Centre de Calcul
de l'Institut National de Physique Nucléaire
et de Physique des Particules

IPv6-only : a feedback from the LHC VOs

Aresh VEDAEE & Sebastien GADRAT (CC-IN2P3)

Drivers for moving to IPv6-only

- **Scarcity of routable IPv4 addresses**
  - top-level IPv4 addresses exhaustion occurred on 31 January 2011 (IANA)
- **Use IPv6 addresses for external public networking (avoid NAT)**
- **To simplify operations and the resolution of problems**
  - Dual-stack infrastructure is the most complex
- **Increase security**
  - LHCOPN links directly connecting Tier-1 data-centres often bypass border firewalls. Removing one protocol would decrease the attack surface
- **US Federal Government updated directive in Nov 2020 to be "IPv6-only"**
  - The US government mandates 50% of systems IPv6-only in 2024 and 80% in 2025, affecting FNAL and BNL
- **Opportunistic CPU resources offers could arrive and/or favour IPv6-only**
  - Several commercial cloud providers (e.g. AWS Lightsail, Hostynet, Servinga) offer discounted IPv6-only instances.
- **IPv6-specific capability**
  - the SciTag packet marking initiative using the IPv6 Flow Label
- **IPv6 routing is more efficient than IPv4 (17-30% saving in used electric power)**
  - Fixed size packet header – easier to forward
  - IPv4 networking includes lots of NAT – uses power
  - IPv4 ARP more power hungry than IPv6 SLAAC and Ras
  - Dual-stack networking is the worst for power used

CONCLUSION: Debate "IPv6 readiness" by the HEPIX Community (September 2010) => HEPIX IPv6 working group since April 2011

https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf
https://indico.cern.ch/event/1225423/contributions/5155111/attachments/2735664/4809865/IPv6%20compute%20deployment.pdf
https://indico.cern.ch/event/1307338/contributions/5663197/attachments/2749306/4785227/DC24-IPv6-project-1.pdf
https://doi.org/10.1051/epjconf/202024507045 ("IPv6-only networking on WLCG"[CHEP2019])

HEPiX IPv6 working group roadmap

- **Phase 1 - 2011-2016 - analysis, investigations, testbed, fix storage**
- **Phase 2 - 2017-2023 - WLCG Management Board approved following plan to deploy dual-stack storage on WLCG**
  - The 4 VOs must have the central service (CVMFS, FTS, Panda, Rucio, Pilot factories, etc.) in dual-stack mode by April 2017
  - Other services not directly used by jobs (certificate authorities, software repositories, the GOCDB/OIM, GGUS, VOMS, the ETF test infrastructure and the BDII) ready by April 2018.
  - All Tier1 storage services in IPv6/IPv4 dual-stack mode by April 2018
  - Large number of dual-stack Tier-2s storage services (and perfSonar if installed) by end of 2018
  - Deploy IPv6 on CPU services and worker nodes by 30 June 2024 at the latest (launched on 28 November 2023) => still ongoing!
- **Phase 3 - 2023-onwards - decommission legacy IPv4**
  - A large remaining source of IPv4 traffic data transfers between WNs and remote storage systems (services/clients outside WLCG Tier-1/Tier-2 not yet addressed: Tier-3, Public/Commercial Clouds, Analysis facilities, HPC Nvidia Skyway Ethernet/Infiniband gateway is IPv4-only, non WLCG use cases like BELLE II and JUNO, etc.)
  - Use of new or evolving technologies (new CPU architectures like GPU and non-x86, container orchestration, …) not yet tested or tracked

https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf
https://indico.cern.ch/event/1534556/contributions/6581016/attachments/3149362/5592220/Kelsey-IPv6-7oct25.pptx
https://indico.cern.ch/event/1307338/contributions/5663197/attachments/2749306/4785227/DC24-IPv6-project-1.pdf
https://hepix-ipv6.web.cern.ch/sites-ipv6-connectivity
https://hepix-ipv6.web.cern.ch/wlcg-applications
https://indico.cern.ch/event/1608816/contributions/6779058/attachments/3207872/5716748/20260128_hepixWG_IPv6-1.pdf
https://indico.cern.ch/event/1608816/contributions/6779055/attachments/3208816/5714712/WLCG-HEPIX6-20260128-Plan-for-IPv6-only-LHCOPN.pdf

What do we mean by IPv6-only?

Choices (one or more of):

- WLCG T1/T2 site services are IPv6-only (CE, SE, …)
- Other WLCG central services (e.g. Rucio, FTS, CVMFS, etc.) are IPv6-only
- LHCOPN and/or LHCONE networks are IPv6-only
- All WAN WLCG traffic is IPv6-only, including over the general R&E Internet

What does the IPv6 HEPIX working group propose we aim to achieve?

- All WLCG services (site and central) dual-stacked
- All WLCG WAN traffic on LHCOPN and LHCONE is IPv6-only
- If IPv4 is needed locally for other Research Communities – fine
- And data transfers over general R&E internet can still use IPv4

- **WLCG strategy** is to have **IPv6-only data-transfers** for HL-LHC
  - Discontinue IPv4 for data transfers over:
    - LHCOPN before DC27
    - LHCONE before DC29

https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf
https://indico.cern.ch/event/1598655/contributions/6988057/attachments/3260536/5821909/WLCG%20Operations%20Coordination%20towards%20HL-LHC_%20Current%20Focus%20Areas%20and%20Near-Term%20Plans%20-%20HEPiX%20Spring%202026%20(1).pdf

Dual-Stack: Storage deployment status

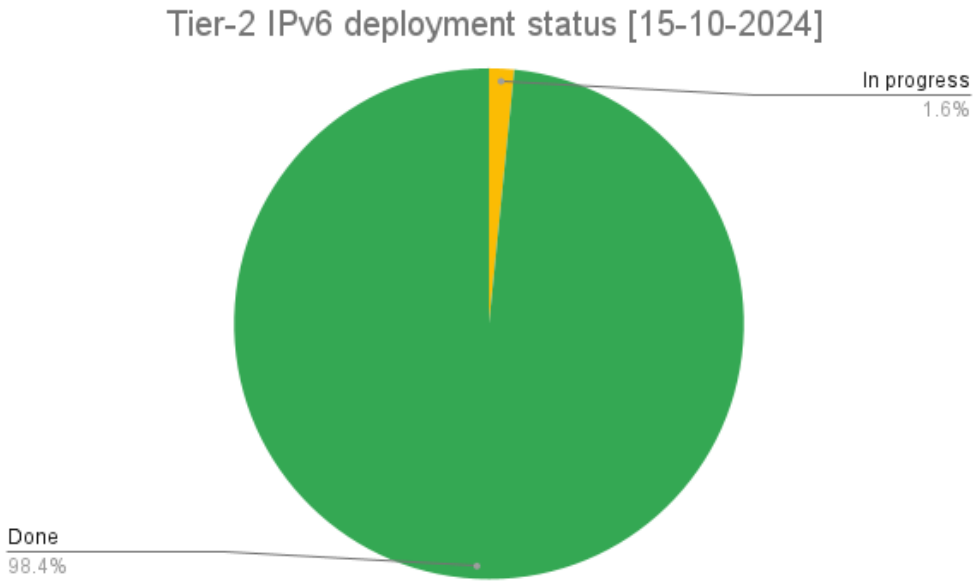

98% is dual stack

https://twiki.cern.ch/twiki/bin/view/LCG/WlcgIpv6

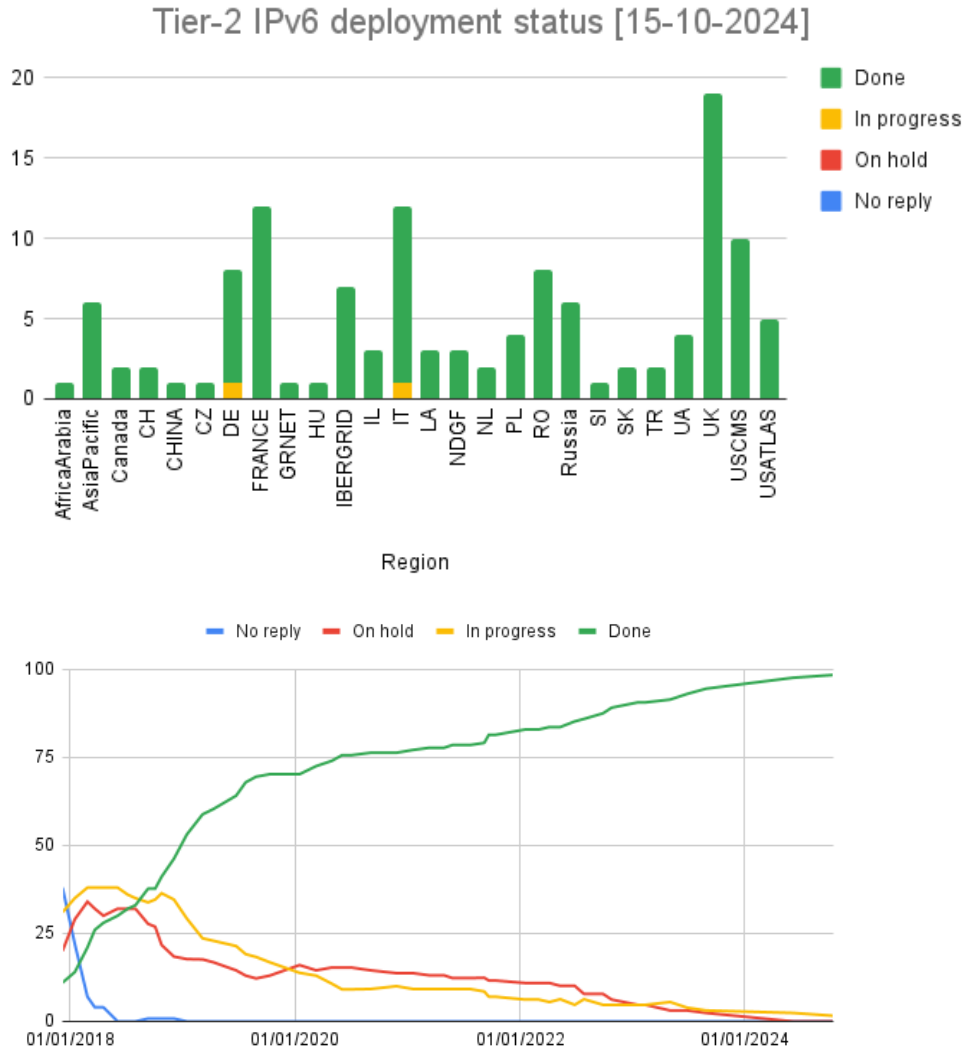

Dual-Stack: T1/T2 CPU deployment status

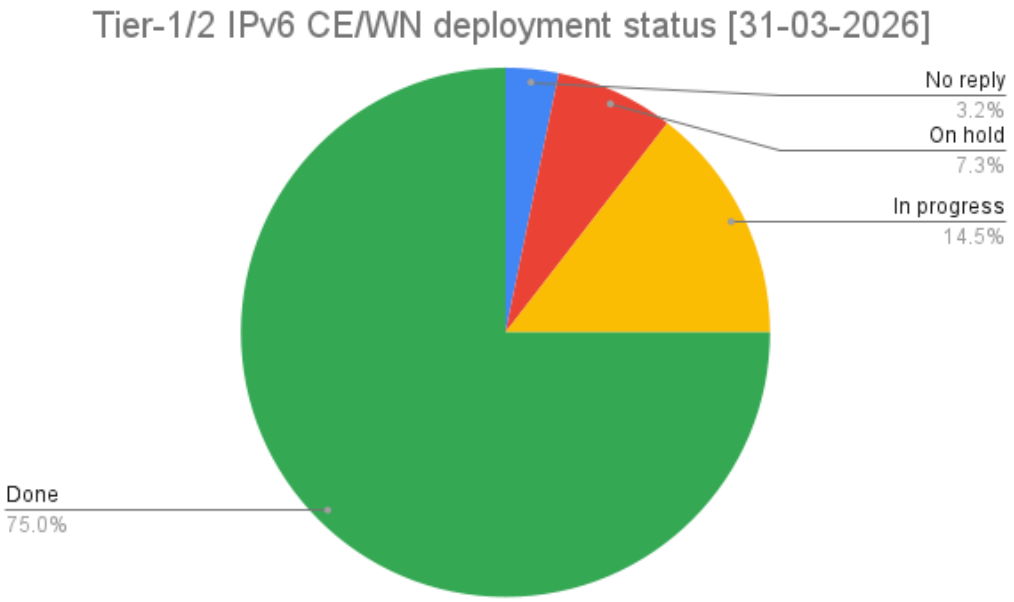

75% is dual stack

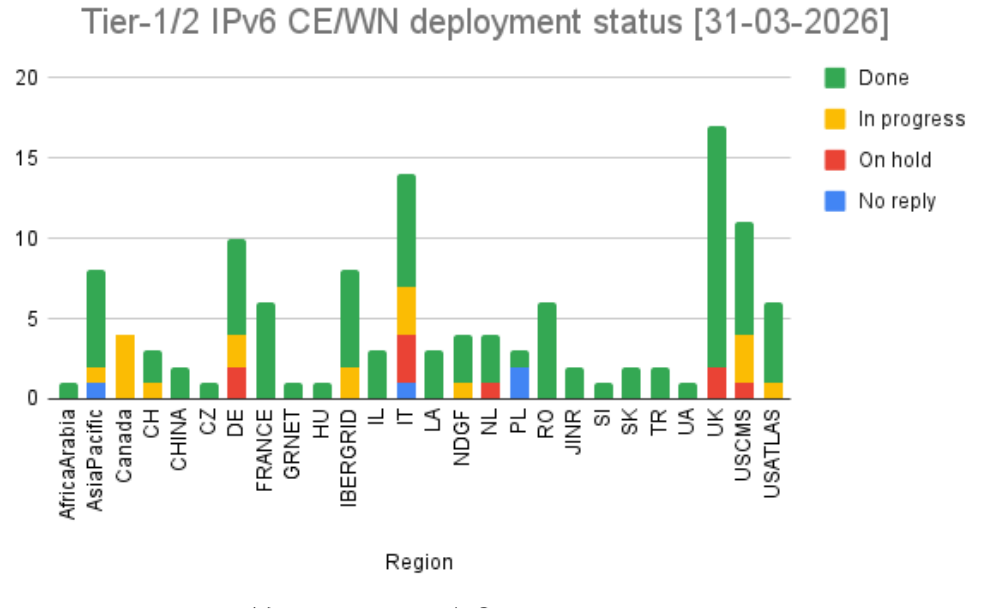

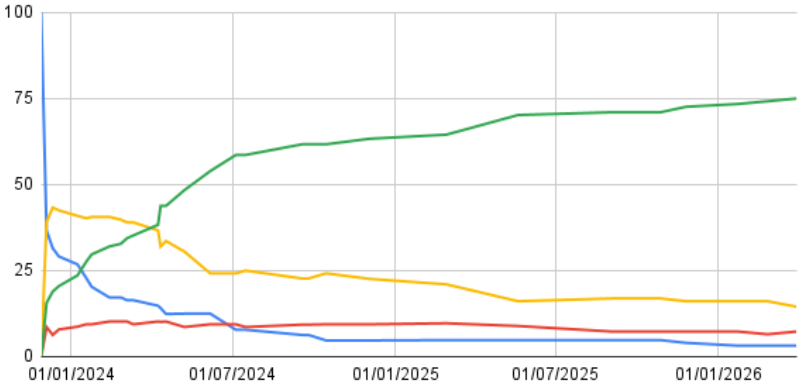

https://twiki.cern.ch/twiki/bin/view/LCG/WlcgIpv6

LHCOPN IPV6 traffic status

LHCOPN % of traffic (T0-T1s & T1s-T1s) which is IPv6 - last 2 years. (85% to 98%)

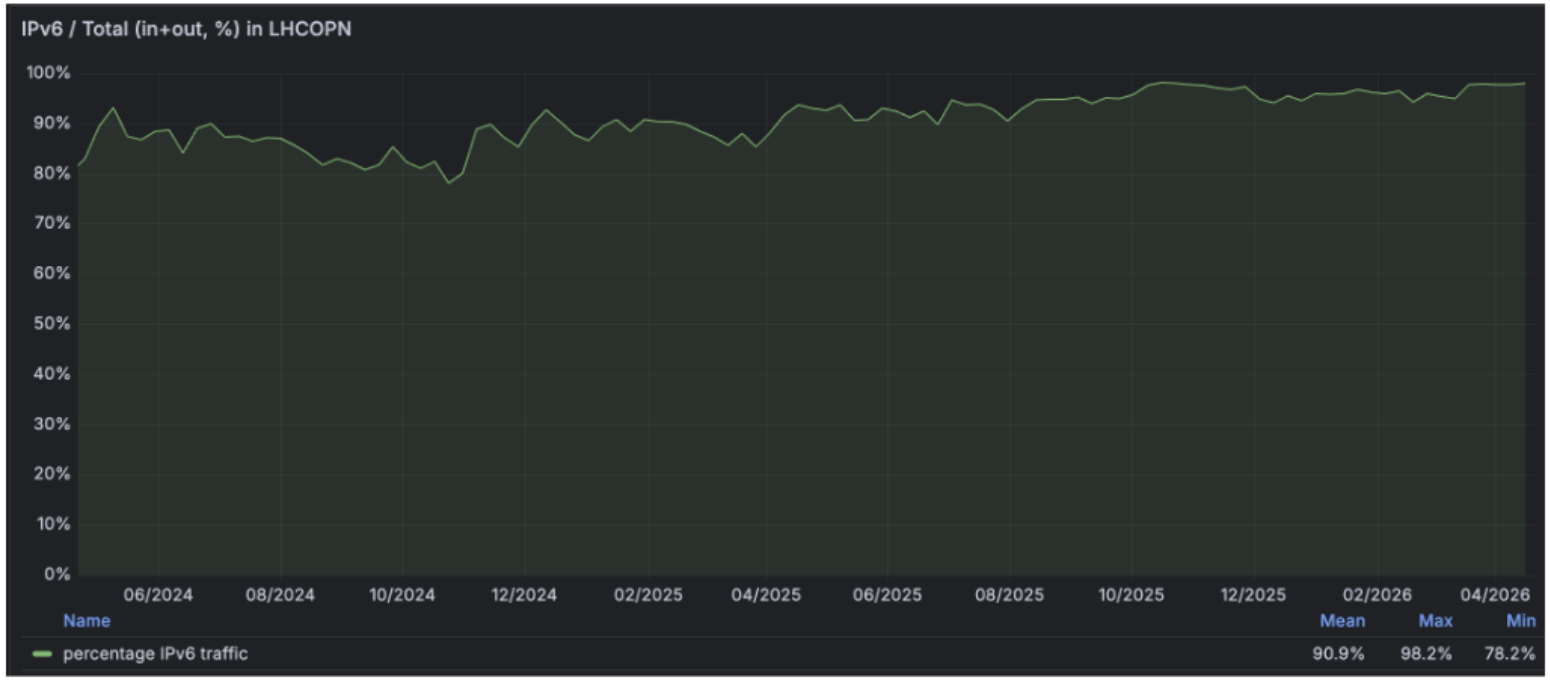

https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf

IPv6-only Network traffic: news from Tier1 sites

- **BNL and FNAL:**
  - LHCOPN IPv6 only mini-challenge (Sep. 2025) (link)
  - LHCONE IPv4 blackout mini-challenge (Feb. 2026) (link)

  For ATLAS and CMS there was no noticeable impact
  - Any remaining IPv4 Traffic re-routed onto LHCONE or R&E networks
  - Few non-WLCG experiments that were using LHCOPN and now use LHCONE

- **PIC successfully dropped IPv4 on LHCOPN on 23 Feb 2026 (link)**
  - No noticeable impact on ATLAS, CMS, LHCb traffic
- **New Poland Tier1 PL-NCBJ (link)**
  - Disabled IPv4 on LHCOPN link (13 April 2026)
  - Part of storage (40%) run on IPv6-only

  Issues with LHCb do to CNAF nodes not supporting IPv6 (overlooked by LHCb Dirac configuration) and issues with some CMS sites like JINR (due to their IPv6 misconfiguration)
- **Next Candidates to remove IPv4 from LHCOPN**
  - CN-IHEP Tier1 (link)

  Once IPv4 is removed in LHCOPN, the traffic should be switched to LHCONE, but the reality is that (for ALICE traffic from IHEP to KISTI) KISTI didn't advertise any address to LHCONE and no INTERNET connection
  - Nordic Tier1 (link)

  The plan was to:
  - Stop routing IPv4 over our LHCOPN connections to CERN
  - The few things that still needed IPv4 should fallback over LHCONE and general-purpose internet
    - Data transfers not noticeably affected
    - But jobs had issues (e.g. ALICE job management frameworks need IPv4 still), so we turned it back on again a bit over a week later

ATLAS views on IPv6-only (1/2)

- **VO Acceptance & Strategy**: ATLAS heavily relies on a highly mesh-connected global infrastructure via Rucio and FTS and they can not lose connectivity to storage. However, they are fully prepared for sites to go IPv6-only provided that central services and other sites can communicate over IPv6.
- **Application & Central service compatibility**: The ATLAS core software (e.g. Athena, Rucio clients) is IPv6-ready. Core services (Panda, Rucio, Harvester, CRIC etc) have been on dual stack for 2 years. Squids are purely local now, while some varnish instances and university DNSs servers are not dual-stack.
- **Endpoint compatibility**:
  - Still many site service endpoints not dual-stack:
    - Dual-stack Storage: 83% => (mostly tier3s and just 2 Tier2s)
    - Dual-stack Compute: 79%
  - IPv6 site tracking is followed in Jira ticket (https://its.cern.ch/jira/browse/ADCINFR-277)

ATLAS views on IPv6-only (2/2)

- **Monitoring**:
  - Hammer Clouds PFT [link] and AFT [link] testing for IPv6
  - Still a problem trying to monitor ipv6/ipv4 on the FTS dashboard

French sites look ok

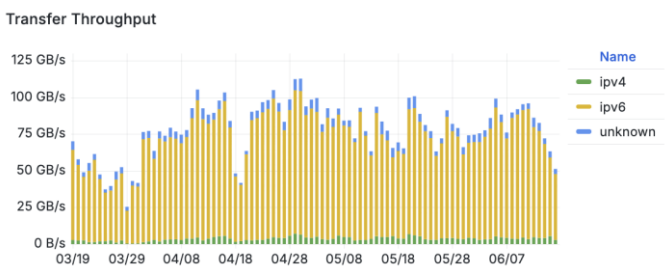

- **Operational risk**: Medium-low. A small fraction of ATLAS Tier-2/Tier3 storage is not yet IPv6-accessible. If a French site goes IPv6-only, it may lose connectivity to those specific storage endpoints until they upgrade. This is a known issue and sites are being pushed to complete deployment.
- **Future prospects**: all LHCOPN and LHCONE traffic on ipv6 by DC27
  - LHCOPN doesn't have the VO/activity information
    - fireflies/scitag monitoring is promising => Goal: 80% WLCG traffic monitored for DC27.
    - Unfortunately it needs to be enabled on the storages => Will need instructions and a campaign
    - Requires also sites to add their network routes to WLCG CRIC

https://indico.cern.ch/event/1608816/contributions/6779058/attachments/3207872/5716748/20260128_hepixWG_IPv6-1.pdf

ALICE views on IPv6-only (1/2)

- **VO Acceptance & Strategy**: ALICE operates a strictly federated storage model. Because their computing model forces jobs to strictly interact with the grid environment via their custom middleware, a full IPv6 storage deployment across the grid is practically required before individual sites can cleanly run IPv6-only WNs without edge cases.
- **Central services & Application compatibility**: modern jAliEn framework (Java-based) and current ROOT6 / XRootD versions are fully IPv6-compliant. But ALICE has a specific constraint: re-analysis of older data (e.g., LHC Run 2) sometimes requires legacy software that may not be IPv6-capable. This is flagged as a hard problem by the working group. But in practice it's a moot point.
- **Endpoint compatibility**:
  - Still many site service endpoints not dual-stack:
    - Dual-stack Storage: 94%
    - Dual-stack Compute: 75%

ALICE views on IPv6-only (2/2)

- **Monitoring:**
  - VOBOX and Storage endpoints: http://alimonitor.cern.ch/ipv6/
  - CE endpoints: http://alimonitor.cern.ch?6213
  - Connections on IPv6 to the central services: http://alimonitor.cern.ch?6214

For some French sites look there is reported a lack of DNS address or connectivity (see backup slides)

- **Operational risk** : Medium. ALICE is traditionally more sensitive to network topology gaps. For routine operations, ALICE is largely ready. However, in a couple of cases dropping IPv4 was incomplete and this has created asymmetrical routing issues.
- **Future prospects**: LHCOPN/LHCONE should be ok to drop IPv4 as long as IPv4 on the general internet remains available.
  - T1 storages are dual stacked for raw data replication as 3rd party XRootD transfers
  - Concerning compute nodes, if for any reason the jobs have to fall back to it (accessing a file on the 6% of the IPv4-only storage volume) then they require some connectivity (NAT, firewalled etc).

https://indico.cern.ch/event/1608816/contributions/6779058/attachments/3207872/5716724/IPv6%20status%20@%202026-01-29.pdf

LHCb views on IPv6-only (1/2)

- **VO Acceptance & Strategy:** LHCb data management is tightly controlled but highly automated via DIRAC. LHCb accepts and encourages IPv6 deployment, provided it passes their rigorous DIRAC data management testing
- **Central services & Application compatibility:** DIRAC (LHCb's workload management system) was made IPv6-compatible years ago. The LHCb simulation and analysis software (Gaudi ecosystem) handles data access cleanly over IPv6 via standard ROOT/XRootD plugins => LHCb can run on IPv6-only CPUs since 2017
- **Endpoint compatibility**:
  - Most site service endpoints in dual-stack:
    - Dual-stack Storage: 100%
    - Dual-stack Compute: 80%

LHCb views on IPv6-only (2/2)

- **Monitoring:** no particular monitoring for IPv6. ETF tests for storage not IPv6 compliant yet. Monitoring with FTS still with “unknowns”

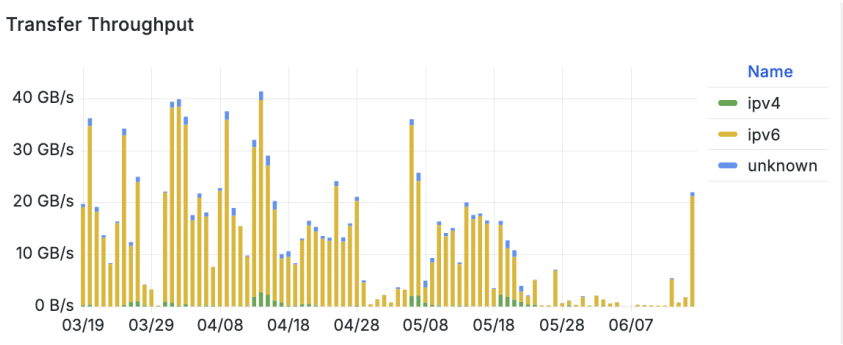

- **Operational risk:** Low. LHCb is in a similar position to CMS—fully deployed and tested. If a French site transitions to IPv6-only, DIRAC will properly schedule jobs there, but it will prioritize matching jobs to data that resides on dual-stack or IPv6-only storage to prevent failures.
- **Future prospects**: in favour of testing LHCONE IPv6-only for DC27

https://indico.jlab.org/event/459/contributions/11618/attachments/9663/14087/Kelsey11may23-IPv6.pdf
https://indico.cern.ch/event/1598655/contributions/6988152/attachments/3261920/5825358/IPv6-Kelsey-23Apr26-V3b.pdf

CMS views on IPv6-only (1/2)

- **VO Acceptance & Strategy:** CMS uses a highly flexible data management system, i.e. AAA/XRootD federation which requires both storage and compute endpoints to be IPv6 compliant.

- **Central services & Application compatibility:**
  - The CMS submission infrastructure based on GlideInWMS and HTCondor is fully IPv6 compliant
  - The official soft CMSSW is now natively IPv6 (tested and validated)
    - However some legacy CMSSW versions still require IPv4

- **Endpoint compatibility**:
  - Full IPv6-only workflows validated at several sites (lately, CERN, KIT and Bristol)
  - New T1_PL_NCBJ will provide soon a IPv6-only storage
  - Most sites are providing dual stack

Main resources (2 following talks as well as some internal discussions regarding operational issues and prospects)
- F2F IPv6 WG @Manchester January 2026
- HEPiX IPv6 WG report @Lisbon May 2026

CMS views on IPv6-only (2/2)

- **Monitoring:**
  - FTS monitoring has been improved, though still some unknowns (common issue for other VOs)
  - dCache and XRootD now use ipver (like FTS), should be ok, but need to be validated
- **Operational risk:** Low
  - Some sites use network low latency switches that does not support IPv6 (notably CNAF)
  - Some HPC resources don't provide IPv6, and the Nvidia Skyway IPv6 incompatibility (not a real bug, but need to be worked out)
    - CNAF, Nebraska, USCMS and Aachen
  - Most remaining issues coming from the network
    - Some sites' ISP don't support (yet) IPv6
    - Issues with firewalls
    - Network IPv6 peering sometimes break (especially for asian sites)
  - No particular issues for the french sites

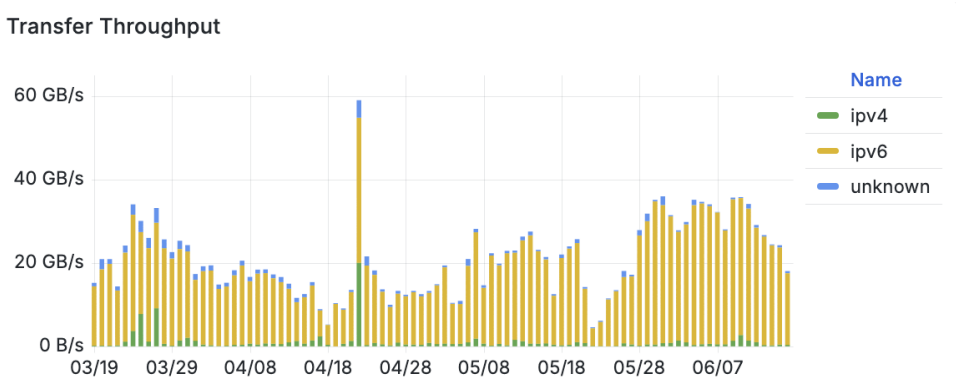

- **Future prospects (plan):**
  - Although CMS is a big supporter of enabling IPv6, they don't push sites to go IPv6-only for the time being
  - Want to validate the IPv6-only storage at T1_PL_NCBJ
  - Then discuss the few remaining issues and legacy IPv4 resources

General recommendations for sites

- **Dual-stack first:** The WLCG guidance is to run dual-stack until all peer sites and services are confirmed IPv6-ready. IPv6-only is not yet endorsed for production across WLCG—dual-stack remains the operational mode.
- **IPv6-only Storage/Compute:** turning off IPv4 on your Worker Nodes may need to implement NAT64/DNS64 [https://www.researchgate.net/publication/346917619_IPv6-only_networking_on_WLCG?utm_source=chatgpt.com], in case a job attempts to contact an external, non-WLCG IPv4-only service (such as GitHub, an external database, a minor Linux package repository, cloud APIs, etc.). It is highly recommended to keep Storage Elements (SEs) dual-stacked for a bit longer if they serve external, non-LHC communities that lag behind on IPv6 adoption.
- **IPv6-only network traffic:** LHCOPN is the first candidate network to disable IPv4 peering, but this hasn't happened yet. Ensure that if your LHCOPN link drops IPv4, your network perimeter is configured so that any accidental or legacy IPv4 traffic safely reroutes via LHCONE or the general Research & Education (R&E) or general public internet rather than being dropped entirely.
- **Coordinate with your VO:** Before switching to IPv6-only, confirm with experiment operations that all critical endpoints (especially data sources for reprocessing campaigns) are reachable. And legacy analysis software (like for ALICE) should be negligible.

SUMMARY

- **Dual-stack site deployment:**
  - Storage ended but for CPU and WN's still underway
- **IPv6-only network traffic:**
  - Data transfers deadlines:
    - from LHCOPN before DC27
    - from LHCONE before DC29
  - Fallback routing is needed for remaining IPv4-traffic
  - Monitoring (FTS, LHCOPN) needs improvements
    - Monitor network traffic locally and support central service monitoring
- **IPv6-only storage/compute:**
  - Do not oppose IPv6-only CPU resources anymore but some risk remains
  - The experiments somewhat differ in how sensitive they are to incomplete IPv6 coverage.

BACKUP

Refresh - **US Government IPv6 Mandate**

FY23	• All new federal systems to be IPv6 enabled at deployment. • **20%** of all networked federal systems **IPv6-only**
FY24	• **50%** of all networked federal systems **IPv6-only**
FY25	• **80%** of all networked federal systems **IPv6-only** • Identify, plan, schedule retirement/replacement of remaining networked systems that cannot be converted to IPv6-only

- US National Labs (T1s) are included; **but** university-run T2s are not subject to the mandate

Fermilab

https://www.ipv6.org.uk/wp-content/uploads/2022/10/11_GridPP-WLCG_IPv6_28nov22.pdf

VO	T2 storage on IPv6 (%)
ALICE	94
ATLAS	98
CMS	100
LHCb	100
WLCG	98

(checked on 15-10-2024)

https://twiki.cern.ch/twiki/bin/view/LCG/WlcgIpv6

TOTAL	54 SEs in 36 sites		47 have IPv6 DNS records (31 sites)	45 are reachable on IPv6	T0: 219.3 PB / 219.3 PB (100%) T1: 65.21 PB / 65.21 PB (100%) T2: 57.97 PB / 77.45 PB (74%) All: 342.5 PB / 362 PB (94%)	109 CEs in 65 sites		66 have IPv6 DNS records (39 sites)	62 can reach the world on IPv6

http://alimonitor.cern.ch/ipv6/

Site	SE					CE			
	Name	Endpoint	IPv6 address	Reachable	Size	Name	Endpoint	IPv6 address	Outgoing connectivity
9. CCIN2P3	CCIN2P3::TAPE	ccxroot.in2p3.fr	[2001:660:5009:84:134:158:239:103]			CCIN2P3::HTC	ccwlcgalice01.in2p3.fr	[2001:660:5009:304:134:158:235...	
	CCIN2P3::SE	ccxrdralice.in2p3.fr	[2001:660:5009:84:134:158:239:108]		9.379 PB	CCIN2P3::HTC_2	ccwlcgalice02.in2p3.fr	[2001:660:5009:304:134:158:234...	
18. GRIF	GRIF::EOS	eos.grif.fr	[2001:660:302c:23:134:158:132:49]		2.482 PB				
19. GRIF_IPNO						GRIF_IPNO::HTC	ipnvobox2.in2p3.fr	[2001:660:3024:200:134:158:78:...	
						GRIF_IPNO::IJCLAB	grid-79231.ijclab.in2p3.fr	No IPv6 DNS addr	
20. GRIF_IRFU						GRIF_IRFU::LCG	node09.datagrid.cea.fr	[2001:660:3031:110:4:0:0:9]	
13. CLERMONT						Clermont::LCG	clrvoboxalicce1.in2p3.fr	No IPv6 DNS addr	
						Clermont::ARC	clrvoboxalice02.in2p3.fr	No IPv6 DNS addr	
57. STRASBOURG_IRES	STRASBOURG_IRES::SE2	sbgse6.in2p3.fr	[2001:660:4705:d002:134:158:150:228]		652.5 TB	Strasbourg_IRES::LCG	sbgvoboxalice.in2p3.fr	No IPv6 DNS addr	
26. IPNL						IPNL::LCG	lyogrid08.in2p3.fr	[2001:660:5009:a013:0:0:0:f6]	

http://alimonitor.cern.ch/ipv6/

`curl -6` connectivity to CERN

	Series	Last value	Min	Avg	Max
1.	CCIN2P3	100	100	100	100
2.	Clermont	97.93	50	97.53	100
3.	GRIF_IPNO	98.01	10.77	99.77	100
4.	GRIF_IRFU	100	98.59	99.99	100
5.	Strasbourg_IRES	100	0	97.67	100

http://alimonitor.cern.ch?6215

PFT for IPv6

Site	S	R	C	F	Eff
IN2P3-CC_16	0	0	6	0	1.00
IN2P3-CC	0	0	6	0	1.00
IN2P3-LAPP	0	0	6	0	1.00
IN2P3-CPPM-TEST	0	0	6	0	1.00
IN2P3-CPPM	0	0	6	0	1.00
GRIF-LPNHE	0	0	5	0	1.00
GRIF-LAL	0	0	5	0	1.00
GRIF-IRFU	0	0	5	0	1.00

https://hammercloud.cern.ch/hc/app/atlas/test/20332824/

AFT for IPv6

Site	S	R	C	F	Eff
IN2P3-CC	0	0	6	0	1.00
IN2P3-LAPP	0	0	6	0	1.00
IN2P3-CPPM	0	0	6	0	1.00
GRIF-LPNHE	0	0	4	1	0.80
GRIF-LAL	0	0	5	0	1.00
GRIF-IRFU	0	0	5	0	1.00

https://hammercloud.cern.ch/hc/app/atlas/test/20332831/

ATLAS French Cloud on IPv6

```
$ cat list.txt | grep ipv6 | grep "\.fr"
67 ipv6 cccondorce01.in2p3.fr
68 ipv6 cccondorce02.in2p3.fr
69 ipv6 cccondorce03.in2p3.fr
70 ipv6 cccondorce04.in2p3.fr
71 ipv6 cccondorce05.in2p3.fr
130 ipv6 clrarcce01.in2p3.fr
131 ipv6 clrarcce02.in2p3.fr
132 ipv6 clrarcce03.in2p3.fr
148 ipv6 grid-72065.ijclab.in2p3.fr
181 ipv6 lapp-ce04.in2p3.fr
182 ipv6 lapp-ce05.in2p3.fr
183 ipv6 lapp-ce06.in2p3.fr
192 ipv6 llrcondorce.in2p3.fr
193 ipv6 lpnhe-ce02.in2p3.fr
194 ipv6 marcce01.in2p3.fr
195 ipv6 marcce02.in2p3.fr
197 ipv6 node16.datagrid.cea.fr
205 ipv6 pre7230.datagrid.cea.fr
5 ipv6 ccperfsonar1.in2p3.fr
6 ipv6 ccperfsonar2.in2p3.fr
7 ipv6 clrperf-bwctl.in2p3.fr
8 ipv6 clrperf-owamp.in2p3.fr
41 ipv6 lapp-ps01.in2p3.fr
42 ipv6 lapp-ps02.in2p3.fr
55 ipv6 llrpsonar1.in2p3.fr
56 ipv6 llrpsonar2.in2p3.fr
57 ipv6 lpnhe-psb.in2p3.fr
58 ipv6 lpnhe-psl.in2p3.fr
61 ipv6 marperf01.in2p3.fr
62 ipv6 marperf02.in2p3.fr
63 ipv6 nanperfs01.in2p3.fr
64 ipv6 nanperfs02.in2p3.fr
100 ipv6 perfsonar01.datagrid.cea.fr
104 ipv6 perfsonar02.datagrid.cea.fr
155 ipv6 psonar1.lal.in2p3.fr
156 ipv6 psonar2.lal.in2p3.fr
7 ipv6 ccdavatlas-tape.in2p3.fr
8 ipv6 ccdavatlas.in2p3.fr
10 ipv6 ccxrootdatlas.in2p3.fr
13 ipv6 clrlcgse01.in2p3.fr
27 ipv6 eos.grif.fr
47 ipv6 lapp-dccentral01.in2p3.fr
61 ipv6 marsedpm.in2p3.fr
7 ipv6 clrsquid00.in2p3.fr
8 ipv6 clrsquid01.in2p3.fr
22 ipv6 jcsquid02.ijclab.in2p3.fr
24 ipv6 lapp-squid03.in2p3.fr
25 ipv6 lapp-squid04.in2p3.fr
33 ipv6 lpnhe-cache.in2p3.fr
34 ipv6 marsq01.in2p3.fr
35 ipv6 node24.datagrid.cea.fr
36 ipv6 node25.datagrid.cea.fr
2 ipv6 ccatlasvrn01.in2p3.fr
```

```
$ cat list.txt | grep ipv4 | grep "\.fr"
21 ipv4 grid06.lal.in2p3.fr
7 ipv4 ccatlassquid.in2p3.fr
25 ipv4 login5.frontera.tacc.utexas.edu
26 ipv4 marsq02.in2p3.fr
27 ipv4 marsq03.in2p3.fr
```

https://its.cern.ch/jira/secure/attachment/712773/260223-all-ipv4-ipv6-none

VO deployment status

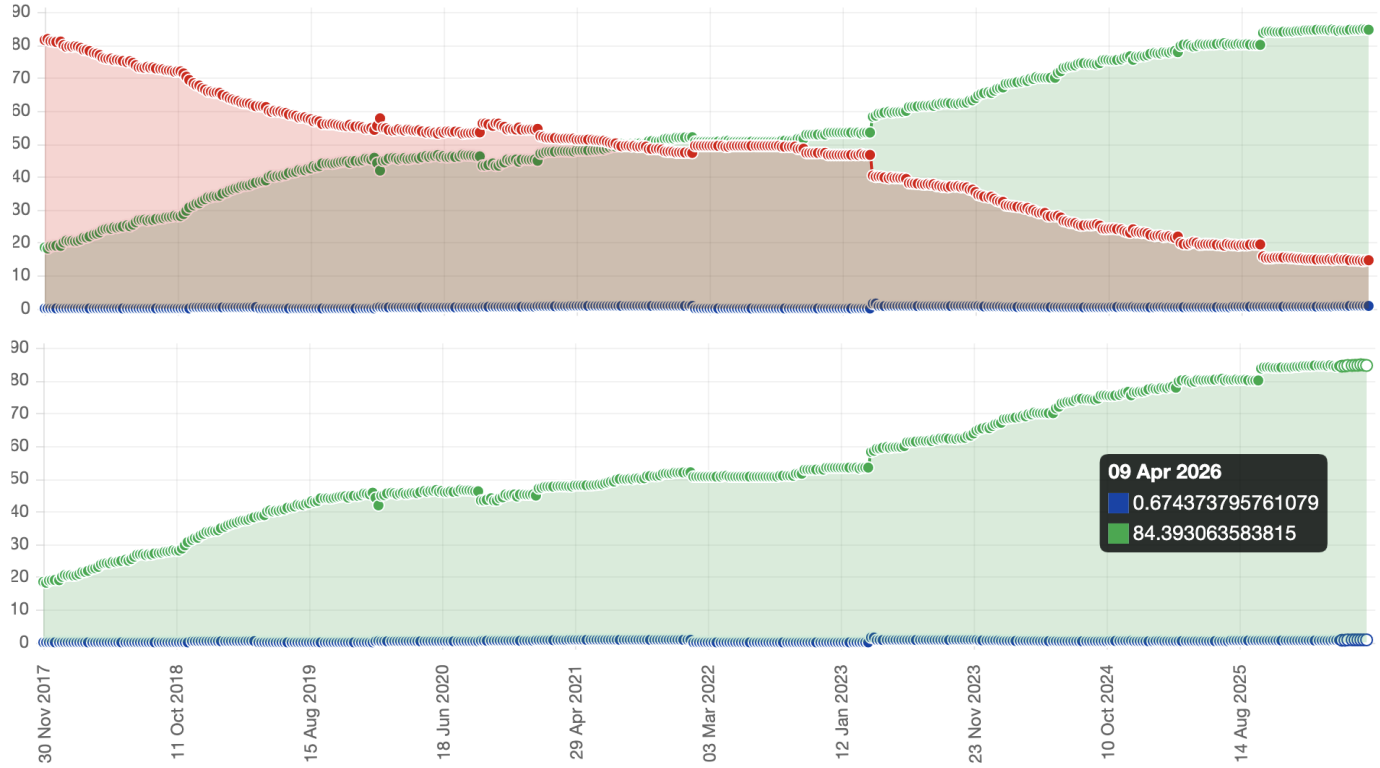

- Percentage of IPv6-only endpoints
- Percentage of dual-stack endpoints
- Percentage of IPv4-only endpoints

These graphs record, on a weekly basis, the fraction of service endpoints listed in the VO Feeds of the 4 major LHC experiments (Alice, Atlas, CMS, LHC-B)
=> Currently: **85% is dual stack**

Non-IPv6 traffic among French Sites