

# WLCG Experiments Test Framework (ETF)


Marian Babik, CERN  
HEPiX IPv6 WG meeting



# Motivation

## Current Service Availability Monitoring (SAM) structure

- **WLCG Experiments Test Framework (ETF)**
  - WLCG testing middleware (running so called SAM tests)
  - Active testing of the sites services and reporting back to SAM3/MONIT
  - Common to all experiments
  - Main source for WLCG Availability/Reliability Reports (different from EGI/ARGO)
- **SAM3/MONIT**
  - Aggregation (via custom algos), visualisation and reporting
  - Support for multiple sources of metrics (e.g. ALICE storage tests, ATLAS ASAP)
- **A generic test framework remains fundamental for WLCG monitoring**
  - Keeping track of sites availability/reliability
  - Running deployment campaigns (IPv6, HTTP, etc.)
  - Provides means of isolation when debugging site/experiments issues
    - Middleware bugs, site setup/configuration, latency sensitive issues/timeouts, etc.
  - Contributing to the operational toolchain of the experiments



# Overview

Generic test middleware based on open source

- Checkmk, Nagios core and Messaging (ActiveMQ)

Focuses on functional testing (atomic)

- Direct job submissions, worker node env. testing
- Core storage operations
- Remote API testing and/or network testing (ping/icmp)

~ 150 sites, 1200 hosts monitored

~ 10 metrics/host

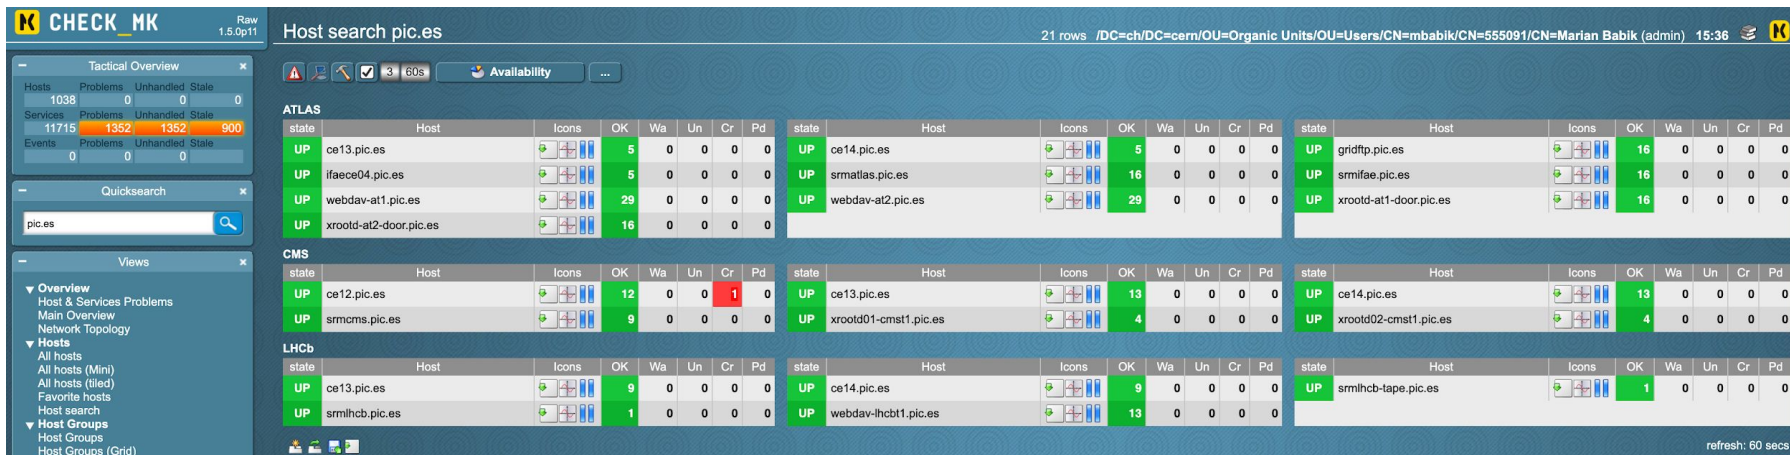
~ 1M metrics/day

High-level functional testing

Plugins conforming to Nagios standard

Configurable schedule for test execution

Checkmk dashboard to show results



# Architecture

## ETF Core Framework

- Frontend API, configuration, scheduling, alerts

## Plugins (probes/tests)

- Range of available plugins to test broad range of services
- Contributed by experiments, PTs, TFs and open source projects (Checkmk), etc.
- Python library to help write plugins (**python-nap**)

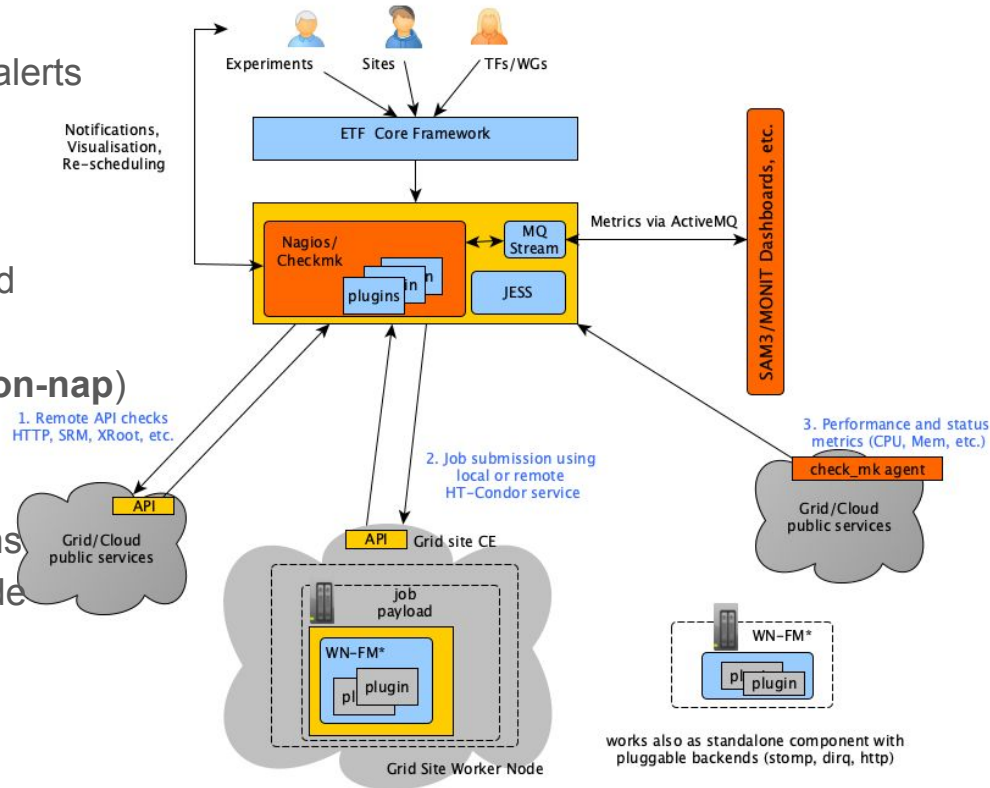
## MQ Stream for publishing results

## Job Submission Framework (JESS)

- Framework to write job submission plugins (submit/manage jobs, retrieve worker node results, etc.)

## Worker Node Framework (WN-FM)

- Micro-scheduler to run tests on the WNs (configure and execute WN tests, collect results)




# Deployment and Operations

Experiment instances @CERN (IPv4-only/IPv6-only in QA, IPv4-only in PROD)  
perfSONAR infrastructure monitoring @OSG

## ETF now runs in containers and is integrated with gitlab CI

- Each experiment has its own container/image and gitlab repository
  - Full control over packages and versions to be deployed
- ETF can be deployed in the experiment-specific environment if needed
- Faster development cycle - changes propagated to QA upon each commit
  - Each commit triggers container rebuild and deployment to QA, one-click deploy to prod
- Simplified deployment - auto-deployed directly from gitlab
  - Easy to rollback



# Plugins/Tests

Plugins	Users/Experiments	Maintained by
<b>Job Submission</b>		
CREAM, ARC, HTCONDOR-CE <i>JESS**</i>	LHCb, ALICE, ATLAS, CMS	ETF
<b>Worker Nodes</b>		
ATLAS (3), CMS (11), LHCb (7)	ATLAS, CMS, LHCb	ATLAS, CMS, LHCb
<b>Storage</b>		
GFAL2 (SRM, gsiftp, XRoot, HTTP)	ATLAS	ATLAS
GFAL2 (SRM)	CMS	CMS
XRoot**	CMS	CMS
HTTPs/WebDAV**	HTTP TF*	HTTP TF*
<b>Network</b>		
perfSONAR infrastructure**	WLCG Network Throughput WG	OSG, WLCG

\*\*Uses **new library for writing plugins** ([python-nap](#)) \*Probe is still supported by GFAL2 team

# Summary

- ETF is a container-based application combining open source software with a set of frameworks and APIs to provide flexible testing suite
- Easy to extend, re-locate and support new experiments and technologies
- Supported as part of the CERN IT Monitoring Stack
- Currently deployed at CERN for all four experiments
  - Supporting IPv4-only and IPv6-only monitoring
  - Experiments contacts have access in case they need to debug and/or follow up on issues
  - Central instance provides a site-level view (one place to see results from all experiments)
- Additional deployment at OSG for perfSONAR infrastructure monitoring
  - Strong interest from other communities to have this available as a generic tool
- SAM reporting ready for IPv6
  - We're able to aggregate IPv6 results and compute either IPv6-only profiles or IPv6/IPv4
  - Experiments need to add IPv6 metrics to the A/R profile

# Questions ?

Docs: <https://etf.cern.ch/docs/latest/>

Central instance: [https://etf.cern.ch/etf/check\\_mk/](https://etf.cern.ch/etf/check_mk/)

Instances (access requires IGTF/x509 cert loaded in the browser):

<a href="#">CMS production</a>	<a href="#">CMS QA IPv6</a>	<a href="#">CMS QA</a>	Code: <a href="#">CMS gitlab</a>
<a href="#">ATLAS production</a>	<a href="#">ATLAS QA IPv6</a>	<a href="#">ATLAS QA</a>	Code: <a href="#">ATLAS gitlab</a>
<a href="#">LHCb production</a>	<a href="#">LHCb QA IPv6</a>	<a href="#">LHCb QA</a>	Code: <a href="#">LHCb gitlab</a>
<a href="#">ALICE production</a>		<a href="#">ALICE QA</a>	Code: <a href="#">ALICE gitlab</a>
<a href="#">pS production</a>		<a href="#">pS QA</a>	Code: <a href="#">pS gitlab</a>

ETF framework

[ETF core containers](#) [ETF Job Submission \(Jess\)](#)

[ETF nagios plugins lib. NAP](#) [ETF rule-based configuration \(ncgx\)](#)

ETF support channels: GGUS: Grid Monitoring or [etf-support@cern.ch](mailto:etf-support@cern.ch) (SNOW)



# perfSONAR Update

S. McKee<sup>1</sup>, B. Bockelman<sup>2</sup>, R. Gardner<sup>3</sup>,

I. Vukotic<sup>3</sup>, M. Babik<sup>4</sup>, D. Weitzel<sup>5</sup>, M. Zvada<sup>5</sup>, E. F. Hernandez<sup>6</sup>,

<sup>1</sup> University of Michigan, <sup>2</sup> Morgridge Institute of Research, <sup>3</sup> University of Chicago, <sup>4</sup> CERN,

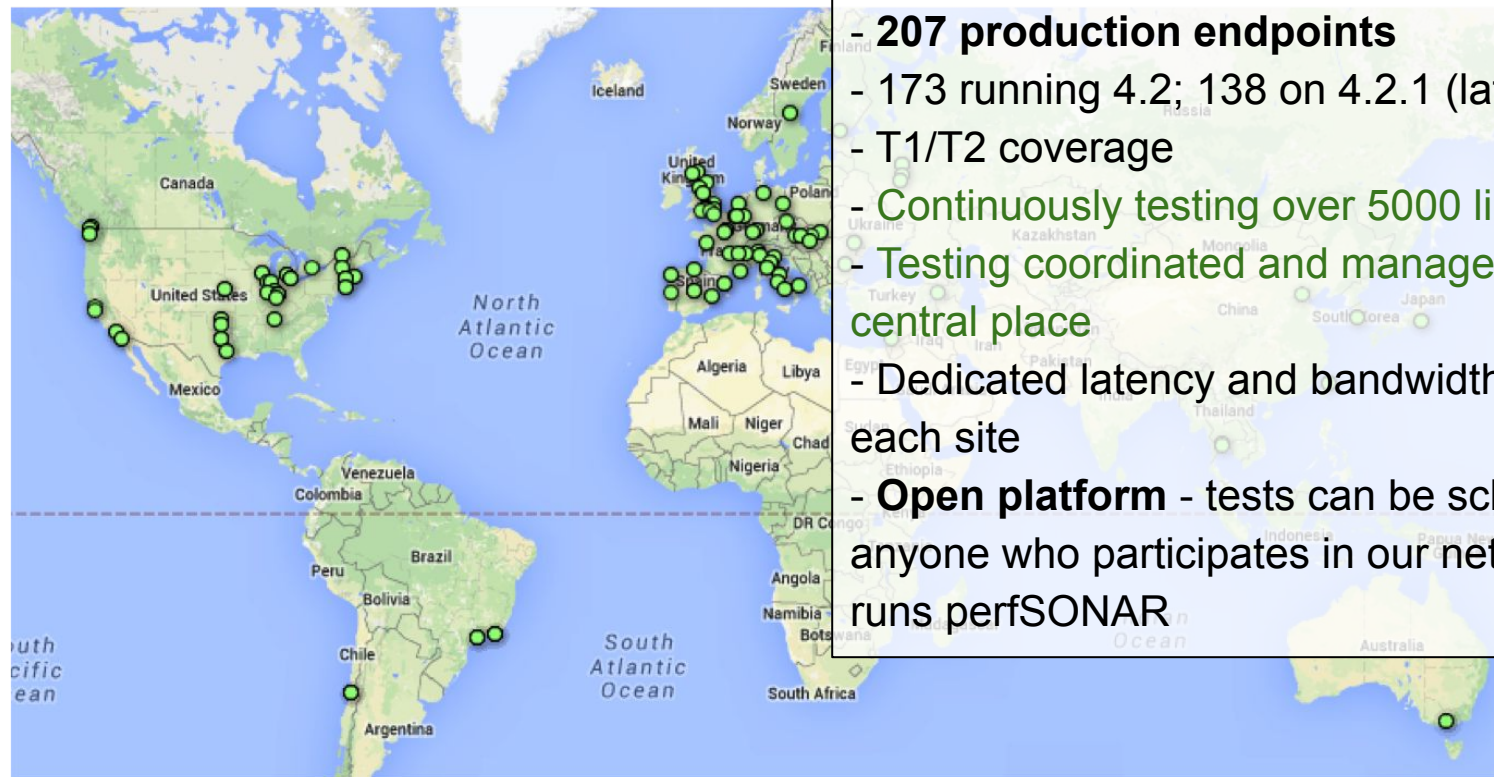
<sup>5</sup> University of Nebraska, <sup>6</sup> UCSD



perfSONAR 4.2 was released (4.2.2 is the latest release)

- **New plugins**
  - GridFTP plug-in - Significant interest from NRP community and others.
  - **Test schedule pre-emption** - Easier for manual 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
- 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
- Docker support (for “testpoint” deployment)
- **SL6 no longer supported**
  - Our recommendation: **reinstall** with CentOS7 ASAP; **don't worry about saving data**

# perfSONAR deployment

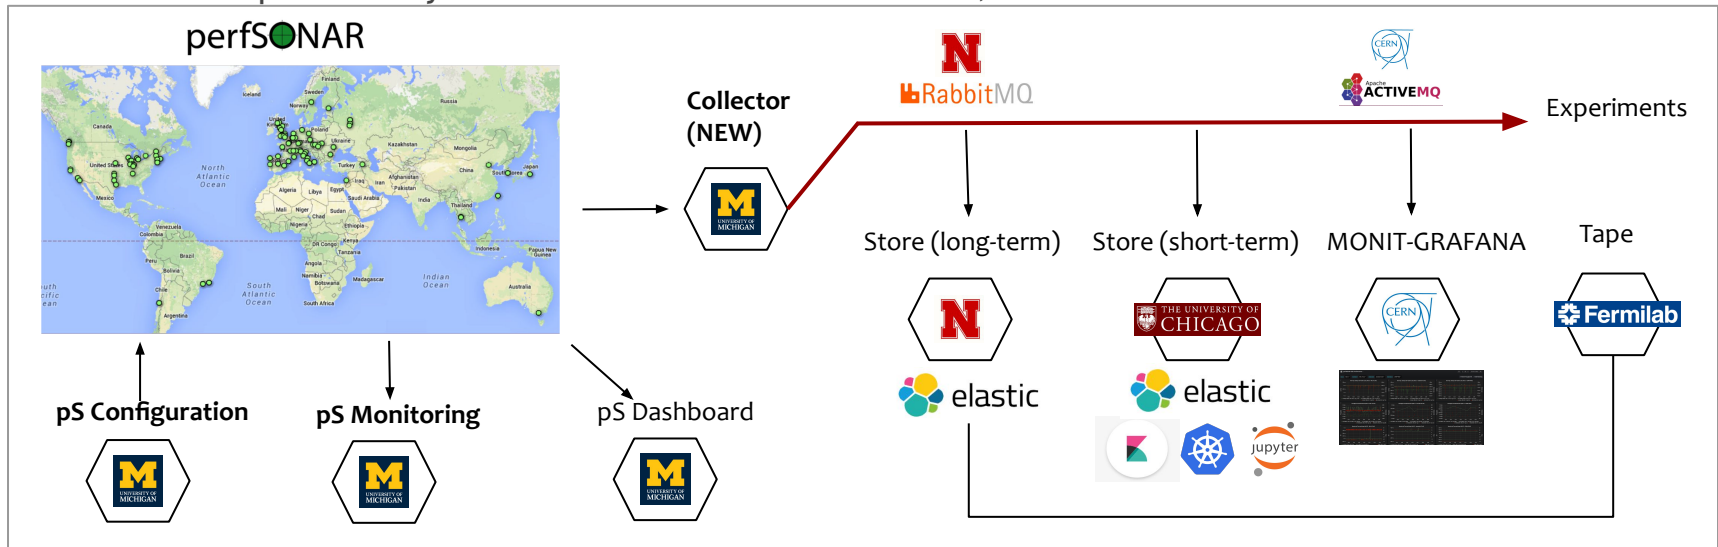


**261 Active** perfSONAR instances


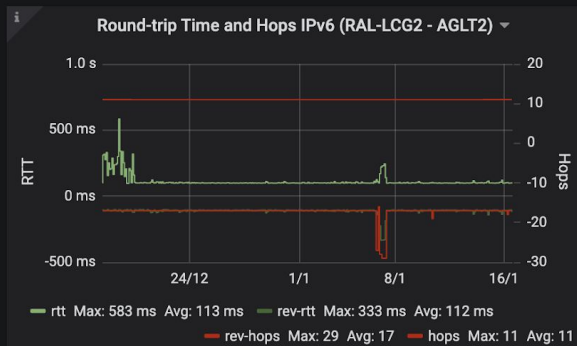




- **207 production endpoints**
- 173 running 4.2; 138 on 4.2.1 (latest)
- T1/T2 coverage
- Continuously testing over 5000 links
- Testing coordinated and managed from central place
- Dedicated latency and bandwidth nodes at each site
- **Open platform** - tests can be scheduled by anyone who participates in our network and runs perfSONAR

# Platform Overview


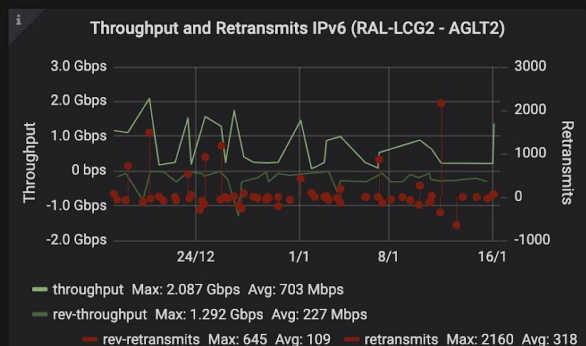




- Collects, stores, configures and transports all network metrics
  - Distributed deployment - operated in collaboration
- All perfSONAR metrics are available via **API, live stream or directly on the analytical platforms**
  - Complementary network metrics such as ESNNet, LHCOPN traffic also via same channels



# MONIT perfSONAR IPv6 dashboard



# MONIT perfSONAR IPv6 dashboard



# Network Analytics Activities

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

- **Machine learning**: At **Chicago** we have **Sushant Bansal** (Master's student)
- **Path Analysis**: At **Michigan** we have **Manjari Trivedi** (Undergraduate), **Yuan Li** (Graduate student; graduated)
- **R&E Network Analytics**: In **Bulgaria** we have **Petya Vasileva** (PhD student)

The students have worked independently over summer 2019 learning about the data we have and the analytics platform itself

For Fall 2019, the goal was 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

# Collaboration with MEPhI on Network Visualization


Containerized Version running at UC <https://perfsonar.uc.ssl-hep.org/graph/viewer>

## Network Traces Graph Visualization

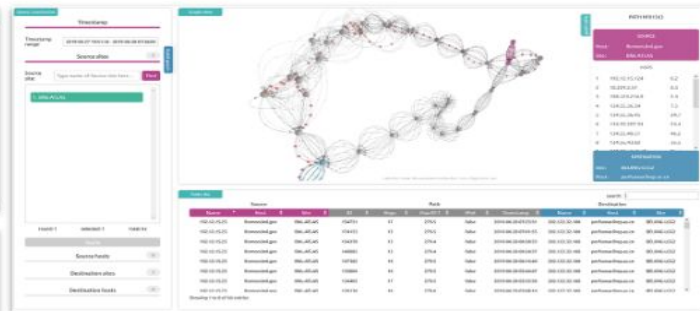

MEPhI Team

### Application being Updated


- Rebuilding **Django** server that serves as API for Elasticsearch
  - Due to Cross-Origin Resource Sharing (CORS) and security reasons
- Moving front-end part of application from **jQuery** to **ReactJS**
  - allows to replace heavy modules with light React components
- Resolving issues related to unique path Identification
  - makes possible to monitor paths but not single records about path



Source self-directed paths



Lost nodes trace





# 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

- **Analytical studies and diagnostics**

- Providing derived metrics that will make it easier to diagnose issues
- Alerting and notifications
- Advanced visualisation of network paths

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

- **perfSONAR** part of the standard benchmarking tools, developed as part of EU OCRE

- **Collaboration with GridPP and CCSTAR communities**

- Common platform for configuration, collection and storage of network measurements

- **HEPiX IPv6 WG** - now testing bandwidth and paths over IPv6

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

# Plans

- Providing new features that will **make it easier to integrate with other communities**
- EU projects ARCHIVER and ESCAPE plan to use our infrastructure
- ALICE, StashCache and CC\* communities will continue to evolve
- **100G and experimental testing** - planning 100G testbed and evaluations
  - If interested please join perfsonar-100g mailing list (subscription link is in the references)
- Working closely with the **SAND** (<https://sand-ci.org/>) project on analytics
  - SAND plans to continue hardening the data pipeline while also continuing to work on extracting value from the various network metrics being gathered
  - The students associated with the project are working on cleaning, annotating and using the data to provide useful tools to better understand our networks and localize problems.
  - The SAND project goals in its final year are to provide near real-time network problem identification, path correlation and problem localization through the use of analytics and visualization and this is well-aligned with OSG and WLCG needs.

# 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 will contribute to the operations and R&D in the network area
- Further analytical studies are planned to better understand our use of networks and how it could be improved

# References

- OSG/WLCG Networking Documentation
  - <https://opensciencegrid.github.io/networking/>
  - <https://toolkitinfo.opensciencegrid.org/>
- perfSONAR Stream Structure (API)
  - [http://software.es.net/esmond/perfsonar\\_client\\_rest.html](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](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/>

