

Jisc perfSONAR Activities

Tim Chown (Jisc)

tim.chown@jisc.ac.uk

5th perfSONAR User Workshop

Paris, 6 May 2026

Topics

What are we currently working on?

- High level engagement
 - Hosting meshes for Janet communities
 - Time synchronisation
 - Low latency networking
 - Working towards easy-to-deploy small nodes
 - perfSONAR for performance testing and tuning
 - 100G+ testing
 - Testing MTU – towards jumbo frames (9000 MTU)
-
- We submit many issues to the [perfSONAR Git repo](#) – many of those are still open
 - Reviewing issues for a spring clean with the developers

High level view

Activities

- Engagement with Jisc members and communities
 - Promoting use of perfSONAR
 - Inclusion in events, communities like Research Network Engineering (RNE)
- GÉANT GN5-2 WP6 participation
 - Tim, Raul and Duncan (total 30% FTE)
- Jisc-WLCG-dev meetings
 - Monthly calls focused on GridPP requirements – supporting big science data transfers
 - Understanding (throughput) performance issues
- Monitoring low latency networking use cases
 - Recently a focus on latency for networked music performance (NMP)
 - Nice to support arts use cases too

Hosted meshes

For a variety of use cases

- Includes
 - UK test – open to any Janet site, mainly GridPP – main mesh
 - Time synchronisation – improving OWD accuracy
 - South Africa – originally to check against EBI/EMBL Globus performance
 - XLZD – servers in Boulby mine
 - FE college performance case - small nodes on Janet MRS routers
 - ...
- See <https://ps-mesh.perf.ja.net/grafana/dashboards/f/ddl3qr8ufgagwe/perfsonar-psconfig>

UK test mesh

Open to any Janet-connected site

- Mainly GridPP (UK CERN experiment) sites, but can also be HPC, IRIS, ...
- A mix of 10G and 100G servers, two are on LHCONE
- Each maintained by the local admin (who have other priorities too)
- Includes two Jisc reference servers, ps-london (100G) and ps-slough (10G)
- Largely in a good shape in terms of %age of results being published
- Many 'unexpected' results to drill down into with the server admins
 - This is currently our focus
- See <https://ps-mesh.perf.ja.net/grafana/d/aa15bb77-25d6-53d1-b3c7-bf51b6417afa/jisc-uk-tests?orgId=1&from=now-24h&to=now&timezone=browser>

Time Synchronisation Mesh

Working to improve accuracy (and avoid negative latencies)

- [Mesh](#) open to those interested in accurate time sync
 - How can we get good synchronisation
 - NTP – e.g., chrony config, bearing in mind server-side load
 - PTP – how much better is it?
 - If we get good synchronisation, what does it tell us?
 - Includes Asiera (formerly HEANet), SURF
 - Others welcome
 - Just email netperf@jisc.ac.uk



Asiera - effect of PTP on one node




Chrony Changes

```
server ntp5.ja.net minpoll 5 maxpoll 5 iburst
server ntp6.ja.net minpoll 5 maxpoll 5 iburst
lock_all
# 2^5 = 32s
vm_swappiness=1 # sysconfig
```

```
server ntp5.ja.net iburst
server ntp6.ja.net iburst
```

default $2^{10}s = 1024s$



Low latency networking

"It's not just megabytes, it's milliseconds" – Claudio Allocchio

- Recent specific example is networked music performances (NMP)
- End-to-end round-trip delay ideally under 30ms
- **Need to monitor to allow both technical staff and non (network) technical users to understand how the network affects the quality of their sessions**
 - perfSONAR gives measurements over time, not just when issues are experienced
- Benefits from accurate time sync
- **Would benefit from a new perfSONAR test – iperf2 rpm (responsiveness under load) test**
 - <https://datatracker.ietf.org/doc/draft-ietf-ippm-responsiveness/>
 - Sadly, perfSONAR only includes a very old iperf2 version
- Currently having monthly calls with NMP users (RAM, Glasgow, Trieste)
 - Have an NMP session at Jisc Networkshop and a network latency BoF at TNC next month

Towards easy-to-deploy small nodes

We want to offer a cheap, lightweight perfSONAR option for members

- Implies using a small form factor node (as per PMP project) or container
 - 1G throughput tests enough for many cases, 10G desirable
- We just ordered AMD NUCs (Ryzen 7, 32GB RAM, 512G SSD), also looking at RPi specs
 - NUCs support USB3.2 for 10G; will test, need to check we have copper switch interfaces
- Ideal for drop-in at sites
 - Can plug directly into our Managed Router Service (MRS) router ports
 - **We will want them deployed inside member sites for NMP monitoring**
- **BUT** need to consider security
 - Potential issues with a Jisc-configured device deployed inside a member's network
 - **Thus want a simple set of build instructions for Linux + perfSONAR as a minimal testpoint**
 - We provide the config file and then host the results archive and Grafana views

perfSONAR for performance testing and tuning

Leveraging pscheduler options

- Interest for many use cases in determining best tuning for servers
 - *Focused on optimising data transfers, especially with high RTT paths*
 - Can test with CLI, or change config and look at plots over time
- Lots of very useful **pscheduler** options:
 - *Congestion Control Algorithm (CCA)*: --congestion - Reno, CUBIC, H-TCP, BBR,...
 - *Maximum Transmission Unit (MTU)*: --mss (actually TCP maximum segment size)
 - Currently a CERN campaign for sites to use 9000 MTU
 - *TCP window size*: --window size – important for long fat pipes
 - *Number of streams*: --parallel
 - *Pacing*: --bandwidth, e.g., --b 10G
- See https://docs.perfsonar.net/pscheduler_ref_tests_tools.html

100G+ testing

Higher-speed testing – setting 400G+ state-of-the-art

- There is a 100G mesh hosted by the WLCG
 - Has been dormant - we are talking to the WLCG team about rebooting this activity
- Raul has been testing driving links at 400G-800G in collaboration with GÉANT
 - **More on this tomorrow**
 - It seems we should get that experience now, before 400G is more common
 - Both single stream capability testing, and 'fill the link' testing
- May need different / additional tools – Raul has looked at:
 - iperf3 - <https://github.com/esnet/iperf>
 - qperf - <https://github.com/linux-rdma/qperf/> (support RMDA)
 - ib-write - <https://github.com/linux-rdma/perftest> (ditto)
- **These would need integration into perfSONAR** – some work is underway

Testing MTU

Towards jumbo frames: 9000 MTU

- Interest from WLCG/GridPP in using 9000 MTU
 - Will be important for data rates when HL-LHC arrives
- Raul has patched **tracemtu** – new code is now shipping in 5.2.4
 - We can now run tests between sites
 - **Requires UDP ports 33434-33634 to be open** ([see pS docs](#))
- Next step is to create an MTU mesh
 - Include sites who believe they are configured for 9000 MTU
 - Green = 9000, Orange = 1501-8999, Yellow = 1500, Red = not working
 - **Raul talking to Andy Lake (ESnet) on how to extract and present the necessary data**

Open issues – spring clean time 😊

We submit quite a few issues, and then try to help devs to address them

- Adding a legend for meshes, as we had with MaDDash ([issue](#))
- perfSONAR-toolkit-sysctl and network-manager clash ([issue](#))
- Recording IP version used ([issue](#))
- MTU tooling – in part done with new tracemtu ([issue](#))
- Hiding unused test types from results pages, e.g., for latency mesh ([issue](#))
- PTP support in perfSONAR ([issue](#))
- Adding more tools – a more advanced http test ([issue](#)), udpst ([issue](#)), iperf2 rpm ([issue](#))
- Powstream tool runaway bug (still happening?) ([issue](#))
- Trace mesh showing green boxes when the test fails ([issue](#))
- Services aborting due to database error (still happening?) ([issue](#))
- Loss charts fixed at 0-100% even when the loss is lower ([issue](#))

Thank you

Tim Chown

Network Development Manager

tim.chown@jisc.ac.uk

 help@jisc.ac.uk

 0300 300 2212

 jisc.ac.uk



Except where otherwise noted, this work is licensed under CC-BY-NC-ND.



 [@jisc.bsky.social](https://twitter.com/jisc.bsky.social)

 [@jiscsocial](https://www.instagram.com/jiscsocial)

 [linkedin.com/company/jisc](https://www.linkedin.com/company/jisc)

Jisc