



Janet end-to-end performance initiative (e2epi)

MSN 2015

Cosener's House, 2nd July 2015

Aims of the Janet ezepe project

- » Identify existing and emerging communities seeking to run high performance applications (high throughput, low latency, ...)
- » Determine and share best practices
- » Foster discussion between research communities, network operators (inc. Janet & end sites) and **network researchers**
- » Help set expectations, and raise awareness of issues
- » Raise 'high water marks'
- » Offer specialist support & tools (2 FTEs being advertised soon)

Who are the communities?

» Existing

- › Astrophysics (DiRAC), Particle Physics (GridPP), Human Genome Project (Hinxton), Environmental Science (CEDA), Oceanography,...
- › Typical data rates through Janet of 1-6Gbit/s
- › LHC is just restarting

» Emerging

- › Cyber security, 'Data science', Met Office, UK-RDF (23PB), ...
- › New Funding Council requirements to archive data
- › +...?

Various shifts happening in university/HE network use

- » Traditionally 'computing service' functions were all on-site
- » Increased outsourcing
 - › Cloud-based services – raises many issues, including performance
- » Off-site data centres
 - › My own university Southampton has a new DC in Fareham
 - › Jisc shared DC at Slough - <https://www.jisc.ac.uk/shared-data-centre>
 - Through Infinity - has six anchor tenants, with HPC facilities
 - › Possible new shared Jisc DC in the north – looking for tenants

Articulating current/future requirements?

- » Would like to understand end user community requirements
 - › Inform future Janet network planning
- » Some communities better at articulating this than others
 - › Familiar with their data, not with the network
 - “I need to copy this 100TB data set to RAL within 2 days”
- » Example: LHC Network Future Look, 2014
 - › Tier-1 likely to rise to ~20Gbit/s by 2016
 - › Tier-2 likely to need at least 10Gbit/s

What data transfer tools are being used?


- » Existing communities are using fairly common sets of tools to transfer large data sets
- » Includes:
 - › GridFTP - <http://toolkit.globus.org/toolkit/docs/latest-stable/gridftp/>
 - › BBCP - <https://www2.cisl.ucar.edu/docs/transfer/bbcp>
 - › FDT - <http://monalisa.cern.ch/FDT/>
 - › HPN ssh/scp - <http://www.psc.edu/index.php/hpn-ssh>
- » Can we improve on these? And make them easy to use

Network monitoring

- » What are the most appropriate tools to measure performance?
 - › Preferably as close to the application end points as possible
- » Currently, perfSONAR is widely used - <http://www.perfsonar.net>
 - › Includes some 'standard' tools, like owamp, bwctl, iperf
 - › Can script tests and create performance matrix across a community
- » But what other systems are appropriate, or better?
- » Some interesting examples of late
 - › RIPE Atlas Project - <https://atlas.ripe.net/>

Janet network visualisation?

- » Janet has its Netsight system - <http://netsight.ja.net/>
 - » Includes data on link capacities and link usage
 - » Therefore can produce weathermap views
- » Question - how to make views available to users?
 - » Would like to help in expectation management, and also understanding ongoing issues
 - » Some concerns about making full weather map public
 - » Could we create tailored, per-community views?



What methods are available to improve performance?

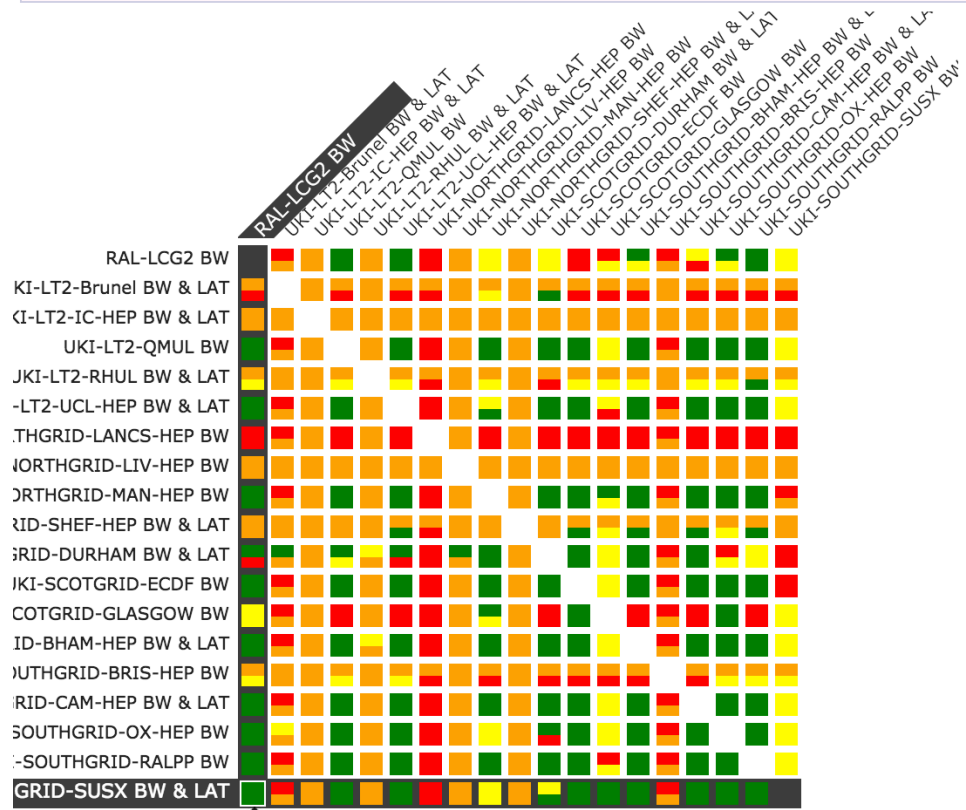
- » These include...
- » Engineering the backbone and access networks
 - › Site access links, appropriate external peerings
- » Tuning end systems
 - › Classic TCP tuning, disk system tuning, ...
 - › Some very good guidance at <https://fasterdata.es.net/>
- » Site network architectures
 - › Many problems lie within the end sites – switches, middleboxes, ...

Example site issue - Firewall impact

- » Some firewall architectures not suited to high throughput flows
 - › Or may be applying IDS/etc to those flows
- » This can often place a cap on performance
 - › Result is that sites increasingly 'bypassing' their firewalls
 - › Recent examples – Sussex, Durham (now filling their link)
- » ESnet approach in US is to promote a 'Science DMZ'
 - › Sites altered architectures to provide high performance 'onramp'
 - › See <https://fasterdata.es.net/science-dmz/>



UK Config - UK Cloud BWCTL Mesh Test



Average throughput is 0.941Gbps
Average throughput is 0.928Gbps

UK Cloud BWCTL Mesh Test



Source: **lcgps02.gridpp.rl.ac.uk** - 130.246.176.110
Capacity: Unknown MTU: Unknown


Destination: **grid-perfsonar.hpc.susx.ac.uk** - 139.184.80.18 [\[traceroute\]](#)
Capacity: Unknown MTU: Unknown

[Link to this chart](#)

Zoom: 1d 3d 1w 1m 1y

Wed Jul 2 11:33:33 2014 -- Thu Jul 2 11:33:33 2015

[Previous 1y](#)



How can you get involved?

- » Carrot: the ezepi project may be an opportunity for network researchers to find willing application domains for their work
- » Share any success stories you have, & new applicable research
- » Join the ezepi Jiscmail list:
 - › www.jiscmail.ac.uk/cgi-bin/webadmin?Ao=E2EPI
- » Come to the open ezepi workshop (date tbc – check the list 😊)
 - › Likely to be in London in Sep/Oct- a broad audience very welcome
- » Email me – tjc@soton.ac.uk or tim.chown@jisc.ac.uk