



The HEPiX IPv6 working group

David Kelsey (STFC UKRI)
HEPiX IPv6 WG meeting, CERN, 26 Oct 2022

On behalf of all in the HEPiX IPv6 working group



Active in HEPiX IPv6 Working Group – last 12 months

- M Babik (CERN), M Bly (RAL), N Buraglio (ESnet), T Chown (Jisc), D Christidis (U Texas/ATLAS), J Chudoba (Prague), C Condurache (EGI.eu), P Demar (FNAL), J Flix (PIC), C Grigoras (CERN/ALICE), B Hoeft (KIT), H Ito (BNL), D P Kelsey (RAL), E Martelli (CERN), S McKee (U Michigan), R Nandakumar (RAL/LHCb), K Ohrenberg (DESY), F Prelz (INFN), D Rand (Imperial), A Sciabà (CERN/CMS), E Simmonds (FNAL), T Skirvin (FNAL)
- Many more in the past, and others join from time to time
- *and thanks also to WLCG operations, WLCG sites, LHC experiments, networking teams, monitoring groups, storage developers...*



Overview of “Introduction”

- HEPiX community – all of HEP & similar research communities
- A short history of the IPv6 working group
 - Phase 1 2011-2016 Investigation, testing and fixing software
 - Phase 2 2016-2020 Dual-stack IPv6/IPv4 deployment on WLCG
 - Phase 3 2020-? Planning for IPv6-only WLCG
- Today’s agenda

Why IPv6?

- Survey of 18 major HEP sites (Sep 2010) – IPv6 readiness
 - NRENs ready, **Universities and Labs not ready**
 - Some reported lack of IPv4 address space, including CERN
- HEPiX meeting – Nov 2010
 - IANA projecting IPv4 address exhaustion
 - Sep 2010 – memo from US Federal CIO to all Exec depts (incl DOE) – dual-stack
- Opportunistic CPU resources could be IPv6-only
- Recognition: much middleware, software and technology **not yet IPv6 capable**
- HEPiX decided to create a working group (started April 2011)

HEPiX IPv6 Working Group (2011-16)

- **Phase 1 – full analysis of work to be done**
 - Applications, system and network tools, operational security
 - Create and operate a distributed test-bed
 - No interference with WLCG production data analysis!

In 2012

- Active HEPiX IPv6 test-bed with ~ 12 sites
 - engagement of all 4 LHC experiments
- Testing regular data transfers across the testbed
- Testing dual-stack services (production) at Imperial College London
- Concluded not able to support **IPv6-only clients** until **at least 2014**

2015

- At CHEP conference in April 2015
 - 75% of WLCG Tier-1 sites are IPv6-ready (but only 20% of Tier2)
 - 10% of sites now reporting lack of IPv4 addresses
- Most important IPv6-only use case
 - Sites, Clouds providing CPU (virtual machines)
 - Opportunistic resources may be IPv6-only
- Storage technology
 - Making IPv6-capable taking much time – but good collaboration!
- **Need dual-stack federated storage services**
 - And dual-stack central WLCG and Experiment services


Phase 2: IPv6 deployment (2016-20)

Plan approved by WLCG Management Board

- **April 2017** – support for IPv6-only CPU client starts
 - Tier-1s to provide dual-stack storage testbed
- **April 2018**
 - Tier-1 dual-stack storage in production mode
- By end of LHC Run 2 (**end 2018**)
 - A large number of Tier-2s to migrate storage to dual-stack IPv6/IPv4
- **Monitoring is important**

Turning on IPv6 on CERN Tier-0 disk storage (EOS) in Jan 2018

Non-LHCOPN/non-LHCONE traffic




IPv6 traffic at CERN


LHCOPN and LHCONE IPv4 and IPv6 traffic volumes seen at CERN Tier0

Edoardo Martelli

IPv4 and IPv6 traffic volumes month by month



Percentage of IPv6 traffic over the total



IPv6 traffic on LHCOPN & LHCONE as seen at CERN

- > 50% of all traffic is IPv6
- From June 2019 onwards


[LINK](#) to these plots

Tier-2s: GGUS tickets to all Tier-2 sites

Andrea Sciaba



- The deployment launched in November 2017
- Steady progress (Oct 2019) ([status](#))
 - About 70% of T2 sites have storage on dual stack

Tier-2 IPv6 deployment status [11-10-2019]



FTS data transfers – Jan 2020 (>60%)

WLCG FTS IPv6 traffic



Phase 3: IPv6-only networking (2020-)

- Still need to be ready for use of (opportunistic) IPv6-only CPU
- **BUT there are other drivers for IPv6-only:**
 - lack of public IPv4 addresses in Data Centres
 - Use only IPv6 addresses for external public networking?
 - multiONE (different communities)
 - multiple overlay networks
 - sites likely have lack of sufficient IPv4 address space
 - SCITAG – packet marking
 - USA Federal Government – directive on IPv6-only

WLCG - from dual-stack to IPv6-only (CHEP2019) <https://doi.org/10.1051/epjconf/202024507045>

- Planning for an **IPv6-only** WLCG
- To **simplify** operations
 - Dual-stack infrastructure is the most complex
 - Dual-stack is less secure
- Large infrastructures (e.g. Facebook, EE/BT) use IPv6-only internally
- The plan - the goal we are working towards
 - IPv6-only for the majority of WLCG services and clients
 - With ongoing support for IPv4-only clients where needed
- Timetable to be defined

Summary



- WLCG is ready to support use of IPv6-only clients
- Tier-1s all have production storage accessible over IPv6
- Tier-2s ~90% sites are done
- Monitoring data transfers and configuration is essential
- Why do two dual-stack endpoints use IPv4 between them?
 - A priority for 2022
- Phase 3 – we are planning for move to IPv6-only services
 - Dual-stack is NOT the desired end-point!
- ***message to new research communities - build on IPv6 from start***

IPv6 WG – agenda 26 Oct 2022

- *Not our standard meeting agenda*
- Updates from
 - Sites
 - Experiments
 - WLCG Tier 2s - status
 - FNAL/ESnet
- Monitoring (CERN MONIT & FTS)
- Discovery and analysis of ongoing use of IPv4
- IPv6-only plans
- We allow for WLCG Data Challenge meeting (13:00 to 15:00)