



Queen Mary
University of London

IPv6 at QMUL for HPC/HTC

Christopher J. Walker
<C.J.Walker@qmul.ac.uk>

Research IT
Queen Mary University of London



Overview

- Background
- Motivation
- Deployment
- Conclusions




QMUL in numbers

- 25,000 students
- 4,500 staff
- £428m annual income (£144m research)
- 4 campuses in East and Central London
- Russell group member




- Science & Engineering
- Human & Social Sciences
- Medicine & Dentistry
- QMUL physicists discovered Proxima Centauri B
- Research into Tamoxifen breast cancer treatment
- Hosts Genomics England

Overview of an HPC cluster



Motivation

- Collisions 25ns
 - 100PB/year
- QMUL
 - Small fraction
 - Currently 5PB storage



Motivation

- Worldwide project
- CERN running out of IPv4 space
 - VMs
- CERN collaborators running out of IPv4
 - Cloud
- QMUL capable of IPv6 (in 2012)

WLCG Sites

GridPP

- 19 UK Universities
- CERN comp




Timeline

- 2001(ish) IPv6 capable routers
- 2012 Network upgrade money (GridPP)
 - 10Gbit WAN upgrade
 - Perfsonar
 - Bandwidth and Latency and test
 - IPv6 and jumbo frames

QMUL Network Traffic

Queen Mary and Westfield College, Mile End Road [2]



	To Site	From Site
Mean	969Mbit/s	121Mbit/s
Max	985Mbit/s	761Mbit/s
Min	716Mbit/s	9.62Mbit/s
Last	985Mbit/s	11.9Mbit/s

Times on graphs are expressed in your local timezone.
Queen Mary University of London, Mile End Road



	To Site	From Site
Mean	2.65 Gbit/s	7.2 Gbit/s
Max	51 Gbit/s	18.1 Gbit/s
Min	0.05 Gbit/s	0.02 Gbit/s
Last	0.528 Gbit/s	11.2 Gbit/s
Current Link Speed	20 Gbit/s	

www.qmul.ac.uk



/QMUL



@QMUL

Queen Mary University of London

Migration strategy

- New VLAN
 - Dual stack (and jumbo frames)
- Network test (Perfsonar and RIPE atlas)
- WAN facing hosts
 - Dual stack when redeployed

RIPE atlas probe

- V1 probe - #4481
 - Deployed 2012
- V3 probe
 - For HPC in Slough
 - Deployed 2016



Issues encountered

- | Routes
 - | May be different to IPv4
 - | Geneva ->QMUL via New York (fixed)
- | Software (IPv6) / ASIC (IPv4)
 - | Older routers may give poor performance
- | Preferred over IPv4
 - | If IPv6 address (AAAA record) in DNS, it will be used by machines that think they are IPv6 connected.
- | Blocked differently by firewalls

Issues Encountered

- Routes
 - IPv4 may be different to IPv6
 - IPv6 traffic to CERN via New York
- Preferred over IPv4
 - If IPv6 address (AAAA record) in DNS, it will be used by machines that think they are IPv6

Apocrita move to Slough

- Network IPv6 ready
- Everything suddenly has an IPv6 address
 - SLAAC
 - Multi-homed hosts have two
 - External
 - Internal



Current Status

- GridPP
 - 100% dual stack
 - 30% of traffic IPv6
- University HPC
 - 100% IPv6 capable
 - Not advertising AAAA yet

Conclusions

- Two clusters, two different risk appetite
 - GridPP
 - Other sites to fall back on
 - 30% of traffic IPv6
 - Apocrita
 - Institutional HPC cluster

