



OSG/WLCG Networking Update

Marian Babik, Shawn McKee



Outline

- Activity and News
- Collaborating Projects
- Platform and Services Updates/Additions
- Platform Use
- Plans
- Summary




OSG/WLCG Networking Activities

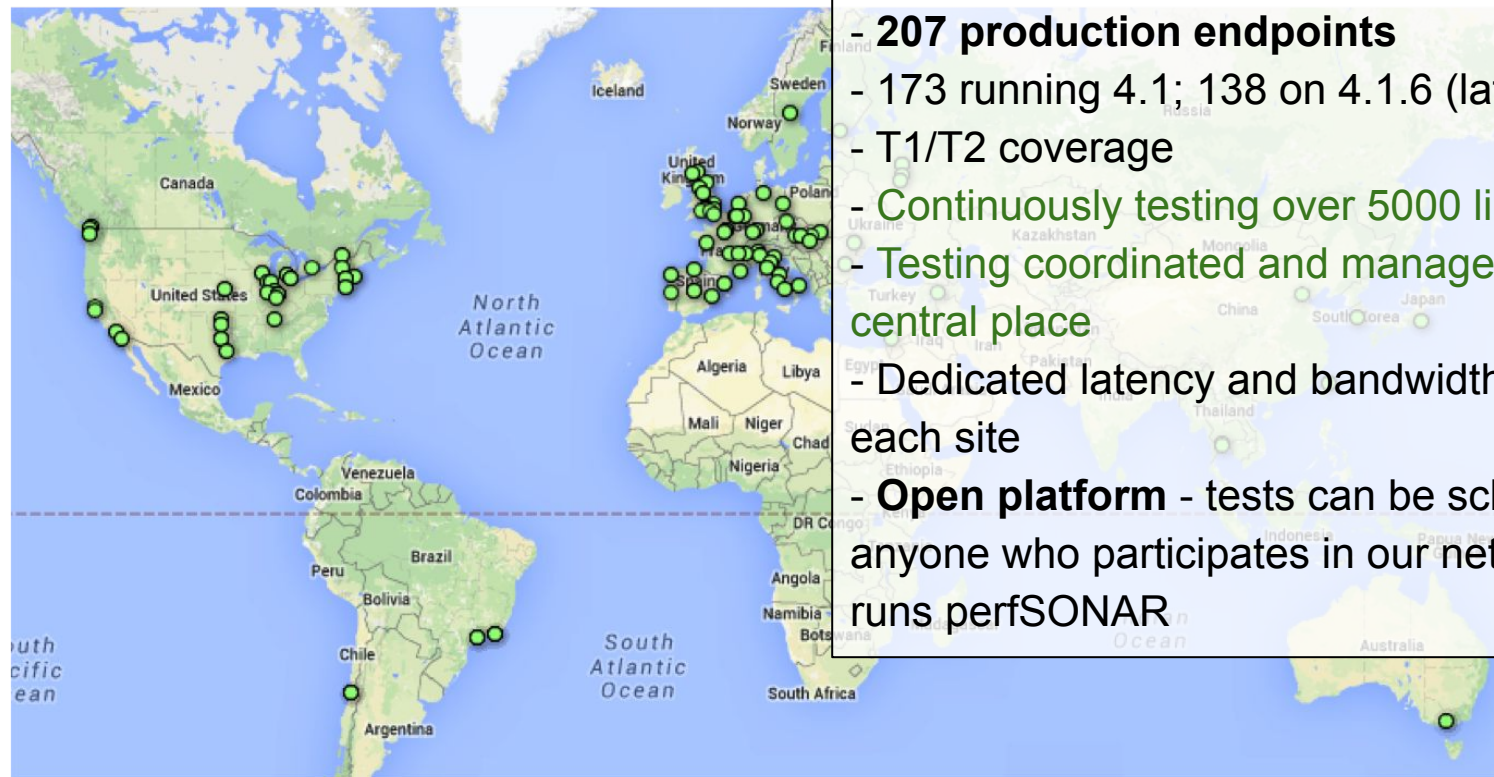
- OSG has entered its 7th year of supporting WLCG/OSG networking:
 - Assisting its users and affiliates in identifying and fixing network bottlenecks
 - **Developing and operating a comprehensive Network Monitoring Platform**
 - Improving our ability to manage and use network topology and network metrics 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 is 4.1 was released at the end of August last year

- **New plugins**
 - Network traffic capture (via 'snmp')
 - TWAMP (two-way active measurement protocol) - more accurate round trip measurements than the ones from ping, can test devices not running perfSONAR
- **New configuration**
 - PWA/PSCONFIG - **new central web interface and toolkit configuration mechanism**
 - Brings a lot more options and better use of pScheduler
- **pScheduler adds preemptive scheduling support**
 - **Retires BWCTL** - still installed but no longer configured
 - pScheduler requires port 443 to be open to all (potential) testing nodes
 - Latest version **4.1.6** fixes a known issue with duplicated testing - please update ASAP
- **Docker support (for “testpoint” deployment)**
- **Drops SL6 support which is the OS for most of our instances**
 - Our recommendation: **reinstall** with CentOS7; **don't worry about saving data**



perfSONAR deployment



261 Active perfSONAR instances

- **207 production endpoints**

- 173 running 4.1; 138 on 4.1.6 (latest)

- 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


Campaign to Upgrade perfSONAR

- We have recently begun a campaign encouraging sites to upgrade their perfSONAR deployments, both **hardware** and **software**
 - Many sites deployed their perfSONAR systems >5 years ago and the hardware is often just at the minimum (or even below) what is required to run the tests we need
 - With perfSONAR 4.1, all sites running CentOS 6.x need to reinstall using CentOS 7.x since perfSONAR no longer support CentOS 6.x
- It is possible to get robust reliable network metrics using perfSONAR 4.1+ reasonable hardware.
 - Duncan Rand has really helped get the UK sites in shape:

UK Mesh Config - UK IPv4 Bandwidth - Throughput




⚠ Found a total of 4 problems involving 3 hosts in the grid




UK Mesh Config - UK IPv4 Traceroute - Path Count




⚠ Found a total of 1 problem involving 1 host in the grid



UK Mesh Config - UK IPv4 Latency - Loss



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

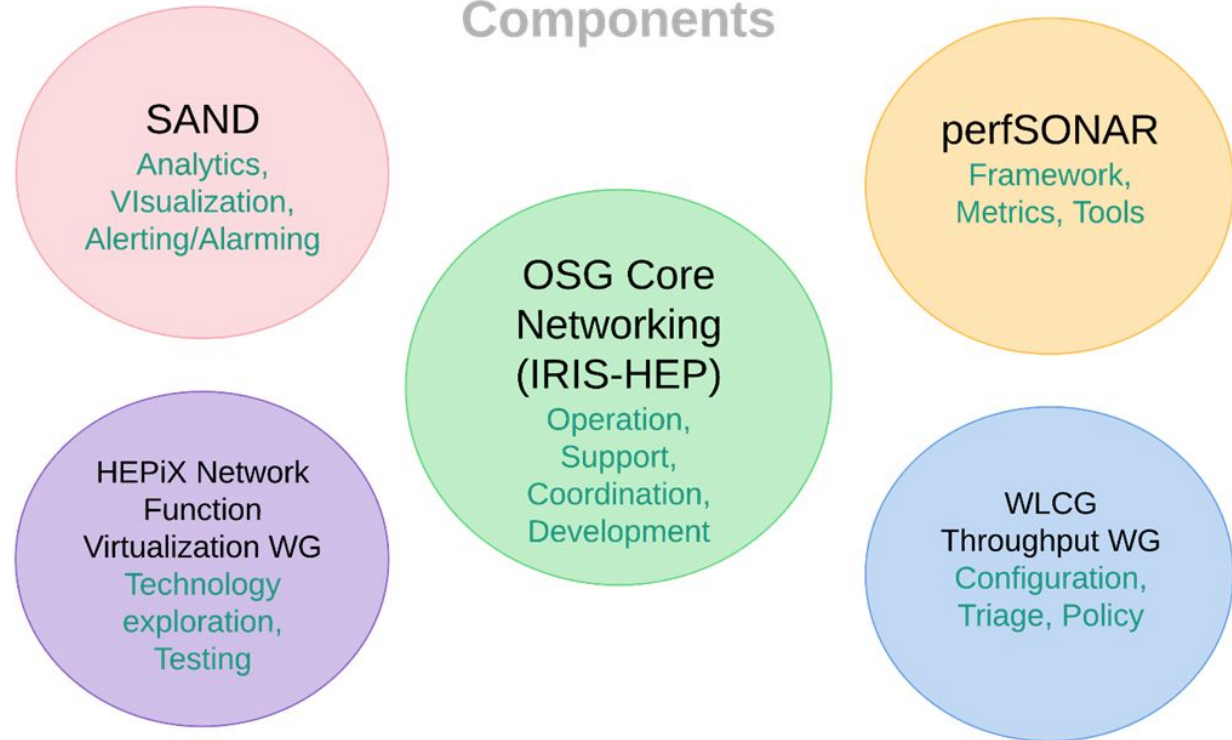


Networking Projects

There are now 4 coupled projects around the core **OSG Network Area**

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

OSG Networking Components



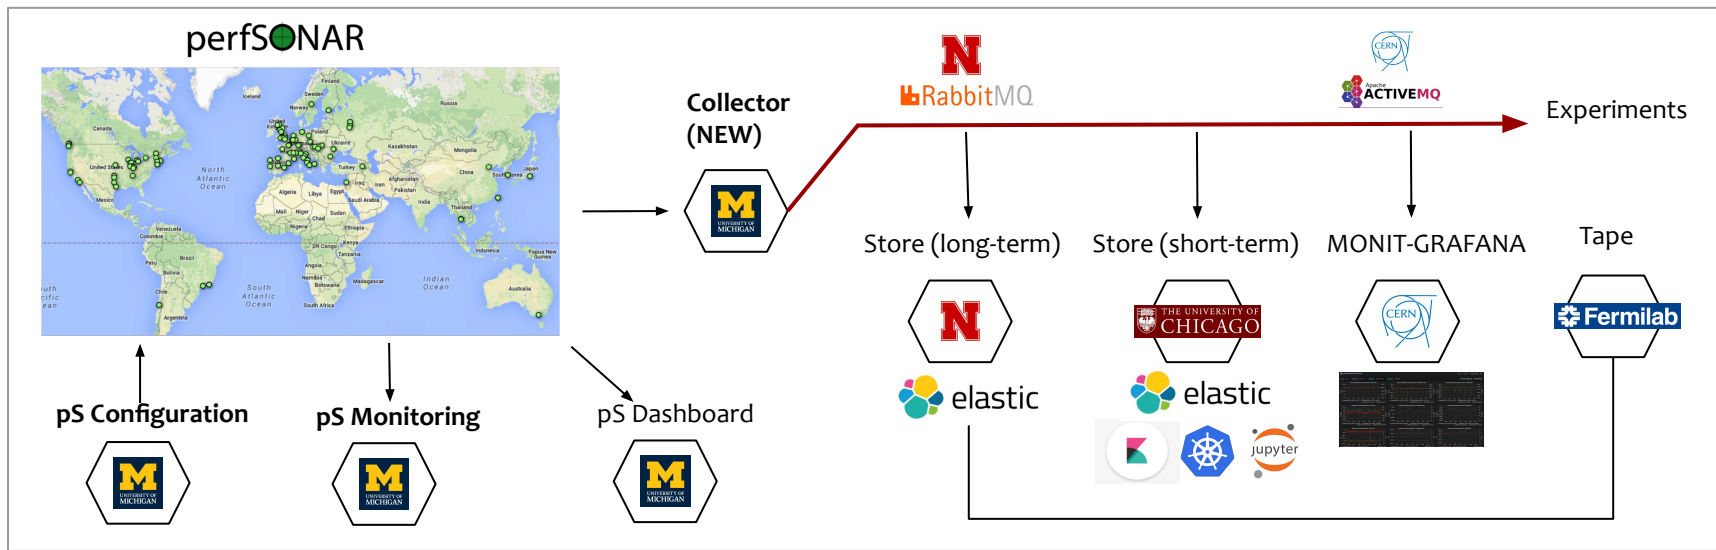
'New' Collaborating Projects: IRIS-HEP and SAND

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. This institute funds the **LHC part of Open Science Grid, including the networking area** and is creating 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 **Service Analysis and Network Diagnosis (SAND)** project, funded by the National Science Foundation as well (**CC* INTEGRATION 1827116**). SAND started September 1, 2018. The SAND project coordinates network monitoring and diagnostics from over a hundred application-level measurement points and aggregates them into a measurement archive. This curated measurement archive provides access to historical and quasi-real-time network performance data, allowing for higher-level diagnostics and analyses. For more info see <https://sand-ci.org/>

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



Recent Pipeline Updates

- As part of SAND, a **new collector** gathering perfSONAR measurements from all toolkits worldwide has been implemented and put in production
 - Significantly reduces latency and improves performance
 - Grafana monitoring at <https://gracc.opensciencegrid.org/dashboard/db/perfsonar-collector?orgId=1>
 - Improved alerting and notifications for the pipeline operations
- **Streaming** now available on both **RabbitMQ** and **ActiveMQ**
- **New central configuration** put in production (**PWA**)
 - All meshes updated to test throughput and traces over both **IPv4** and **IPv6**
- **Monitoring** updates
 - Support for sites and host groups/mesh notifications
 - **New target version now 4.1.6** - this is motivated primarily to check if sites have enabled auto-updates - **very important for security and bug-fixing**

New Toolkit Info Web Page

At a prior OSG All-hands there was a request to provide a “front-end” web page the could help toolkit owners in managing and fully utilizing the various resources and services OSG provides.

We now have a **prototype** running that we plan to evolve based upon feedback:

<https://toolkitinfo.opensciencegrid.org/>

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: **lhcmon.bnl.gov**

Customized Web links for **lhcmon.bnl.gov**

- [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... 1 row /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=mc

Availability

Local site etf

| state | Host | Icons | OK | Wa | Un | Cr | Pd |
|-------|----------------|-------|----|----|----|----|----|
| UP | lhcmon.bnl.gov | | 15 | 0 | 0 | 0 | 0 |

refresh: 60 secs

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

NST

Toolkit Info Web Page (2)

You can select any of the currently registered perfSONAR toolkit instances to get a set of customized links specific to that instance.

If you know part of the DNS name, you can start typing in the box to narrow the selection list.

The screenshot shows the 'The perfSONAR Toolkit Information Page' on the Open Science Grid. A search box labeled 'Select toolkit:' contains the text 'perf'. A dropdown menu is open, listing various toolkit instances such as 'ccperfsnar1.in2p3.fr', 'clrperf-bwctl.in2p3.fr', and 'lhcmon.bnl.gov'. Below the search box, there are links for 'Customized Web' and 'Monitoring of th...'. On the right side of the page, there is a table showing the status of the 'lhcmon.bnl.gov' instance. The table has columns for 'state', 'Host', 'Icons', 'OK', 'Wa', 'Un', 'Cr', and 'Pd'. The 'state' is 'UP' and the 'OK' value is '15'. Below the table, there is a 'refresh: 60 secs' button. The page also features logos for 'Open Science Grid' and 'WLCG Worldwide LHC Computing Grid'.

The perfSONAR Toolkit Information Page

Select toolkit: perf Submit

OSG Network Services Analytics and Dashboards

perfSONAR Toolkit is: **lhcmon.bnl.gov**

Host sea... 1 row /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=mc

Availability 1 60s

| state | Host | Icons | OK | Wa | Un | Cr | Pd |
|-------|----------------|-------|----|----|----|----|----|
| UP | lhcmon.bnl.gov | | 15 | 0 | 0 | 0 | 0 |

refresh: 60 secs

NSF Search & Innovation High Energy Physics

Toolkit Info Web Page (3)

There are additional menus setup to provide one-stop shopping to relevant services, documentation and dashboards.

We are also implementing “hover-over” text boxes to help describe the various links.

Please try it out and email me with your feedback!

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**

perfSONAR Collector Mc Packet Loss Alarm OSG Networking OSG/WLCG MaDDash WLCG Grafana Network Dashboards

perfSONAR Ingest R Packet Loss Alarm perfSONAR Toolk OSG/WLCG pSConfig (P OSG Analytics Platform

Status of ITB Servic Packet Loss BarG ESnet FasterData Home (Network/system Tu OSG perfSONAR Infrastructure details on a Kibana Dashboard

Status of Production Se Packet Loss BarG ESnet Network Troubleshooting Guide

refresh: 60 secs

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

NSF

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
- Currently working on adding **perfSONAR** as part of standard benchmarking activities

- 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



Network Operations


- Sites experiencing network issues should **first** contact their local network team or directly their regional and backbone (R&E) providers
- A group focusing on helping sites and experiments with network performance using perfSONAR - **WLCG Network Throughput**
 - Please include as many details as possible (include any existing tickets with R&Es)
 - For list of existing and recently resolved issues see this [link](#)
- **LHCONE operations** - support for establishing and operating LHCONE infrastructure - regular meetings and support mailing list
- **LHCOPN/LHCONE community** - organised bi-annually - good place to meet R&Es and discuss architecture and plans
 - Next one is in Umea, Sweden, <https://indico.cern.ch/event/772031/>)

Operations

Architecture,
Infrastructure

perfSONAR near-term releases

- **perfSONAR 4.2** (Q1 (Q2?) 2019)
 - **GridFTP plug-in** - Significant interest from NRP community and others.
 - Measurement pre-emption - Easier for diagnostic tests to get a slot on busy hosts
 - **Additional pSConfig utilities** - Continuing to make meshes easier to build and manage through command-line and graphical interface
 - Lookup Service improvements - Bulk renewals and record signing
- **perfSONAR 4.3** (Q3 2019)
 - User Interface and Visualization Strategy - Seek to improve user experience and operational efficiency within development team by consolidating code
 - **pScheduler Resource Pooling** - Better management of resources like ports, potential gains in environments like Kubernetes where ports may be constrained
 - **Esmond Updates** - Option to run using pure postgresql (no cassandra)



Reminder about perfSONAR Deployment

We need network visibility to understand performance, find problems and enable orchestration


- **All sites should have deployed perfSONAR and have a plan to keep the hardware and software updated**
 - The recommendation is to provide two instances: **latency** and **throughput** (which could be the same server with at least two NICs)
 - The perfSONAR instances should be **(co)located** with your sites STORAGE, network-wise
 - The **throughput** instance should use the same NIC capacity as your storage servers
 - Additional perfSONAR instances can be helpful for identifying LAN issues
- <https://opensciencegrid.org/networking/perfsonar/installation/#perfsonar-installation-guide>

WG near-term Plans

- Complete campaign to update perfSONARs to **CC7 and latest release**
 - We have already contacted all T1s and are moving on to T2s
- Re-organise **LHCONE** mesh
 - Create a better fit for the underlying infrastructure
 - Add uni-directional testing to perfSONARs hosted by R&Es on LHCONE
- **100G deployments**
 - CERN now running 40Gbps and plans to also add 100Gbps support
 - SARA now running 100Gbps, BNL running 80Gbps
- Working with the data in **Elasticsearch** to **correlate** and **visualize** **traceroute paths with their related network metrics (packet-loss, latency, bandwidth)**
- We will be working closely with the **SAND** (<https://sand-ci.org/>) project to:
 - Improve the **robustness** and **efficiency** of the data pipeline
 - Create new analytics capabilities
 - Tune-up the **alerting components** that users can subscribe to

Summary

- OSG in collaboration with WLCG are operating a comprehensive network monitoring platform
- Platform has been used in a wide range of activities from core OSG/WLCG operations to Cloud testing and IPv6 deployment
- Providing feedback to LHCOPN/LHCONE, HEPiX, WLCG and OSG communities
- Next version of perfSONAR will enable additional functionality as well as improve overall stability and performance
- IRIS-HEP and SAND started and will contribute to the R&D in the network area
- Further analytical studies are planned to better understand our use of networks and how it could be improved
- More on networking technology in Rolf's tech watch talk Thursday



References

- **New toolkit info page:** <https://toolkitinfo.opensciencegrid.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/>

