The SAND Project at the Halfway Point

Shawn McKee / University of Michigan
on behalf of the SAND Collaboration

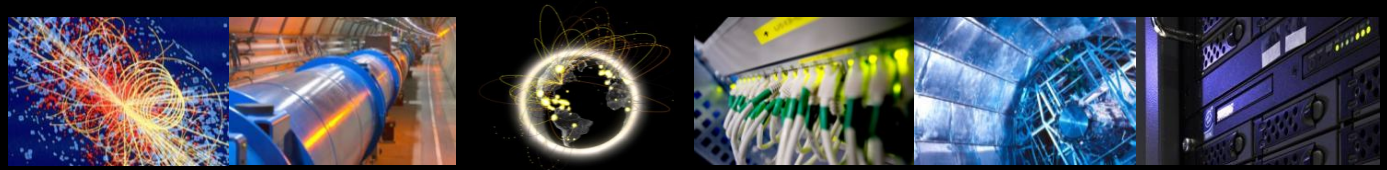

WLCG Worldwide LHC Computing Grid

HEPiX, Amsterdam, October 2019

Open Science Grid

Outline

- Overview of WLCG/OSG Networking; The Motivation for SAND
- The SAND project and Goals
- Project Activities and Results so far
- Next steps

The Motivation for SAND

- OSG is in its 7th year of supporting WLCG/OSG networking focused on:
  - Assisting its users and affiliates in identifying and fixing network bottlenecks
  - Developing and operating a comprehensive Network Monitoring Platform
  - **Improving our ability to manage and use network topology and network metrics for analytics**
- WLCG Network Throughput Working Group was established to ensure sites and experiments can better understand and fix networking issues:
  - Oversees the **WLCG perfSONAR infrastructure**
  - **Coordinates WLCG network performance incidents** - runs a dedicated support unit which involves sites, network experts, R&Es and perfSONAR developers
- **After this long-term work, we find ourselves with a very rich dataset in need of exploitation but without the effort available to deliver for our community...thus….**

The NSF SAND Project

Service **A**nalysis and **N**etwork **D**iagnosis

a NSF funded project (award #1827116) focusing on combining, and analyzing disparate network monitoring and service logging data.

(**GOAL**: capitalize on our rich network dataset!!)

Website
https://sand-ci.org/
(Project started in September 2018 and has 2 years funding)

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

Brian Bockelman
Research Assistant Professor
University of Nebraska
Computer Science & Engineering
bockelman@unl.edu

Shawn McKee
Research Scientist
University of Michigan Physics
smckee@umich.edu

Rob Gardner
Senior Scientist
University of Chicago Physics
rwg@hep.uchicago.edu

SAND Project Vision

It will **extend** and **augment** the **OSG networking** efforts with a **primary goal** of extracting useful insights and metrics from the wealth of network data being gathered from perfSONAR, FTS, R&E network flows and related network information from HTCondor and others.

Shown on the top diagram to the right is the logical **SAND** data flow from source to analytics.

The bottom diagram to the right shows the potential power of the extensive network tomography we have by continuously measuring thousands of R&E network paths. In this example, 3 host-pairs see differing packet loss on intersecting paths. **We can infer a solution!**

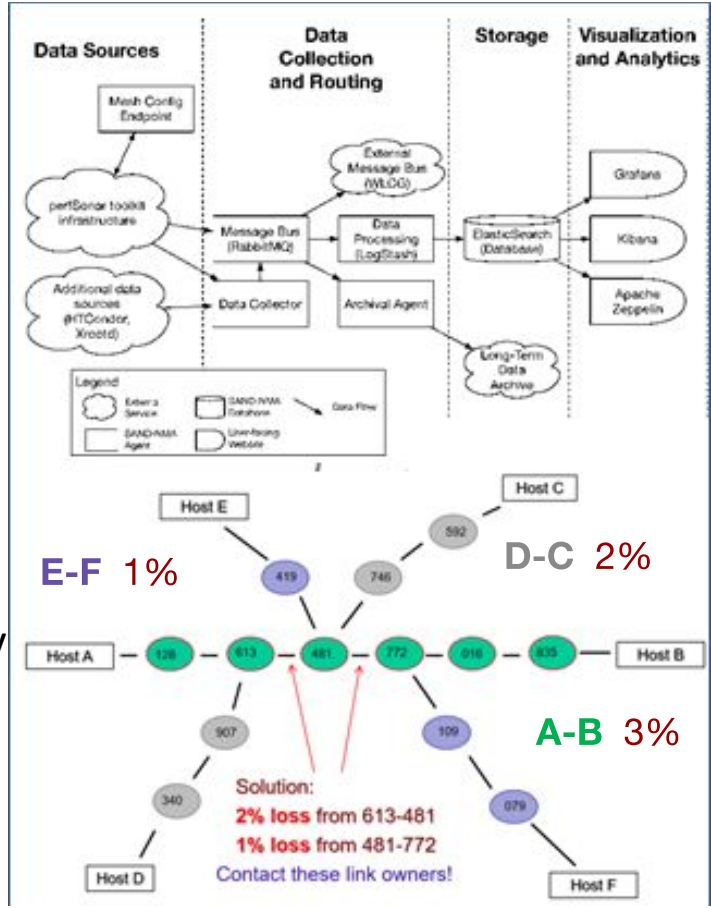

SAND Activities to Date

Initial efforts targeted improving the network data pipeline from OSG

- OSG was using an infrastructure called RSV to gather perfSONAR data
  - There were issues with reliability and latency
  - With help from the SAND project, a new collector was created that has much lower latency, more complete monitoring and is significantly more robust.
- The network metrics being collected were only going to an ELK-based analytics platform in Chicago
  - We added a new “long-term” ELK destination in Nebraska
  - We also added tape backup of the data at FNAL (tested and successfully used this year!)
- Initial planning for a new push-based (from each toolkit) model is ready
  - Push based version tested in August and working with manual configuration
  - BUT we need a way to centrally configure the pS RabbitMQ archiver. Targeting ~Dec

We have also been working with the collected data and have identified challenges that we need to address to make it more useful

- As part of the we augmented the traceroute with ASN info.

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

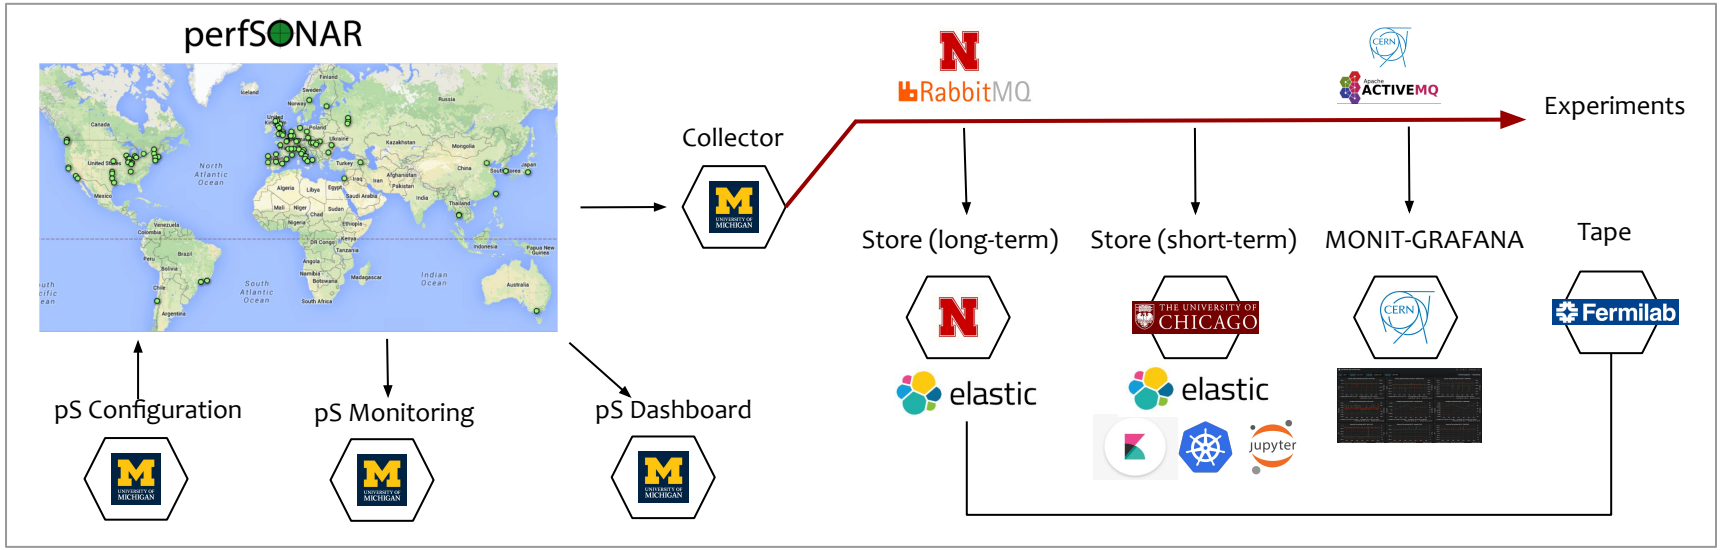

Available Data Overview

SAND and OSG/WLCG are gathering a number of potentially very useful metrics

- **perfSONAR** data from over 260 instances all over the world
- **ESnet** network traffic (snmp counters)
- **WLCG** data transfers (FTS)
- **LHCOPN** data (from CERN networking)

This data is being transferred using message bus technologies (RabbitMQ (OSG) and ActiveMQ (CERN)) and ends up in two different Elasticsearch instances (University of Chicago analytics platform and University of Nebraska)

This data could provide powerful insights into our R&E network infrastructure by using the **temporal** and **spatial** information we have available.

perfSONAR Data Details

We are collecting a number of different types of data from perfSONAR which are sent to different “topics” on the RabbitMQ bus and put into their own index in Elasticsearch:

- **ps_alarms** : These are generated alarms based on other ps indices
- **ps_meta :** Tracks toolkit version, host info, various metadata
- **ps_owd :** One-way Delay measurements from perfSONAR (latency)
- **ps_packet_loss :** The percentage of packets lost in latency testing (10 Hz)
- **ps_retransmits :** During throughput testing, tracks retransmits
- **ps_status :** Tracks status of measurements (coverage, efficiency)
- **ps_throughput :** Measures throughput via iperf
- **ps_trace :** Measures the layer-3 network path via traceroute

You can explore the details via Kibana:

https://atlas-kibana.mwt2.org/s/networking/app/kibana#/discover?_g=()

SAND Collaboration Meeting Details

Our face-to-face collaboration meeting was held June 17-18, 2019 at U Chicago

Main topic areas discussed day 1

- Network pipeline
- Monitoring tools
- Containerizing perfSONAR
- Engaging with and enabling a broader community
- Topology and data cleaning

The second day was a “hackathon” were we worked on items from day 1.

The “**Team**”
Picture credit: **Rob Gardner** (that’s why he’s missing)

SAND Near-term Plans

- A goal of SAND is to create new analytics, visualizations and user-interfaces to extract value from the perfSONAR (and related) network metrics
- **Initial architecture**: Data-pipeline to ELK stack, visualizations via Kibana, Grafana and perhaps other tools, analytics via Jupyter notebooks and creation of "architecture plugins" to leverage this framework.
  - Examples:
    1. Alarming dashboards that show Top-N problem links (SRC-DEST with largest packet loss in last N hours, SRC-DEST with most routes in last N hours, SRC-DEST with largest change in measured throughput in last N hours, SRC with most average packet loss averaged over all DEST, DEST with most average packet loss averaged over all SRC)
    2. Route correlation: Identify SRC-DEST pairs with similar behavior changes at a point in time and analyze common hops in their routes
    3. Alerting system based upon alarming and route work. Users subscribe to various alerts using SRC, DEST, packet-loss, change in BW, etc

SAND Near-term Plans (2)

As noted we just completed our first face-to-face collaboration meeting in June and we have a few items on our list:

- Network topology - cleaning, re-organizing, visualizing.
- On-demand perfSONAR (containerized variants for specific use-cases)
- Engaging the broader NSF research community (CC* grant recipients)
- Improving end-users ability to find networking information
- Transitioning from a “pull” data model to a secure “push” model

The next few slides will cover these plans

Network Topology

Whenever we identify a possible network problem, the first question is: **what path is being measured**?

- Internet paths are designed to change in response to network changes
- Knowing the path in place when a problem is identified is critical
  - We need this path to know where to look for issues.
  - A change in the path could actually be the cause of the problem.

It should be noted that having many paths continuously monitored is a very powerful tool for both identify network issues and localizing them!

- Gedanken experiment: at approximately the same time, 5 host-pairs show an increase in packet loss. **What is the inference we can make by correlating their paths?**

Fortunately, we are scheduling regular “traceroute” tests between our perfSONAR measurement end-points

Unfortunately, the output of traceroute can be problematic in many ways!

Issues with Traceroute and Network Paths

While we regularly try to measure the network paths between our hosts (and by proxy, between our sites), the traceroute tools has some limitations

- It sometimes doesn't reach the destination
- Hops along the way can fail to respond in time, leaving "holes" in the path
- The trivial variations in traceroutes can lead to 10's of thousands of routes
- The "route" it delivers can be false
  https://www.cellstream.com/reference-reading/tipsandtricks/403-ecmp-linux-paristr

For all these reasons, we have **challenges** in trying to use our traceroute results to understand the network topology

The SAND project is planning to work on cleaning things up

- We are trying to identify logical paths to contain trivially varying physical paths to simplify things
- We need to identify when multiple links might exist at L2
- We have added "AS" number to the traceroute data to simplify understand when a major route change happens.
- **We are working on ways to visualize, compare and understand our network paths**

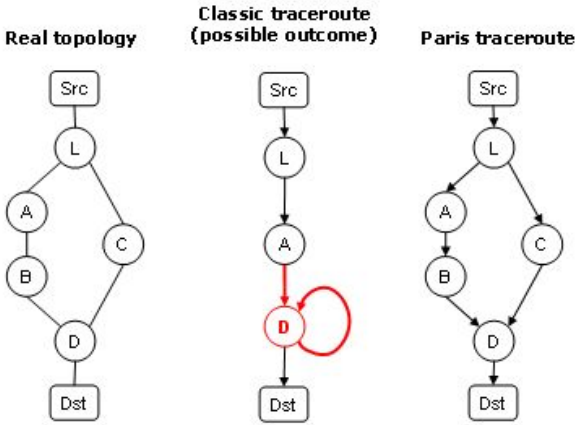

Students Working with SAND

During the spring of 2019 we engaged a group of students to work on analysis and visualization of our network metrics

- At **Chicago** we have **Sushant Bansal** (Master's student in UC Computer Science) focusing on machine learning
- At **Michigan** we have **Manjari Trivedi** (Undergraduate) and **Yuan Li** (recent Master's graduate UM School of Information) focusing on path analysis
- In **Bulgaria** we have **Petya Vasileva** (PhD student) working on R&E network analytics

The students have worked independently over the summer learning about the data we have and the analytics platform itself

For this Fall, the goal is to **clean up** and **annotate** the path information, filtering out bad or incomplete traceroute measurements and then work on analyzing, organizing and displaying path information with corresponding network metrics like packet-loss, throughput or delay

Prototype path display using network metrics from ES

Visualizing NSF CC* Institutions

The NSF has had a very successful series of Campus Cyberinfrastructre (CC) solications, and all require recipients to deploy perfSONAR

SAND wants to make it easy for these sites to be seen by simply adding a ‘CCSTAR’ community to their perfSONAR toolkits https://display.sand-ci.org/

Of course showing them on the map is just a first step

Our next step is to provide a very easy way for sites to “opt-in” to **SAND** so the we can begin to gather their perfSONAR data and provide our analytics, alerting and monitoring for them.

Finding Relevant Information

So far I have shown a few different links. Another area the SAND team would like to improve is to make it easier to find all the relevant tools, docs and data We have setup a web server at: https://toolkitinfo.opensciencegrid.org/toolkitinfo/

Our intent is to continue to maintain and add-to the various menus available to allow a broad range of users to easily find and access network data and analytics results.

We will be adding info on any future containerized perfSONAR, new topology capabilities and links on adding your site data to SAND.

SAND, Machine Learning and a Network Database

Given the scope and duration of the SAND project, we must be limited in what we try to undertake. There are **two areas** that we feel could be valuable to pursue but will take more effort than the project may have to spend:

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. Requires cleaned, annotated data to make progress (significant effort)
   b. Petya and Sushant are both interested in working in this area
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. Would require continuous real time updates as metrics arrive
   b. Could quickly identify problematic links

Collaboration with MEPhI on Network Visualization

Containerized Version running at UC https://perfsonar.uc.ssl-hep.org/graph/viewer **(Try it!)**

Network Traces Graph Visualization

MEPhI Team

Data Source

ElasticSearch: **Chicago ES**
Index name: **ps_trace**

Approach

Django + PostgreSQL

JavaScript libraries:

- three.js,
- d3-force-3d,
- jquery

Current status & plans

- **Django + ElasticSearch (~~PostgreSQL~~)**
- Adaptive design
- Adding advanced search and new features ...

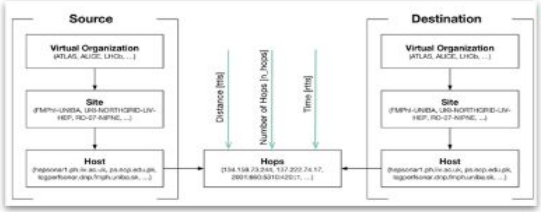

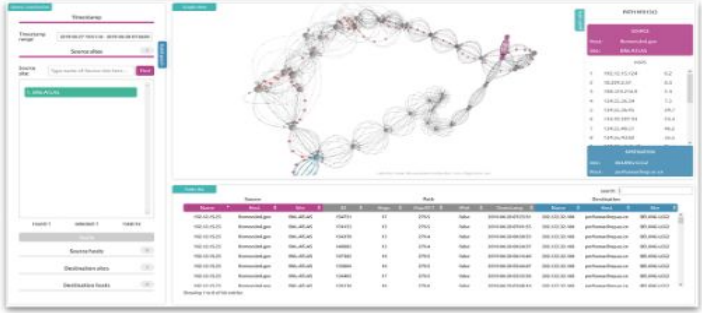

Source self-directed paths

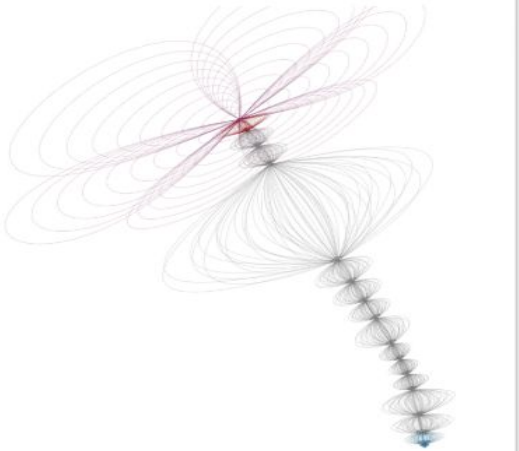

Lost nodes trace

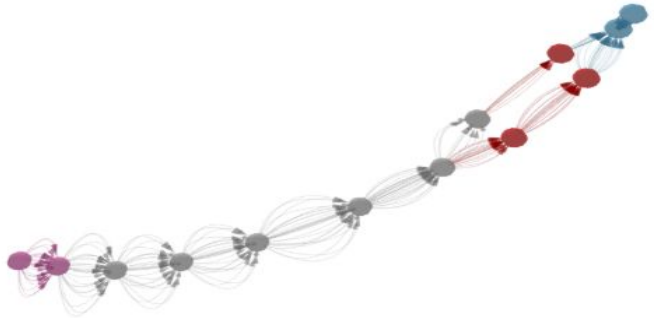

Summary

- The **SAND project** is working to
  - Maintain an effective, efficient metrics pipeline
  - Provide an infrastructure to monitor our networks and analyze various metrics
  - Extract new insights from measurements of our existing, complex global infrastructure.
- **The primary goal for SAND is to better extract “value” for our Scientists, Site and Network Administrators from the extensive network metrics OSG/WLCG is gathering.**
- We are looking for collaborators with an interest in any of the topics I covered. Contact us if you or your group are interested.

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:

- SAND: NSF CC* OAC-1827116

Questions, Comments?

This material is based upon work supported by the **National Science Foundation** under Grant No. **1827116**.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- SAND webpage
  - http://sand-ci.org
- 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/
- Grafana dashboards
  - http://monit-grafana-open.cern.ch/
- SAND GitHub Repository
  - https://github.com/sand-ci

Backup Slides

Providing Easy-to-Use-and-Deploy perfSONAR

- **Working in collaboration with the SLATE team (https://slate-ci.io) we want to develop some easy to deploy, containerized perfSONAR instances.**
- **perfSONAR "toolbelt"**: Use case "*I have a questionable network endpoint and would like to understand the performance from that endpoint to known remote endpoint XYZ. Perform the tests (oneshot) and tell me about my network. Email me a text document with the results.*" In terms of usefulness:
  - "**docker run perfsonar-toolbel**t": results in a text document summarizing the performance from localhost to fixed known endpoint.
  - "**docker run perfsonar-toolbelt foo.example.com**": results in a text document summarizing performance from localhost to foo.example.com.
  - "**docker run sand-toolbelt**": With secrets provided and a registered endpoint in the mesh config, sets up a continuous set of tests that reports to SAND-NMA.

Some Context: IRIS-HEP

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/

Grafana - perfSONAR dashboard

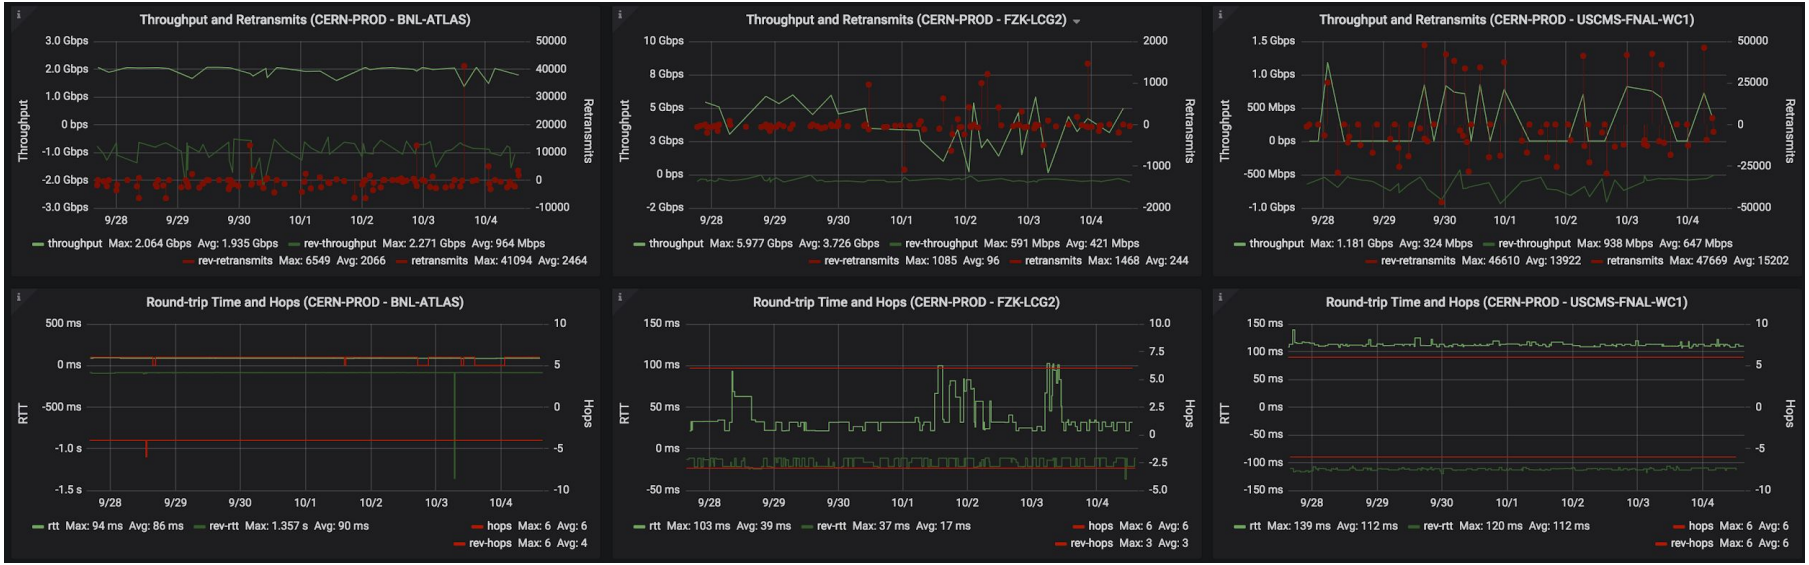

- Now includes all WLCG sites that run perfSONAR
  - Additional work needed to better filter production nodes
- Added additional row that tracks RTT and number of hops as reported by traceroute/tracepath
- Can you spot the network issue(s) above ?

Grafana - IPv6 dashboard

- Added IPv6 dashboard
  - Side-by-side comparison btw. IPv4 and IPv6 performance
- Due to performance limitations it was agreed that won't configure IPv6 latency tests

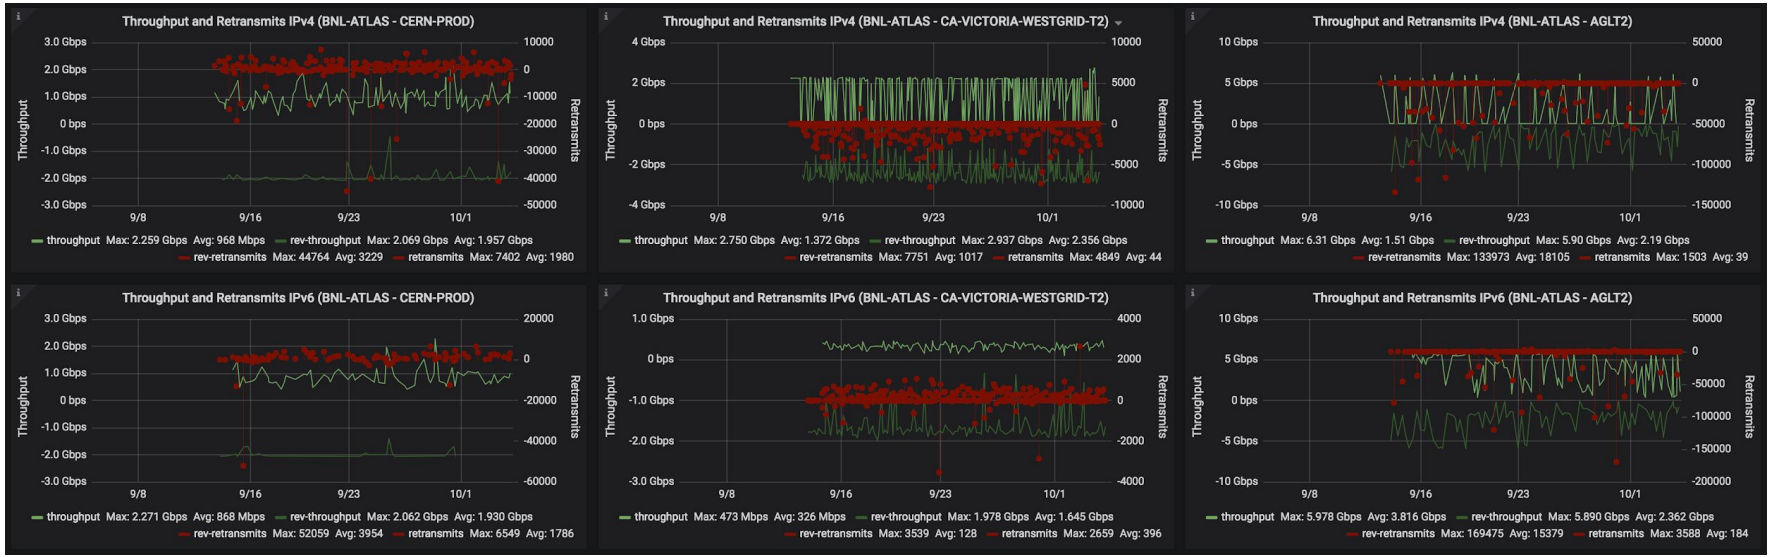

See more Grafana dashboards at http://monit-grafana-open.cern.ch/