

IPv6 WG status

David Kelsey

RAL, STFC, UK Research and Innovation

(on behalf of the HEPiX IPv6 Working Group)

WLCG GDB, CERN, 13 Sep 2023

On behalf of all members of the HEPiX IPv6 working group - (many thanks all!)



M Babik (CERN), M Bly (RAL), N Buraglio (ESnet), T Chown (Jisc),
D Christidis (CERN/ATLAS), J Chudoba (FZU Prague), 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), C Misa Moreira (CERN),
R Nandakumar (RAL/LHCb), K Ohrenberg (DESY), F Prelz (INFN), D Rand
(Imperial), A Sciabà (CERN/CMS), T Skirvin (FNAL)

- Many more in the past, and members join/leave from time to time
- *many thanks also to WLCG operations, WLCG sites, LHC experiments, networking teams, monitoring groups, storage developers...*

Outline

- IPv6 traffic growth
- Drivers for IPv6
- Deployment of IPv6/IPv4 dual-stack storage (good news!)
- Importance of monitoring
- Plans for IPv6-only WLCG
- Obstacles found and fixed
- More good news - IPv6 traffic levels on LHCOPN & LHCONE
- Obstacles to IPv6 still to be addressed
- Summary


IPv6 traffic continues to grow

Google

IPv6 Adoption


We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.

Native: 44.98% 6to4/Teredo: 0.00% Total IPv6: 44.98% | Sep 9, 2023



WLCG Data Transfers

%IPv6 - CMS FTS - GSIFTP & SRM



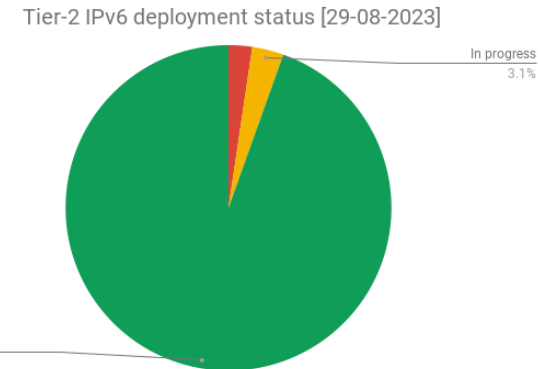

Stops when CMS moved away from GSIFTP & SRM

Reminder - Drivers for use of IPv6


- Sites running out of routable IPv4 addresses (avoid NAT)
 - Use IPv6 addresses for external public networking
- To be ready to support use of IPv6-only CPU clients
- There are **other drivers** for IPv6:
 - scitags.org – packet marking (in header of IPv6 packets)
 - Research Networking Technical Working Group ([RNTWG](#))
 - USA Federal Government – [directive](#) on “IPv6-only” (Nov 2020)
 - multiONE (several LHCONE’s for different communities)
 - either, the services must be in different IP LANs (suggests use of IPv6)
 - or use the scitags in IPv6 header flow label for policy based routing

Good news! - IPv6/IPv4 at Tier-1/2 sites

- Tier-1 complete
- Tier-2 deployment from Nov17
- ([status](#)) shows >94% T2 sites
 - **97%** of Tier-2 storage is dual



Experiment	Fraction of T2 storage accessible via IPv6
ALICE	91%
ATLAS	95%
CMS	100%
LHCb	100%
Overall	97%



Importance of monitoring

- We must monitor
 - deployment of IPv6-capable services
 - fraction of data transfers taking place over IPv6
- Monitoring implementations used for IPv6
 - perfSONAR
 - ETF - experiment test framework
 - FTS (File Transfer Service)
 - Network utilisation and traffic plots
 - e.g. IPv6 versus IPv4 on LHCOPN/LHCONE
- But in recent years some existing **monitoring stopped working**
 - FTS over WebDAV not tracking IPv6 (GSIFTP and SRM was instrumented)
 - work is ongoing to fix this problem

Our aim: IPv6-only on WLCG (CHEP2019)

<https://doi.org/10.1051/epjconf/202024507045>

- The end point of the transition from IPv4 is an **IPv6-only** WLCG core network
- To **simplify** operations
 - Dual-stack infrastructure is the most complex
 - Dual-stack has more security threat vectors
- Large infrastructures (e.g. Facebook, Microsoft,...) use IPv6-only internally
- The goal we are still working towards
 - “IPv6-only” for the majority of WLCG services and clients
 - With ongoing support for IPv4-only clients where needed/possible
- **Timetable** still to be defined (before LHC Run 4 should be very possible)

“Obstacles” to IPv6

There are many reasons stopping the full use of IPv6/IPv4

- Dual stack is an essential step on the journey to IPv6-only

The Obstacles that we have been addressing:

1. **WLCG Sites not yet deployed IPv6 networking**

~done

2. **Sites have IPv6 but Tier-2 has no dual-stack storage**

~done

3. **IPv6 monitoring not available or broken**


see next slide

4. **Service is dual-stack but IPv4 being used**

see next slide

- no time to describe all the obstacles we found and fixed

Some obstacles fixed (#3 and #4)



Some FTS **monitoring** now able to distinguish IPv6 from IPv4


ATLAS & CMS HTTP transfers into CERN (last year)

– IPv6 showing from August 2022 onwards

IPv6 is yellow

Data transfers into USA/ATLAS Great Lakes Tier 2 (AGTL2) Found to **use IPv4** even when both ends dual-stack (dCache/WebDAV) **java.net.preferIPv6Addresses (default: false) - Now set to “true”** Fixed at 17:00 on 14 Feb 2022 (confirmed in the plot!) This fix is essential for all dCache instances - fixed in v7.2.11

Many other uses of IPv4 have been investigated
htcondor, xRootD, FTS and Top-100 talkers LHCOPN using IPv4
Many problems identified and fixed.




WorkerNode migration to IPv6 at KIT

- <https://indico.jlab.org/event/459/contributions/11661/>
- Built an IPv6-only testbed
 - found problems with DNS, installation, squids, CVMFS, monitoring...
- So, took a different approach - migrate WN farm towards IPv6
- Needed detailed monitoring of ALL WN network traffic
 - packetbeat on all nodes storing to OpenSearch and analysed with Kibana
- Initially a small subset of WN, then the whole farm
 - 0.5 TB of data in 6 days
- at start the batch system was IPv4 then dual-stack deployed
- Apr22 - 28% IPv6; Dec22 - 67% IPv6
- Ongoing detailed work (many applications) to keep improving
- Shows how effective monitoring the details can be



Good news (IPv6 on WLCG) after removing several “obstacles” during the last year

LHCOPN network (at CERN) ~95% IPv6 during 30 days before CHEP




Storage - Tier-1 (100%)
and Tier-2 (93%)

LHCONE network at
CESNET (CZ)
- last 30 days
Ingress ~93% IPv6
Egress ~90% IPv6

Good news (2) - %IPv6 on LHCONE (Imperial College London)



dCache storage preference set to IPv6




Since Feb 2022
~90% IPv6

Obstacles to IPv6 - still to be addressed

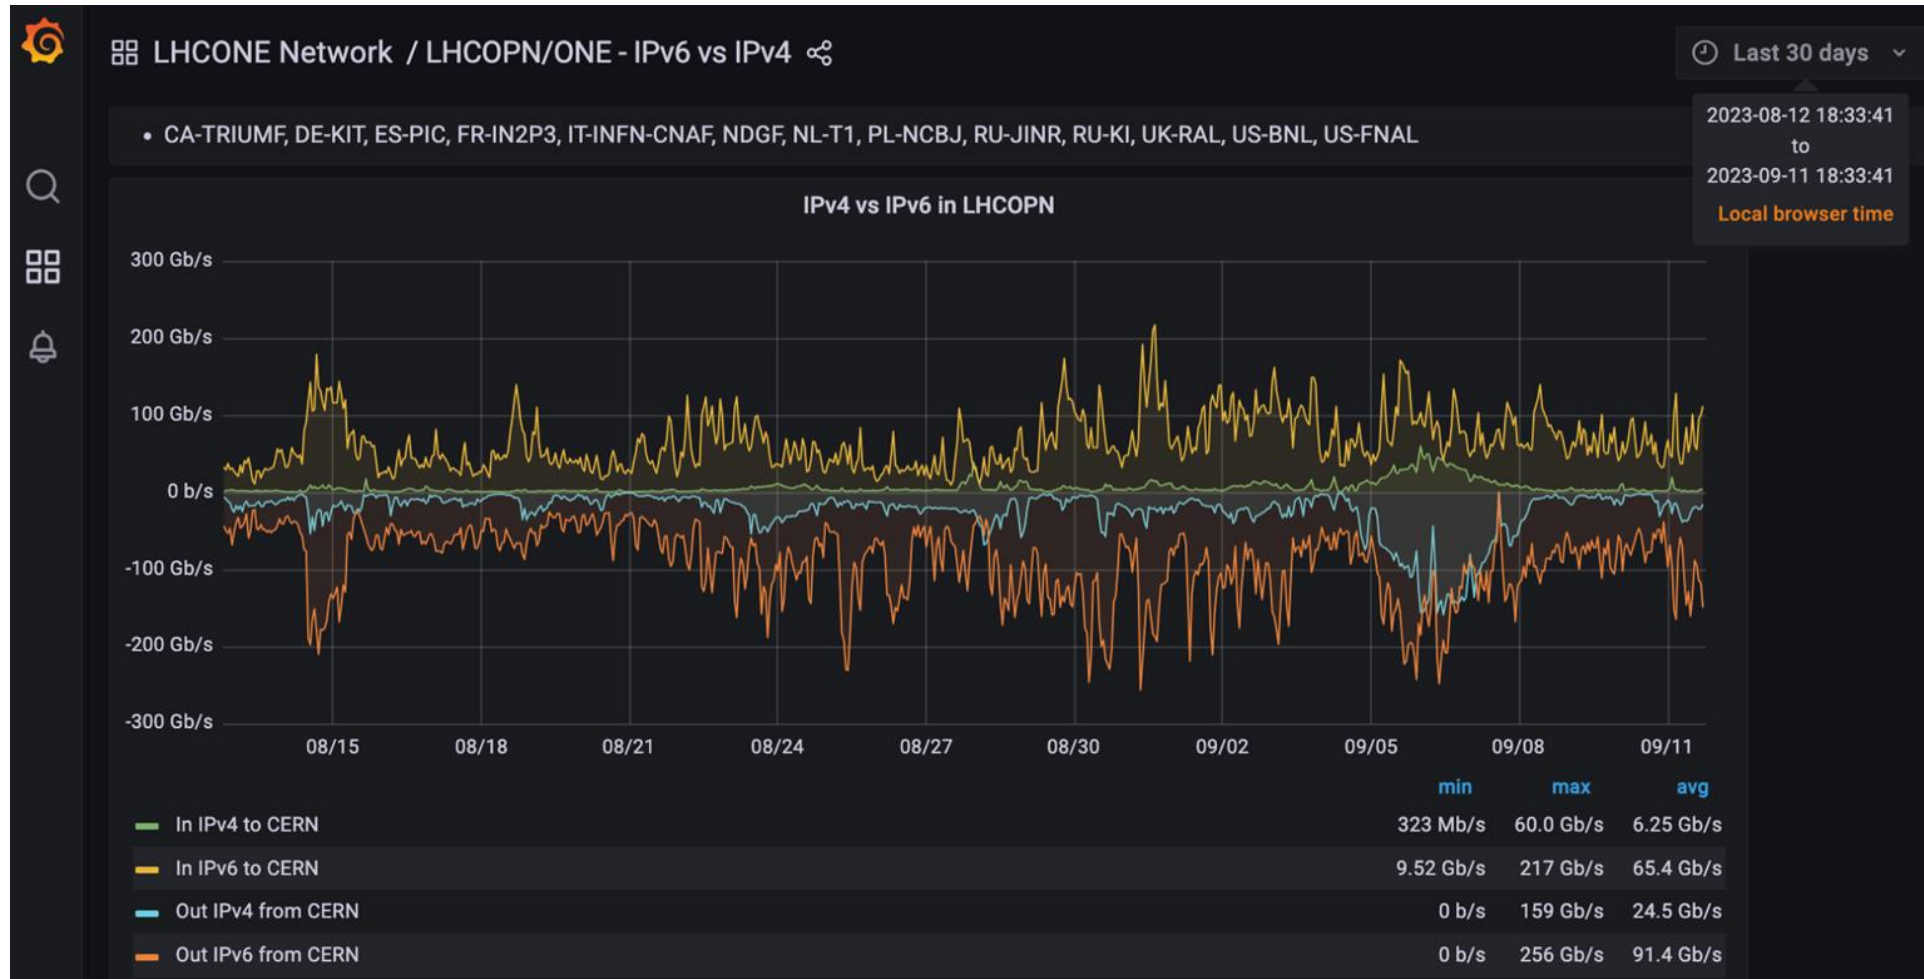
5. **Non-storage services not yet dual-stack**
 - a. ~60% of all WLCG services are dual-stack today
6. **WLCG client CPU (worker nodes, VMs, containers) many IPv4-only**
7. **Services/clients outside of WLCG Tier-1/Tier-2 not yet considered**
 - a. Tier-3, Public/Commercial Clouds, Analysis facilities, Experiment portals...
8. **Use of new or evolving technologies not yet tested or tracked**
 - a. New CPU architectures (GPU, non-x86, ...), container orchestration, ...
9. **“People” can be the obstacle**
 - a. they do not consider use of IPv6 or refuse to deploy!
10. **Analysis of old data using old software**
 - a. e.g. ALICE analysing Run 2 data with IPv4-only version of xRootD
 - b. very possible that all experiments have such a requirement

We will try to fix where possible - but much is outside of our control!

LHCOPN traffic - %IPv6 (large drops) August/September 2023



IPv4 and IPv6 traffic IN/OUT CERN (same timeslot)



IPv4 vs IPv6 traffic same timeslot (a T1)



Summary

- WLCG now supports IPv6-only clients
- Tier-1s: complete; Tier-2s: 97% storage is IPv6 capable
- Monitoring data transfers is essential - was broken and being fixed
- We have concentrated on removing obstacles to IPv6
 - LHCOPN/LHCONE is 90-95% IPv6 (after obstacles removed) - **but not always!**
- We continue to address more obstacles to IPv6 in WLCG
- **End point is still IPv6-only services (can we do this before Run 4?)**
- ***Message to WLCG sites and LHC experiments:***
 - *Deploy dual-stack on all services & CPU clients and prefer IPv6*
- ***Message to new research communities - build on IPv6 from start***

Questions, Discussion?