

The Jisc logo is a white text 'Jisc' inside an orange square. The background of the slide features a blurred image of a man with glasses and a beard, wearing a blue sweater, looking down at a device. There are also decorative wavy lines at the bottom of the slide.

Jisc

Networkshop_50

Using RARE/FreeRtr for the Janet network performance test facility

Tim Chown

Network development manager, Jisc

50

Agenda for this session

P4 / programmable networking in campuses

We have two talks

- I'll give a brief overview of Jisc's work towards launching a **100Gbps network performance test facility** which we are hosting behind a white box running RARE/FreeRtr with P4
- Prof Nick Race will give some general information about P4 and then talk about **research work at Lancaster University** in applying programmable networking


Supporting network performance testing

Jisc has 10G-capable test systems you can use today

- Hosted at our Slough shared DC
- Juniper EX4600 in front of physical servers running Linux
- This provides you with three capabilities:
 - **iperf** - *iperf-slough-10g.ja.net*
 - [perfSONAR](#) - *ps-slough-10g.ja.net* and *ps-slough-1g.ja.net*
 - DTN application tests, including [Globus endpoint](#)
- Email netperf@jisc.ac.uk for further info on these
- Or to suggest we add other tools that would help you

Example: perfSONAR mesh for UK GridPP community


UK Mesh Config - UK IPv4 Bandwidth - Throughput



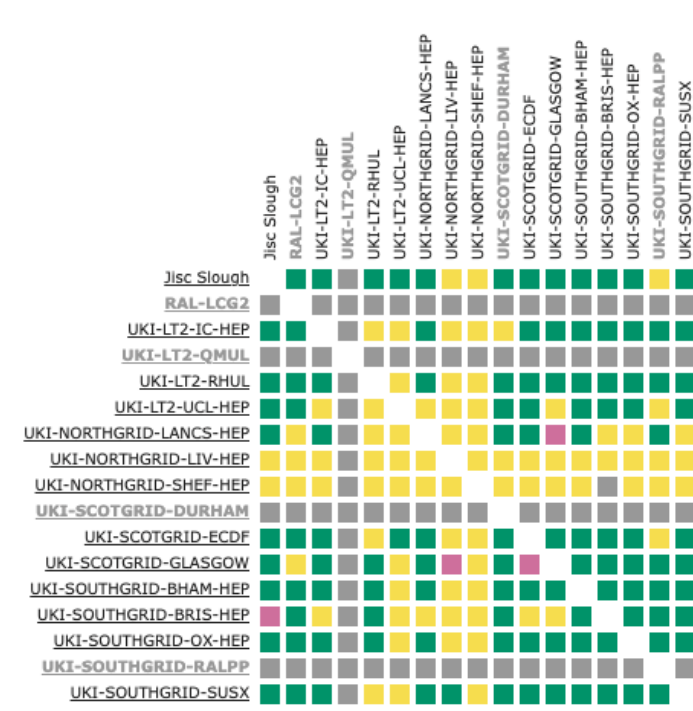
! Found a total of 15 problems involving 12 hosts in the grid



UK Mesh Config - UK IPv4 Latency - Loss



! Found a total of 4 problems involving 4 hosts in the grid



See <https://psmad.opensciencegrid.org/maddash-webui/index.cgi?dashboard=UK%20Mesh%20Config>

Drilling down to specific measurements over time

perfSONAR test results - [documentation](#)

[Share/open in new window](#)

Source

ps-slough-10g.ja.net
194.81.18.227,2001:630:3c:f803:0:0:0:6
[Host info](#) ▾

Destination

lt2ps00-bw.grid.hep.ph.ic.ac.uk
146.179.233.220,2a0c:5bc0:c8:2:1e34:daff:fe75:3496
[Host info](#) ▾


Report range

← Custom Range →

2022-04-17 15:18 2022-05-17 15:18 **Submit**

Sun, 17 Apr 2022 14:18:42 GMT to
Tue, 17 May 2022 14:18:42 GMT

Show/hide chart rows Throughput Packet Loss Latency Application Response Time



Extending the test capability

Going beyond 10G

- Need to be able to test at the capacity of our member sites
 - Several now connected at 100G, one at 2x100G
 - And certainly to test at > 10G
- Want to be cost-efficient
 - Traditional vendor 100G interfaces are not cheap
- Interested in exploring new alternatives in production
 - The test facility is important, but not mission critical to Janet
- **Would be nice to run a flexible platform where we can more easily add or test features the community want**

Enter RARE/FreeRtr

An open source routing platform supporting P4

- Combines FreeRtr control plane with P4 (or other) data plane
- **Feature rich, supports very high performance on low-cost white box hardware**
- Built in GÉANT GN4-3 project, validated on GÉANT P4 Lab
- RARE docs: <https://wiki.geant.org/display/RARE/Home>
- RARE/FreeRtr docs: <http://docs.freertr.org/>
 - Includes install for Wedge 100BF-32X (which Jisc is using)
 - ONIE image, CI/CD deployment, Nix package manager
- Mail list: <https://lists.geant.org/sympa/info/rare-users>



Why RARE/FreeRtr?

Much to like

- Runs on relatively cheap but performant 100G hardware
 - Edgecore Wedge 100BF-32X, 6.4Tbit/s capacity, cost £6,500+VAT
 - Variant of the switch used by Facebook in their DCs
- Demonstrated performant in [2021 Data Mover Challenge](#)
 - Achieved >85Gbps from Amsterdam to Singapore and Australia
- Open platform, with **flexibility to add new features**
 - e.g., [PoIKA](#), policy routing, AMT and bier, WLCG packet marking
- Built by members of the NREN community

Are there downsides to this approach?

There are potential issues

- It's not a Juniper/Cisco/<insert vendor here> box
 - So there's a new environment to learn
 - There is no 'commercial' support for RARE/FreeRtr
- The Wedge hardware is fast, but limited in some ways
 - In particular the buffers are only 16MB
 - The impact of this may depend on what we try to run
 - If necessary, we will consider a higher spec white box or revert to more traditional hardware

When will the 100G facility be available?

Status

- New rack has been set up in London
 - Hardware is being installed
 - Many thanks to David Richardson in our ops team 😊
- We will need to do some internal testing
 - Then will look for community testers
- If all goes well, expect a fuller launch in July/August

- Any questions, please email me or netperf@jisc.ac.uk

Summary

And what next?

- Important that Jisc can provide test facilities to our members
- We see value in getting experience with low cost white box hardware with P4 support
 - Deploying at the edge in a non-critical use case initially
- RARE/FreeRtr is an excellent platform for doing that
- Also keen to explore other new approaches in the facility
 - Streaming telemetry is one example
 - Please feel free to suggest other tools we might provide

Thank you

Tim Chown

Network development manager

tim.chown@jisc.ac.uk

T 01235 822106

help@jisc.ac.uk

jisc.ac.uk



Workshop_50

50