



Science and  
Technology  
Facilities Council

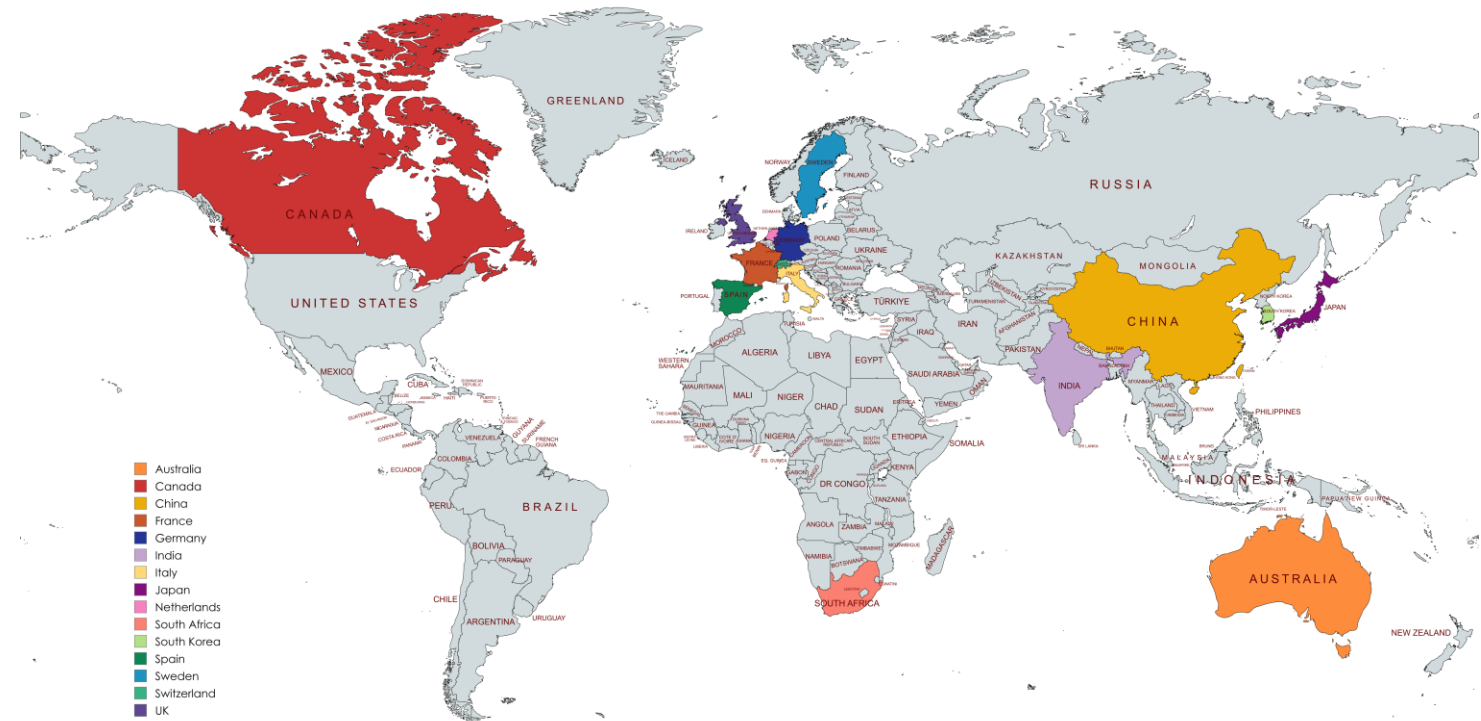
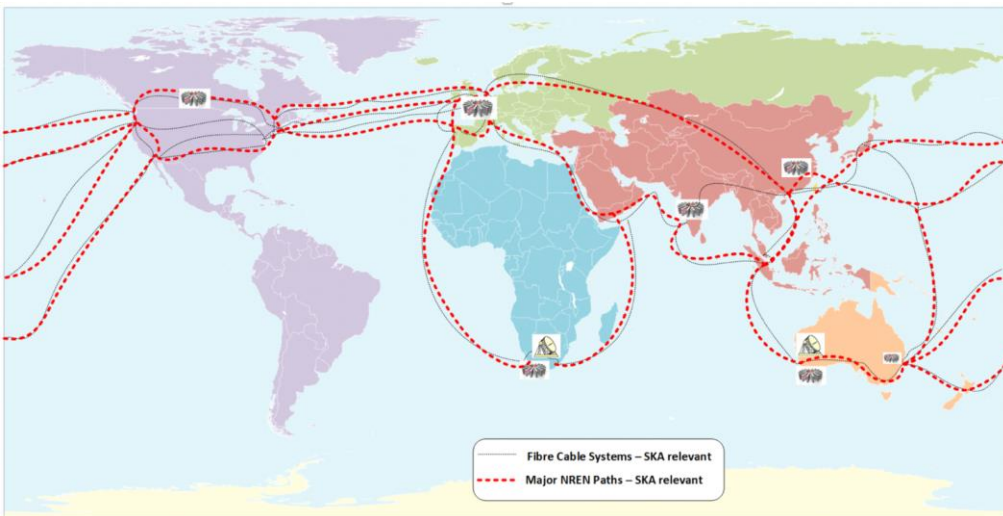
Scientific Computing

# SKA perfSONAR mesh

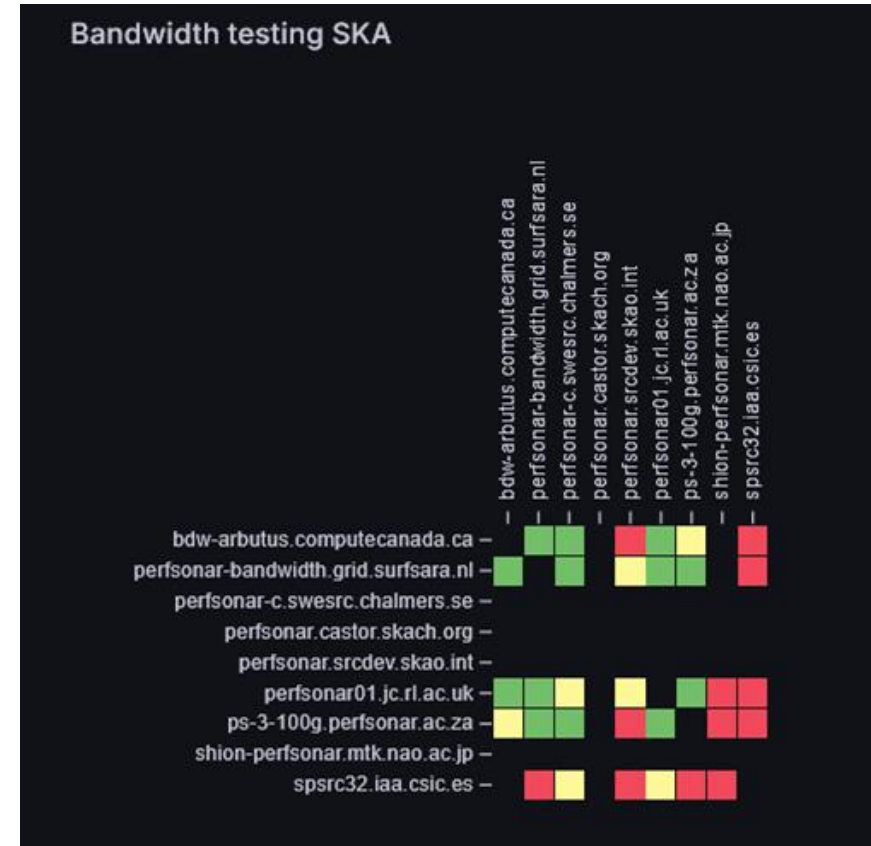
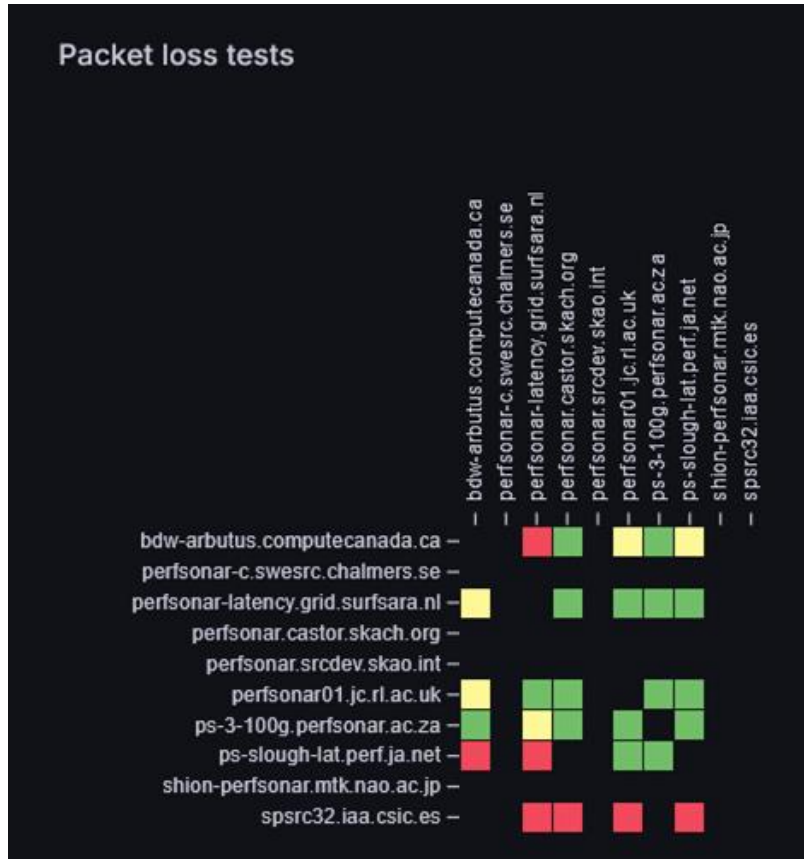
LHCONE meeting 18/03/2025

# SKA Regional Centres

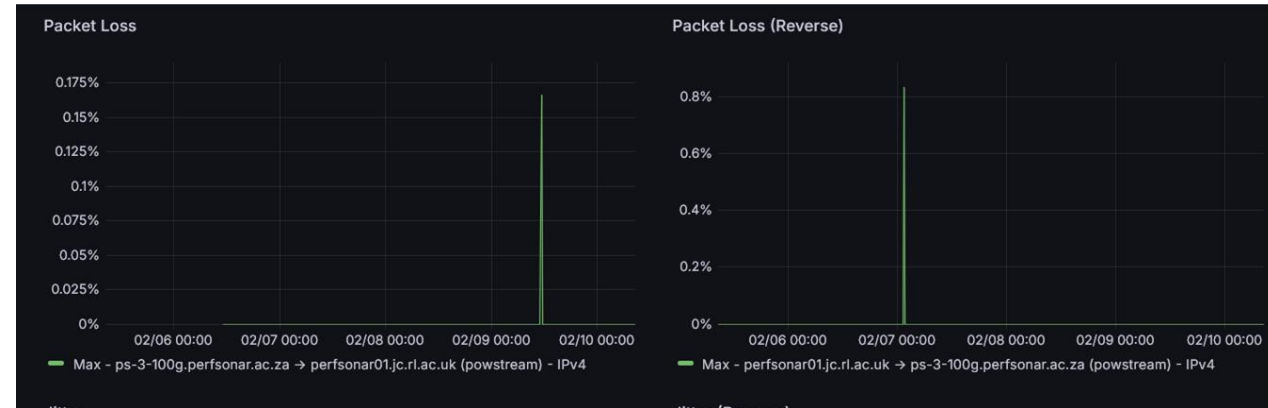

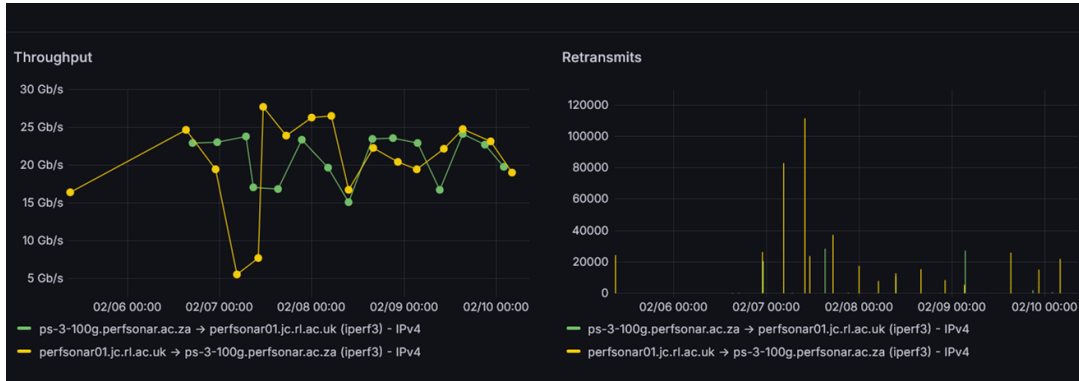
Square Kilometre Array project split into regional centres:



# Current Mesh

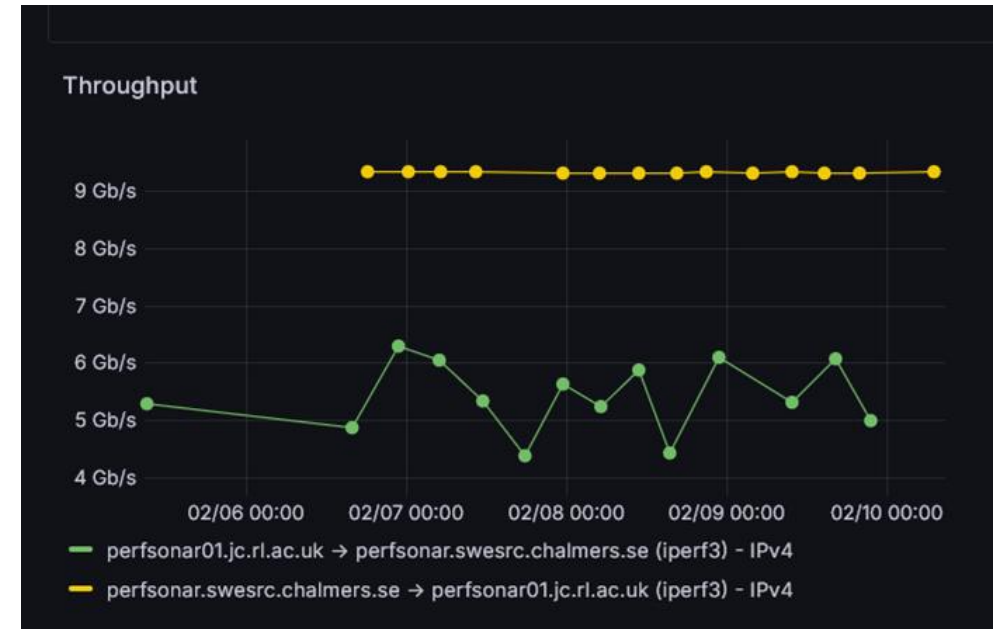
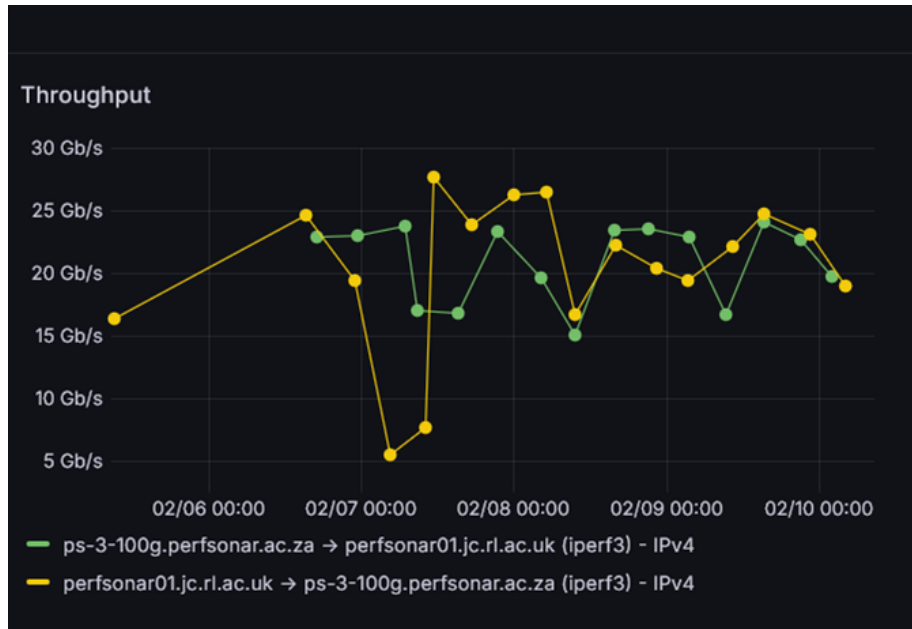


# Current Mesh



# Different Capacities

Different nodes have different capacities  
Some 100Gb, others 10Gb etc

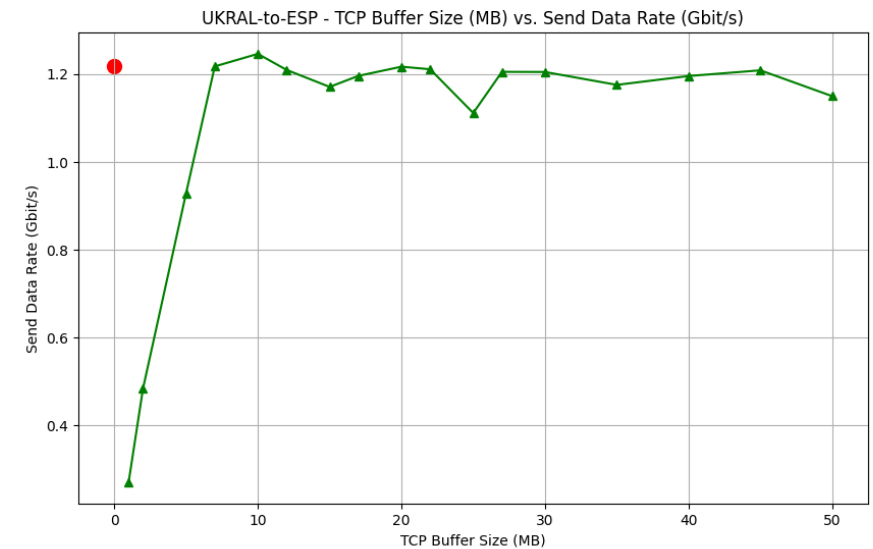
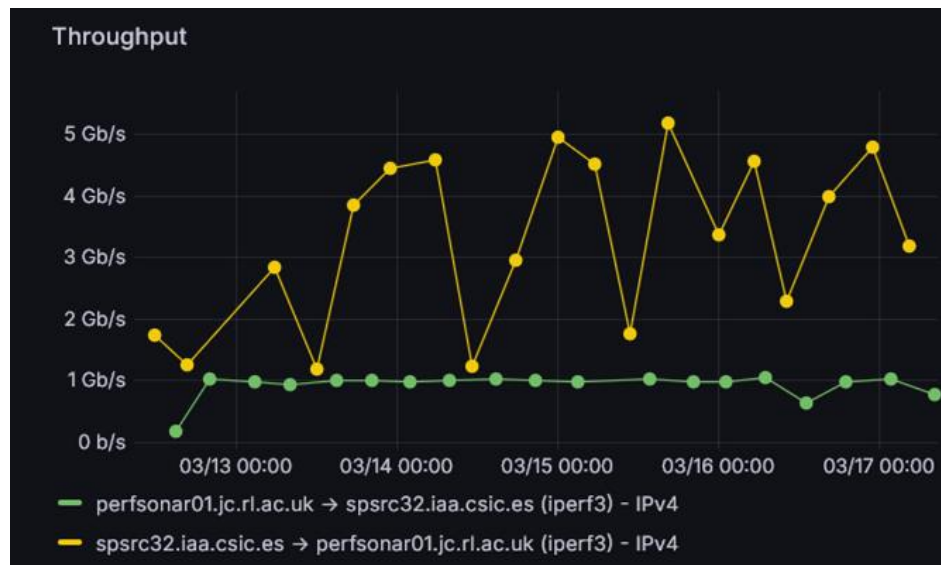


# Comparing tests

Tests using TCPmon compared to single stream iperf3 for some lower capacity nodes.

Some nodes have not enough hardware currently for deployment, using cloud architecture temporarily.

Sanity checking perfSONAR iperf results with different methods:



# SKA regional centres (SRC) nodes $\neq$ sites

SRC nodes are country members of the SRCnetwork

Each of these can have different sites, such as the UK node having >3 sites.


Sites will be setup with a separate mini mesh currently...

(I broke the dashboard prior to this, currently 3 nodes, RAL, Cambridge, Manchester)

Extraordinarily huge help from JISC in general

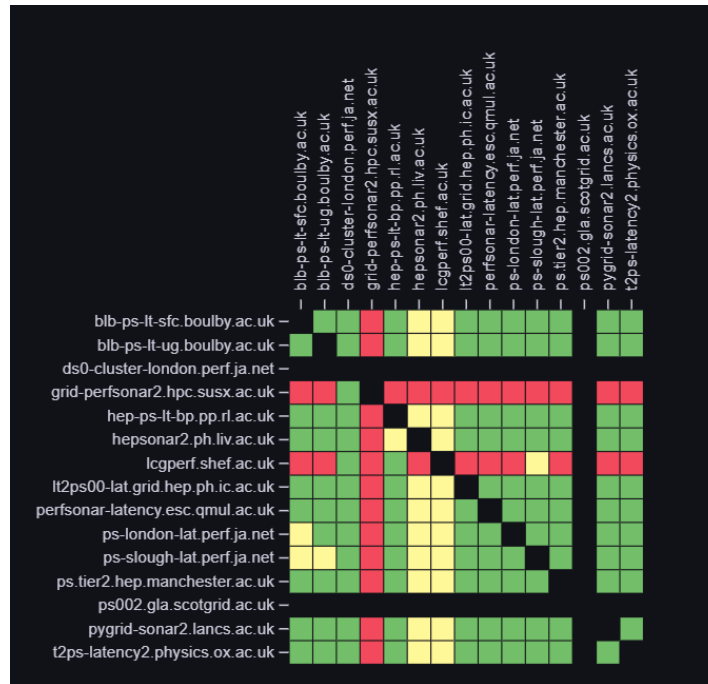


Scientific Computing



# JISC – UK tests

Jisc perfonar mesh has large number of sites  
 Some overlap between UKSRC sites  
 This allows measurements despite changes to SRCnet mesh



# Changes to the SRCnet mesh

Current setup used Esnet recommendations for host tuning ~ 100Gb/s:

<https://fasterdata.es.net/host-tuning/linux/100g-tuning/>

E.g. TCP buffer size 2GB

NIC buffer size max

Disabled SMT

Fair Queing

Additionally bonded dual 100Gb interface  
With 9000 MTU

# Changes to the SRCnet mesh

To do list:

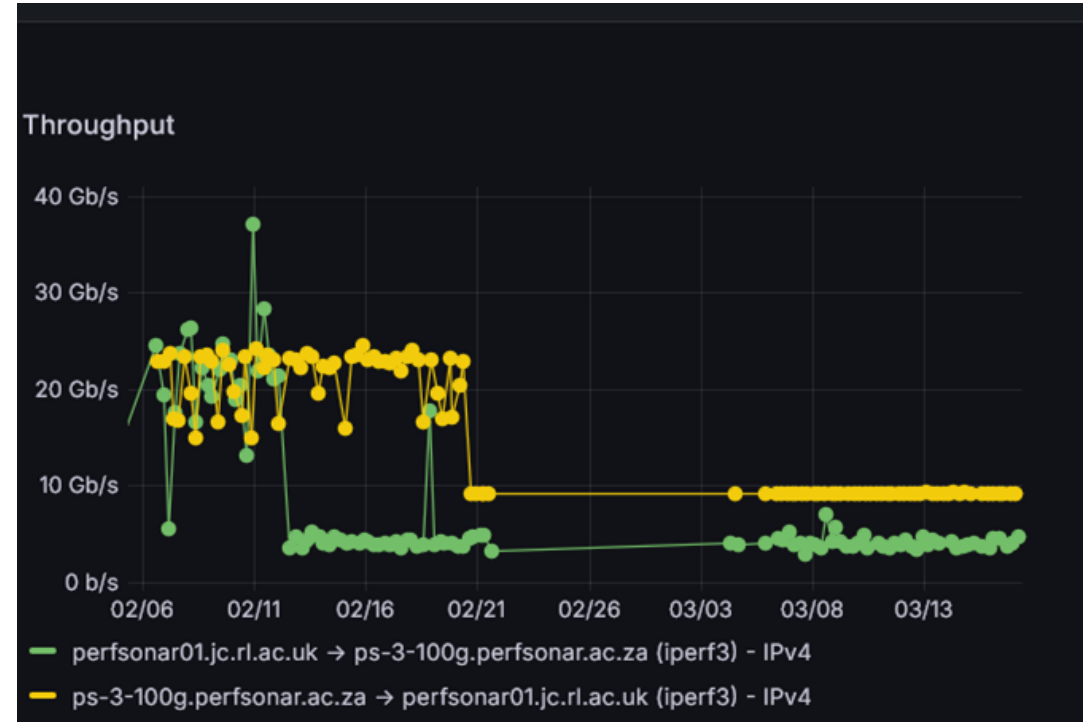
Split interfaces

IPv6

Ongoing improvement from nodes (as more hardware gets installed)

Full representation from SRCs

... Needs more looking into so I can \*not\* break the networking... Apologies to networking teams at RAL




# Changes to the SRCnet mesh

Other sites increasing hardware for deployments:

Changing from 10Gb interfaces to 100Gb interfaces:

**Old Node:**

**New Node:**



```
Summary
Interval      Throughput      Retransmits      Receiver Throughput
0.0 - 10.0    37.13 Gbps      110522           34.94 Gbps

No further runs scheduled.
```

```
Summary
Interval      Throughput      Retransmits      Receiver Throughput
0.0 - 10.0    5.30 Gbps       8                5.15 Gbps

No further runs scheduled.
```

# Summary

Work in progress...

Nodes updating and upgrading hardware

Configuration changes being made to existing nodes

Sites coming online and adding to mesh

Different deployment methods (currently)

Mesh config still to be determined

Needing IPv6 / split interface setup for RAL

SRCnet still developing

# Acknowledgements/thanks

Images from SKA website (and SKA confluence... )

SRCnet program team,  
SRC teams (especially Spain, Sweden,  
Manchester/HQ, Switzerland)

UK RAL SKA team

JISC

RAL Networking team

perfSONAR developers

Richard Hughes Jones for TCPmon script



Scientific Computing



# Questions?