



perfSONAR IPv6 Monitoring Update

Shawn McKee, Marian Babik
on behalf of WLCG Network Throughput WG



Outline

- OSG/WLCG Network Monitoring and WLCG Network Throughput WG
- perfSONAR community updates
- perfSONAR infrastructure status
- Plans
- Summary





OSG/WLCG networking projects

There are 4 coupled projects around the core **OSG Net Area**

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

OSG Networking Components



- The [4.3 release](#) moved to Python3, added PWA support for archivers and some new tools
- **4.3.4 bug fix** followed
 - Fixes security issue with SSRF (server side request forgery) found by BNL & SLAC
 - Has fix to memory leak in RabbitMQ archiver plus other minor fixes
- [4.4 release](#) is in beta testing since May
 - Adds loopback tests, AS capability in limits and updates RMQ archiver
 - Fixes number of bugs in pscheduler, psconfig and PWA
- Traces were updated in most WLCG meshes to now run every 10 minutes
- perfSONAR session at [TechExtra Nov 2nd](#)
 - Demo of the ELK stack capabilities, PWA, WIFI and Cloud monitoring
- Michael Johnson, perfSONAR developer & responsible for PWA, **has left the project in Feb 2021**. John Grigutis will take over.
- RNP (the Brazilian research network) [joined the perfSONAR collaboration](#)




perfSONAR deployment


243 Active perfSONAR instances - **207 production endpoints** - T1/T2 coverage
- Continuously testing over 5000 links - testing coordinated and managed from central place

243 Registered perfSONARs
243 Active perfSONARs


perfSONAR Toolkit NIC Speed Histogram




perfSONAR Toolkit Host Memory Histogram



NIC speed




perfSONAR distributions



pS NIC MTU



pS NIC Speed vs Time



Our global toolkit deployment has a range of systems in terms of age and capability

[Dashboard in ELK](#)


Sites should consider upgrading not just the perfSONAR software but also the underlying **hardware**, as nodes become too old or are unable to test at the site storage speed.

LHCOPN 23rd March 2021

OPN Mesh Config - OPN IPv6 Bandwidth - Throughput

Throughput ≥ 1 Gbps Throughput < 1 Gbps Throughput $\leq .5$ Gbps


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



OPN Mesh Config - OPN Latency - Loss

Loss rate is $\leq 0.001\%$ Loss rate is $> 0.001\%$ Loss rate

! Found a total of 2 problems involving 2 hosts in the grid




LHCONE - 23rd of March 2021

LHCONE - LHCONE Bandwidth IPv4 - Throughput

Throughput ≥ 1 Gbps Throughput < 1 Gbps Throughput $\leq .5$ Gbps Unable to find

! Found a total of 41 problems involving 29 hosts in the grid

Partner



Test **Aust->ASGC**
initiated/stored by **Aust**

ATLAS Great Lakes Tier-2 UM
Australia-ATLAS

Test **ASGC->Aust**
initiated/stored by **Aust**




100Gbps Testing 23rd Mar 2021

WLCG 100G Mesh - WLCG 100G IPv4 Bandwidth - Throughput

Throughput >= 1Gbps Throughput < 1Gbps Throughput <= .5Gbps Unable to fill

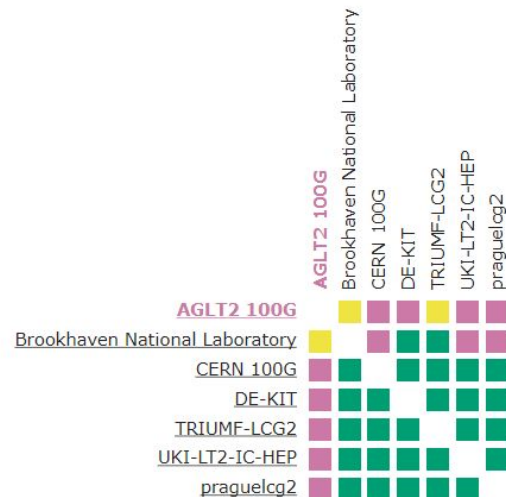
! Found a total of 2 problems involving 1 host in the grid



WLCG 100G Mesh - WLCG 100G IPv6 Bandwidth - Throughput

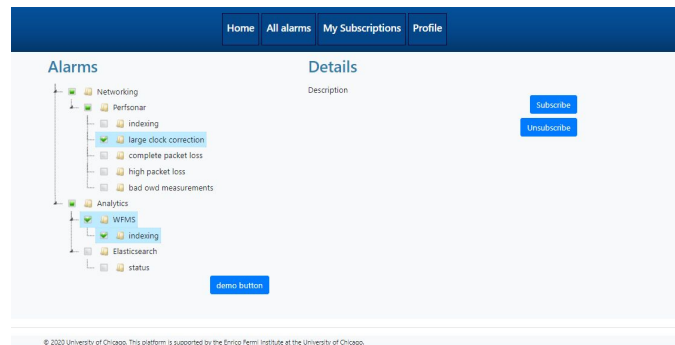
Throughput >= 1Gbps Throughput < 1Gbps Throughput <= .5Gbps Unable to fill

! Found a total of 2 problems involving 1 host in the grid




Plans and Near-term Activities

- The focus is on analytics and getting issue identification into the production pipeline
 - Students from IRIS-HEP and SAND are working on this
 - Prototype user-alarm-subscription developing ->
- We already have Kibana dashboards looking at
 - [Bandwidth](#)
 - [Traceroute](#)
 - [Packetloss](#) / [Latency](#)
 - [Infrastructure](#)
- To organize access to all the various resources we have developed we created <https://toolkitinfo.opensciencegrid.org> (Try it; give feedback!)
- **Upcoming: WLCG Data Challenges!!** To support these we need additional network monitoring (near realtime net use by site).
 - We are working with the CRIC team to add a NetMonitoringURL sites can provide



Summary

- OSG in collaboration with WLCG are operating a comprehensive network monitoring platform
- Providing feedback to LHCOPN/LHCONE, HEPiX, WLCG and OSG communities
- IRIS-HEP and SAND are already contributing to the R&D in the network area
- Current focus is on analytical studies that would lead to a production level network alerting service
- We also expect 100 Gbps mesh to grow with potential to contribute to the network data challenges and network performance studies



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](#)

References

- OSG/WLCG Networking Documentation
 - <https://opensciencegrid.github.io/networking/>
- perfSONAR Infrastructure Dashboard
 - <https://atlas-kibana.mwt2.org:5601/s/networking/goto/9911c54099b2be47ff9700772c3778b7>
- perfSONAR Dashboard and Monitoring
 - <http://maddash.opensciencegrid.org/maddash-webui>
 - https://psetf.opensciencegrid.org/etf/check_mk
- perfSONAR Central Configuration
 - <https://psconfig.opensciencegrid.org/>
- Toolkit information page
 - <https://toolkitinfo.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/>

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. (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**