



HEPiX IPv6 Working Group

David Kelsey
GDB, CERN
11 Jan 2012




Background

- IPv4 address exhaustion started (see next slides)
- But no firm requirement for WLCG IPv6-only systems
 - N.b. migration could take 2-3 years – must prepare!
- CERN and DESY (and others?)
 - site plans to support IPv6 well advanced
- US public-facing services (general, not Grid) must implement native IPv6 by 30 Sep 2012
- When can WLCG move to dual stack services?
 - To support IPv6-only systems
- More complex than just routing, naming, addressing
 - Need survey of all essential software and tools
 - Development work may be required
 - Must perform sufficient testing, including performance
 - And plan for operational impact, including training staff



IPv4 Free Addresses (/8 blocks)



<http://en.wikipedia.org/wiki/File:Ipv4-exhaust.svg>



IPv4 Addresses

- From Geoff Huston (<http://ipv4.potaroo.net>)
- IANA Unallocated Address Pool (Global)
Exhaustion happened: **03-Feb-2011**
- Projected RIR Address Pool Exhaustion Dates:
 - APNIC: **19-Apr-2011** (Asia Pacific - happened)
 - RIPENCC: **26-Jul-2012** (Europe)
 - ARIN: **19-Jul-2013** (North America)
 - LACNIC: **20-Jan-2014** (South America)
 - AFRINIC: **20-Oct-2014** (Africa)



HEPiX IPv6 WG

Created in April 2011 with aims:

- Phase 1 to consider whether and how IPv6 should be deployed in HEP
 - especially for WLCG
- Readiness and Gap analysis
 - HEP applications, middleware, security issues, system management and monitoring tools, end to end network monitoring tools
- Run a distributed HEP testbed
 - to help explore all the above issues
- Initial report at end of 2011



WG membership

- Active during 2011:
 - CERN, DESY, EPFL, FNAL, FZU, GARR, Glasgow, INFN, KIT, Manchester, RAL, SARA, SWITCH, Umea, USLHCNet (Caltech)
 - CMS
- Nearly 50 on the mail list
 - Majority are networking experts



HEPiX IPv6 WG Meetings

- <http://indico.cern.ch/categoryDisplay.py?categId=3538>
- 3 face to face meetings (June, Sep and Dec)
- 8 video meetings
- Next face to face meeting
 - 15/16 March 2012 at CERN



Working Group admin

- Mail list
 - ipv6 AT hepix DOT org
- Wiki
 - <https://w3.hepix.org/ipv6-bis/>
- For IPv6 testing we have created a new Grid Virtual Organisation
 - ipv6.hepix.org
 - VOMS services provided by INFN



IPv6 testbed

- <https://w3.hepik.org/ipv6-bis/doku.php?id=ipv6:testbed>
- Distributed Testbed well established
 - CERN, DESY, FZU, GARR, INFN, KIT, USLHCnet (Caltech)
 - Others plan to join soon
- At least one dual stack test node on public IPv4 and IPv6 networks at each site
 - Running SL5 (with valid X.509 host cert)
- IPv4 name, IPv6 name and common name in DNS
- Standard connectivity tests (ping6 etc.)
 - Script on wiki to validate the existence and config
- Recently achieved full $n \times n$ mesh of GridFTP transfers



IPv6 testbed (2)

- GridFTP tests – full $n \times n$ mesh tested
 - Usual list of configuration issues
 - Firewalls, certificates, clocks etc
- Next work
 - CMS will perform more extensive data transfer tests with PhEDEx
- In parallel with this install FTS
- Then DPM?



Asset Survey

- Perform full Asset survey
 - ~Complete list by end April 2012
- Establish IPv6 readiness of WLCG applications, middleware and tools (ongoing during 2012)
 - By code survey, testing, etc
 - Identify show-stoppers and quantify effort and resources required to fix
- What is in scope?
 - All WLCG services run on Tier 0/1/2
 - Experiment applications and services
 - Management and monitoring tools
 - Batch systems, storage systems, etc.
- What does IPv6-ready mean?
 - Works when contacted by a IPv6-only device



EGI, EMI and IPv6

- <https://wiki.egi.eu/wiki/IPv6>
- IPv6 survey of NGIs recently completed
- EGI operations (SA1) and EGI tools for operations (JRA1) have started IPv6 related activities
- EMI scans for IPv6 compliance during builds
 - We will provide EMI access to testbed
- IPv6 readiness of the tools is at https://wiki.egi.eu/wiki/EGI-JRA1_IPv6_Readiness
- All agree that we should work together



Other activities in 2011

- Started to consider security implications
- Dissemination: Submitted abstracts to 2012 conferences
 - ISGC (Feb), EGI Community Forum (Mar), CHEP(May), TNC(May)



Current status & recommendations

- Should we deploy IPv6? Answer: Yes! When we are ready.
- Focus on WLCG services (outward facing)
- Aim to implement Dual Stack on all WLCG services
 - Avoid tunnels, proxies, gateways, IPv6-only services
- Expand testbed gradually during 2012
 - continue to work with EGI and EMI
 - All WLCG services
 - Perform more extensive functionality and performance tests
- Consider operational impact, including security and monitoring
- Start to work on implementation plans
- Plan several HEP IPv6 “Days” during 2013
 - turn on dual stack for 24 hours on production infrastructure and test/observe



Issues for GDB

- Comments welcome on our plans for 2012
- Resources will be required during 2012
 - Engage with more Sites, particularly Tier 1 and non-EU
 - Need more sys admin and grid expertise
 - CPU, storage and effort for testbed
 - Engage with the other LHC experiments
 - CMS already active (Andreas Pfeiffer and Tony Wildish)
 - at least one member of the Working Group per experiment
 - initial work is the Asset Survey
- Sites will need to make their own migration plans
- Decision on implementation timescale - later
- WLCG cannot support IPv6-only systems before at least 2014