

This

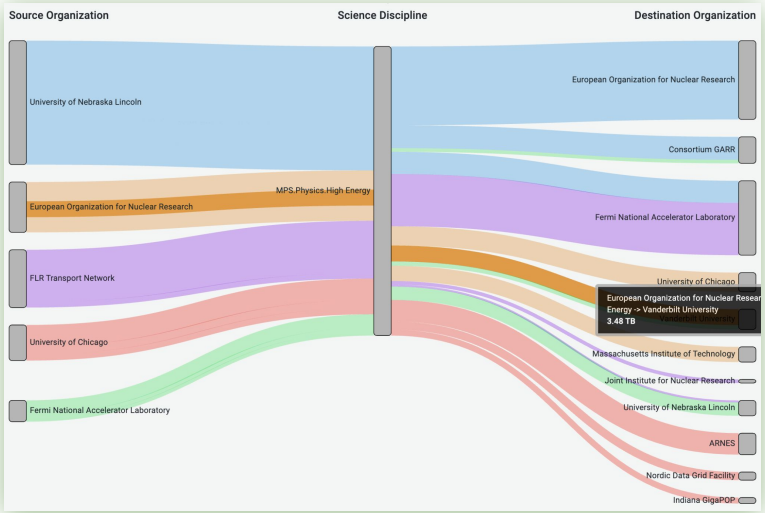

Latest Developments

- SNMP data when production Q4 2021, legacy Netbeam system to be retired end of March 2022

- Science Engagement Map
  - An interactive network display tailored to external users interested in value proposition and science engagement

- Integration with experimental data sets
  - High Touch - packet analysis done in hardware
  - LHC Firefly packets - additional flows

- More data sets coming soon
  - perfSONAR
  - L2 flow records

ESnet

Live Demo

Stardust Flow Dashboards

Breakdown by Application:

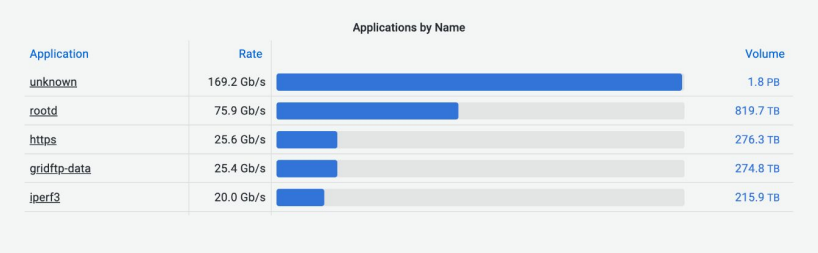

Breakdown by Interface:

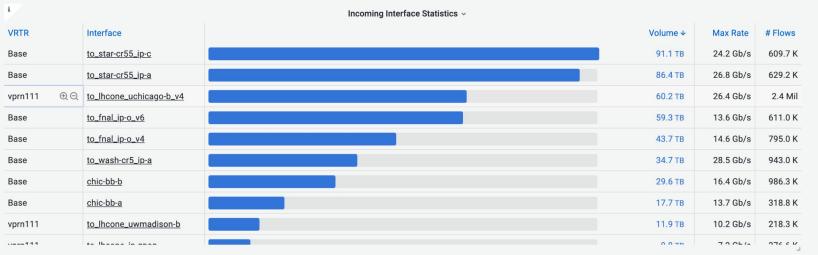

Breakdown by Peer and AS Path:

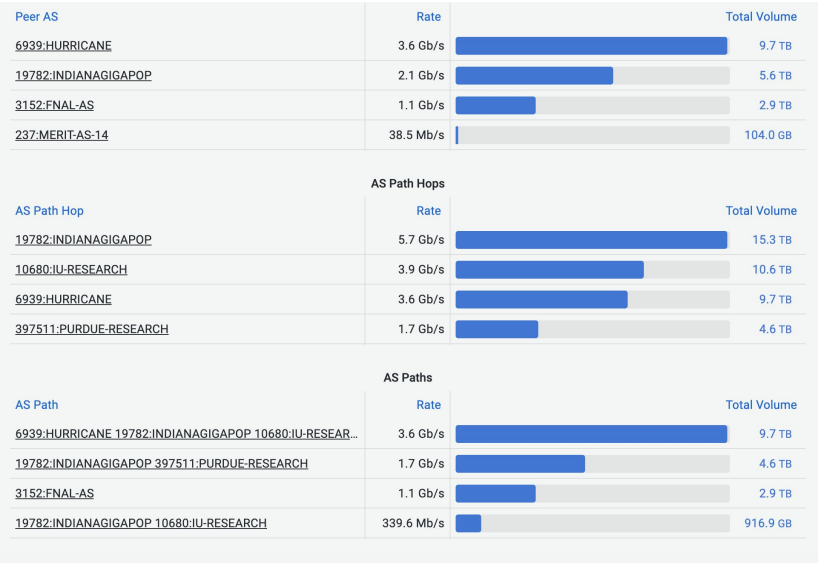

net

Stardust SNMP Dashboards

Mapping router internals

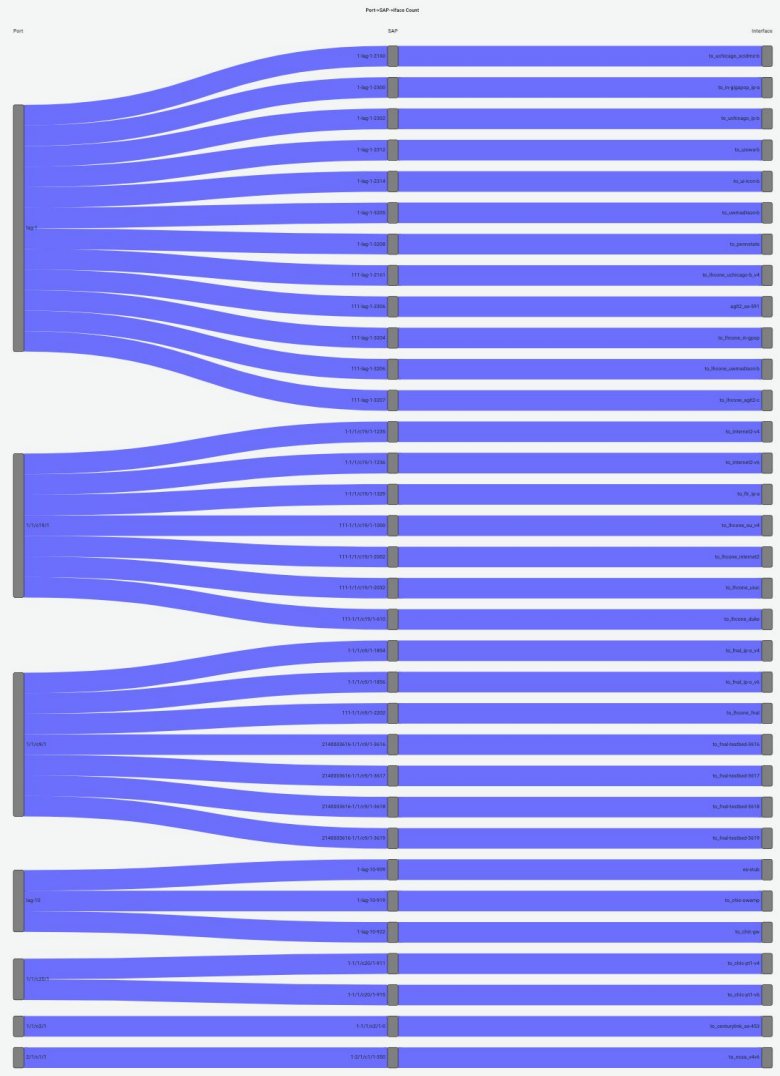

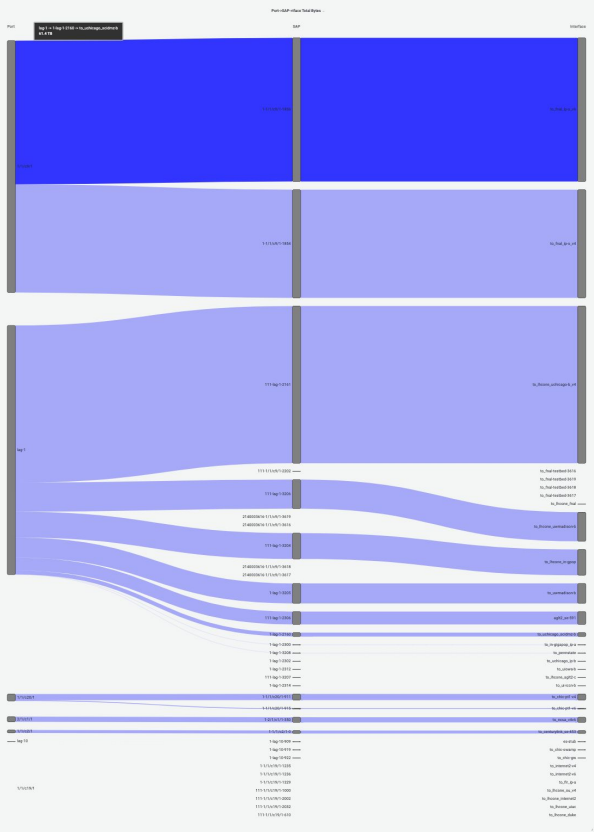

ESnet

Stardust Transponder Dashboards

Metrics Across Network:

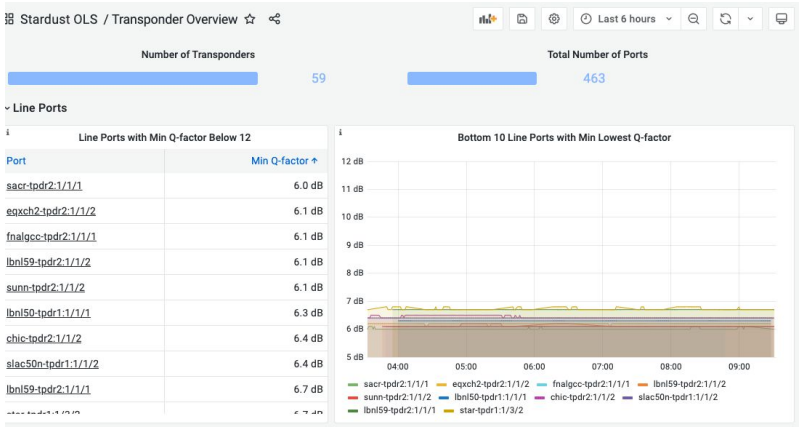

Port Statistics:

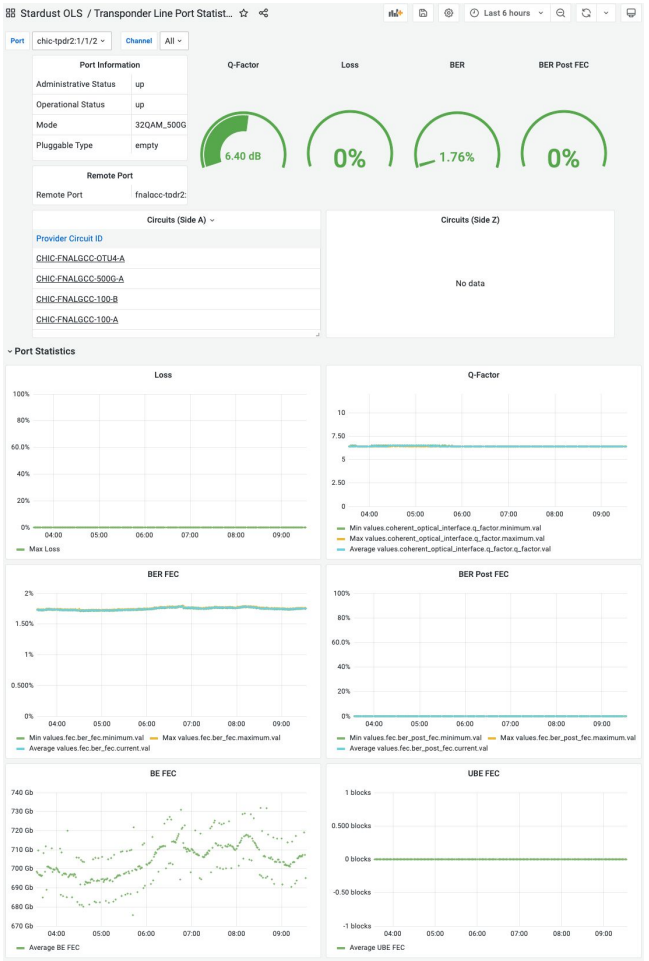

Circuit Statistics:

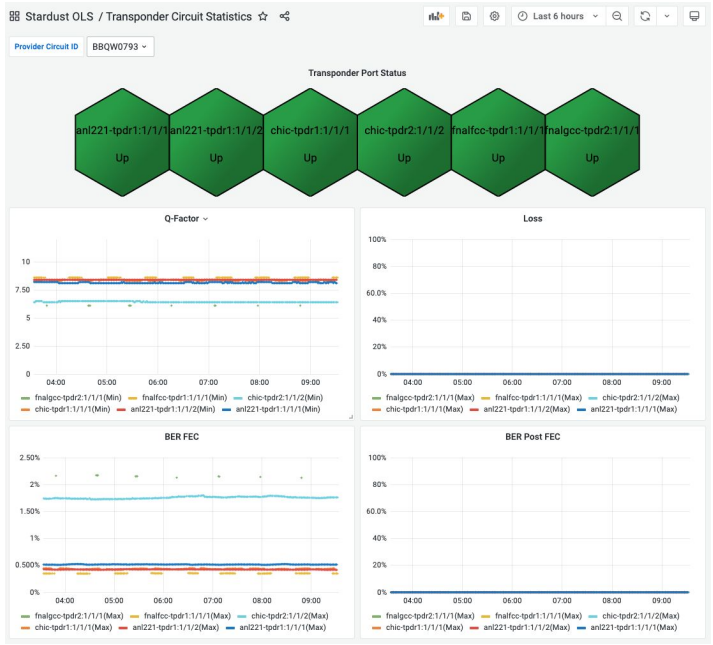

ESnet

Engagemap:

Network Example

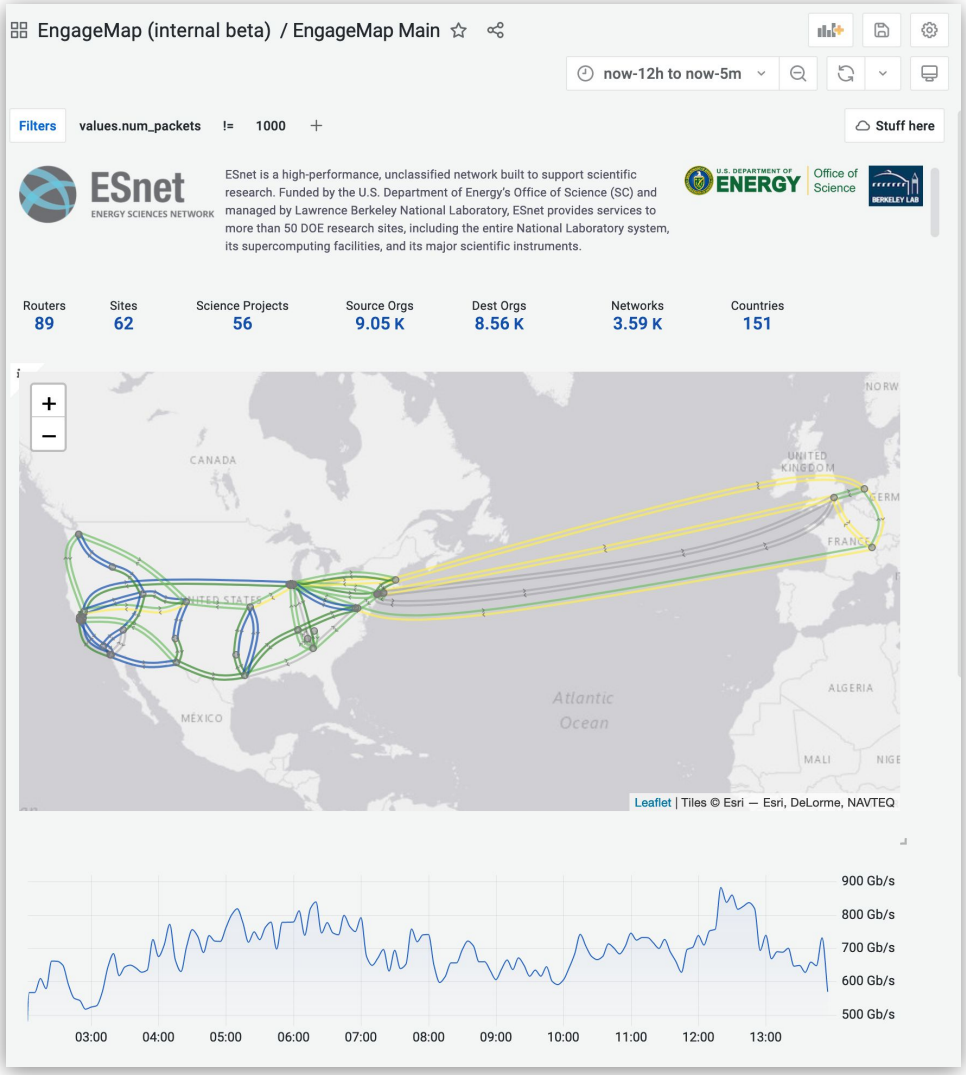

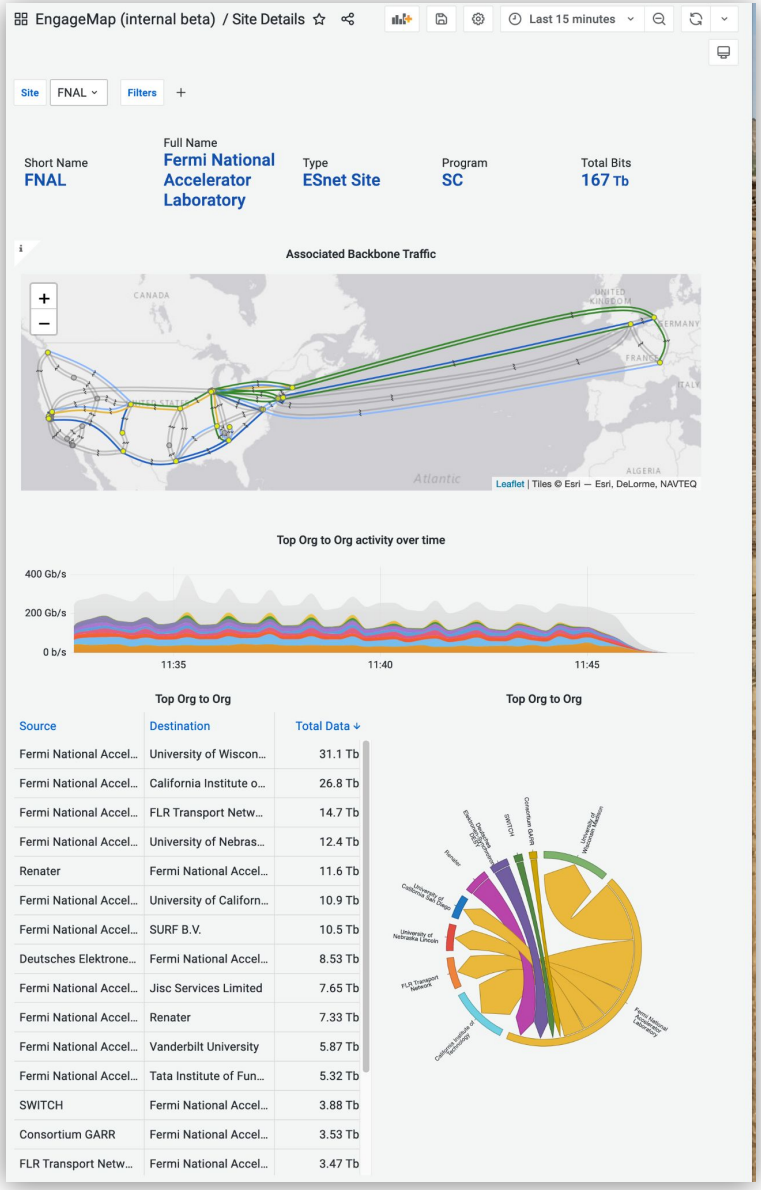

ESnet

Engagemap:

Non-network example

- https://science.osti.gov/User-Facilities/User-Statistics/By-Institution
- Explore ability to import and display OSTI User Survey data
- Eventually we want explore useful ways to combine this with our science registry and network traffic data

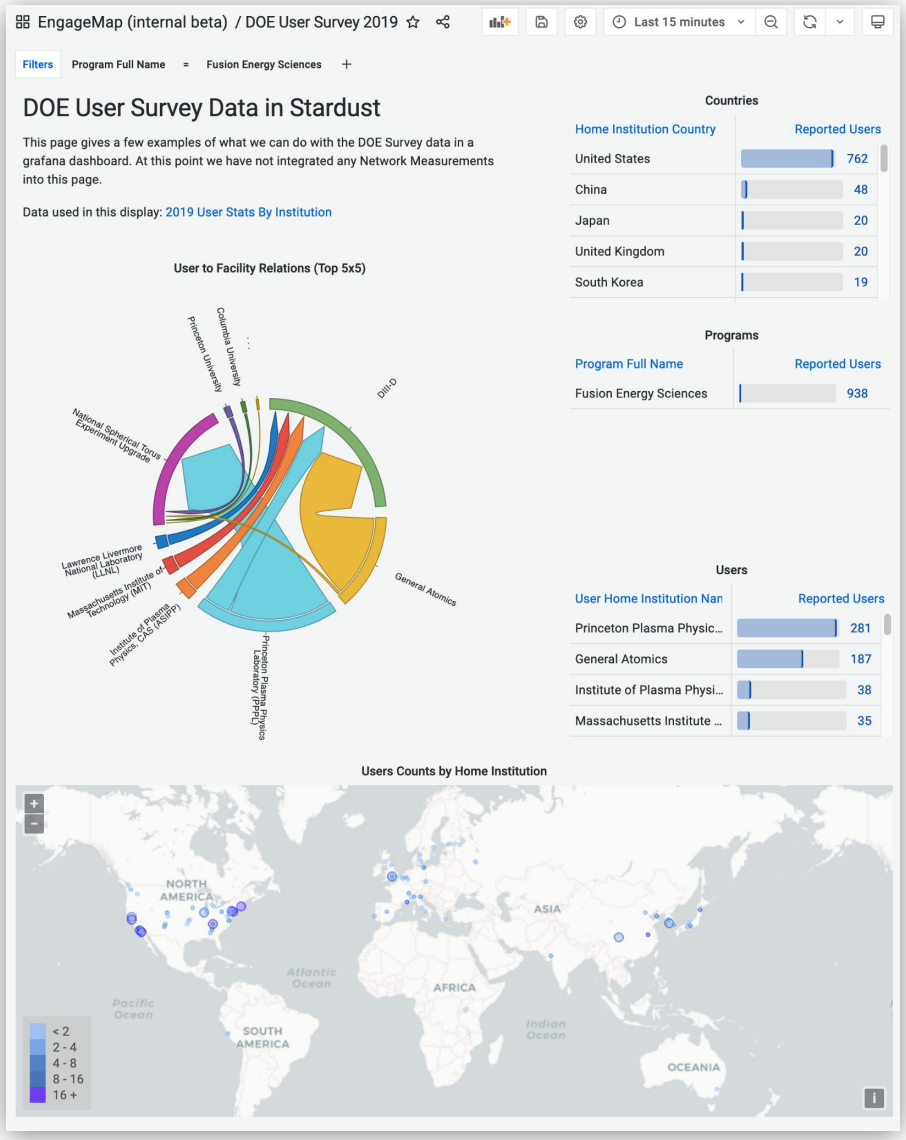

ESnet

LHC Data Challenge

- First of 4 “data challenges” conducted from Oct 4-7, 2021 as readiness test for LHC high luminosity era
- With BNL (ATLAS) and FNAL (CMS) as US Tier 1 sites, ESnet a critical part of network path
- Specific ask was to provide a view of network metrics relevant to LHC sites
- Dashboards here

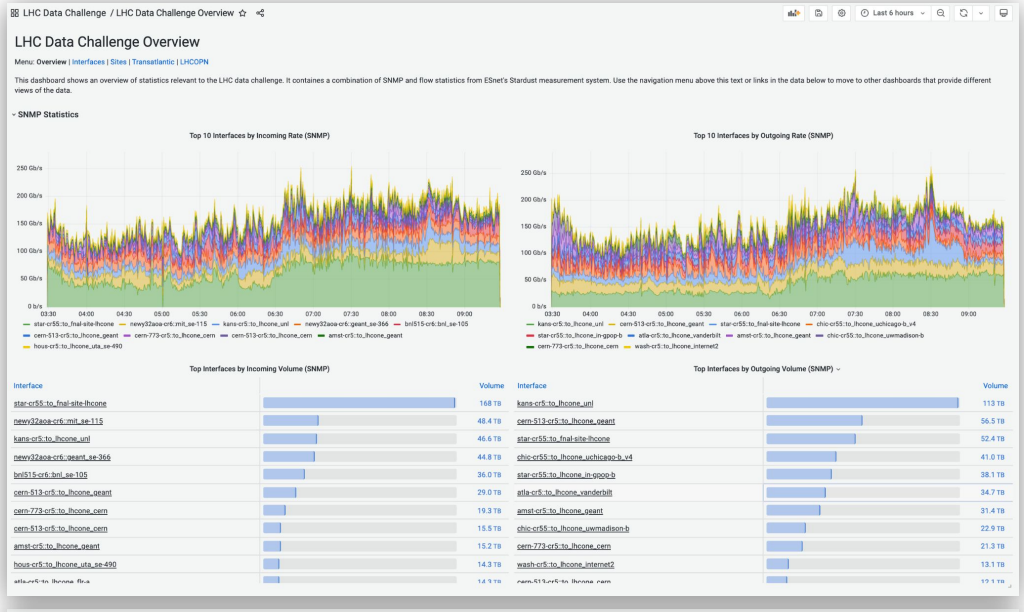

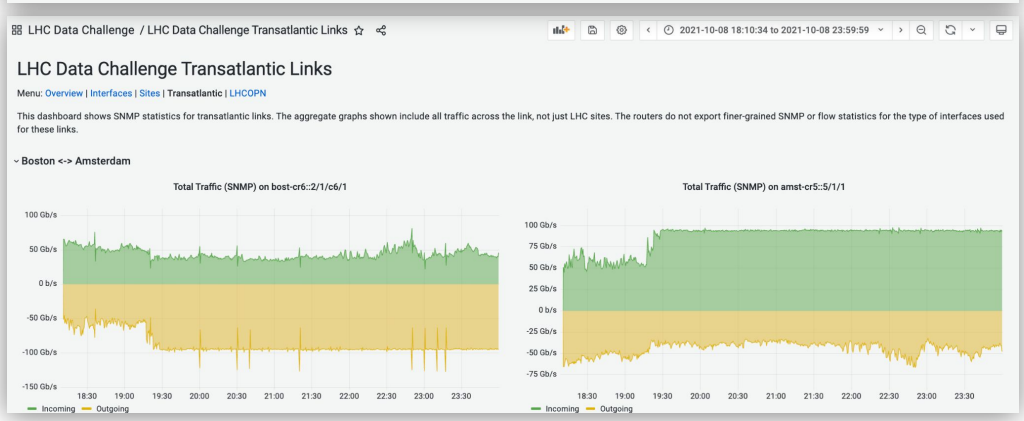

ESnet

Future Plans

- Refine Stardust capabilities
  - Provide full coverage for established measurement data types
  - Expand our data sets when opportunities arise and it makes sense to
  - Continue to contribute innovative visualizations
    - Network Map, Sankey, Chord, Slope and other diagrams.

- Expand R&E community collaborations and technical leadership
  - Jointly contribute to relevant metadata registries
  - Share our code, patterns and lessons learned
  - Support ML/AI research through *appropriately shared* network data

- Improve support for domain sciences that rely on ESnet
  - Provide means to appropriately share relevant data bidirectionally
  - Create project focused reporting / dashboards

ESnet