The Service Analysis and Network Diagnostic (SAND) Project

MORGRIDGE
INSTITUTE FOR RESEARCH
CORE COMPUTATION

SAND

FEARLESS SCIENCE

Service Analysis and Network Diagnosis

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

- **GOAL**: capitalize and expand our rich network dataset!
- Website https://sand-ci.org
- Project started in September 2018 and will last 2 years.

Brian Bockelman
Morgridge Institute

Shawn McKee
University of Michigan-Ann Arbor

Rob Gardner
University of Chicago

The "**Team**" at our F2F at U. Chicago
Picture credit: **Rob Gardner**

The Community Context

What is the context we are embedded in? Where did SAND come from?

- 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** Many issues are potentially resolvable within the working group.

Building / Refining a Data Pipeline

The first SAND priority was to overhaul the data collection pipeline to the one below.

- Collects, stores, configures and transports all network metrics
- Distributed deployment – operated across UC, Michigan, and Nebraska.

All perfSONAR metrics are available via API, live stream or directly on the analytical platforms.

- Includes additional input data, such as HTCondor file transfer metrics

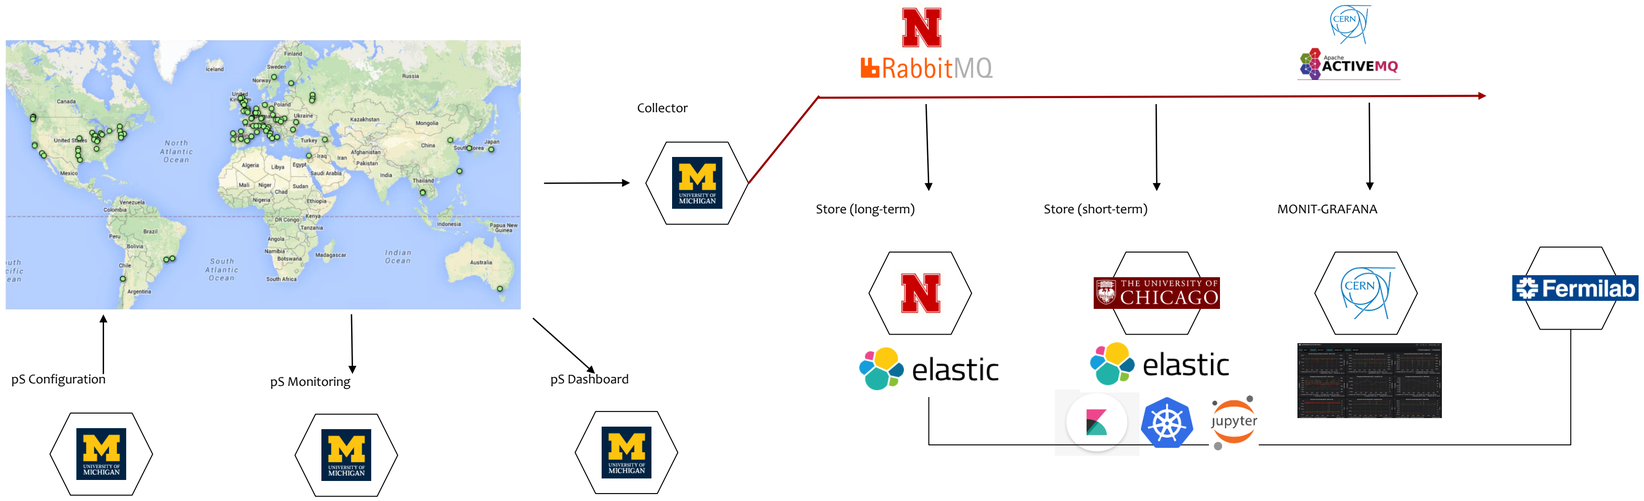

Year 1 Accomplishments

In the last year we:

- Overhauled the data collection pipeline to reduce latency and complexity / overlapping functionality.
  - Improved monitoring and alerting of the infrastructure itself.
  - Integrates the data with a long-term archive, online analytics platform, and the CERN data archive.
- Collected over 1.5B test records into the archive – and gather about 2.5M new records a day.
- Started to build simple dashboards to visualize the data and understand the collection process.
- Demonstrated the ability to switch to a “push” model instead of a “pull”.

Global Tests in the last 24 Hours	
1,903,360	Latency
13,028	Throughput
460,561	Traceroutes
2,376,949	Total
Global Tests in the last Year	
1,471,960,268	Latency
4,502,822	Throughput
143,696,469	Traceroutes
1,620,159,559	Total

Platform Use

perfSONAR - Kibana

atlas-kibana.mwt2.org/s/networking/app/kibana#/dashboard/perfSONAR?_g=(filters:!(),refreshInterval:(pause:!t,val...

Bookmarks | Apple | News | CMS Global Pool | CmsXrootdArchite... | Xrootd Nagios | Building a CMSS...

Dashboard / perfSONAR

source

AGLT2 ×

destination

MWT2 ×

Apply changes | Cancel changes | Clear form

packet loss measurements

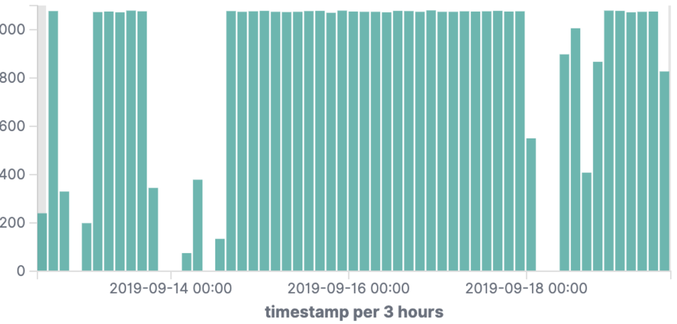

pS - packet loss average

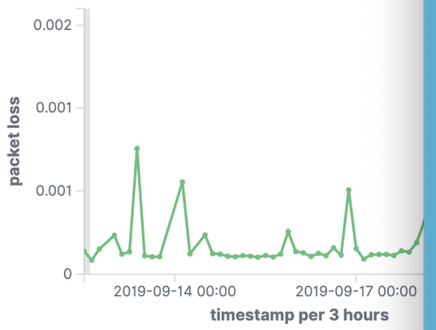

owd measurements

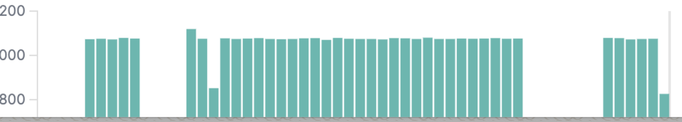

pS - owd average

HTCondor Transfers - Kibana

gracc.opensciencegrid.org/kibana/app/kibana#/dashboard/AWC4YIw_ZzwVQT7jgpwE?_g=(refreshInterval:(display:...

Apps | Bookmarks | Apple | News | CMS Global Pool | CmsXrootdArchite... | Xrootd Nagios | Building a CMSS... | Other Boo

Average Retransmissions by source and destination

Source	Destination	Unique IPs	Average Retransmissions	Average Rordering	Sum of bytes
login.duke.ci-connect.net	colorado.edu	1	1380.5	83.7	5.332GB
login.uscms.org	rwth-aachen.de	2	928	3	77.518MB
login.duke.ci-connect.net	syr.edu	4	567.111	13.778	2.421GB
xd-login.opensciencegrid.org	syr.edu	3	548.222	34.111	3.973GB
login.uscms.org	ac.be	2	481.75	83.583	432.231MB
login03.osgconnect.net	cancercomputer.com	1	243.667	3	395.117MB
login.uscms.org	ac.uk	15	216.182	45.091	938.814MB
login02.osgconnect.net	fsu.edu	2	211.667	85.667	300.061MB
login.uscms.org	ultralight.org	1	140.2	32.6	406.016MB
login.uscms.org	infn.it	12	114.182	49.606	482.842MB

Export: Raw | Formatted

1 2 3 4 5 ...13

HTCondor Ratio of Retransmissions to Gigabytes

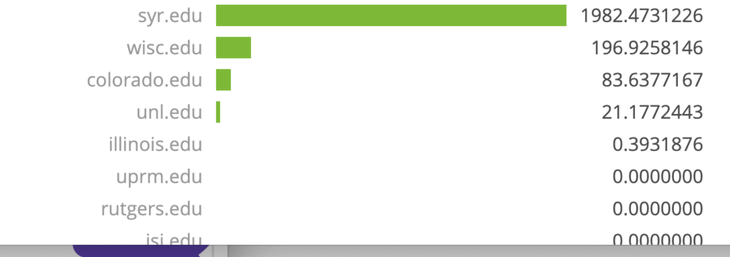

HTCondor Retransmissions and Bytes by Destination

Submit Host	Destination	Total Retransmissions	Total Bytes
login02.osgconnect.net	syr.edu	62706	17.21
login02.osgconnect.net	wisc.edu	2258	5.274
login02.osgconnect.net	gsu.edu	2207	342.1
login02.osgconnect.net	colorado.edu	1446	902.9

SAND

MORGRIDGE INSTITUTE FOR RESEARCH

Students Working on 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 approaches to understand the dataset.
- At Michigan we have **Manjari Trivedi** (Undergraduate) and **Yuan Li** (recent Master's graduate UM School of Information) focusing on path analysis.

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** 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

Where to now? Making the platform more useful

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.
- Finish transition from a “pull” data model to a “push” model.

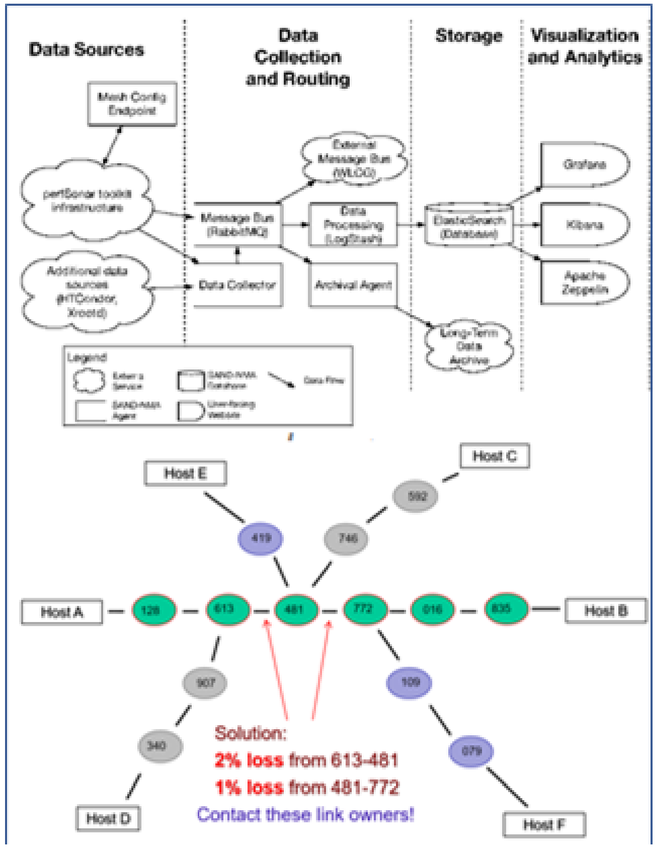

Where to now? New Data Sources

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 (see 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

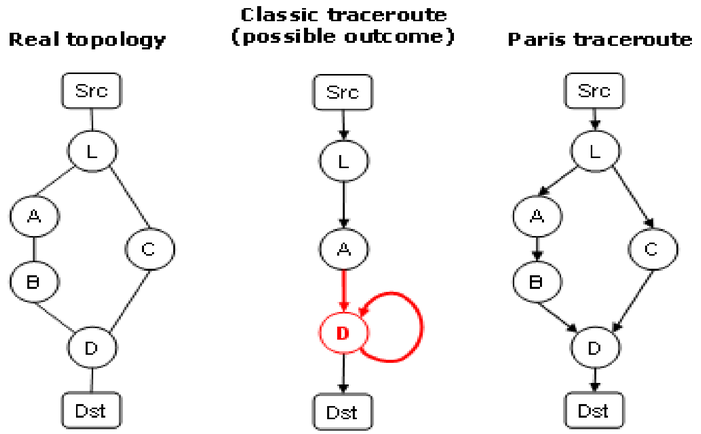

SAND and CC* Sites – We need your help!

We would like to build a map of all the NSF CC* recipients utilizing perfSonar to monitor their networks.

- Would like to ask everyone here who identifies as part of the CC* team – and using perfSonar – to follow the directions to the right.
- In a follow-up, we'll ask if you want to join the open data collection and contribute your endpoint's data to our dataset.

https://sand-ci.org/join-community/

SAND and CC* Sites – Snapshot of prototype display!

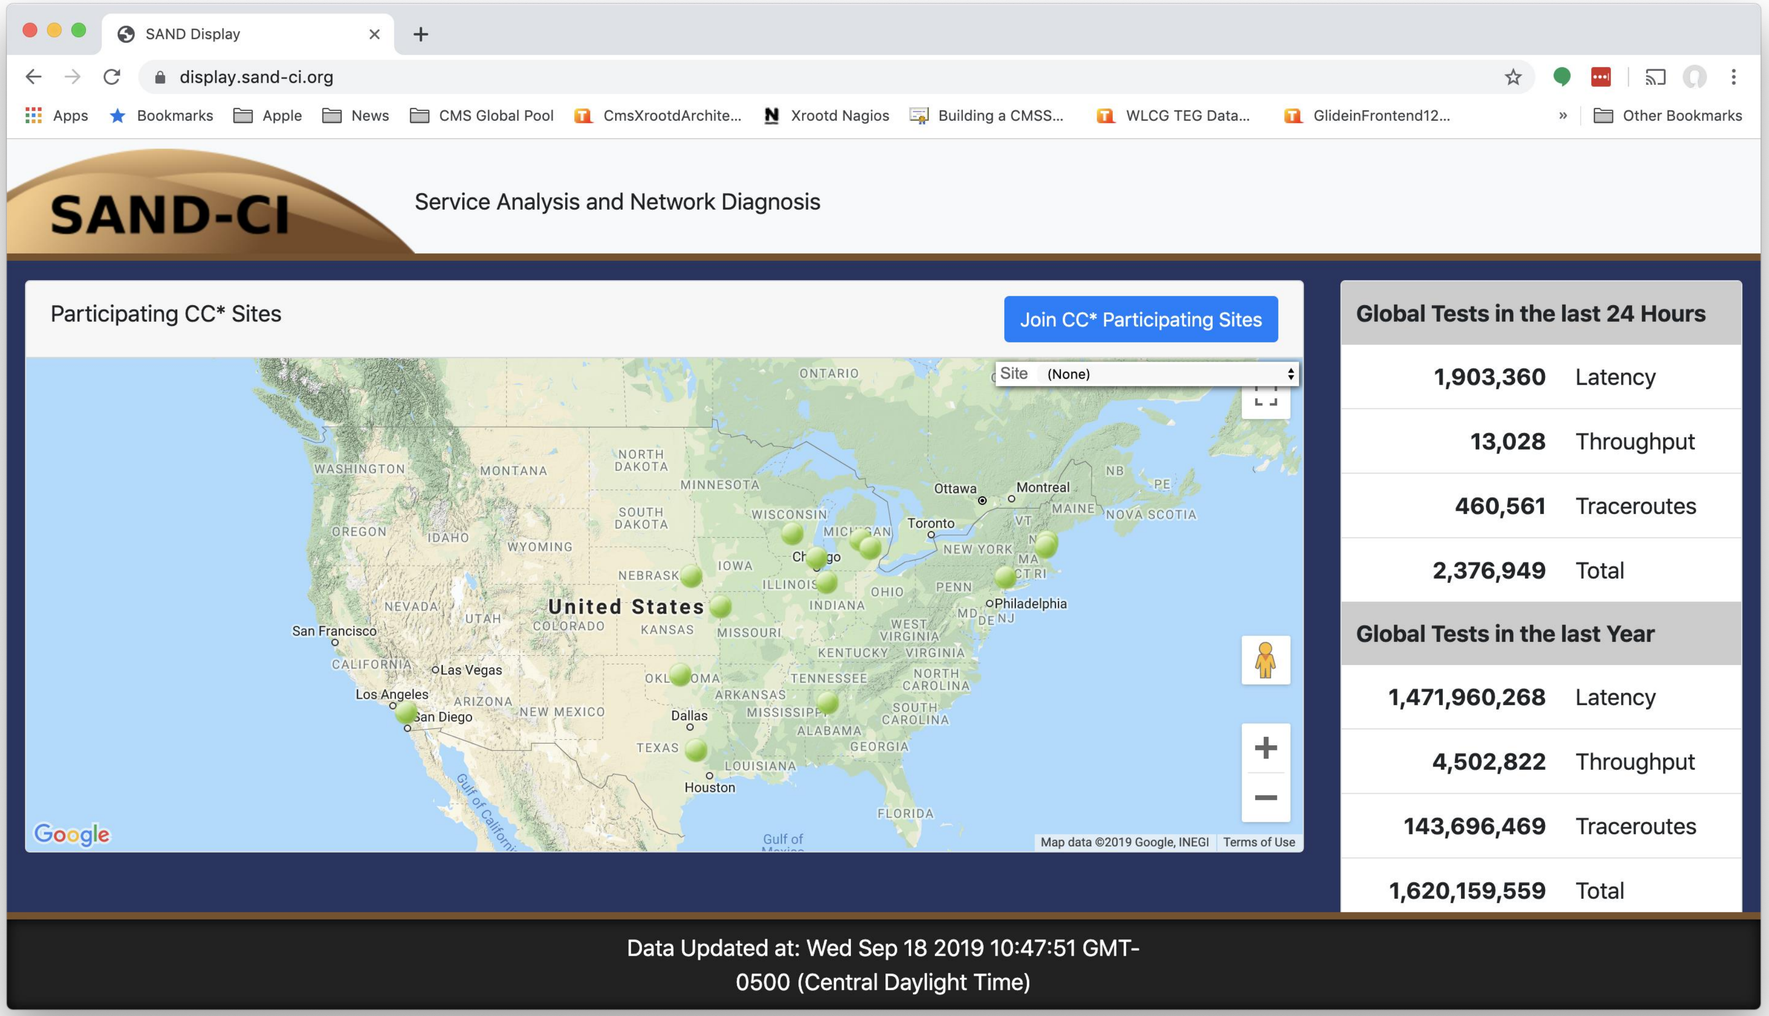

View the in-progress visualization at https://display.sand-ci.org!

MORGRIDGE

INSTITUTE FOR RESEARCH

CORE COMPUTATION

morgridge.org

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.

FEARLESS SCIENCE