


IPv6 deployment on WLCG compute services

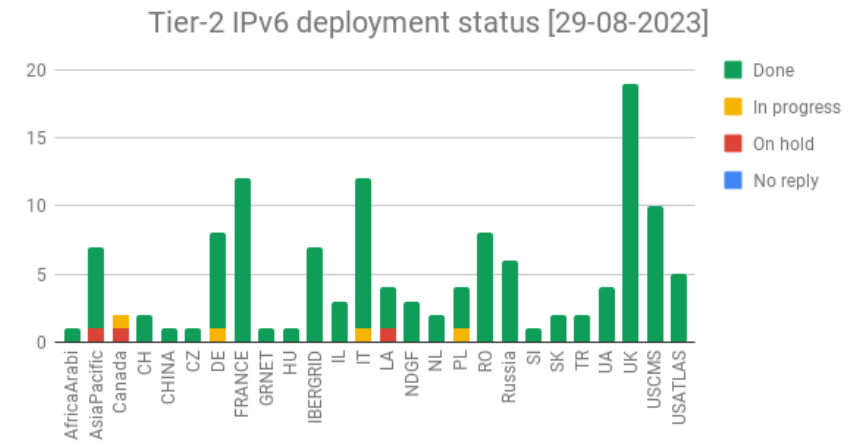
A. Sciabà on behalf of the HEPiX IP6 WG and WLCG IPv6 TF

WLCG MB meeting, 17 October 2023



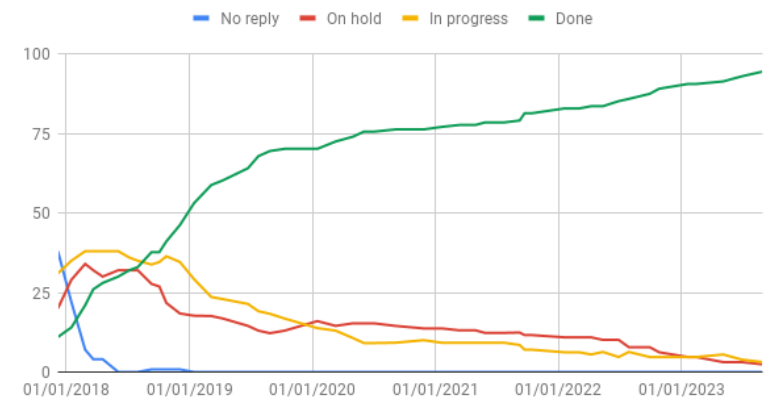
Status of IPv6 in WLCG

- WLCG's agreed end-point is to have IPv6-only sites and services
 - Discussed at many MB meetings, last time in February
 - Removes the complexity of operating dual stacks and improves security
 - Few IPv6-only clusters already there, e.g., at CERN
 - Dual stack is just a necessary step along this path
- The campaign to deploy IPv6 on storage services has practically ended
 - Its goal was to support IPv6-only WNs
 - Took very long mainly due to difficulties convincing the hosting institute to deploy IPv6
- Today, almost all WLCG sites have IPv6



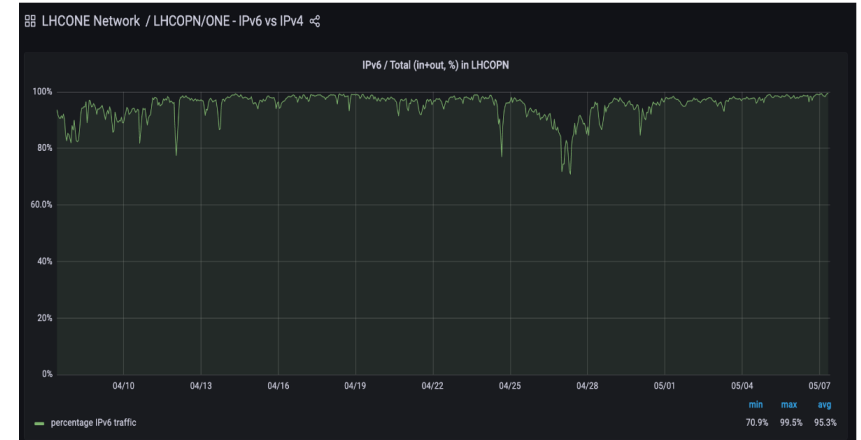
| Region | VO | T2 storage on IPv6 (%) |
|--------|-------|------------------------|
| | ALICE | 91 |
| | ATLAS | 95 |
| | CMS | 100 |
| | LHCb | 100 |
| | WLCG | 97 |

Status vs. time

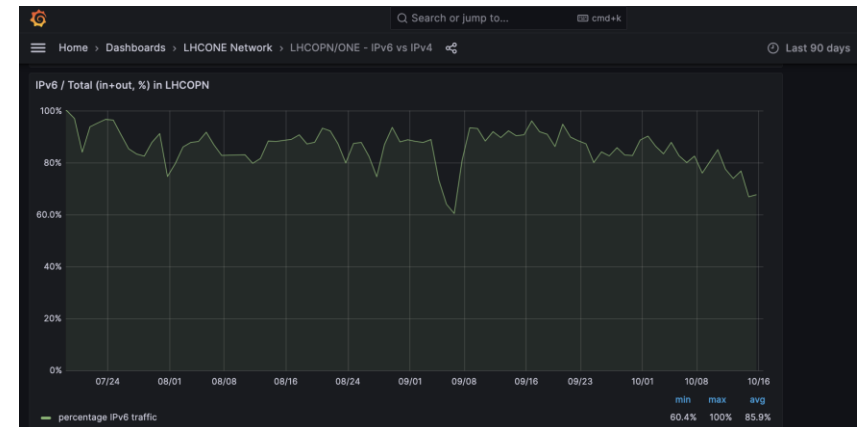


IPv6 on Compute Services: general motivations

- IPv4 addresses are a very scarce resource
 - May limit site expansion and creation of new sites
 - Want to avoid use of NAT
- Strong push by some governments to fully embrace IPv6
 - The US government mandates 50% of systems IPv6-only in 2024 and 80% in 2025, affecting FNAL and BNL
- Packet marking works only for IPv6
- All LHC experiments and Grid middleware fully support IPv6 in production
 - Has been the case for quite some time
- Almost all sites have IPv6 in their intranet
 - The biggest obstacle has been already removed
 - A large fraction of CEs and WNs has IPv6 now
- Eliminate the biggest remaining source of IPv4 traffic
 - Data transfers between WNs and storage systems



LHCPN traffic to/from CERN



ALICE

- Current status
 - Storage
 - 95% of the distributed disk space is dual stacked (only a handful of T2 endpoints are still IPv4-only)
 - Worker nodes
 - 49% of the job slots (150k jobs running) are connected to the ALICE central services over IPv6 (preferred by default)
 - 45% of the machines pass an explicit 'curl -6' test
- Requirements
 - ALICE is asking all remaining sites to dual stack their resources to enable sites that run out of IPv4 addresses to run jobs efficiently (by direct connection to data sources without NAT boxes or proxies)
 - For the same reason sites are encouraged to dual stack their entire infrastructure

ATLAS

- Current status
 - IPv6 support for storage already a requirement, mostly met
 - The long tail has proved troublesome
 - However, 50 % of the requests to the data-management service (Rucio) are done over IPv4
- Requirements
 - ATLAS recommends that all Grid-facing services and resources be requested to deploy IPv6

CMS

- Current status
 - The CMS submission infrastructure is using IPv6 when available
 - About 52% of the CEs already dual stack
 - About 25% of the WNs in the global pool and 74% in the CERN pool have IPv6
- Requirements
 - Need to meet the US government mandate
 - The CMS pilot factories in the US have to become IPv6-only!
 - Proposing deadline on 30 June 2024
 - Chosen to coincide with the CentOS 7 end-of-life

LHCb

- Current status
 - 100% of services and storage already dual stacked
 - Pure IPv6 nodes thoroughly tested with no loss of performance or capability
- Requirements
 - Support full scale deployment on CEs and WNs

DUNE

- Requirements
 - No objections to IPv6 (even only) on WNs
 - Fulfils the US government requirement
 - Move to IPv6-supported hardware Rucio services (FTS)
 - opendsciencegrid.org lacks IPv6!
 - Some sites (e.g., BNL) will need to dual stack internal services (Puppet, Condor) to dual stack
 - Likely timescale to be comfortable with pure ipv6 on WNs: ~12 - 18 months

Proposal submitted for approval to the MB

- All WLCG sites should offer IPv6 connectivity on their compute services (CEs and WNs) by June 30, 2024
 - Switching off IPv4 is not requested nor recommended: sites wishing do to it must discuss it with the supported experiments
- Progress would be tracked by launching a GGUS ticketing campaign, exactly as it happened for the storage services and perfSONAR
 - ETF could be used to test full functionality of the job management system over IPv6

Conclusions

- Almost all WLCG sites already adopted IPv6 for storage and a large fraction provide it on compute
- All LHC experiments fully support IPv6 in their computing operations
- The next logical step for the WLCG network strategy is to reach 100% adoption
 - Low risk, relatively short completion time expected in most cases
 - Leverage the resources of the HEPiX IPv6 working group and WLCG operations