

The Service Analysis and Network Diagnostic (SAND) Project

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



Outline

- Overview of WLCG/OSG Networking; The Motivation for SAND
- Existing data pipeline and dashboards
- Information on the “context” for SAND
- The SAND project and Goals

Reminder: OSG/WLCG Networking Activities

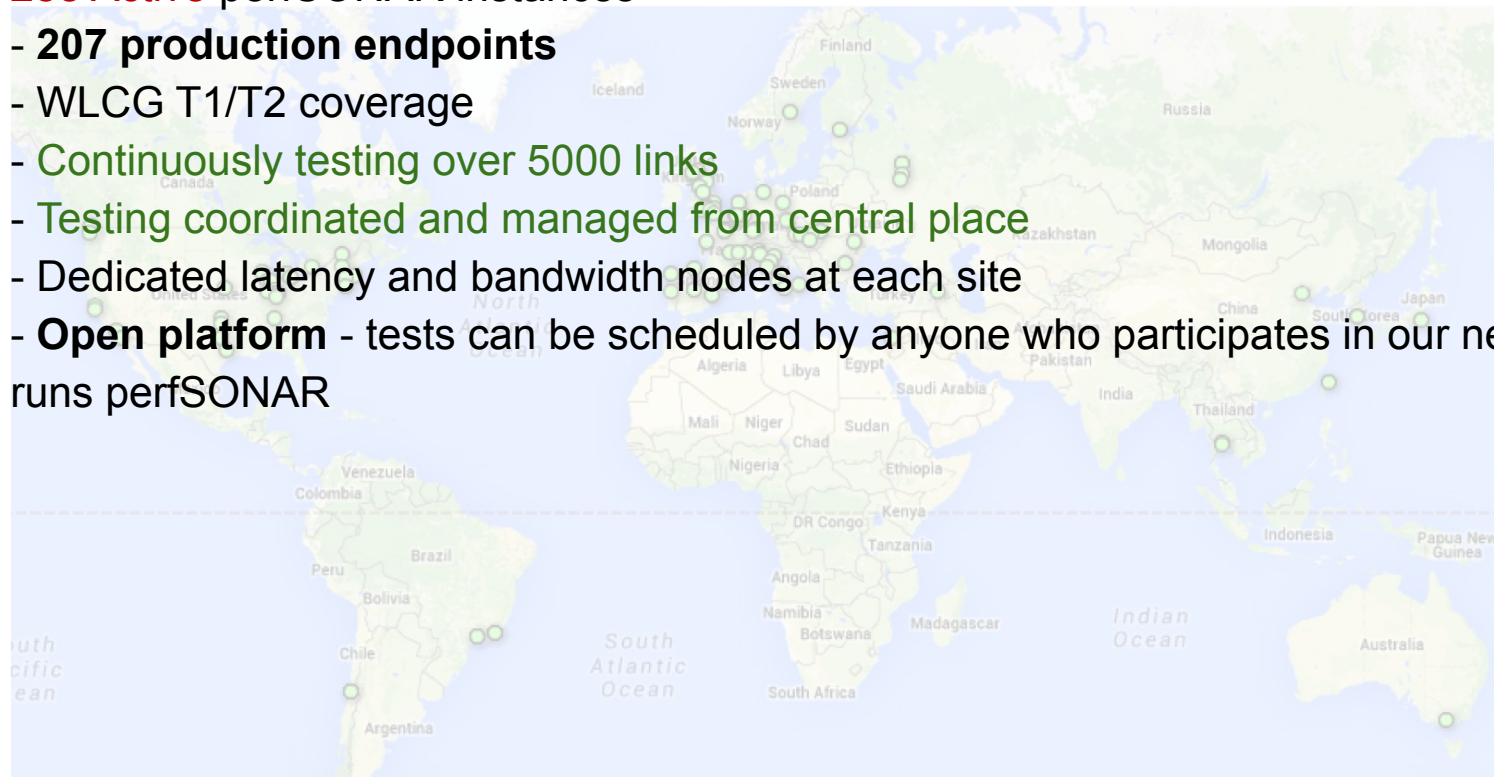
- OSG is in its 7th year of supporting WLCG/OSG networking focused on:
 - Assisting its users and affiliates in identifying and fixing network bottlenecks
 - **Developing and operating a comprehensive Network Monitoring Platform**
 - Improving our ability to manage and use network topology and network 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**
 - Core infrastructure for taking network measurements and performing low-level debugging activities
 - **Coordinates WLCG network performance incidents** - runs a dedicated support unit which involves sites, network experts, R&Es and perfSONAR developers
 - Many issues are potentially resolvable within the working group



perfSONAR deployment

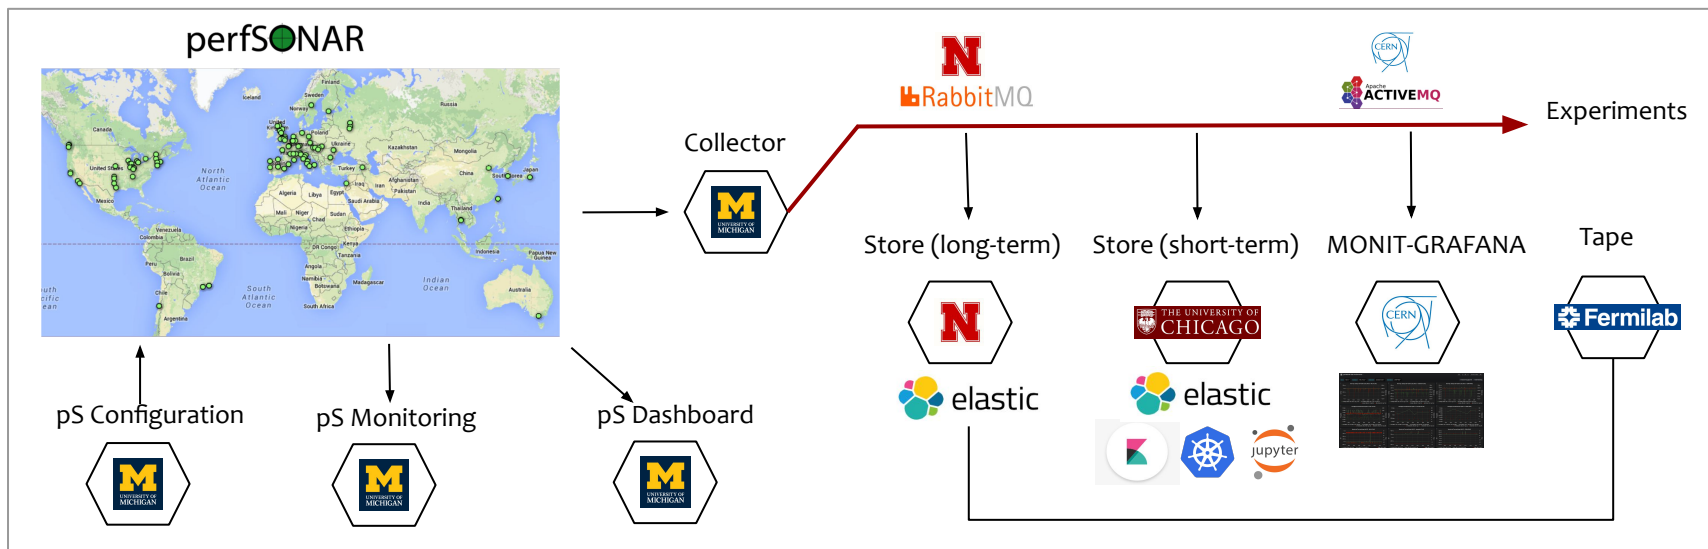
288 Active perfSONAR instances

- **207 production endpoints**
- WLCG 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** - tests can be scheduled by anyone who participates in our network and runs perfSONAR




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 ESNNet, LHCOPN traffic also via same channels




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 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/>

Current Platform Use

- **WLCG and OSG operations**

- Baseline testing and interactive debugging for incidents reported via support unit
- Regular reports at the WLCG operations coordination and WLCG weekly operations
- Providing **Grafana dashboards** that help visualise the metrics

- Enabling analytical studies - data stored in the ATLAS Analytics platform

- Providing an important source for network metrics (bandwidth, latency, path)

- **Cloud testing - HNSciCloud** - testing commercial cloud providers

- Baselining and evaluating network performance: critical to evaluate effectiveness fo LHC
- Relevant to Matt's presentation this morning

- HEPiX IPv6 WG

- Now testing bandwidth and paths over IPv6

- **Collaboration with other science domains deploying perfSONAR**

- E.g., US Universities, Pittsburgh Supercomputer Center, European Bioinformatics Institute
- Also close collaboration with (N)RENs who provide LHCONE perfSONAR coverage

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/>



The NSF SAND Project



SAND: Service Analysis and Network Diagnosis

This is a NSF funded project (award #1827116) focusing on combining, visualizing, 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 will last 2 years)

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

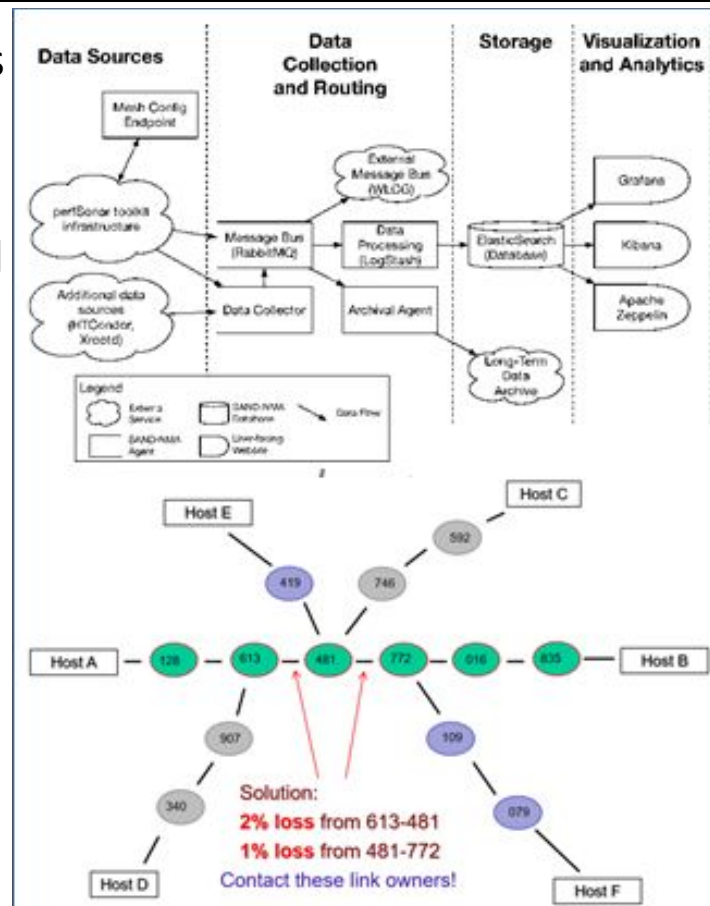
Senior Scientist
University of Chicago
Physics
rwg@hep.uchicago.edu

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 been 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**




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
 - Planning to have a few instances running “push” based in August (after pS 4.2 is out)

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 (See later details)



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

SAND Collaboration Meeting Details

Our first face-to-face collaboration meeting was held June 17-18 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”
where we worked on items from day 1.



The “Team”

Picture credit: **Rob Gardner** (that’s why he’s missing)



SAND Near-term Plans

- Main 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 is 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



Providing Easy-to-Use-and-Deploy perfSONAR

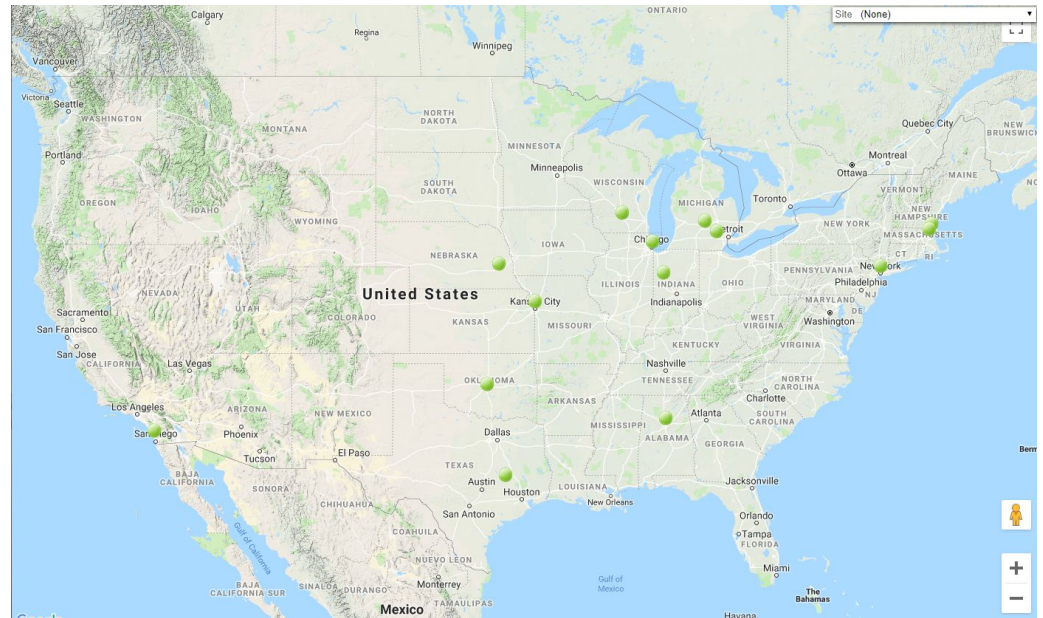
- Working in collaboration with the SLATE team (<https://slate-ci.io>) we want to develop some easy to deploy, containerized perfSONAR instances.
- perfSONA “toolbelt”: Use case *“I have a questionable network endpoint and would like to understand the performance from that endpoint to known remote endpoint 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-toolbelt”**: 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.

Visualizing NSF CC* Institutions

The NSF has had a very successful series of Campus Cyberinfrastructure (CC) solicitations, 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

We want to then 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/>

The goal 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 to adding your site data to SAND.

The perfSONAR Toolkit Information Page

Open Science Grid

WLCG
Worldwide LHC Computing Grid

Select toolkit: Submit

OSG Network Pipeline Pipeline Alarms Documentation OSG Network Services Analytics and Dashboards

Your selected perfSONAR Toolkit is:

Customized Web links for

- [This toolkit's web interface](#)
- [This toolkit's timeline of service availability](#)
- [Monitoring of this toolkit's services/configuration](#)
- [Testing instructions for this toolkit \(JSON\)](#)
- [This toolkit's settings and status](#)

Host sea... 300 rows /DC=ch/DC=cern/OU=Organic Units/OU=Users/C...

state	Host	Icons	OK	Wa	Un	Cr	Pd	state
UP	alice14.spbu.ru		40	0	2	4	0	UP atlas-
UP	atlas-npt2.bu.edu		12	0	0	2	0	UP atlas-

iris hep
Institute for Research & Innovation
in Software for High Energy Physics

NSF

[Contact us about this webpage](#)

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. Have a possible PhD student in Bulgaria who may be working in this space
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

Summary

- The SAND project is working to
 - Maintain an effective, efficient metrics pipeline
 - Provide an infrastructure to monitor our infrastructure 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.

Questions?



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/>
- ATLAS Analytics Platform
 - <https://indico.cern.ch/event/587955/contributions/2937506/>
 - <https://indico.cern.ch/event/587955/contributions/2937891/>

