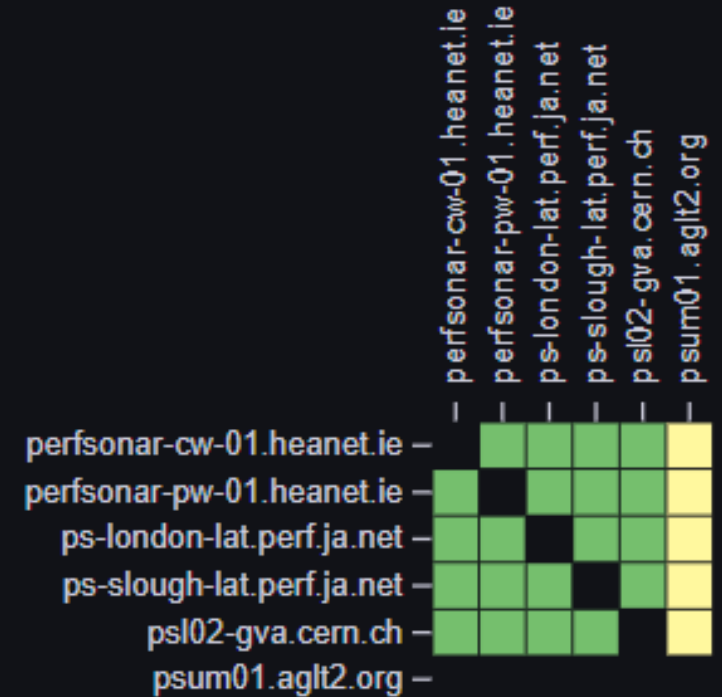


Time Synchronisation and perfSONAR

Christopher J. Walker
 Raul Lopes
 Duncan Rand
 Tim Chown


IPv4 Loss Tests



Overview

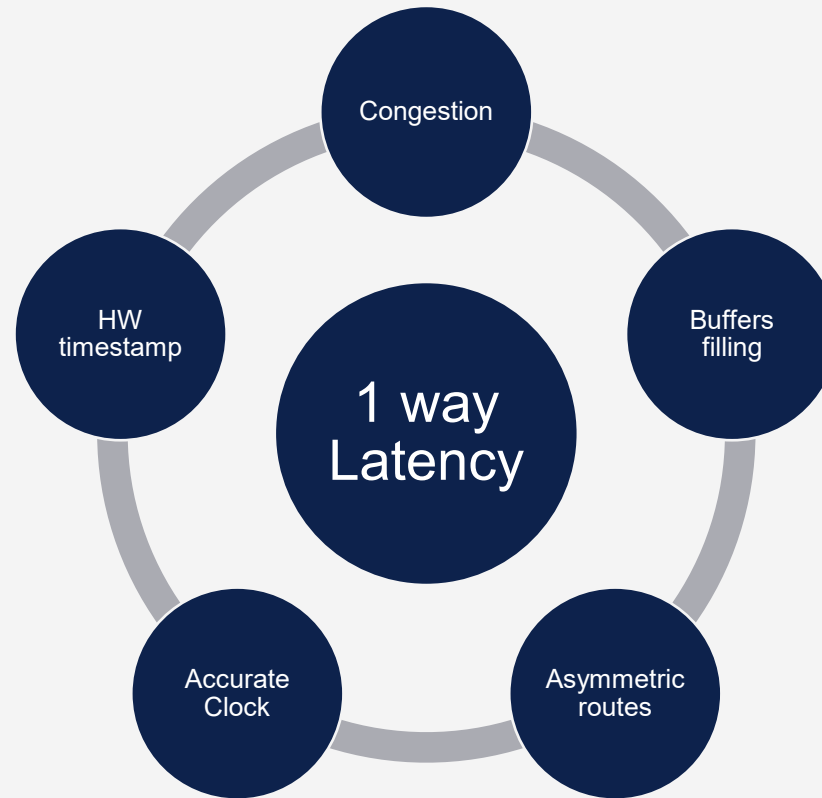
Time synchronisation mesh: motivation

- What can Latency tell us?
 - Do we *need* bandwidth tests to all?
- Why is clock drift a problem?
- How can we improve clock drift
 - NTP (Jisc)
 - Accurate clock
 - Hardware Timestamping
 - PTP (HEANET, CERN, AGLT2)
- OWAMP
 - Will hardware timestamping make a difference



Motivation for accurate timing

Measure Network Performance



What accuracy is required

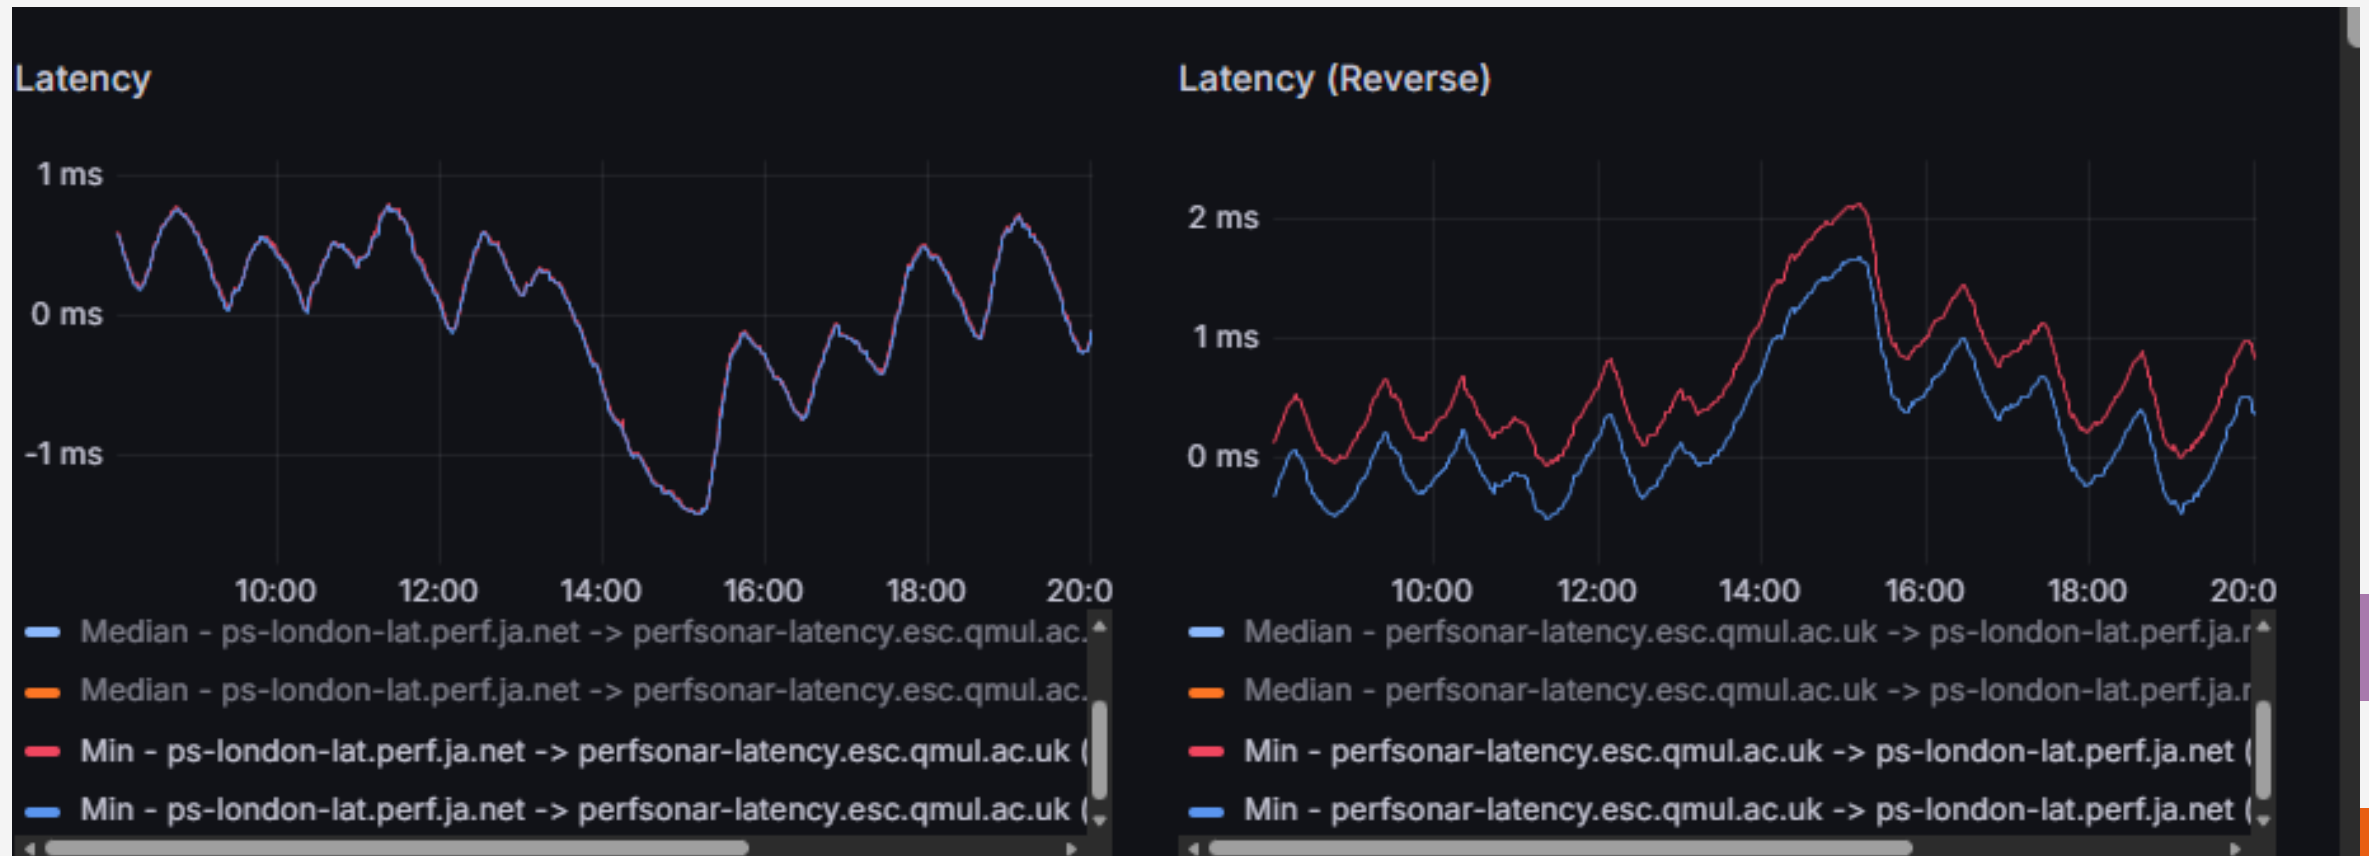
Consider 1250 Byte=10,000bit packet

- Packet quantisation time
 - 10Gbit : 1 μ s
 - 100Gbit : 0.1 μ s
- Clock stability:
 - ~ 500 μ s (0.5ms)
 - 500 packets at 10Gig
 - 5000 packets at 100Gig
- What accuracy do we need?

One Way Latency - Effect of clock drift

- Jisc(London) -> QMUL
 - a mirror image of
- QMUL -> Jisc(London)

- Clock Drift
- IPv4 and IPv6 different routes



Effect of clock offset

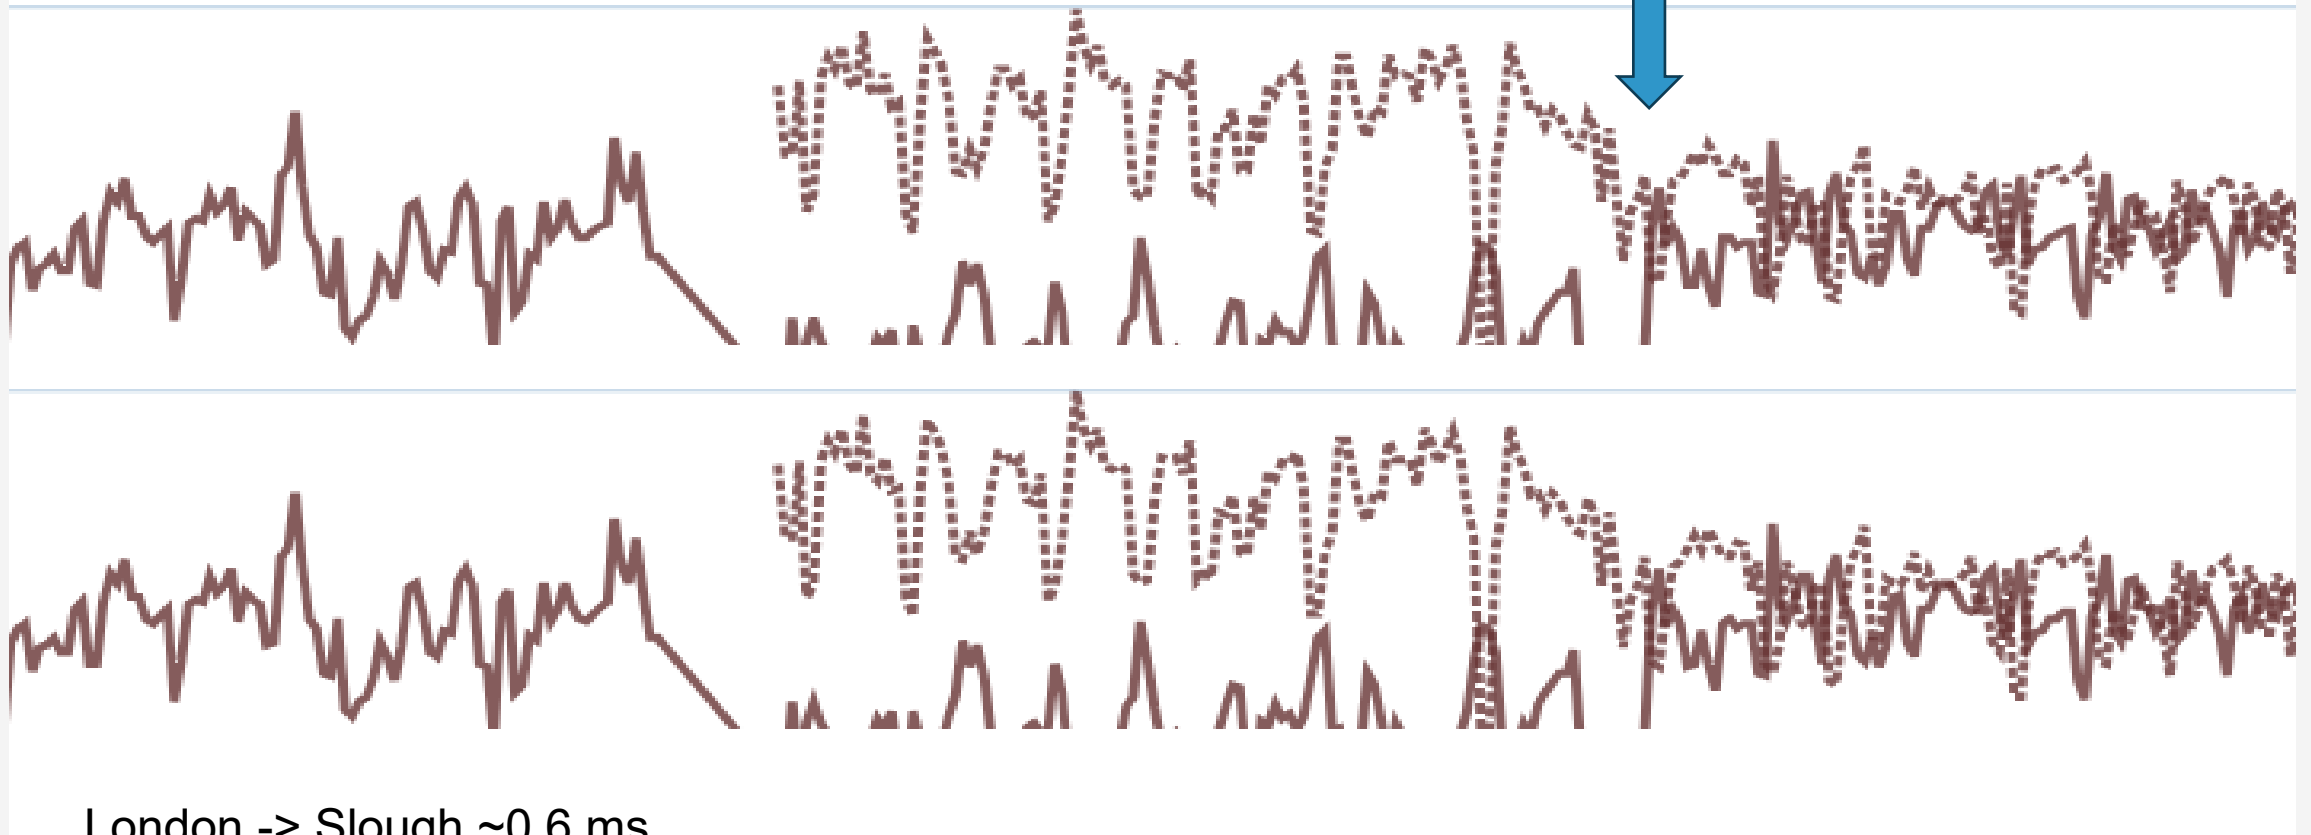
- Latency = $T_2 - T_1$
- Observed latency
 - = $T_2(B) - T_1()$
 - = $T_2 - T_1 + \text{Clock offset.}$



Perfsonar one way latency measurements

Jisc London to/from Slough (1 month)

HW timestamping enabled



Network Card timestamping

RHEL 9 “Basic System Settings - Chrony”

- Network card support?

```
# ethtool -T enp175s0f0np0
```

- Output

```
Time stamping parameters for enp175s0f0np0:
```

```
Capabilities:
```

```
    hardware-transmit
```

```
    hardware-receive
```

```
    hardware-raw-clock
```

```
PTP Hardware Clock: 0
```

```
Hardware Transmit Timestamp Modes:
```

```
    off
```

```
    on
```

Network Card timestamping (2)

- Enabling in **/etc/chrony.conf**

```
# Enable hardware timestamping on all interfaces that
support it.
hwtimestamp *
```

- Checking

```
# chronyc ntpdata
...
Offset      : -0.000132777 seconds
Peer delay   : 0.005870986 seconds
Peer dispersion : 0.000003835 seconds
Response time : 0.000374075 seconds
...
TX timestamping : Hardware
RX timestamping : Hardware
```

Chrony changes

A

- 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

B

- server ntp5.ja.net iburst
- server ntp6.ja.net iburst
- # default $2^{10}s = 1024s$ (2

Latency




— Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4
— Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6
— Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6
— Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4

Latency (Reverse)




— 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4
— 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
— 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
— 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4

Latency



- 95th percentile - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4
- 95th percentile - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4
- 95th percentile - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6
- 95th percentile - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6

Latency (Reverse)



- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4

Packet Loss

Latency



- Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4
- Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv4
- Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6
- Min - ps-london-lat.perf.ja.net -> ps-slough-lat.perf.ja.net (powstream) - IPv6

Packet Loss (Reverse)

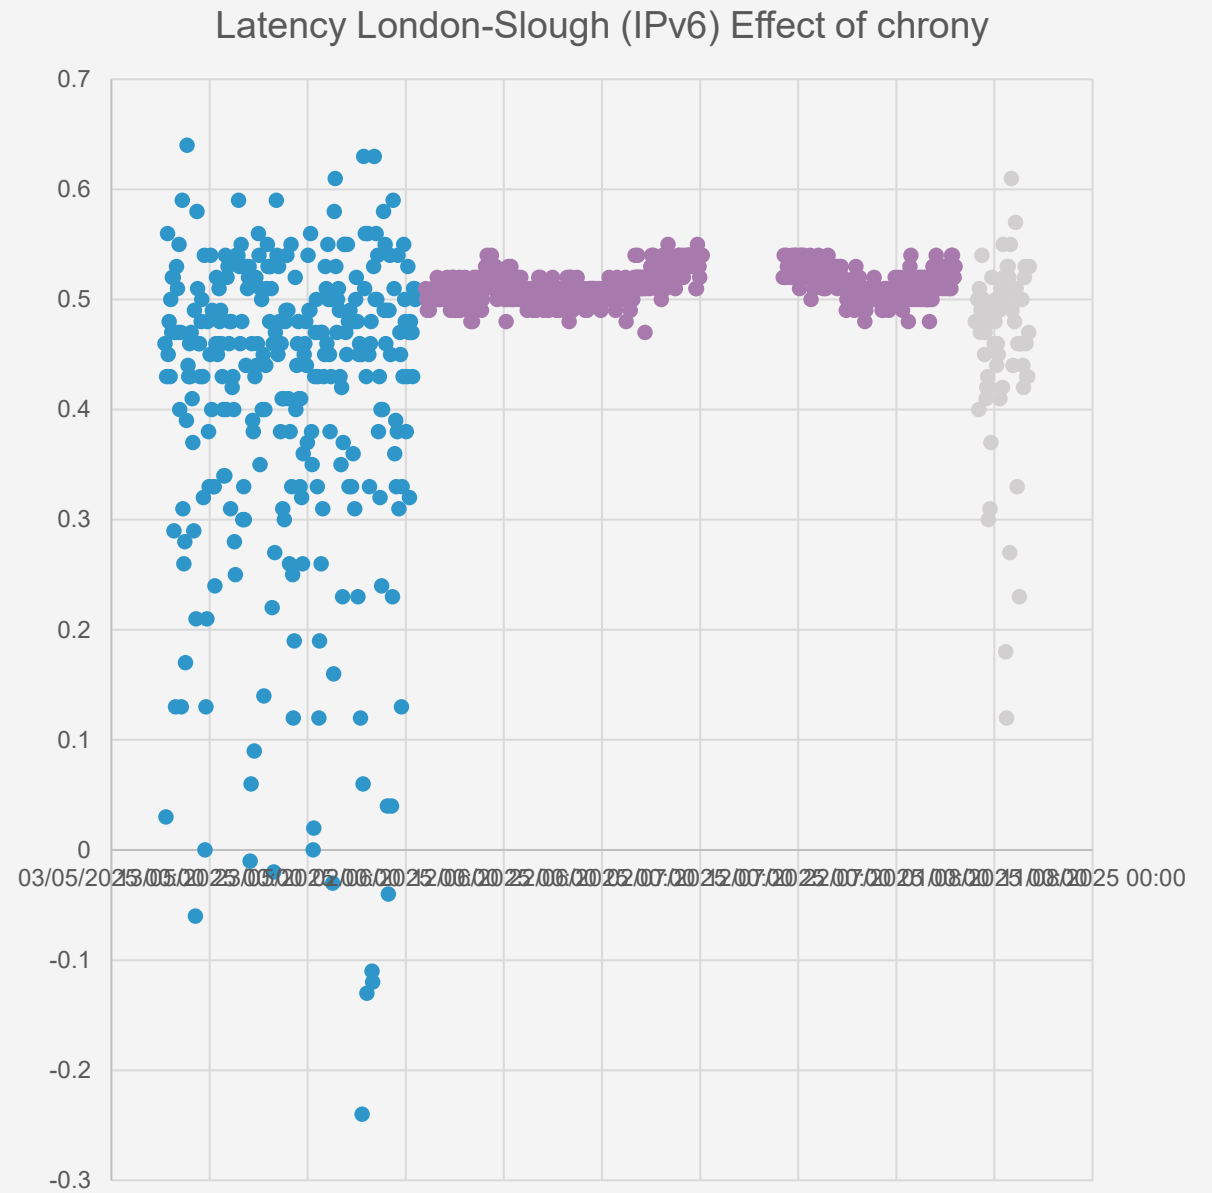
Latency (Reverse)



- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv6
- 95th percentile - ps-slough-lat.perf.ja.net -> ps-london-lat.perf.ja.net (powstream) - IPv4

Minimum latency

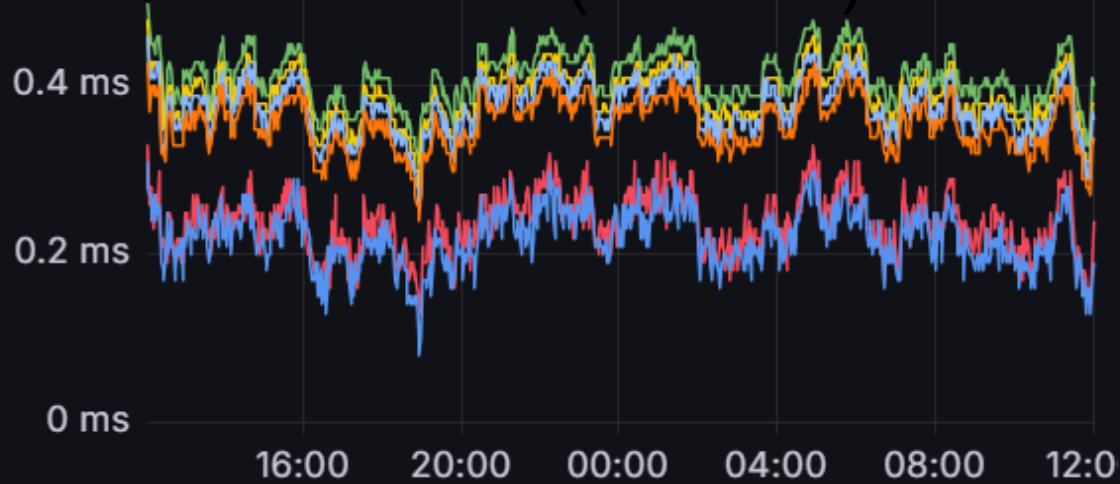
	Mean	Stdev
Set 1	0.407605	0.149712
Set 2	0.512782	0.01399
Set 3	0.456567	0.086595



HEAnet

- Pw-01 -> cw-01
 - No indication of congestion

• QMUL -> Jisc(London)



- 95th percentile - perfsonar-pw-01.heanet.ie -> perfsonar-cw-01.heanet.ie
- 95th percentile - perfsonar-pw-01.heanet.ie -> perfsonar-cw-01.heanet.ie
- Median - perfsonar-pw-01.heanet.ie -> perfsonar-cw-01.heanet.ie
- Median - perfsonar-pw-01.heanet.ie -> perfsonar-cw-01.heanet.ie

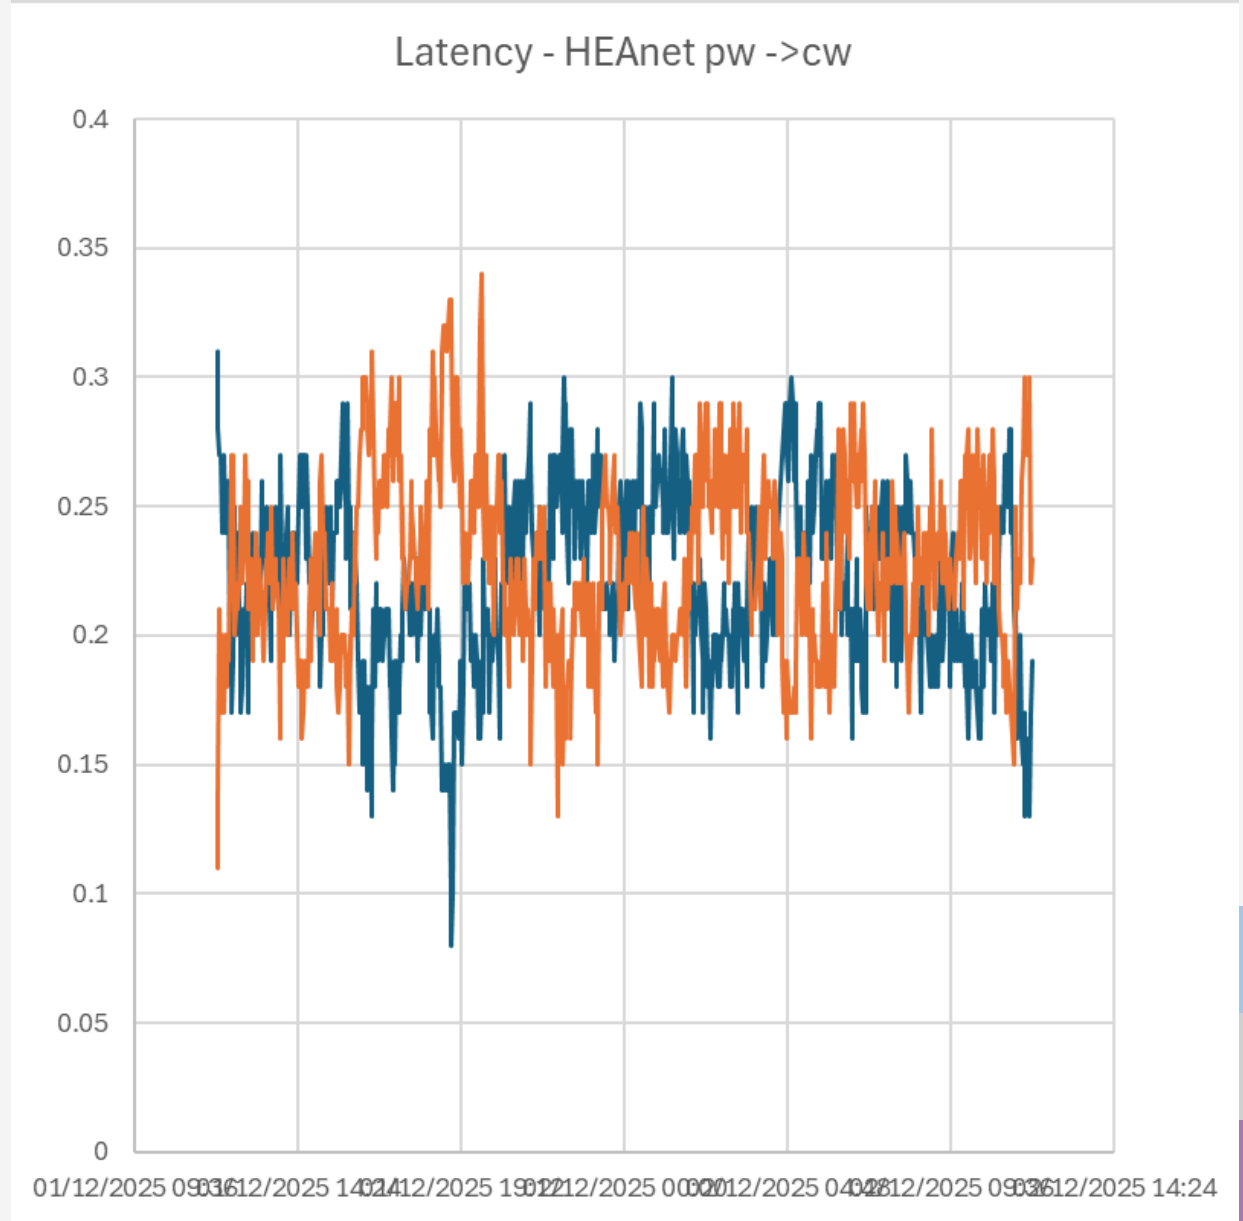
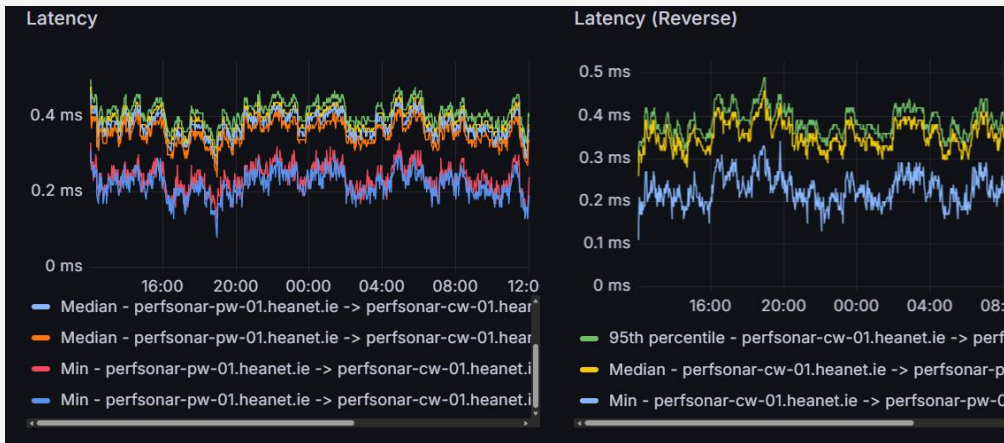


- 95th percentile - perfsonar-cw-01.heanet.ie -> perfsonar-pw-01.heanet.ie
- Median - perfsonar-cw-01.heanet.ie -> perfsonar-pw-01.heanet.ie
- Min - perfsonar-cw-01.heanet.ie -> perfsonar-pw-01.heanet.ie
- Min - perfsonar-cw-01.heanet.ie -> perfsonar-pw-01.heanet.ie

HEAnet latency

perfSONAR pw ->cw

- Latency pw ->cw 0.220+-0.034ms



Conclusions

Accurate time synchronisation needed for one way delay measurement

- Latency spikes (likely due to buffers filling)
 - London (100G) -> Slough (10G) ~75ms
 - Slough (10G) -> London (100G) ~ 2ms
 - Is bufferbloat an issue?
 - Does it impact performance
- Time synchronisation
 - Chrony
 - Increasing polling frequency increases stability
 - Hardware timestamping – makes a difference (probably)
 - PTP
 - How much will it improve things?

What next?

What next?

- PTP
 - Can it give better synchronisation
- Chrony
 - What advice can we give
 - Is 32s polling acceptable
 - Out to edge networks?
 - How to roll out?
- OWAMP
 - Is hardware timestamping necessary?

Measurements


- Measurements possible?
 - With current synchronisation
 - Network music performance
 - With better synchronisation
 - What accuracy do we need?
- Are we measuring the right thing?
 - Is min / 95th percentile the right measurement?


Thank you

Christopher J. Walker

Research Support Senior Manager

Christopher.Walker@jisc.ac.uk

 help@jisc.ac.uk

 0300 300 2212

 [jisc.ac.uk](https://www.jisc.ac.uk)



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



 [@jisc.bsky.social](https://bsky.app/profile/jisc.ac.uk)

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

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

 Jisc

Polling options - [chrony – chrony.conf\(5\)](#)

Minpoll poll

- This option specifies the minimum interval between requests sent to the server as a power of 2 in seconds. For example, **minpoll 5** would mean that the polling interval should not drop below 32 seconds. The **default is 6** (64 seconds), the minimum is -7 (1/128th of a second), and the maximum is 24 (6 months). Note that intervals shorter than 6 (64 seconds) should generally not be used with public servers on the Internet, because it might be considered abuse. A sub-second interval will be enabled only when the server is reachable and the round-trip delay is shorter than 10 milliseconds, i.e. the server should be in a local network.

Maxpoll poll

- This option specifies the maximum interval between requests sent to the server as a power of 2 in seconds. For example, **maxpoll 9** indicates that the polling interval should stay at or below 9 (512 seconds). The **default is 10** (1024 seconds), the minimum is -7 (1/128th of a second), and the maximum is 24 (6 months).

Effect of Clock offset

Latency

$$= T(\text{Host B}) - T(\text{Host A})$$

Measured latency

$$= T(\text{Host B}) - T(\text{Host A}) + \text{Clock offset}$$

Bristol Corn exchange dual time clock