

perfSONAR Monitoring Update

Shawn McKee / U Michigan, Marian Babik / CERN
on behalf of WLCG Network Throughput WG

2024


At the #52 LHCOPN/LHCONE Meeting, Beijing, China

<https://indico.cern.ch/event/1410638/>



Outline


- News and Updates
- Network Measurement Platform
- New Infrastructure Monitoring (psetf)
- New dashboards (maddash)
- Analytics



perfSONAR News

- **perfSONAR 5.1.3** is the latest release
 - Number of bug-fixes since 5.0; bi-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 Alma/Rocky 8 and 9, Debian 11/12, Ubuntu 20/22
 - **CentOS7 is no longer supported, sites should update ASAP**
- **perfSONAR session at LHCOPN/LHCONE tomorrow**

3:45 PM	perfSONAR News and Updates (TBC) Speaker: Szymon Trocha (Poznan Supercomputing and Networking Centre (PCSS))	🕒 20m
4:05 PM	Experiences with perfSONAR in Jisc/UK Speaker: Tim Chown	🕒 20m
4:25 PM	perfSONAR Network Analytics through Machine Learning Speaker: Petya Vasileva (University of Michigan (US))	🕒 20m




perfSONAR Infrastructure


181
Active perfSONARs

68
Communities in Use


NIC speed



perfSONAR distributions



pS NIC MTU



- 64% toolkits now on 5.1 and EL9
 - 25% still on EL7 !
- Core deployments are still on 10Gbps, but we have about 20% with 100Gbps
 - For WLCG/OSG testing purposes 10Gbps is still sufficient
 - Important to refresh HW along with the update to EL9
- MTU - around 36% on jumbo frames (9000), rest is on standard frames (1500)
- We have small testbed with about 10 perfSONARs with BBRv3 enabled


perfSONAR Infrastructure Evolution

181

Active perfSONARs

68

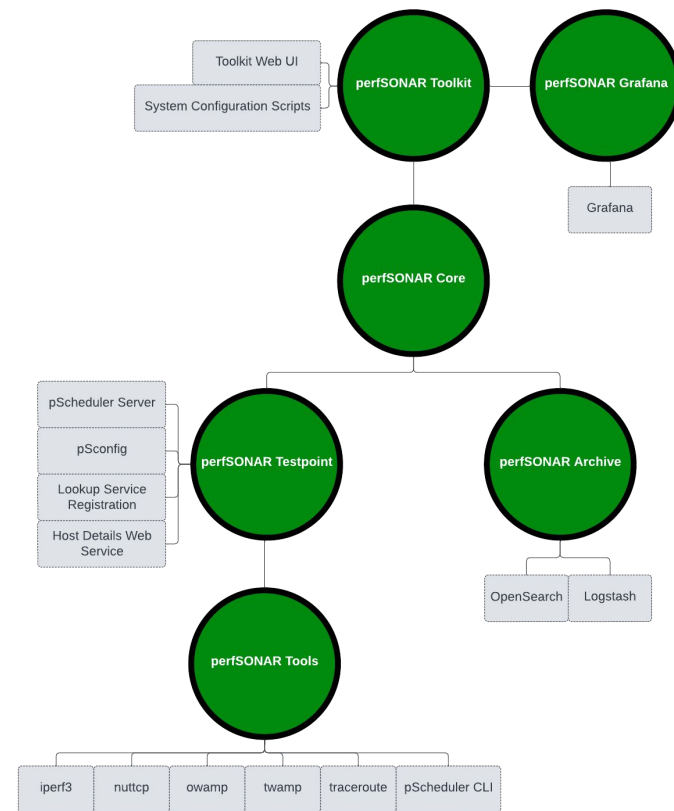
Communities in Use



- Long-term trend of perfSONAR releases adoption shows a steady downtrend
 - Only reachable “toolkit version” installations are shown
 - 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 ES)
 - 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)

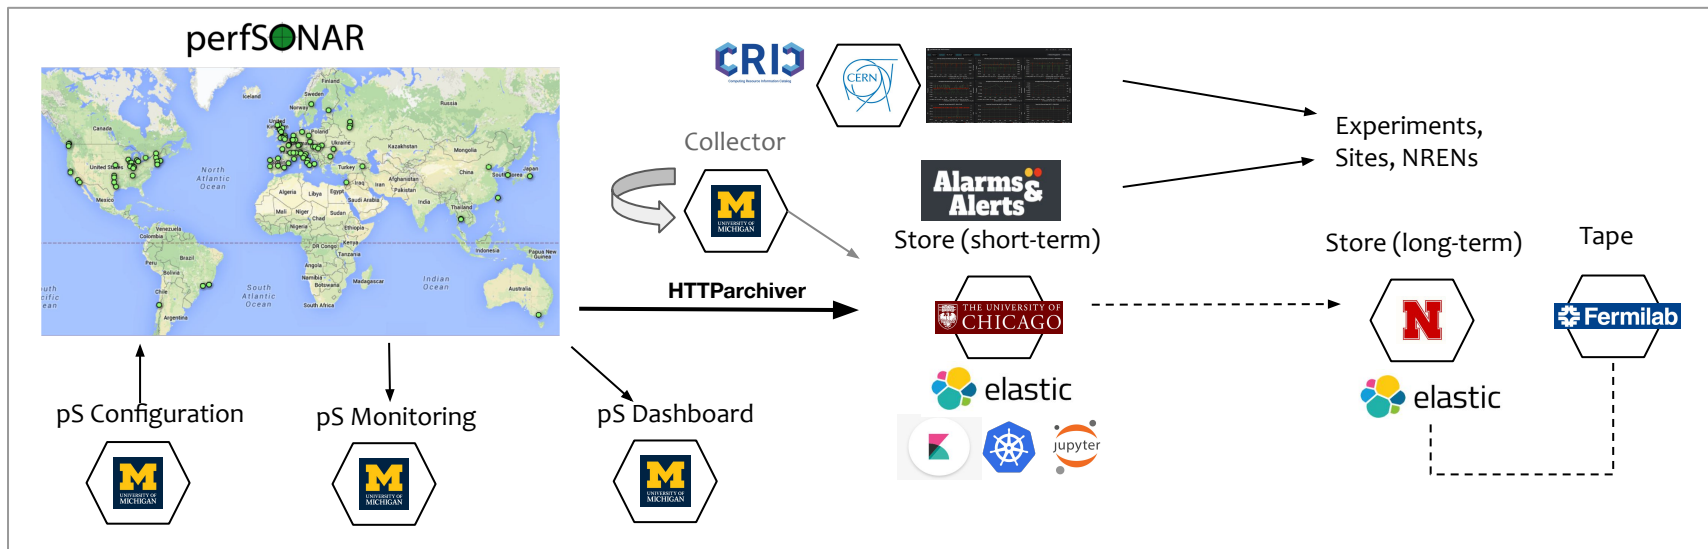
perfSONAR Testpoint

- perfSONAR Testpoint:
 - Automatically runs tests on a regular schedule
 - Participates in a centrally managed set of tests
 - Publishes the existence of a measurement node
- Can run in a container, i.e. starting testpoint as easy as
 - `docker run -d --name perfsonar-testpoint --net=host perfsonar/testpoint`
- Main concerns
 - No local cache that would sustain an outage of central services (48-72 hours)
 - No way to remotely monitor - needs extra packages/containers to run `node_exporter` and `perfsonar exporter`
 - Missing auto-updates and persistent configuration
 - Performance and integration with OS needs testing
- Looking for volunteers willing to deploy, test and develop the missing components



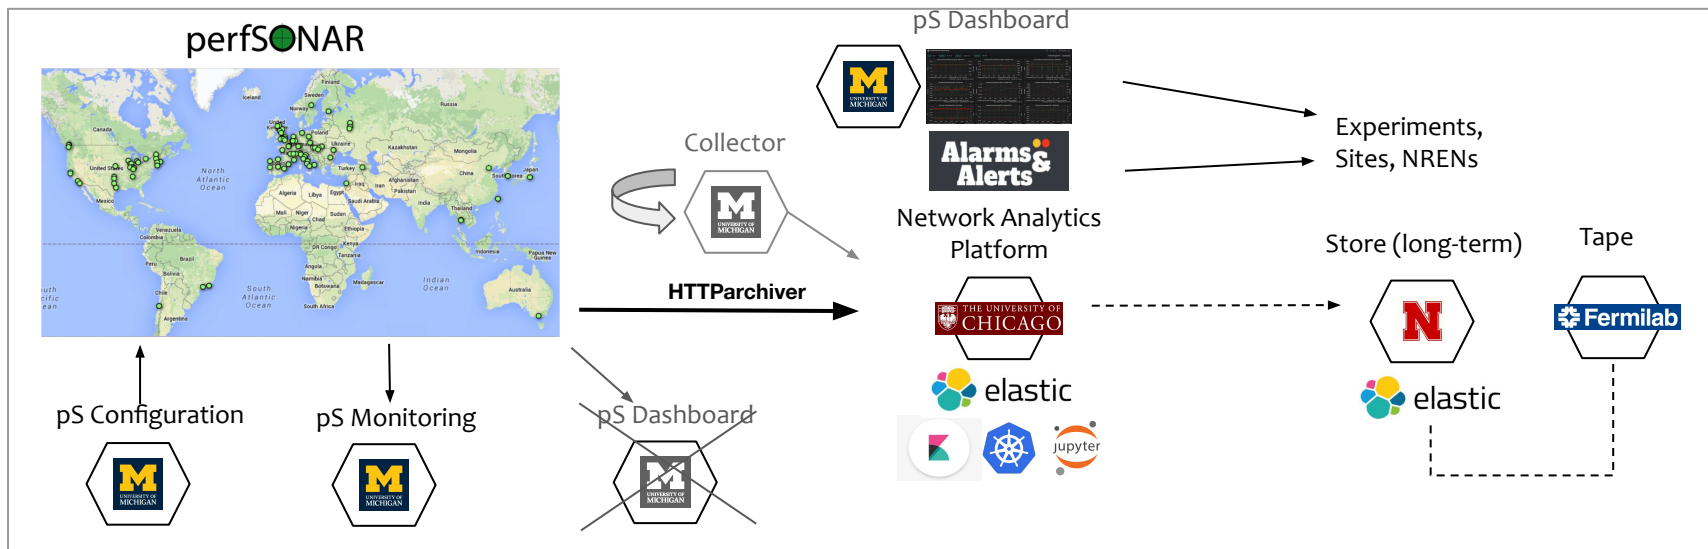
Network Measurement Platform Status

- 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



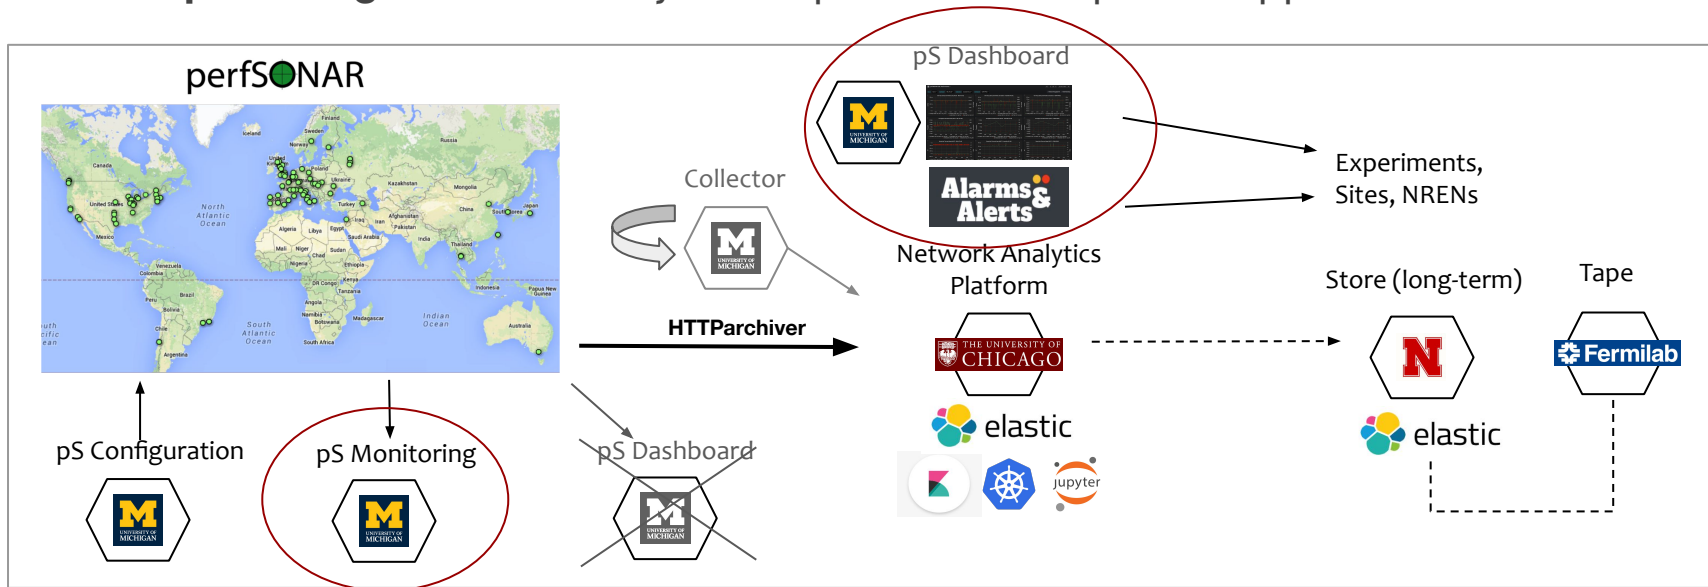
Network Measurement Platform Plans

- 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** - new tool has been announced



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 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 ElasticSearch
- Now in pre-production at psetf-itb.aglt2.org

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
CRIT	personar.dur.scotgrid.ac.uk	🚫	191 s	connect to address personar.dur.scotgrid.ac.uk at 443: Connection timed out
CRIT	personar.nersc.gov	🚫	8 m	TCP CRITICAL - Invalid hostname, address or socket address: Connection timed out
CRIT	ps-development.bnl.gov	🚫	23 m	connect to address ps-development.bnl.gov and port 861: Connection timed out
CRIT	ps-latency.clumeq.mcgill.ca	🚫	31 m	connect to address ps-latency.clumeq.mcgill.ca at 861: Connection timed out

Service Problems (unhandled)

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

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-phys2.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: 308 Unhandled p. 0
- Services: 5015 Unhandled p. 3161
- Events: 0 Unhandled p. 0

Bookmarks

Master control

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

Quicksearch

perfSONAR Infrastructure Monitoring



Services of host ps-latency.lhcmon.triumf.ca

Monitor > Overview > All hosts > ps-latency.lhcmon.triumf.ca > Services of host



Monitor



Customize



Setup

Commands Host Services Export Display Help ↕

Acknowledge problems

Schedule downtimes

Filter

Show checkboxes

ps-latency.lhcmon.triumf.ca

ps-latency.lhcmon.triumf.ca

State	Service	Icons	Summary	Age	Checked	Perf-O-Meter
OK	perfSONAR configuration: contacts	☰	OK - Contact and organization found	2024-07-02 17:26:53	47 m	
OK	perfSONAR configuration: location	☰	OK - Location: -122.5428/49.1646	2024-07-02 17:26:53	47 m	
OK	perfSONAR configuration: meshes	☰	OK - Auto-URL configured correctly	2024-07-02 17:26:53	47 m	
WARN	perfSONAR ES freshness: owamp	☰	WARNING - ES 52.78 % complete for event type: histogram-owdelay	6 h	6 h	
CRIT	perfSONAR esmond freshness: owamp	☰	CRITICAL - esmond 0.0 % complete for event type: histogram-owdelay	2024-07-01 16:22:33	112 m	
OK	perfSONAR hardware check	☰	OK - CPU:1/24cores/3700.0000Mhz RAM:62GB NIC:10.0Gbps/9000MTU/IPv6 enabled	2024-07-02 17:26:53	47 m	
OK	perfSONAR json summary	☰	OK - Toolkit metadata successfully retrieved	2024-07-02 17:26:50	47 m	
OK	perfSONAR services: ntp	☰	OK - NTP synchronized	2024-07-02 17:26:53	47 m	
OK	perfSONAR services: owamp	☰	TCP OK - 0.056 second response time on ps-latency.lhcmon.triumf.ca port 861	2024-08-25 15:14:32	7.15 s	
OK	perfSONAR services: pscheduler	☰	OK - pscheduler stats retrieved	2024-07-02 17:35:25	39 m	
WARN	perfSONAR services: pscheduler diags	☰	Performing basic troubleshooting of ps-latency.lhcmon.triumf.ca.	19 h	19 h	
OK	perfSONAR services: regular testing/pscheduler	☰	OK - pscheduler is enabled and running	2024-08-31 15:26:54	47 m	
OK	perfSONAR services: versions	☰	OK - Toolkit version found: 5.0.8-1.e17	2024-07-02 17:26:53	47 m	
OK	perfSONAR services: web/https IPv6	☰	TCP OK - 0.056 second response time on ps-latency.lhcmon.triumf.ca port 443	2024-08-25 15:14:41	14 m	

perfSONAR Infrastructure Monitoring: Freshness

Service perfSONAR ES freshness: owamp, ps-latency.lhcom.triumf.ca
Monitor > Overview > All hosts > ps-latency.lhcom.triumf.ca > Services of host > Service

Commands Service Host Export Display Help

Monitor: Acknowledge problems Schedule downtimes Filter Show checkboxes Services of host

Customize: 1 row

Site alias	Local site etf
Host name	ps-latency.lhcom.triumf.ca
Service description	perfSONAR ES freshness: owamp
Service labels	
Service icons	☰ 🐾
Service state	WARN
Summary	WARNING - ES 52.78 % complete for event type: histogram-owdelay Time-range: 3600 Even-types checked: histogram-owdelay Mesh (Event-type): Canadian Tier1-2 Latency (histogram-owdelay) Missing destinations: 1 Missing destinations: lcg-bw.sfu.computecanada.ca Mesh (Event-type): LHCONe Latency (histogram-owdelay) Missing destinations: 6 Missing destinations: lat.scinet.utoronto.ca, perfsonar-ps-01.desy.de, wash-ps-ps-lat.es.net, atlas-npt1.bu.edu, perfsonar2.ultralight.org, perfsonar-latency-v471-portal.gridka.de Mesh (Event-type): OPN Latency (histogram-owdelay) Missing destinations: 2 Missing destinations: lhccperfmon.bnl.gov, t1-ps01-lt.gridpp.rl.ac.uk
Details	Mesh (Event-type): WLCG ATLAS Latency (histogram-owdelay) Missing destinations: 28 Missing destinations: sonar2.itim-cj.ro, ps1.oceph.ou.edu, perfsonar01.lcg.cscs.ch, perfsonar-lt.tier2.hep.manchester.ac.uk, psonar1.lal.in2p3.fr, mwt2-ps01.campuscluster.illinois.edu, perfsonar2.na.infn.it, perfsonar2.cc.kek.jp, ps-latency.atlas.unimelb.edu.au, ps-latency.clumeq.mcgill.ca, perfsonar01.datagrid.cea.fr, pygrid-sonar2.lancs.ac.uk, perfsonar-ps.ndgf.org, lapp-ps02.in2p3.fr, ps0002.m45.ihep.su, ps-latency.hepnetcanada.ca, t2ps-latency2.physics.ox.ac.uk, hepsonar2.ph.liv.ac.uk, lcg-pslat.uw.computecanada.ca, t1-pfsn2.jinr-t1.ru, perfsonar2.ihep.ac.cn, lpnhe-psl.in2p3.fr, tech-ps.hep.technion.ac.il, perfsonar02-lep-grid.saske.sk, btw-lat.grid.kiae.ru, uct2-net1.mwt2.org, iut2-net09.iu.edu, echolot1.zeuthen.desy.de, ccperfsonar2.in2p3.fr, perfsonar1.icepp.jp, btw-lat.t1.grid.kiae.ru, psi01-gva.cern.ch, perfsonar2-de-kit.gridka.de, perfsonar-ow.cnaf.infn.it, psonar9-191-2.fnal.gov, picperfsonar-latency.pic.es Working destinations: argus-11.net2.mghpcc.org, perfsonar-latency.grid.surfsara.nl, perfsonar01.ft.uam.es, perfsonar2.nipne.ro, ps01-l.farm.particle.cz, psum02.itep.ru, llrpsonar1.in2p3.fr, ps-gsd01.sdfarm.kr, marperf02.in2p3.fr, pmsu01.aglt2.org, psum01.aglt2.org, ps002.gla.scotgrid.ac.uk, perfsonar01.datagrid.cea.fr, pygrid-sonar2.lancs.ac.uk, perfsonar-ps.ndgf.org, lapp-ps02.in2p3.fr, ps0002.m45.ihep.su, ps-latency.hepnetcanada.ca, t2ps-latency2.physics.ox.ac.uk, hepsonar2.ph.liv.ac.uk, lcg-pslat.uw.computecanada.ca, t1-pfsn2.jinr-t1.ru, perfsonar2.ihep.ac.cn, lpnhe-psl.in2p3.fr, tech-ps.hep.technion.ac.il, perfsonar02-lep-grid.saske.sk, btw-lat.grid.kiae.ru, uct2-net1.mwt2.org, iut2-net09.iu.edu, echolot1.zeuthen.desy.de, ccperfsonar2.in2p3.fr, perfsonar1.icepp.jp, btw-lat.t1.grid.kiae.ru, psi01-gva.cern.ch, perfsonar2-de-kit.gridka.de, perfsonar-ow.cnaf.infn.it, psonar9-191-2.fnal.gov, picperfsonar-latency.pic.es

Documentation for this check can be found at [www](#)

perfSONAR Infrastructure Monitoring: pscheduler

checkmk

Service perfSONAR services: pscheduler diags, psb02-gva.cern.ch

Monitor > Overview > All hosts > psb02-gva.cern.ch > Services of host > Service

30

Commands Service Host Export Display Help

Monitor

Customize 1 row

Setup

Site alias	Local site etf
Host name	psb02-gva.cern.ch
Service description	perfSONAR services: pscheduler diags
Service labels	
Service icons	☰
Service state	OK
Summary	Performing basic troubleshooting of psb02-gva.cern.ch.
Details	<pre>psb02-gva.cern.ch: Checking that host "psb02-gva.cern.ch" resolves... 2001:1458:301:86::100:b Looking for pScheduler... OK. Fetching API level... 6 Checking clock... Unsynchronized (Not considered fatal) Exercising API... Archivers... Contexts... Tests... Tools... OK. Fetching service status... OK. Checking services... Ticker... Scheduler... Runner... Archiver... OK. Checking limits... OK. Last run scheduled... 5 seconds ago Last run completed... 6 seconds ago Server Statistics: Archiving Backlog 0 Upcoming 0 HTTP Queue Backlog 0 Length 1 Runs Pending 17767 On Deck 0 Running 28 Cleanup 3 Finished 91390 Overdue 0 Missed 1 Failed 16064 Preempted 0 Non-Starting ... 84 pScheduler appears to be functioning normally.</pre>

Help

User

perfSONAR Dashboard

- New dashboard service replacing previous maddash
- Based on Grafana - based on perfSONAR 5 code base
 - Modified to use the central ElasticSearch, which uses different schema
 - Generation of dashboards required some new code which is now upstream




perfSONAR Dashboard

OPN IPv4 Traceroute



OPN IPv6 Traceroute



perfSONAR Dashboard

Q Search or jump to...

Home > Dashboards > perfSONAR pSConfig > OPN-ALL ☆

- Home
- Starred
- Dashboards**
- Playlists
- Snapshots
- Library panels
- Public dashboards
- Explore
- Metrics
- Alerting
- Alert rules
- Contact points
- Notification policies
- Silences
- Groups
- Admin
- Connections
- Add new connection
- Data sources
- Administration
- General
- Stats and license
- Default preferences

OPN Ping

	btw-bw.grid.kiae.ru	btw-lat.t1.grid.kiae.ru	ccperfsnar1.in2p3.fr	ccperfsnar2.in2p3.fr	lhcmn.bnl.gov	lhcperfmn.bnl.gov	perfsnar-bandwidth.grid.surfsara.nl	perfsnar-de-kit.gridka.de	perfsnar-ow.cnaif.infn.it	perfsnar-ps.cnaif.infn.it	perfsnar-ps.ndgf.org	perfsnar-ps2.ndgf.org	perfsnar2-de-kit.gridka.de	perfsnar2.lhep.ac.cn	plcperfsnar-bandwidth.plc.es	plcperfsnar-latency.plc.es	ps-bandwidth.lhcmn.triumf.ca	ps-gsd01.sdfarm.kr	ps-gsd02.sdfarm.kr	ps-latency.lhcmn.triumf.ca	psb02-gva.cern.ch	psl01-gva.cern.ch	psonar9-191-2.fnal.gov	psonar9.fnal.gov	t1-plsn1.jinr-t1.ru	t1-plsn2.jinr-t1.ru	t1-ps01-bw.gridpp.rl.ac.uk	t1-ps01-lt.gridpp.rl.ac.uk		
btw-bw.grid.kiae.ru																														
btw-lat.t1.grid.kiae.ru																														
ccperfsnar1.in2p3.fr																														
ccperfsnar2.in2p3.fr																														
lhcmn.bnl.gov																														
lhcperfmn.bnl.gov																														
perfsnar-bandwidth.grid.surfsara.nl																														
perfsnar-de-kit.gridka.de																														
perfsnar-ow.cnaif.infn.it																														
perfsnar-ps.cnaif.infn.it																														
perfsnar-ps.ndgf.org																														
perfsnar-ps2.ndgf.org																														
perfsnar2-de-kit.gridka.de																														
perfsnar2.lhep.ac.cn																														
plcperfsnar-bandwidth.plc.es																														
plcperfsnar-latency.plc.es																														
ps-bandwidth.lhcmn.triumf.ca																														
ps-gsd01.sdfarm.kr																														
ps-gsd02.sdfarm.kr																														
ps-latency.lhcmn.triumf.ca																														
psb02-gva.cern.ch																														
psl01-gva.cern.ch																														
psonar9-191-2.fnal.gov																														
psonar9.fnal.gov																														
t1-plsn1.jinr-t1.ru																														
t1-plsn2.jinr-t1.ru																														
t1-ps01-bw.gridpp.rl.ac.uk																														
t1-ps01-lt.gridpp.rl.ac.uk																														

perfSONAR Dashboard

The screenshot displays the perfSONAR Dashboard interface. The main content area is titled "WLCG ATLAS Latency" and features a heatmap where rows represent various sites and columns represent different test configurations. The heatmap cells are color-coded: green for good performance, yellow for average, and red for poor performance. A red circle highlights a specific column in the heatmap, which corresponds to the "perfSONAR pSConfig" entry in the sidebar menu.

The sidebar menu on the left includes the following items:

- Home
- Starred
- Dashboards
- Playlists
- Snapshots
- Library panels
- Public dashboards
- Explore
- Metrics
- Alerting
- Alert rules
- Contact points
- Notification policies
- Silences
- Groups
- Admin
- Connections
- Add new connection
- Data sources
- Administration
- General
- Stats and license
- Default preferences
- Settings
- Organizations
- Plugins and data

The top navigation bar includes a search field "Search or jump to...", a share button, and a refresh button. The bottom right corner of the dashboard shows the text "#53 LHCOPN/LHCONE Mtg".

Analytics

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, 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.)

Summary

- Updates to perfSONAR and OSG/WLCG network measurement platform
 - perfSONAR 5.2 is coming with new features.
 - New infrastructure monitoring and dashboards will appear in production soon
 - Simplified deployment models are being evaluated
- 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.

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](#)



Open Science Grid



WLCG
Worldwide LHC Computing Grid




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>

perfSONAR Infrastructure Evolution

181
Active perfSONARs

68
Communities in Use



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

Backup Slides Follow

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	Personar	bad owd measurements	*
Networking	Personar	large clock correction	*
Networking	Personar	complete packet loss	*
Networking	Personar	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

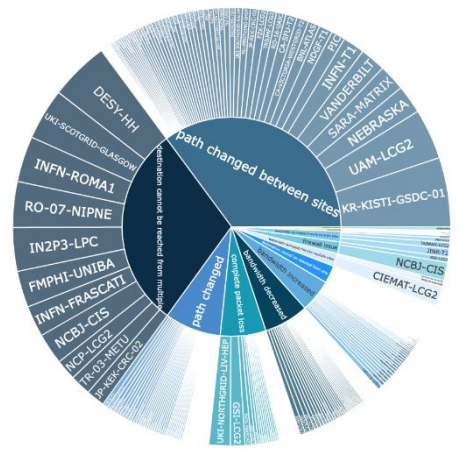
Packet loss

- ▶ 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



Search & Explore the Networking Alarms

03/04/2023 → 03/05/2023

Rounded to the day

Search for a site

Search for an event type


List of alarms

BANDWIDTH DECREASED

from	to	src_site	dest_site	ipv	ipv6	last3days_avg	%change	alarm_link
2023-02-12 04:08	2023-03-05 04:08	AGLT2	RRC-KI-T1	ipv4	false	96	-64	VIEW
2023-02-12 04:08	2023-03-05 04:08	IN2P3-CC	FZK-LCG2	ipv4	false	1008	-74	VIEW

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



DC24

WLCG 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.

Tools and Applications for Network Data

- To organize access to all the various resources we recommend using our Toolkitinfo page: <https://toolkitinfo.opensciencegrid.org/>
- Reminder: we already have Kibana dashboards looking at
 - [Bandwidth](#)
 - [Traceroute](#)
 - [Packetloss](#) / [Latency](#)
 - [Infrastructure](#)
- For this meeting we want to update our recent work towards a user subscribable alerting and alarming service
 - User interface to subscribe is **AAAS** (ATLAS Alerting and Alarming Service)
 - Tool to explore alerts is **pS-Dash** (Plotly base perfSONAR dashboard UI tool)



Alarms & Alerts Service

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

(Uses EDUGain/InCommon)

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



Open Science Grid



WLCG
Worldwide LHC Computing Grid



The Alerting and Alarming Tools Subscription Interface

Alarms

- Analytics
- Networking
 - Perfsonar
 - Infrastructure
 - RENs
 - path changed
 - Sites
 - destination cannot be reached from multiple
 - destination cannot be reached from any
 - bandwidth increased from/to multiple sites
 - bandwidth decreased from/to multiple sites
 - bandwidth increased
 - bandwidth decreased
 - source cannot reach any
 - high packet loss on multiple links
 - high packet loss
- Virtual Placement
- SLATE
- WFMS

Heartbeats

- SLATE

Current Subscriptions

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

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

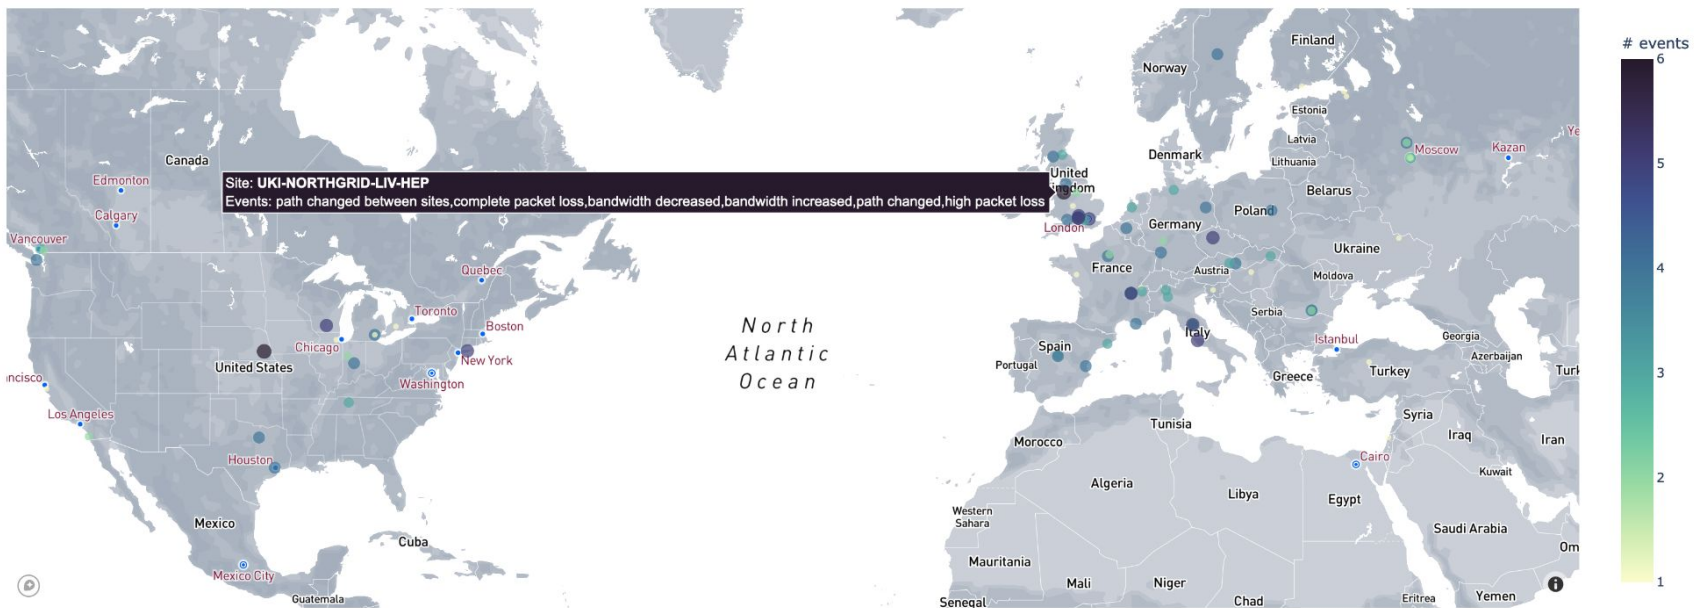
- ▶ bandwidth decreased (from/to multiple sites)
- ▶ bandwidth increased (from/to multiple sites)


pSDash (perfSONAR Dashboard)



Selected site: UKI-NORTHGRID-LIV-HEP

Alarms reported in the past 24 hours (2023-03-05 19:00 UTC)





Search & Explore the Networking Alarms

03/04/2023 → 03/05/2023
 Rounded to the day

Search for a site

Search for an event type

List of alarms

BANDWIDTH DECREASED

from	to	src_site	dest_site	ipv	ipv6	last3days_avg	%change	alarm_link
2023-02-12 04:08	2023-03-05 04:08	AGLT2	RRC-KI-T1	ipv4	false	96	-64	VIEW
2023-02-12 04:08	2023-03-05 04:08	IN2P3-CC	FZK-LCG2	ipv4	false	1008	-74	VIEW

WLCG perfSONAR Path Statistics

We uniquely identify each traceroute (route **IP** path) with a SHA1 hash.

route-sha1	
count	19995.000000
mean	19.911678
std	43.373343
min	1.000000
25%	2.000000
50%	4.000000
75%	12.000000
max	377.000000

5264 links tested Link="hop" (IP-to-IP)

4415 traversed nodes Node="router"

Statistics on the left concern all the "paths" we are tracking with about 20K unique paths found

About 50% of src-dest pairs have 4 or less paths.

AS (Autonomous System) Path Changed

NOTE: Paths denoted by route IP are too noisy; instead use AS number

ASN sequence

[7896, 7896, 293, 293, 293, 293, 293]

[7896, 7896, 293, 293, 293, 293, 293, 43]


[7896, 7896, 7896, 7896, 57, 57, 57, 293, 293, 293, 293, 293, 43]

Reduced ASNs

[7896, 293]

[7896, 293, 43]

[7896, 57, 293, 43]



Example: LHCOPN/LHCONE Load Balancing

Source
INFN-T1
 2001:760:4205:254::11

Destination
SARA-MATRIX
 2001:610:108:203a::31

Total number of traceroute measures: 280
 Other networking alarms: None found

BASELINE PATH

Taken in 50% of time Always reaches destination: NO

137 513 1103

ALTERNATIVE PATHS


Taken in 50.0% of time Always reaches destination: NO

137 20965 1103

Taken in 0.0% of time Always reaches destination: NO

137 20965

AS numbers for every hop and the frequency of their occurrences at each position (source and destination not included)



At position	Typically goes through	Changed to
3	513 CERN, CH	20965 GEANT The GEANT IP Service, NL

At position	Typically goes through	Changed to
4	513 CERN, CH	20965 GEANT The GEANT IP Service, NL

At position	Typically goes through	Changed to
5	513 CERN, CH	20965 GEANT The GEANT IP Service, NL

At position	Typically goes through	Changed to
6	1103 SURFNET-NL SURFnet, The Netherlands, NL	20965 GEANT The GEANT IP Service, NL

Example: LHCOPN Alternate via ESnet

USCMS-FNAL-WC1->CERN-PROD

Source

USCMS-FNAL-WC1

131.225.205.23

Destination

CERN-PROD

128.142.208.134

Total number of traceroute measures: 248

Other networking alarms: None found

BASELINE PATH

Taken in 99% of time

Always reaches destination: NO

3152 513

ALTERNATIVE PATHS

Taken in 0.0% of time

Always reaches destination: NO


3152

Taken in 1.0% of time

Always reaches destination: NO

3152 291 293 513

AS numbers for every hop and the frequency of their occurrences at each position (source and destination not included)



At position	Typically goes through	Changed to
3	513 CERN, CH	291 ESNET-EAST, US

At position	Typically goes through	Changed to
4	513 CERN, CH	293 ESNET, US

At position	Typically goes through	Changed to
5	513 CERN, CH	293 ESNET, US

Example: FNAL Incident (BW drop)


USCMS-FNAL-WC1 to NDGF-T1

Source
USCMS-FNAL-WC1

Destination
NDGF-T1

Change: -100%

Total number of throughput measures: 71
Other networking alarms
| High packet loss: 2 | High packet loss on multiple links: 1 |



push	MA	src	dest	src_host	dest_host	ipv6	src_site	src_VO	dest_site	dest_VO	src_production	dest_production	timestamp	throughput	pair	dt	MBps
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1661561263000	1048161726	131.225.205.23->109.105.124.88	2022-08-27T00:47:43	1048.16
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1661618925000	952946516	131.225.205.23->109.105.124.88	2022-08-27T16:48:45	952.95
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662626239000	1045220096	131.225.205.23->109.105.124.88	2022-09-08T08:37:19	1045.22
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1661678463000	1072068304	131.225.205.23->109.105.124.88	2022-08-28T09:21:03	1072.07
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662439905000	1072905581	131.225.205.23->109.105.124.88	2022-09-06T04:51:45	1072.91
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1661659004000	1073324325	131.225.205.23->109.105.124.88	2022-08-28T03:56:44	1073.32
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662672411000	1074163359	131.225.205.23->109.105.124.88	2022-09-08T21:26:51	1074.16
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662902418000	1071231326	131.225.205.23->109.105.124.88	2022-09-11T13:120:18	1071.23
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662093921000	1085912472	131.225.205.23->109.105.124.88	2022-09-02T04:45:21	1085.91
false	131.225.205.23	131.225.205.23	109.105.124.88	psonar3.fnal.gov	perfnsonar-ps2.ndgf.org	false	USCHS-FNAL-WC1	UNKNOWN	NDGF-T1	ATLAS	true	true	1662696230000	1068710540	131.225.205.23->109.105.124.88	2022-09-09T04:03:50	1068.71

Example: Fail-over to Commodity Network

Nebraska -> RAL-LCG2

Source

Nebraska

2600:900:6:1102:2eea:7fff:fef5:d140

Destination

RAL-LCG2

2001:630:58:1820::82f6:b06d

Total number of traceroute measures: 280

Other networking alarms: None found

BASELINE PATH

Taken in 99% of time

Always reaches destination: NO

7896 11537 20965 786

ALTERNATIVE PATHS

Taken in 1.0% of time

Always reaches destination: NO

7896 6939 786

Taken in 0.0% of time

Always reaches destination: YES

7896 7029 6461 3257 1299 786

AS numbers for every hop and the frequency of their occurrences at each position (source and destination not included)




At position	Typically goes through	Changed to
4	11537 INTERNET2-RESEARCH-EDU, US	7029 WINDSTREAM, US

At position	Typically goes through	Changed to
7	11537 INTERNET2-RESEARCH-EDU, US	6939 HURRICANE, US

At position	Typically goes through	Changed to
7	6461 ZAYO-6461, US	6939 HURRICANE, US

Challenges and Ongoing Work

Paths differ significantly




Correlating Tests with Paths: Two Timescales

- Latency tests: **every minute**
- Throughput tests: **6-24 hours**
- ◆ Traceroute tests: **every 10 min**


The problem: All tests run independently

How do we combine them?



Connecting Throughput to Traceroute

Can we consider the trace routes closest in time to the throughput records?




Our starting choice: Use **both** tracepaths (just before; just after) as valid paths and attribute BW to both.



Have to see if this is superior to just using the last measured route before the measurement...

Attaching Throughput Results to Sets of Routers/Links

Each colored box represents a specific router along the path



Example Throughput Attribution by Router




Each router on the path gets the closest (in time) throughput values

Checking Router Results vs Time

Max bandwidth seen: 16 307 Mb



Max bandwidth seen: 1868 Mb




Look for a down trend
(threshold below 10% of the max throughput)


Is everything OK
with/around that router?

Initial Example Result: One Router; Throughput vs Time

2001:630:0:9011::189




Each **point** represents the throughput values collected when the node was on the path




Other Activities / Plans

- Working to organize and annotate our data for ML/AI work (Petya Vasileva)
- Working with the RNTWG (see previous RNTWG update talk) on identifying and monitoring network traffic details via the SciTags initiative.
- Exploring other network monitoring activities in the perfSONAR space including [ARGUS](#)
- Planning to augment [WLCG-CRIC](#) (yesterday's discussion) network meta data (which paths/networks are LHCOPN / LHCONE / Research&Education / Commercial)




Distributions of Throughput


Distribution of the minimum



Distribution of the average



Distribution of the maximum



Distribution of the mode




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. (See coming [talk](#) :))



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**

