

WLCG and IPv6

David Kelsey (STFC-RAL)
LHCOPN/LHCONE, Rome
28 Apr 2014

Thanks to ...

Thanks to all colleagues in the HEPiX IPv6 WG

S. Campana (CERN), K. Chadwick (FNAL), G. Chen (IHEP),
J. Chudoba (FZU), A. Dewhurst (STFC), M. Elias (FZU), S. Fayer
(Imperial), T. Finnern (DESY), C. Grigoras (CERN), B. Hoelt (KIT), T.
Kouba (FZU), F. Lopez Munoz (PIC), E. Martelli (CERN),
M. Mitchell (Glasgow), K. Ohrenberg (DESY), A. Pfeiffer (CERN), F.
Prelz (INFN), F. Qi (IHEP), D. Rand (Imperial), S. Rozsa (Caltech), A.
Sciaba (CERN), R. Voicu (Caltech), C. J. Walker (QMUL), T. Wildish
(Princeton)

... and others not named here!


Outline

- The HEPiX IPv6 working group
- IPv4 address exhaustion
- The HEPiX testbed
- Dual stack services
- Software and tools survey
- Network monitoring
- Future plans

HEPiX IPv6 Working Group

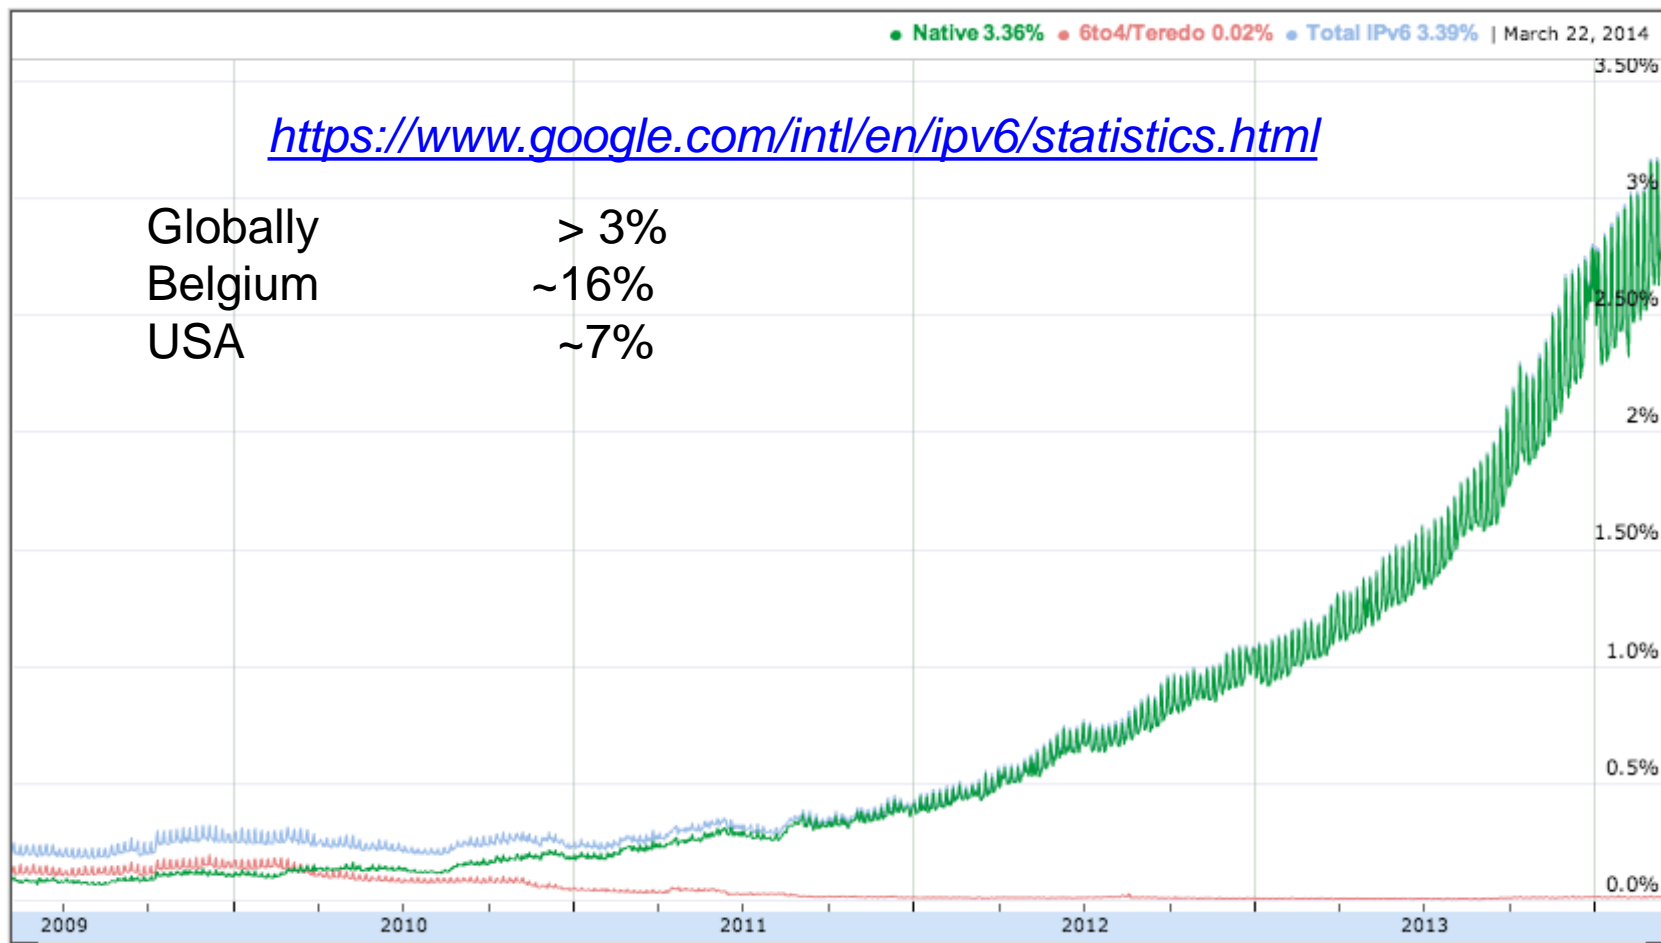
- Started back in 2011 (after regular discussions for years!)
- Consider how IPv6 should be deployed in HEP
 - Worldwide Large Hadron Collider Grid (WLCG)
- Readiness and Gap analysis
- HEP applications, middleware, security issues, system management and monitoring tools, end to end network monitoring tools, operations
 - One reason why communities have not adopted IPv6
- Run a distributed HEP testbed
 - to help explore all the above issues
- Not driving every site's implementation of IPv6!
 - But can share experiences and information

IPv4 free addresses - /8



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

Google accessed by IPv6 clients



IPv4/IPv6 at CERN

- IPv4 addresses predicted to run out at CERN during 2015
 - New Wigner computer centre
 - Many virtual machines
- LHC Run 2 starts soon
- Pressure for WLCG to deploy dual-stack services
 - At least those needed by IPv6-only worker nodes at CERN
- CERN is ready to support full use of IPv6 in weeks
- Note: many other sites have not started yet :=(

File transfer tests over IPv6 ...

The HEPiX IPv6 Testbed

- A dedicated distributed testbed
- 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
- Caltech, CERN, Chicago, DESY, FNAL, FZU, Glasgow, Imperial, IHEP, INFN, KIT, PIC, QMUL
- <http://hepix-ipv6.web.cern.ch/content/hepix-ipv6-distributed-testbed>

IPv6 file transfer tests

- Tony Wildish (CMS)
- Simple mesh of simultaneous data transfers (IPv6)
- Transfers a 1 GB file using GridFTP (globus-url-copy)
 - Measures time to transfer
 - Records any errors
- Uses UberFTP to confirm arrival and then delete
- Have been running for last year

GridFTP IPv6 data transfer mesh

- > 2 PB data transferred in 6 months
- Success rate > 87%
- Very High!
 - Operating at risk



Time to transfer max
500 s

Data transfer lessons learned

- CMS PhEDEx is production ready on IPv6
- The testing is very useful for finding site IPv6 problems
 - E.g. allows some level of stress testing
- Over the months we have experienced some breaks, performance problems, asymmetric routing, ...
 - E.g. cpu utilisation in FNAL border router reaches 60% with IPv6 traffic flows of 200-300 Mbps
- The reason for ongoing data transfer tests
- It is useful for a site to be in the IPv6 testbed
- Now testing other storage systems and protocols

Dual-Stack services at Imperial

- Tier 2 Imperial College, London
- Configured a subset of their service to be IPv4/IPv6 dual stack
 - DNS, SSH, NFS, EMI 2 and EMI 3 CREAM CEs, EMI 2 Worker Nodes, ARC CE and dCache (headnode, SRM component only), all BDII services
 - Stateless address autoconfiguration (SLAAC)
- No problems observed
 - Did not need to turn IPv6 off

Software and Tools survey ...

IPv6 compliance of WLCG applications

Software Component	Type	Used by Experiment	Version	IPv6 Compliance
AiiEN	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		YES
Dashboard Google Earth	Monitoring	ALICE, ATLAS, CMS, LHCb		Claimed

Software survey (2)

dCache	Middleware	ALICE, ATLAS, CMS, LHCb	2.6.19	Claimed
dCache	Middleware	ALICE, ATLAS, CMS, LHCb	1.9.12	NO
DIRAC	LHC Experiment Application	LHCb		Unknown
DPM	Middleware	ALICE, ATLAS, CMS, LHCb		YES with caveats
EGI Accounting Portal	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown
EOS	Middleware	ALICE, ATLAS, CMS, LHCb		NO
Experiment Dashboards	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown
Frontier	LHC Experiment Application	ATLAS, CMS		NO
FTS	Middleware	ATLAS, CMS, LHCb	FTS 3	YES
Ganglia	Monitoring			YES
GFAL/lcg_util	Middleware	ALICE, ATLAS, CMS, LHCb		Claimed
glideinWMS	Middleware	CMS		Claimed
globus toolkit	Middleware	ALICE, ATLAS, CMS, LHCb	5.2.5	YES
GOCdb	System Management Tool	ALICE, ATLAS, CMS, LHCb		Unknown
Gratia Accounting	Monitoring			Unknown
Gridsite	Middleware			Claimed
Gstat	Monitoring			Unknown

Software survey (3)

GUMS	Middleware	ATLAS, CMS		Unknown
Hadoop	Middleware	ATLAS, CMS		NO
HammerCloud	LHC Experiment Application	ATLAS, CMS, LHCb		Unknown
HTCondor	Middleware	ATLAS, CMS		YES with caveats
iCMS	LHC Experiment Application	CMS		Unknown
LFC	Middleware	ATLAS, LHCb	1.8.7	YES
MonALISA	Monitoring			Claimed
MyOSG	Monitoring			Unknown
MyProxy	Middleware		EMI 1	Unknown
MyWLCG	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown
Nagios	Monitoring	ALICE, ATLAS, CMS, LHCb		Claimed
OpenAFS	Other Application	ALICE, ATLAS, CMS, LHCb	Last	NO
PanDA	LHC Experiment Application	ATLAS, CMS		Unknown
perfSONAR	Monitoring	ATLAS, CMS		YES
PhEDEx agents	LHC Experiment Application	CMS		YES
Puppet	System Management Tool	ALICE, ATLAS, CMS, LHCb		YES
REBUS	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown

Software survey (4)

REBUS	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown
SAM	Monitoring	ALICE, ATLAS, CMS, LHCb		Unknown
Scientific Linux	Operating System	ALICE, ATLAS, CMS, LHCb	5.7	YES
SDT IB and QA pages	LHC Experiment Application	CMS		Unknown
Squid	Other Application	ALICE, ATLAS, CMS, LHCb	2.8	NO
Squid	Other Application	ALICE, ATLAS, CMS, LHCb	3.1	YES
StoRM	Middleware	ALICE, ATLAS, CMS, LHCb	1.11.2	YES with caveats
Ticket system -- GGUS	System Management Tool	ALICE, ATLAS, CMS, LHCb		Unknown
various D web tools	LHC Experiment Application	CMS		Unknown
VOMS	Middleware	ALICE, ATLAS, CMS, LHCb	2.0.11	YES
VOMS-Admin	Middleware	ALICE, ATLAS, CMS, LHCb	2.0.11	YES with caveats
WLCG squid monitoring	Monitoring	ATLAS, CMS		Claimed
WMAgent	LHC Experiment Application	CMS		Unknown
WMS	Middleware	ALICE, ATLAS, CMS, LHCb	EMI 2	Unknown
xrootd	Middleware	ALICE, ATLAS, CMS	3.1.0	NO
xrootd	Middleware		4.0.0	Claimed

Future plans and next steps ...

Next testing

- KIT, PIC, FZU, NDGF
 - PhEDEx/SRM/dCache tests (on testbed)
- Then move to use of SRM/FTS/PhEDEx
 - With various **production** endpoints
 - Volunteer sites
- Working on clearly defined use case (from expts)
 - CERN Worker Node/Virtual Machine IPv6-only
 - What WLCG services outside CERN need to be dual-stack?

IPV6 Network Monitoring

- PerfSONAR tests over IPv6 have been successful
 - Imperial, Oxford, FZU, ...
- BUT – not possible to monitor both IPv4 and IPv6 on the same endpoints (control)
 - Would like to maintain full dual histories
- Discussions ongoing with Shawn McKee
- Do we need separate monitoring?
 - We should decide the best approach soon

Next steps

- Perform survey of all WLCG sites
 - When will they run out of IPv4 addresses
 - When will they be capable of IPv6 support
- One-day pre-GDB meeting on IPv6
 - 10 June 2014 at CERN
 - Aimed primarily at Tier 1 sites (others welcome)
 - Encourage IPv6 support
 - Move to dual-stack WLCG services
 - Testing needs to move to production services

Further info

- 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 accepted for publication in proceedings of CHEP2013

Message to networking

- We need good network monitoring in dual-stack environment
 - Should be deployed in advance of much production use
- We do need to encourage sites to deploy IPv6
 - The 10 June 2014 workshop should help
- Any support you (LHC networking) can give would be good
- VOLUNTEERS to join the group are always welcome
 - please contact me!

Questions ?