TNC 2026:
Network Telemetry BoF

Helsinki, 9 June 2026

Tim Chown (Jisc)
Claudio Allocchio (GARR)
Lætitia Delvaux (PCSS)
Andy Lake (ESnet)

It’s (partly) about queues!

Welcome!

This ‘Birds of a Feather’ will look at a variety of perspectives on network latency.

Please feel free to ask any questions as we go – we want to maximise discussion time

(Rough) Agenda

Topic	Speaker
Introduction *A 'lie of the land' mentimeter*	Tim Chown
Which use cases are we interested in? *Example: networked music performance*	Claudio Allocchio
What tools do we currently use to measure latency? *A review of tooling – perfSONAR, TimeMap and more* *The importance of accurate time*	Lætitia Delvaux Tim Chown
Aggregating measurements *How do we present data to both expert and 'lay' users?*	Andy Lake
Latency-related issues *Bufferbloat, latency spikes/jitter, re-routing, …*	Tim Chown
Open discussion *Which gaps do we need to address? How might we?*	All

A quick menti…

Just three questions

- What type of organisation are you from?
- Which tools do you use to measure network latency?
- What latency-related issues have you experienced?

First talk: Networked music performances

Claudio Allocchio (GARR)

Other use cases? Discuss…

Jisc

Second talk: perfSONAR for measuring latency

Lætitia Delvaux (PCSS)

Related: accurate time

Time synchronisation

An example of what can go wrong…

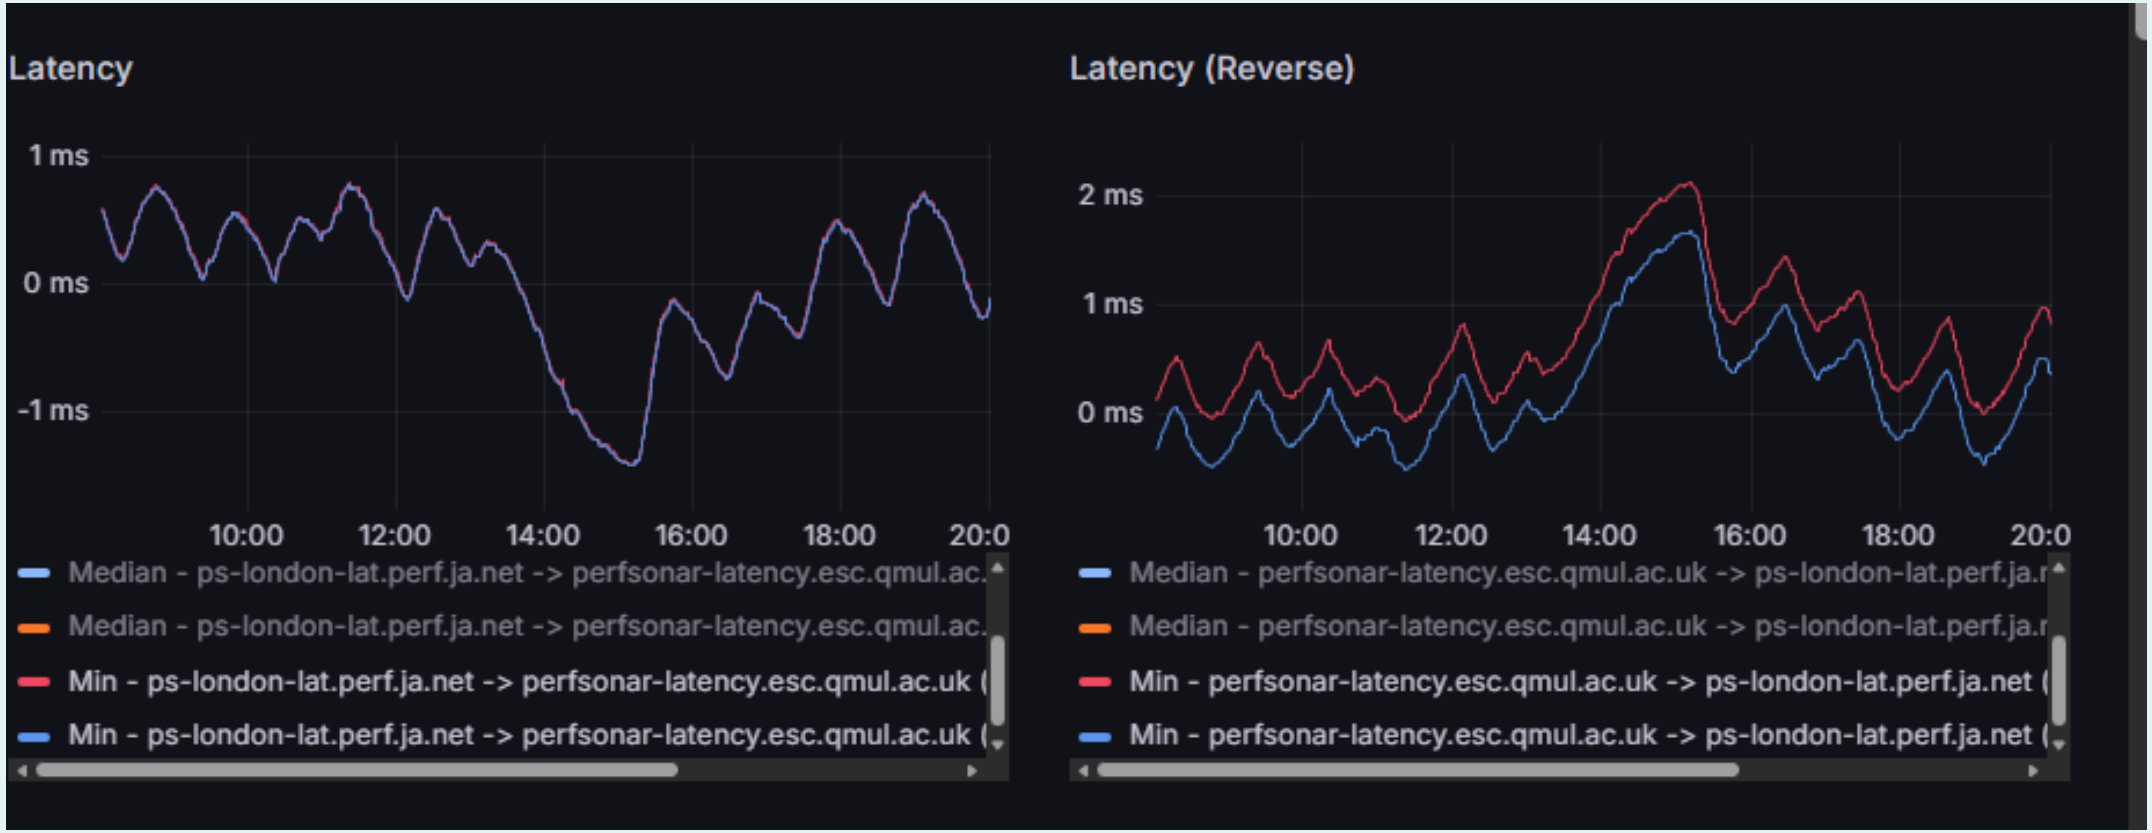

Tweaking chronyd

~20 minute polling

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

default polling at 2^10s = 1024s (17 mins)

Reduce to 32 second polling

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

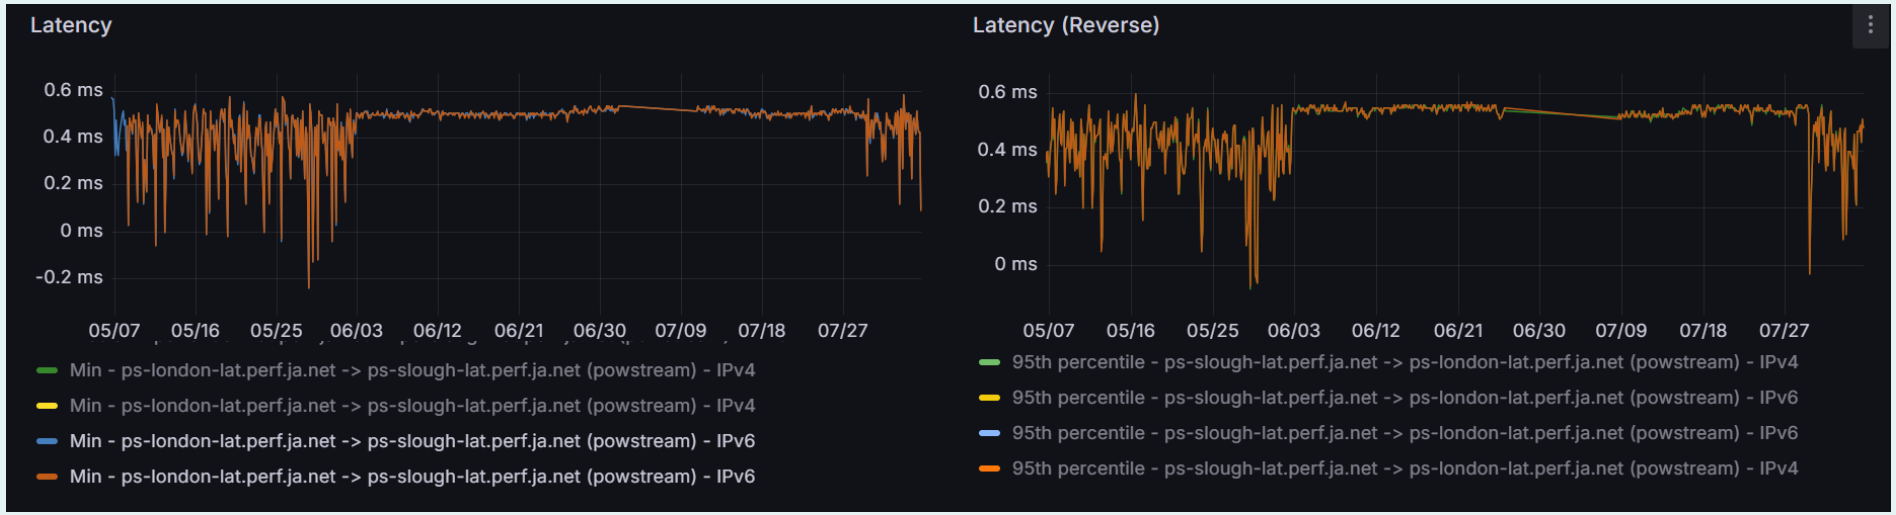

BUT… this will increase load on your servers

Jisc time sync perfSONAR mesh

Exploring accurate time

- Jisc has been chatting to some NRENs about time sync
- Curious as to how far you can go with tuning NTP and using PTP
  - Not looking at ultra-precise/stable T&F
- Running a perfSONAR mesh of interested parties
  - Coordinated by Chris Walker at Jisc
  - Using Jisc perfSONAR archive and viewer
  - If interested email netperf@jisc.ac.uk

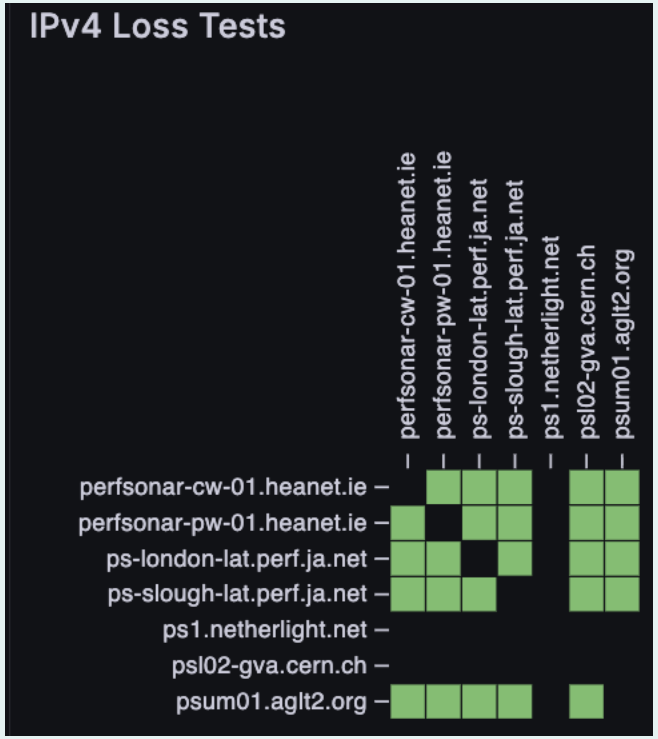

Other latency monitoring tools

TimeMap

Per-hop latency

- Developed within the GÉANT project
  - https://network.geant.org/timemap/
  - https://timemap.geant.org/
- Open source collector
- Requires TWAMP configuration on network routers
- Can be run on NRENs or in campuses
- Multi-domain analysis needs some form of data sharing/aggregation

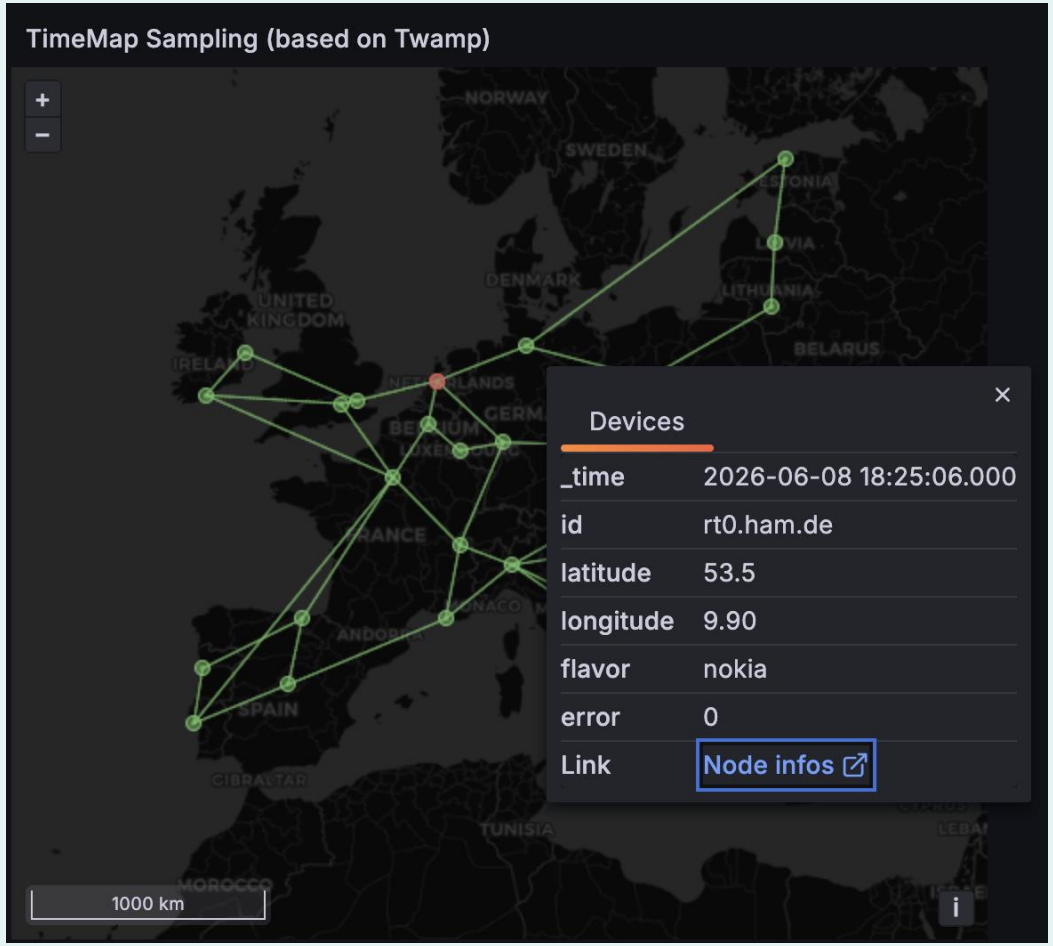

TimeMap: drill-down view

IETF “responsiveness under working conditions”

Recent IETF activity

- A new responsiveness test
- IETF draft
  - https://datatracker.ietf.org/doc/draft-ietf-ippm-responsiveness/
  - Measures responses per minute (RPM)
- Builds link up to “working conditions”
- Then runs HTTP-based tests
- See https://bufferbloat.org/test
- RPM test also available in latest iperf2
  - `iperf -c <server> -i 1 --bounceback`

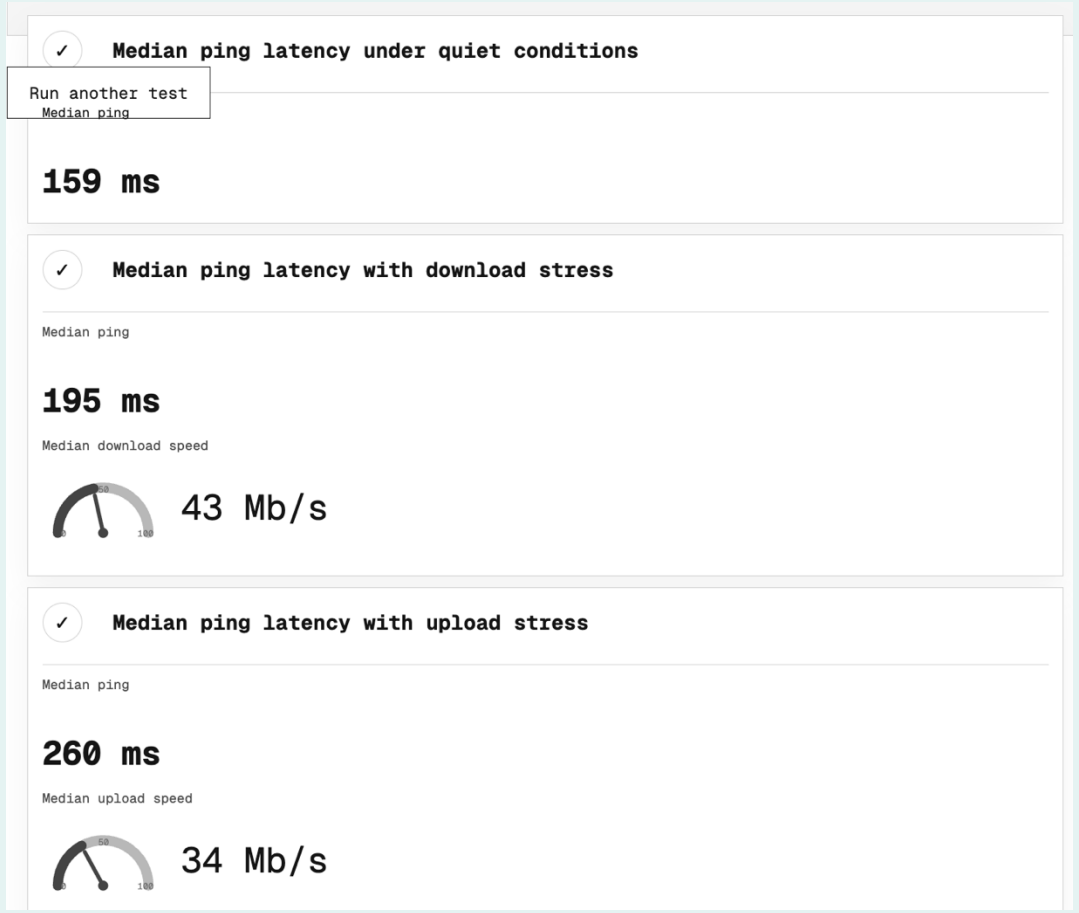

iperf2

```
$ iperf -c iperf-slough-10g.perf.ja.net -i 1 --bounceback
------------------------------------------------------------
Client connecting to iperf-slough-10g.perf.ja.net, TCP port 5001 with pid 3526621 (1 flows)
Write buffer size:  100 Byte
Bursting:  100 Byte writes 10 times every 1.00 second(s)
Bounce-back test (size= 100 Byte) (server hold req=0 usecs & tcp_quickack)
TOS set to 0x0 and nodelay (Nagle off)
TCP window size: 64.0 KByte (default)
------------------------------------------------------------
[  1] local 194.81.18.227%ens2f0 port 59692 connected with 194.82.175.101 port 5001 (bb w/quickack len/hold=100/0) (sock=3) (icwnd/mss/irtt=87/8948/1307) (ct=1.37 ms) on 2026-06-08 17:25:00.603 (BST)
[ ID] Interval       Transfer    Bandwidth       BB cnt=avg/min/max/stdev       Rtry  Cwnd/RTT    RPS
[  1] 0.00-1.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.602/1.271/3.983/0.837 ms    0  174K/1309 us    624 rps
[  1] 1.00-2.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.347/1.293/1.385/0.034 ms    0  209K/1318 us    742 rps
[  1] 2.00-3.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.325/1.278/1.395/0.041 ms    0  209K/1300 us    755 rps
[  1] 3.00-4.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.349/1.257/1.414/0.043 ms    0  209K/1319 us    742 rps
[  1] 4.00-5.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.357/1.287/1.396/0.031 ms    0  209K/1332 us    737 rps
[  1] 5.00-6.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.356/1.291/1.406/0.039 ms    0  209K/1329 us    737 rps
[  1] 6.00-7.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.318/1.256/1.363/0.041 ms    0  209K/1299 us    759 rps
[  1] 7.00-8.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.302/1.263/1.345/0.026 ms    0  209K/1292 us    768 rps
[  1] 8.00-9.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.307/1.290/1.332/0.013 ms    0  209K/1286 us    765 rps
[  1] 9.00-10.00 sec  1.95 KBytes  16.0 Kbits/sec    10=1.289/1.231/1.328/0.035 ms    0  209K/1280 us    776 rps
[  1] 0.00-10.01 sec  19.5 KBytes  16.0 Kbits/sec    100=1.355/1.231/3.983/0.268 ms    0  209K/2973 us    738 rps
[  1] 0.00-10.01 sec BB8(f)-PDF: bin(w=100us):cnt(100)=13:26,14:71,15:2,40:1 (5.00/95.00/99.7%=13/14/40,Outliers=0,obl/obu=0/0)
```

Third talk: Aggregating measurements: MetrANOVA

Andy Lake (ESnet)

Latency-related Issues

Jisc

Bufferbloat

Causes

- Large device buffers
- Congestion happens before loss
- Hosts/applications not aware of on-path buffering

- See https://bufferbloat.org/

Solutions

- Use devices with shallow buffers
  - Loss detected sooner, but with traditional TCP affects throughput badly

  => use delay-aware TCP, e.g., BBR

- Low Latency, Low Loss, and Scalable Throughput (L4S)– uses Explicit Congestion Notification (ECN)
  - Included in iOS

Latency spikes / jitter

Causes

- Congestion
- Other packet loss
- Interface mismatches

Solutions?

- Can we apply “Science DMZ” principles to low latency scenarios?

Science DMZ principles

Should we apply these to low latency?

- Principles:
  - A network architecture with science network distinct from the business network
  - Dedicated systems for data transfer – software, tuning, and hardware
  - Persistent network characteristic monitoring
  - Security policies and enforcement mechanisms tailored for high performance science environments
- See https://fasterdata.es.net/science-dmz/

Challenges?

- In the science data movement case, there are usually a small number of dedicated data transfer nodes on a campus
- But many devices may need low latency
- Might put a limited number of devices such as NMP devices in a “low latency DMZ”?

Rerouting

Causes

- Link failures – failover to resilient link
- BGP changes
- ECMP or similar tools

Solutions?

- Tricky
- Ensure resilient paths have similar latency properties – minimize ‘surprises’

Other issues?

What else should we add?

- Thoughts?

Solutions?

- ?

Open discussion

Improving NREN latency-related services

What gaps remain?

What should we work on together?

Thoughts?