WLCG Network Throughput Working Group Update

Marian Babik (CERN), Shawn McKee (UM)
With contributions from SAND, IRIS-HEP & OSG Network Area

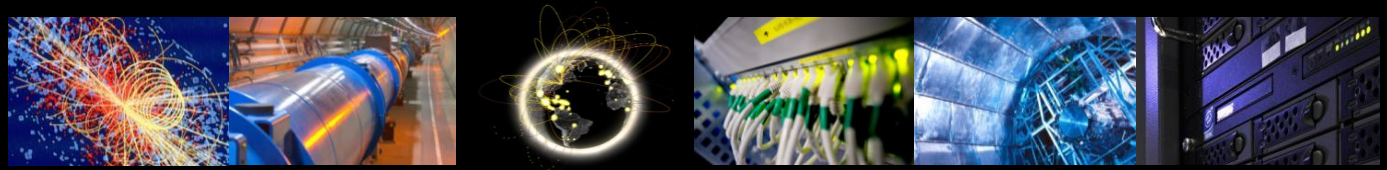

WLCG Worldwide LHC Computing Grid

Open Science Grid

Outline

- Scope
- perfSONAR infrastructure status
- Testing tools and coverage
- Network Platform Updates
- Plans and Challenges
- Summary

OSG/WLCG networking projects

There are 4 coupled projects around the core OSG Net Area

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

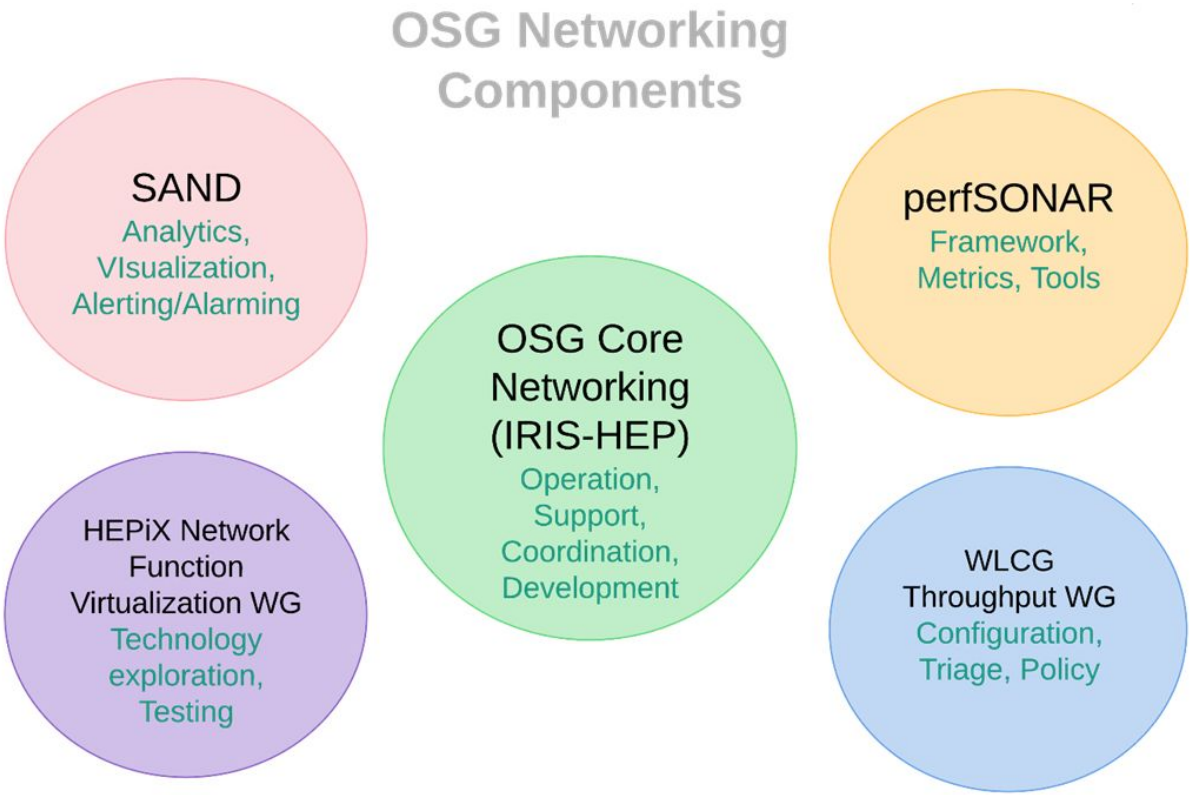

perfSONAR deployment

288 Active perfSONAR instances

- **207 production endpoints**
- T1/T2 coverage
- Continuously testing over 5000 links
- Testing coordinated and managed from a central location
- Dedicated latency and bandwidth endpoints at each site
- **Open platform** - tests can be scheduled by anyone who participates in our network and runs perfSONAR

perfSONAR capabilities

perfSONAR has significantly improved over the past few years

- Platform became quite popular - deployed at almost 2000 domains
  - Open platform gives sites the possibility to debug both local and remote network issues
- Hardware support
  - Support to run latency/bandwidth from a **single node** with multiple NICs
  - **Mobile/low-cost nodes** supported; **100Gbps nodes** relatively inexpensive now (spec)
- Growing range of available testing tools
  - Migration to pscheduler enabled easy integration of new testing tools
  - **Throughput**: iperf3, iperf2, nuttcp, ethr, s3, globus-url-copy
    - iperf2, ethr support multi-stream tests with multithreading capability
    - iperf3 can report TCP retransmits - shows when TCP gets congested
  - **Latency**: ping, owamp, twamp
  - **Path**: traceroute, tracepath, paris-traceroute
  - **Traffic**: snmp-get
- Improved security and support
  - pScheduler limits system: restricts what tests can be run, by whom, with what specs

perfSONAR News

- 4.3.2 - latest release (release notes)
  - Python3 support
  - PWA/psconfig support for multiple archivers
    - Enables publishing directly to RabbitMQ using central configuration tool
  - Integrated ethr and S3 benchmark test tools
- Traces now run in full mesh for all our endpoints
- Throughput and traces now performed for both IPv4 and IPv6
- Communities that have joined recently or plan to use our infrastructure in the near-term:
  - WLCG ALICE, EU ESCAPE, EU ARCHIVER, DUNE, SLATE
- perfSONAR F2F developers meeting took place in June
  - Plan to move to ELK stack (ElasticSearch/Grafana)
- perfSONAR session took eplace @TechExtra Nov 2nd
  - Will discuss what's in the pipeline for 4.4 release
  - Demo of the ELK stack capabilities, PWA, WIFI and Cloud monitoring

WLCG/OSG Network Testing

Tests	Communities	Notes

Throughput:
`iperf3 - 30 sec; one stream; freq: 23h`
`both ipv6 and ipv4`
Latency:
`owamp - continuous @10Hz`
`dual-stack (either ipv6 or ipv4)`
Path:
`traceroute/tracepath; freq: 30 min`
`both ipv6 and ipv4`
	ATLAS, CMS, LHCb, DUNE,	Full mesh - all sites test to all sites

	EU ESCAPE, ARCHIVER	EU projects; EU ARCHIVER plans to test to commercial clouds (based on HNSciCloud experiences) OCRE test framework

Throughput:
`iperf3 - 30 sec; one stream`
`(first 5 sec omitted); freq: 6h`
`both ipv6 and ipv4`
`Rest is the same as above`
	OPN, Belle II, ALICE, CMS T0/T1/US T2, StashCache, US ATLAS	Full mesh

	LHCONE	Disjoint mesh (Sites -> R&E providers only)
	UK, FR, CA, US CCSTAR	Full mesh regional
Throughput: `iperf3 - 30 sec; one stream; freq: 6h`	WLCG 100Gbps, US-LHC 100 Gbps	Selected sites with 100 Gbps capability

Network 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

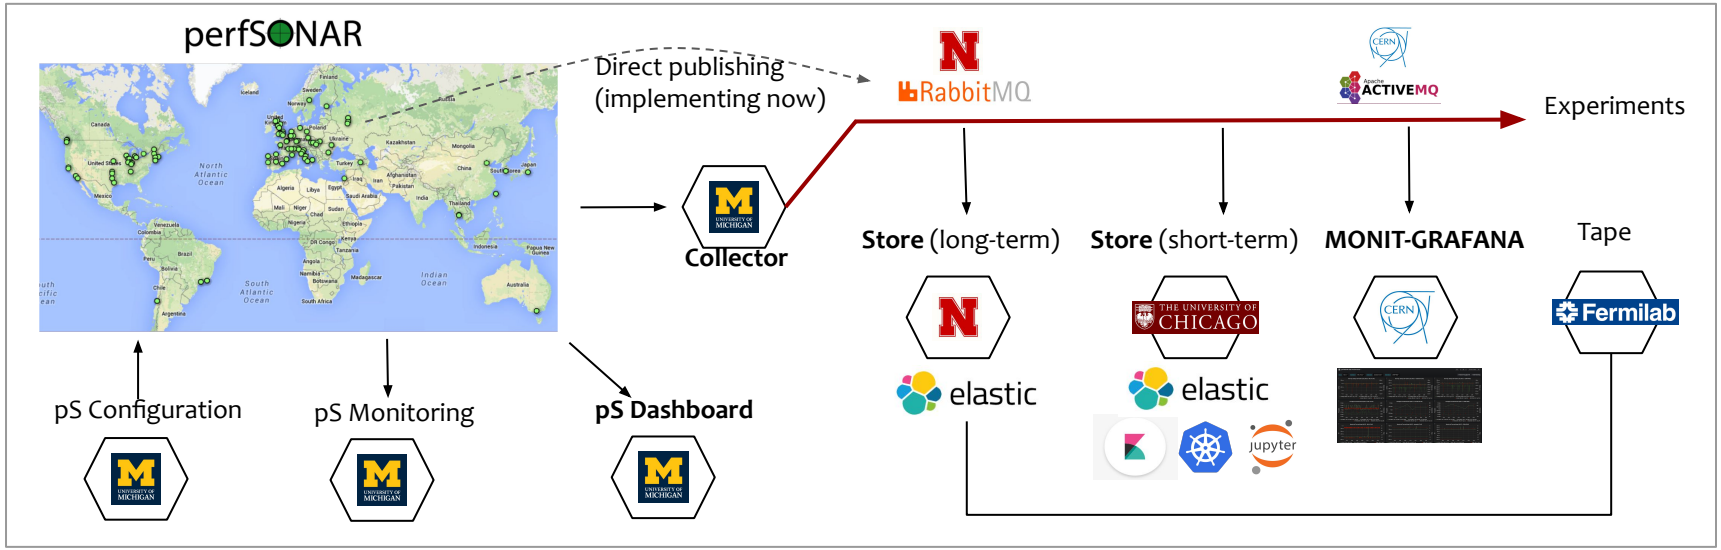

Network Platform Capabilities

Collect

- Currently working on a push model - perfSONAR publishing results directly to the message bug - enables real-time analytics, better support for docker test point, etc.
- Flexible collector - will adapt and try to pull results if they're not available on the bus
- Possibility to replay results from tape and re-collect missing periods

Store

- Both short-term and long-term storage is available via ElasticSearch (ES)
  - In collaboration with ATLAS Analytics Platform
- Ingesters are used to enrich/clean/filter data before inserting in ES
  - Topology integration (sites), statistics, etc.

Visualise

- Integrated with CERN MONIT and ATLAS Analytics infrastructures
- Multiple dashboards available:
  - High-level overview/matrix view from psmad; path visualisation from TrACER
  - High-level overview (graphs, tables) from kibana at UC ES
  - Detailed graphs in CERN MONIT Grafana

Network Throughput Support Unit

- Network issues are rare, but usually quite difficult to resolve
- **WLCG Network Throughput** was established to help sites and experiments with network performance issues using perfSONAR
  - A list of existing and resolved issues available at link
  - Gained a lot of experience over the years, now able to quickly resolve most of the common issues
  - Working in collaboration with GEANT's EDUPERT and GlobalNOC
- Common issues
  - **Maximum Transmission Unit** (MTU and MTU+load balancing)
    - WG was setup in LHCOPN/LHCONE to follow this up (MTU)
  - **Routing** - mostly LHCONE (or commercial routes instead of R&E)
    - Issues with major R&Es very rare, if so then usually during/after interventions; some small fraction due to IPv6 routes
  - **Site equipment** - insufficient buffers to support higher latencies; site uplink capacities; network design and equipment issues

Plans

Notifications and alerting

- In collaboration with ATLAS Analytics Platform
- Working on triggering events based on specific criteria
  - Low throughput for certain period, high packet loss, route changes, etc.
- Aggregating such events into alarms/notifications that sites/experiments can subscribe to

Analytics

- Path analytics - Identify common hops between multiple src and dst pairs; detect path symmetry; find all unique paths across a period (Manjari Trivedi/UM)
- Site summaries and reports of problematic node pairs (Petya Vasileva/UM) - blacklisting the problematic nodes (outliers) from the analytics; general data cleansing
- Use machine-learning to lay the foundations for a performance anomaly detection service (Edris Qarghah/UC)

Visualisation

- Consolidate existing dashboards and better integrate them into existing platforms

Challenges

- Infrastructure
  - Diverse deployment and hardware used across sites - some deployments use powerful bare metals, others use 1 CPU VMs - **important to set realistic expectations**
  - **New technologies such as containers/cloud infrastructures**
    - Main challenge is a very dynamic environment
  - Low-level technologies are also advancing (e.g. TCP BBRv2)
- Platform
  - **Topology** - missing good source of topology that would map perfSONARs to sites and vice versa (including R&E and non-WLCG sites)
  - Self-subscription user interfaces for notifications/alerts
  - Machine readable **LHCOPN/LHCONE topology** so that we can better analyse paths
- Network Throughput
  - Most of the issues during LHC runs, very few during LS
  - **Data lakes** will bring additional layer of complexity
  - **Clouds and containers** pose a great challenge (NAT, network virtualisation, traceability)
    - Both combined will require not only R&D in testing tools, but also new infrastructure

Summary

- Currently operating one of the biggest R&E perfSONAR networks
- OSG in collaboration with WLCG are operating a comprehensive network monitoring platform to collect, store and visualise network metrics
- Platform has been used in a wide range of activities from core OSG/WLCG operations to Cloud testing and IPv6 deployment
- Number of communities are now using our infrastructure:
  - ATLAS, CMS, LHCb, StashCache, Belle II, CCSTAR, Jisc (UK), GRIF/IN2P3 (FR), LHCOPN/LHCONE, ALICE, EU ESCAPE, EU ARCHIVER, DUNE and SLATE
- IRIS-HEP and SAND contributing to the R&D in the network area
- Current focus is on analytical studies that would lead to a production level network alerting/notification service
- We also expect 100 Gbps mesh to grow with a potential to contribute to the network data challenges and network performance studies

References

- OSG/WLCG Networking Documentation
  - https://opensciencegrid.github.io/networking/
- perfSONAR Stream Structure
  - http://software.es.net/esmond/perfsonar_client_rest.html
- 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 Analytics Platform
  - https://indico.cern.ch/event/587955/contributions/2937506/
  - https://indico.cern.ch/event/587955/contributions/2937891/

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
- SAND: NSF CC* OAC-1827116

Backup Slides Follow

100Gbps Testing

- LHCOPN/LHCONE 100Gbps mesh

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

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps | Unable to find test data

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

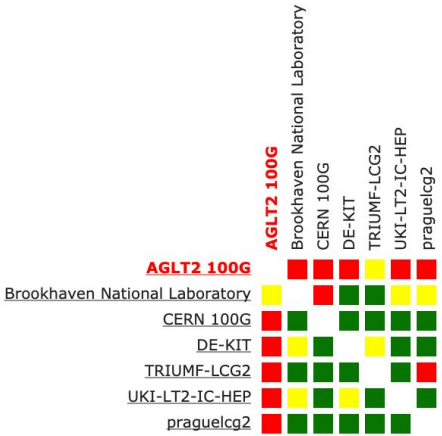

WLCG 100G Mesh - WLCG 100G IPv4 Bandwidth - Throughput

Throughput >= 1Gbps | Throughput < 1Gbps | Throughput <= .5Gbps | Unable to find test data

Found a total of 3 problems involving 2 hosts in the grid

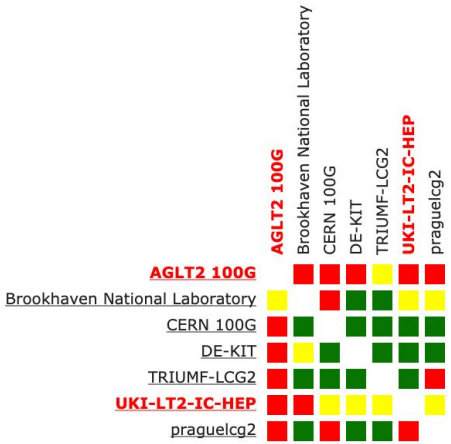

- US-LHC 100Gbps mesh currently being developed
  - 100Gbps mailing list

Network Platform News

- Updated set of dashboards created in UC ES/Kibana
  - https://atlas-kibana.mwt2.org/s/networking/goto/20dd25907d61df98a0b85b1dfaed54e1
  - Provides high-level overview of latency, loss, traces and throughput
  - Implementation of the result of student projects to provide new user tools in our production system
    - Path analytics - Identify common hops between multiple src and dst pairs; detect path symmetry; find all unique paths across a period (Manjari Trivedi/UM)
    - Site summaries and reports of problematic node pairs (Petya Vasileva/UM) - blacklisting the problematic nodes (outliers) from the analytics
    - Use machine-learning to lay the foundations for a performance anomaly detection service (Edris Qarghah/UC)
- TRACer - Path visualisation tool developed by MEPHi
  - Still in Beta, but already provides very interesting views into perfSONAR traces
  - Video of next release available at https://yadi.sk/i/tyhiA-e3GGKqDQ
- Updated toolkit info web page
  - https://toolkitinfo.opensciencegrid.org/toolkitinfo (Tommy Shearer/IRIS-HEP)

Collaboration with MEPhI on Network Visualization

Containerized Version running at UC

We have been collaborating with MEPhI (Moscow Engineering Physics Institute) on network path visualization and have developed **TRACer** to visualize our perfSONAR path measurement data

TRACer Allows us to explore traceroute measurements from our OSG/WLCG perfSONAR measurements

NOTE: We have LOTS of traceroute measurements and TRACer can overwhelm your browser. When TRACer starts, it will load a very small interval of recent data, allowing the use to apply additional filters of interest (selectings sources, destinations and time windows). If you try to load too much data it will be sloooow. We are working on improving this behavior.

For more information, we have a preprint of a paper on TRACer available at **https://doi.org/10.6084/m9.figshare.12724865.v1**

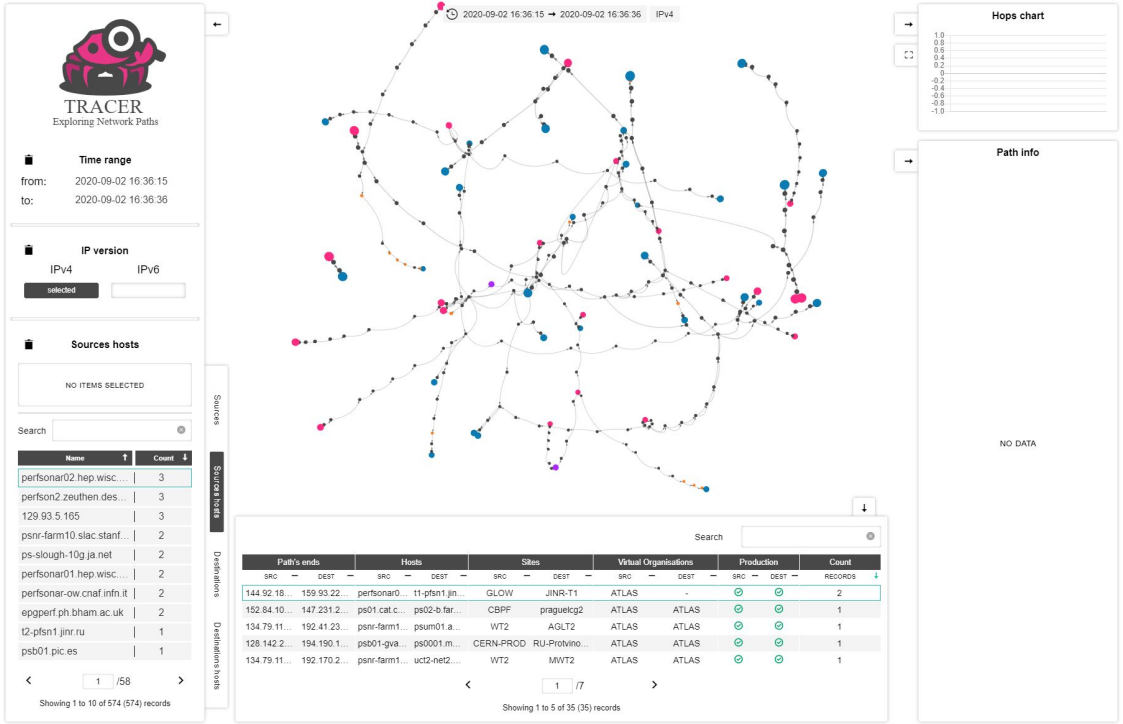

Network Platform News

- Added **new data sources** to the messaging
  - Re-established SNMP from ESNet and added HTCondor transfer metrics
- Directly publishing (push-mode) from perfSONAR toolkit
  - Most of the developments are ready, testing has started (US region for now)
  - Developed in collaboration with perfSONAR developers
- Infrastructure monitoring - updates and new capabilities coming this fall:
  - Updates to **CheckMK version 1.6.x**, including plugin updates
  - Improvements to the configuration system to make it useful for other communities (dependent only on the perfSONAR ecosystem)
  - **Self-subscription for alerts** to toolkit contact, based on joining a particular community (using perfSONAR lookup service/information system)
  - **Migration to CRIC for topology** as well as working with CRIC team on potential improvements related to mapping perfSONARs to WLCG experiment site names
  - LHCONE Grafana dashboards based on newly established LHCONE mesh
  - New metric to track overall toolkit utilisation (for throughput tests)

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 :))

IRIS-HEP (http://iris-hep.org)

The Institute for Research and Innovation in Software in High Energy Physics (**IRIS-HEP**) project has been funded by National Science Foundation in the US as grant OAC-1836650 as of 1 September, 2018.

The institute focuses on preparing for **High Luminosity (HL) LHC** and is funded at **$5M** / year for 5 years. There are three primary development areas:

- Innovative algorithms for data reconstruction and triggering;
- Highly performant analysis systems that reduce `time-to-insight' and maximize the HL-LHC physics potential;
- Data organization, management and access systems for the community's upcoming Exabyte era.

The institute also funds the **LHC part of Open Science Grid, including the networking area** and will create a new integration path (the Scalable Systems Laboratory) to deliver its R&D activities into the distributed and scientific production infrastructures. **Website for more info**: http://iris-hep.org/

The NSF funded SAND Project

SAND: Service Analysis and Network Diagnosis

This is a newly funded NSF project (award #1827116) f
visualizing, and analyzing disparate network monitoring an

It will **extend** and **augment** the **OSG networking** efforts wit
extracting useful insights and metrics from the wealth of n
from perfSONAR, FTS, R&E network flows and related netw
HTCondor and others.

Website https://sand-ci.org/ (Project just started in Septer
years)

PI: **Brian Bockelman**, **Co-PIs**: **Shawn McKee, Rob Gardner**

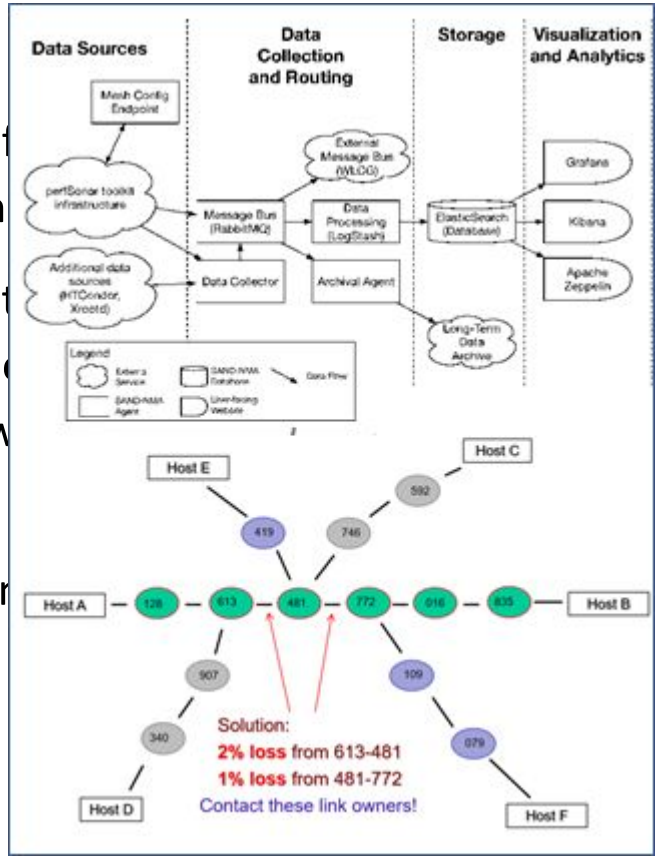

Grafana - Inter-Regional Latency Dashboard

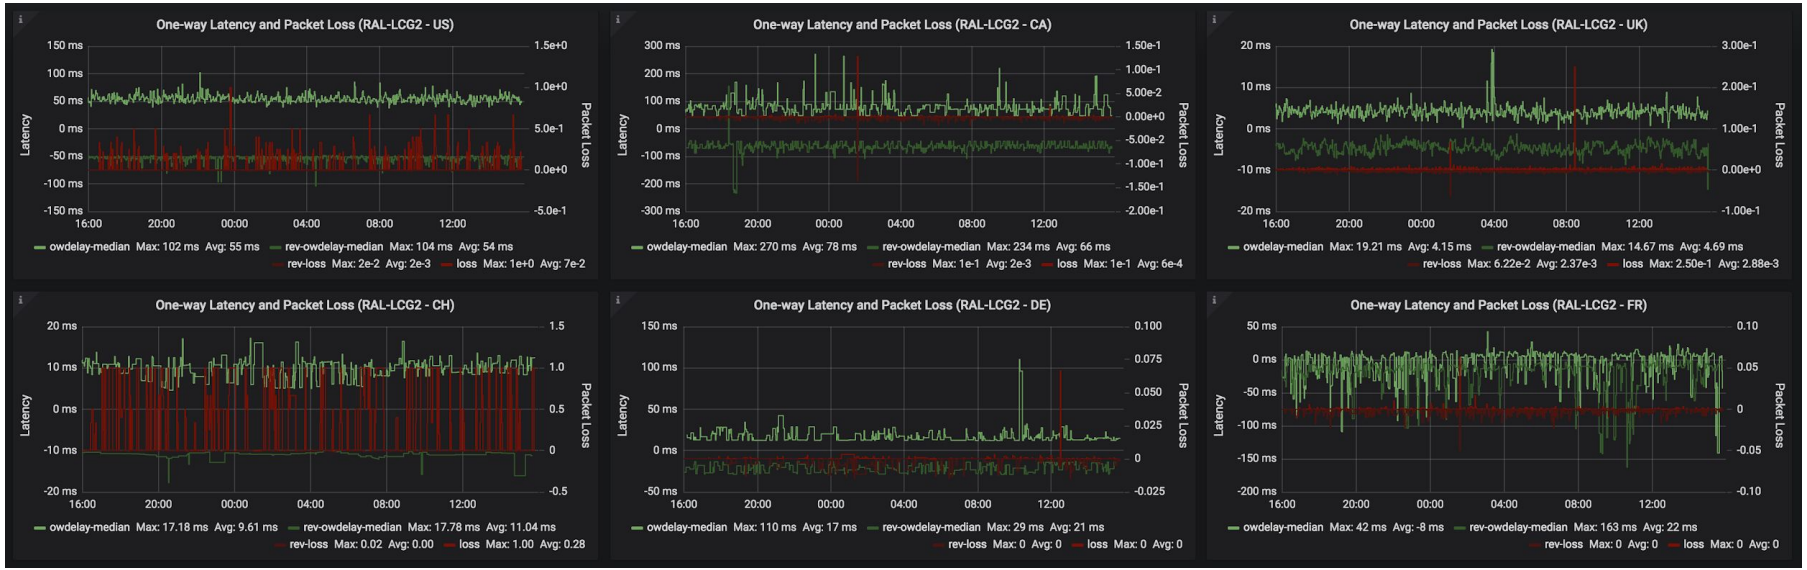

Networking Challenges

There are number of challenges in the networking, which will require improved collaboration with other sciences as well as HEP-focused R&D:

- Capacity/share for data intensive sciences
  - No issues wrt available technology, however
  - What if N more HEP-scale science domains start competing for the same resources ?
- Remote data access proliferating in the current DDM design
  - Promoted as a way to solve challenges within experiment's DDM
  - Different patterns of network usage emerging
    - Moving from large streams to a mix of large and small frequent event streams
- Integration of Commercial Clouds
  - Impact on funding, usage policies, security, etc.
- Technology evolution
  - Software Defined Networking (SDN)/Network Functions Virtualisation (NFV)

Network Evolution Areas

The following are some of the key areas for HEP Networking R&D:

- Improving efficiency of data transfers
  - TCP BBR - version 2 is in the works with promising improvements
  - Exploring alternative protocols for transfers (UDP)
- Caching
  - Data caches co-located with network hubs in a similar way as on commercial CDNs
- Federations/Clouds
  - Overlay networks spanning multiple domains
  - Multi-clouds - expanding DC networking via L3VPNs
- Technology
  - SDN/NFV approaches - currently looked at by HEPiX NFV WG
    - Compute - Agile service delivery on Cloud Infrastructures (OpenStack, Kubernetes)
    - Data Transfers - Network resource optimisation - dynamically optimising the network based on its load and state (more in Shawn/Ilija)
  - SD-WAN approaches - https://www.mode.net/

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**