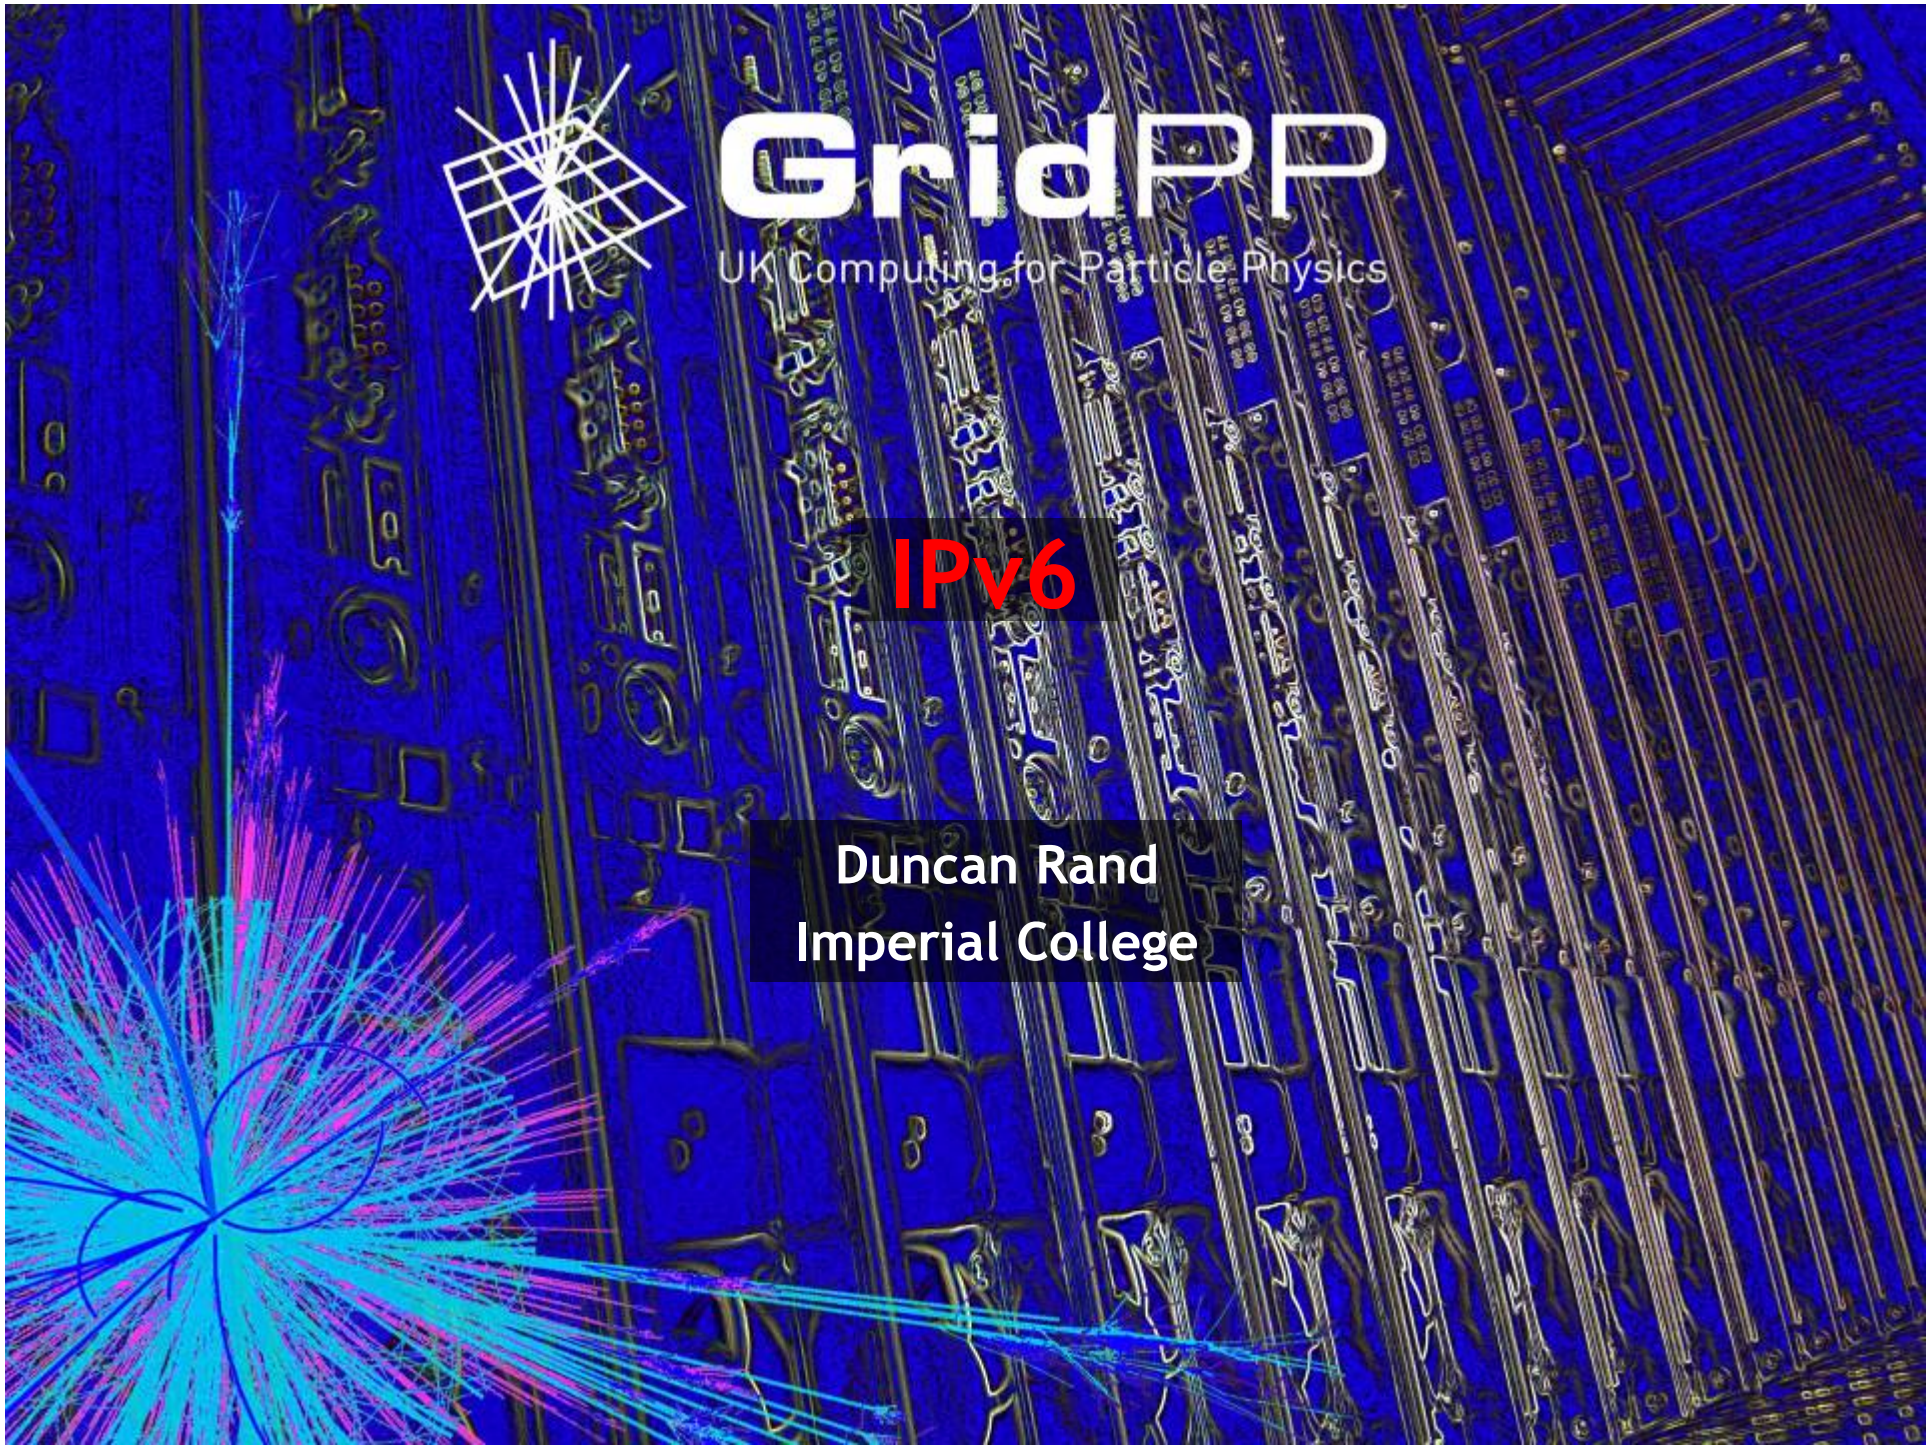
GridPP
UK Computing for Particle Physics
IPv6
Duncan Rand
Imperial College

- Reporting work done by Simon, Chris, Ewan, Raul, Alastair, Sam, Gareth, David and others
- Summary of where we are and where we are going
- What I'm aware of, there maybe other work ongoing
- Aim to raise points for discussion

IPv4 addresses are running out

IPv4 Address Report

This report generated at 07-Feb-2014 08:57 UTC.

IANA Unallocated Address Pool Exhaustion:
03-Feb-2011

Projected RIR Address Pool Exhaustion Dates:

RIR	Projected Exhaustion Date	Remaining Addresses in RIR Pool (/8s)
APNIC:	**19-Apr-2011** (actual)	0.8062
RIPE NCC:	**14-Sep-2012** (actual)	0.8435
LACNIC:	**01-Jan-2015**	1.2013
ARIN:	**18-Mar-2015**	1.4164
AFRINIC:	**31-Aug-2021**	3.2226

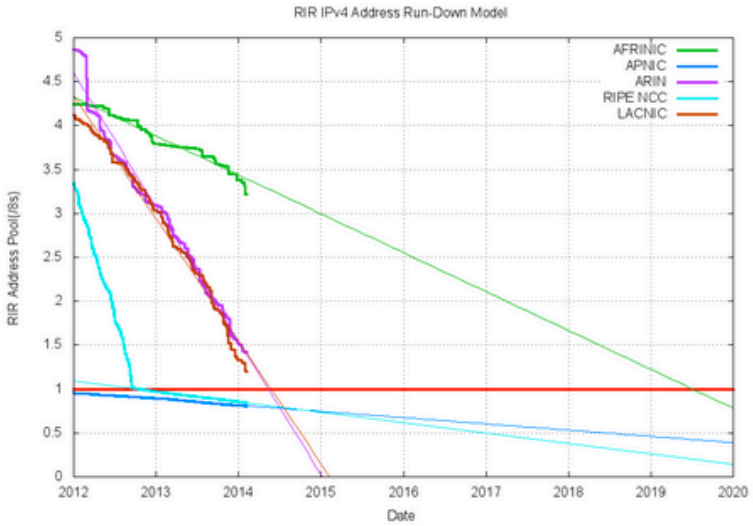

Projection of consumption of Remaining RIR Address Pools

Situation

- CERN's Tier-0 needs 150K virtual machines (worker nodes) and there aren't that number of IPv4 addresses available
- Current virtual machine adoption plan will cause IPv4 depletion during 2014-2015
- Considering two options
  - VMs with only public IPv6 addresses
  - VMs with private IPv4 and public IPv6 addresses
- More details from Edoardo's slides:

  https://indico.cern.ch/getFile.py/access?contribId=2&sessionId=0&resId=1&materialId=0&confId=251056
- Other IPv6 resources may also become available
  - opportunistic resources or grid sites from areas where IPv4 addresses have run out

GridPP
UK Computing for Particle Physics

Google IPv6 traffic

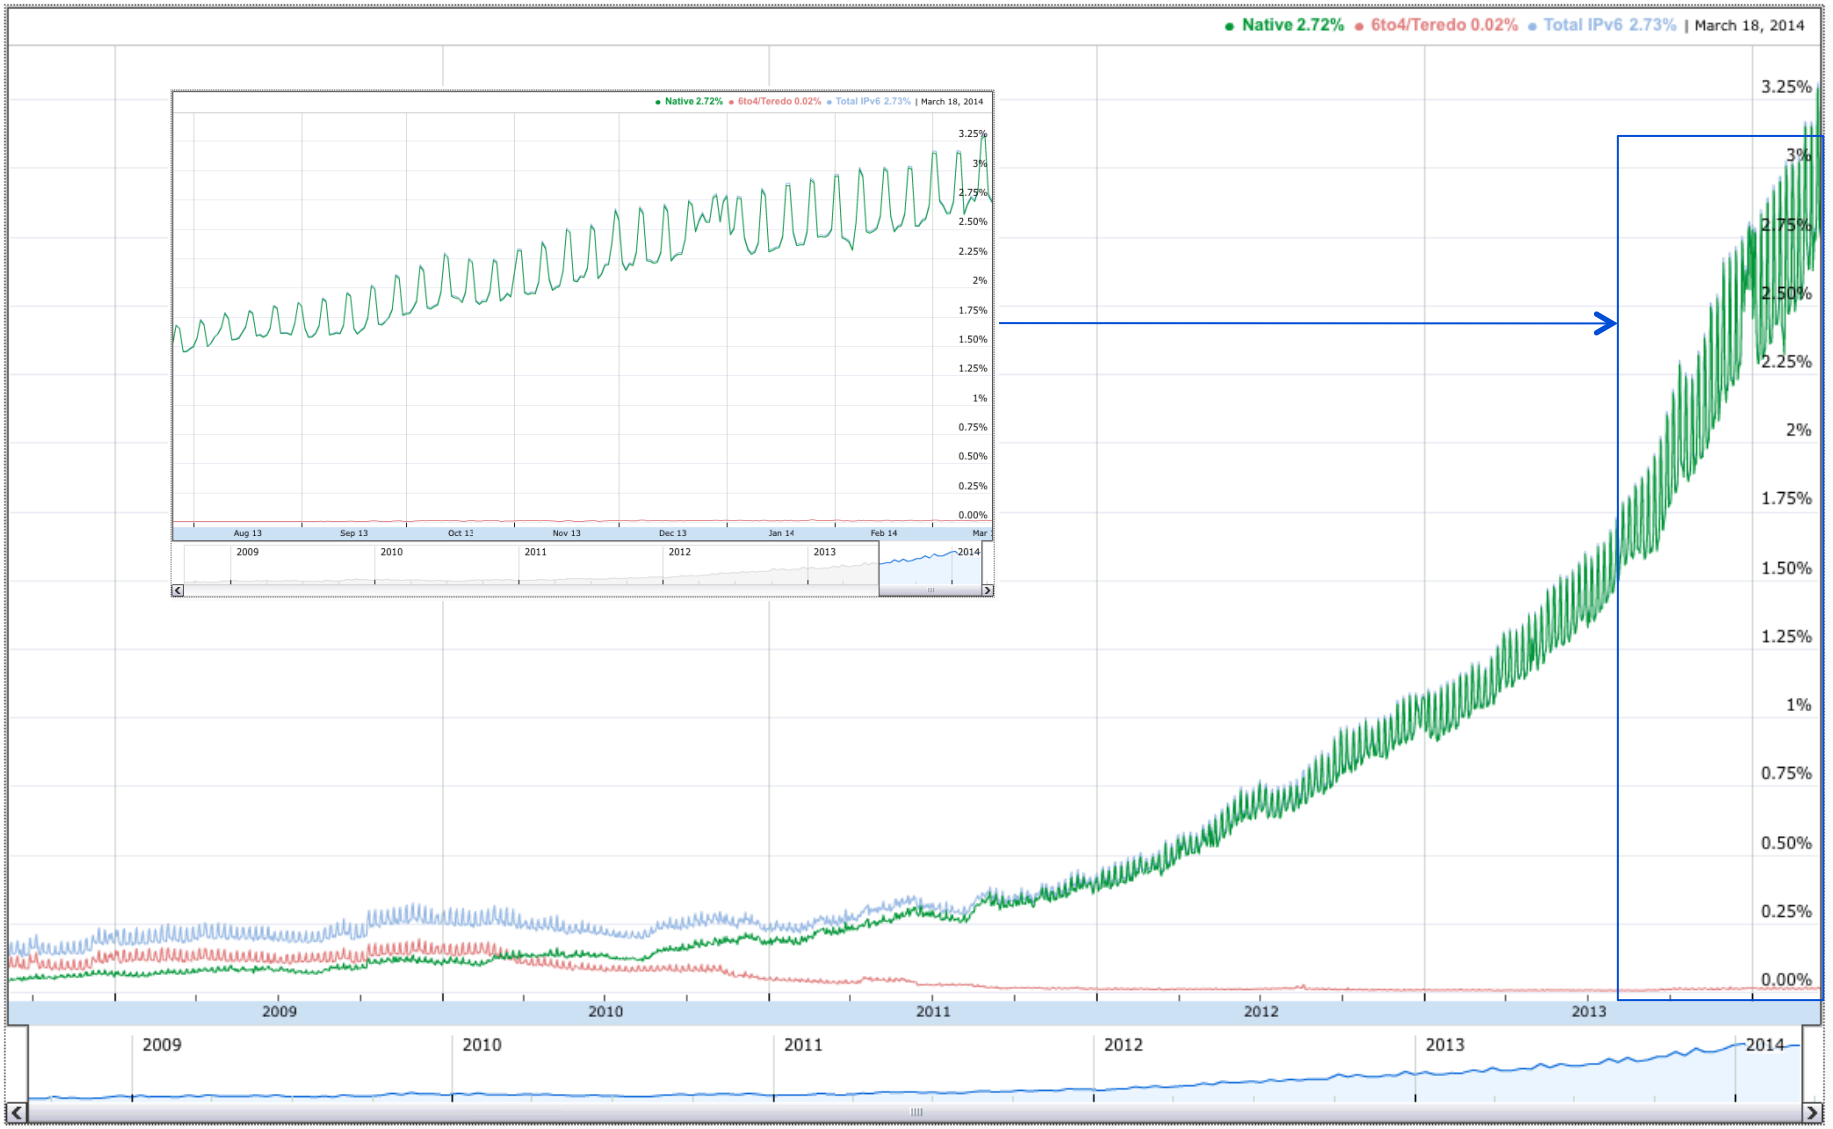

HEPiX IPv6 working group

- Meets monthly by Vidyo

- Quarterly F2F meetings at CERN
- Next F2F meeting 10/11 April 2014

- Oxford have now joined and we now have a strong UK Tier-2 contribution

- Alastair Dewhurst is the ATLAS representative
- More participants welcome

- Join email list via e-groups or ask Dave Kelsey
- Pre-GDB in June

- http://hepix-ipv6.web.cern.ch/

HEPiX WG testing

- GridFTP testbed fell slightly into disrepair, transfers failing for various reasons (now mostly fixed)
- Likely because the hosts are not production hosts
- Led to a discussion on what to do next
- Decided the emphasis should be more on production services with volunteer sites
- Some sites concerned about compensation if their availability suffers whilst testing IPv6
  - Simone Campana preparing a paper for WLCG MB (analogy with middleware testing)

IPv6 compliance of WLCG applications

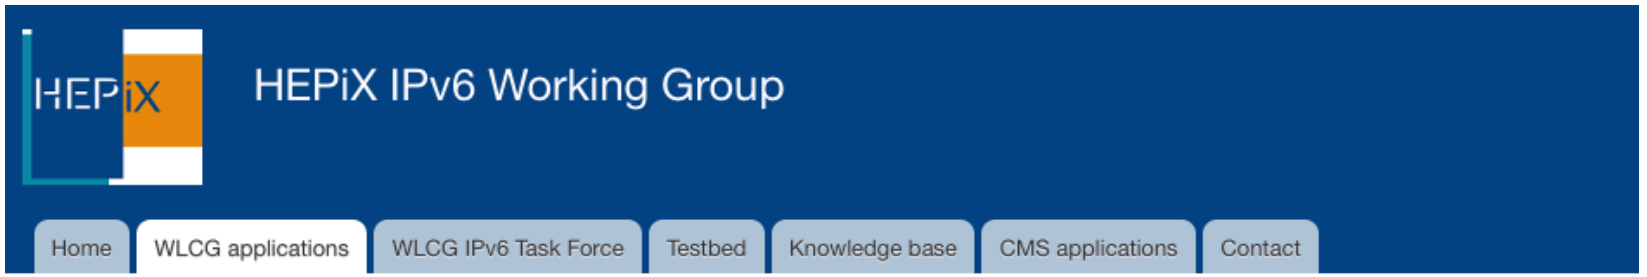

Home

IPv6 compliance of WLCG applications

Software Component	Type	Used by Experiment	Version	IPv6 Compliance
AliEN	LHC Experiment Application	ALICE		
ARC CE	Middleware	ATLAS, CMS		YES
ARGUS	Middleware	ALICE, ATLAS, CMS, LHCb		Unknown
BDII	Middleware	ATLAS, CMS, LHCb	EMI 2	YES
BestMAN	Middleware	ATLAS, CMS		Unknown
CASTOR	Middleware	ALICE, ATLAS, CMS, LHCb		NO
cfengine	Monitoring			Unknown
CMS Tag Collector	LHC Experiment Application	CMS		Unknown
CMSSW	LHC Experiment Application	CMS		Unknown
cmsweb	LHC Experiment Application	CMS		Unknown
CRAB 2	LHC Experiment Application	CMS		Unknown
Cream CE	Middleware	ALICE, ATLAS, CMS, LHCb	1.16.2	YES

http://hepix-ipv6.web.cern.ch/wlcg-applications

WLCG IPv6 Task Force

- Formed by Andrea Sciaba and Simone Campana, I am on it
- Working closely with HEPix working group
- Developed a set of scenarios and use cases
- Reports to WLCG operations coordination meeting
- Web page: https://twiki.cern.ch/twiki/bin/view/LCG/WlcgIpv6

ATLAS and IPv6

- Alastair Dewhurst has recently joined HEPiX WG/ WLCG TF for ATLAS
- Raising IPv6 awareness within ATLAS Distributed Computing (ADC)
- "We will have to use some IPv6 resources in Run 2"
- Does ATLAS software work with IPv6 machines?
  - Panda client, pilots, squid, storage, FAX
- Aiming to make CERN-based ADC development servers dual-stack so as to start testing compatibility
- 3 IPv6 SEs added to AGIS:
  - UKI-SCOTGRID-GLASGOW_ATLASIPV6TEST, UKI-SOUTHGRID-OX-HEP_ATLASIPV6TEST, prague-lcg2_ATLASIPV6TEST
- 2 IPv6 CEs (with WN behind):
  - Imperial College and Fermilab (CMS sites, but Imperial runs ATLAS MC).
- "Goal: by September 2014 to have run a task at a Grid site using IPv6 WNs"
- Talk: http://tinyurl.com/nefb79d (ATLAS restricted)
- Peter Love also involved with his ADC Dev and APF developer hats on

Imperial status and plans

- Most hosts within the group both grid and non-grid are now dual-stack
- Production grid services running dual-stack include
  - Top BDIIs
  - Site BDIIs
  - CEs (CREAM and ARC)
  - WNs
  - Squid for CernVM-FS
  - UIs
  - WMSs
  - etc.
- Considering how we might make a single WN or group of WN IPv6-only within our production cluster…

Imperial status and plans

- dCache
  - partial dual-stack (head node but not disk servers)
  - waiting for the IPv6 patch for 2.6.19 to be released
- EMI-3 DPM: v6se00 – simple IPv6-only DPM host, no pool nodes
- EMI-3 Storm: v6se01 – simple IPv6-only Storm host, no pool nodes
- Might replace the latter with a test dCache node now v2.8.0 has been released

Imperial status and plans

- EMI-3 CREAM CE: v6ce00
  - nominally IPV6-only
  - actually dual-stack but IPv4 access only allowed from WN
  - SGE communicates via (public) IPv4 addresses

- EMI-3 WN: v6wn00
  - IPv6 only
  - gridftp now works provided:
    - gridenv_set "GLOBUS_FTP_CLIENT_IPV6" "TRUE"
    - gridenv_set "GLOBUS_IO_IPV6" "TRUE”
  - needs patched uberftp also
  - CernVM-FS: can list files but not tested more than that

Imperial status and plans

- Dual-stack EMI-3 UI: v6ui00
- Recently installed HTCondor and APF (auto pilot factory used by ATLAS for pilot job submission)
  - Initially could submit jobs to dual-stack CREAM CE (Imperial) and IPv4 CE (Brunel) but not IPv6-only CREAM CE (v6ce00)
  - the cream grid manager looked old (1.12)
  - built HTCondor with cream_gahp from source loading existing UI globus libraries dynamically
  - Could submit to the IPv6 CREAM CE
  - Fixes due for release in HTCondor v8.1.5, probably in a month or two

Openstack

- We at IC tried to install Havana Openstack service with dual-stack and IPv6 only VMs
- Difficult, concluded Havana is some way from being IPv6 compatible
- CERN cloud roadmap to IPv6
  - #1. Unblock IPv6 traffic
  - #2. CERN IPv6 integration
  - #3. Neutron deployment
  - Luis Fernandez Alvarez “IPv6, CERN Cloud and OpenStack” http://tinyurl.com/q8v7yd6

QMUL status and plans

- IPv6 VLAN
  - with /48 allocated
- Dual-stack perfSONAR test node (perfsonar03)
- IPv6 ATLAS RIPE probe
- Plans:
  - Storm SE
  - IPv6-only version of Steve Lloyd's network tests (probably using the ipv6.hepix.org VO)

Brunel status and plans

- Installed a dual-stack perfSONAR host
- Dual-stack DPM/DMLite server with one pool node
  - webdav enabled
- Plans
  - build a dual-stack cluster
  - investigate DPM and xrootd when IPv6 compatible xrootd4 comes out
  - cloud storage and IPv6
  - performance evaluation of IPv6
    - potential gains for data transfers by multicasting?
    - potential gains in parallel computing due to improved MPI performance?
  - IPv6 and security

Oxford status and plans

- A dual stack UI (accounts available to known individuals who want them)
- IPv6 DPM
- Dual-stack StoRM server that's dual stacked on the same basis as the ARGUS
- Dual stack site BDII (mainly so that it can reach the node BDII on the DPM)
- A dual stack Squid 3.x that forwards to our IPv4-only 'Frontier' squid (mostly used at the moment for package updates on the IPv6-only DPM).
- Cream CE/torque server with a couple of worker nodes, all of which are dual stack.

Oxford status and plans

- Dual stack perfSONAR boxes, with protocol specific suffixed DNS entries
  - t2ps-bandwidth.physics.ox.ac.uk
  - t2ps-bandwidth-v4.physics.ox.ac.uk
  - t2ps-bandwidth-v6.physics.ox.ac.uk
- A dual stack DNS resolver so that our v6-only machines can access the University's v4-only DNS servers
  - we can publish AAAA records from the University's servers, but updating them means emailing central IT to hand edit the zone files, so they don't like us doing that too often.
- The site ARGUS server is dual stacked in that it has an IPv6 address
  - no idea yet whether it's meaningfully working over IPv6, I haven't tried it.

Glasgow status and plans

- IPv6 DPM taking part in the DPM testing
- Part of the HEPiX GridFTP testbed
- Other miscellaneous machines for internal software testing (Configuration Management/Provisioning/etc..)

- Had some issues with Reverse DNS lookups for IPv6, these need to be discussed with the Central University team
  - Central Team initially only gave us forward lookups (were asked for reverse DNS).
  - Failed reverse DNS Caused both DPM and Gridftp testbed to break periodically.
  - Unsure of production status of IPv6 DNS at present -

- IPv6 plans currently on hiatus with Mark leaving
  - Need to re-align priorities.

IPv6 network monitoring

- Some sites e.g. CERN, google etc have monitoring
- It would be nice to have monitoring of the level of IPv6 traffic within the UK or within our individual sites

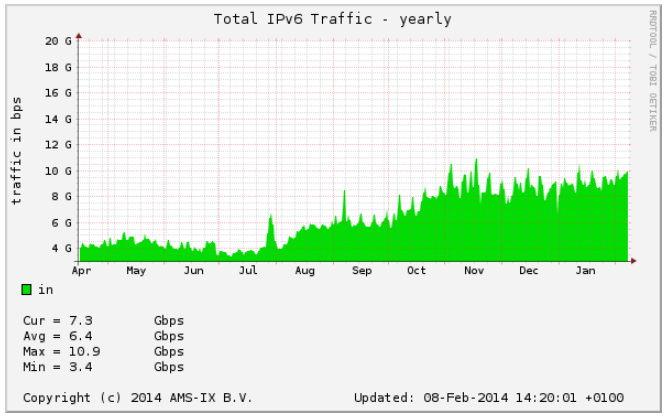

Amsterdam Internet Exchange

https://www.ams-ix.net/technical/statistics/sflow-stats/ipv6-traffic

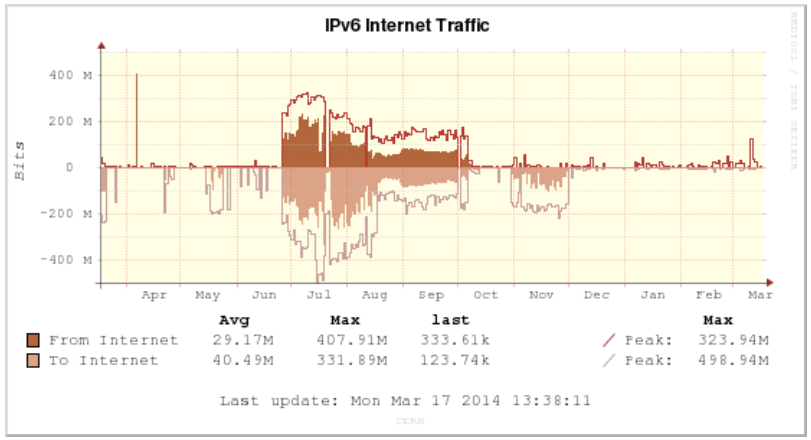

CERN IPv6 traffic

http://tinyurl.com/ok773t2

Prototype IPv6 perfSONAR dashboard

HEPiX IPv6 sites - IPv4 throughput test

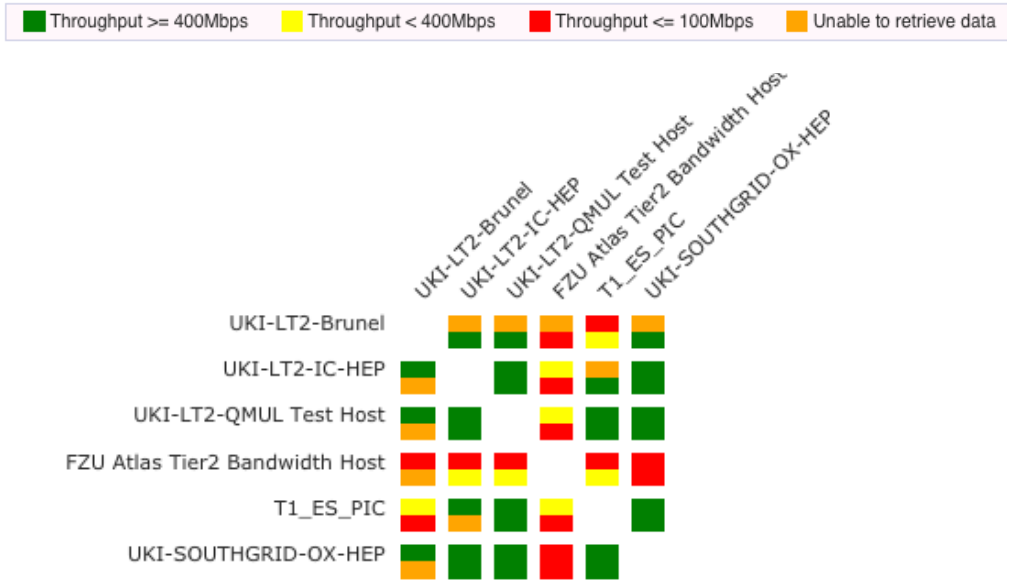

HEPiX IPv6 sites - IPv6 throughput test

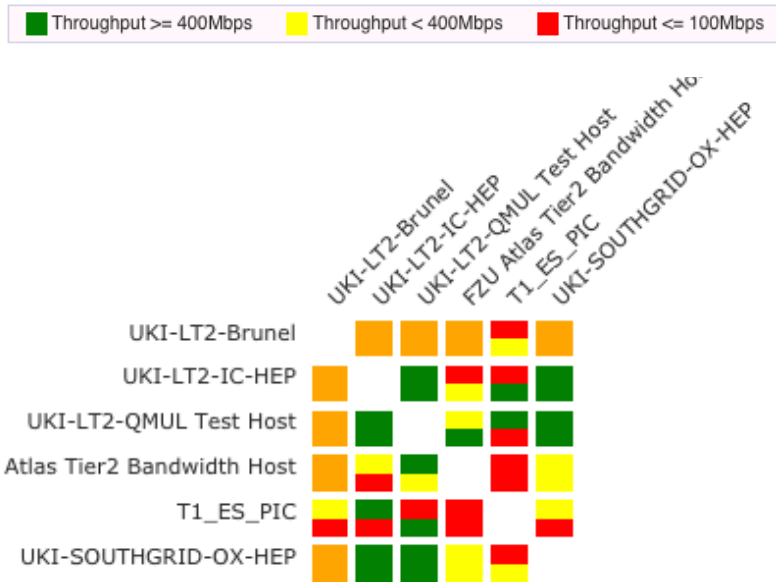

HEPiX IPv6 sites - IPv4 latency test

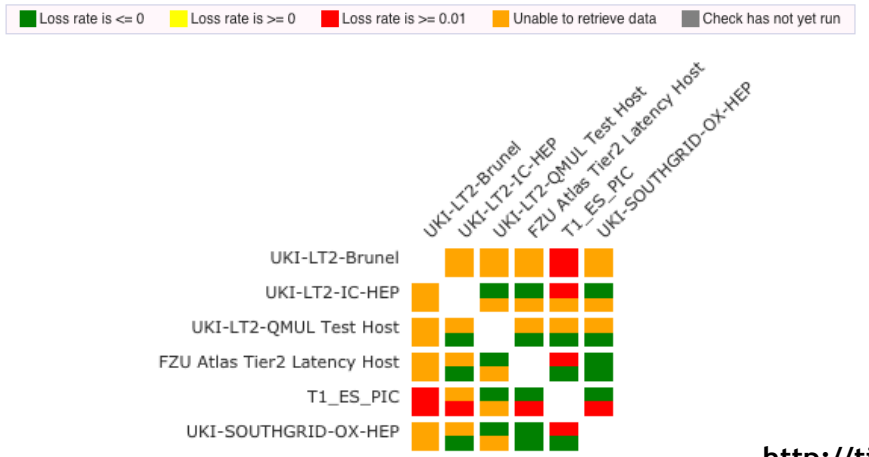

HEPiX IPv6 sites - IPv6 latency test

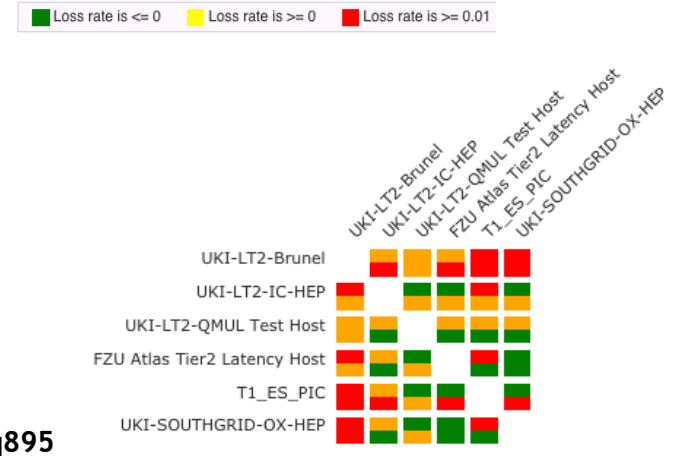

http://tinyurl.com/npsq895

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

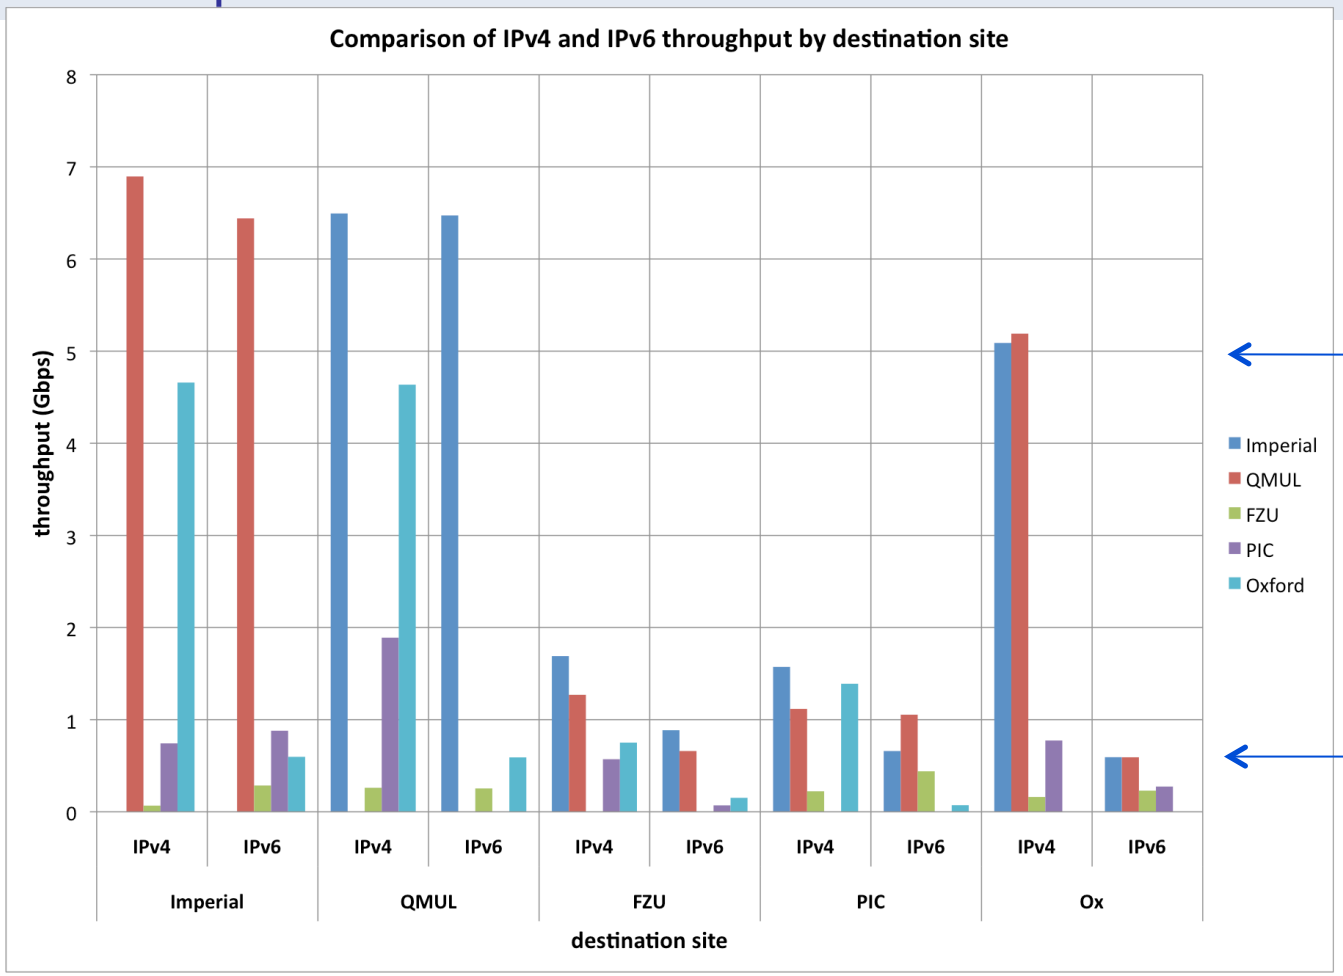

Oxford IPv4

Oxford IPv6

Dual-stack perfSONAR

- Dual-stack perfSONAR hosts favour IPv6 by default
- Oxford to PIC suffered an artificial drop in rate measured as PIC switched their perfSONAR host from IPv4-only to dual-stack

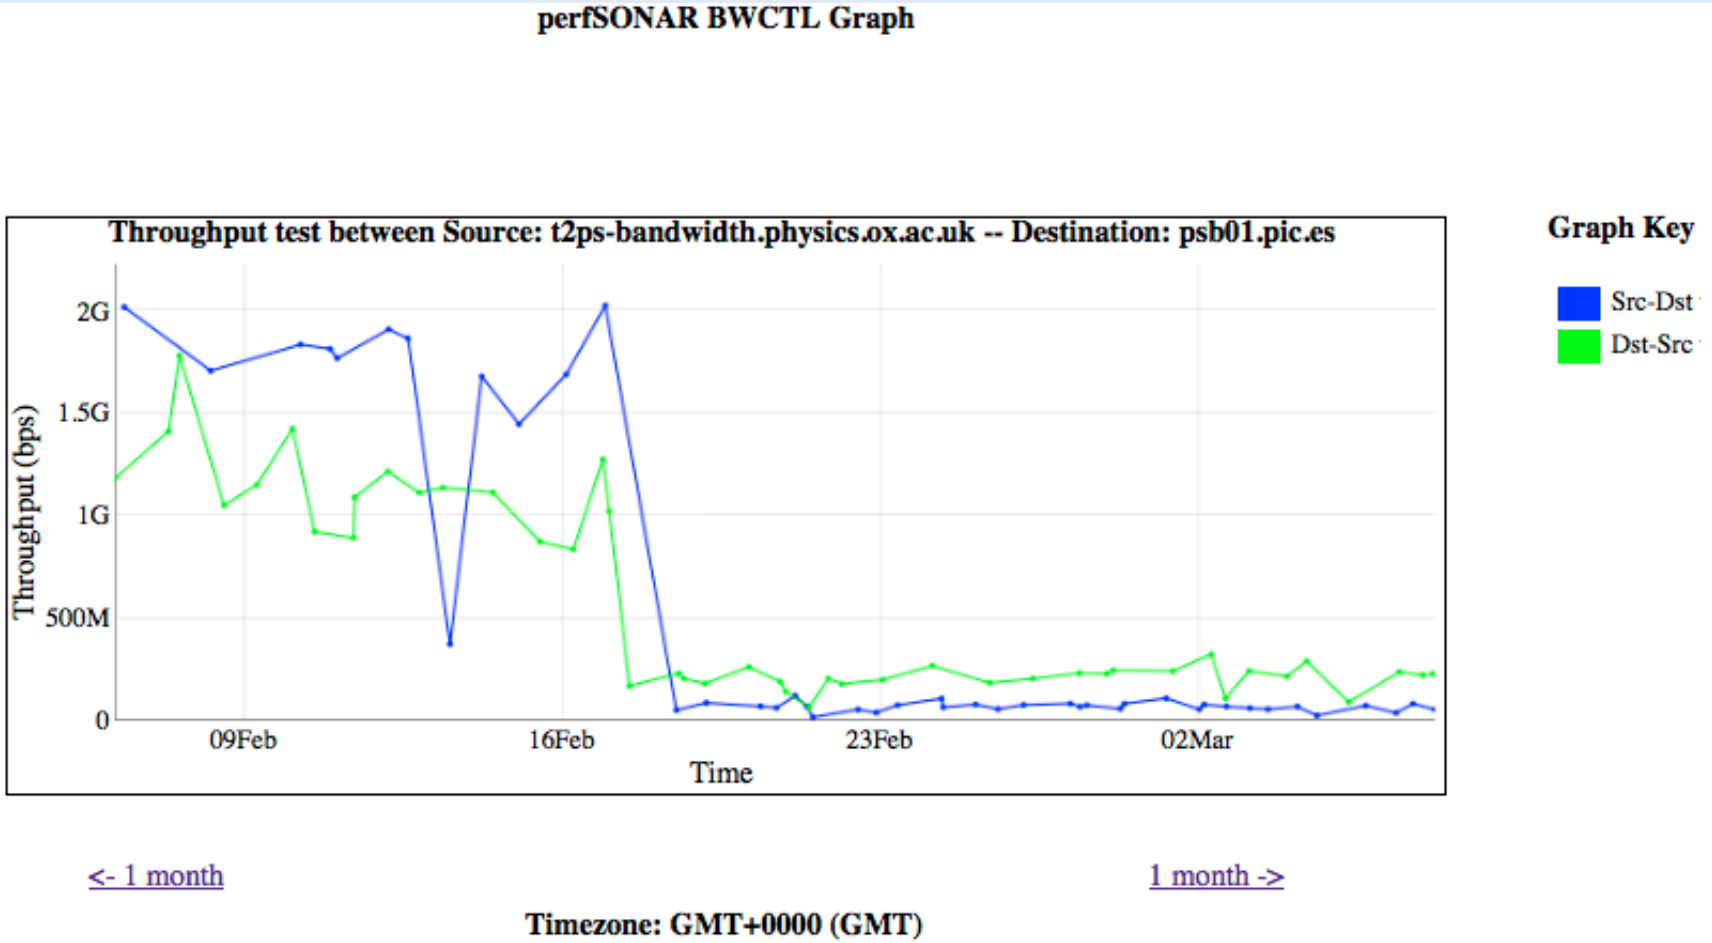

http://tinyurl.com/obzdcgu

GridPP Institutes' IPv6 status

Estimating IPv6 & DNSSEC External Service Deployment Status

NIST IPv6 and DNSSEC Deployment Monitor

- Detailed IPv6 & DNSSEC Service Interface Statistics for 20140207 -

Domain	Organization	DNS	Mail	Web
uk.ac.stfc.	RAL	[4] 1/1/1 [M]	[4] 0/0/0 [O]	[1] 0/0/0 [I]
uk.ac.brunel.	Brunel University	[5] 2/1/2 [M]	[4] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.imperial.	Imperial College London	[4] 3/3/3 [P]	[4] 4/4/0 [P]	[1] 1/1/1 [I]
uk.ac.qmul.	Queen Mary University of London	[4] 0/0/0 [P]	[2] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.rhul.	Royal Holloway University of London	[2] 0/0/0 [I]	[3] 0/0/0 [O]	[1] 0/0/0 [I]
uk.ac.ucl.	University College London	[4] 1/1/1 [M]	[6] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.lancs.	University of Lancaster	[2] 0/0/0 [I]	[5] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.liv.	University of Liverpool	[4] 0/0/0 [M]	[2] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.manchester.	University of Manchester	[4] 1/1/1 [M]	[2] 0/0/0 [I]	[1] 0/0/0 [P]
uk.ac.shef.	University of Sheffield	[3] 2/2/2 [M]	[5] 4/4/0 [O]	[1] 0/0/0 [I]
uk.ac.dur.	University of Durham	[4] 1/1/1 [M]	[4] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.ed.	University of Edinburgh	[6] 0/0/0 [I]	[3] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.gla.	University of Glasgow	[3] 0/0/0 [I]	[2] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.bham.	University of Birmingham	[2] 0/0/0 [I]	[6] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.bris.	University of Bristol	[3] 1/1/1 [M]	[4] 0/0/0 [O]	[1] 0/0/0 [I]
uk.ac.cam.	University of Cambridge	[6] 4/4/4 [M]	[1] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.ed.	University of Edinburgh	[6] 0/0/0 [I]	[3] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.ox.	University of Oxford	[4] 1/1/1 [M]	[1] 0/0/0 [I]	[1] 0/0/0 [I]
uk.ac.stfc.	RAL PPD	[4] 1/1/1 [M]	[4] 0/0/0 [O]	[0] 0/0/0 [-]
uk.ac.susx.	University of Sussex	[4] 0/0/0 [M]	[5] 0/0/0 [I]	[1] 0/0/0 [I]

For each domain and service, report:

1. number of IPv4 interfaces
2. number with IPv6 assigned
3. number of IPv6 addresses pingable
4. number of interfaces actually running over IPv6

A number of GridPP institutes apparently have some IPv6 ready DNS servers…

http://fedv6-deployment.antd.nist.gov/

Future

- Openstack: we would like to make the GridPP Cloud able to run dual-stack or IPv6 only VMs
  - monitoring CERN's activities in this area
- Security: do we want to get involved in this somehow?
- For sites interested in IPv6 the general consensus from sites with experience is that installing a dual-stack perfSONAR host might make a good starting point…
- Then consider making a few services such as site BDII, CE and a few WN dual-stack, as Imperial and Oxford have
- The WLCG VOs are starting to show more interest
- At the start of Run2 we will have less time for this than we do now

Summary

- The UK is playing a significant role within the HEPiX working group
- Imperial, Oxford, QMUL, Brunel and Glasgow have an IPv6 presence
- Critical mass of experience building up
- Alastair and Pete raising awareness of IPv6 within ATLAS Distributed Computing
- Prototype perfSONAR dashboard allows differences between IPv4 and IPv6 throughput to be identified
- For sites, perfSONAR hosts are a good place to start