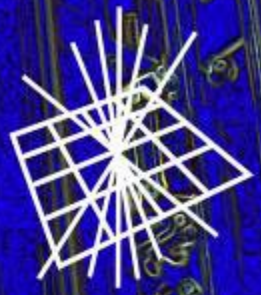

GridPP

UK Computing for Particle Physics

UK IPv6 testing

Christopher J. Walker, Duncan Rand

Queen Mary
University of London

Overview

- Need for IPv6
- IPv6 Perfsonar testing
- The QMUL Experience
- UK IPv6 testing
  - Imperial
  - Glasgow
  - Brunel
  - Oxford
- Advert for Pre-GDB
- Conclusions

Why IPv6?

- Why do we need IPv6?
  - IPv4 addresses running out
- But we've got plenty of IPv4 addresses?
  - CERN doesn't
    - Expect to run out in 2015
  - Some other Tier-2 sites short
- Solution: Dual stack for world facing services
  - Notably storage

QMUL Experience

- College running IPv4 in 6 tunnel from China
  - Increase bandwidth to QMUL for BUPT students
- 2013 Grid cluster asks for some IPv6 space
  - Takes ~2 weeks
    - Allocated a /48 (from college /29)
  - Deploy RIPE atlas probe (http://atlas.ripe.net)
  - Perfsonar dual stack
- 2014 College backbone network migration
  - IPv6 moves after IPv4
    - then doesn't work to within college
  - Can we turn IPv6 off and think about how to do it properly!!!
    - NO!!!
- Current
  - Storm testing
  -

GridPP
UK Computing for Particle Physics

Atlas.ripe.net

- 5500 probes, 4413 public
  - QMUL, Oxford and CERN + many others
  - “Free”

IPv6 testing with Perfsonar

- What to test?
  - Dual stack?
    - Default
      - Test IPv6 if available, IPv4 if not
    - Both (Better)
      - Currently issues in the database
      - Workarounds in place
- Hepix mesh config – IPv4/IPv6 explicitly specified.
  - Duncan.rand@imperial.ac.uk

Maddash Latency summary

- http://netmon02.grid.hep.ph.ic.ac.uk:8080/maddash-wel

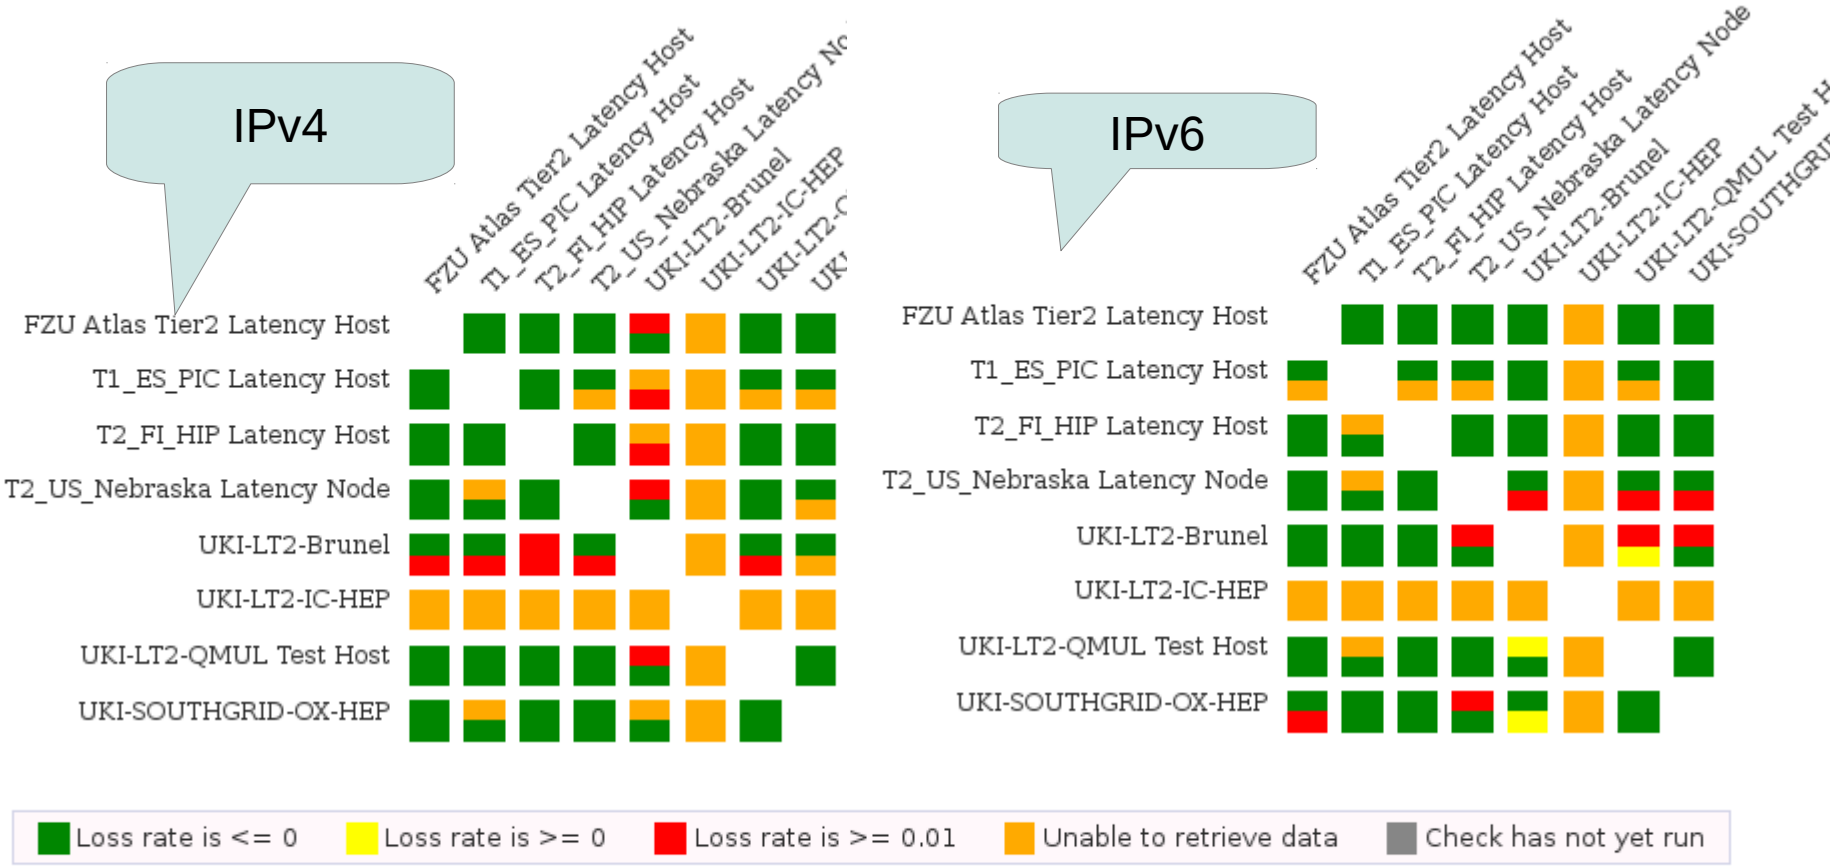

Maddash Bandwidth summary

- http://netmon02.grid.hep.ph.ic.ac.uk:8080/maddash-wel

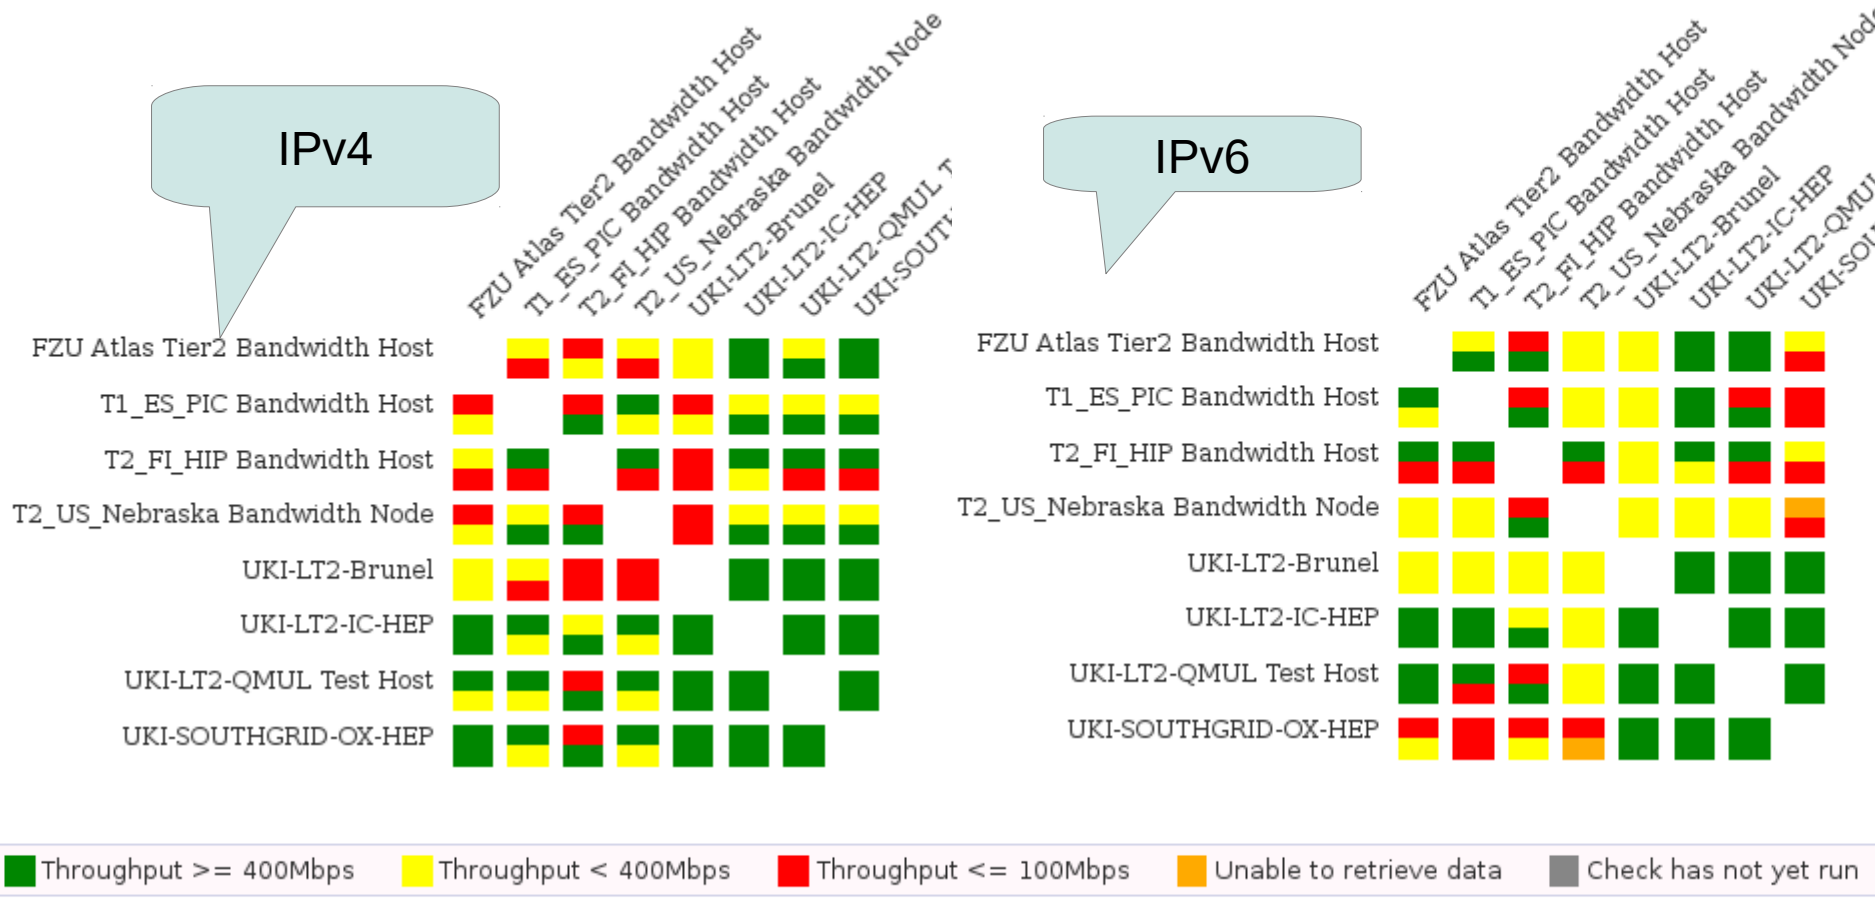

IPv4 versus IPv6 throughput

- Most sites have similar performance with IPv4 and IPv6
- Oxford is a notable but not unexpected exception: IPv6 rates around 0.6 Gbps

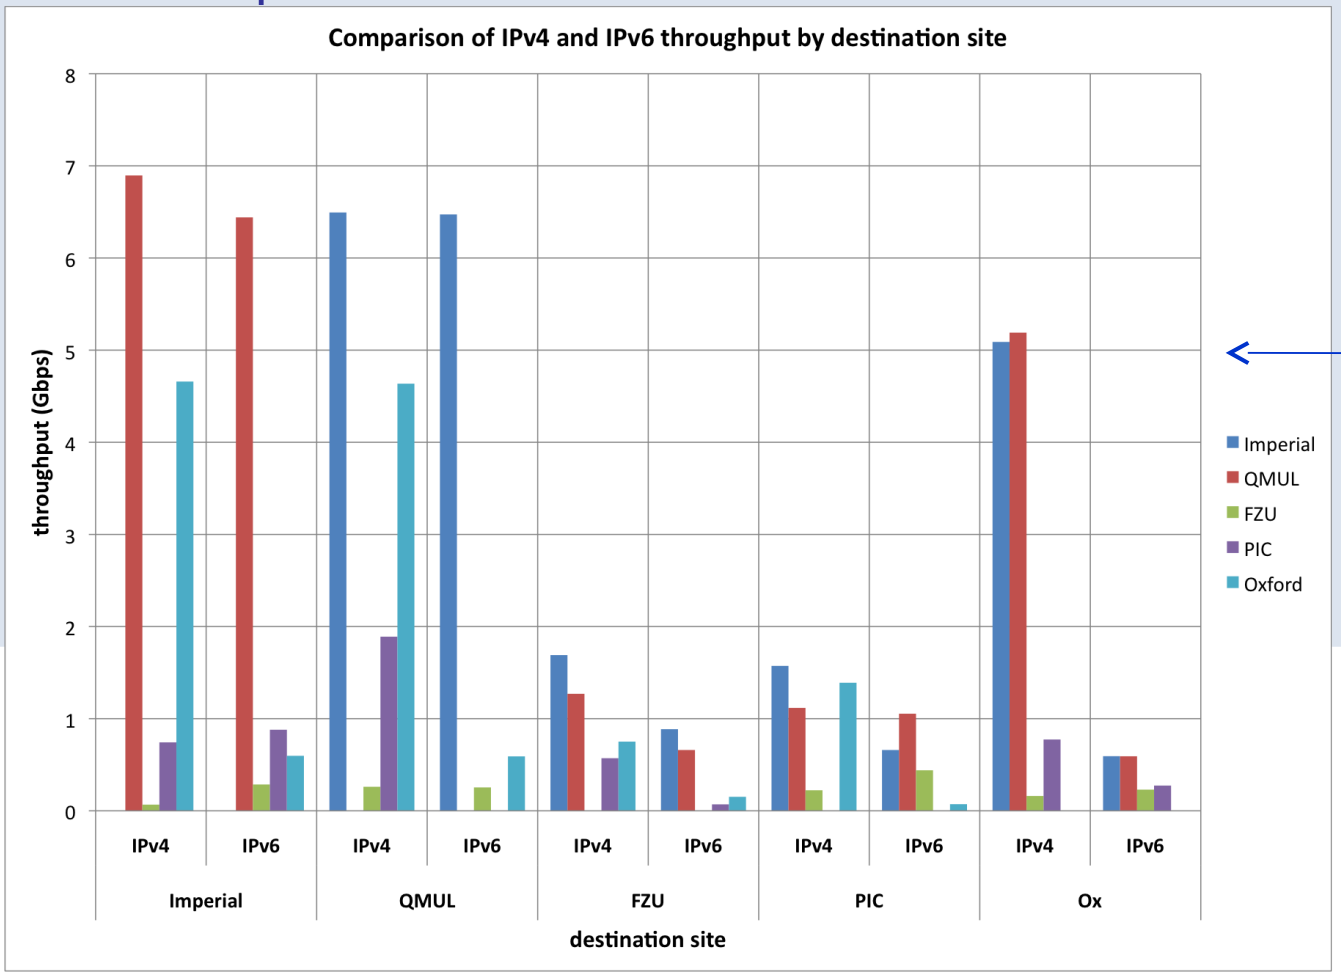

Oxford IPv4

Oxford IPv6

Issues found (IPv6 in software)

- IPv6 lower bandwidth than IPv4
- QMUL → Oxford
  - 4.431 Gbps (IPv4)
  - 0.575 Gbps (IPv6)
- Oxford → QMUL
  - 4.419 Gbps (IPv4)
  - 0.58 Gbps (IPv6)
- Oxford
  - Known issue
  - IPv4 in hardware, IPv6 in software on router.
  - Will be fixed on hardware refresh.

Issues Found - Routing

- QMUL → FZU
  - IPv4 Jumbo frames
  - IPv6 Jumbo frames
- FZU → QMUL
  - IPv4 Jumbo frames
  - IPv6 – MTU=1500
- IPv6 routing different from IPv4 (and suboptimal)
  - Geant/Czech NREN issue (now fixed)

ICMPv6 blocked

- QMUL (MTU=9000) → Brunel (MTU=1500)
  - ICMPv6 blocked by site firewall
  - Performance improves when ICMP unblocked

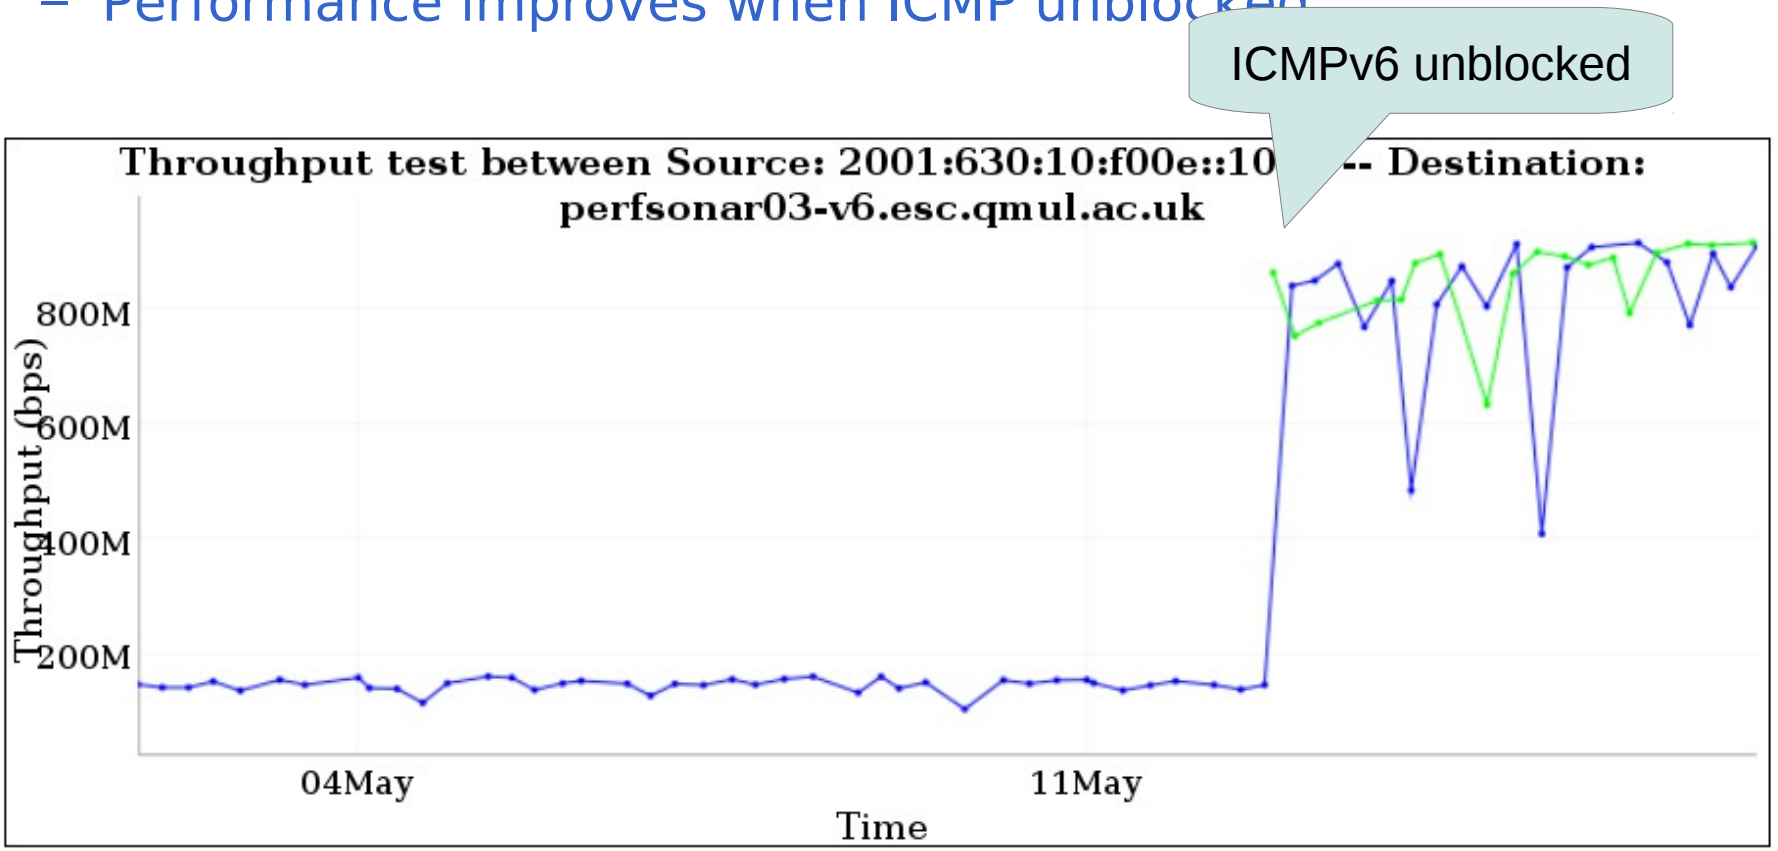

UK site status

- Different sites differently ready:
  - Not yet asked for IPv6 addresses
  - “On the local networking team's vast to do list. “
  - All services dual stack except dCache (Imperial)
- IPv6 allocation
  - /64 (One subnet)
  - /48 (65,536 subnets)
- IPv6 testing
  - Imperial (Simon Fayer and Duncan Rand)
  - QMUL (Chris Walker)
  - Brunel (Raul Lopes)
  - Oxford (Ewan McMahon)
  - Glasgow (Mark Mitchell)
  - RAL (Alastair Dewhurst – Atlas, Dave Kelsey - Hepix)
    - No IPv6 connectivity yet,
- https://www.gridpp.ac.uk/wiki/IPv6_site_status

Site testing

- Imperial
  - All production services except dCache dual stack
  - Some IPv6 only nodes
  - Test StoRM and DPM
  - Openstack Havana – attempted, but failed
- QMUL
  - Perfsonar, RIPE probe,
  - VM host with StoRM test instance (xrootd and gridftp soon?)
- Brunel
  - Perfsonar, Test DPM
  - Test cluster planned – keen to see if security/performance benefits
- Oxford
  - Test DPM, StoRM, Login node
  - Squid

GridPP
UK Computing for Particle Physics

QMUL IPv6 status

- IPv6 VLAN
- Dual-stack perfSONAR test node (perfsonar03)
- ATLAS RIPE probe
- Plans to install IPv6 test SE and CE VMs
- Contemplating an IPv6 only version of Steve Lloyd’s network tests (possibly using the ipv6.hepix.org VO)

Pre-GDB

- Pre GDB 10 June
  - https://indico.cern.ch/event/313194/
  - Aimed at Tier-1s, but everyone welcome
  - Please sign up
- Agenda
  - Set up IPv6 on your laptop (hands on)
  - IPv6 and IPv6 for sysadmins
  - File transfer
  - Experiment plans
- Organised by Hepix IPv6 working group
  - http://hepix-ipv6.web.cern.ch/
  - More people welcome

Conclusions

- Significant UK involvement in IPv6 testing
  - Contributing to Hepix working group
- Lots more to do
  - Hope to be able to help other UK sites
- Perfsonar
  - Good place to start
- Things unlikely to “just work”
  - Starting sooner rather than later a good idea