

# Identifying required steps towards future IPv6-only operation

*Duncan Rand, Tim Chown, Bruno Heinrich Hoeft, David Kelsey, Edoardo Martelli*

# History of IPv6 on WLCG

- Survey of HEPiX Community (September 2010) – “IPv6 readiness”
  - National NRENs are ready; universities and labs are not ready
- Some lack of IPv4 address space, including CERN
- IANA projecting imminent IPv4 address exhaustion
- Sep 2010 – memo from US Federal CIO to all depts including Department of Energy (HEP national labs) - Deploy dual-stack!
- Our middleware, software, technology and tools are not yet IPv6 capable
- This will take lots of time to fix - so started a working group in April 2011!
- (Offers of opportunistic CPU resources could arrive and be IPv6-only...)

# IPv6 Deployment on WLCG (Phase 2)

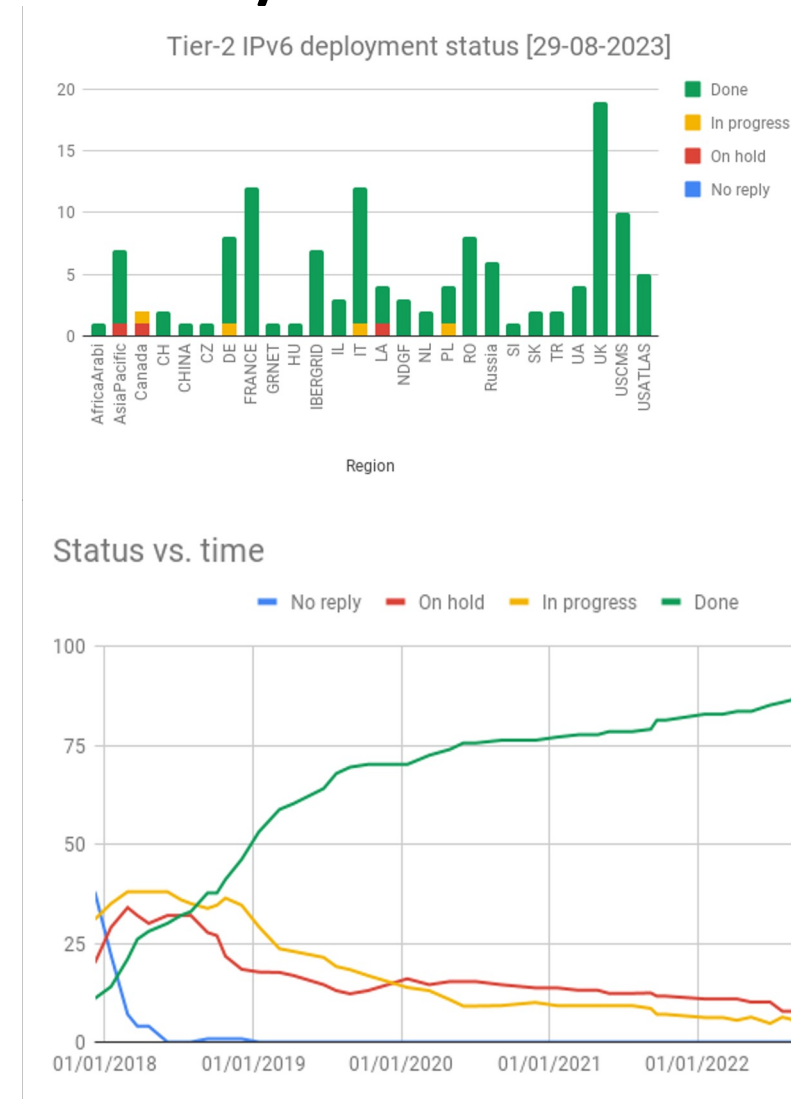
2017 onwards (as approved by WLCG Management Board)

- All Tier1 storage services in IPv6/IPv4 dual-stack mode from April 2018
- Aim for large number of dual-stack Tier-2s storage services **by end 2018**

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

- Tier-1 complete
- Tier-2 deployment from Nov 2017
- ([status](#)) shows >94% Tier-2 sites
  - **97%** of Tier-2 storage dual stack

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



Can work very well...

# Imperial London - LHCONe - 100 Gbps on IPv6

<https://shapingthefutureofjanet.jiscinvolve.org/wp/uncategorized/100gbps-of-cern-data-over-ipv6-on-the-janet-network/>




Figure 2 — The traffic levels seen in the network view correspond to those seen by the WLCG File Transfer Service (FTS) visualization tools.

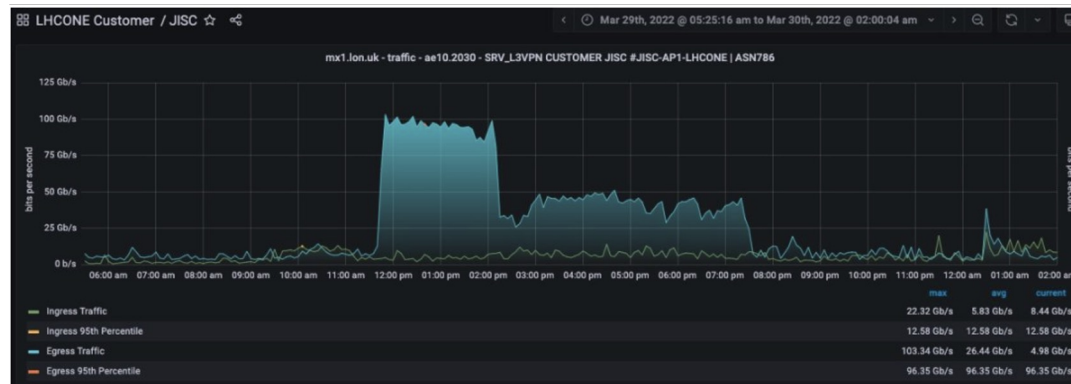


Figure 3 — It was also interesting to see this traffic reflected in the monitoring platform for the GÉANT pan-European research and education backbone network.



Figure 1 — Imperial monitoring shows the two-hour period where the 100G link was filled and where 100% of the LHCONe traffic was IPv6.

# Current drivers for use of IPv6 in 2023

- 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
  - Offers of opportunistic CPU resources could arrive and be IPv6-only
- US Federal Government updated [directive](#) in Nov 2020 to be “IPv6-only”
- New driver using an IPv6-specific capability:
  - the SciTag packet marking initiative using the IPv6 Flow Label
  - being coordinated by the WLCG Research Networking Technical Working Group ([RNTWG](#))

# Use of the IPv6 Flow Label for WLCG Packet Marking

Dale W. Carder - LBNL / ESnet (presenter)

Tim Chown - Jisc


Shawn McKee - University of Michigan

Marian Babik - CERN

**draft-cc-v6ops-wlcg-flow-label-marking**

*IETF 117, San Francisco, 25 July 2023*

- However, we still have a proportion of traffic going over IPv4
- For example: Tier-1 storage has been dual-stack for some time and yet ~14% of traffic over the last 6 months on LHCOPN is IPv4





# Identifying remaining sources of IPv4 traffic

- Likely that there are several sources of this IPv4 traffic
- Project goal is to identify and act on (or if not identify why we cannot act on) such IPv4 cases
- The end game being once everything is IPv6 we can turn off IPv4
  
- Part of this traffic will probably be between WN and storage systems
- The HEPiX IP6 WG and WLCG IPv6 TF have a recent proposal to make worker nodes dual-stack
- Presented by Andrea Sciabà to WLCG Operations Coordination meeting on 2 November
  - <https://indico.cern.ch/event/1341866/#1-ipv6-deployment-campaign>

# Worker Node migration to IPv6 at KIT

*(Bruno Hoeft)*

- <https://indico.jlab.org/event/459/contributions/11661/>
- migrate the CPU (Worker Node) farm towards IPv6
- 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
- Apr22 - 28% IPv6; Dec22 - 67% IPv6
- Ongoing detailed work (many applications) to keep improving
- Shows how effective monitoring and fixing can be
- Idea is to extend this approach to the OPN link between CERN and KIT

# 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, Microsoft,...) use IPv6-only internally
- The goal we are working towards
  - IPv6-only for the majority of WLCG services and clients
- Timetable still to be defined and agreed with Management Board
- Summarised in “WLCG - from dual-stack to IPv6-only” (CHEP2019)  
<https://doi.org/10.1051/epjconf/202024507045>

# DC24 Project: Description

- At the time of writing no specific link has been identified; this will be the first phase of the work through discussion with site operators. An LHCOPN link may be a good starting point
  - initially CERN & KIT
- Benefit: verify use of IPv6 on specific WLCG link(s), with a view to identifying IPv4 activity and determining how this may be moved to IPv6
- Success evaluation: identification of causes of IPv4 traffic on a link, and identification of steps required to migrate to IPv6-only

# DC24 Project: Work

- Identify specific link(s) to study: **initially CERN & KIT**
- Inspect traffic on the agreed link(s) to determine remaining use of IPv4
- Identify how remaining use of IPv4 can be removed, and all traffic be IPv6
- Establish a proposal for making specific link(s) IPv6-only in the future
- Include additional links where volunteered
- Note: the goal is not to migrate specific links to IPv6-only for DC24, but to investigate current status and what would be required