# WLCG Network Monitoring and Analytics Update

Shawn McKee, Marian Babik

Fall HEPiX (virtual), Oct 28, 2021

on behalf of WLCG Network Throughput WG

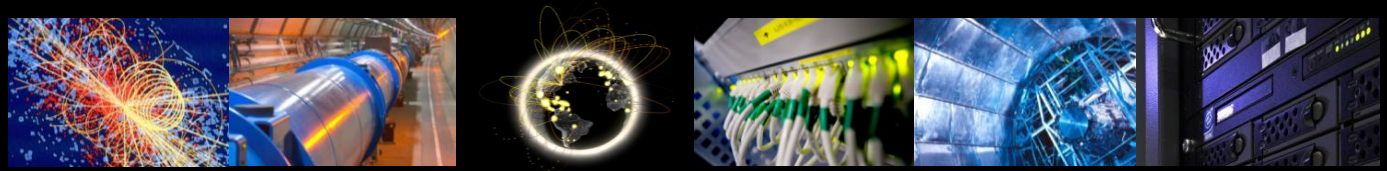

WLCG Worldwide LHC Computing Grid

Open Science Grid

# Outline

- Overview/review of past work
- Current status and tools
- New applications and analytics
- Next steps (CRIC, site diagrams, site-net-mon URLs, etc)

# Reminder: LHCONE Map & Toolkits (Complex Nets!!)

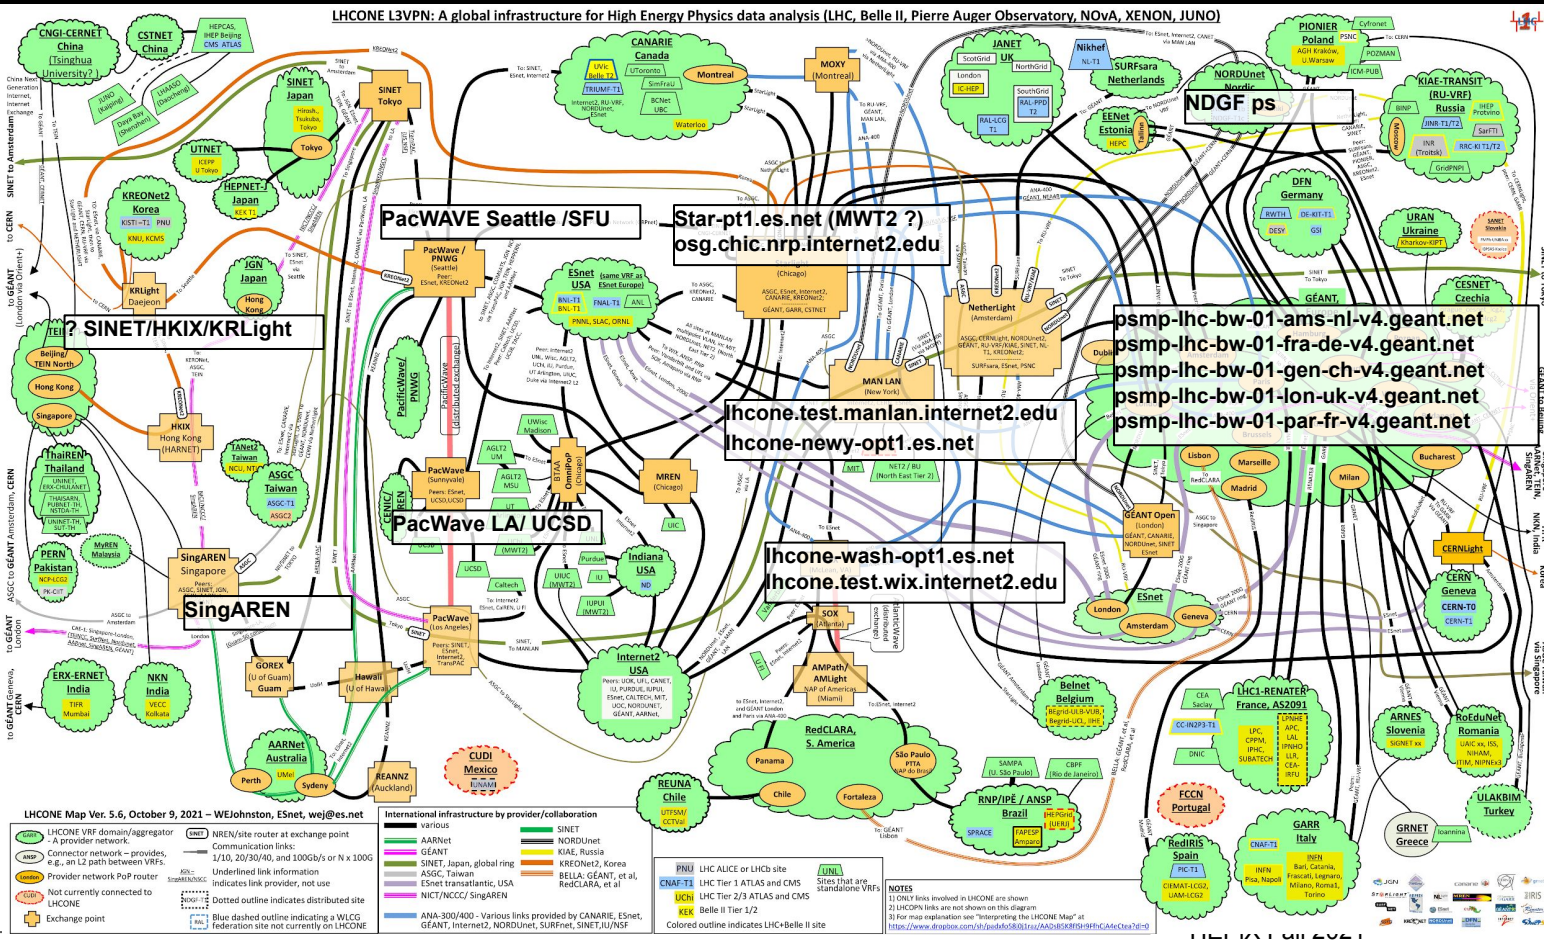

# WLCG Network Throughput Support Unit

We continue to maintain a support channel where sites and experiments can report potential network performance incidents:

- Relevant sites, (N)RENs are notified and perfSONAR infrastructure is used to narrow down the problem to particular link(s) and segment. Also tracking past incidents.
- Provides feedback to WLCG operations and LHCOPN/LHCONE community

**Most common issues**: MTU*, MTU+Load Balancing, routing (mainly remote sites), site equipment/design, firewall, workloads causing high network usage

*The LHCOPN/LHCONE working group has made a recommendation regarding MTU, at least for the networks involved in WLCG. Sites are of course free to use whatever MTU they want but they should understand the trade-offs.

# Review of WLCG Network Monitoring Activities

The focus in WLCG Network Monitoring for many years has been the deployment and configuration of our global **perfSONAR** monitoring framework

- The **deployment** and **operations** of **perfSONAR** has been stable
- The network pipeline has been hardened and is operating reliably.

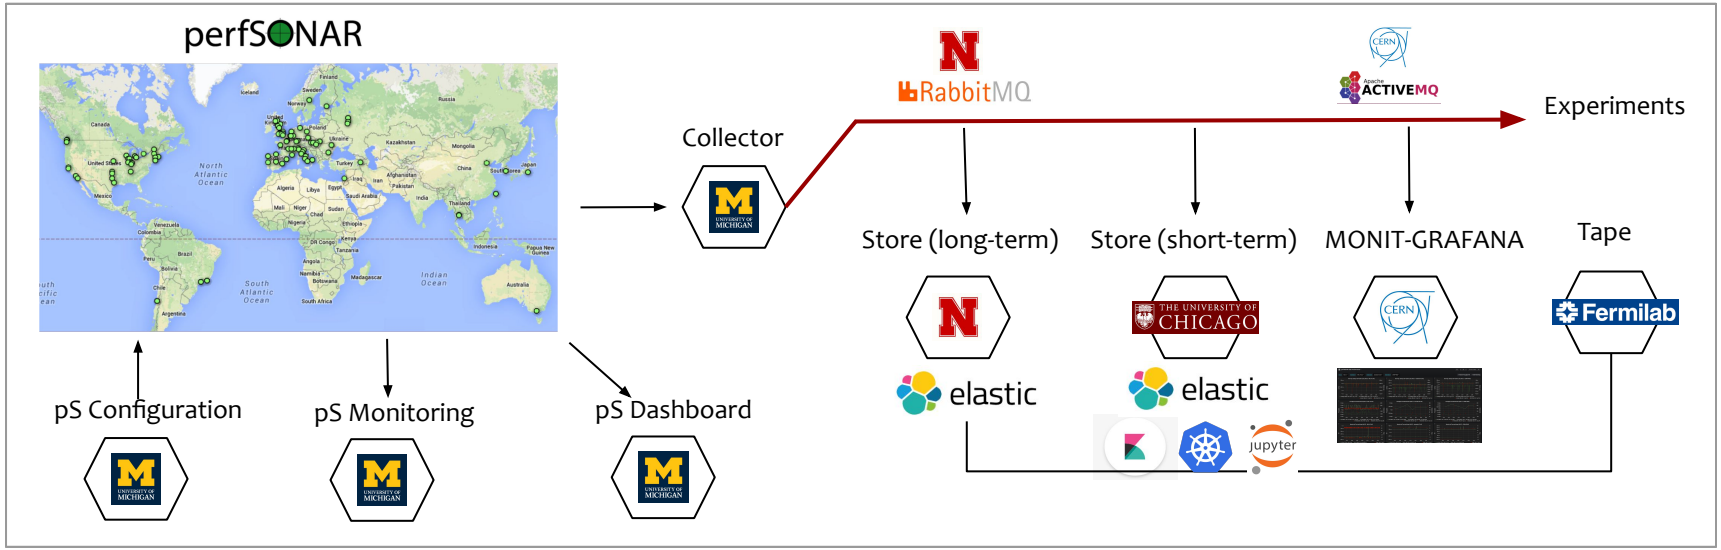

# Network Data Volume via SAND (OSG/WLCG) Pipeline

**The table below shows the amount of data transferred across our pipeline in 3 years of operations**

| SAND Pipeline | Data Type | Total Tests | Tests/day | Storage Size |
|---|---|---|---|---|
| **perfSONAR** | Latency | 6.91B | 7.95M | 3.1TB |
| | Packet Loss | 7.00B | 8.08M | 2.4TB |
| | Retransmits | 14.7M | 18.8k | 6.3GB |
| | Throughput | 15.6M | 19.2k | 7.0GB |
| | Network Path | 1.28B | 2.14M | 1.5TB |
| **ESnet** | Traffic | 1.1B | 44.7M | 1.74TB |
| | Interfaces | 3.2M | 11.8k | 530MB |
| **HTCondor** | Job Transfers | 734M | 446k | 610GB |
| **Total** | | **17.1B** | **64.4M** | **9.4TB** |

# Current Status and Tools

# Existing Tools and Applications for Network Data

With all this data we need tools that organize and make sense of it all

- We have Kibana dashboards looking at
  - Bandwidth
  - Traceroute
  - Packetloss / Latency
  - Infrastructure
- With the completion of the **SAND** project, we have a few prototype tools that help us analyze and utilize our net data
  - We have a new perfSONAR focused dashboard: **ps-dash**
  - We have added a self-subscribe tool for network alarms call **AAAS**
  - ***Will show these later***
- To organize access to all the various resources we have developed we created https://toolkitinfo.opensciencegrid.org (Try it; give us feedback!)

# Access and Visualization

In the next few slides we show some examples of how the data is accessed and visualized.

For **exploring the data**, perhaps the easiest means is by using our Kibana interface: https://atlas-kibana.mwt2.org/s/networking/goto/3c4e90deab15fd901289035a9227de62?auth_provider_hint=anonymous1

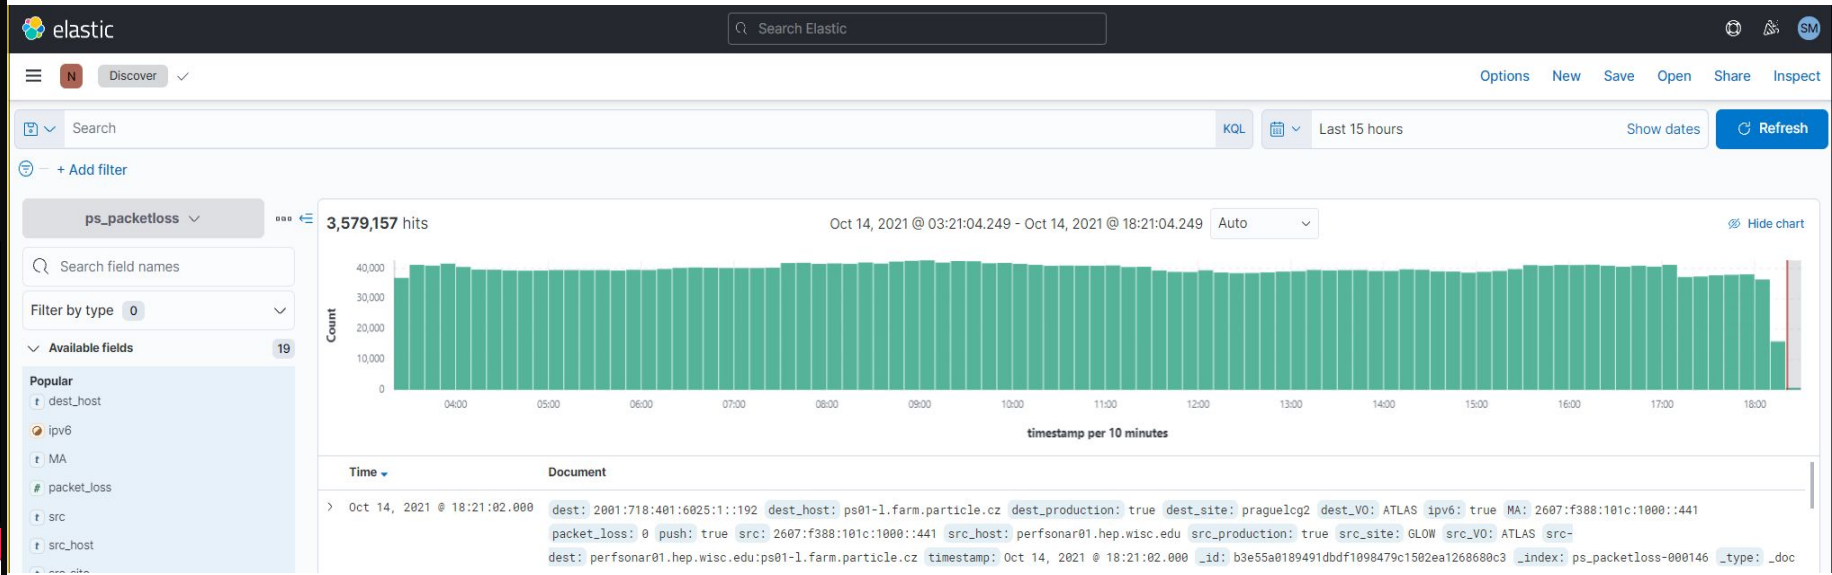

# Example Kibana Dashboard for Traceroute Data

For specific data we have various Kibana dashboards. Below is one for **traceroute data** we process through our infrastructure. Similar dashboards exist for packet-loss, throughput, latency and even one on our perfSONAR infrastructure.

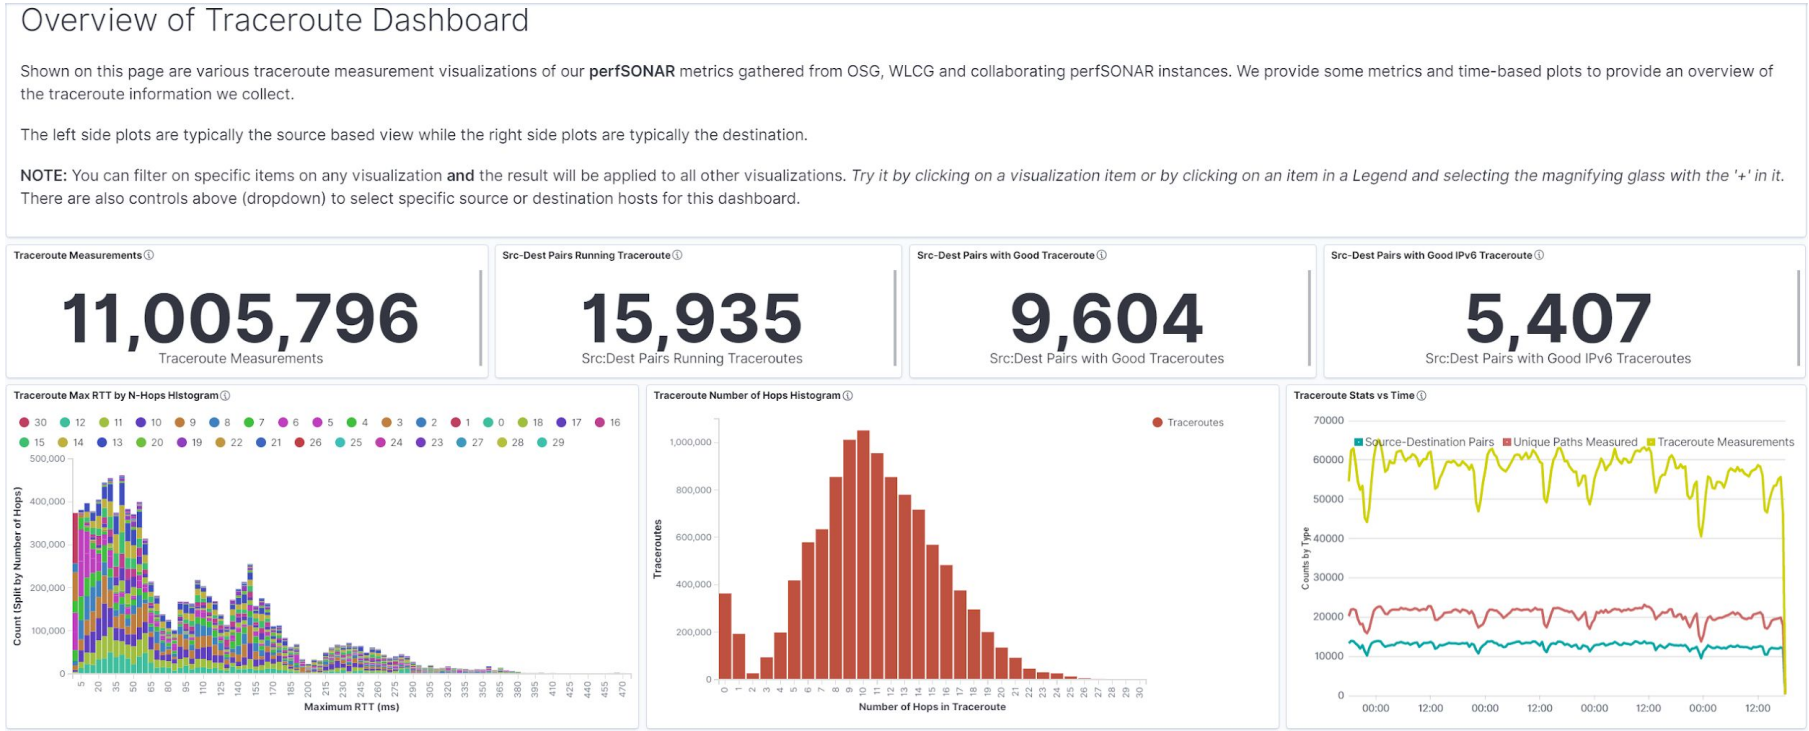

# Example perfSONAR Infrastructure Dashboard

240 Active perfSONAR instances - **207 production endpoints** - T1/T2 coverage

- Continuously testing over 5000 links - testing coordinated and managed from central place
- Dedicated latency and bandwidth nodes at each site - **Open platform (testing and data)**

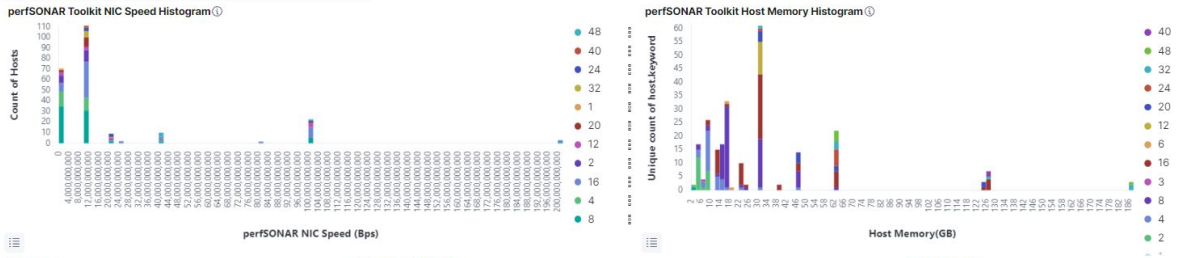


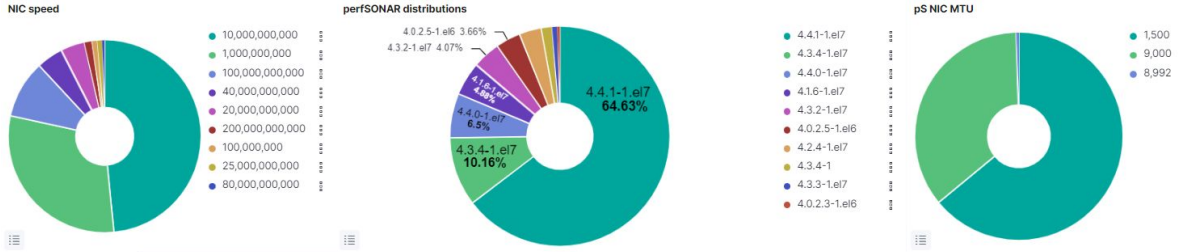


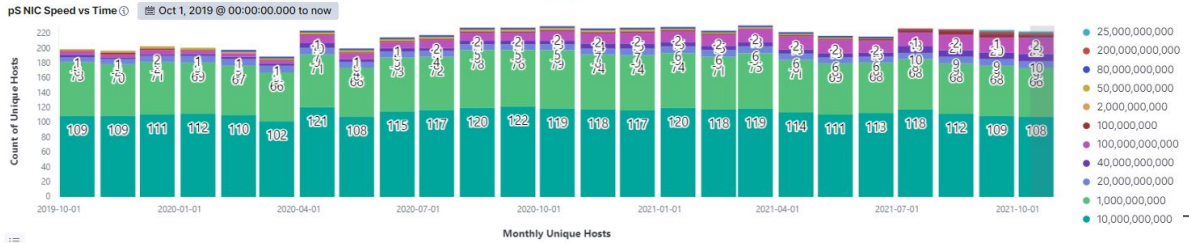


Our global toolkit deployment has a **range** of systems in terms of age and capability.

Dashboard in ELK

**Sites** should remember to not only upgrade perfSONAR software but also the underlying **hardware,** as nodes become too old or are unable to test at the site storage speed.

# Grafana - perfSONAR dashboard at CERN

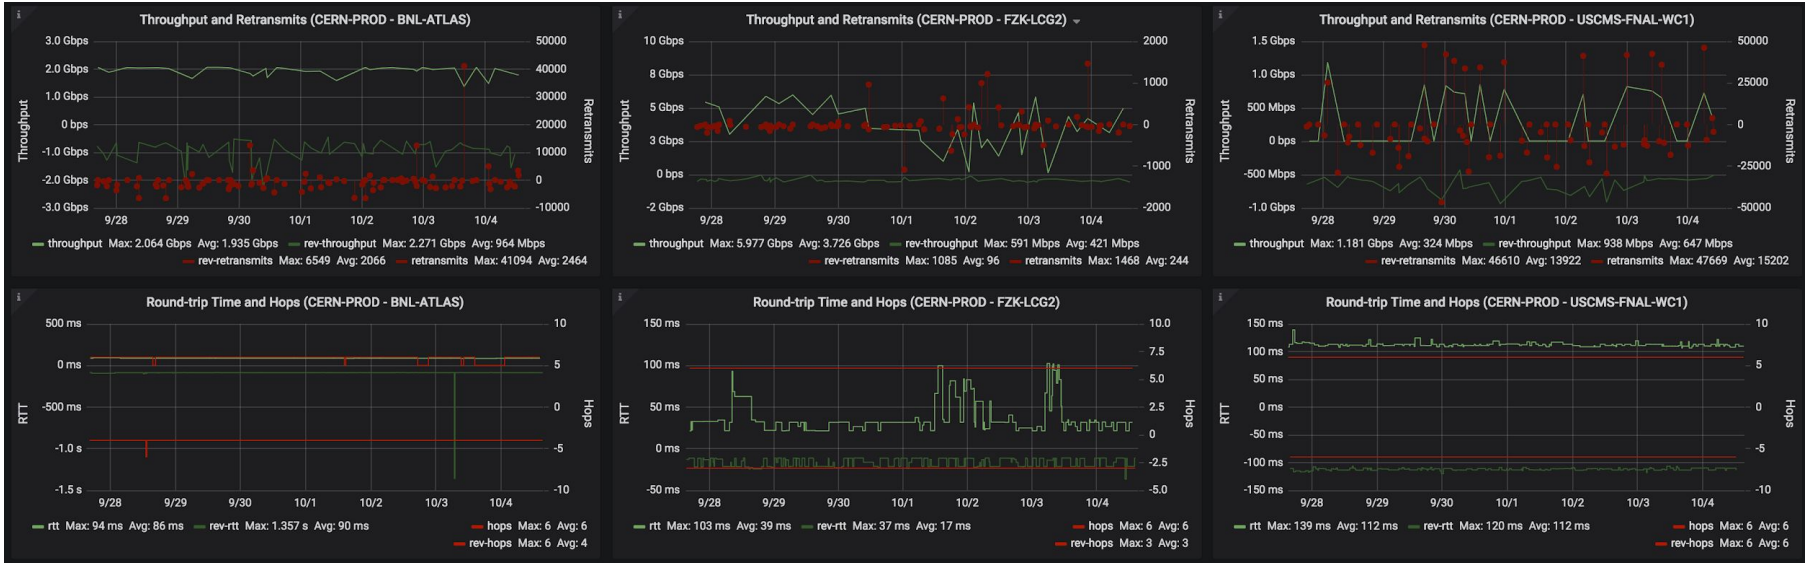


- **CERN** hosts [Grafana dashboards] for our network metrics
- Now includes all WLCG sites that run perfSONAR
  - Additional work needed to better filter production nodes
- Added additional row that tracks RTT and number of hops as reported by traceroute/tracepath
- We have some work to reorganize a few of the dashboards showing summaries

# Update on WLCG Network Data Challenge

**WLCG data challenge was Oct 04-08**

Goal was 240 Gbps from T0 to T1's and from T1's to T2's involving primarily ATLAS and CMS

The network was NOT the bottleneck in general!

**Still doing the analysis/post-mortem**

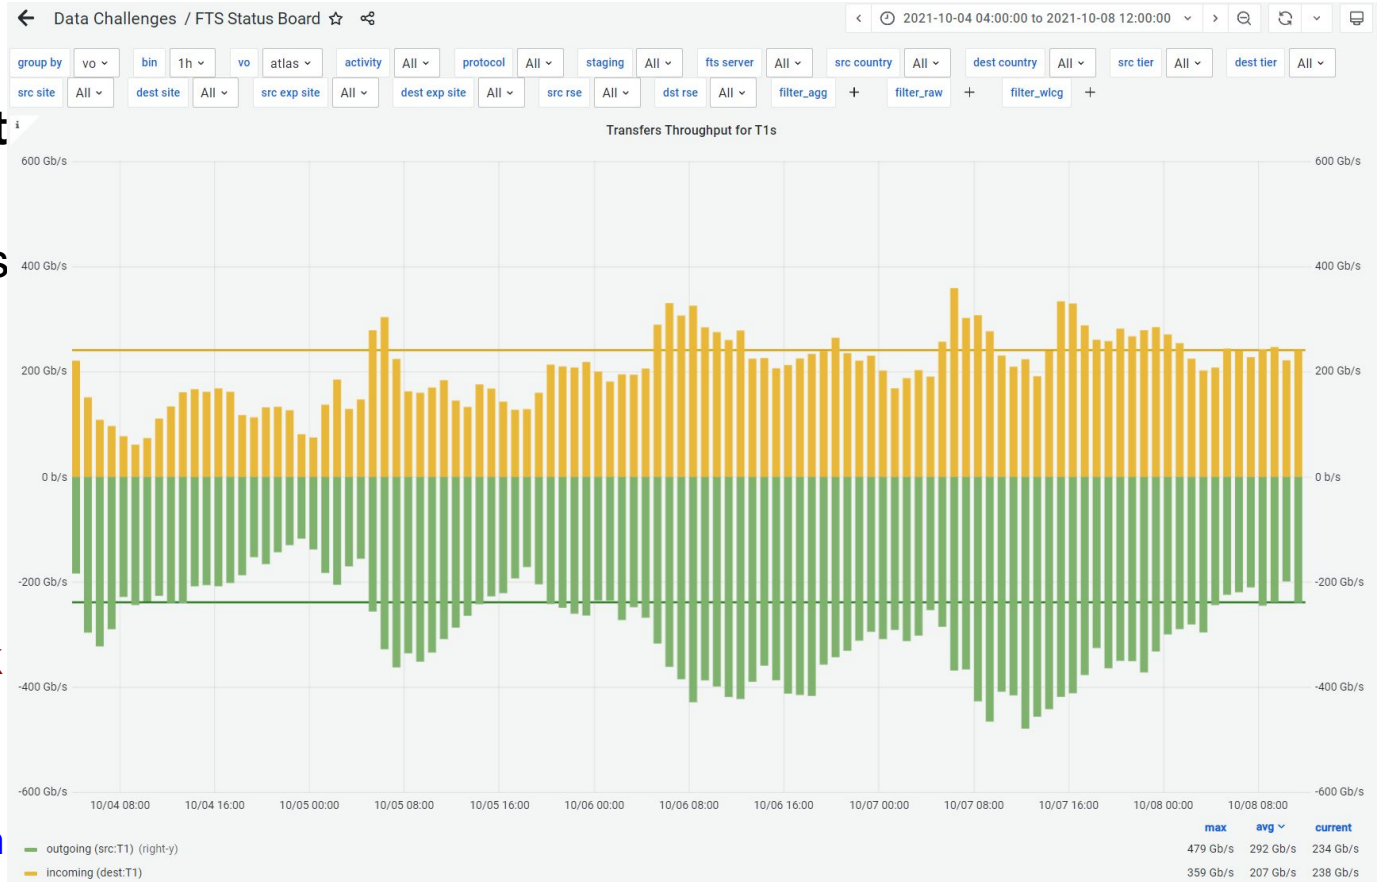

# ESnet Monitoring for WLCG Data Challenge

**ESnet** created a very nice [monitoring dashboard] which is now part of our infrastructure

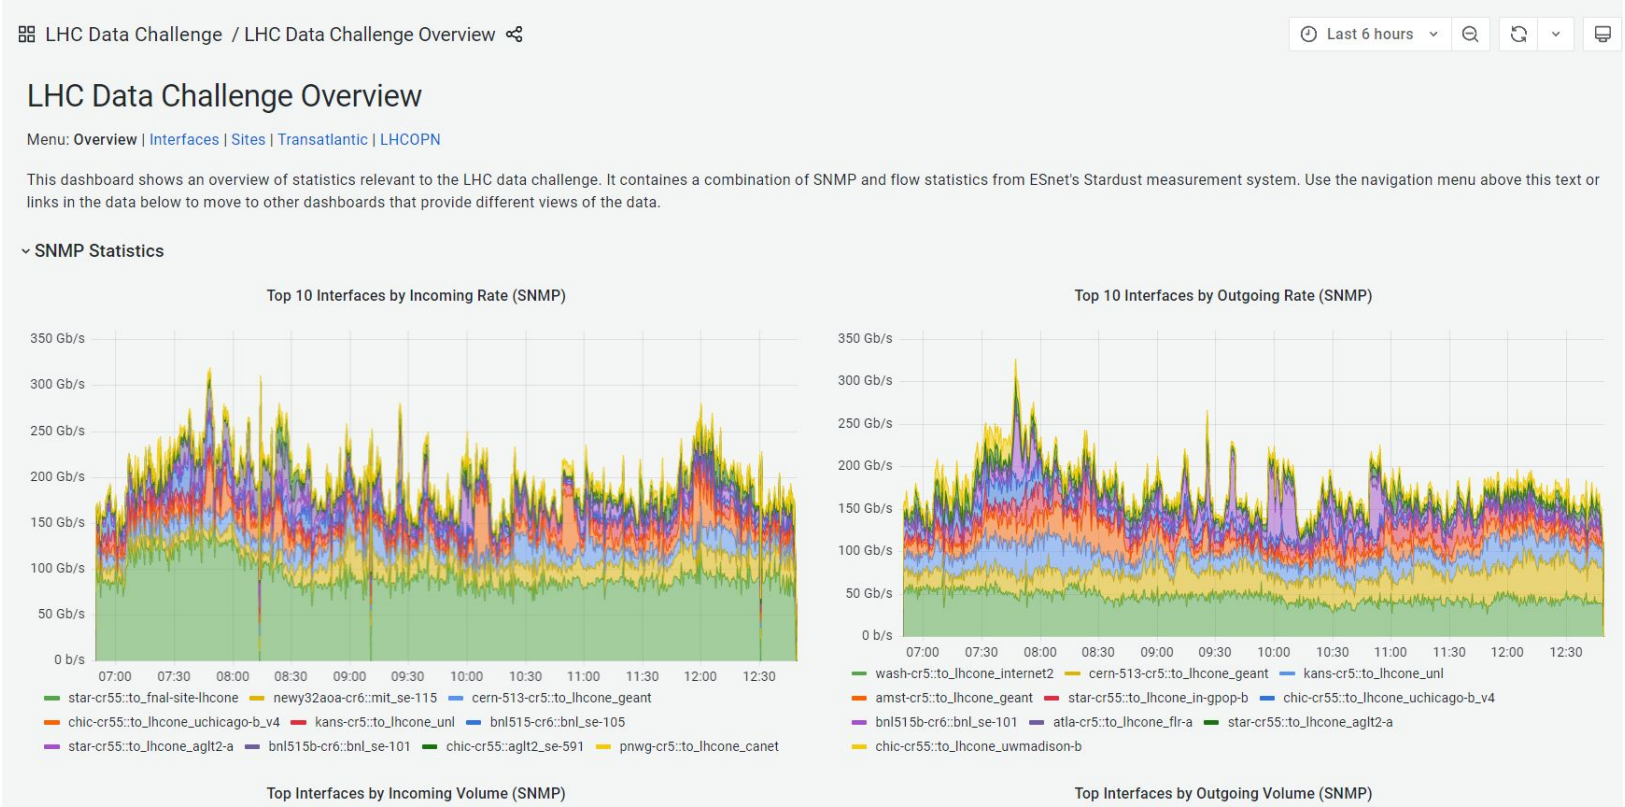

# New Applications and Analytics

# Motivation and History

As we have constructed our network monitoring infrastructure we have gotten significant feedback from our users.

One of the most frequent requests is a mechanism to alert users when network issues involving their infrastructure happen.

- This is more challenging than you might think
- We want to alert users to issues **they** can address
- Need to ensure it is really a **network problem** and have it **localized**

One of the most important components delivered by **SAND** is a new Alerting and Alarming Service which tries to address this need.

# New Services

With the **AAAS** service at https://aaas.atlas-ml.org/ we have a working prototype for network-related alerting

- Users **login** with their **institutional credentials** and can choose which types of events to subscribe to and with what filter and how often to be notified.
- While we still have a significant amount of work to do to tune sensitivity vs noise and add new types of alarms, we feel it has reached the “beta” stage (feedback would be welcome!).

We also have a new dashboard application called **pS-Dash** that includes data analysis and ranking and serves as the “alert” development platform for AAAS

These applications were developed in collaboration with students who have been working on this from **SAND** and **IRIS-HEP/OSG-LHC**

# Tommy Shearer and Toolkitinfo, AAAS and pS-Dash

**IRIS-HEP Fellow** Working on OSG Network Data and user interfaces
Fellowship dates: May-August 2020
Home Institution: University of Michigan, Ann Arbor

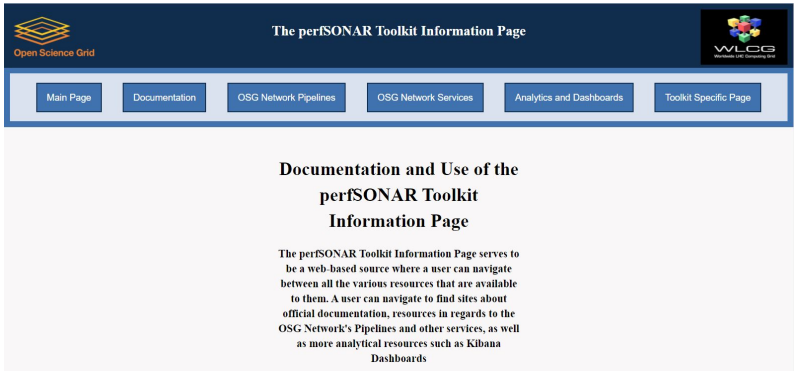


## Improving the User Interface to OSG-LHC Network Metrics

The OSG-LHC Network Monitoring collects and manages many different types of metrics related to network performance. My work has created a system of user interfaces that are organized by data type, and filterable by the parameters of interest. Additionally, my work on the OSG Toolkit Information Page (http://toolkitinfo.opensciencegrid.org) has allowed users to easily maneuver these interfaces dynamically with queried toolkit hosts, organization by name or distance, and dynamical linking to user interfaces. Furthermore, I also am working on the user interface for a new application that will provide customizable contacting for Alarms and Alerts (**AAAS**) within the OSG-LHC Network. I work with many different people and users to ensure that my work will be the most optimized for the largest amount of possible use-cases.

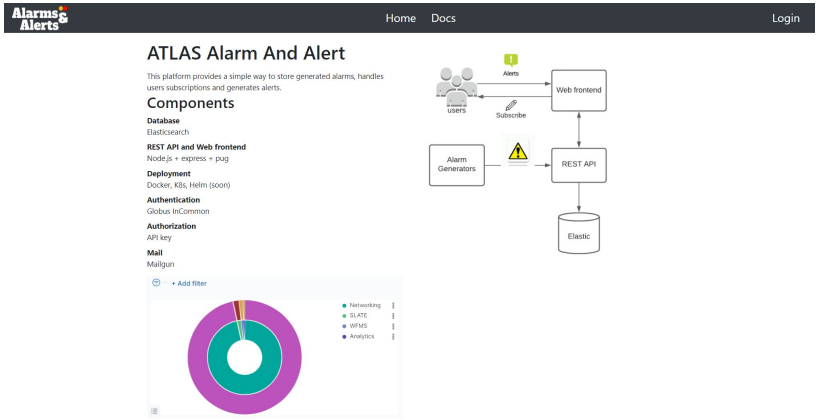

# Petya Vasileva and pS-Dash

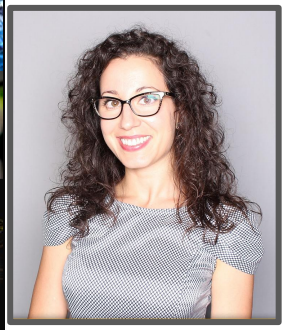

I am a PhD student @ University of Plovdiv, Bulgaria

Currently, my project is focused on implementing a method to find problems on the network by using the data from perfSONAR

I have created a Plotly Dash (https://ps-dash.uc.ssl-hep.org/) application which

- ➔ Ranks Sites based on their measures
- ➔ Blacklists nodes and highlights the most problematic ones
- ➔ Provides plots for various selections
- ➔ Prototypes alerting functions for **AAAS**

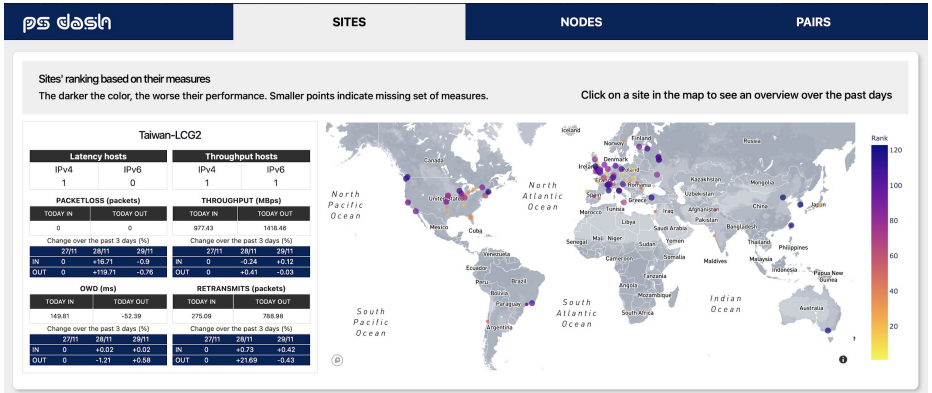


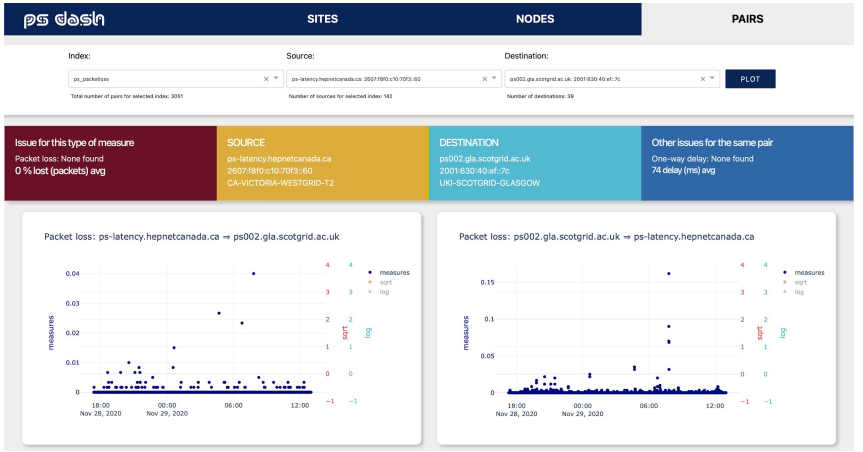

# ATLAS Alarms & Alerts Service

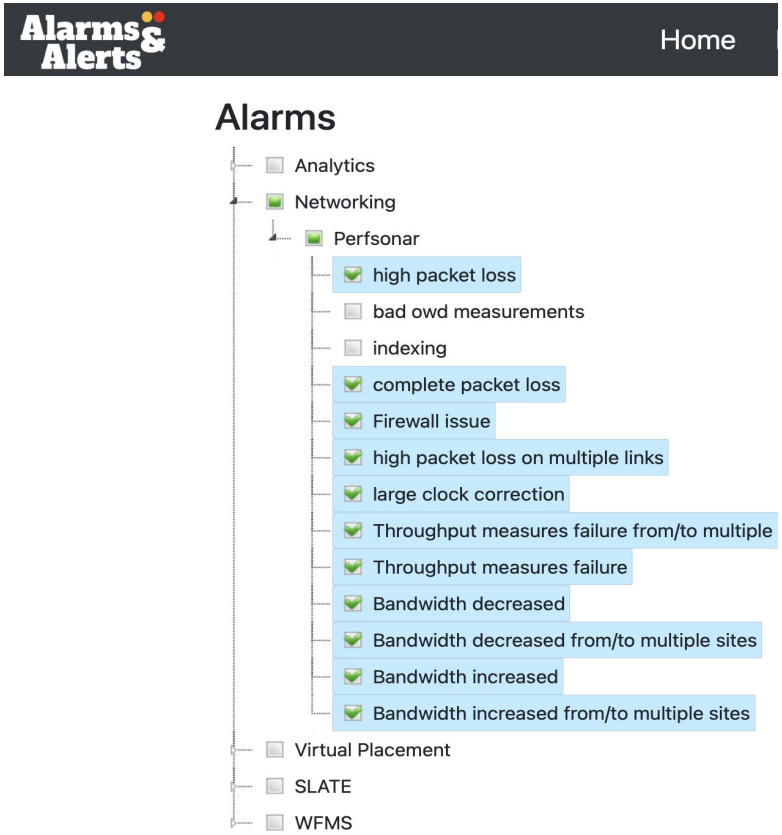


https://aaas.atlas-ml.org/

**Purpose**: provides user-subscribable alerting for specific types of network issues found by analyzing perfSONAR data

Currently available:

- Main packet loss issues
- Main throughput issues

**Future plans:**

- Add traceroute alarms:
  - Destination never reached
  - Path changes too often
  - Node causes issues with multiple sites

# pS (perfSONAR) Dash

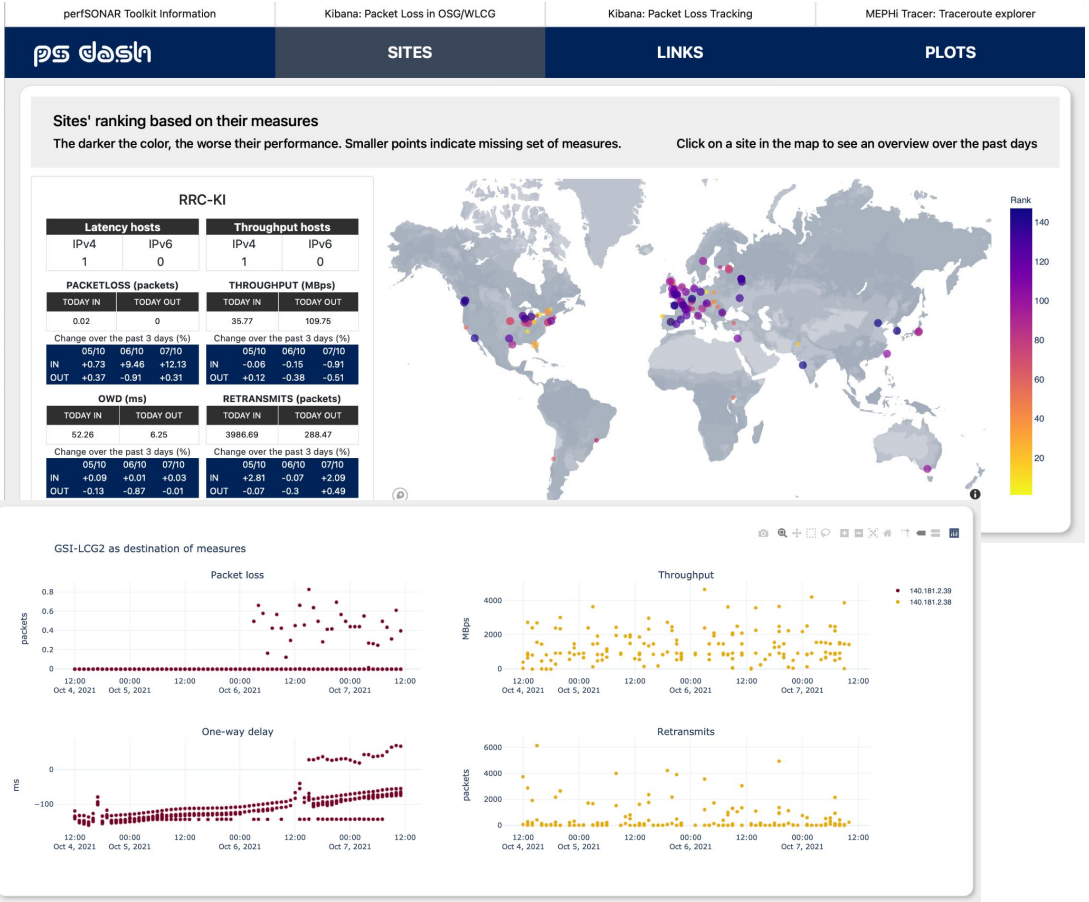


https://ps-dash.uc.ssl-hep.org/

**Purpose**: provides a user dashboard to explore analyzed and summarized perfSONAR data.

Currently:

- Allows users to monitor their sites
- Provides tools for detecting basic problems

**Future plans:**

- Add today's Alarms
- Add traceroute data & plots
- Refine ranks
- Deduce causes for found issues

# Next Steps

# Network Information In CRIC (1/2)

Edoardo Martelli asked for updates to CRIC to register networking information. Plan is for this to become the **source** of LHCONE/LHCOPN details for networking. Extensions are planned to track perfSONAR as well. See dev example https://atlas-cric-dev.cern.ch/core/netsite/list/

Export | Filter | Reload | Columns 9/10

NetworkRoute list

Show 100 entries

| Site | NetworkRoute | primary NetSite | spare NetSite | ASN | Subnets | LHCONE limit | LHCONE limit | collaborations |
|---|---|---|---|---|---|---|---|---|
| AGLT2 | AGLT2_PTP | AGLT2_UM | AGLT2_MSU | 229 | 192.41.238.0/28 | 0 | -1 | US-ATLAS, WLCG, US-CMS |
| AGLT2 | AGLT2_Shared_Alloc | AGLT2_UM | AGLT2_MSU | 237 | 2001:48a8:68f7:4000::/50 | 0 | -1 | US-ATLAS, WLCG |
| AGLT2 | MSU | AGLT2_MSU | | 229 | 192.41.236.0/23 | 0 | -1 | US-ATLAS |
| AGLT2 | MSUv6 | AGLT2_MSU | | 237 | 2001:48a8:68f7:8000::/50 | 0 | -1 | US-ATLAS |
| AGLT2 | UM | AGLT2_UM | | 229 | 192.41.230.0/23 | 0 | -1 | US-ATLAS, WLCG |
| AGLT2 | UM-2 | AGLT2_UM | | 237 | 2001:48a8:68f7::/50 | 0 | -1 | US-ATLAS, WLCG |
| MWT2 | IU | IU | | 10680 | 149.165.224.0/24, 149.165.225.0/24, 149.165.236.64/29 | 100 | -1 | US-ATLAS |
| MWT2 | UC | UC | | 160 | 192.170.224.0/19, 2605:9a00:10::/48 | 0 | -1 | DUNE |
| MWT2 | UIUC-route | UIUC | | 38 | 72.36.81.64/26, 72.36.96.0/24, 2620:0:e01:4800::/56 | 0 | -1 | US-ATLAS, DUNE |
| Site | NetworkRoute | primary NetSite | spare NetSite | ASN | Subnets | LHCONE limit | LHCONE limit | collaborations |

Showing 1 to 9 of 9 entries

Previous 1 Next

# Network Information in CRIC (2/2)

Having network information in **CRIC** can help us better organize and identify our sites and related networks.

**CRIC** may eventually host our perfSONAR registrations, currently in OSG topology or GOCDB.

**CRIC** is also intended to provide a location where site's can register a site-specific URL that provides network monitoring, diagrams and details.

It is very important we have a **single source of truth** for networking details!

Once this is production, we will have a how-to guide about filling in data.

# Machine Learning and a Network Database

Given the scope and duration of the SAND project, we were limited in what we could undertake. There are **two specific areas** that we feel would be valuable next steps:

1. **Machine Learning** (**ML**): Identifying network issues in "noisy" data, using packet loss measurements to understand achievable bandwidth and looking for complex interactions in network traffic are all areas that might benefit from ML
   a. **Requirement:** cleaned, annotated data to make progress (significant effort)
   b. **Benefit:** Anomaly detection could be a very fruitful area using throughput changes, routing changes, packet loss and the associated time and space (topology) information.
2. Constructing and maintaining a **network "Link" database**: The full set of R&E network paths use by our community is tractable (~50K links). It would be a powerful resource to have each link recorded with **owner** of each end, associated **IPs**, **AS** numbers, **contact information** AND **dynamic information** about **min**, **max** and **average** traffic seen on the link.
   a. **Requirement**: Would require continuous real time updates as metrics arrive
   b. **Benefit**: Could quickly identify problematic links and track their history and performance.

# Additional Future Work Areas

- More work on network topology - cleaning, analyzing and visualizing.
- Add **site specific monitoring** (from CRIC URLs)
  - Data will need to be “harvested” and integrated into our pipeline
- Improving the latency and reliability of our pipeline by transitioning from a primarily “pull” data model (collectors) to a secure “push” model.
- Incorporating flow / packet marking accounting/monitoring (Marian’s talk)
- Exploring **packet pacing** and **traffic shaping** to better match resources
- Broadening the engagement with the Global research community
  - Access to R&E monitoring and integration with our tools and datastores
- Others? Please suggest additional ideas :)

# Summary and Conclusion

The **collaboration of OSG, WLCG and various research projects** have created an extensive, reliable infrastructure to monitor our networks, gather metrics and provide analytics and visualization.

While making progress in evolving our tools and infrastructure, we need to continue to monitor and maintain what we have built.

There are a number of challenges remaining, including the difficult meta-challenge of making effective insights available for everyone to easily use.

**Questions or Comments?**

# Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented, especially Ilija Vukotic, Derek Weitzel, John Thiltges, Petya Vasileva and Tommy Shearer

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- OSG: NSF MPS-1148698
- IRIS-HEP: NSF OAC-1836650
- SAND: NSF CC* INTEGRATION 1827116

# Useful URLs

- OSG/WLCG Networking Documentation
  - https://opensciencegrid.github.io/networking/
- perfSONAR Infrastructure Dashboard
  - https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7
- perfSONAR Dashboard and Monitoring
  - http://maddash.opensciencegrid.org/maddash-webui
  - https://psetf.opensciencegrid.org/etf/check_mk
- perfSONAR Central Configuration
  - https://psconfig.opensciencegrid.org/
- Toolkit information page
  - https://toolkitinfo.opensciencegrid.org/
- Grafana dashboards
  - http://monit-grafana-open.cern.ch/
- ATLAS Alerting and Alarming Service: https://aaas.atlas-ml.org/
- The pS Dash application: https://ps-dash.uc.ssl-hep.org/
- ESnet WLCG DC Dashboard: https://public.stardust.es.net/d/IkFCB5Hnk/lhc-data-challenge-overview?orgId=1

# Backup Slides Follow

# LHCONE - 23rd of March 2021

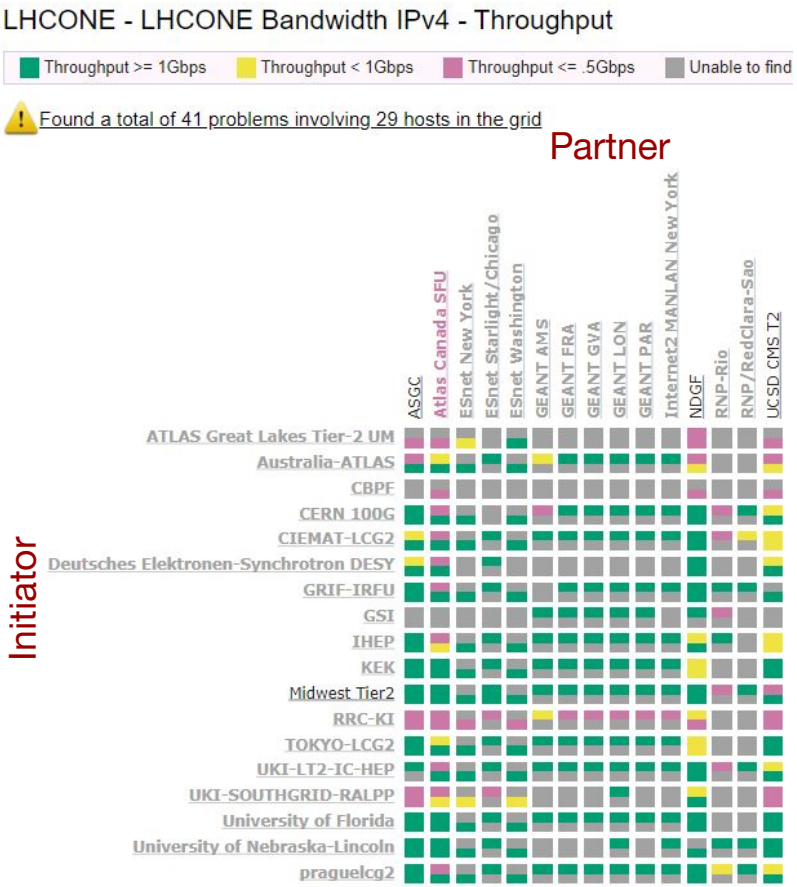


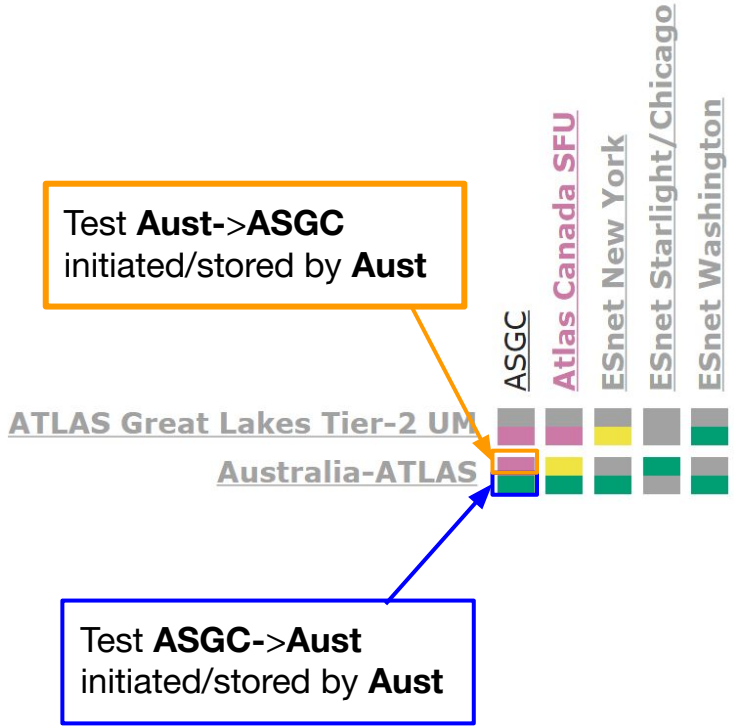

# LHCONE 10th Oct 2021

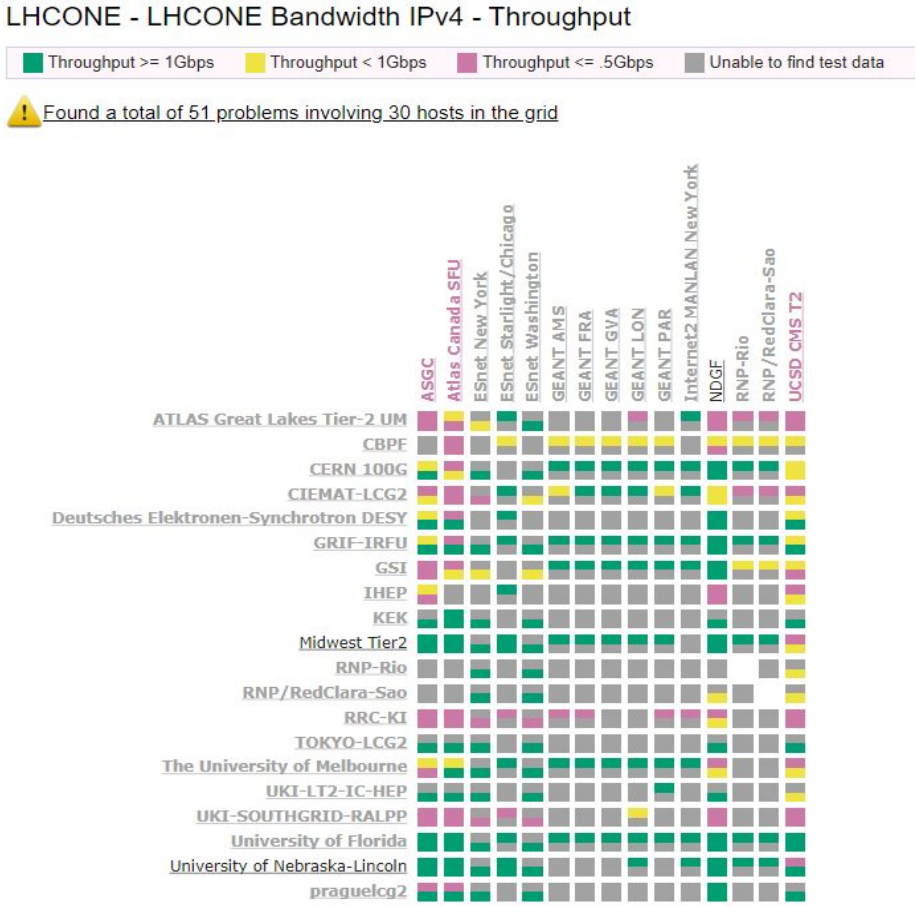


Lot's of missing test results in this mesh, even more than in March!!

We have identified some issues with the 4.4.1 perfSONAR toolkits and hitting thread limits after some time operating normally.

The perfSONAR developers have been made aware of this and are working to identify and fix the problem