LHCOPN/LHCONE Monitoring Update

Shawn McKee, Marian Babik

Fall LHCONE/LHCOPN (#47) Meeting, Oct 11, 2021

on behalf of WLCG Network Throughput WG

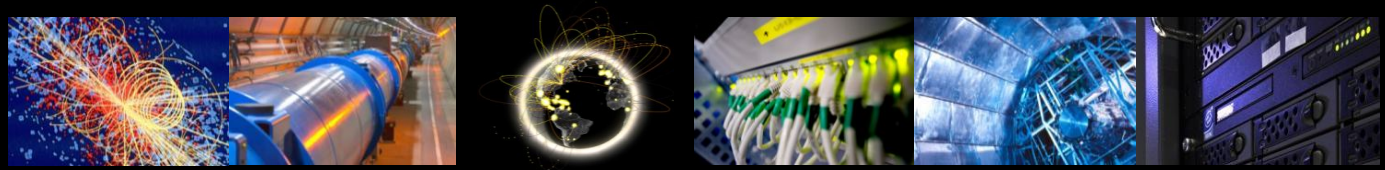

WLCG Worldwide LHC Computing Grid

Open Science Grid

Outline

- OSG/WLCG Network Monitoring and WLCG Network Throughput WG
- perfSONAR community updates
- LHCOPN/LHCONE perfSONAR infrastructure status
- New Tools and the WLCG Data Challenge
- Summary

OSG/WLCG networking projects

There have been 4 coupled projects around the core OSG Net Area

1. **SAND** (NSF) project for analytics (ended)
2. **HEPiX** NFV WG (finished work)
3. **perfSONAR** project
4. **WLCG Network** Throughput WG

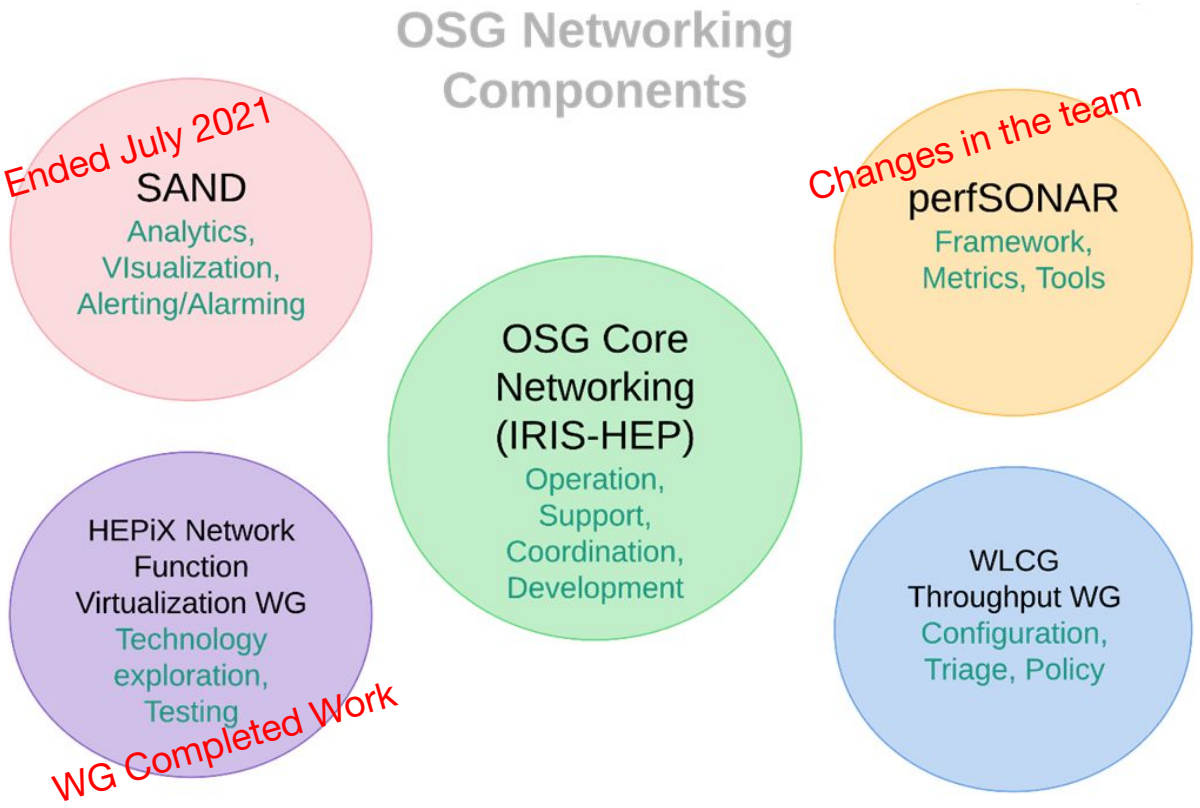

Reminder: Network Measurement Platform Overview

- Collects, stores, configures and transports all network metrics
  - Distributed deployment - operated in collaboration
- All perfSONAR metrics are available via API, live stream or directly on the analytical platforms
  - Complementary network metrics such as ESNet, LHCOPN traffic also via same channels

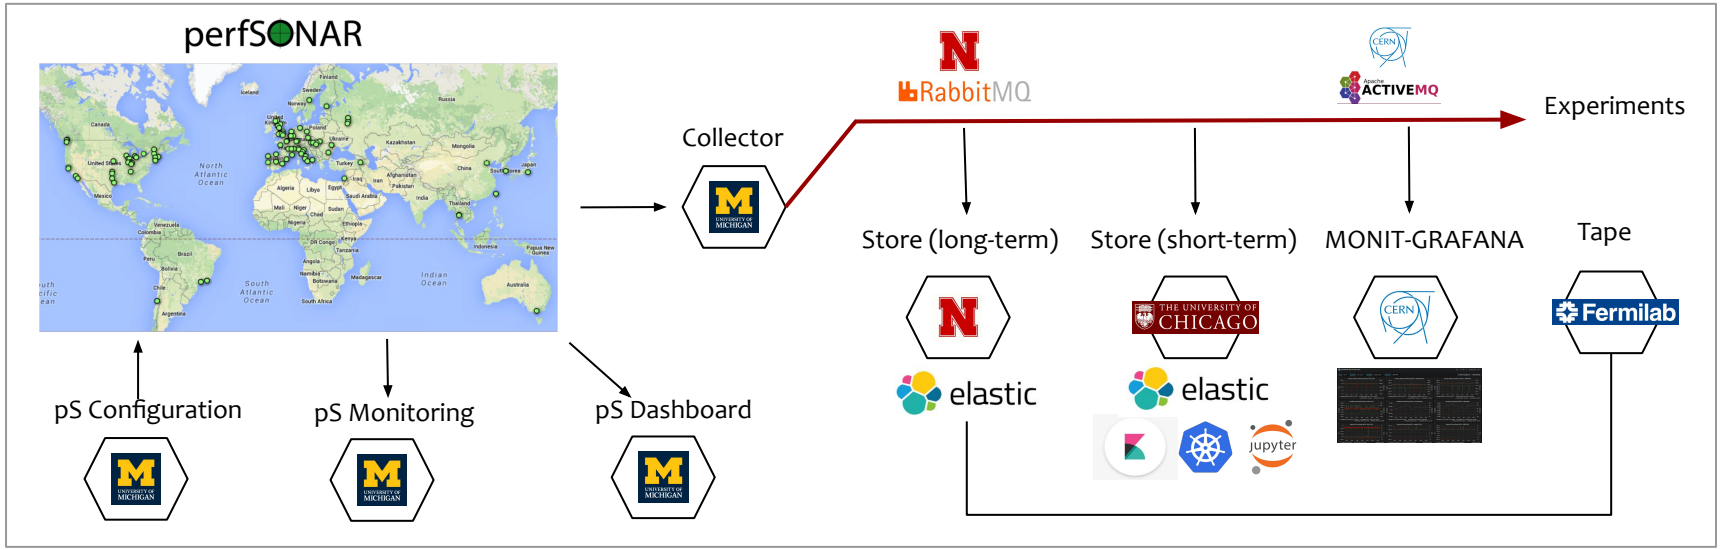

perfSONAR News

- The 4.4.0 release on July 8, 2021 added a number of new features
  - Support for loopback tests, graphs archive whitelist and a new RabbitMQ archiver
  - We identified some issues for our WLCG use-cases.
- 4.4.1 bug fix released September 7, 2021
  - Numerous pScheduler fixes to improve reliability and resource use
- We have been working on having some instances write directly to the RabbitMQ bus but this has been a bit challenging.
- Currently seeing issues with 4.4.1 nodes and problems hitting thread limits on busy nodes after running for a long time.
- New hire John Grigutis has taken over PWA, **working on many fixes/updates** that have been identified/requested. Still have issues with Auto-URL

perfSONAR deployment

238 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)**

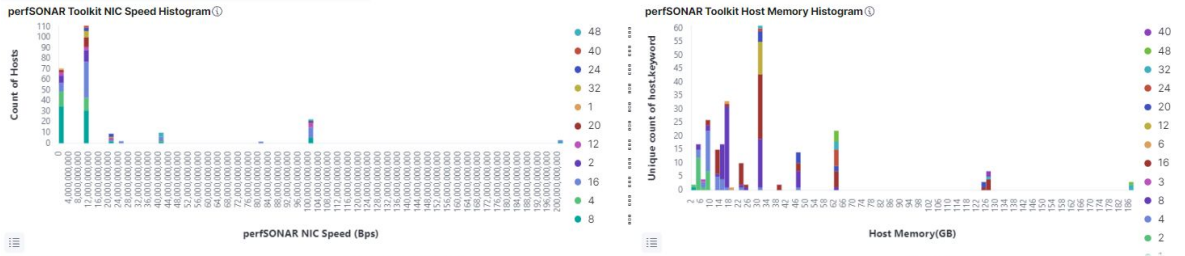

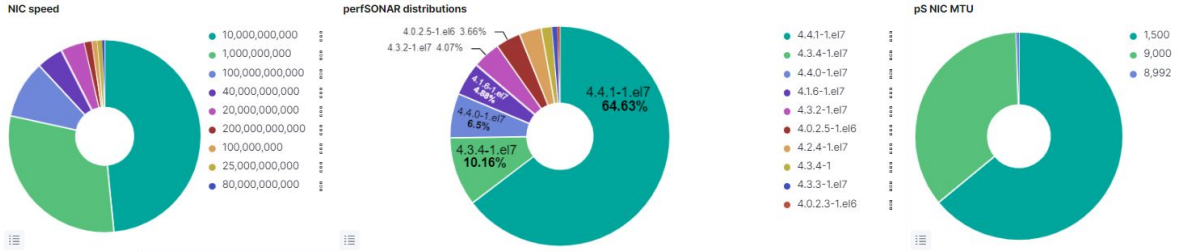

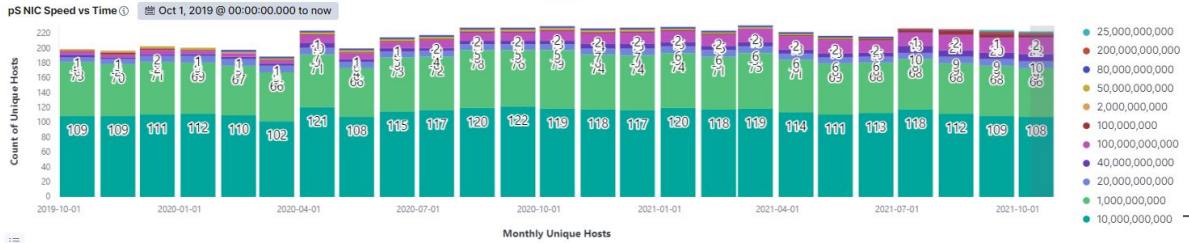

Our global toolkit deployment has a range of systems in terms of age and capability

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.

LHCOPN 23rd March 2021

OPN Mesh Config - OPN IPv6 Bandwidth - Throughput

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps

Found a total of 4 problems involving 3 hosts in the grid

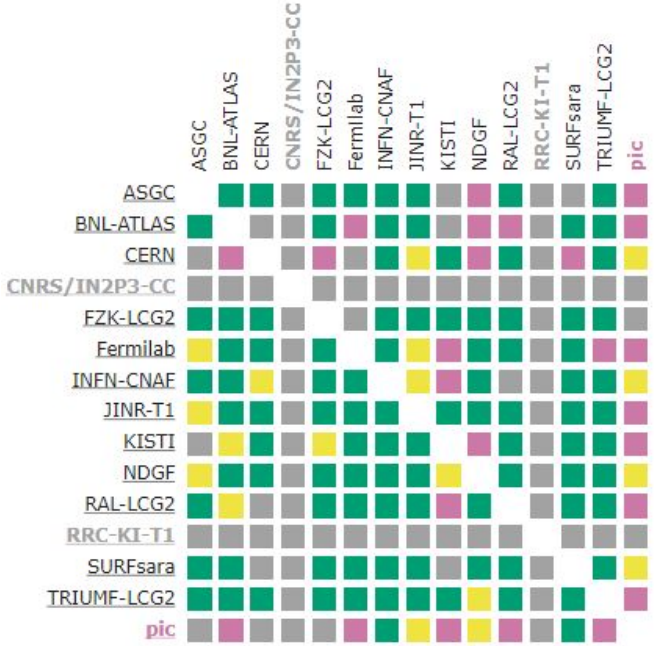

OPN Mesh Config - OPN Latency - Loss

Loss rate is <= 0.001% | Loss rate is > 0.001% | Loss rate

Found a total of 2 problems involving 2 hosts in the grid

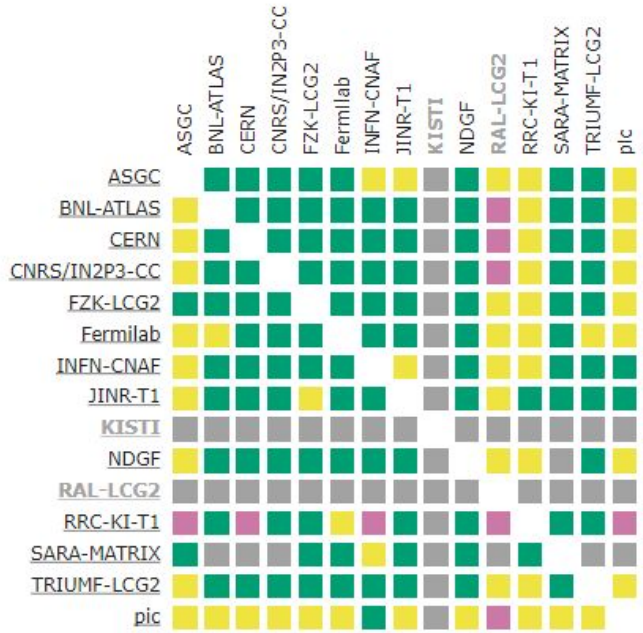

LHCOPN 10th Oct 2021

OPN Mesh Config - OPN IPv6 Bandwidth - Throughput

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps | Unable

Found a total of 8 problems involving 6 hosts in the grid

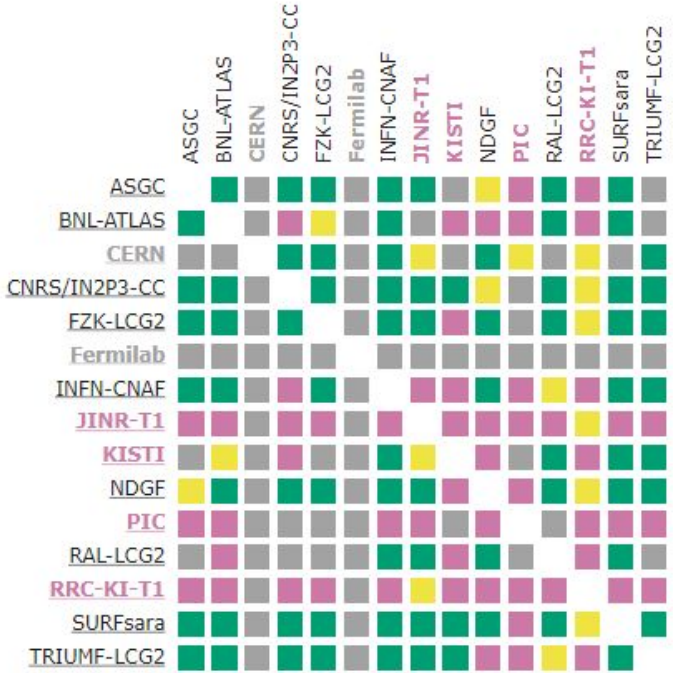

OPN Mesh Config - OPN Latency - Loss

Loss rate is <= 0.001% | Loss rate is > 0.001% | Loss rate is >= 0.1%

Found a total of 5 problems involving 5 hosts in the grid

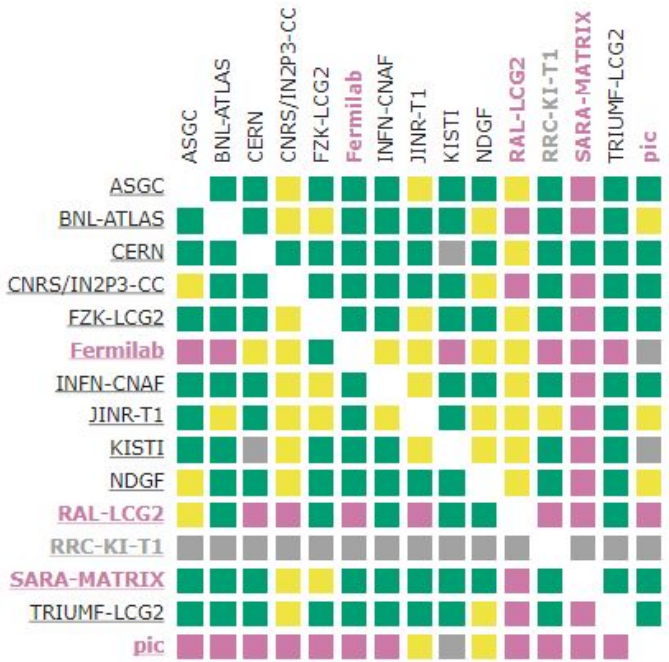

Reminder: LHCONE mesh Added (strategic locations)

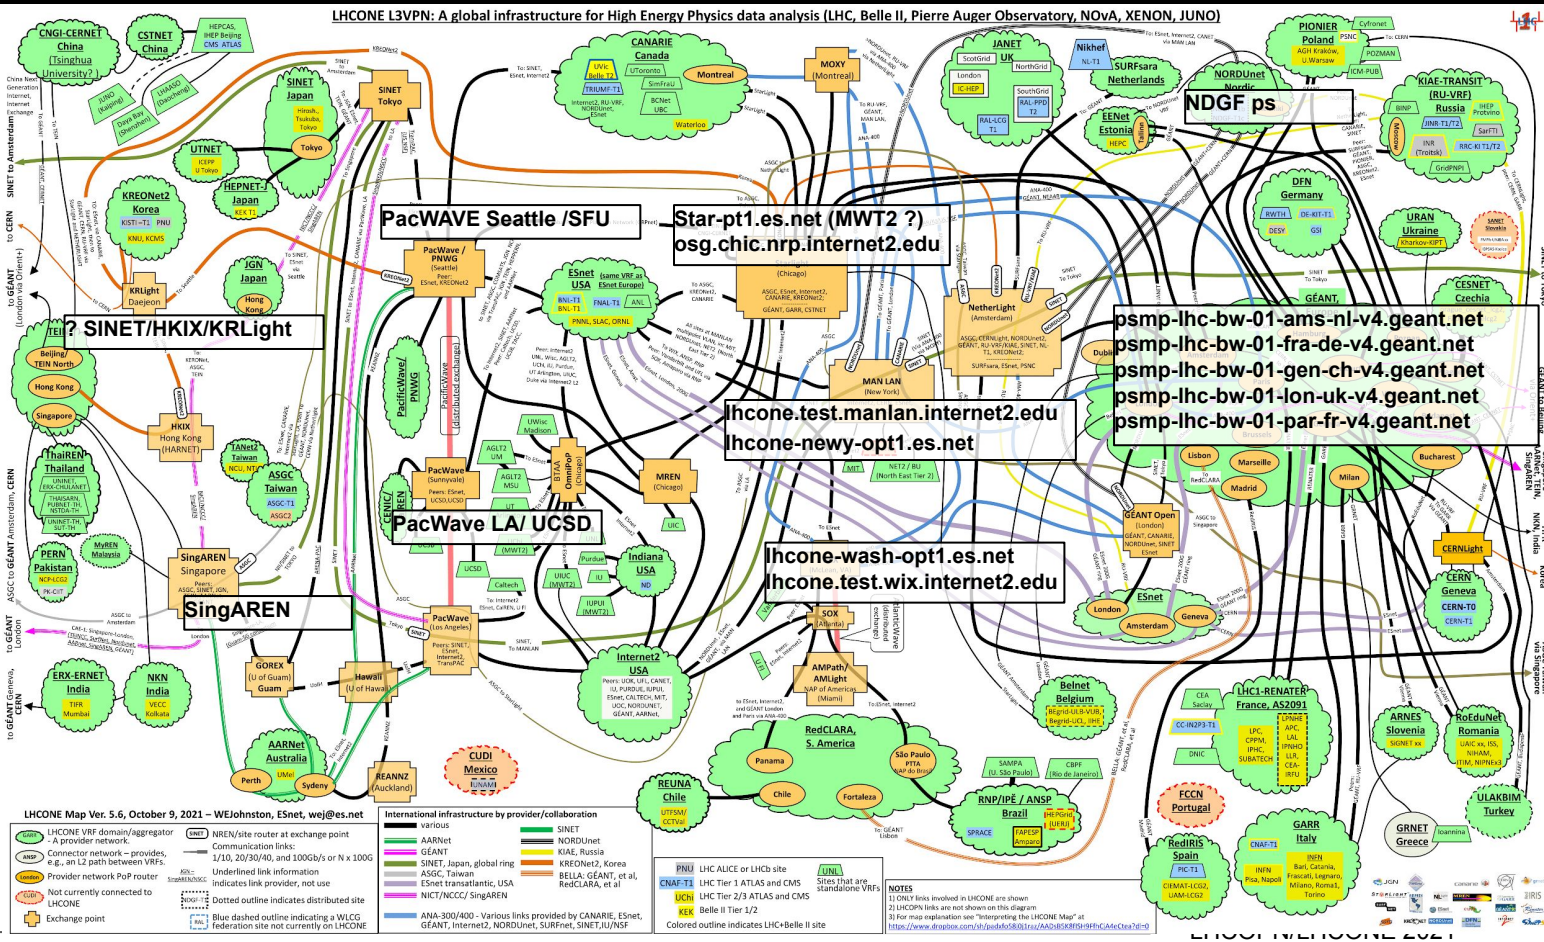

LHCONE - 23rd of March 2021

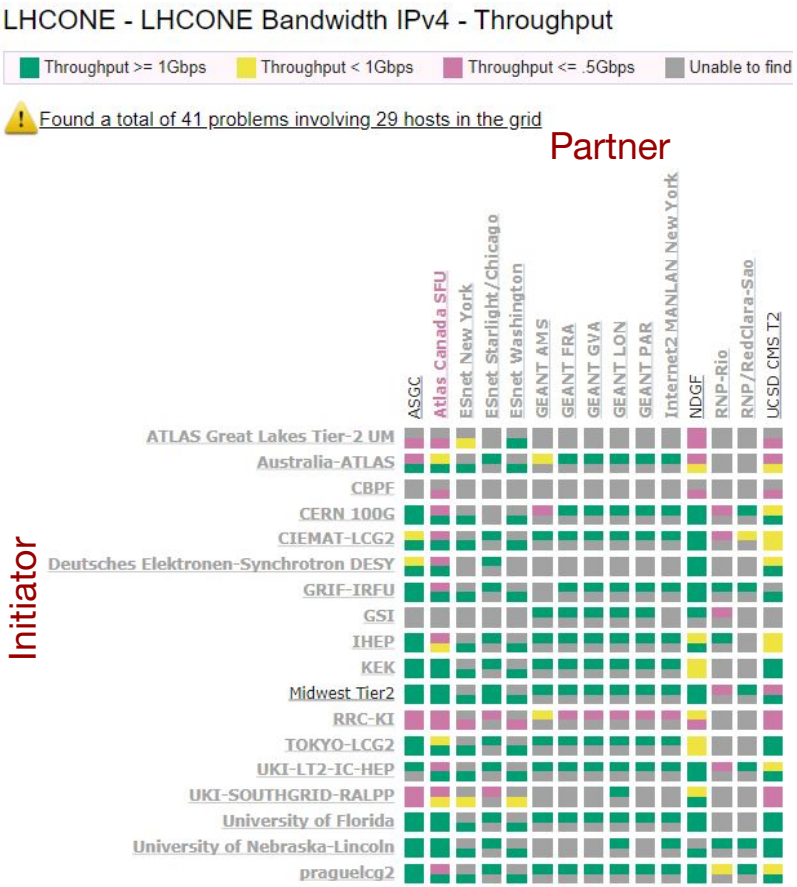

LHCONE 10th Oct 2021

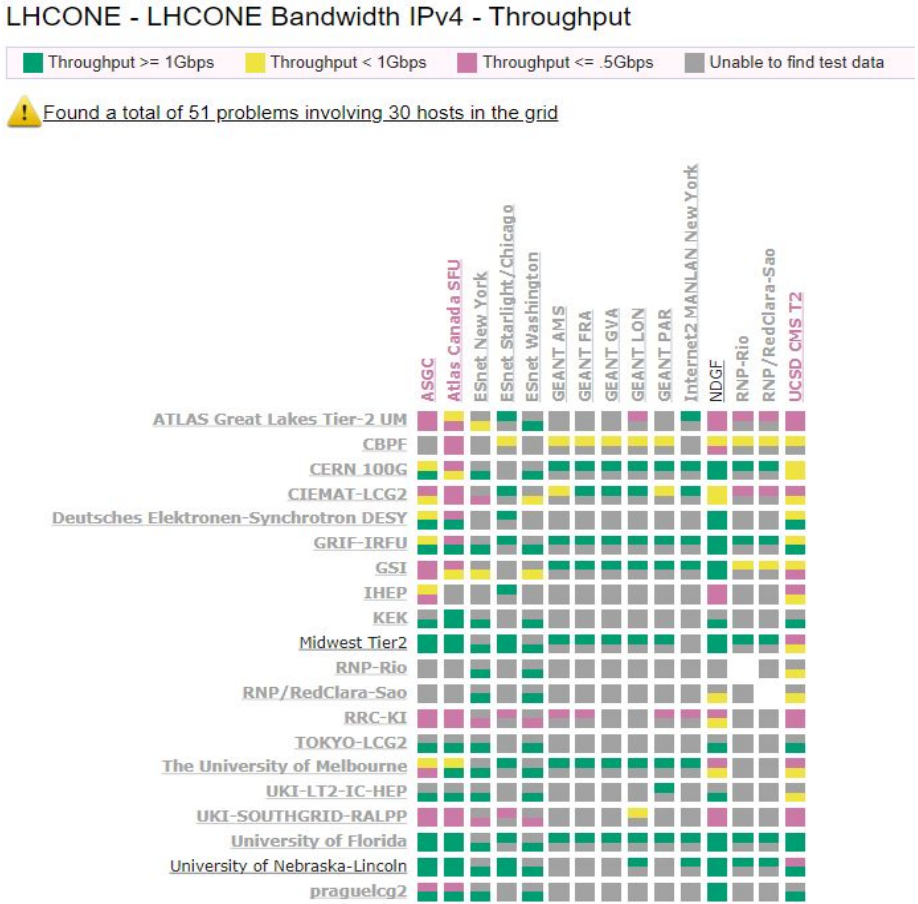

Lot's of missing test results in this mesh, even more than in March!!

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

100Gbps Testing 23rd Mar 2021

- LHCOPN/LHCONE 100Gbps mesh looked good in spring.

WLCG 100G Mesh - WLCG 100G IPv4 Bandwidth - Throughput

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps | Unable to fi

Found a total of 2 problems involving 1 host in the grid

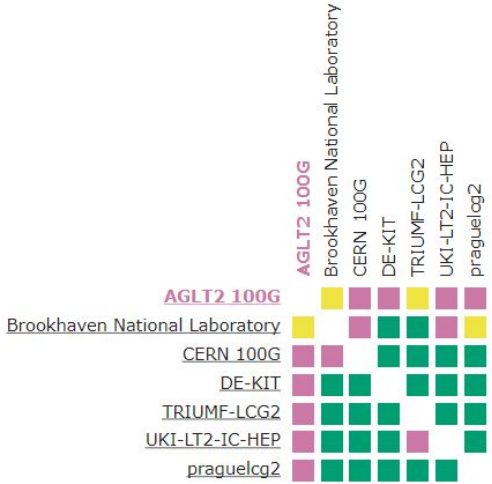

WLCG 100G Mesh - WLCG 100G IPv6 Bandwidth - Throughput

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps | Unable to fi

Found a total of 2 problems involving 1 host in the grid

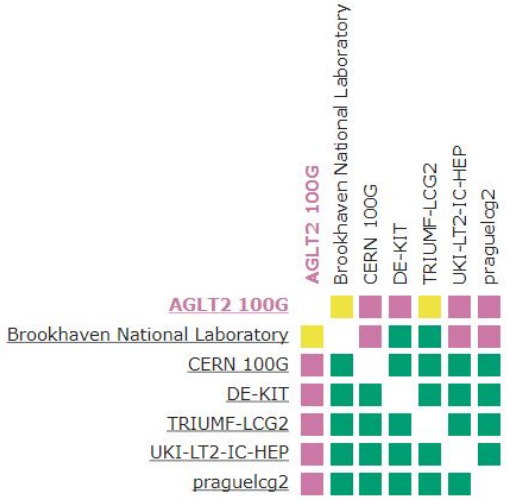

100Gbps Testing 10th Oct 2021

- LHCOPN/LHCONE 100Gbps mesh not looking as good.
- Could be some of the 4.4.1 issues we have been seeing

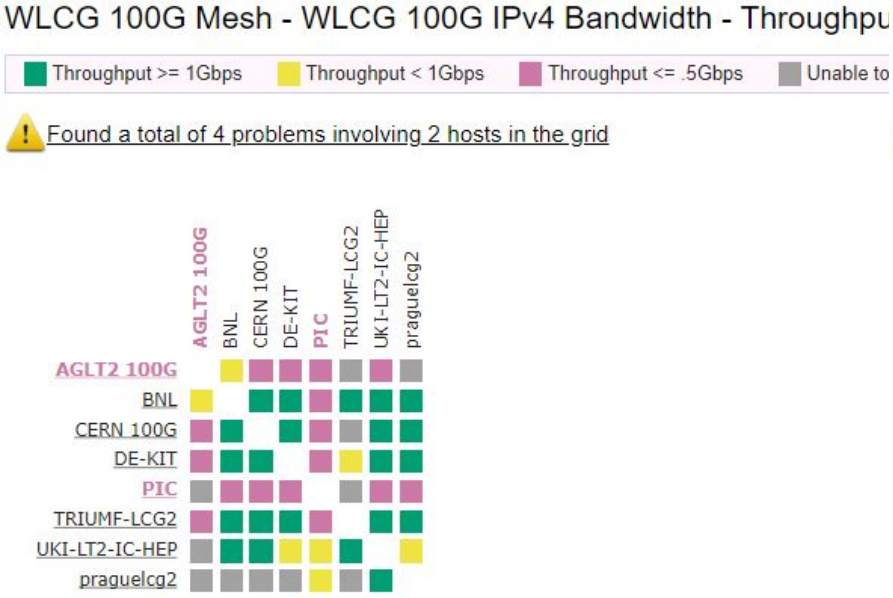

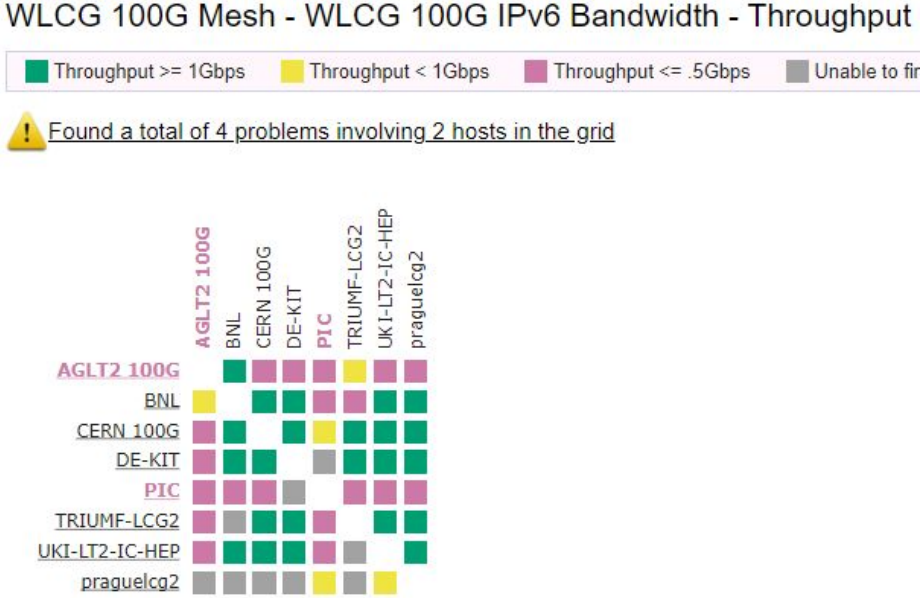

Tools and Applications for Network Data

- We already 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**
  - ***Next two pages have the details on these two apps***
- To organize access to all the various resources we have developed we created https://toolkitinfo.opensciencegrid.org (Try it; give us feedback!)

pS (perfSONAR) Dash

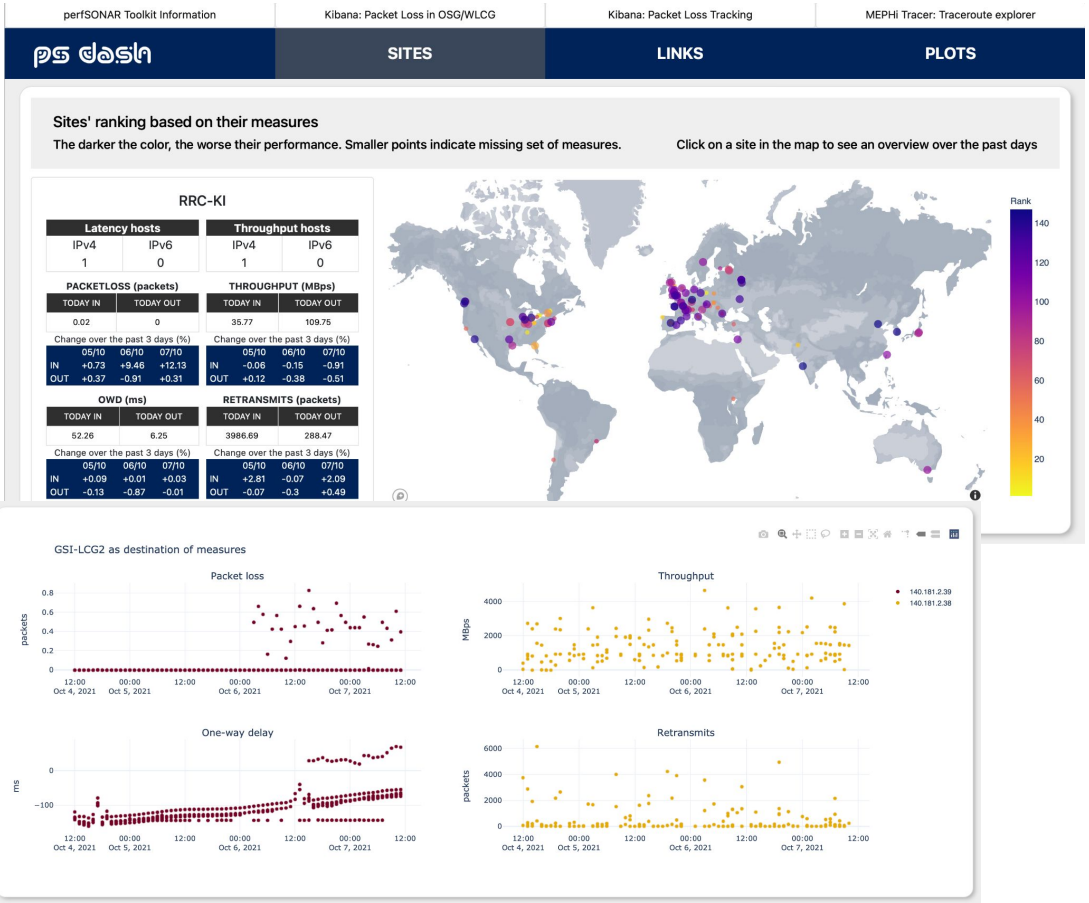

https://ps-dash.uc.ssl-hep.org/

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
- Deduct possible cause for found issues

ATLAS Alarms & Alerts Service

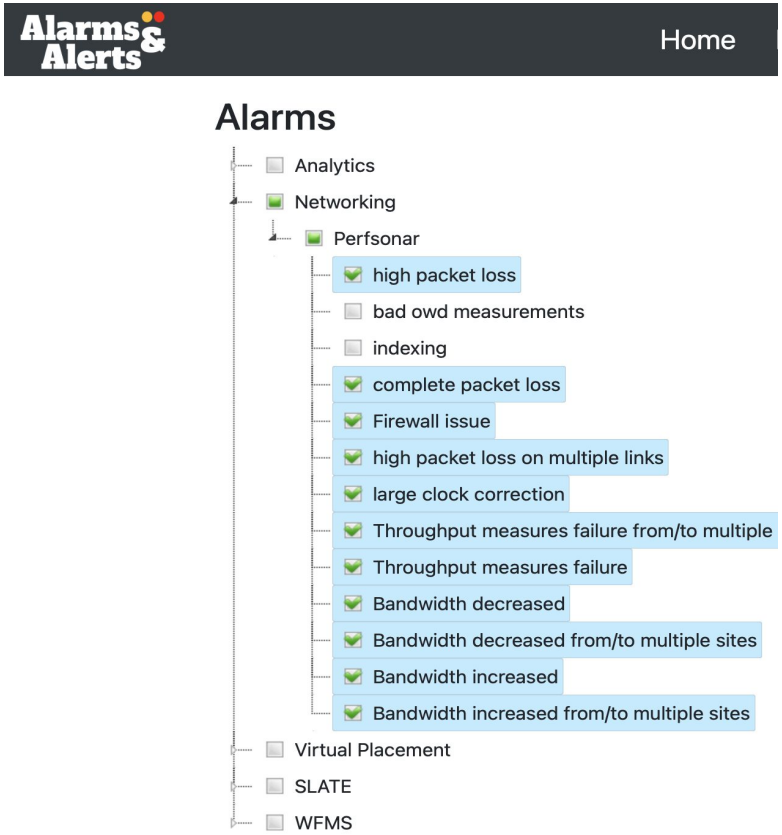

https://aaas.atlas-ml.org/

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

Update on WLCG Network Data Challenge (1 / 2)

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!

This week is the Tape challenge!

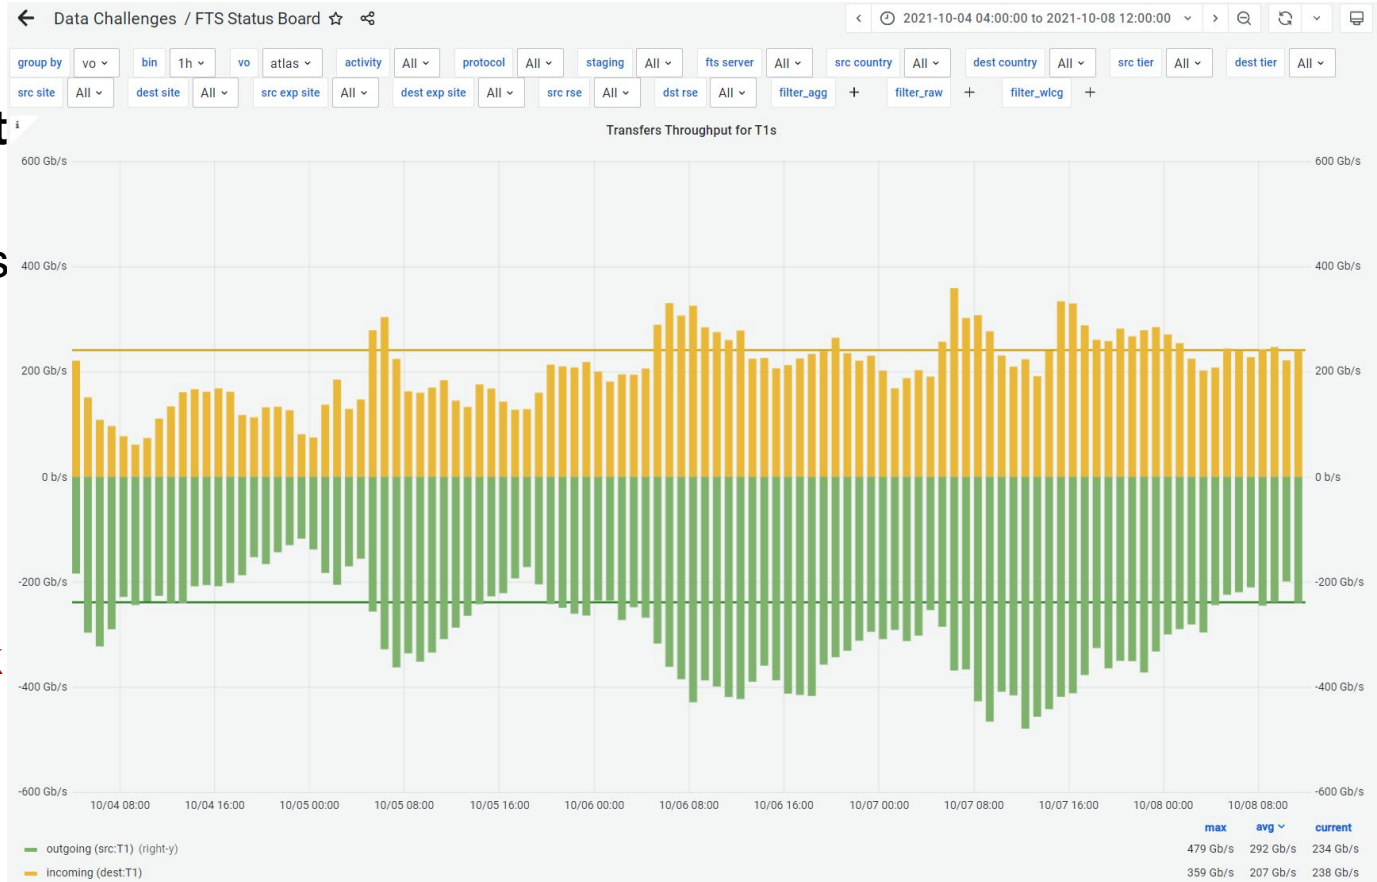

ESnet Monitoring for WLCG Data Challenge

ESnet created a very nice [monitoring dashboard]

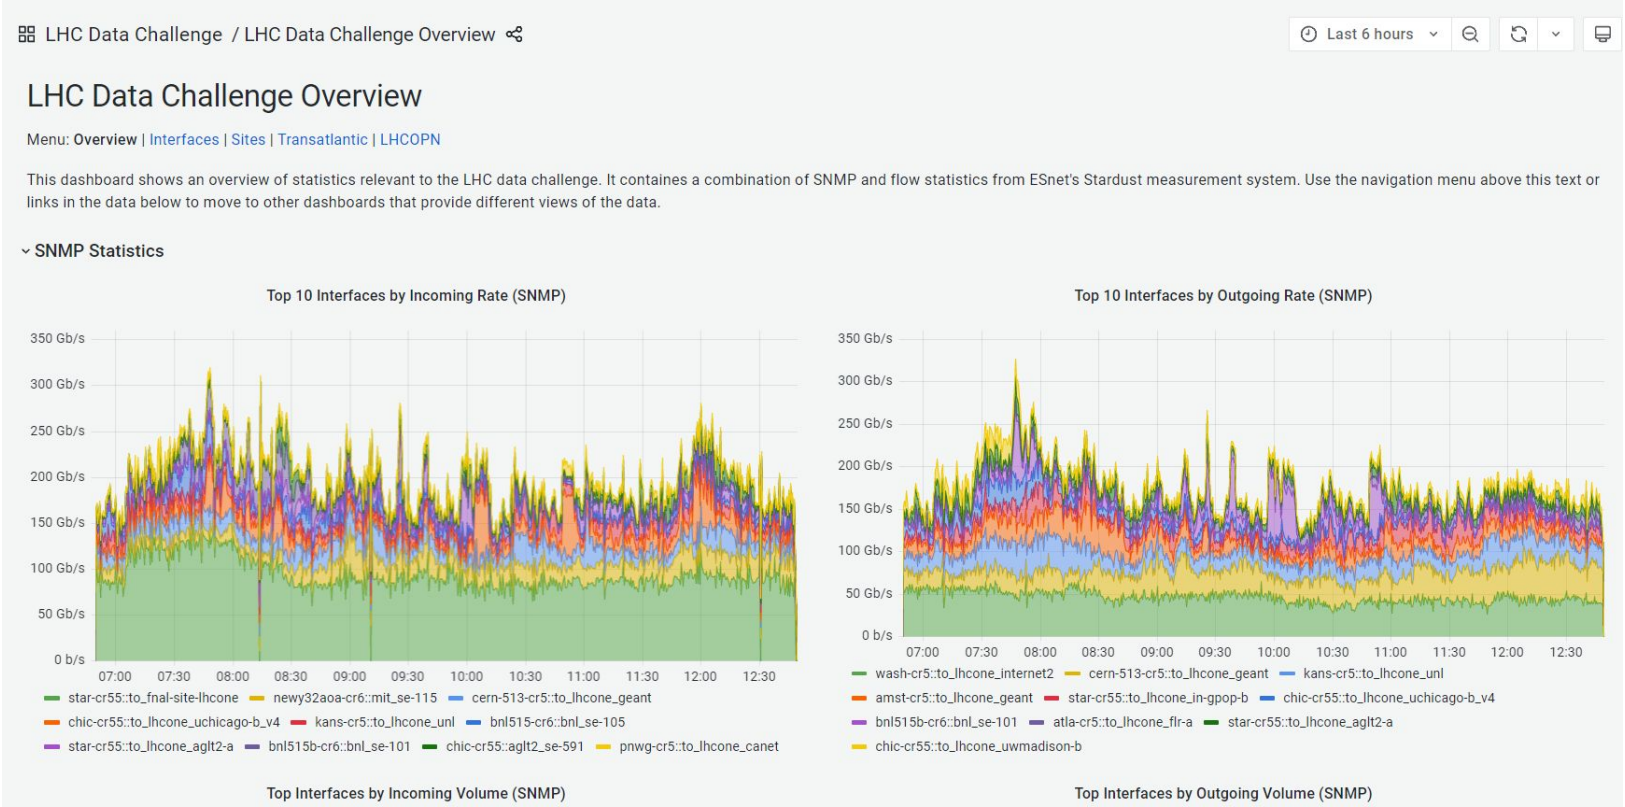

Summary

- OSG in collaboration with WLCG operates a comprehensive network monitoring platform
  - Provides data and feedback to LHCOPN/LHCONE, HEPiX, WLCG and OSG communities
- The IRIS-HEP and SAND projects have produced some new tools for exploring and utilizing our network data
- The WLCG Data Challenge was a nice opportunity to look for bottlenecks and will take a while to analyze all the results. Network doesn't **seem** to be a current limit for running at 10% of HL-LHC scale.
- We have to continue to watch our network monitoring infrastructure as it is a complex system with lots of areas for issues to develop.

Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

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

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

WLCG Network Throughput Support Unit

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.
- 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

As there is no consensus on the MTU to be recommended on the segments connecting servers and clients, LHCOPN/LHCONE working group was established to investigate and produce a recommendation. (See coming talk :))

Importance of Measuring Our Networks

- End-to-end network issues are difficult to spot and localize
  - Network problems are multi-domain, complicating the process
  - Performance issues involving the network are complicated by the number of components involved end-to-end
  - Standardizing on specific tools and methods focuses resources more effectively and provides better self-support.
- **Network problems can severely impact experiments workflows and have taken weeks, months and even years to get addressed!**
- perfSONAR provides a number of standard metrics we can use
  - Latency, Bandwidth and Traceroute
  - These measurements are critical for network visibility
- **Without measuring our complex, global networks we wouldn't be able to reliably use those network to do science**