



Jisc: Using perfSONAR

Dr Tim Chown, Network Development Manager, Jisc

**Building high-performing
campus infrastructures for research**

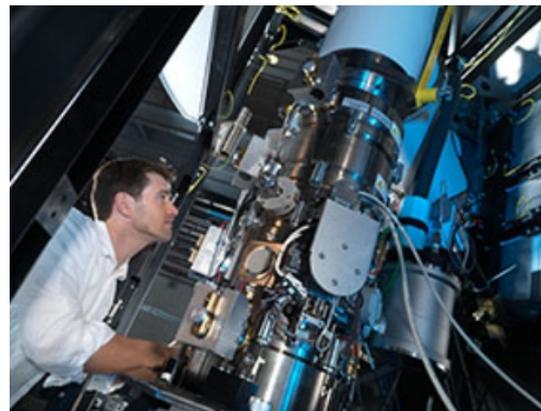
26 Sep 2018

» Our high-level goals:

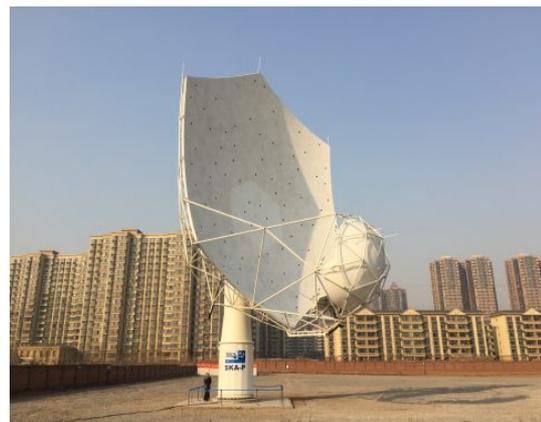
- › “Supporting user communities in making optimal use of the Janet network for high-performance network applications”
- › “We will identify, document and share best practices on high performance networking, to raise awareness amongst Janet network connected communities of the issues and factors affecting end-to-end network performance”



- » Growing interest in moving large volumes of research data
 - › Captured or generated data to remote computing facility
 - › Remote visualisation
 - › Data replication / distributed storage / backups
 - › To / from cloud
- » Data set volumes are increasing
 - › 100 TB is no longer 'very large'
 - › But moving 100 TB takes 10 Gbit/s of throughput for 24 hours
- » Examples:
 - › Astrophysics, genomics, environmental sciences, ...
 - › The new Titan Krios cryo-EM/ET microscope at Diamond
 - › The Square Kilometer Array (SKA) project



www.diamond.ac.uk



www.skatelescope.org

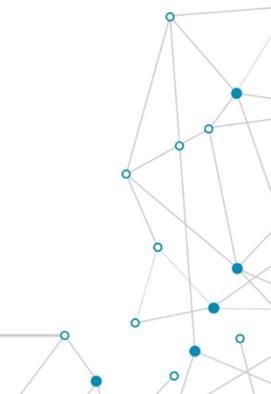


» An ongoing activity within Jisc:

- › Engaging with existing and data-intensive research communities and identifying emerging communities who need to move data around Janet (and beyond)
- › Creating dialogue between Jisc, computing services, and research communities
- › Holding workshops, facilitating discussion on e-mail lists, etc.
- › Helping researchers manage expectations
- › Sharing best practices in identifying and rectifying causes of poor performance
- › Promoting good practices in campus network engineering, esp. 'Science DMZ'
- › **Promoting deployment of performance measurement tools, esp. perfSONAR**

» More information:

- › <https://www.jisc.ac.uk/rd/projects/janet-end-to-end-performance-initiative>

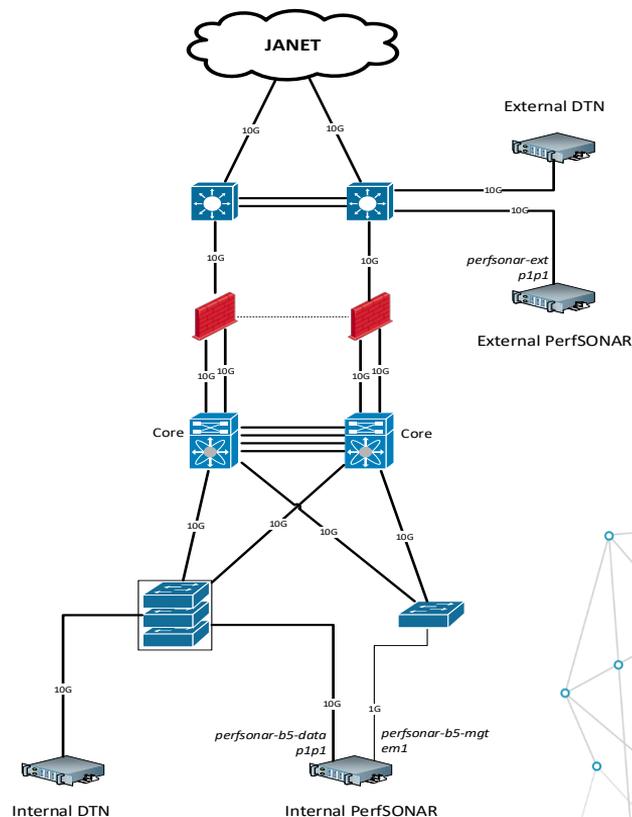


- » An example of data still being moved by physical media
 - › Southampton μ -VIS X-Ray Imaging Centre
 - › Has local facilities, but takes samples to Diamond Light Source ~6 times a year
 - › Might gather 10-40 TB of experimental result data per visit
 - › One data set typically a ~50 GB file, plus up to 5,000 8-25 MB files
 - › Tried using network and *rsync*; obtained ~30 MB/s (240 Mbit/s)
 - › Would take 4 days to copy 10 TB home over Janet, best case

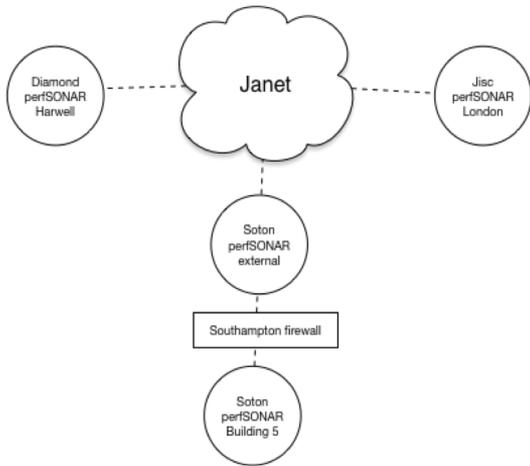
- » We ought to be able to do better...
 - › Diamond end has already deployed Science DMZ
 - › Southampton has a 10 Gbit/s campus link to Janet
 - › A target of 2 Gbit/s would allow ~1 TB per hour



- » Met with Diamond and Soton IT & research staff
- » Agreed a phased plan of action:
 - › Change to using Globus software tools
 - › Deploy perfSONAR to measure network characteristics
 - › Engineer 10 Gbit/s link to research file store, internal to campus firewall
 - › Pilot a 10 Gbit/s DTN at the campus edge
- » Outcome:
 - › External data transfers achieving 2-4 Gbit/s
 - › Potential to transfer their most recent 12 TB data set in 6-12 hours (overnight)



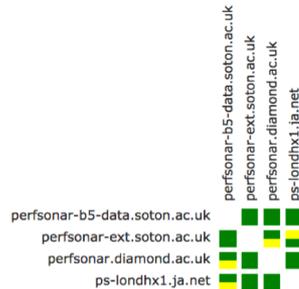
- » Jisc has deployed two perfSONAR servers – one at a London PoP, one at our Slough DC
- » We set up a perfSONAR mesh for the Southampton case study (on a Jisc VM mesh server)
- » Used measurement points at Diamond, Janet (London), and two at Southampton
- » See - <http://ps-dash.dev.ja.net/maddash-webui/index.cgi?dashboard=SES>



SES - Traceroute

■ Number of Paths is <= 1
 ■ Number of Paths is >:

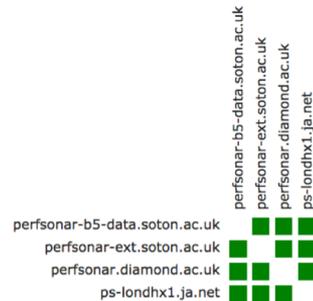
✓ No problems found in grid



SES - Throughput Testing

■ Throughput >= 900Mbps
 ■ Throughput < 900Mbps

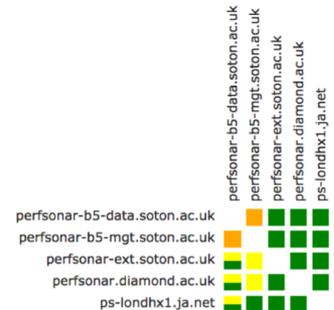
✓ No problems found in grid



SES - Latency Testing

■ Loss rate is <= 0
 ■ Loss rate is >= 0
 ■ Loss rate is :

✓ No problems found in grid



Janet London pS node to internal pS node

Source

ps-londhx1.ja.net
194.83.97.209,2001:630:3c:f800:0:0:0:209

[Host info](#) ▾

Destination

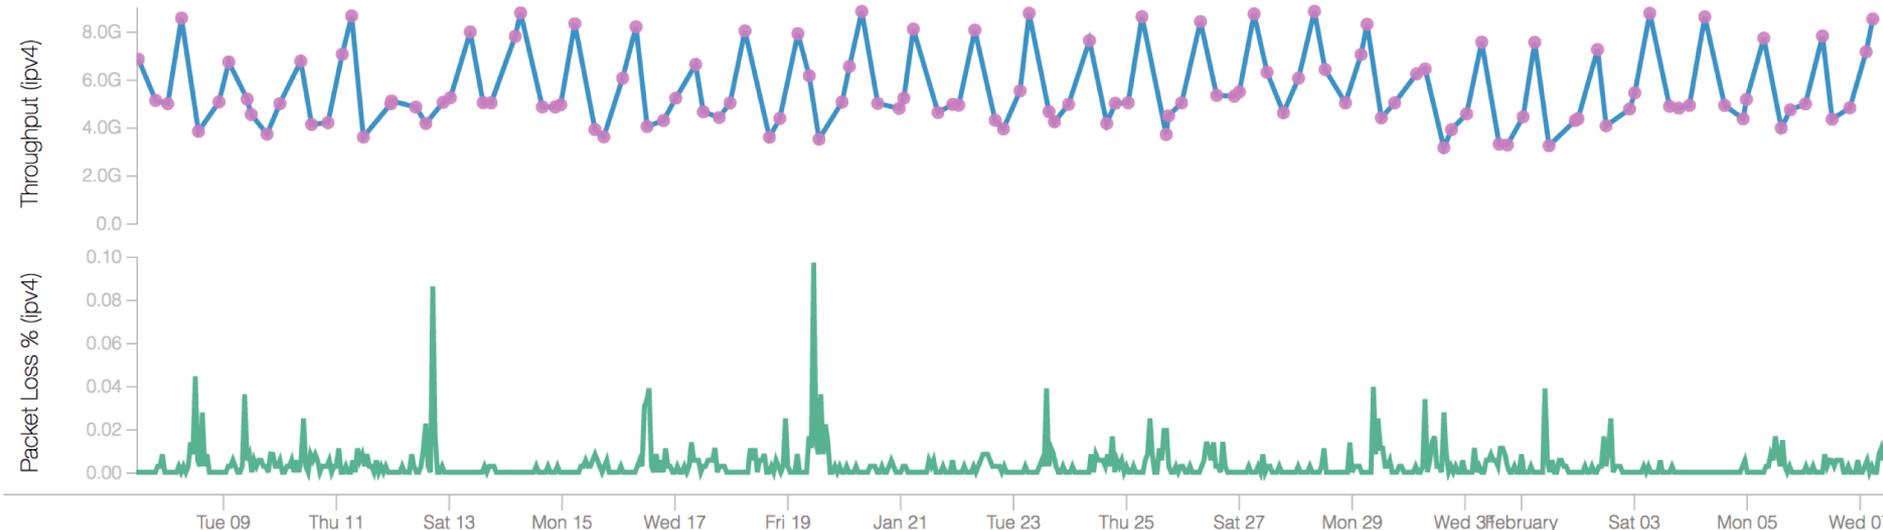
perfonar-b5-data.soton.ac.uk
152.78.176.16

[Host info](#) ▾

Report range

← 1 month →

Sun 01/07/2018 to Wed 02/07/2018
10:57:37 (GMT+0) 10:57:37 (GMT+0)



Janet London pS node to external pS node

Source

ps-londhx1.ja.net
 194.83.97.209,2001:630:3c:f800:0:0:0:209
[Host info](#) ▾

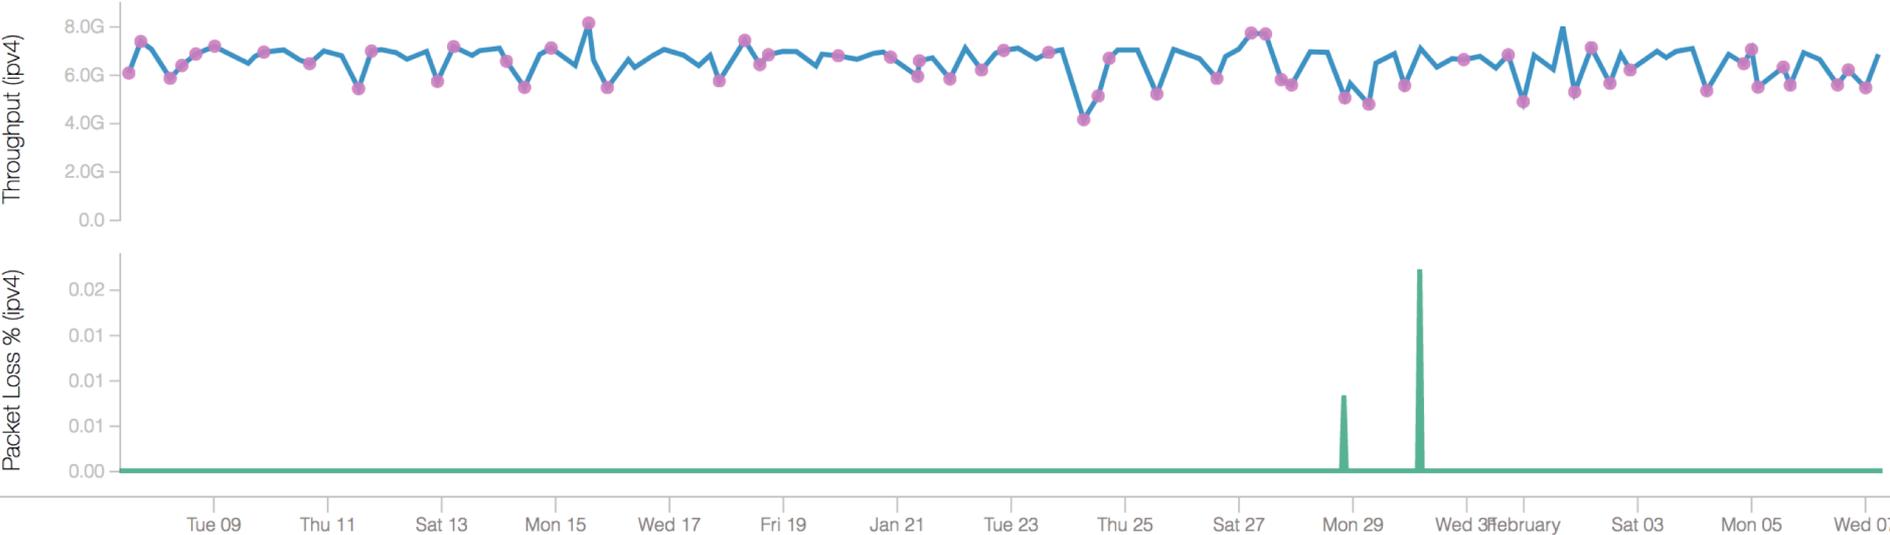
Destination

perfsonar-ext.soton.ac.uk
 152.78.1.2
[Host info](#) ▾

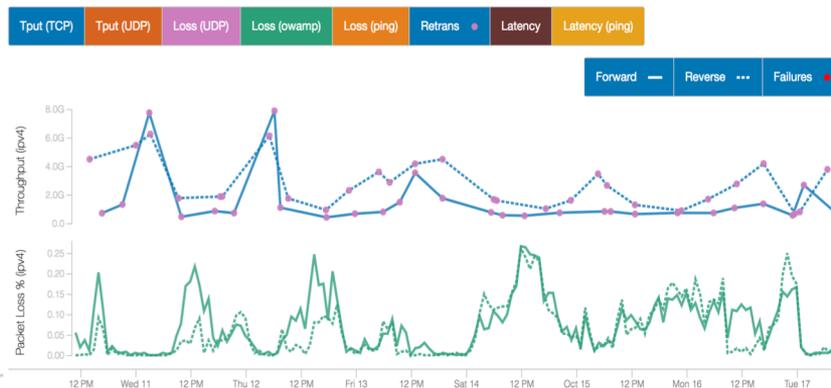
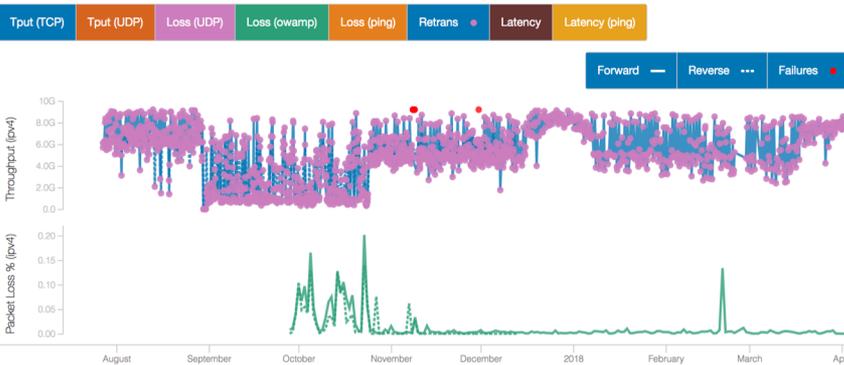
Report range

Sun 01/07/2018 to Wed 02/07/2018
 10:59:12 (GMT+0) to 10:59:12 (GMT+0)

Tput (TCP)
Tput (UDP)
Loss (UDP)
Loss (owamp)
Loss (ping)
Retrans ●
Latency
Latency (ping)
Forward —
Reverse ...
Failures ●



- » Slight persistent packet loss after a routine update of the Southampton firewall
- » Resulting throughput issues not reported by users, or observable with Jisc Netsight view
- » But highlighted by perfSONAR; clear drop in throughput, with higher loss (up to 0.3 %)
- » Also gives interesting insight into traffic characteristics over a year-long period



- » Working with Imperial College and SingAREN
- » New genomics project, needs to send/receive up to 200 TB of data between sites
- » perfSONAR highlighted one-way issue on Singapore -> Janet path; faulty hardware
- » Resolved with TEINCC/CERNET, now get 2.5 Gbit/s single stream, both ways

Source

bwctl-10g-ps.singaren.net.sg
203.30.39.13,2001:df0:21a:0:f6e9:d4ff:fea4:6432

[Host info](#) ▾

Destination

ps-londhx1.ja.net
194.83.97.209,2001:630:3c:f800:0:0:0:209

[Host info](#) ▾

Report range

← 1 month ▾ →

Sat 03/10/2018 to Tue 04/10/2018
11:25:07 (GMT+0) 12:25:07 (GMT+1)

Tput (TCP)

Tput (UDP)

Loss (UDP)

Loss (owamp)

Loss (ping)

Retrans ●

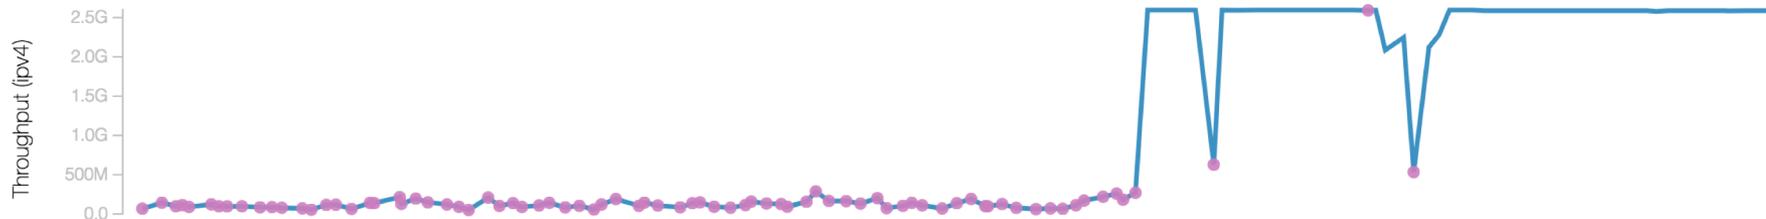
Latency

Latency (ping)

Forward —

Reverse ...

Failures ●



```
$ pscheduler task --slip PTH throughput --source=bwctl-10g-ps.singaren.net.sg --dest=ps-londhx1.ja.net -t 30 --ip-version 4
```

```
* Stream ID 5
```

Interval	Throughput	Retransmits	Current Window
0.0 - 1.0	2.67 Mbps	0	81.40 KBytes
1.0 - 2.0	9.97 Mbps	0	333.74 KBytes
2.0 - 3.0	64.37 Mbps	0	2.27 MBytes
3.0 - 4.0	314.58 Mbps	0	9.68 MBytes
4.0 - 5.0	1.22 Gbps	0	25.23 MBytes
5.0 - 6.0	723.46 Mbps	1	12.58 MBytes
6.0 - 7.0	602.85 Mbps	2	6.59 MBytes
7.0 - 8.0	402.02 Mbps	0	6.40 MBytes
8.0 - 9.0	314.56 Mbps	0	6.65 MBytes
9.0 - 10.0	335.56 Mbps	1	3.63 MBytes
10.0 - 11.0	115.32 Mbps	2	1.20 MBytes
11.0 - 12.0	62.92 Mbps	0	976.80 KBytes
12.0 - 13.0	52.43 Mbps	0	1.21 MBytes
13.0 - 14.0	83.89 Mbps	0	1.89 MBytes
14.0 - 15.0	83.87 Mbps	1	1.22 MBytes
15.0 - 16.0	62.92 Mbps	0	1.27 MBytes
16.0 - 17.0	83.89 Mbps	0	1.69 MBytes
17.0 - 18.0	94.37 Mbps	0	2.58 MBytes
18.0 - 19.0	178.27 Mbps	0	4.18 MBytes
19.0 - 20.0	251.65 Mbps	0	6.39 MBytes
20.0 - 21.0	419.18 Mbps	0	9.95 MBytes
21.0 - 22.0	566.56 Mbps	0	13.96 MBytes
22.0 - 23.0	765.42 Mbps	0	18.87 MBytes
23.0 - 24.0	1.05 Gbps	1	11.84 MBytes
24.0 - 25.0	566.22 Mbps	0	11.90 MBytes
25.0 - 26.0	461.40 Mbps	1	6.00 MBytes
26.0 - 27.0	293.60 Mbps	0	6.17 MBytes
27.0 - 28.0	262.15 Mbps	1	3.26 MBytes
28.0 - 29.0	167.77 Mbps	0	3.35 MBytes
29.0 - 30.0	167.76 Mbps	0	4.01 MBytes

```
Summary
```

Interval	Throughput	Retransmits
0.0 - 30.0	325.88 Mbps	10



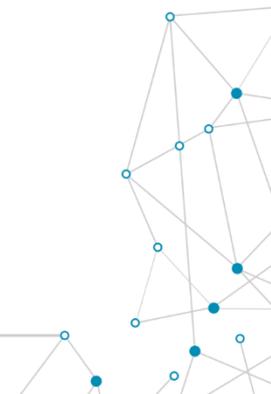
```
$ pscheduler task --slip PT1H throughput --source=bwctl-10g-ps.singaren.net.sg --dest=ps-londhx1.ja.net -t 30 --ip-version 4
```

```
* Stream ID 5
```

Interval	Throughput	Retransmits	Current Window
0.0 - 1.0	2.15 Mbps	0	81.40 KBytes
1.0 - 2.0	8.34 Mbps	0	260.48 KBytes
2.0 - 3.0	45.42 Mbps	0	1.14 MBytes
3.0 - 4.0	145.17 Mbps	0	4.89 MBytes
4.0 - 5.0	995.93 Mbps	0	33.80 MBytes
5.0 - 6.0	2.58 Gbps	0	50.40 MBytes
6.0 - 7.0	2.42 Gbps	0	50.40 MBytes
7.0 - 8.0	2.80 Gbps	0	50.40 MBytes
8.0 - 9.0	2.42 Gbps	0	50.40 MBytes
9.0 - 10.0	2.62 Gbps	0	50.40 MBytes
10.0 - 11.0	2.61 Gbps	0	50.40 MBytes
11.0 - 12.0	2.41 Gbps	0	50.40 MBytes
12.0 - 13.0	2.81 Gbps	0	50.40 MBytes
13.0 - 14.0	2.42 Gbps	0	50.40 MBytes
14.0 - 15.0	2.69 Gbps	0	50.40 MBytes
15.0 - 16.0	2.53 Gbps	0	50.40 MBytes
16.0 - 17.0	2.46 Gbps	0	50.40 MBytes
17.0 - 18.0	2.76 Gbps	0	50.40 MBytes
18.0 - 19.0	2.42 Gbps	0	50.40 MBytes
19.0 - 20.0	2.80 Gbps	0	50.40 MBytes
20.0 - 21.0	2.43 Gbps	0	50.40 MBytes
21.0 - 22.0	2.56 Gbps	0	50.40 MBytes
22.0 - 23.0	2.66 Gbps	0	50.40 MBytes
23.0 - 24.0	2.42 Gbps	0	50.40 MBytes
24.0 - 25.0	2.80 Gbps	0	50.40 MBytes
25.0 - 26.0	2.42 Gbps	0	50.40 MBytes
26.0 - 27.0	2.65 Gbps	0	50.40 MBytes
27.0 - 28.0	2.58 Gbps	0	50.40 MBytes
28.0 - 29.0	2.41 Gbps	0	50.40 MBytes
29.0 - 30.0	2.81 Gbps	0	50.40 MBytes

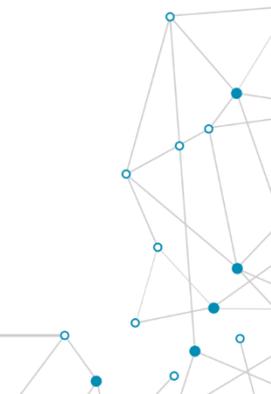
```
Summary
```

Interval	Throughput	Retransmits	Receiver Throughput
0.0 - 30.0	2.19 Gbps	0	2.17 Gbps



- » Janet E2EPI project page
 - › <https://www.jisc.ac.uk/rd/projects/janet-end-to-end-performance-initiative>
- » JiscMail E2EPI list (approx 100 subscribers)
 - › <https://www.jiscmail.ac.uk/cgi-bin/webadmin?Ao=E2EPI>
- » Campus Network Engineering for Data-Intensive Science workshop slides
 - › <https://www.jisc.ac.uk/events/campus-network-engineering-for-data-intensive-science-workshop-19-oct-2016>

- » Fasterdata knowledge base
 - › <http://fasterdata.es.net/>
- » eduPERT knowledge base
 - › <http://kb.pert.geant.net/PERTKB/WebHome>



A close-up photograph of a person's hand holding a white business card. The hand is positioned at the top and bottom edges of the card, with fingers visible. The background is a plain, light color.

Dr Tim Chown

Network Development Manager

tim.chown@jisc.ac.uk

jisc.ac.uk