

The HEPiX IPv6 Working Group

David Kelsey (STFC-RAL)

GridPP34, QMUL

30 Apr 2015

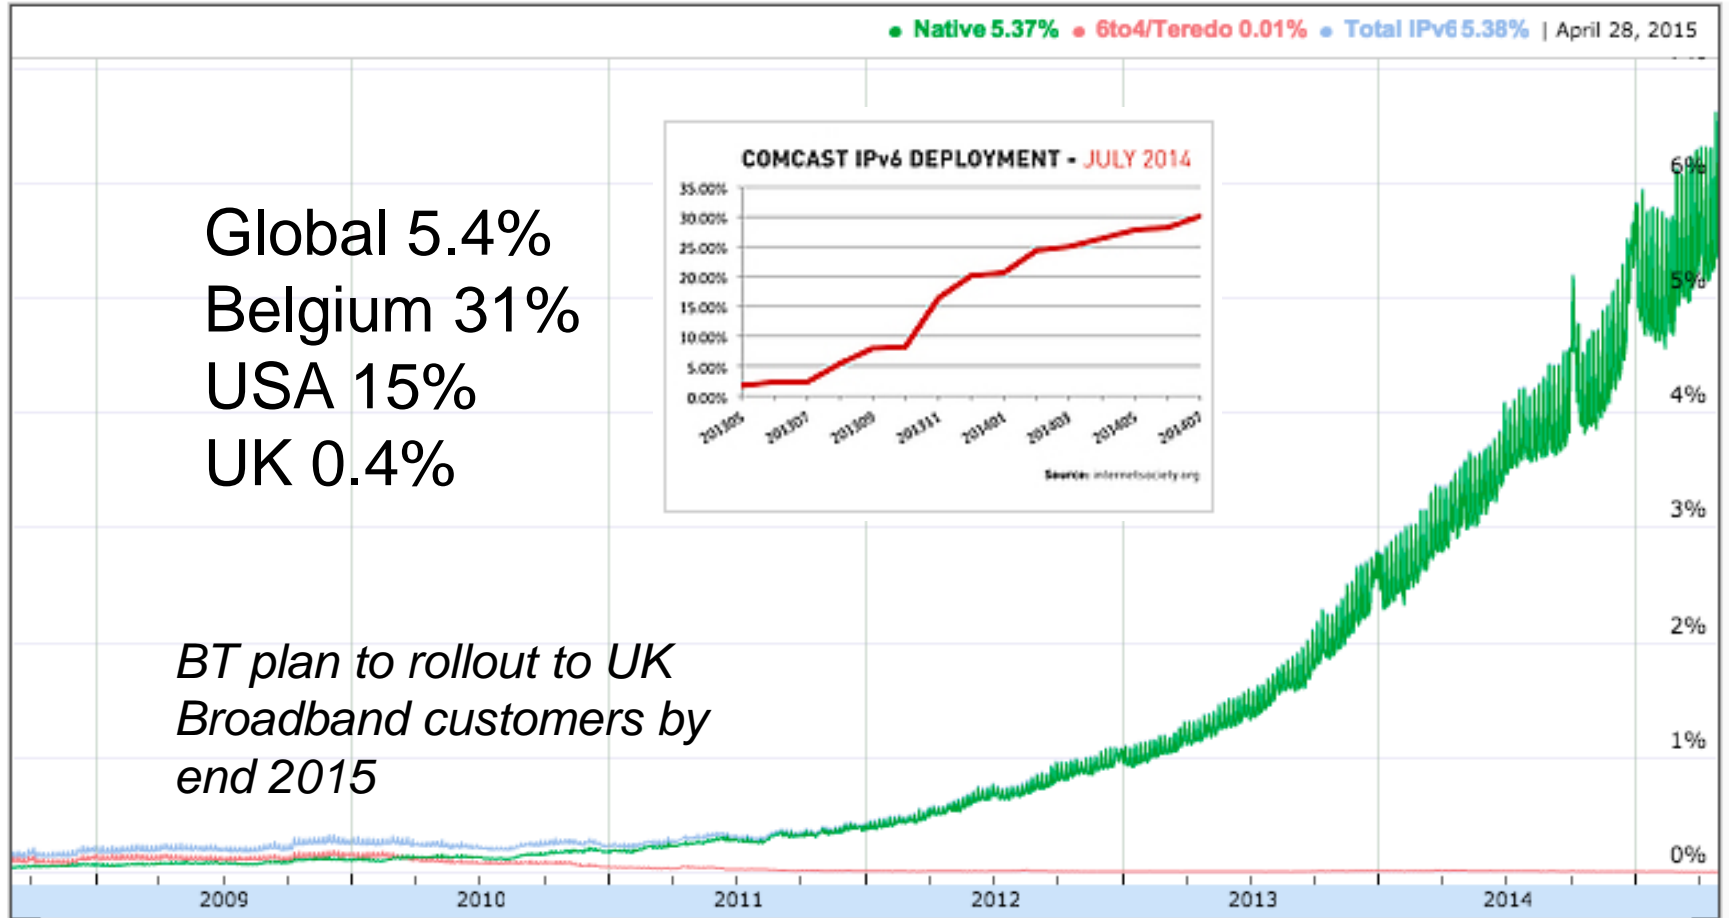
HEPiX IPv6 Working Group

- Members of the group – ***Thanks to them!***

J. Bernier (IN2P3), S. Campana (CERN), K. Chadwick (FNAL), J. Chudoba (FZU), A. Dewhurst (RAL), M. Elias (FZU), S. Fayer (Imperial), T. Finnern (DESY), C. Grigoras (CERN), T. Hartmann (KIT), B. Hoeft (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. Walker (QMUL), T. Wildish (Princeton)

And many others in earlier times

Google IPv6 stats



WLCG Sites - IPv6 readiness

Reminder

- World running out of IPv4 addresses
 - RIRs Asia/Pacific, Europe, Latin America/Caribbean run out of IPv4
- We did a WLCG site survey in summer 2014
- **Tier 1s: ~75% are IPv6 ready. Rest within 1 year**
- **Tier 2s: ~20% are IPv6 ready. ~20% within 1-2 year**
 - ~ 60% have no plans yet
- **~10% of sites report lack of IPv4 addresses** (in next 1-2 years)
 - Including CERN
- Growth of Virtual Machines and/or Containers
- **This is a challenge!**

LHC Experiments & IPv6

Most important IPv6-only use case to address

Site providing IPv6-only CPU - Virtual Machines/Worker Nodes

- Need dual-stack **federated storage** services
- And dual-stack **central services** too
 - Central experiment job submission infrastructure
 - VOMS, catalogues, info etc

LHC Experiments & IPv6 (2)

ATLAS Request (WLCG GDB Oct 2014)

- Tier 1s provide dual stack perfSONAR (April 2015)
- T2Ds provide dual stack perfSONAR (August 2015)

CMS Request (WLCG GDB Sep 2014)

- Substantial fraction CMS data accessible via AAA through IPv6 by end 2015

ALICE runs many dual-stack central services (works)

- testing all their frameworks

LHCb job submission (DIRAC)

- Tests started on lxplus-ipv6 (does not break IPv4)
- About to test on DIRAC v6r13

IPv6 testbed – data transfers

- Caltech, CERN, DESY, FNAL, FZU, Glasgow, IHEP, KIT, INFN, Imperial, Lyon, NDGF, PIC, Wisconsin
- Simple GridFTP mesh of tests over IPv6
- Ran for 706 days (ended Feb 2015)
 - Total data transferred ~ 6.8 PB (10 TB per day)
- Useful for: Verifying network stability, routes etc.

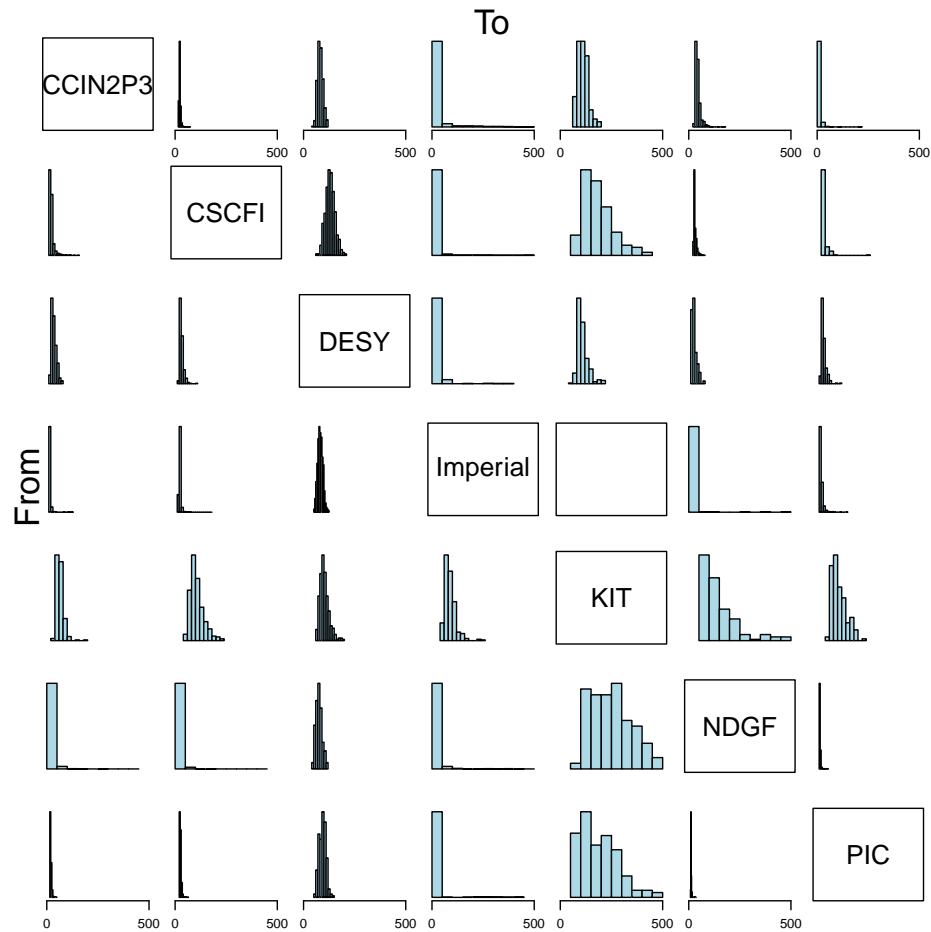
FTS3/dCache data transfers

A new testbed mesh via dual-stack FTS3 (KIT)

- CCIN2P3, CSCFI, DESY, KIT, Imperial, NDGF and PIC
- All running dCache (V2.10.x or later)
 - Some dual-stack, some IPv6-only
- Transfers 1 GB files over SRM/gsiftp
- Now works – but performance differences and failures
- Problems along the way included
 - LHCONE/LHCOPN IPv6 peering
 - DNS problems
 - gfal2 – (V2.8.1 does not accept TURLs with IPv6 number)
 - dCache configuration (for hosts with DNS aliases)

FTS3 data transfer mesh

Time to transfer (s)



IPv6 readiness of WLCG software

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

Software Component	Type	Used by Experiment	Version	IPv6 Compliant
AliEN	LHC Experiment Application	ALICE		
ARC CE	Middleware	ATLAS, CMS		
ARGUS	Middleware	ALICE, ATLAS, CMS, LHCb		
BDII	Middleware	ATLAS, CMS, LHCb	EMI 2	YES
BestMAN	Middleware	ATLAS, CMS		YES
CASTOR	Middleware	ALICE, ATLAS, CMS		NO
cfengine	Monitoring			Unknown
CMS Tag Collector	LHC Experiment Application			Unknown
CMSSW	LHC Experiment Application			Unknown
cmsweb	LHC Experiment Application			YES
CRAB 2	Middleware			Unknown
Cream	Middleware	ALICE, ATLAS, CMS, LHCb	1.16.2	YES
data	Middleware	ALICE, ATLAS, CMS, LHCb	2.1.15	YES
data	Monitoring	ALICE, ATLAS, CMS, LHCb		Claimed
d	Middleware	ALICE, ATLAS, CMS, LHCb	2.6.19	NO
dCa	Middleware	ALICE, ATLAS, CMS, LHCb	2.9.4	YES with caveats
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	1.9.12	NO

Much progress during 2014-15.
Many storage technologies now support IPv6

IPv6 & “Data”

- **DPM** – Yes – need config changes & MySQL 5.5+
- **dCache** – Yes – V2.9.x (2.9.4 several fixes)
- **StoRM** – Yes – need config changes – V1.11.2 tested
- **XRootD** – Yes – V4.0.0 (dual-stack), V4.1+ (IPv6 only)
- **FTS** – Yes – Version 3 and beyond
- **CASTOR** - No
- **EOS** – No – but will when built on xRootD 4

Many IPv6 improvements during last year


IPV6 Network Monitoring

perfSONAR monitoring

- As we turn on production dual-stack
 - We need to collect IPv6 monitoring data
- Since perfSONAR V3.4 Dual-stack is recommended
 - Add AAAA record to DNS
 - If both ends are dual-stack
 - IPv4 & IPv6 tests will both be run

perfSONAR dashboard

Dual-Stack Mesh Config - IPv6 Bandwidth Test



<http://maddash.aglt2.org/>

Dual-Stack Mesh Dashboard
(does not yet show IPv6 correctly!)

LHCOPN/LHCONE & IPv6

- To get same network performance
 - Need IPv6 data to take the same route as IPv4
- **LHCOPN** sites should be peering with IPv6
 - By April 2015
- CERN and 5 Tier 1s completed
 - More to come
- **LHCONE** – IPv6 peering by August 2015

2015 Deployment plan

- LHCOPN/LHCONE IPv6 routing
- perfSONAR dual-stack monitoring
- Gradually deploy production dual-stack data services
 - When LHCOPN/LHCONE peering ready
 - When testing proves functionality and performance
- Important central services to dual-stack
 - Including dual-stack SAM
- Currently ~2% of WLCG services are dual-stack
- Need more work on guidance, e.g. security

Summary

- Deployment of WLCG dual-stack is starting
 - LHCOPN/LHCONE routing
 - perfSONAR monitoring
- Controlled deployment of dual-stack storage
 - And central services
- New members in WG/Testbed always welcome

Links

- HEPiX IPv6 web

<http://hepixonweb.cern.ch>

- 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>

- Paper published in proceedings of CHEP2013
 - For earlier work

Questions?