


WLCG Network Monitoring and Analytics Updates

Shawn McKee / U Michigan, Marian Babik / CERN
Petya Vasileva / U Michigan, Ilija Vukotic / U Chicago
on behalf of WLCG Network Throughput WG

2024

Spring 2024 HEPiX, Paris, France

<https://indico.cern.ch/event/1377701/timetable/#20240417>



Outline

- Infrastructure Status and Updates
- DC24 Review
- Network Analytics



Open Science Grid



WLCG
Worldwide LHC Computing Grid

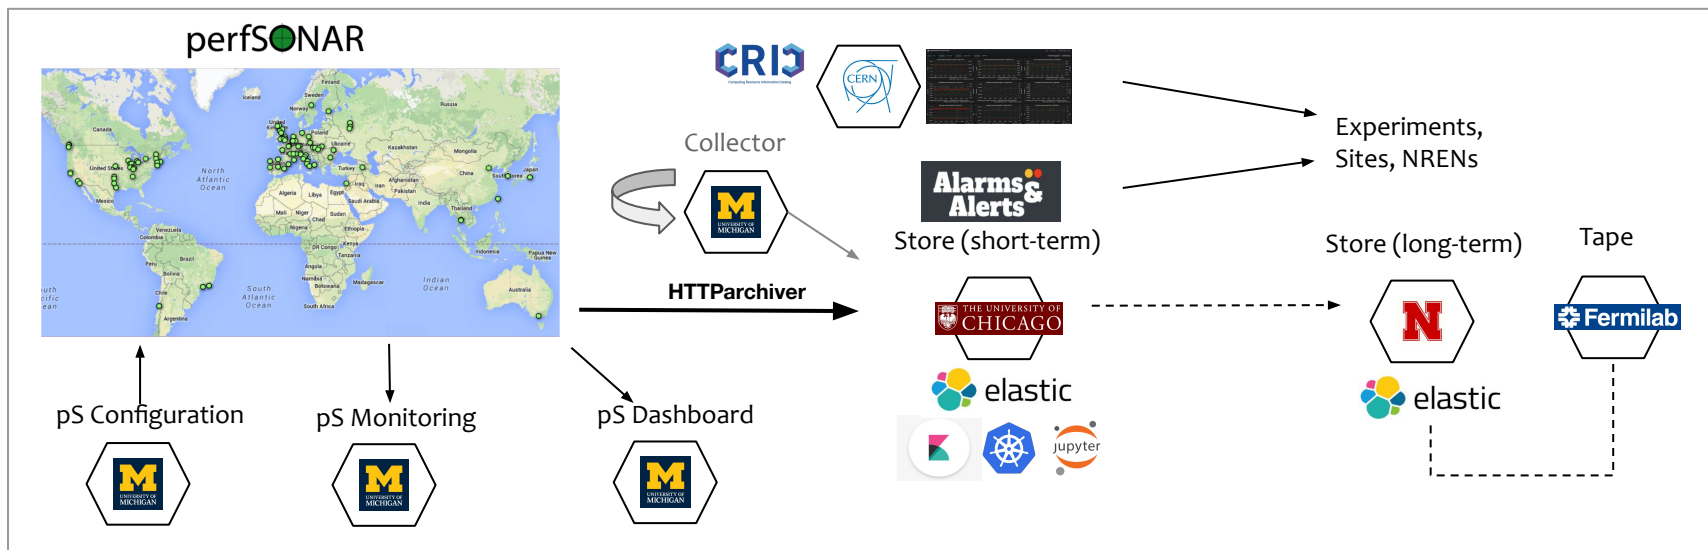


- perfSONAR 5.0.8 is the latest release
 - Number of bug-fixes since 5.0
 - Weekly meetings with the developers
 - Update campaign in WLCG
 - Various issues, mostly archiving, but also e.g. legacy limits configuration ([fix](#))
 - Toolkit support for latest CC7 and Alma/Rocky 8 and 9 compatible systems (Alma).
 - **Sites should plan to update by June (end-of-life for CentOS7)**
- **5.1 Beta Release**
 - New Grafana interface - **replacing** toolkit and maddash graphs
 - Threaded iperf3 support
 - Enhanced instrumentation, better troubleshooting of archiving issues
 - **OS support:** Alma/Rocky 9, Debian 11/12, Ubuntu 20/22 (updated docker),
No support for CentOS 7, Debian 10, Ubuntu 18




Network Measurement Platform Status

- Our platform collects, stores, configures and transports all network metrics
- Evolution based on the perfSONAR 5 already partially implemented.
 - Now directly publishing results from perfSONARs to ES@UC
 - Collector used **only as a fallback**;
 - WLCG CRIC now used for topology



Network Measurement Platform Plans

- Evolution based on the perfSONAR 5 already partially implemented.
 - 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 (retire maddash)
 - **ps Configuration** - clarify development roadmap and support




perfSONAR Infrastructure


210
Active perfSONARs

Number of Communities in Use
70


NIC speed



perfSONAR distributions



pS NIC MTU



- Core deployments are still on 10Gbps, we have about 17% with 100Gbps
 - For WLCG/OSG testing purposes 10Gbps is still sufficient
 - Important to **refresh HW** along with the update to 5.1
- Most of the infrastructure is on 5.0.8, but significant fraction still on **CentOS 7**
- MTU - around 40% on jumbo frames (9000), rest is on standard frames (1500)
- We have small testbed with about 10 perfSONARs with BBRv3 enabled
 - Enabled testing TCP congestion algorithm benefits and jumbo frame trade-offs
 - Open for participation

DC24

WLCG Data Challenge 2024 took place in Feb 2024; targeting 25% of HL-LHC

Our DC24 plans included the following:

- Update and utilize **perfSONAR** to clean up links and fix problems before DC24.
- Instrument and document **site networks**, for at least our largest sites.
- **Network planning**: we need to make sure our sites and their local and regional networks are aware of our requirements and timeline and are planning appropriately
- **IPv6** should be enabled everywhere not just because of packet marking, but because it will allow us to get back to a single stack sooner!



psDash Network Status

psDash

SITES OVERVIEW

SEARCH ALARMS

EXPLORE PATHS

MAJOR ALARMS

Status of all sites in the past 24 hours

Summary

0

28

104

3

Highest number of alarms from site

**KR-KISTI-GSDC-1-LHCOPNE
(KR): 75**

Highest number of alarms from country

United Kingdom: 108


SITE	STATUS	NETWORK	INFRASTRUCTURE	OTHER	URL
filter data...					
KR-KISTI-GSDC-1-LHCOPNE	●	75	0	0	See latest alarms
TR-ULAKBIM-LHCONE	●	1	1	0	See latest alarms
INFN-LNL-2-LHCONE	●	1	3	0	See latest alarms
FMPHI-UNIBA	●	2	1	0	See latest alarms
BUDAPEST	●	2	1	0	See latest alarms
BEGRID-ULB-VUB	●	2	1	0	See latest alarms
TECHNION-HEP	●	2	1	1	See latest alarms
UKI-LT2-RHUL	●	2	2	0	See latest alarms
NCBJ-LHCOPN	●	2	2	0	See latest alarms



Network Status dashboard - part of Network Analytics platform - shows network performance based on perfSONAR measurements. Status (ok/warning/critical/unknown) aggregates network and infrastructure metrics.

Site Network Utilisation

WLCG Site Network Input/Output



Site Network Utilisation - computed from aggregated utilisation (SNMP counters) provided by sites via simple API. Screenshot shows network utilisation during DC24 as seen by the sites.

Analytics

Alarms & Alerts Interface

Components

Database

Elasticsearch

REST API and Web frontend

Node.js + express + pug

Deployment

Docker, K8s, Helm (soon)

Authentication

Globus InCommon

Authorization

API key

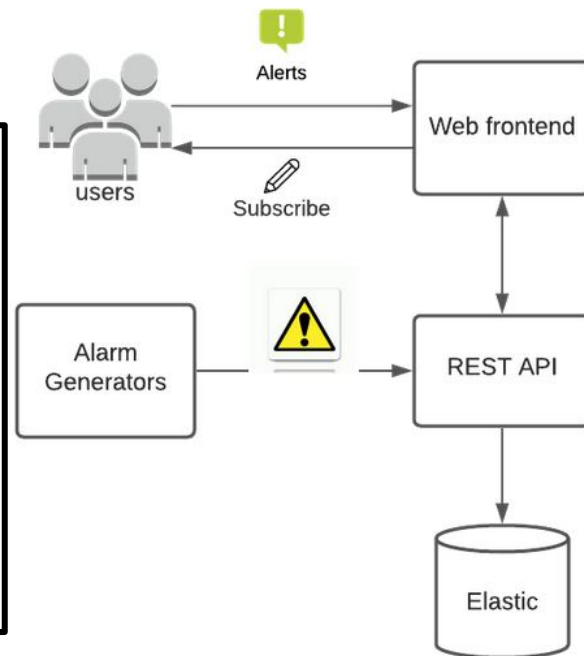
Mail

Mailgun

<https://psa.osg-htc.org>

(Uses EDUGain/InCommon)

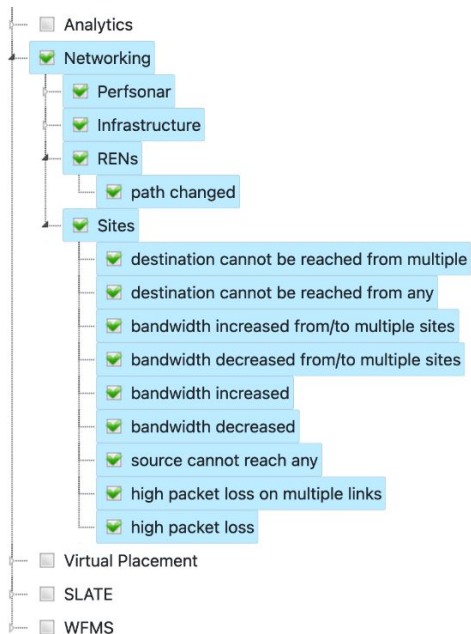
Purpose: provides user-subscribable alerting for specific types of network issues found by analyzing perfSONAR data



Two main improvements needed: **Acknowledging alerts** that are being worked on and **adding user notification mailing lists**

Subscription Interface

Alarms



Heartbeats

- Virtual Placement
- SLATE

Current Subscriptions

Category	Subcategory	Event	Tags
Networking	Perfsonar	bad owd measurements	*
Networking	Perfsonar	large clock correction	*
Networking	Perfsonar	complete packet loss	*
Networking	Perfsonar	firewall issue	MWT2
Networking	Infrastructure	indexing	*
Networking	Sites	destination cannot be reached from multiple	*
Networking	Sites	destination cannot be reached from any	*
Networking	Sites	high packet loss on multiple links	BNL-ATLAS
Networking	Sites	source cannot reach any	*
Networking	Sites	high packet loss	*
Networking	Sites	bandwidth decreased from/to multiple sites	*
Networking	Sites	bandwidth decreased	*

Alarm Types and Relation to perfSONAR Data

All based on perfSonar data

One-Way Delay

- ▶ bad owd measurements
- ▶ large clock correction

Traceroute

- ▶ path changed
- ▶ destination cannot be reached
- ▶ source cannot reach any

Packet loss

- ▶ complete packet loss
- ▶ firewall issue
- ▶ high packet loss (on multiple links)

Throughput

- ▶ bandwidth decreased (from/to multiple sites)
- ▶ bandwidth increased (from/to multiple sites)

psDash Alarms Dashboard

perSONAR Toolkit Information

Kibana: Packet Loss in OSG/WLCG

Kibana: Packet Loss Tracking

MEPHi Tracer: Traceroute explorer


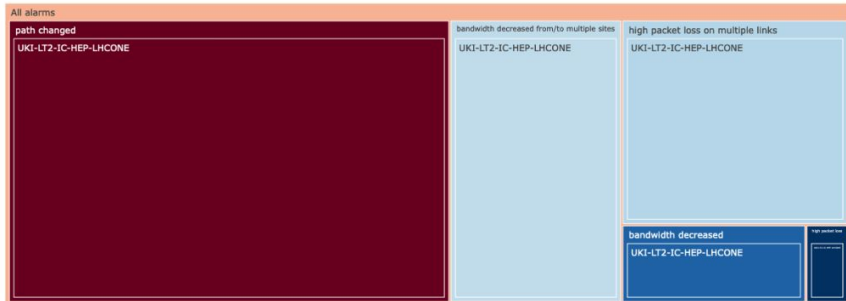


SITES OVERVIEW

SEARCH ALARMS

EXPLORE PATHS

MAJOR ALARMS



Search & Explore the Networking Alarms

03/14/2024 → 04/16/2024
Rounded to the day

UKI-LT2-1C-HEP-LHCONE

path changed | bandwidth decreased | bandwidth decreased from/to multiple sites | high packet loss on multiple links | high packet loss

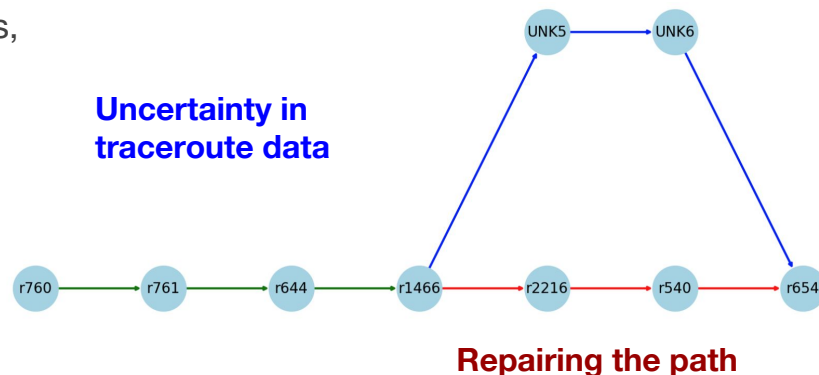
List of alarms

BANDWIDTH DECREASED

from	to	src_site	dest_site	ipv	ipv6	last3days_avg	change	alarm_link
2024-03-20 04:08:02.000Z	2024-04-10 04:08:02.000Z	UKI-LT2-1C-HEP-LHCONE	BEIJING-LCG2-LHCONE	ipv4	false	41	-61	VIEW IN A NEW TAB
2024-03-19 04:08:03.000Z	2024-04-09 04:08:03.000Z	UKI-LT2-1C-HEP-LHCONE	BEIJING-LCG2-LHCONE	ipv4	false	40.5	-61	VIEW IN A NEW TAB
2024-03-13 04:08:02.000Z	2024-04-03 04:08:02.000Z	UKI-LT2-1C-HEP-LHCONE	BEIJING-LCG2-LHCONE	ipv4	false	38	-61	VIEW IN A NEW TAB
2024-03-13 04:08:02.000Z	2024-04-03 04:08:02.000Z	UKI-LT2-1C-HEP-LHCONE	BEIJING-LCG2-LHCONE	ipv6	true	39	-60	VIEW IN A NEW TAB
2024-03-12 04:08:02.000Z	2024-04-02 04:08:02.000Z	UKI-LT2-1C-HEP-LHCONE	NDGF-T1-LHCOPNE	ipv4	false	312	-75	VIEW IN A NEW TAB
2024-03-06 04:08:03.000Z	2024-03-27 04:08:03.000Z	UKI-LT2-1C-HEP-LHCONE	UKI-NORTHGRID-LANCS-HEP	ipv6	true	912	-66	VIEW IN A NEW TAB
2024-03-06 04:08:03.000Z	2024-03-27 04:08:03.000Z	UKI-LT2-1C-HEP-LHCONE	UKI-NORTHGRID-MAN-HEP	ipv6	true	912.33	-71	VIEW IN A NEW TAB

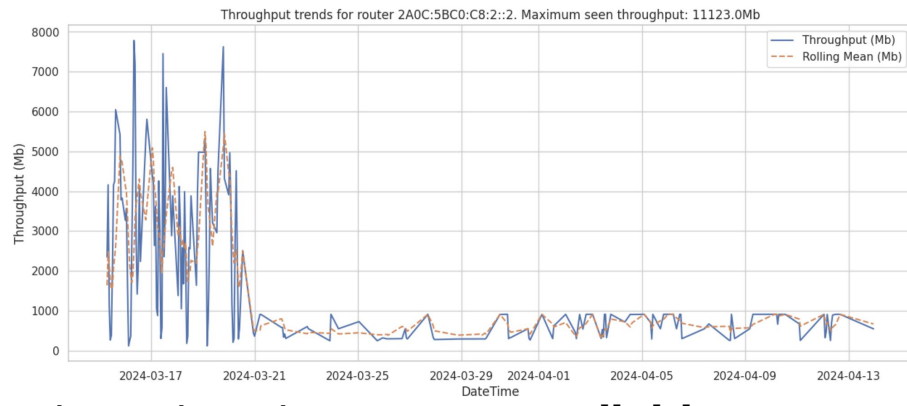
Network Analytics R&D

- Investigate ML models/methods to process network measurements
- 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 (which ones ?), Bayesian/probabilistic approaches,
 - Detect anomalies in network paths and bandwidth measurements
 - Compare with the existing heuristic algorithms that we have developed
- Correlate 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.)



Plans for the Analytics Platform


- Production of the anomaly detection based on Bayesian inference
 - Uses RTT, traceroutes, TTLs as input and detects anomalies
- Continue working on the neural network models that correlate throughputs and traceroutes
 - Generating real-world model of our entire network (all routers)
 - Not only detecting anomalies, but also trying to pinpoint the location of the issue →
- Improve infrastructure alarming to the point where we can **reliably differentiate infrastructure and network issues**
- Network availability dashboard in production



Summary

- Updates to perfSONAR and OSG/WLCG network measurement platform
 - perfSONAR 5.1 is coming with new features and will require all sites to update OS.
 - Plan to adapt the network measurement platform to benefit from changes in 5.1
- 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 preparing monthly meetings with site network teams:**
 - Discuss how sites are deploying, managing and planning for WLCG networking requirements
 - Next meeting April 18th 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.

Questions / Discussion?



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/>
- perfSONAR Infrastructure Dashboard
 - <https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7>
- WLCG DOMA DC24 plans
 - <https://indico.cern.ch/event/1225415/contributions/5155042/attachments/2593516/4476291/Data%20Challenge%202024.pdf>
- perfSONAR Central Configuration
 - <https://psconfig.opensciencegrid.org/>
- Toolkit information page
 - <https://toolkitinfo.opensciencegrid.org/>
- Grafana dashboards
 - <http://monit-grafana-open.cern.ch/>
- ATLAS Alerting and Alarming Service: <https://psa.osg-htc.org/>
- The perfSONAR Dashboard application: <https://ps-dash.uc.ssl-hep.org/>
- ESnet WLCG Stardust Dashboard:
<https://public.stardust.es.net/d/XkxDL5H7z/esnet-public-dashboards?orgId=1>

Backup Slides Follow


WLCG Network Throughput Support Unit

Support channel where sites and experiments can report potential network performance incidents:

- Relevant sites, (N)RENs are notified and perfSONAR infrastructure is used to narrow down the problem to particular link(s) and segment. Also [tracking past incidents](#).
- Feedback to WLCG operations and LHCOPN/LHCONE community

Most common issues: MTU, MTU+Load Balancing, routing (mainly remote sites), site equipment/design, firewall, workloads causing high network usage

As there is no consensus on the MTU to be recommended on the segments connecting servers and clients, LHCOPN/LHCONE working group was established to investigate and produce a recommendation.



Importance of Measuring Our Networks

- **End-to-end network issues are difficult to spot and localize**
 - Network problems are multi-domain, complicating the process
 - Performance issues involving the network are complicated by the number of components involved end-to-end
 - Standardizing on specific tools and methods focuses resources more effectively and provides better self-support.
- **Network problems can severely impact experiments workflows and have taken weeks, months and even years to get addressed!**
- **perfSONAR provides a number of standard metrics we can use**
 - Latency, Bandwidth and Traceroute
 - These measurements are critical for network visibility
- **Without measuring our complex, global networks we wouldn't be able to reliably use those network to do science**

