

The HEPiX IPv6 Working Group

David Kelsey

STFC-RAL

23 Oct 2012

Outline

- The global IPv4 address exhaustion
- The HEPiX IPv6 working group
- HEP/WLCG site IPv6 readiness survey
- IPv6 testing news
- News from elsewhere (EGI, EMI, USA)
- Future plans

IPv4 Address exhaustion

- From Geoff Huston (<http://ipv4.potaroo.net>)
- IANA Unallocated Address Pool (Global)
Exhaustion happened: **03-Feb-2011**
- Regional (RIR) Address Pool Exhaustion Dates:
 - APNIC: **19-Apr-2011** (Asia Pacific - happened)
 - RIPENCC: **14-Sep-2012** (Europe - happened)
 - ARIN: **25-Aug-2013** (North America)
 - LACNIC: **07-Jun-2015** (South America)
 - AFRINIC: **21-Sep-2019** (Africa)

HEPiX IPv6 Working Group

Created in April 2011 with aims:

- Consider whether/how IPv6 should be deployed in HEP
 - especially WLCG (Worldwide Large Hadron Collider Grid)
- 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
- We meet face to face 3 or 4 times a year
 - And by video conference in between
 - Next F2F meeting at CERN, January 24-25

HEP Site IPv6 Survey

- Questionnaire sent on 26 Sep 2012
- Survey of all WLCG/HEP sites
 - WLCG GDB & CB, HEPiX-users, lcg-rollout
- Questions
 1. Is your site already offering connectivity, routing and naming services for IPv6 end systems (i.e. is your site already connected to the worldwide IPv6 network and ready to support IPv6 end systems)?

IPv6 site survey (2)

2. If so, have you already enabled IPv6 on some (all?) of the systems you manage - please provide some details (e.g. "yes, our external web server is running dual-stack IPv4/IPv6")
3. If your site is not yet IPv6-ready are there plans for this? If so what are the timelines?
4. Does your site currently have problems with lack of IPv4 addresses? Or foreseen in the near future?
5. Are there other issues you are aware of? Or interesting work going on in this area?
6. Any other comments?

IPv6 site survey (3)

- 43 sites replied to date, in 21 countries
 - Europe plus USA, Canada, China, South Africa, Brazil
- Q1: 15 sites are already IPv6 enabled
 - Some others report that University is, but HEP is not
- Q2: Applications enabled include DNS, web servers, email servers, CAs, wireless, windows domain, testbeds
 - Two sites report extensive use of dual-stack
 - Lots of end-systems enabled by default, e.g. Windows 7 and associated tunnels (& perhaps rogue Router Adverts)

IPv6 site survey (4)

- Q3: another 10 sites have plans for the future
 - Usually within next 12 months
 - Others say: no pressure, no need, other more urgent tasks, why rush?...
- Q4: Two sites (South Africa, FZU) report problems with lack of IPv4 addresses now. CERN predicts possible problems in 2014
 - Several note lots of growth from virtualisation
 - “Would be good to have real network addresses”

IPv6 site survey (5)

- Q5 and Q6
 - Concerns about IP Address Management
 - DHCPv6 vs. Autoconfiguration (ND)
 - DNS servers: Router advertisement (RA) vs. DHCPv6
 - Currently no support to configure DNS servers in DHCPv6
 - Concerns about IPv6 security
 - Memories of DECnet Phase V transition
 - We succeeded there, we can do it again!
 - Several concerns about applications or tools not being ready

Observations from survey

- IPv6 is happening – working group should continue!
 - End systems and core networks are ready
 - Applications, tools and sites are not, but are working on it
 - Cannot ignore IPv6
 - E.g. need to manage IPv6 on Windows clients and look for RAs
- Some evidence of IPv4 address shortage at sites
 - But no great pressure reported

HEPiX IPv6 Testing

- We have deployed a distributed testbed
 - CERN, DESY, FZU, GARR, INFN, KIT and USLHCnet
 - Should expand soon, e.g. UK joining
- Connected to IPv6 and IPv4 networks
 - IPv6-only/IPv4-only names also registered in DNS
 - e.g. hepix-v6.desy.de & hepix-v4.desy.de
- <https://w3.hepik.org/ipv6-bis/doku.php?id=ipv6:testbed>
- Have tested several middleware components and tools, mainly in data management

Tests for IPv6

- Does the service break/slow down when used with IPv4 on a dual-stack host with IPv6 enabled?
- Will the service *try* using (connecting/binding to) an IPv6 address (AAAA record), when available from DNS?
- Will the service *prefer* IPv6 addresses from DNS, when preferred at the host level?
 - Does this need to be configured?
 - How?
- Can the service fall back gracefully to IPv4 if needed?

Data tests – summer 2012

- Already reported last time that GridFTP and globus_url_copy works
- FTS can be made to work
 - see Francesco Prelz talk at HEPiX April 2012)
- Several FTS channels defined and used over IPv6
- Successfully installed IPv6 DPM on several nodes
- Successfully transferred data to a dual-stack DPM server over IPv6

Software & Tools IPv6 Survey

- An “Asset” survey is underway
 - A spreadsheet to be completed by sites and the LHC experiments
 - Includes **all** applications, middleware and tools
 - Tickets to be entered for all problems found
- If IPv6-readiness is known, can be recorded
- Otherwise we will need to investigate further
 - Ask developer and/or supplier
 - Scan source code or look for network calls while running
 - Test the running application under dual stack conditions

IPv6 problems found

Problems include...

- OpenAFS, dCache, UberFTP
- FTS & globus_url_copy (but can be made to work)
- MyProxy (now fixed)
- ISC dhcp on Scientific Linux 5 (works on SL6)
- ARNES/Slovenia – EGI testing
 - No LRMS system works (ARC) – SLURM, Torque, PBS, ...
- Many IGTF CA CRLs not available on IPv6
- netboot/dhcp/pxe
- Work ongoing – lots not yet assessed (many storage solutions)

IPv6 readiness *an example*

Component	Version	Component	Version
Argus	1.4.0	LFC	1.8.3
GridFTP		DPM	1.8.3
BDII	1.0.2	StorRM	1.8.2
		UI	3.2.11-1
CASTOR	2.1.13-5	VOBOX	3.2.15-1
CVFMS	2.0.13	WMS	3.3.7
dCache	1.9.12	WN	3.2.12-1
FTS	2.2.8		

Works

Rebuild or reconfigure

Does not work

Unknown

News from EGI

- Presentation at IPv6 WG 5th Oct 2012
 - Mario Reale (GARR, Italy)
- Testing middleware and tools in a dual-stack environment
 - All client-server permutations in dual stack/IPv6 only
- Testbed with 4 sites, 15 servers, ARC, gLite, UNICORE, IGE Globus
- Concentrating on Site Computing and InfoSys
- <https://wiki.egi.eu/wiki/IPv6>
- Test reports on <https://wiki.egi.eu/wiki/IPv6TestReports>
- Top BDII in place
 - Can be used to join EGI and HEPiX testbeds

News from EMI

- Presentation at IPv6 WG 5th Oct 2012
 - Andrew Elwell & Dusan Klinek (CERN)
- Testbed (virtual machines) at CERN
- Test BDII, VOMS, MyProxy, UI, Nagios, Argus, DPM, LFC, FTS, CE, WN (torque)
- BDII needed mods
- They modified NAGIOS to probe both V4 address and V6 address in dual-stack

News from the USA

- Presentation at IPv6 WG 4th Oct 2012
 - Phil DeMar (FNAL)
- US Government mandate: Public facing services to support IPv6 by 30 Sep 2012
 - But no enforcement element – certainly not yet complete
 - DOE National Labs are not in scope (but have made good progress)
 - Does not include scientific computing
- <http://fedv6-deployment.antd.nist.gov/cgi-bin/generate-gov>
- <http://my.es.net/sites/ipv6>
- Next steps: concentrate on client IPv6 issues
 - Auto-config, tunneling, dual-stack, unique local addresses (ULA)

HEPiX IPv6 testing plans for coming months

- CMS tests
 - Install and test PhEDEx between several end points
- LHCb tests
 - Glasgow & RAL
 - Test job submission and CVMFS
- Ongoing data component testing
 - dCache, STORM, CASTOR?, ...
- Merge HEPiX and EGI/EMI Testbeds
 - Install BDII (just one top-level or also site BDII?)
- Expand site resources
 - Install WNs
 - Install other Grid services
 - Test WMS and batch systems
 - using standard Grid installation

Future plans

- Continue WLCG/HEP site readiness survey, asset survey and testbed/testing
 - Review status early in 2013
 - Produce plans for LHC LS1 and/or later
- Need to perform tests on the production infrastructure
 - involve WLCG Tier 1 centres
 - Use/assess configuration mgt tools/procedures used at sites
- Plan several HEP IPv6 “Days” (for LS1?)
 - turn on dual stack for 24 hours on production infrastructure and test/observe
- Earliest date for production support of IPv6-only systems is (currently) Jan 2014

Further info

- HEPiX IPv6 wiki

<https://w3.hepix.org/ipv6-bis/>

- Working group meetings

<http://indico.cern.ch/categoryDisplay.py?categId=3538>

- Next F2F meeting at CERN, January 24-25

Summary

- **MUCH** work still to be done during the next year or three & effort is difficult to find
 - Further volunteers welcome to join
 - France in particular is a missing active member...
 - Volunteer profile: switch to service rather than network experts
 - Please contact David Kelsey
- not able to support IPv6-only systems in WLCG before 2014 (at earliest – could be later?)
 - Decision on timetable to be made during 2013
 - And needs to be jointly made with EGI, NDGF, OSG etc.