

Network Monitoring Update

Shawn McKee, Marian Babik

Spring [LHCONE/LHCOPN \(#54\) Meeting](#), March 18, 2025

on behalf of WLCG Network Throughput WG



Overview

We have a long history of presenting monitoring results from LHCOPN and LHCONE by showing various meshes from MaDDash and tracking changes

For this presentation we will be focusing more on the ongoing and planned transitions

- We have deprecated MaDDash now (see October 2024 presentation).
- The perfSONAR deployments are in need of changes (updating hardware and implementing a new type of deployment).
- Defining and tuning our tests is going to be a focus area during this calendar year.
- New analytics and associated alerting & alarming is under very active development



- **perfSONAR 5.1.4** is the latest release (5.2 by summer?)
 - Number of bug-fixes since 5.0; bi-weekly meetings with the developers
 - Update campaign in WLCG (not really pursued yet)
 - Various issues, mostly archiving, but also e.g. legacy limits configuration ([fix](#))
 - Toolkit support for latest Alma/Rocky 8 and 9, Debian 11/12, Ubuntu 20/22
 - **CentOS7 is no longer supported, sites should update (via redeploy) ASAP**
- **We continue to have issues with resiliency and reliability with our WLCG deployment of perfSONAR toolkits.**
 - We suspect we are just trying to test to many things with some toolkits
 - The transition to deploying OpenSearch(OS) as the MA may also require more resources than our hosts are capable of delivering
 - We are exploring new recommendations we can make for those deploying perfSONAR for WLCG/LHCOPN/LHCONE




perfSONAR Infrastructure


175
Active perfSONARs

65
Communities in Use


NIC speed



perfSONAR distributions



pS NIC MTU




- 95.8% toolkits now on 5.x
 - Still about 4.2% on 4.x :(
 - 35% still on EL7 (all the 4.x and some of the 5.x)
- Core deployments are still on 10Gbps, but we have about 20% with 100Gbps
 - For WLCG/OSG testing purposes 10Gbps is still sufficient
 - We have some **200 Gbps** and **400 Gbps** hosts
 - **Important to refresh HW along with the update to EL9**
- MTU - around 36% on jumbo frames (9000), rest is on standard frames (1500)

perfSONAR Infrastructure Evolution

175
Active perfSONARs

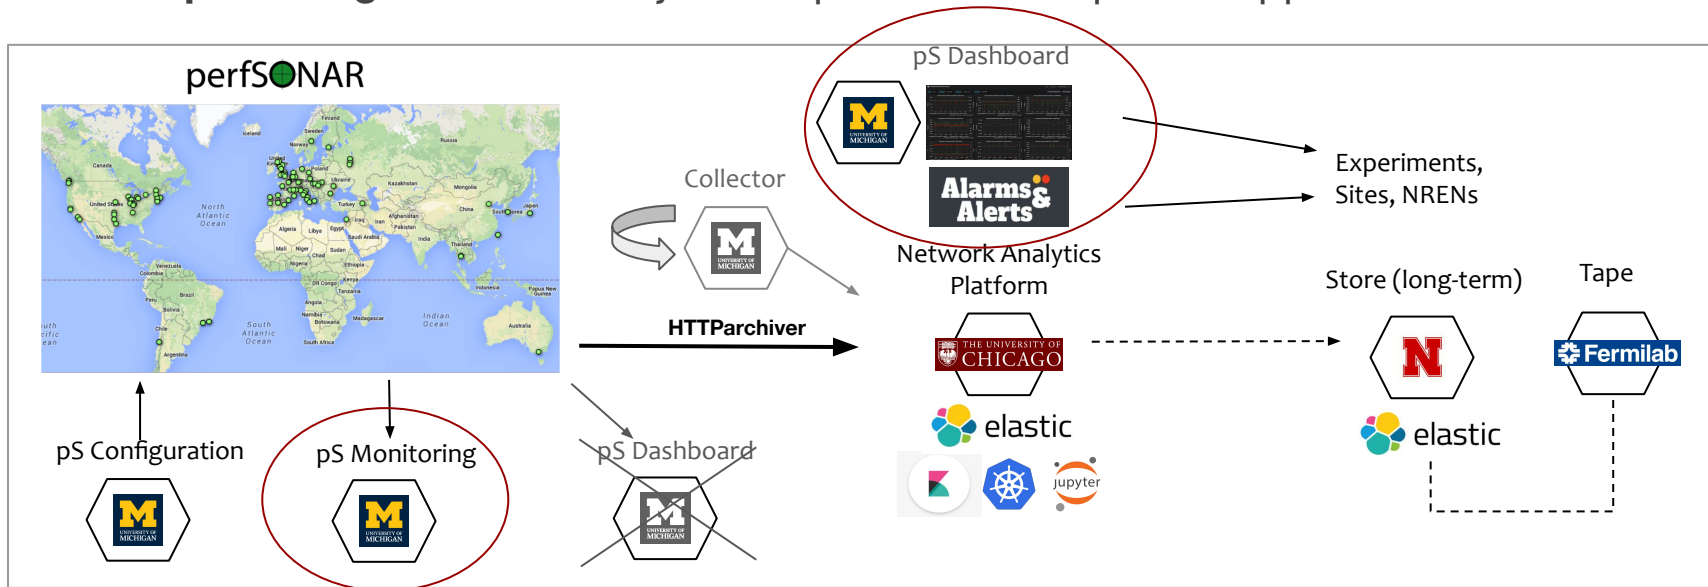
65
Communities in Use



- Long-term trend of perfSONAR releases adoption shows a steady downtrend
 - Only reachable “toolkit version” installations are shown (stacked plot)
 - Stability and reliability of the releases clearly playing a role
- New strategy and deployment models will need to be considered
 - Exploring simple deployments (without complex components such as OS)
 - Providing means to easily reset/restart deployment (to avoid accumulation of issues)
 - Easy to co-locate with storages (or even co-hosted on storages with multiple NICs)

Network Measurement Platform Plans

- Our platform/pipeline is evolving, with some components replaced or retired.
 - Forwarding to UNL and backup to FNAL still to be implemented
 - **pS Monitoring** - update to latest CheckMK and enable SSO authentication
 - **ps Dashboard** - integrate with Analytics Platform/Grafana (MaDDash retired)
 - **ps Configuration** - clarify development roadmap and support



Future perfSONAR Deployment Option

The **perfSONAR Testpoint container** allows a simplified deployment model that is less resource intensive and easier to maintain and update.

- **Key Features and Challenges:**

- Easy to deploy using Docker, e.g., `docker run -d --name perfsonar-testpoint --net=host perfsonar/testpoint`.
- Concerns:
 - Potentially limited local cache to sustain central service outages.
 - Limited remote monitoring capability without additional features or packages.
 - Deployment docs, performance and integration with OS need development & testing (e.g., policy based routing for multi-NIC)

- **Future Development and Operations:**

- We are developing instructions for OS bootstrapping, Docker-compose files, and OS configuration scripts.
- Monitoring improvements rely on Prometheus perfSONAR node exporter and host exporter and these need testing and tuning.
- Focus will be on stability, persistent configuration, caching results, and automated updates via cronjobs.

Target for initial version: ready by end of **summer 2025** (depends on pS developments)

perfSONAR Infrastructure Monitoring

- Updated to CheckMK 2.3.0 (from 1.6.0)
- Integration with CILogon (single-sign on) - moving away from x509 certs
- New tests
 - Node diagnostics based on “pscheduler troubleshoot” command
 - Tracking measurements in central UC ElasticSearch
- Now in pre-production at psetf-itb.aglt2.org (uses OSG CILogon registr.)

Main dashboard
Monitor > Overview > Main dashboard

Dashboard Add Dashboards Display Help

Host statistics

308 Up	0 In downtime	0 Unreachable	0 Down	308 Total
--------	---------------	---------------	--------	-----------

Service statistics

1854 OK	0 In downtime	0 On down host	276 Warning	1224 Unknown	1661 Critical	5015 Total
---------	---------------	----------------	-------------	--------------	---------------	------------

Host Problems (unhandled)

State	Host	Icons	Age	Summary
-------	------	-------	-----	---------

Service Problems (unhandled)

Your query produced more than 1000 results. Repeat query and allow more results. Note: the shown results are incomplete a

State	Host	Service	Icons	Summary
CRIT	personar.dur.scotgrid.ac.uk	perfSONAR services: web/https IPv6	🚫	connect to address personar.dur.scotgrid.ac.uk at 443: Connection timed out
CRIT	personar.nersc.gov	perfSONAR services: web/https IPv6	🚫	TCP CRITICAL - Invalid hostname, address or soc personar.nersc.gov
CRIT	ps-development.bnl.gov	perfSONAR services: owamp	🚫	connect to address ps-development.bnl.gov and p Connection timed out
CRIT	ps-latency.clumeq.mcgill.ca	perfSONAR services: owamp	🚫	connect to address ps-latency.clumeq.mcgill.ca at 861: Connection timed out

Events of recent 4 hours

Time	Host	Service	Summary
191 s	personar1.nipne.ro	perfSONAR services: web/https IPv6	connect to address personar1.nipne.ro and port 443: Connection timed out
8 m	personar1.nipne.ro	perfSONAR services: web/https IPv6	connect to address personar1.nipne.ro and port 443: Connection timed out
23 m	psonartest2.fnal.gov	perfSONAR services: owamp	TCP OK - 127.300 second response time on psonartest2.fnal.gov port 861
31 m	t2-pfsn2.jinr.ru	perfSONAR services: pscheduler	UNKNOWN - Exception caught while executing plugin (invalid literal for int()) with base 10: b'<DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0/EN">
47 m	ps100.farm.particle.cz	perfSONAR configuration: meshes	OK - Auto-URL configured correctly

Overview

Hosts	Unhandled p.
308	0
Services	Unhandled p.
5015	3161
Events	Unhandled p.
0	0

Bookmarks

Master control

- Notifications
- Service checks
- Host checks
- Flap detection
- Event handlers
- Performance data

Quicksearch

Measurements via Grafana

Our **MaDDash** has been replaced with **Grafana + ESnet plugins** that allow us to directly display data already gathered in our Central Measurement Archive.

(see beta version at <https://maddash.aglt2.org/dashboards>)

This is a much better situation than what we had with **MaDDash** because:

- No need to “re”gather data from each perfSONAR instance.
- No latency issues for the displayed data (previously up to tens of hours).
- Provides direct visibility for the centrally gathered data.

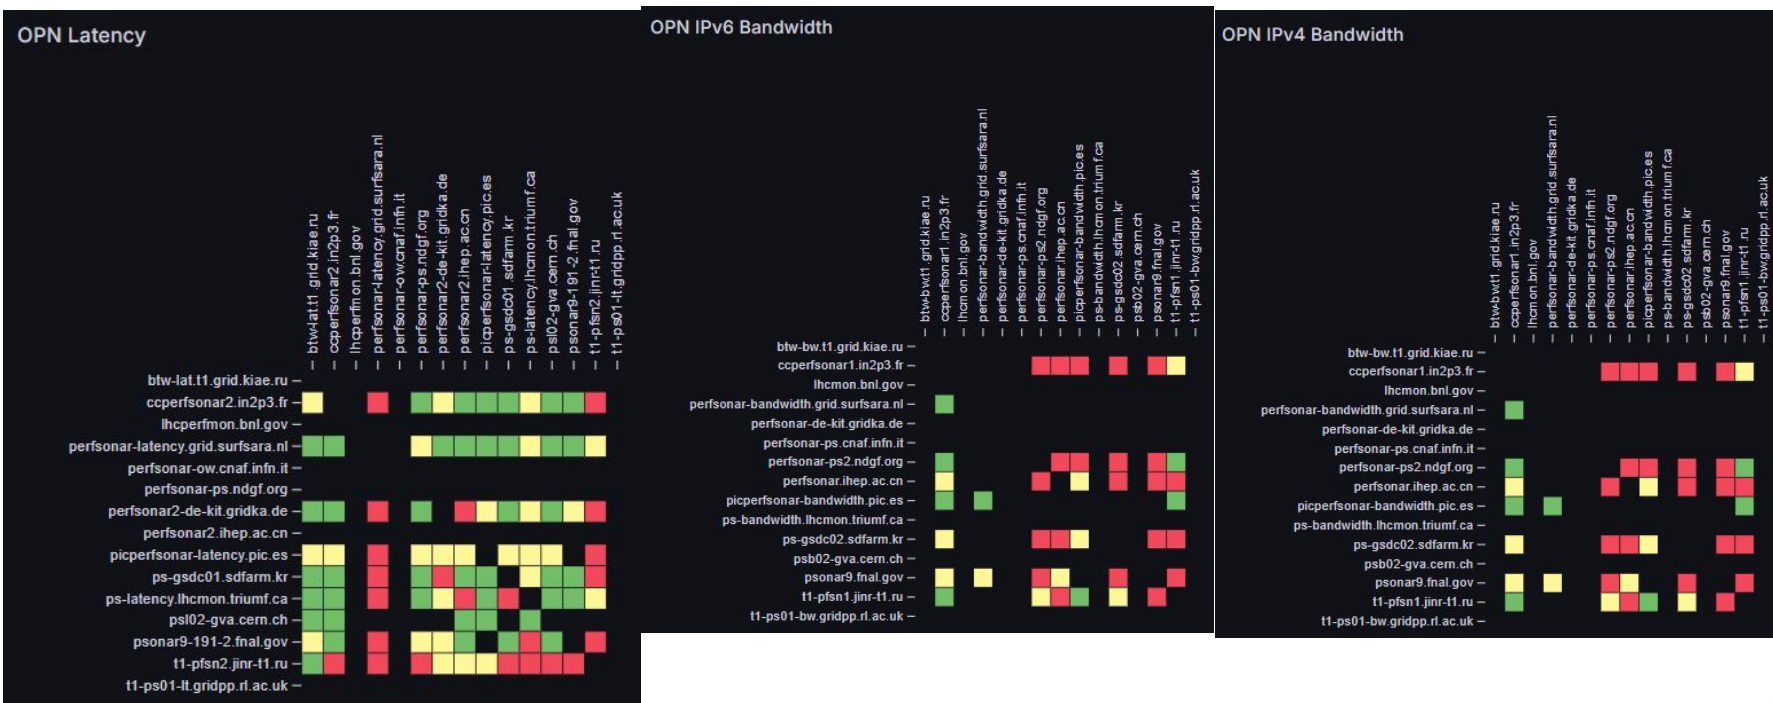
We do have some challenges to address

- The drill-down capabilities (clicking on a mesh box) needs work because our central data format differs from perfSONAR devs version
- We also need to change what data is gathered since we are missing some required metadata for some measurement types.



Examples for LHCOPN from MaDDash Replacement

Below are some examples of current LHCOPN meshes from the MaDDash replacement built upon the new perfSONAR developers Grafana+ESnet plugin




Analytics Summary

Network Analytics has been a long-term effort for our small group. We want to better exploit the large complex network information we gather.

There are four general areas of work here:

- **Data-preprocessing**, e.g.
 - Train neural networks to predict network paths, e.g. help us fill the gaps in traceroute(s)
- Build **model(s)** that represents our network(s)
 - Network measurements are inherently noisy and therefore require robust models
- Use ML models for **anomaly detection** (for alerts & alarms)
 - Neural networks, Bayesian/probabilistic approaches
 - Detect anomalies in network paths and bandwidth measurements
 - Compare with the existing heuristic algorithms that we have developed
- **Correlation** with other data
 - Traceroutes with throughput for example, but also outside of perfSONAR, e.g. FTS
 - New types of data appearing (high-touch, scitags, in-band telemetry, etc.)




Analytics Core Team

We have a number of people who meet weekly to discuss network analytics and our measurement and analysis platforms.

The core group is comprised of Marian Babik, Shawn McKee, **Petya Vasileva** and Ilija Vukotic.

We often have others join our meetings from ATLAS, CMS, R&E networks, the perfSONAR team and site or network admins. (Notes available [here](#))

Through **IRIS-HEP** (NSF funded institute) we have also managed to bring in a few fellows who have worked on specific projects identifying network issues or creating alerts and alarms. This year we have a returning IRIS-HEP fellow, **Yana Holoborodko** working with us at CERN till October.



Analytics Tools

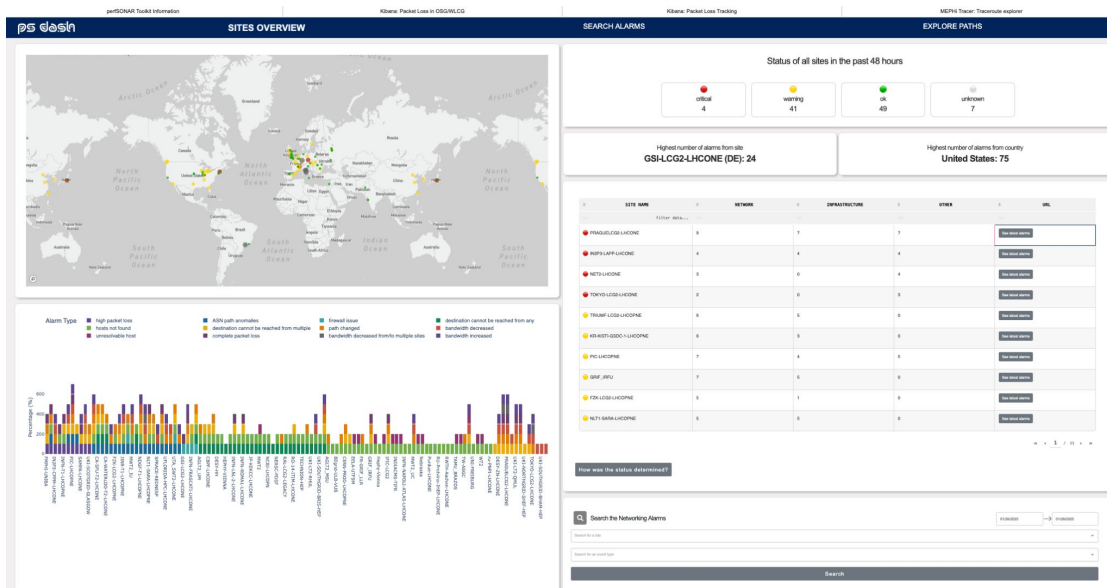
We have developed tools to analyze and alert upon alarms identified.

Petya has given recent [detailed presentation](#) on this work at the February 2025 ATLAS S&C meeting at CERN.

The main user interface is pS-Dash (shown here =>)

This is a web GUI that serves as a front-end to our gathered data and generated alarms.


[Explore it](#) and let us know your comments/suggestions



Summary

- Updates to perfSONAR and OSG/WLCG network measurement platform
 - perfSONAR 5.2 is coming with new features some of which are required for our evolution.
 - New infrastructure monitoring and dashboards are (or will appear) in production.
 - Simplified deployment models are being evaluated and documented.
- Ongoing efforts in network analytics and ML methods for our data
 - Focus on pre-processing (gaps, predictive models) and anomaly detection
 - Opportunity to collaborate on models and data sets
- **We are running monthly meetings with site network teams:**
 - Discuss how sites are deploying, managing their network infrastructure and planning for WLCG networking requirements
 - [Meetings](#) every 3rd week of a month on Thursday at 10am EST (to join mail wlcg-site-net-requests@umich.edu)
- **We have to continue to watch our network monitoring infrastructure as it is a complex system with lots of areas for issues to develop.**

Question or Comments?



Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- [OSG: NSF MPS-1148698](#)
- [IRIS-HEP: NSF OAC-1836650](#)



Useful URLs

- OSG/WLCG Networking Documentation
 - <https://opensciencegrid.github.io/networking/> (old, being updated soon)
- perfSONAR Infrastructure Dashboard
 - <https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7>
- perfSONAR Dashboard and Monitoring
 - <http://maddash.aglt2.org/maddash-webui>
 - (old) https://psetf.opensciencegrid.org/etf/check_mk
 - (new, beta) https://psetf-itb.aglt2.org/etf/check_mk
- perfSONAR Central Configuration
 - <https://psconfig.opensciencegrid.org/>
- Toolkit information page
 - <https://toolkitinfo.opensciencegrid.org/>
- ATLAS Alerting and Alarming Service: <https://aaas.atlas-ml.org/>
- The pS Dash application: <https://ps-dash.uc.ssl-hep.org/>
- ESnet WLCG DC Dashboard:
<https://public.stardust.es.net/d/1kFCB5Hnk/lhc-data-challenge-overview?orgId=1>

Backup Slides Follow