

Plans for dual-stack IPv4/IPv6 services on WLCG

(The HEPiX IPv6 Working Group)

David Kelsey (STFC-RAL)

HEPiX. Lincoln, Nebraska

15 Oct 2014

Outline

- Some IPv4 and IPv6 statistics
- Update on activities of the HEPiX IPv6 working group
- Plans for testing and the deployment of WLCG dual-stack services

- *With MANY THANKS to all my colleagues in the IPv6 working group – they deserve all the credit!*


2014 members of the WG

J. Bernier (IN2P3), K. Chadwick (FNAL), S. Campana (CERN),
J. Chudoba (FZU), A. Dewhurst (RAL), M. Elias (FZU), S. Fayer
(Imperial), T. Finnern (DESY), C. Grigoras (CERN), B. Hoefft (KIT), T.
Idiculla (RAL), D. Kelsey (RAL), F. Lopez Munoz (PIC),
E. MacMahon (Oxford), E. Martelli (CERN), R. Nandakumar (RAL),
K. Ohrenberg (DESY), F. Prelz (INFN), D. Rand (Imperial),
A. Sciaba (CERN), U. Tigerstedt (CSC), R. Voicu (Caltech),
C. J. Walker (QMUL), T. Wildish (Princeton)

And many others in earlier times


Growth in use of IPv6 & IPv4 address exhaustion

IPv6 growth (global Google clients)




Comcast in USA

- The largest US cable and home ISP
- July 22nd 2014
 - Announced completion of work to enable IPv6 on their broadband network



IPv6 traffic at AMS-IX

Total IPv6 Traffic - yearly




CERN ALICE MonALISA

ations

What is this about

IPv6 clients to total request count ratio



Exhaustion of IPv4 addresses

IPv4 Address Report

This report generated at 13-Oct-2014 08:09 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)	0.8177
RIPE NCC:	14-Sep-2012 (actual)	0.9742
LACNIC:	10-Jun-2014 (actual)	0.2169
ARIN:	16-Mar-2015	0.6342
AFRINIC:	12-Jun-2019	3.0109

Microsoft: lack of IPv4 addresses (13 Jun 2014)

IPv4 Exhaustion Gets Real – Microsoft Runs Out Of U.S. Addresses For Azure Cloud – Time To Move To IPv6!

BOOM! IPv4 address exhaustion just hit home really hard for a good number of people. They set up virtual machines (VMs) in a *US region* on Microsoft's Azure Cloud and now suddenly find that when they use those VMs to access other websites they are treated as if they are from a country outside the US. Why?

Because Microsoft RAN OUT OF IPv4 ADDRESSES from its "U.S." blocks of IPv4 addresses!



<http://www.internetsociety.org/deploy360/blog/2014/06/ipv4-exhaustion-gets-real-microsoft-runs-out-of-u-s-addresses-for-azure-cloud-time-to-move-to-ipv6/>



The HEPiX IPv6 working group

HEPiX IPv6 Working Group

- Consider how IPv6 should be deployed in HEP
 - Worldwide Large Hadron Collider Grid (WLCG)
- Readiness and Gap analysis
- HEP applications, middleware, security issues, system management and monitoring tools, end to end network monitoring tools, operations
- Run a distributed HEP testbed
 - to help explore all the above issues
- Working closely with the WLCG IPv6 task force

Survey of WLCG sites: IPv6 readiness

Survey of all WLCG Tier 0/1/2

- Announced on 28 May 2014
- We asked all sites to respond and then reminded twice
- Questions
 - Is your site already offering connectivity, routing and naming services for IPv6 end systems?
 - If so, have you already enabled IPv6 on some of the services you manage?
 - If not, are there plans for this? If so, what are the timelines?
 - Does your site currently have problems with allocating a sufficient number of IPv4 addresses? Or foreseen in the near future?
 - Other work, other comments
- Complete wiki table at

[https://www.gridpp.ac.uk/wiki/2014 IPv6 WLCG Site Survey](https://www.gridpp.ac.uk/wiki/2014_IPv6_WLCG_Site_Survey)

Results to date

- Many thanks to those who have replied
 - Those who have not yet, please do so!
- The table is live
 - Please add new lines or modify your answer when the situation changes (just change the “date” field)
- We had responses from CERN (Tier 0) and
 - 12 Tier 1 sites (2 missing)
 - 96 Tier 2 sites (~60 missing)

Tier 0 and 12 Tier 1 sites

- IPv6 connectivity?
 - 2 yes (CERN, SARA), 7 partial, 4 not yet
- Enabled services?
 - Limited: mainly DNS, web, email etc.
- When will you be ready?
 - 1 now (CERN), 6 within 1 year, 6 not defined
- Lack of IPv4 addresses?
 - Just 1 (CERN)

Tier 2 sites (96 of them)

- IPv6 connectivity?
 - 16 yes, 10 partial, 68 not yet
- Enabled services?
 - Some DNS, web, email etc, but ~14 sites have deployed more widely
 - Some test-bed, some production
- When will you be ready? (for those with no current IPv6)
 - 11 within 1 year, 1 within 2 years, 9 are planning (no timetable)
 - ~ 47 have no plans or timetable is unknown
- Lack of IPv4 addresses?
 - 5 sites now, many say OK now but problems in a few years
 - Many sites note current use of private IPv4 for worker nodes

IPv6 at CERN

- <http://cern.ch/ipv6>
- Deployment of IPv6 (almost) complete
 - 2014, February 18: DHCPv6 leases to any device in the IT buildings 28,31,600
 - 2014, April 1st: DHCPv6 leases to any device in the IT datacentre in building 513
 - 2014, May 6th: DHCPv6 leases to any registered device connected to a portable socket or WIFI
 - 2014, May 8th: dual-stack lxplus instance available at lxplus-ipv6.cern.ch
 - 2014, May 12th: imap, pop, smtp, ldap services dual stack
 - 2014, June 3rd: DHCPv6 leases to any static device in GPN. DHCPv6 deployment completed.
 - 2014, June: ldap service dual stack

LHC Experiments and IPv6

News from Experiments

- All 4 experiments active in the working group
- ALICE – Costin Grigoras
 - Many central services already dual stack
 - 5 site VoBox now support IPv6
 - Working on new AliEN and XRootD V4.0
- ATLAS – Alastair Dewhurst
 - Dual stack pilot factory works to dual-stack CE
 - Will test dual-stack Panda submitting to IPv6-only WN
 - Testing of dual-stack Squid3 soon

Experiments (2)

- CMS – Andrea Sciaba and Tony Wildish
 - GlideinWMS (see OSG)
 - HTCondor will soon allow mixed IPv4 and IPv6 pools
 - Wisconsin running out of IPv4 addresses
- LHCb – Raja Nandakumar
 - Testing dual-stack and IPv6-only UI at Oxford
 - CVMFS access works on dual-stack and IPv6-only
 - As uses dual-stack Squid
 - DIRAC binds by default to just IPv4
 - but easy to update to use IPv6 (and works)
 - Working with developer to patch

IPv6 deployment timescale

- Both ATLAS and CMS recognise the need to be able to support IPv6-only WNs and/or IPv6-only opportunistic resources
- CMS
 - Would like a substantial fraction of CMS data accessible via AAA through IPv6 by end 2015
- ATLAS
 - data will need to be stored on dual stack storage
 - Recognise that this will be a challenge so wish to encourage sites to prepare their infrastructure

IPv6 and ATLAS

“ATLAS would therefore make the following requests to sites so that can be in a position to take advantage of IPv6 resources when they come available:

- Request that all Tier 1s engage with the HEPiX/WLCG IPv6 task force.
- Request that all Tier 1s provide a dual stack PerfSONAR machine by April 2015.
- Request that Tier2Ds provide a dual stack PerfSONAR machine by August 2015.”

IPv6 and ATLAS (2)

“The PerfSONAR machine should be located in their network (close to the storage) to give an accurate performance monitoring of their Storage. We believe that the recently released perfSONAR 3.4 is fully compatible with IPV6 but we (IPV6 task force) will be testing this and hope to confirm this by November.

- ATLAS will support sites that wish to be early adopters and deploy dual stack SEs.
- ATLAS advise sites to consult with their SE's developers before upgrading"

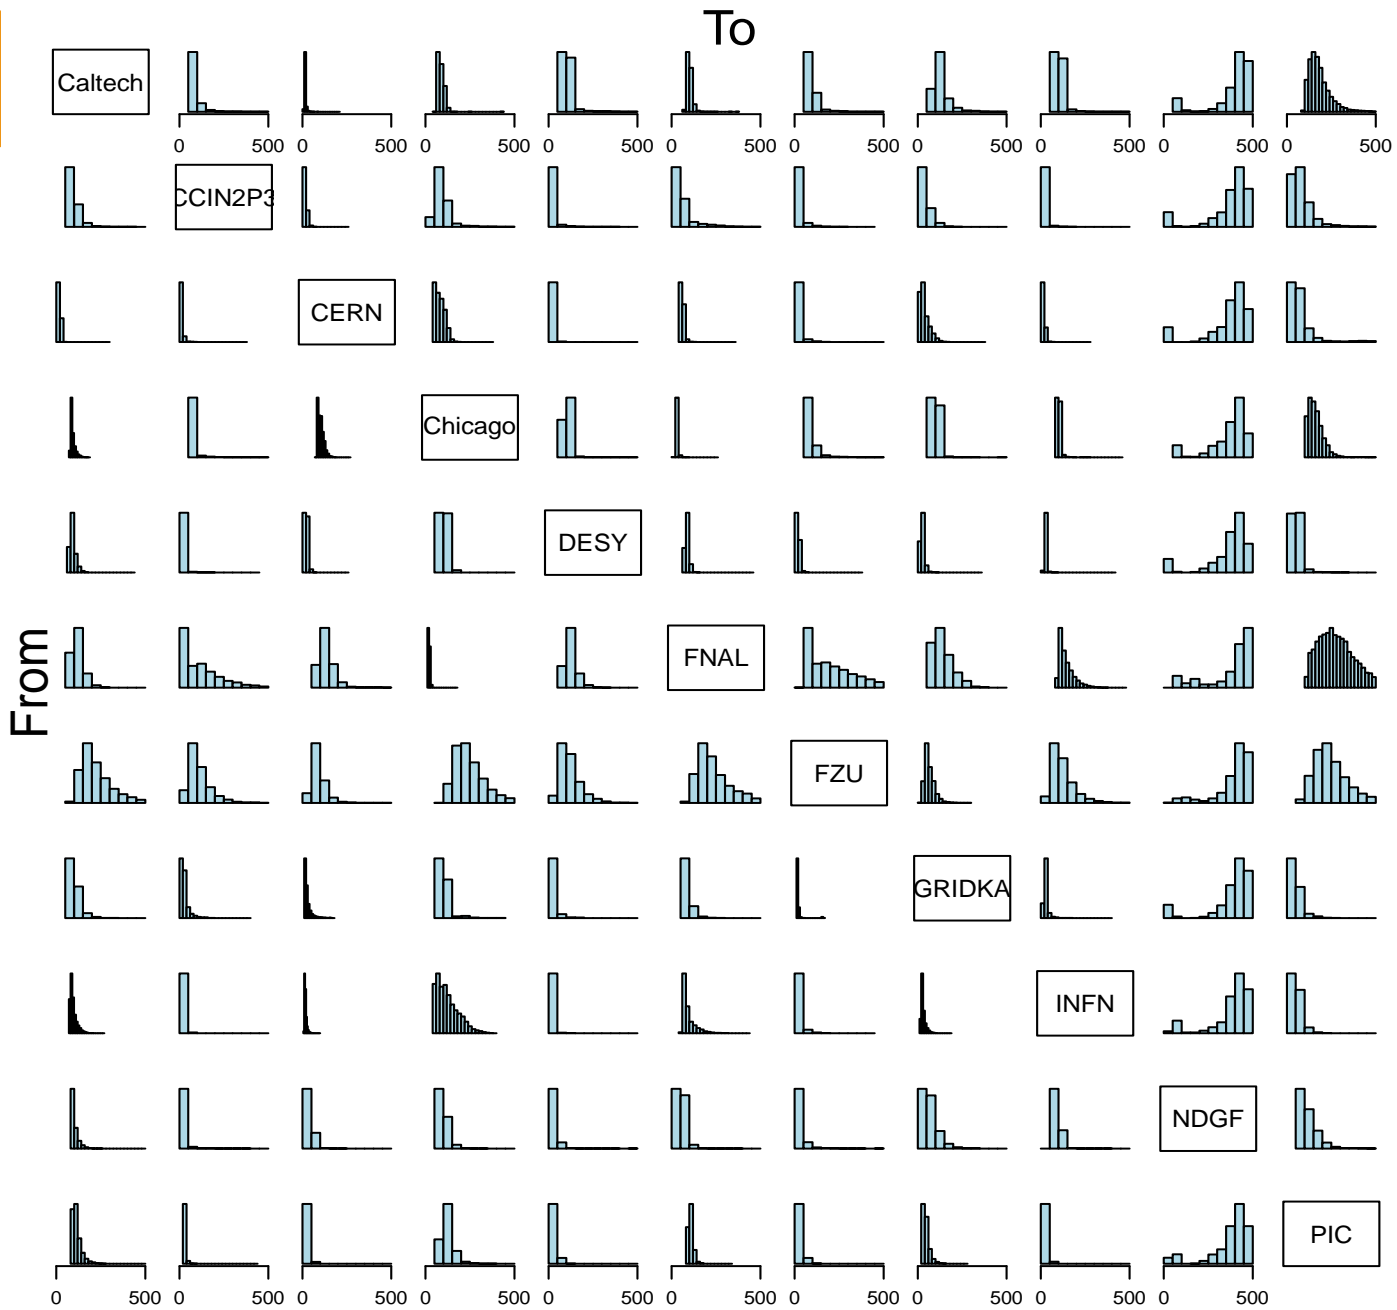
IPv6 Testbed

By the way

<http://test-ipv6.com> is very useful



Testbed – data transfers

- GridFTP tests
 - Running for more than 18 months
 - Each site transfers a 1GB file to every other site using globus-url-copy over IPv6, 3rd-party transfer launched from CERN
 - Filesize is checked and file removed using uberftp/IPv4
- Useful for:
 - Checking health of the testbed
 - Verifying software configurations of servers
 - Verifying network routes etc.



Some IPv6 testing news

- GridFTP mesh will move to PhEDEx and FTS3
- Recent data transfers from dual-stack QMUL (StoRM) to Imperial (dCache 2.10.4)
 - Using IPv6 FTS at Imperial





IPv6 and Operations

Security and IPv6

- A useful discussion with the CERN security team (Oct 2014)
 - Lessons learned during CERN migration
 - Vulnerabilities in new protocols
 - Security tools
- Interesting conference paper
 - <https://www.usenix.org/system/files/conference/woot14/woot14-ullrich.pdf>
- Need to create a top-10 list of security issues for sites deploying IPv6
 - block all IPv6 traffic except ICMP (avoid surprises!)
 - Unblock when requested
 - Firewall rules for IPv4 and IPv6 should be “the same”
- The WG needs to check that logfiles for dual-stack services contain meaningful format IPv6 addresses
 - Challenge sites to confirm ability to deal with malicious IPv6

(some) Software IPv6 readiness

<http://hepixon.ipv6.web.cern.ch/wlcg-applications>

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
CVMFS	Other Application	ALICE, ATLAS, CMS, LHCb	2.1.15	YES
Dashboard Google Earth	Monitoring	ALICE, ATLAS, CMS, LHCb		Claimed
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	2.6.19	NO
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	2.9.4	YES with caveats
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	1.9.12	NO

IPv6 and Storage services

- | | | |
|--------------------------|-----|---|
| • Castor | No | |
| • EOS | No | Needs to be rebuilt on XRootD v4 |
| • DPM (UMD 3.0) | Yes | But needs config tweaks |
| • dCache (v2.10.x) | Yes | |
| • StoRM (v1.11.2) | Yes | But config tweaks |
| • XRootD (v4.0) | Yes | for dual-stack only |
| • XRootD (v4.1 - future) | Yes | to avoid redirection of IPv6-only client to IPv4-only resources |

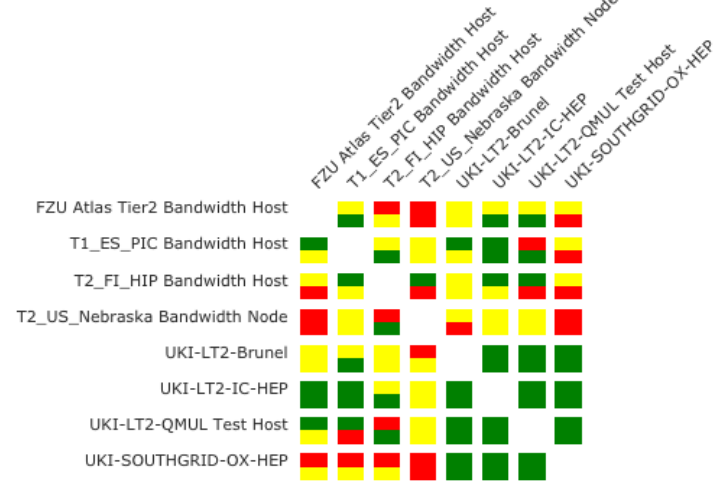

perfSONAR-PS and IPv6



Throughput

IPv4

IPv6



■ Throughput >= 400Mbps
■ Throughput < 400Mbps
■ Throughput <= 100Mbps
■ Unable to retrieve data
■ Check has not yet run

Next steps

Next steps

- Important for **ALL** Tier 1's to now be members of the IPv6 working group
 - On the mailing list, attending meetings, deploying dual-stack PerfSONAR, participating in dual-stack testing of SEs, joining the data transfer testbed
- Move the GridFTP data transfer testbed to PhEDEx and FTS3
- Move SAM tests to dual-stack machines
- During next 6 months (by end of March 2015)
 - We aim to perform enough tests of dual-stack SE services with the sites in the working group to gain confidence

Next steps (2)

- Propose a fuller WLCG SE dual-stack deployment plan (April)
 - For consultation with Ops, Experiments, GDB and MB
 - Together with guidance for sites
- Some central services will also need to be dual-stack
- PerfSONAR remains a great place to start with dual-stack
 - Request Tier 1 to deploy by April 2015
 - Request Tier 2 (if possible) deploy before the long physics run (summer 2015)

Next IPv6 WG meetings

- 13 Nov 2014 – Vidyo – 16:00 – 17:00 CET
- 11 Dec 2014 – Vidyo – 16:00 – 17:00 CET
- Perhaps 21/22 January 2015 - F2F at CERN
– *Still tentative “to be confirmed”*

Further info

- HEPiX IPv6 web

<http://hepixonweb.cern.ch>

- Old HEPiX IPv6 wiki

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

- Working group meetings

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

- WLCG Operations IPv6 Task Force

<http://hepixonweb.cern.ch/content/wlcg-ipv6-task-force-0>

Pre-GDB IPv6 workshop

- 10 June 2014 at CERN
- <https://indico.cern.ch/event/313194/>
- IPv6 technical background
- IPv6 at CERN
- File transfer testing – the IPv6 testbed activities
- Status and configuration of some services
- Experiment testing and plans
- Monitoring
- Site status and experiences
- Next steps – discussion



Questions?