



Janet-hosted test tools

Christopher J. Walker <Christopher.Walker@jisc.ac.uk>

Duncan Rand, Tim Chown, Raul Lopes

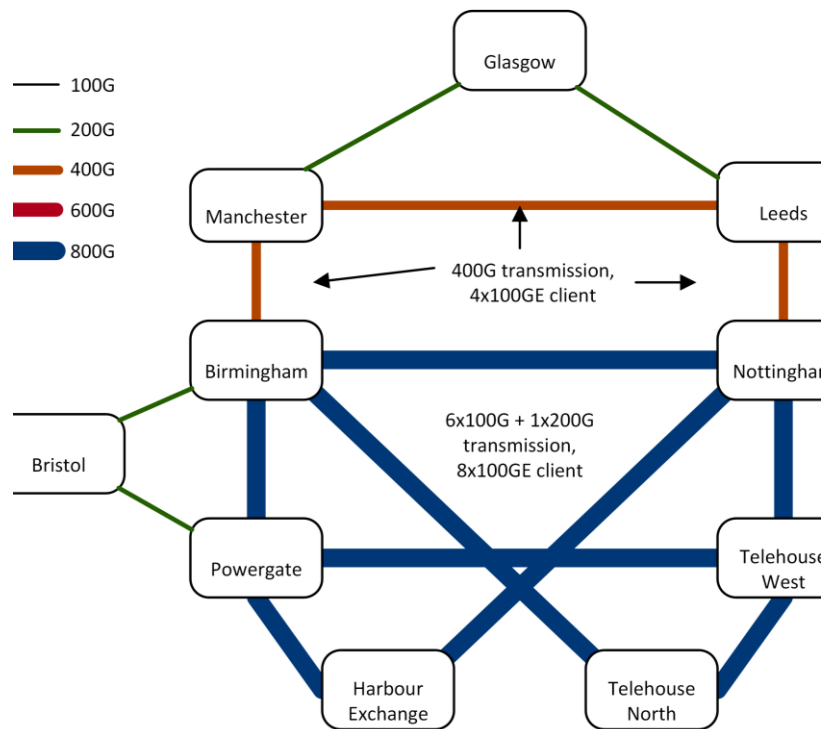
Hepsysman

Oxford, 22nd June 2023

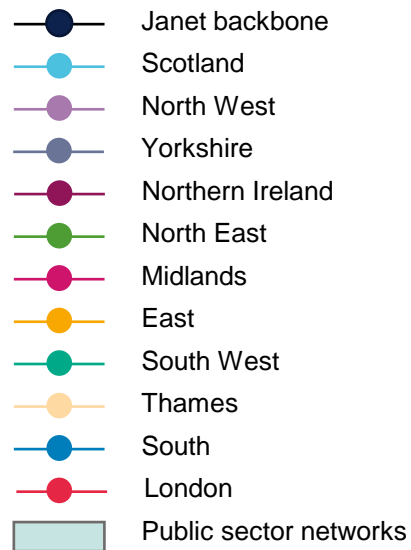
The Janet Network

Current Backbone Status

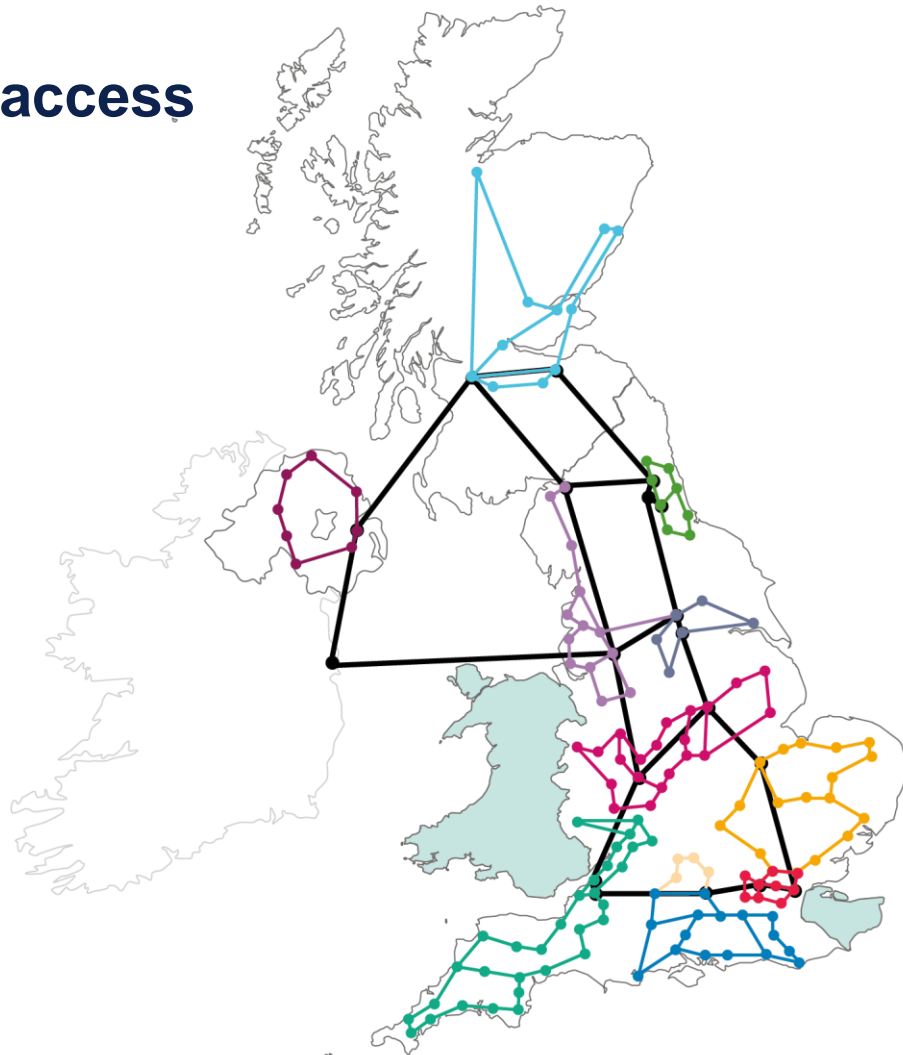
- 8x100Gbps trunks in southern part of the network
- Most traffic enters/exits network in London
- (Mainly) Juniper routing equipment
- Some 400Gbit/s paths in 75GHz channels
 - 4x100GE clients at the moment
 - Looking at 400GE



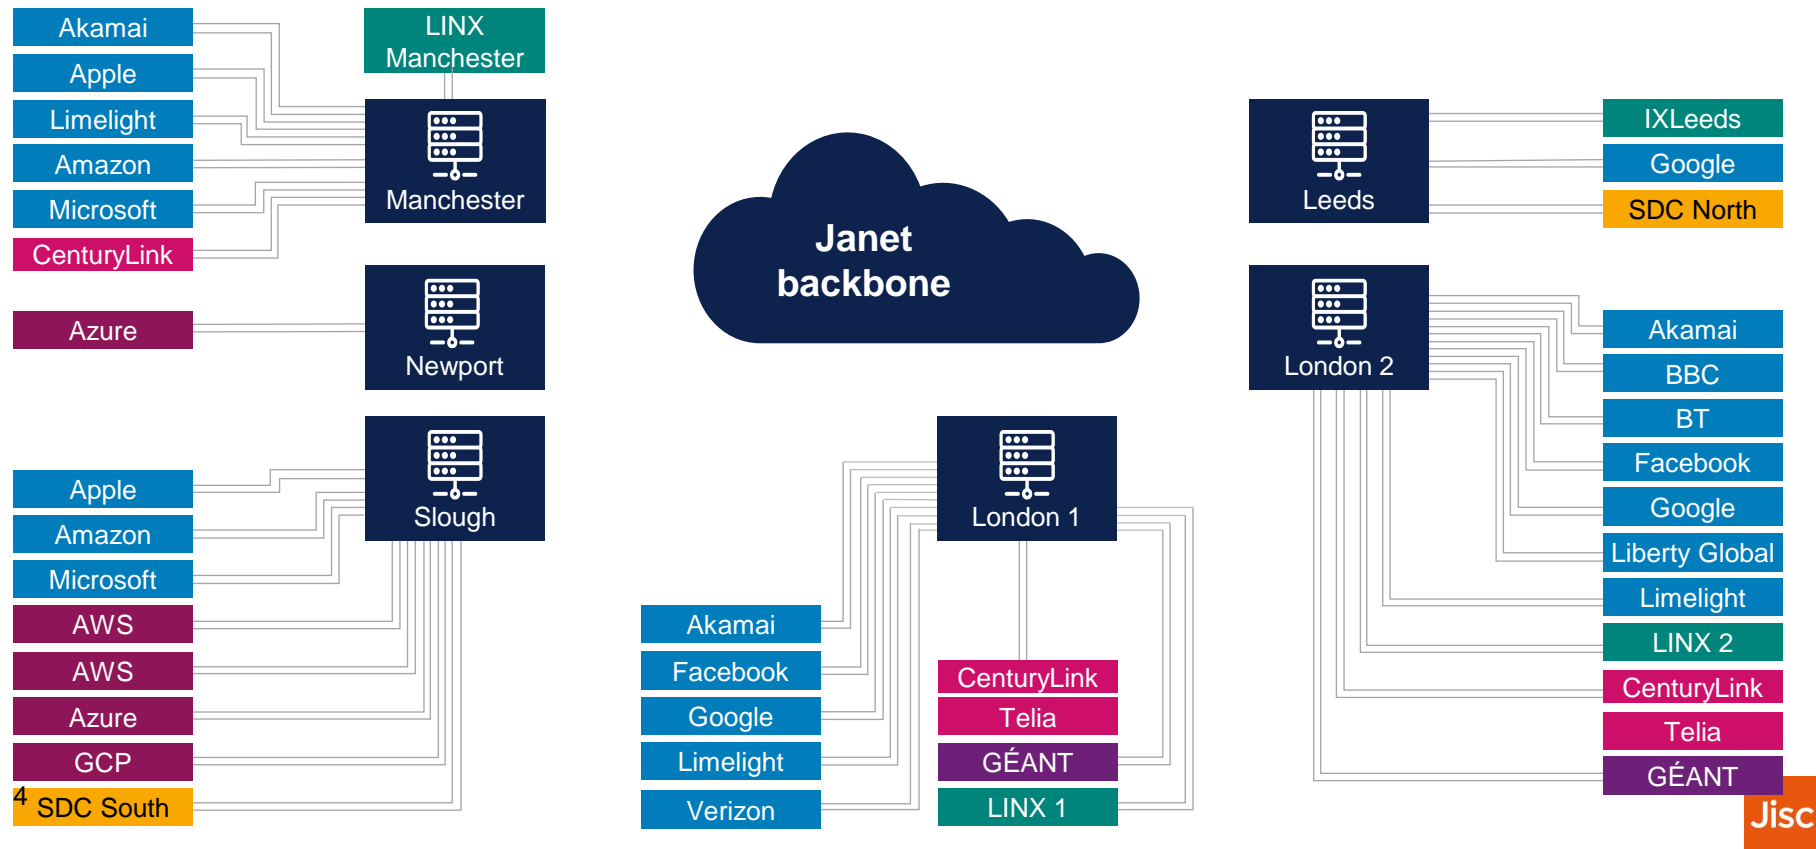
Janet backbone and regional access infrastructure



~1,000 customers and ~1,500 connections.



Jisc interconnectivity (aggregate > 4Tbit/s)



Jisc community support

- Network performance
 - [Using the Janet Network performance test facilities – Jisc](#)
 - [Research network engineering \(RNE\) community group – Jisc](#)
 - Next talk tomorrow (June 23rd)
- Other communities
 - [Digital research community group - Jisc](#)

Sciencemesh.io IDP Test

CERNbox is part of this federated Dropbox like storage

- Would you mind looking at:
 - <https://drive.sunet.se/>
 - Does it say “Account not provisioned.”
 - I’ll collate results

Testing connectivity

- **Ad hoc throughput testing**
 - *Iperf*
 - *ethr*
- **network characteristics over time**
 - perfSONAR
 - A RIPE Atlas anchor
- **Ad hoc disk to disk copying** to/from data transfer nodes (DTN)
 - A RIPE Atlas anchor, for bespoke RIPE Atlas tests
- Note: our NOC have additional tools, such as hardware line testers

Test facility specifics

- 10G iperf and ethr
 - iperf-slough-10g.perf.ja.net
 - ethr-slough-10g.perf.ja.net
- 10G perfSONAR
 - ps-slough-10g.perf.ja.net (bandwidth tests)
 - ps-slough-1g.perf.ja.net (latency tests)
- 100G perfSONAR
 - ps-london-bw.perf.ja.net (bandwidth tests)
 - ps-london-lat.perf.ja.net (latency tests)
- [RIPE Atlas](https://atlas.ripe.net/probes/6695/) (https://atlas.ripe.net/probes/6695/)

Ad-hoc throughput testing: *iperf3*

- There is a 10G *iperf3* server at the Jisc Slough data centre:
iperf-slough-10g.perf.ja.net (accessible over both IPv4 and IPv6)
- Will have a 100G iperf server which will be available on request
- *iperf3* software installation:
 - CentOS7: *sudo yum install iperf3*
- Using from Windows
 - [Microsoft guide to iperf usage](#)
 - [Compiling iperf/iperf3 from source](#)
 - [Example of Windows usage](#)

Example of using *iperf* *--bounceback*

WIFI

\$ walker@heplt024:~/src/iperf-2.1.9/src\$./iperf -c iperf-slough-10g.perf.ja.net --bounceback

[ID]	Interval	Transfer	Bandwidth	BB cnt=avg/min/max/stddev	Rtry	Cwnd/RTT	RPS
[1]	0.00-1.00 sec	1.95 KBytes	16.0 Kbits/sec	10=12.841/10.149/15.303/2.115 ms	0	14K/12121 us	78 rps

WIRED

[1]	0.00-1.00 sec	1.95 KBytes	16.0 Kbits/sec	10=7.268/6.945/7.935/0.262 ms	0	14K/7326 us	138 rps
------	---------------	-------------	----------------	-------------------------------	---	-------------	----------------

Ad-hoc throughput testing: *ethr*

- [ethr](#) is a “comprehensive network measurement tool for TCP, UDP & ICMP” from Microsoft
- Natively cross platform thanks to *golang*
- [Available](#) for CentOS7 with ‘snap’ & installed on our 10G server at Slough

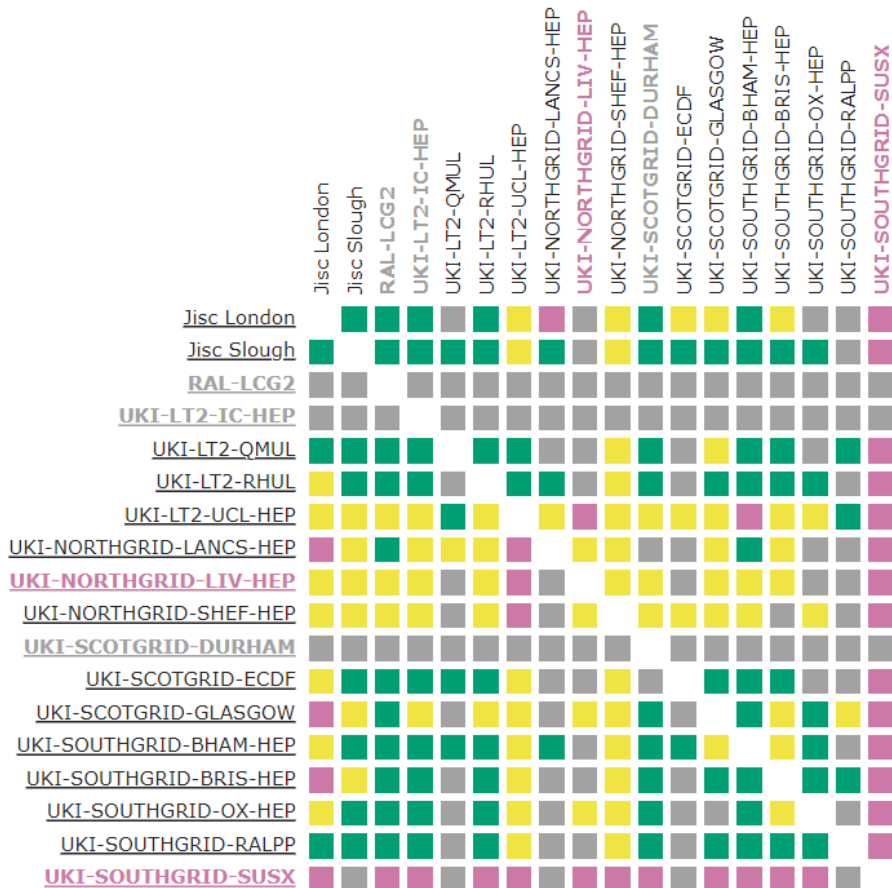
```
$ ethr -d 4s -c iperf-slough-10g.perf.ja.net
Connecting to host [2001:630:3c:f803::12], port 9999
[ 6] local 2001:630:3c:f803::6 port 51706 connected to 2001:630:3c:f803::12 port 9999
-----
[ ID] Protocol  Interval    Bits/s
[ 6]   TCP    000-001 sec   6.67G
[ 6]   TCP    001-002 sec   9.77G
[ 6]   TCP    002-003 sec   7.49G
[ 6]   TCP    003-004 sec   9.45G
Ethr done, duration: 4s.
```

Persistent measurement over time: perfSONAR

- Free, open source: <https://www.perfsonar.net>
- Easy to download and install on CentOS7 (and Debian)
- **Very useful to have persistent testing:** collect history of network characteristics – throughput, loss, latency, path
- Test against our 10Gbps node in the Jisc Slough data centre
 - Bandwidth: <https://ps-slough-10g.perf.ja.net>
 - Latency: <https://ps-slough-1g.perf.ja.net>
- We also are testing 1Gbps small nodes (including RPi) and Docker versions
- We have a GÉANT PMP small node (Gigabyte Brix):
 - <https://pmp-central.geant.org/maddash-webui/>

Persistent measurement over time: perfSONAR

- We have installed a 100G perfSONAR host in London
- Uses 9000 MTU
- <https://ps-london-bw.perf.ja.net>
- <https://ps-london-lat.perf.ja.net>
- Perfsonar 5.0.3 now available
- Supports EL 8 and 9 (beta)
- <https://psmad.opensciencegrid.org/macash-webui/index.cgi>



Perfsonar – services to watch

systemctl start perfsonar-lscachedaemon.service

systemctl start postgresql-10.service

systemctl start httpd

systemctl start pscheduler-archiver.service

systemctl start pscheduler-runner.service

systemctl start pscheduler-scheduler.service

systemctl start pscheduler-ticker.service

systemctl start psconfig-pscheduler-agent.service

systemctl status opensearch.service

/usr/lib/perfsonar/scripts/service_watcher

pscheduler monitor

Perfsonar

Common errors

- ICMPv6 blocked
- Certificates not in browser

10G Data Transfer Node (DTN) & Globus endpoint

- We have a DTN in the Jisc Slough data centre: *dtn-slough-10g.perf.ja.net* (dual-stack)
- A number of files are available: 1M.dat, 2M.dat, 10M.dat, 50M.dat, 1G.dat, 10G.dat, 20G.dat, 100G.dat, 1000G.dat accessible with *globus-url-copy*
- Can copy to /dev/null or to the file system, e.g. copying a 10GB file to /dev/null

```
$ globus-url-copy -vb ftp://dtn-slough-10g.perf.ja.net:2811/space00/10G.dat /dev/null
```
- There is also a directory with 100 x 1GB files for more sustained testing:

```
$ globus-url-copy -r -vb ftp://dtn-slough-10g.perf.ja.net:2811/space00/small/ file:///tmp/
```

100G Data Transfer Node (DTN) & Globus endpoint

- We are installing a 100G DTN facility in the Jisc London centre
- This will use a distributed files system – most likely Ceph or BeeGFS with multiple DTNs
- One goal is to provide a [Globus](#) endpoint to enable Janet connected sites to test their own Globus endpoints up to 100G
- However, other transfer tools can be made available on request to netperf@jisc.ac.uk

Browser-based speed tests

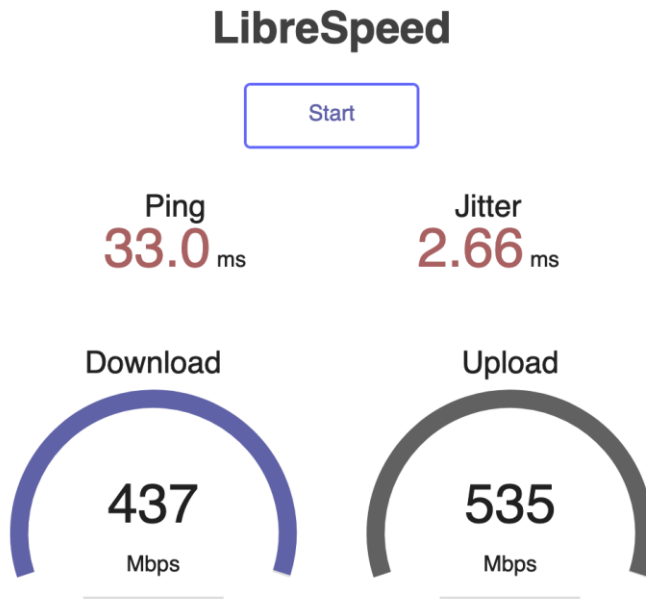
The devil you know!

- Easy to use, familiar to most of us
- Provide a useful first assessment from a browser
 - But often not accurate and generally have a limited performance
- Generally insufficient to demonstrate full connectivity

- We have installed a Librespeed instance as it allows us to offer this service from facilities that we control and understand
 - We will publish it with appropriate text about its limitations
 - People will be directed to contact us for iperf or perfSONAR tests and advice
 - Planning to include a URL for test results, and a CLI/API to the tool

Librespeed instance

Likely to be at <https://speedtest.perf.ja.net>



RIPE Atlas anchor

- See <https://atlas.ripe.net/>
- Supports measurements from RIPE Atlas nodes
 - Hardware (available from RIPE) or software probes
- The RIPE Atlas ecosystem is mature
 - Over 11,000 probes around the world
- Our anchor node is deployed at Slough
 - See <https://atlas.ripe.net/probes/6695/>
 - Useful for loss and latency, but can do more bespoke tests



Open questions

- What performance problems are you facing?
- What help might Jisc be able to provide?
- What other backbone-hosted tools would be useful for you?
- Any other suggestions for areas we should be covering?
 - WiFi performance?
 - Residential quality of experience (to campus and cloud)?
 - Low latency (and jitter) applications?
 - Experiments with (say) TCP BBR, or 9000 MTU?
 - Disk to disk testing at >100G

Thank you

- Please feel free to get in touch
- Email: netperf@jiscmail.ac.uk