



IPv6 deployment on WLCG compute services


A. Sciabà

HEPiX Spring Workshop, Paris, 15-19 April 2024




Status of IPv6 in WLCG

- WLCG's agreed end-point is to have IPv6-only sites and services
 - 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 on site




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



LHCOPN traffic to/from CERN



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
- **The proposal was accepted in October 2023**

ALICE

- Status at the start of the campaign
 - Storage
 - 95% of the distributed disk space is dual stacked (only a handful of T2 endpoints still IPv4-only)
 - Worker nodes
 - 49% of the job slots (150k jobs running) 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

- Status at the start of the campaign
 - 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

- Status at the start of the campaign
 - 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

- Status at the start of the campaign
 - 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

Ticketing campaign

- Launched on 28 November 2023 on all WLCG sites
 - Tier-3 sites not included
- **“Please deploy dual-stack connectivity (IPv4+IPv6) on your computing services (computing elements and worker nodes) as soon as possible and by 30 June 2024 at the latest”**
- Provide estimates for timescale and some details on the necessary steps
- It cannot meet the deadline, explain why

Current status


- Around 44% of sites completed IPv6 deployment
- A few regions (France included) have (almost) all sites on IPv6
 - And a few ones stand out for the opposite reason
- Status always visible from a [twiki page](#)



Time evolution

- About 20% of sites already had IPv6 on compute services at the start of the campaign
- 2.5 months left until the deadline
- A linear extrapolation would have 62% of sites done by the deadline
 - Need to pick up the speed!
 - Not to be too worried, several sites will take their time but still meet the deadline

Tier-1/2 CE/WN IPv6 deployment status vs. time



Analysis of the site replies

Type of answer	Sites
The site already had IPv6	33
The site is silent	25
The site acknowledges the request and adds it to their planning	17
The site promptly reacts and takes necessary action, deployment is completed quickly	15
The site needs to figure out how to do it or planning is not trivial	9
The site promptly reacts and takes necessary action, but deployment takes some time	6
The site needs to wait for some other upgrade	6
The site is silent and suddenly announces that deployment is completed or almost completed	6
The site did something but then went silent	2

Special concerns, problems or requests

- Some sites will couple it to other changes (e.g. new OS, new nodes, change internal routing, etc.)
- Some sites admit that they will not meet the deadline
 - Manpower issues, very complex or special site setup, other constraints, ...
- A few sites asked if they can go IPv6-only or at least test it!
 - Possible but must be arranged with their experiments
- A couple of sites are afraid that WNs (now behind NAT) will become more “exposed”
- A site needed some expert help but eventually managed
- **Overall, sites tend to take the request seriously and act on it**

An example: UAM-LCG2

- The only site which explained in detail what was done to deploy IPv6!
 - https://ggus.eu/index.php?mode=ticket_info&ticket_id=164411

```
Apologies for the delay in the response: this task made it into our TODO list and we wanted to get it done before replying.
```

```
To the extent of our knowledge, all the CEs and WNs at UAM-LCG2 should now be reachable over IPv6. The first step in the process required coordinating with our university's IT staff to make sure the necessary AAAA records were configured. Once they were ready, we began configuring the different machines. This configuration effort has been twofold:
```

- In order to make changes persistent we have modified the interface configuration files under `/etc/sysconfig`` for each machine.
- To avoid rebooting networking on each machine we manually configured the relevant interfaces through `iproute2` (i.e. `ip(8)``) so that there was no service interruption.

```
Once configured, we verified we had outbound connectivity by ping(8)`ing each machine from an external IPv6-enabled subnet: every machine replied as expected.
```

```
The following summarises our site's WNs and CEs:
```

Worker Nodes	IPv6 Address
wn247145.ft.uam.es	-> 2001:720:420:c003::230
wn247146.ft.uam.es	-> 2001:720:420:c003::231
wn247147.ft.uam.es	-> 2001:720:420:c003::232

Conclusions

- Almost all WLCG sites had already adopted IPv6 for storage and a large fraction also provided 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
- Deployment on compute services is proceeding reasonably well
 - Sites very rarely ask for help, they figure it out by themselves!
 - Few sites are pessimistic about their ability to meet the deadline
 - Too many sites (15%) still did not react to the ticket, though
 - 44% of sites have finished!

Discussion

- Desire to still retain NAT for WNs? Discuss
- Are there any sites in the audience which have not yet replied to the ticket? If so, can you comment why?
 - Are you struggling with a technical challenge?
- Anyone in the room who wants to share experiences? Or challenges? Already solved or still to be solved?
- Other issues?