


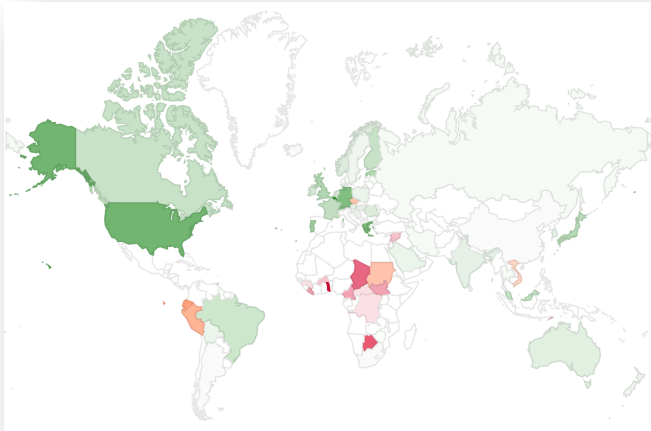


Deployment of IPv6 only CPU resources at WLCG sites

Alastair Dewhurst, on behalf of the HEPiX
IPv6 Working group

Global IPv6 status

- Over 12.5% of Google traffic now goes via IPv6.
- Some commercial hosting companies offer cheaper IPv6 only services.
- June 2016, Apple now mandates all Apps submitted to the App Store must support IPv6-only networking.
- September 2016, Microsoft Azure VMs now available as IPv6.
- For HEP:
 - 11 sites currently provide dual stack storage.
 - 2 sites running IPv6 WN (With NAT64).



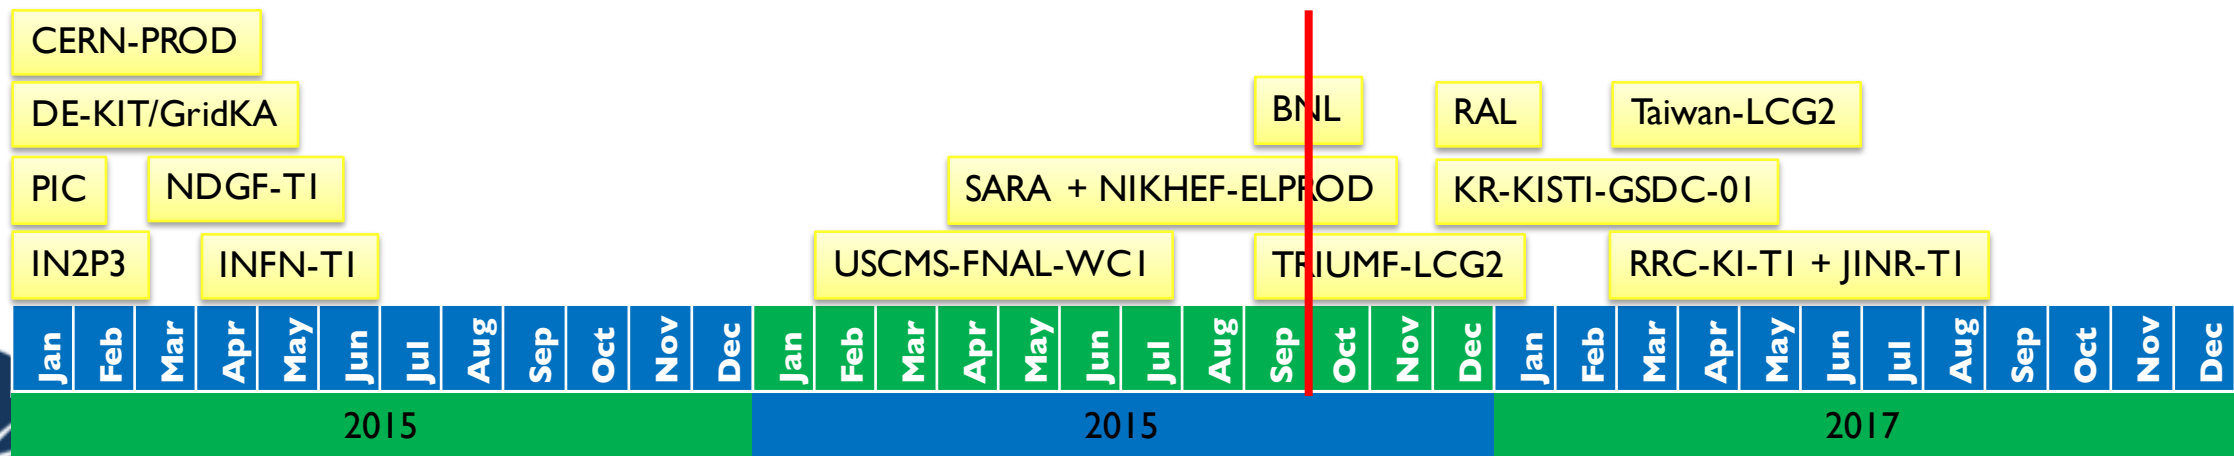
<https://www.google.com/intl/en/ipv6/statistics.html>

Alastair Dewhurst, 11th October 2016



Networking over IPv6

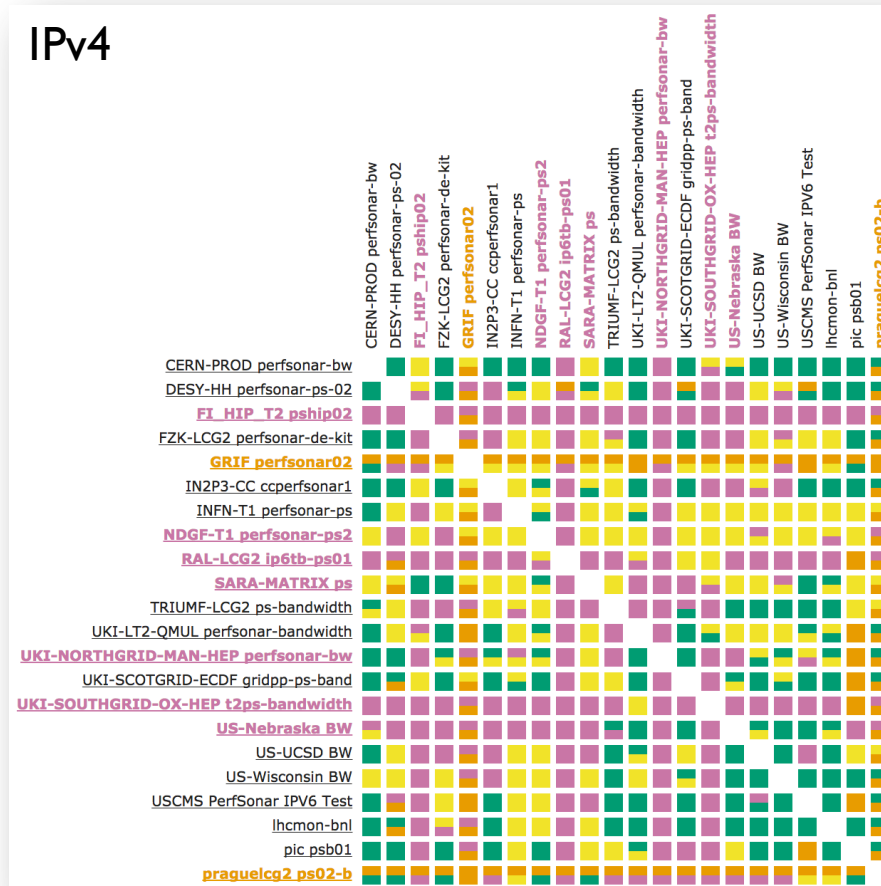
- HEPiX Group working requested:
 - All Tier 1s to offer IPv6 peering to LHCOPN and provide dual stack PerfSONAR machine by April 2015.
 - All Tier 2s to offer IPv6 peering to LHCONe and provide dual stack PerfSONAR machine by August 2015.
- Timeline showing Tier 1 IPv6 peering.




PerfSonar

- Meshes show dual stack PerfSonar endpoints
- <http://maddash.aglt2.org/maddash-webui/index.cgi?dashboard=Dual-Stack%20Mesh%20Config>

IPv4



IPv6



IPv6 only CPU

The HEPiX IPv6 group made a proposal[1] to the WLCG Management Board which has now been accepted to allow sites to migrate their CPU resources to IPv6 only by April 2017.

- All VO's encourage sites to upgrade their storage to dual stack.
- All VO's are working towards making their central services required by Grid jobs dual stack by April 2017.
- Shared central services (e.g. CVMFS) should be accessible via IPv6 by April 2017.
- Tier 1s should provide functional dual stack access to storage (and other services they may run):
 - 1GB/s and 90% availability by April 2017.
 - 10Gb/s and 95% availability by April 2018.

[1] <https://indico.cern.ch/event/467575/contributions/1145552/attachments/1311592/1962831/WLCGIPv6Deployment.pdf>



Software status


- The HEPiX IPv6 group has tested a large amount of software for IPv6 compliance.
- Key storage software and protocols work:
 - dCache, DPM, StoRM
 - XrootD 4, GridFTP, http
- Sites will be expected to provide dual stack storage if they wish to run IPv6 only CPU.

	No Storage	IPv4 only Storage	Dual Stack Storage
IPv4 CPU	✓	✓	✓ (11+)
IPv6 CPU	✓ (0 currently)	✗	✓ (2+ want to)



Central Services

- The FTS service is used to transfer large volumes of data between sites.
 - Even though third-party copy is used, the FTS service needs to be dual stack to allow transfers to go via IPv6.
 - Both sites need to have dual stack storage
 - FTS3 pilot service at CERN supports IPv6, ongoing discussion on how best to upgrade the rest.
- CVMFS is used by all the LHC VOs to distribute software to WN.
 - Assuming site squid is dual stack then stratum 1 does not need to be.



ALICE

- ALICE uses a fully federated storage model based on XrootD.
 - XrootD 4 will need to be provided by the sites
- 100% of data must be available via IPV6 before any WFN can be IPv6 only.
 - As data is duplicated not all sites need to upgrade (but the vast majority will).
- ALICE central services have been dual stack for over a year.
- Around 5% of ALICE sites have upgraded their storage to dual stack.
 - Aim for all sites to upgrade by end of Run 2.



ATLAS

- ATLAS are focusing on making services required by IPv6 only WN dual stack first. These all use https:
- The PanDA servers.
 - Pilots contact the Panda servers to pull in jobs as well as send regular updates.
- The DDM headnodes:
 - These are contacted by the jobs to register files that have been uploaded.
- Some services such as pilot factories have been already made dual stack.



- The job submission middleware, GlideinWMS has been validated as IPv6 compliant.
- The data management system, PhEDEx, uses the Oracle client for communication between local site agents and the central service.
- CMS is upgrading all its central services to dual stack.
 - Focus on the services needed directly by WN - aim to complete by April 2017.
 - All services by end of run 2.



LHCb

11

- LHCb uses the DIRAC framework to submit jobs to the Grid.
- DIRAC officially supports IPv6 and other users are already running with dual stack instances.
- LHCb has 40 VO boxes at CERN.
 - Rationalising these machines down to ~10.
 - All new machines will be dual stack, currently 4 deployed and tested.



Conclusion

12

- There is now a WLCG Management Board endorsed plan to allow sites to run IPV6 only CPU.
- The LHC VOs all have plans to upgrade their infrastructure to meet this goal.
- We expect to review the progress towards IPV6 adoption in 2018 and expand the plan to allow sites to not need IPV4 at all.



Acknowledgements

13

- With thanks to all the people on the HEPiX working group: Andrea Sciaba, Dave Kelsey, Raja Nandakumar, Costin Grigoras, Duncan Rand, Simon Fayer, Dr. Kars Ohrenberg, Daniel Traynor, Bruno Hoeft, Terry Froy, Edoardo Martelli, Kashif Hafeez, Ulf Bobson Severin Tigerstedt, Thomas Finnern, Tiju Idiculla, Marian Babik, Jiri Chudoba, Mr. Fernando Lopez Munoz, Francesco Prelz, Ramiro Voicu.
- IPv6 only CPU proposal that was accepted by the MB:
<https://indico.cern.ch/event/467575/contributions/1145552/attachments/1311592/1962831/WLCGIPv6Deployment.pdf>
- IPv6 Security Poster: <https://indico.cern.ch/event/505613/contributions/2227722/>
- IPv6 Deployment: <https://indico.cern.ch/event/505613/contributions/2227381/>

