The Service Analysis and Network Diagnosis (SAND) Data Pipeline

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

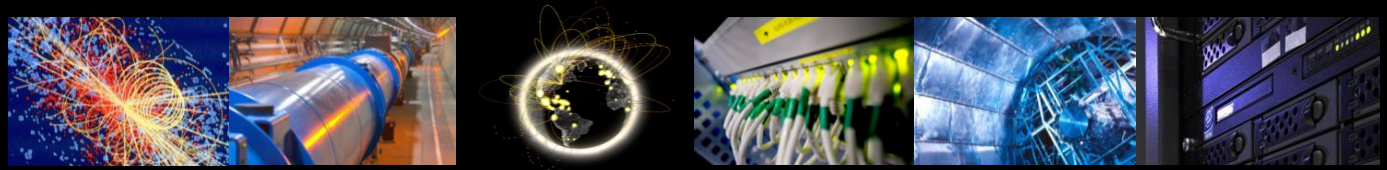

WLCG Worldwide LHC Computing Grid

SC21 Workshop: INDIS, November 15, 2021

Open Science Grid

Outline

- Overview of WLCG/OSG Networking; The Motivation for SAND
- The communities involved
- The pipeline architecture and components
- Data collection details
- Platform use and analytics
- Results and Future Work

Project Problem Space and Motivation

The motivation for the **SAND** project was to help data-intensive, distributed science collaborations make better use of the network by enabling proactive identification of network problems that impact their ability to use the network.

- Network problems are typically **difficult to correctly identify** and often **take a long time to resolve**, in part because there is no single owner or manager of the end-to-end networks used to interconnect resources.
- While the **perfSONAR** project provides tools and a framework that allow both scheduled and on-demand testing of the network from a specific host, much more is needed to identify and localize network problems.

SAND was created to pipeline data from a set of perfSONAR nodes into an analytics platform capable of displaying, analyzing and alerting on identified problems.

SAND Project Vision

It **extends** and **augments** 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!**

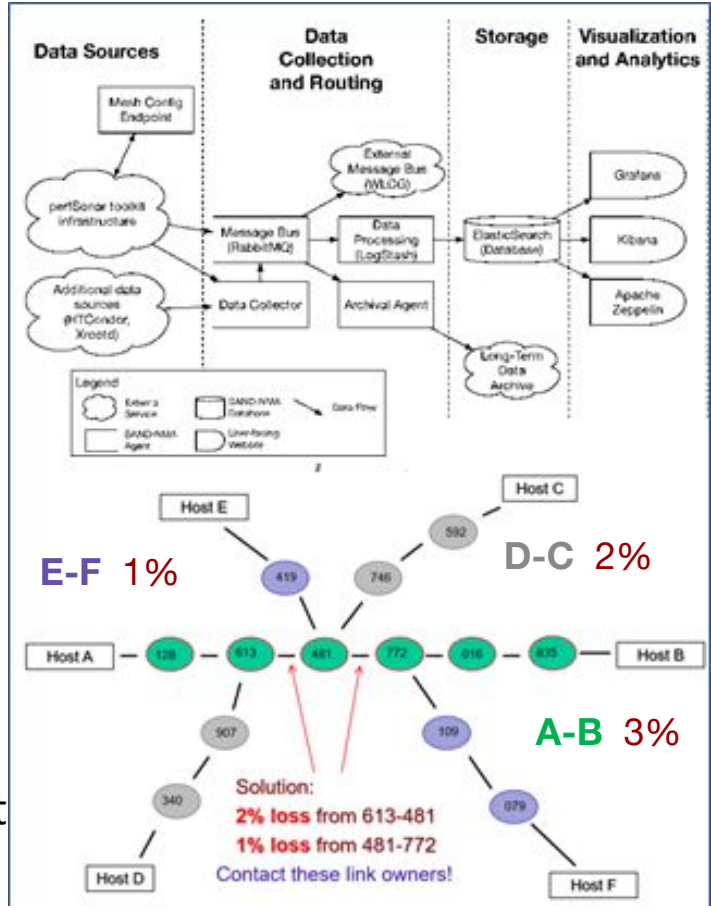

SAND Overview

- SAND has created 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 and Grafana, 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

The Communities Involved

The SAND project has interacted with a number of communities

- The Open Science Grid (**OSG**) as part of its networking area efforts to help its members understand and debug their networks.
- The Worldwide LHC Computing Grid (**WLCG**) to help organize and provide more value from its global perfSONAR toolkit deployment.
- Various **Universities** who are deploying and using perfSONAR.

The SAND data pipeline is gathering data from 288 Active perfSONAR instances

- **207 production endpoints**
- WLCG Tier-1/Tier-2 global 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

The Data Pipeline Architecture

- 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

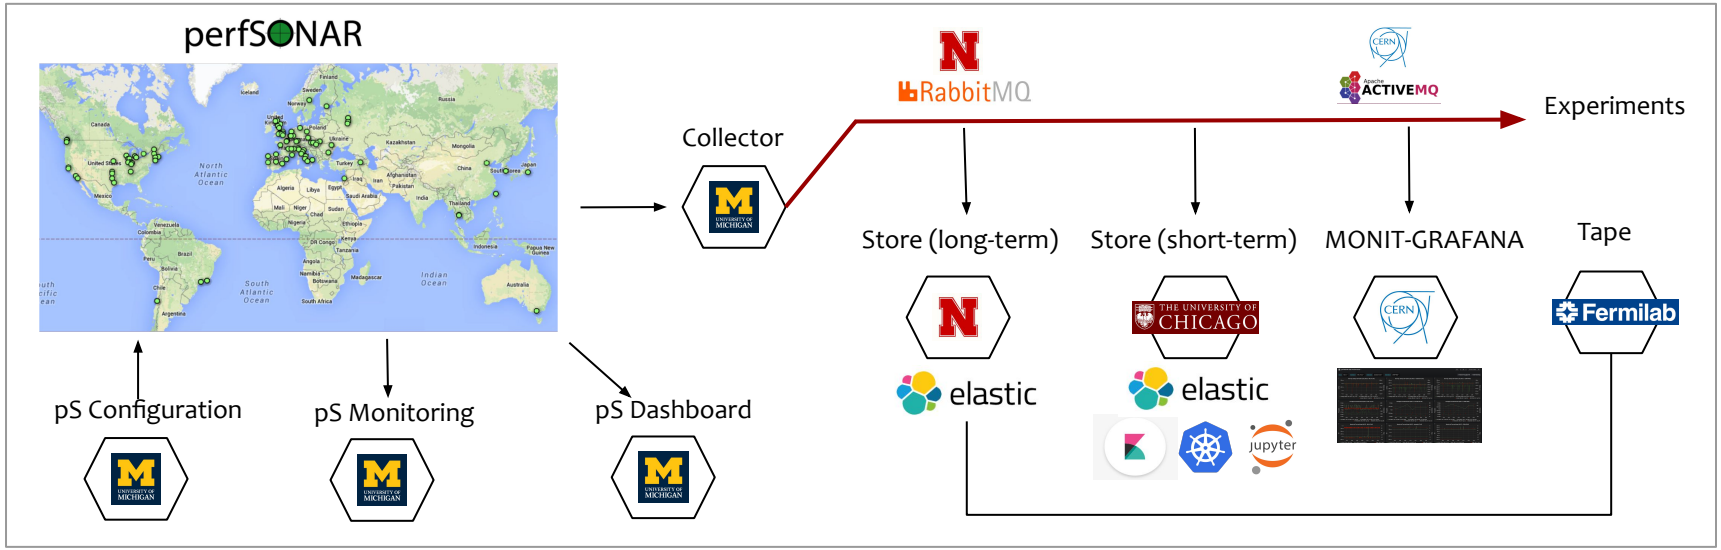

Data Collection Details

SAND is 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)

As shown before, the data pipeline uses collectors and message buses (RabbitMQ (OSG) and ActiveMQ (CERN)) to gather and send data to two different Elasticsearch instances (University of Chicago analytics platform and University of Nebraska)

This data provides 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=()

Platform Use and Associated Analytics

- **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
- Close collaboration with perfSONAR consortium
- 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

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

Current Platform Use

Grafana - perfSONAR dashboard

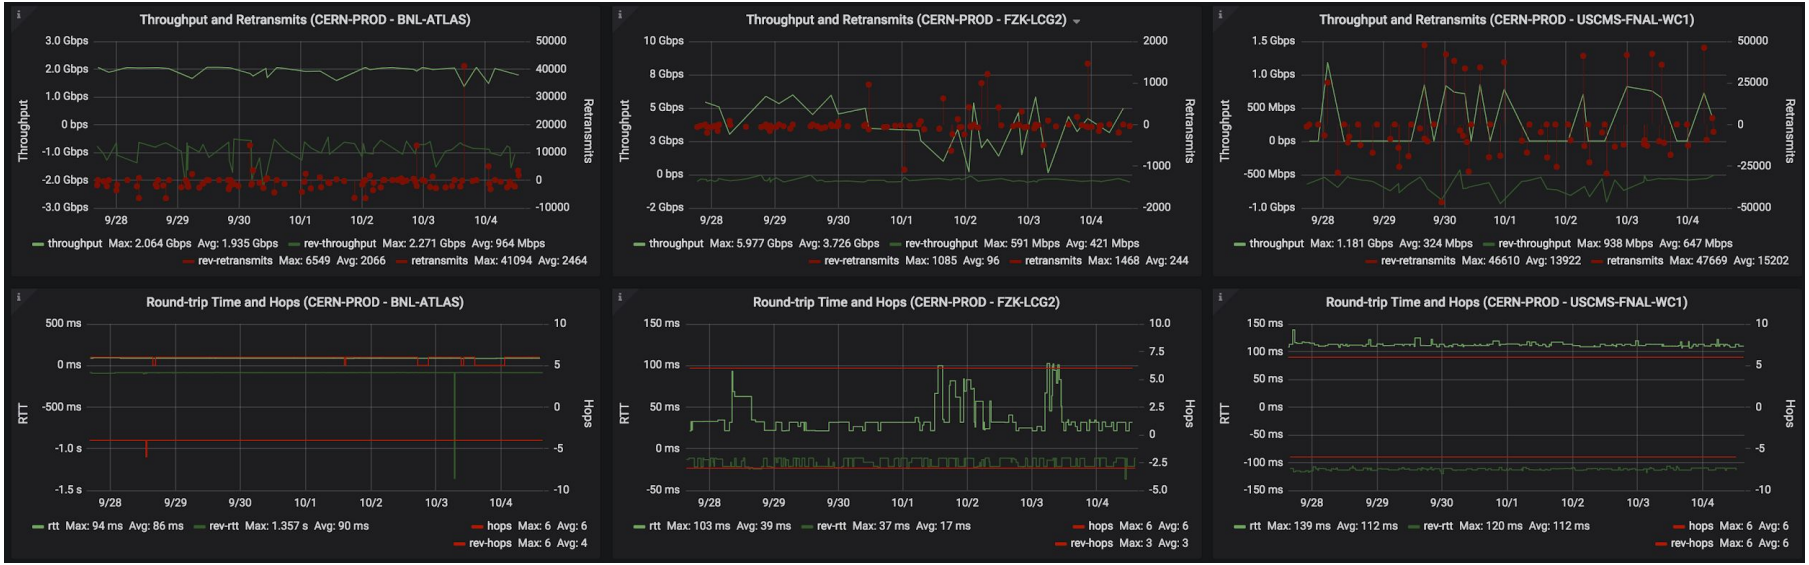

- Now includes all WLCG sites that run perfSONAR
  - Additional work needed to better filter production nodes
- Added additional row that tracks RTT and number of hops as reported by traceroute/tracepath
- Can you spot the network issue(s) above ?

Impact So Far

Future Work

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

Conclusions

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.

Conclusion and Acknowledgements

The SAND project continues to develop and maintain a network data pipeline and analytics platform to better understand and troubleshoot our networks.

Questions?

This material is based upon work supported by the **National Science Foundation** under Grants No. **1827116 and 1836650**.

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

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!