

# HEPiX IPv6 Working Group

David Kelsey (STFC-RAL)

GridPP33 Ambleside

22 Aug 2014

# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- Software readiness survey
- Experiment plans
- Monitoring – PerfSONAR
- Next steps

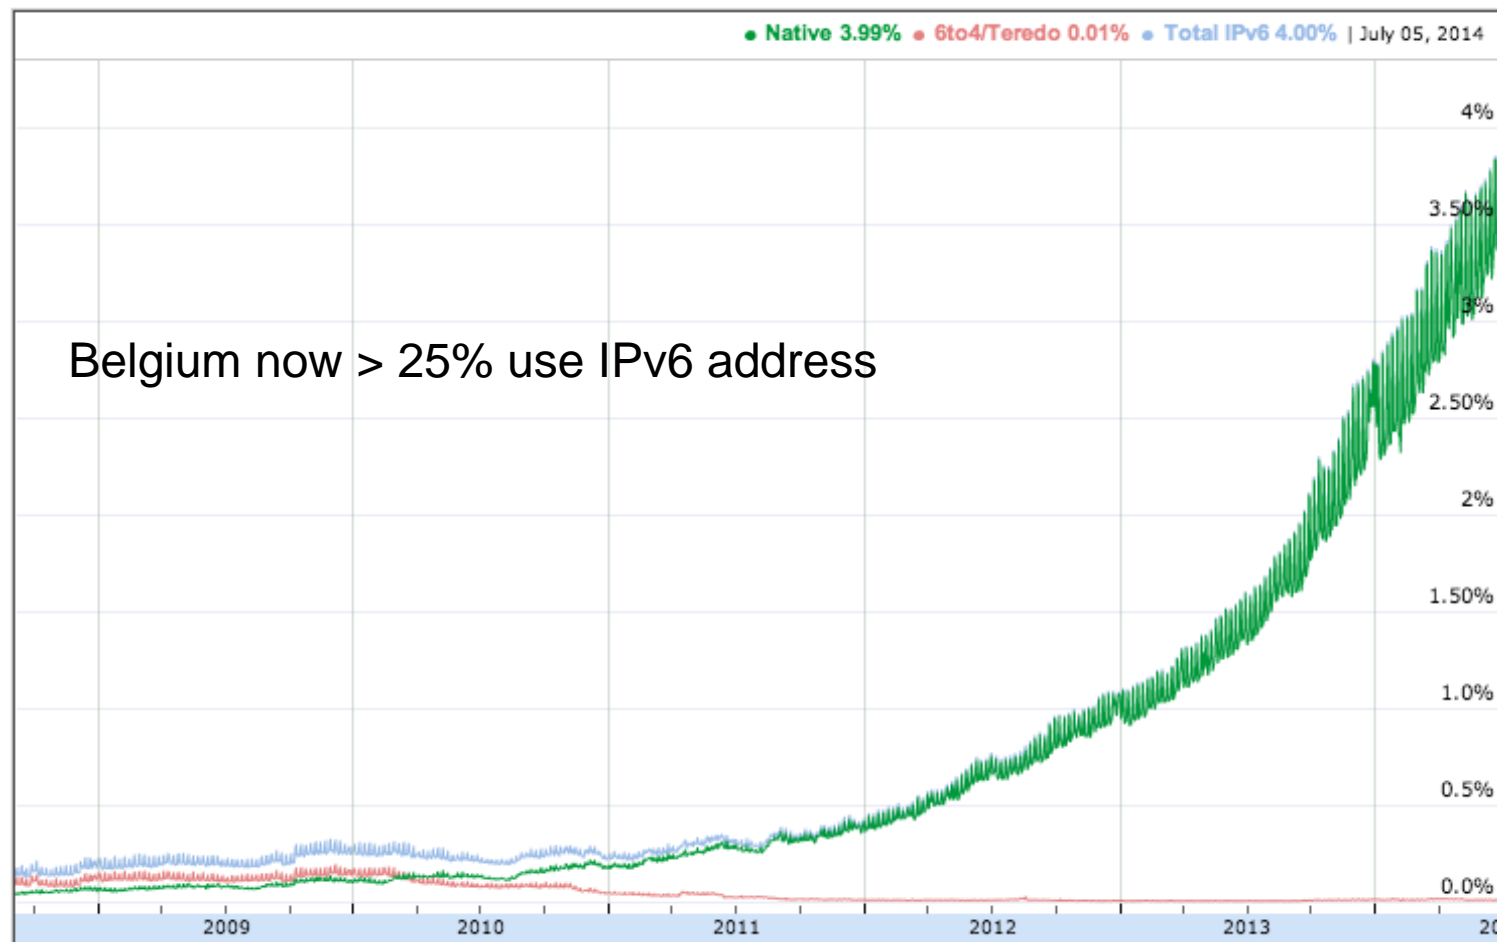
# HEPiX IPv6 Working Group

- 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

# Outline

- The IPv6 working group
- **Use of IPv6 & status of IPv4 address space**
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- Software readiness survey
- Experiment plans
- Monitoring – PerfSONAR
- Next steps

# IPv6 growth (global Google clients)



# 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!

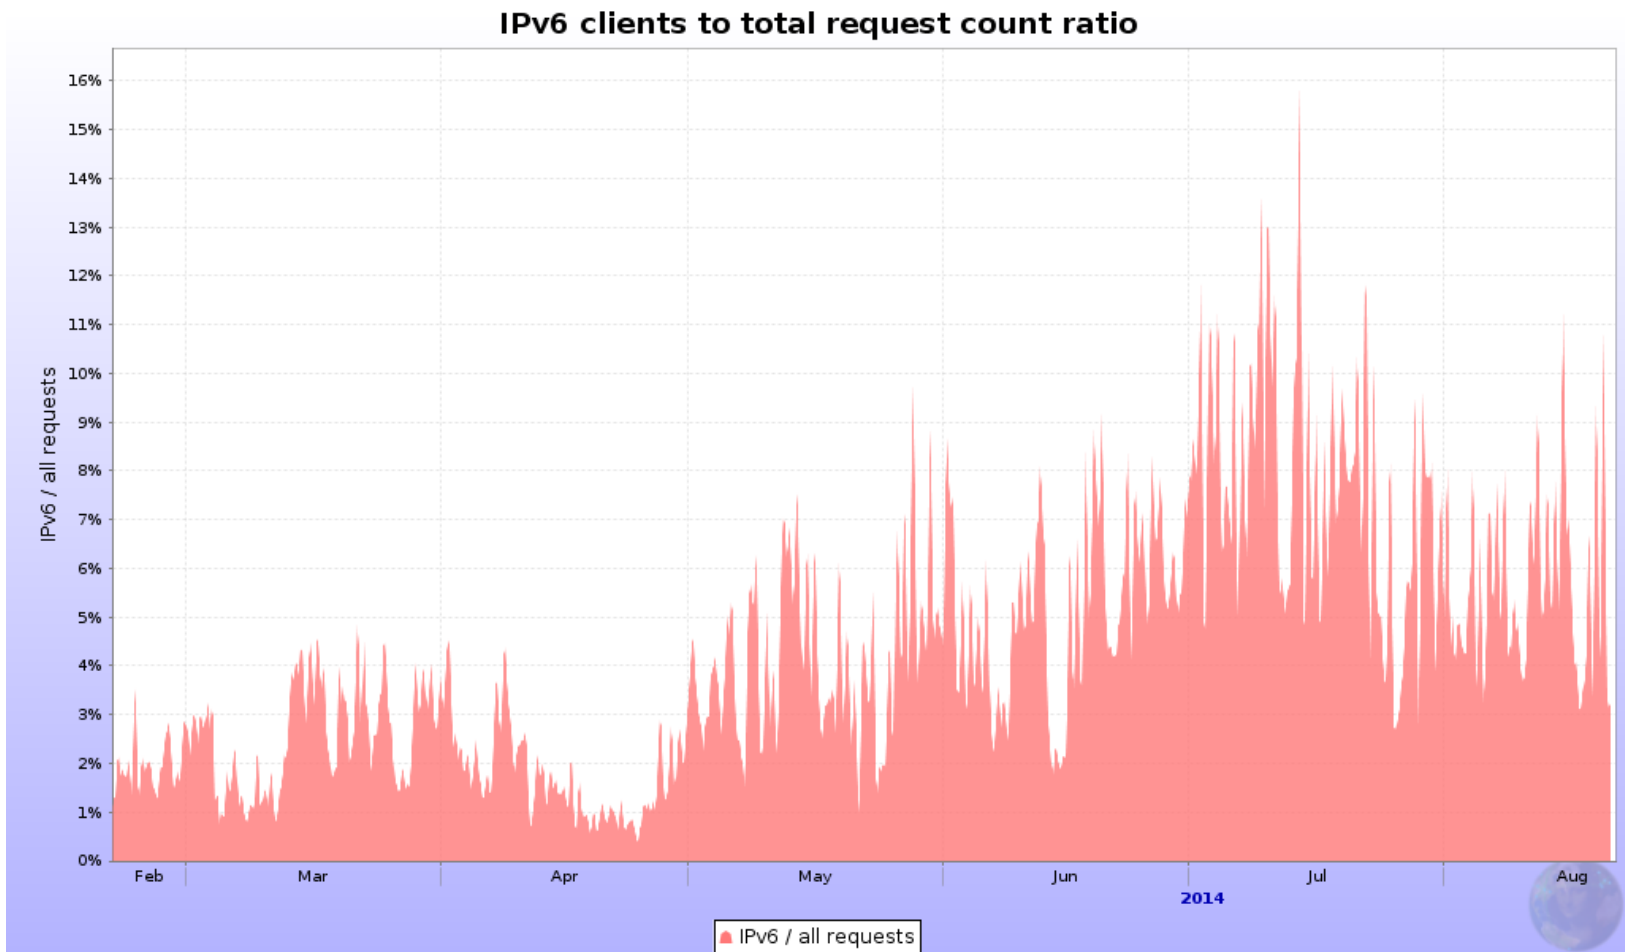
**BOOM!** IPv4 address exhaustion just hit home really hard for a good number of people. They set up virtual machines (VMs) in a *US region* on Microsoft's Azure Cloud and now suddenly find that when they use those VMs to access other websites they are treated as if they are from a country outside the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

# ALICE Central Servers



# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- **WLCG IPv6 site readiness survey**
- IPv6 testbed and data transfers
- Software readiness survey
- Experiment plans
- Monitoring – PerfSONAR
- Next steps



# 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)

# 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)

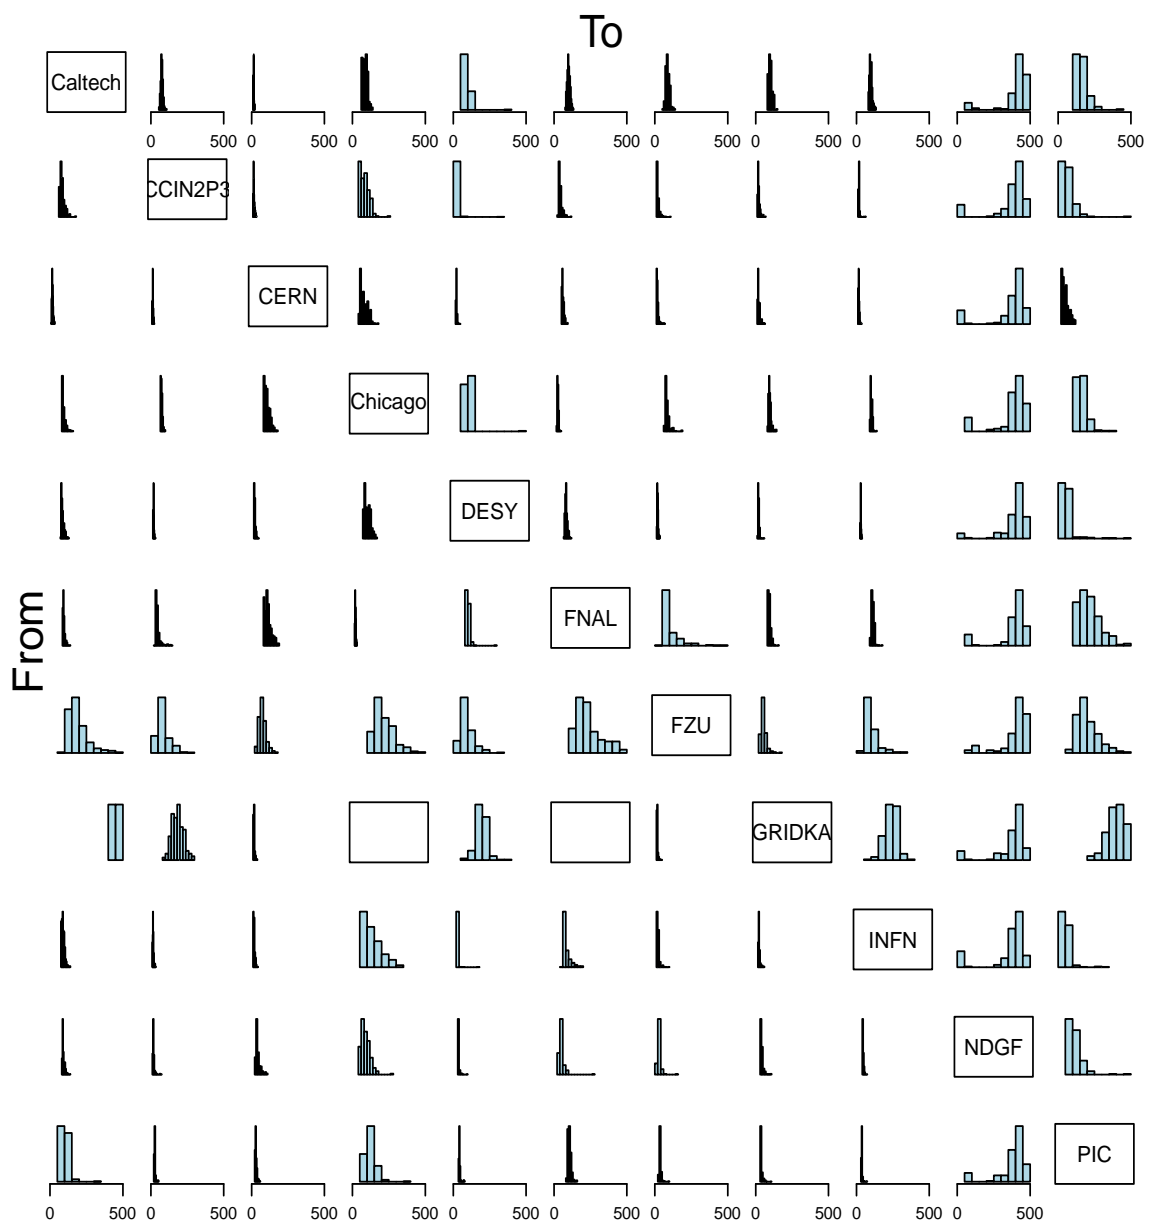
- IPv6 connectivity?
  - 16 yes, 10 partial, 68 not yet
- Enabled services?
  - Some DNS, web, email etc, but ~14 sites have deployed more widely
    - Some test-bed, some production
- When will you be ready? (for those with no current IPv6)
  - 11 within 1 year, 1 within 2 years, 9 are planning (no 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

# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- **IPv6 testbed and data transfers**
- Software readiness survey
- Experiment plans
- Monitoring – PerfSONAR
- Next steps

# Testbed – data transfers

- GridFTP tests
  - Running for more than 1 yr
  - Each site transfers a 1GB file to every other site using globus-url-copy over IPv6, 3<sup>rd</sup>-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.
- Moving soon to use FTS3

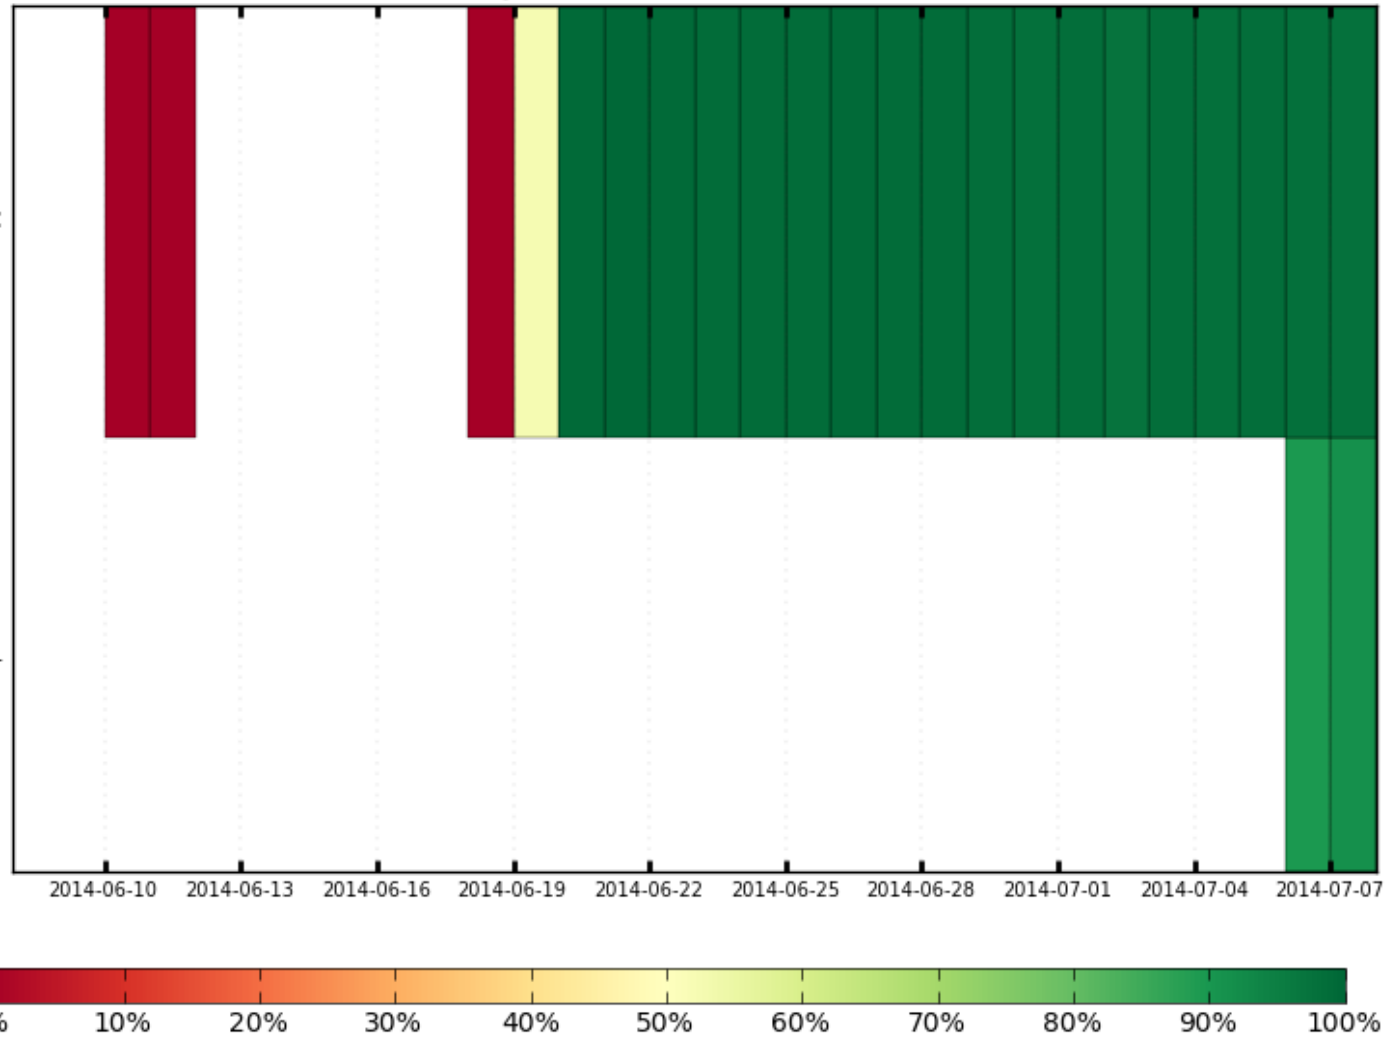


# PhEDEx transfers

- Dual-stack FTS3 server at Imperial
- DPM storage elements at Imperial and Glasgow
- Transfers just worked at production quality!
  - Files deleted automatically on arrival
- Ran for over 2 months, ~120 TB data
- Now tests have restarted
- Recent additions:
  - PIC (dCache) & Nebraska (bestman+hadoop)
    - QMUL (StoRM)

### CMS PhEDEx - Transfer Quality

30 Days from 2014-06-08 to 2014-07-08



From Imperial



# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- **Software readiness survey**
- Experiment plans
- Monitoring – PerfSONAR
- Next steps

<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

# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- Software readiness survey
- **Experiment plans**
- Monitoring – PerfSONAR
- Next steps

# Experiments

- Good engagement with the Experiments in 2014
  - Very important
  - Testing should be driven by their use cases
- ALICE moving to dual-stack and testing all
  - XrootD V4 required
- ATLAS
  - Imperial to test BigPanda instance
  - Panda Dev instance at CERN to be dual-stack

# Experiments (2)

- CMS
  - Test dual-stack glideinWMS
  - Dashboard monitoring
  - Test CRAB3
- LHCb job submission
  - All DIRAC communication needs to be tested
    - on separate VMs
  - Tests started on lxplus-ipv6 (dual-stack)
    - Works for IPv4, & testing on IPv6-only (Oxford)

# Data Management tests

- ATLAS, CMS and LHCb will all use FTS3
  - GridFTP, XrootD and http
- ATLAS already doing IPv4 stress tests of Rucio
  - Could then do IPv6 stress testing
- CMS
  - Test a few endpoints of each storage technology
  - Next – test AAA with IPv6
    - Nebraska enabling IPv6 access to their XrootD
    - Requires update to new XrootD client

# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- Software readiness survey
- Experiment plans
- **Monitoring – PerfSONAR**
- Next steps

- Currently eight sites in the mesh
  - FZU
  - T2\_FI\_HIP
  - PIC
  - Brunel, Imperial, QMUL, Oxford
  - UNL
- Measuring:
  - throughput
  - latency and loss
  - traceroute
  - ping


## perfSONAR Dashboard

Dashboards

### HEPiX IPv6 sites Dashboard

HEPiX IPv6 sites - IPv4 throughput test

■ Throughput  $\geq$  400Mbps
 ■ Throughput  $<$  400Mbps
 ■ Throughput  $\leq$  100Mbps
 ■ Unable to retrieve data



<http://netmon02.grid.hep.ph.ic.ac.uk:8080/maddash-webui>



# IPV6 Network Monitoring

- PerfSONAR monitoring
- As we turn on production dual-stack
  - We need to collect monitoring data
- A good way for a new site to get started
  - Ask Imperial to add to HEPiX IPv6 dashboard
- Encourage this for sites not in the IPv6 group
- NAGIOS and other monitoring
  - Deploy new plugins to test IPv6 services

# Outline

- The IPv6 working group
- Use of IPv6 & status of IPv4 address space
- WLCG IPv6 site readiness survey
- IPv6 testbed and data transfers
- Software readiness survey
- Experiment plans
- Monitoring – PerfSONAR
- **Next steps**

# Next steps

- OSG is now testing – hear more at Nebraska HEPiX mtg
- Continue the file transfer mesh testbed:
  - get more sites, in particular all T1s, involved
  - FTS IPv6 pilot
- More storage options
- ATLAS: BigPanDA dual-stack, Rucio with IPv6
- CMS: dual-stack glideinWMS, IPv6 in AAA
- LHCb: dual-stack DIRAC services
- ALICE: xrootd

# Next steps (2)

- Next milestone
  - dual-stack production services
  - Several sites already did it successfully
  - Need to find more volunteering sites
- Joining perfSonar IPv6 testbed – good first step
- Dual-stack coordinated with WG & experiments
- Test IPv6-only WNs
  - when have enough sites IPv6 ready
  - must not wait for all, prepare to run a in mixed environment

# Timetable for Dual-Stack?

- Impossible to move 100% of services quickly
  - Sites just not ready
- Also have to solve/fix/avoid IPv6 problems
  - Software and tools
- In parallel we need to work on
  - Training for sys admins
  - System management tools and Security logs
- Gradual move during rest of 2014 and Run 2
- Aim for full support of IPv6 WN/VM in Run 3?
- BUT if GridPP can do it sooner – all the better
  - A few lead countries can help others

# 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

# 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?