




Science & Technology
Facilities Council



GridPP
UK Computing for Particle Physics

HEP and IPv6

David Kelsey

(STFC-RAL, UK)

3 Nov 2010

HEPiX, Cornell, Ithaca NY

david.kelsey at stfc.ac.uk



Outline

- Some history
- US Federal directive
- Status of IPv4 address exhaustion
- The HEPiX Questionnaire
- Conclusions
- HEPiX and IPv6? – plans for the future



Some history

- HEP DECnet/OSI Phase V transition
 - European routing migration (early 90's)
 - Phase IV was 16-bit addressing!
- CIDR (93) and NAT (94) saves IPv4
- IPv6 - RFC2460 (Dec 1998)
- IPv6 address allocation starts July 1999
- O/S and router support from ~2000
- Many NRENs support IPv6 (~2003)
- HEPiX Spring 2008 (CERN)
 - Mattias Wadenstein (NDGF) talk on his IPv6 experience



History (2)

- HEPiX Fall 2008 (Taipei)
 - Fred Baker (CISCO)
 - “IPv4/IPv6 Transition: Status and Recommendations”
- April 2010 – Andrew Daviel (TRIUMF)
 - Email to HEPiX list (8 April 2010)
 - **What, if any, plans HEPiX institutions have for IPv6 migration?**
 - Quote: “...*monitoring, security and access control implications are somewhat scary.*”



History (3)

- Lots of discussion followed by email
 - Security, DHCP, monitoring, DNS, ...
- We decided to have a session on IPv6 at the Spring 2010 (Lisbon) HEPiX
 - Then the volcano!
- 15 April 2010 – email from Brian Carpenter (ex Head of CERN Networking)
 - Your IPv6 HEPiX session is mentioned on several IPv6 mail lists!
- Postponed session to this HEPiX meeting



History (4)

- I sent questionnaire at end of Sep 2010
 - To HEPiX Board and later to full list
- 28 Sep 2010
 - Memo from US Federal CIO
 - To all CIOs – Exec Departs & Agencies
- US DOE Labs clearly involved



US Federal transition to IPv6

<http://www.cio.gov/Documents/IPv6MemoFINAL.pdf>

- committed to the operational deployment and use of IPv6
- Enable successful deployment and expansion ... such as Cloud Computing, Broadband, and SmartGrid...
- Reduce complexity and increase transparency of Internet services by eliminating the architectural need to rely on NAT
- ...



Timelines – US Federal

- Upgrade public/external facing servers and services (e.g. web, email, DNS, ISP services, etc) to operationally use native IPv6 by the end of FY 2012
- Upgrade internal client applications that communicate with public Internet servers and supporting enterprise networks to operationally use native IPv6 by the end of FY 2014
- Designate an IPv6 Transition Manager and submit their name, title, and contact information by October 30, 2010

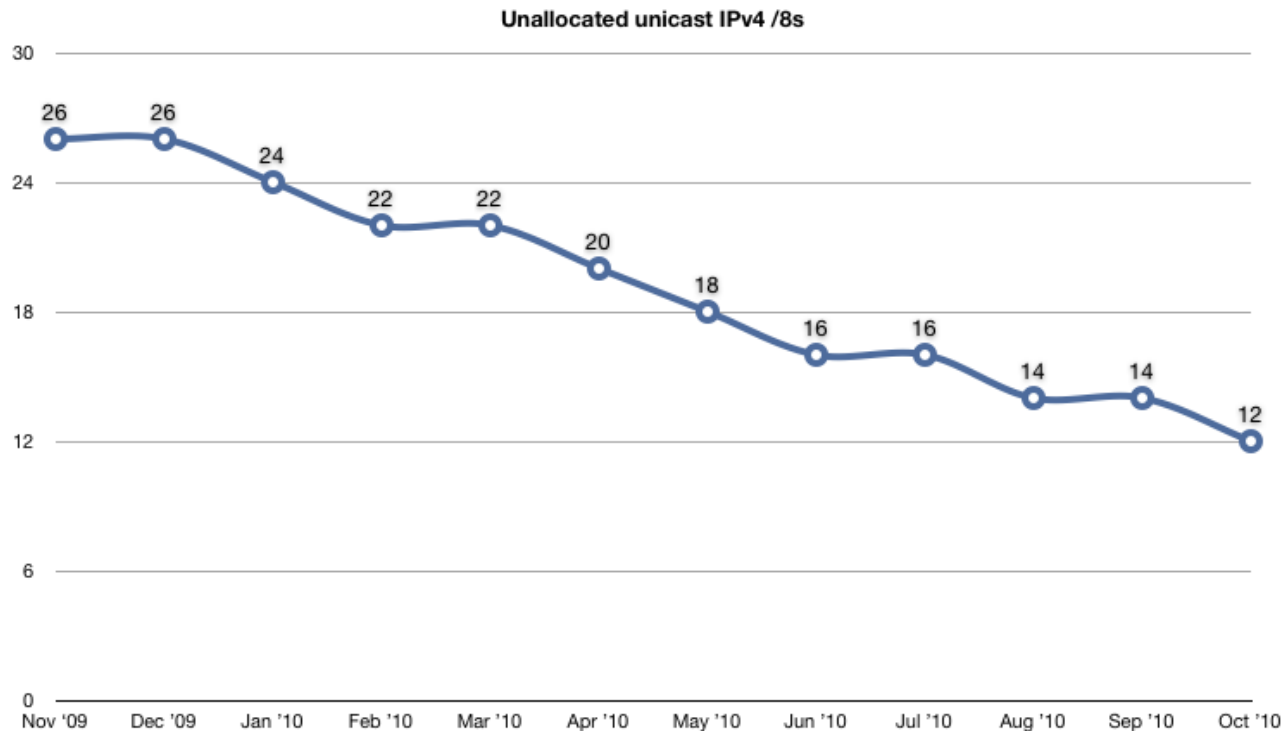


IPv4 address exhaustion

- From Geoff Huston
(<http://ipv4.potaroo.net>)
- **Projected IANA Unallocated Address Pool (global) Exhaustion: 09-Jun-2011**
 - Only 12/256 (/8) addresses are unallocated
 - < 5%
- **Projected RIR Unallocated Address Pool (regional) Exhaustion: 22-Jan-2012**



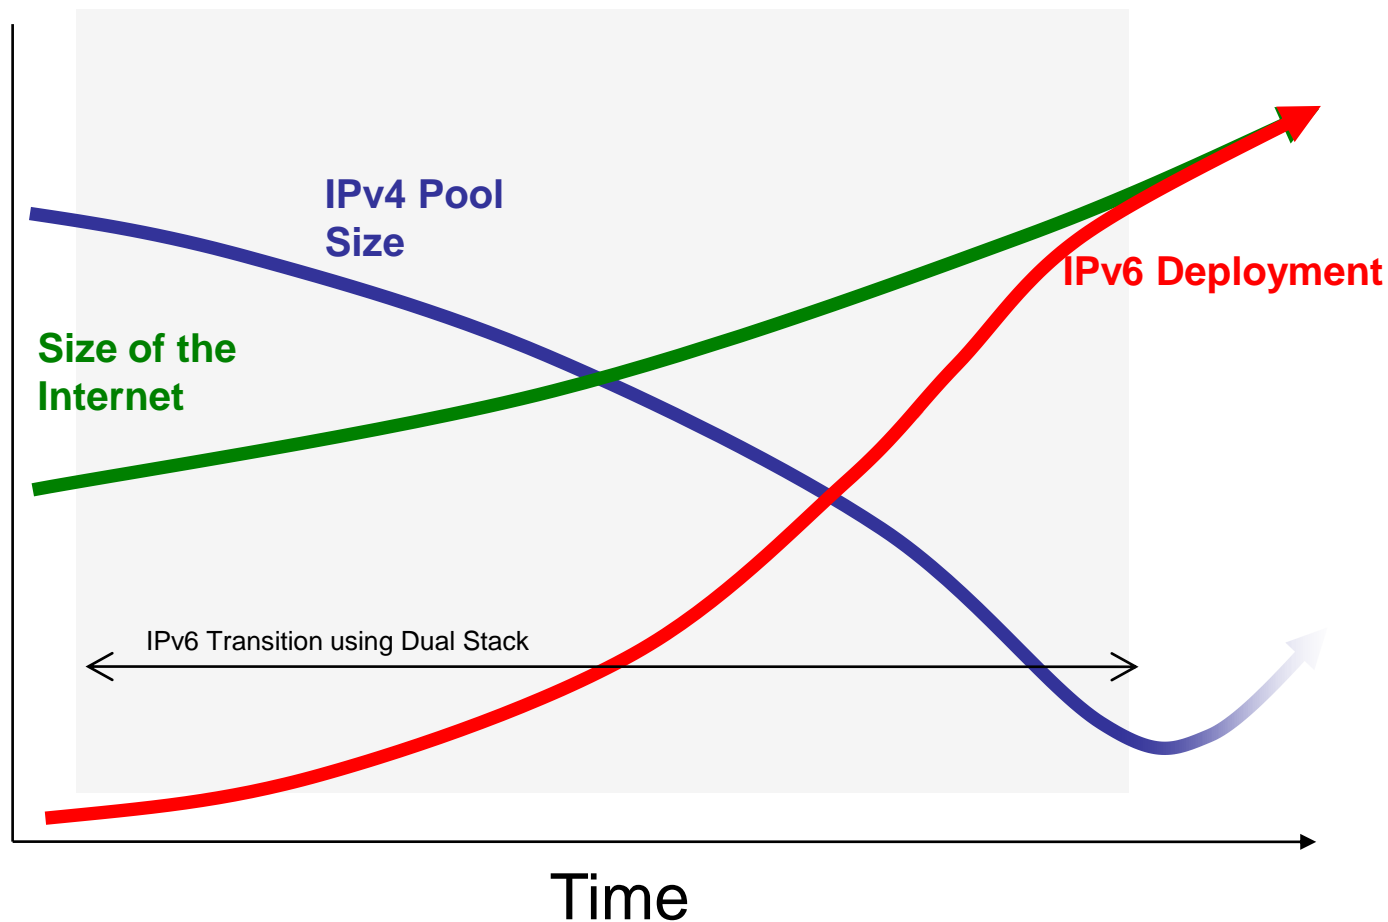
V4 address allocation last year




What's the problem?

This is how we pictured the transition 15 years ago:

*Dave
Wilson,
HEAnet
(Ireland)
TNC2010*



This is where we are **now**:





HEPiX Questionnaire

Many thanks to the 18 responders:

- Guillaume Cessieux (IN2P3), Andrew Daviel (TRIUMF), Phil DeMar (FNAL), Carlos Friacas (Portugal), Denise Heagerty (CERN), Jean-Michel Jouanigot (CERN), Andy Kowalski (JLAB), Pierrick Micou (IRFU), Francesco Prelz (INFN), Fazi Qi (IHEP, China), Sabah Salih (Manchester), Reinhard Strebler (KIT), Robin Tasker (RAL & DL), Steven Timm (FNAL), Rosette Vandenbroucke (Belgium), Mattias Wadenstein (Umea), Stephan Wiesand (DESY), Knut Woller (DESY)



Q1

- *Are you already supporting IPv6 now? If so are you using dual 4/6 stacks? Or perhaps someone is even using just IPV6 and NAT64?*
- Yes: 1 (IHEP, China)
- No: 7
- Testbed(s): 5
- NRENs: 6 Yes, 0 No
- Dual stacks everywhere, no NAT64



Q2

- *If you are not yet supporting IPv6 do you have plans? If so what are the timelines?*
- 7 sites have plans for next 2 years
 - 3 this year, 3 in 2011, 1 by 2012
 - US directive will drive DOE labs
- 3 have no plans
 - waiting for new routers or for useful applications or for user pressure



Q3

- *Is anyone you know of working on the implications of v6 for applications? Have they done a survey of which applications are V6 compliant?*
- Nobody doing anything on HEP applications
 - Concentrating on obvious apps (browser, email etc)
- gLite Grid middleware is 99% compliant (TNC2010)
- Some claim many apps are not V6 compliant
 - e.g. OpenAFS will not be compliant until (at least) Aug 2011
- Condor busy working on V6 compliance
- List of some general application status:
 - <http://www.ipv6-to-standard.org/>



Q4

- *Do you currently have problems with lack of IPv4 addresses? Or foreseen in the near future?*
- Very few problems reported (just 1 said Yes)
 - Manchester HEP has just 2 Class C nets
- Increasing amount of Virtualisation a concern
- Very large clusters may need V6
- Increasing value of V4 addresses may lead to sale of unused addresses



Q5

- *Are there other issues you are aware of? Or interesting work going on in this area?*
- No pressure, lack of effort/funding
- Concerns on monitoring, log-file analysis, security, network management tools, transition mechanisms
- Report that CMS needs at least 1-year to deploy new v6 code
- Planning must take several years
- Lower quality of vendor support
- Lack of maturity of software and tools
- Google and Facebook already publishing some content via IPv6
- Push likely to come from China, India, Africa
- Also from growth of “sensor” networks and “social engineering space”



Conclusions

- IPv6: Infrastructure is the most advanced
 - NRENs have been “ready” for years
- Sites are not seeing any pressure (yet)
 - So, not much is happening
 - But the US directive will make a difference
- The big problem areas
 - Applications
 - 3rd party and homegrown (“we” have to work on these)
 - Very much like Y2K (survey, plan, upgrade, replace)
 - Technical transition details
 - Lack of maturity of tools, other concerns (e.g.security)



Questions?

(then we will look at what to
do next)



HEPiX and IPv6?

- IPv6-only systems will (soon?) arrive
 - Certainly will affect general Grids and Clouds
 - e.g. FermiCloud decided to support IPv6, ...
 - When will WLCG/HEP have to support IPv6-only systems?
- The range of applications in one community (e.g. HEP) is more constrained than trying to solve problems for everyone
- Different from the HEP DECnet/Phase V coordination
 - Then we spent most of our efforts on the routing, addressing and naming issues
 - Now I think we have to concentrate more on applications and tools
 - But we did learn that analysis and planning is essential and takes lots of time!



HEP Coordination

- In addition to applications
 - Work with Site and National networking experts
 - Security concerns
 - Monitoring
 - Network Management (end to end, not backbone)
- Positive benefits of learning together and sharing the work
- It will all take time
 - We are already late



My proposal

- Create a HEPiX IPv6 group
- Phase 1 should be to do a full (gap) analysis of the problem (during 2011)
 - Applications important to HEP
 - Analyse IPv6 compliance
 - Security
 - System and network tools etc.
- Create and operate an HEP distributed testbed
- Propose a timetable for Phase 2
 - Upgrades, implementation, deployment
- Must include an effort/resource requirements analysis