IPv6 & WLCG

update from the IPv6 working group

David Kelsey (STFC, UKRI), Andrea Sciabà (CERN)
on behalf of the HEPiX IPv6 Working Group

WLCG Management Board meeting
27 July 2021

Dual-stack IPv4/IPv6 - reminder

- WLCG MB approved plan for Dual-stack IPv4/IPv6 storage by end LHC Run 2
- Main driver was and still is: “support IPv6-only CPU”
  - Ongoing concern of lack of IPv4 addresses (at some sites)
- Other (more recent) drivers
  - Support packet marking in TCP for monitoring by RNTWG
    - Research Networking Technical Working Group
  - US Federal Government - 80% of services to be IPv6-only by 2025-26
    - come back to this

IPv4/IPv6 and Tier-1 storage

- FTS servers at BNL, CERN, and RAL work in dual stack, while Fermilab is still IPv4-only
  - CMS/Fermilab now use the CERN FTS server, so IPv6 transfers are being used
- FTS transfers do happen via IPv6 at most Tier-1s (but not all transfers)
- IPv6 transfers do not happen at RRC-KI

Experiment	Fraction of T1 storage accessible via IPv6
ALICE	78%
ATLAS	96%
CMS	100%
LHCb	94%
Overall	96%

IPv4/IPv6 deployment at Tier-2 sites

- The deployment campaign was launched in November 2017
- Steady progress (status)
  - About 79% of Tier-2 sites have storage on dual stack

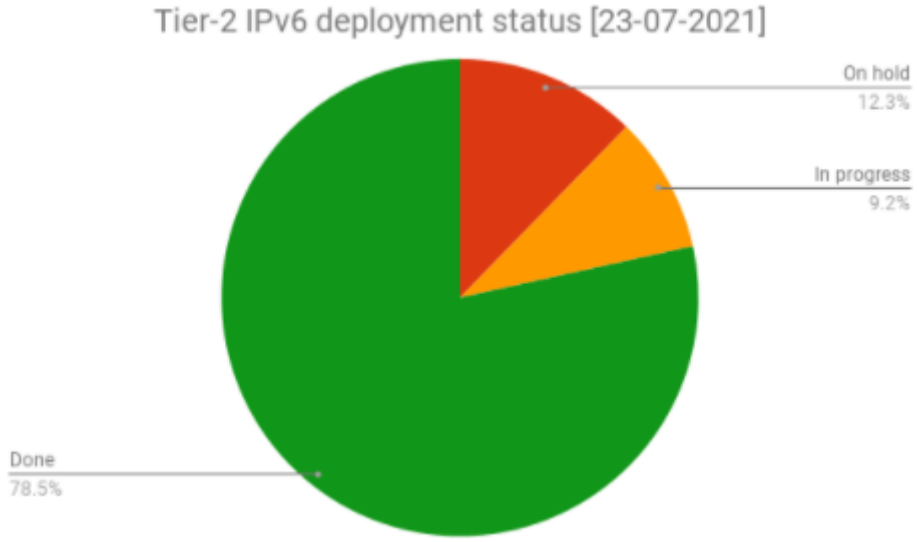

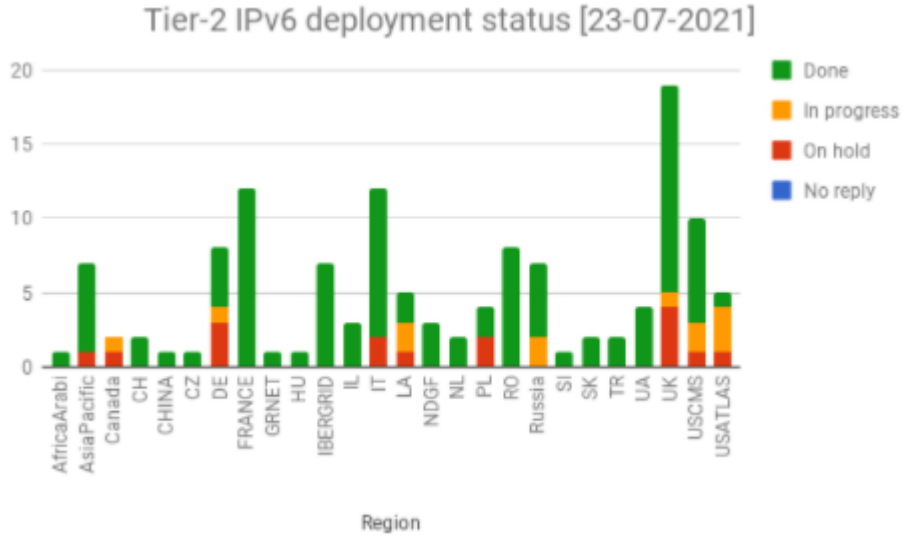

Experiment	Fraction of T2 storage accessible via IPv6
ALICE	86%
ATLAS	62%
CMS	94%
LHCb	75%
Overall	77%

Tier-2 status (cont'd)

- Progress is slower and slower
  - The cause is that remaining sites are the ones having the most difficulties
    - Typical example: IPv6 not a priority for the campus

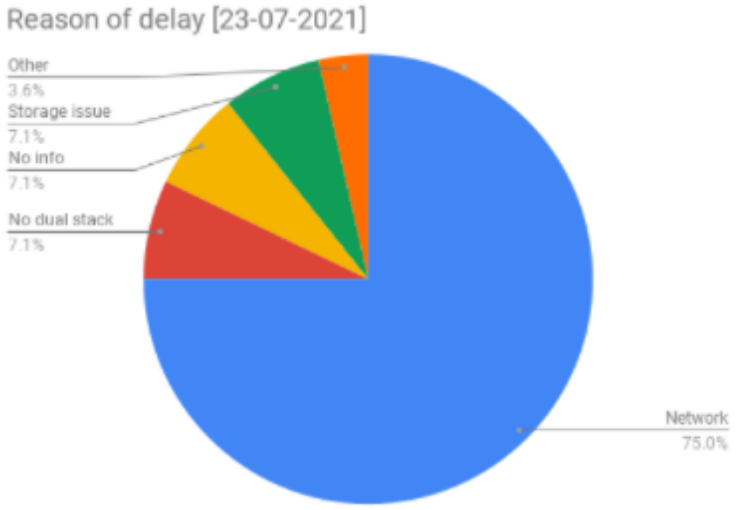

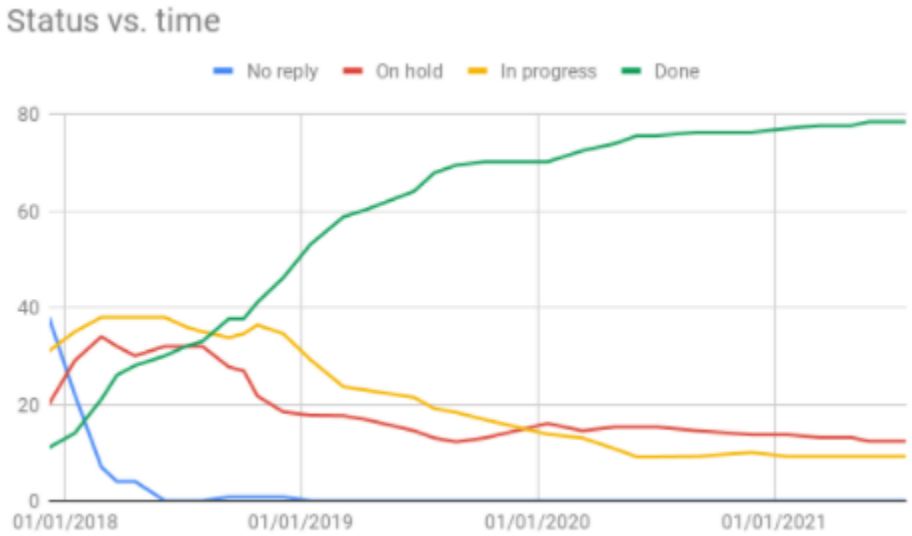

Plans are in place for US ATLAS to complete dual-stack storage in coming months. We will chase sites in UK

Fraction of SRM/GridFTP FTS traffic over IPv6

- Some FTS protocols, e.g. DAVS, not yet instrumented to monitor IPv6
- these are excluded from this plot

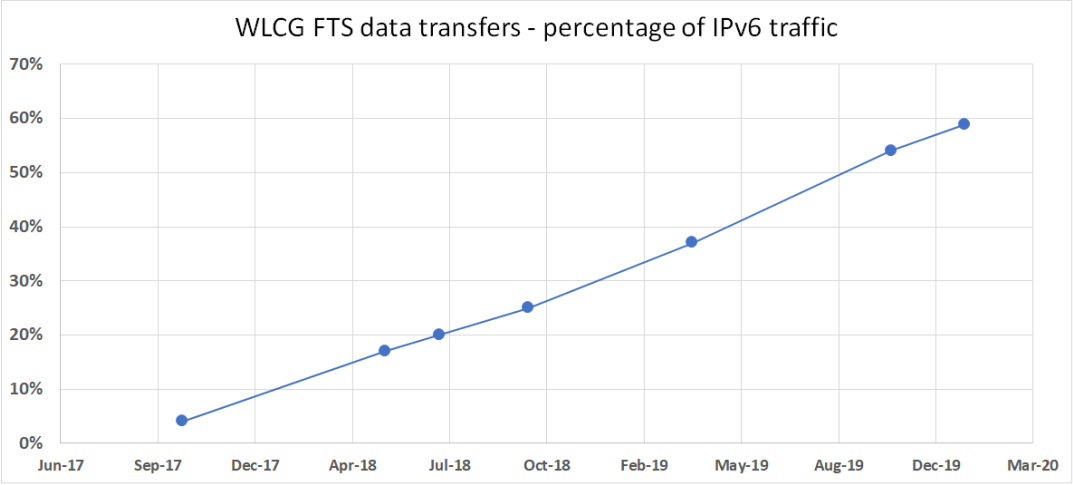

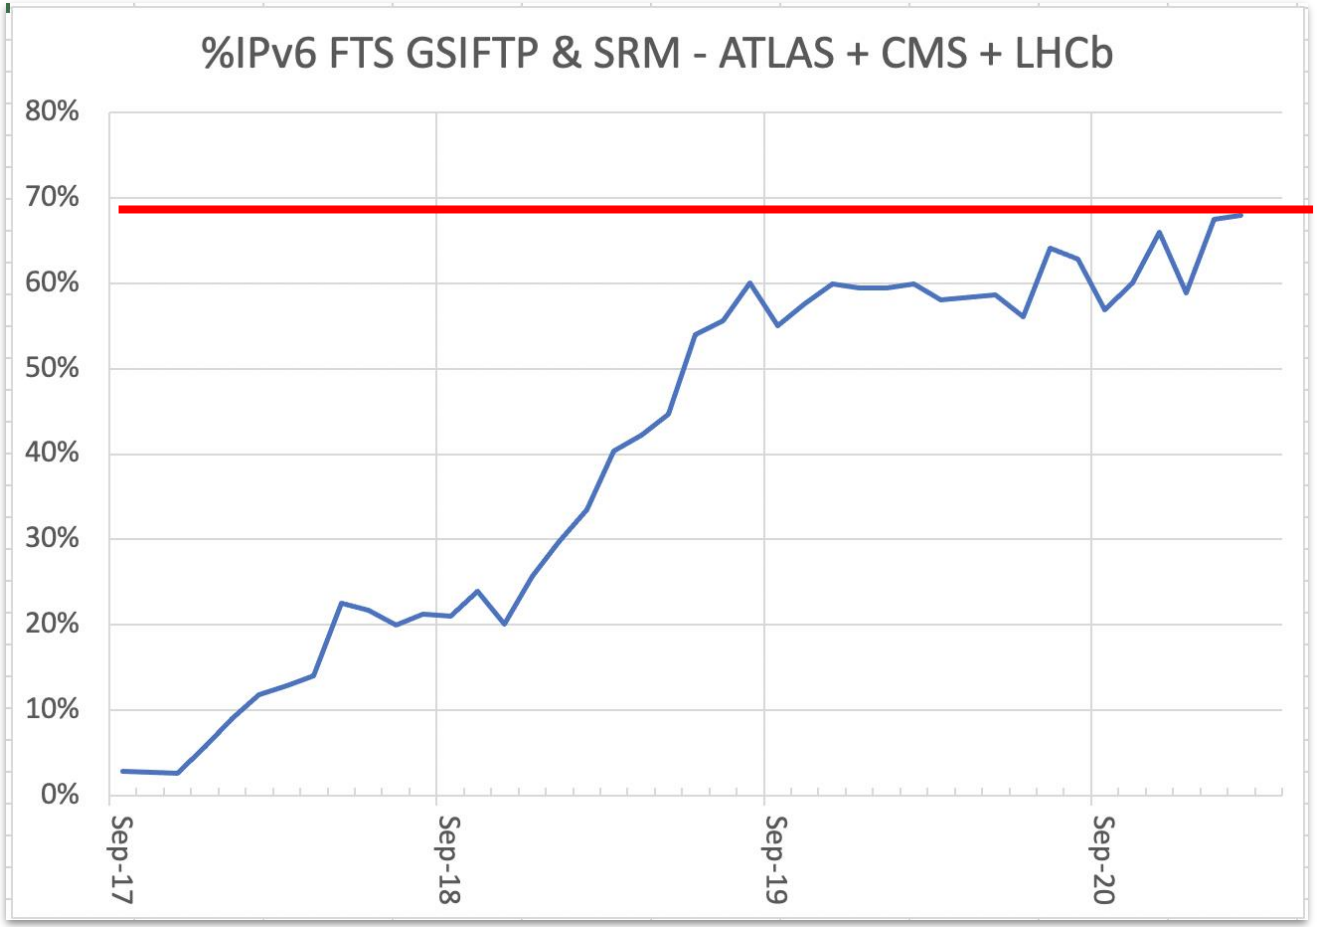

ATLAS and CMS FTS IPv6 traffic - all protocols

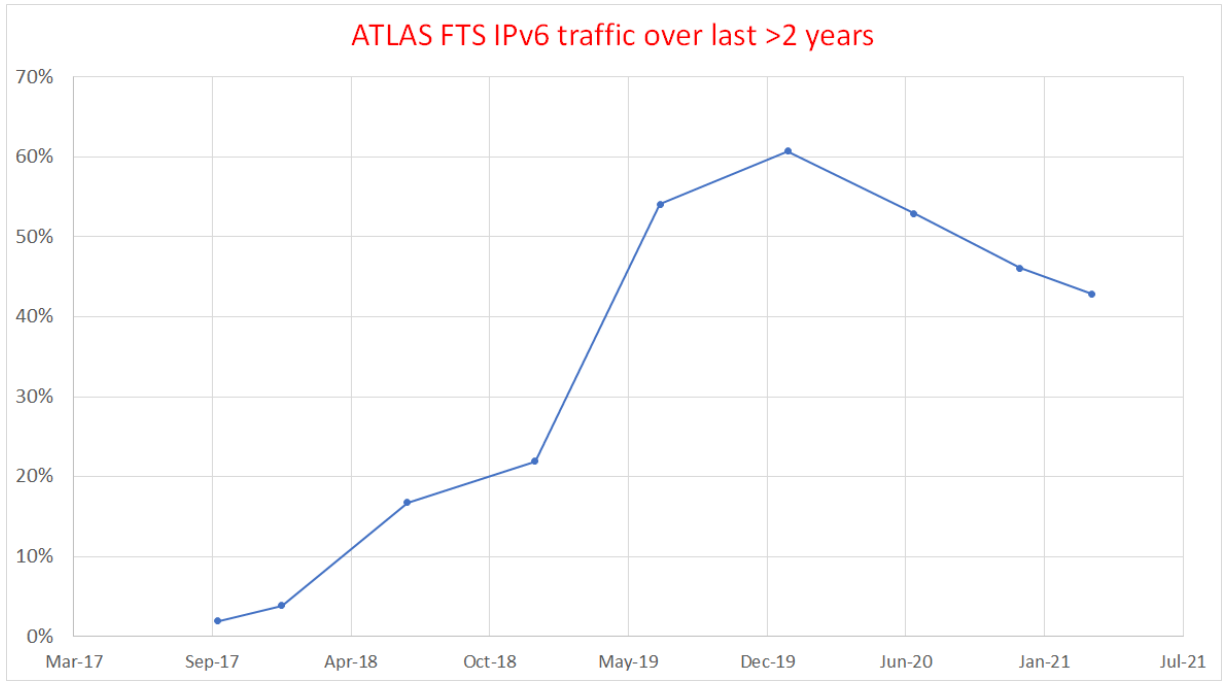

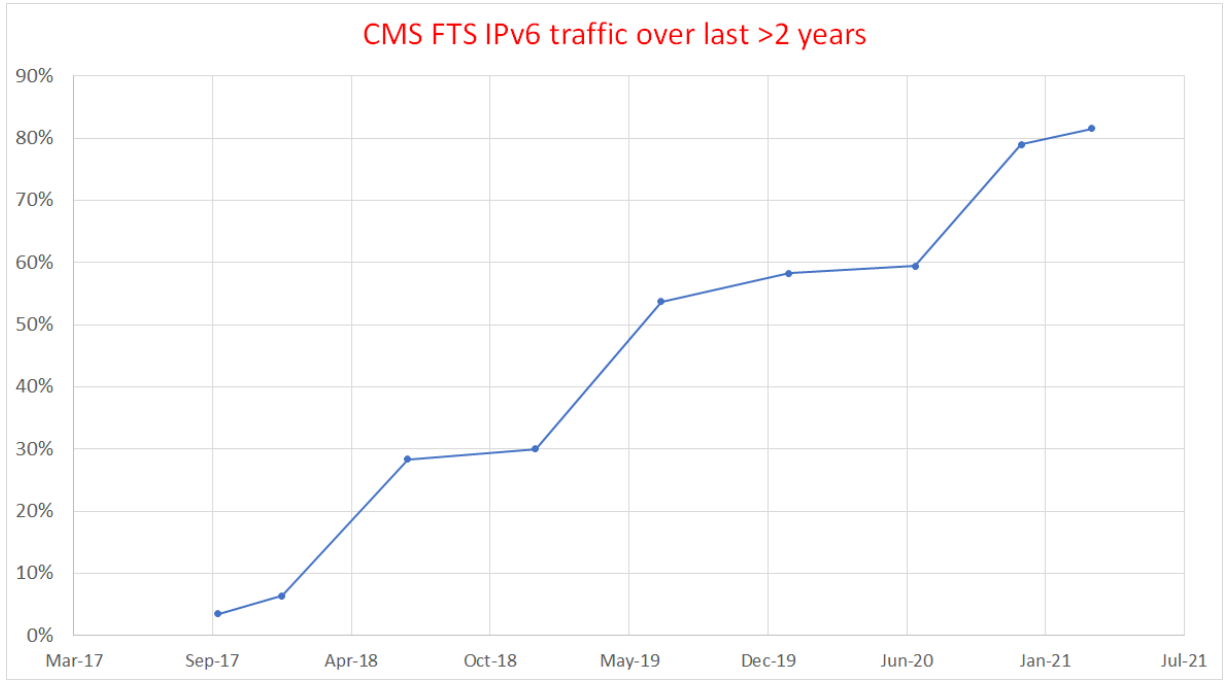

IPv6 traffic on LHCOPN/LHCONE at CERN

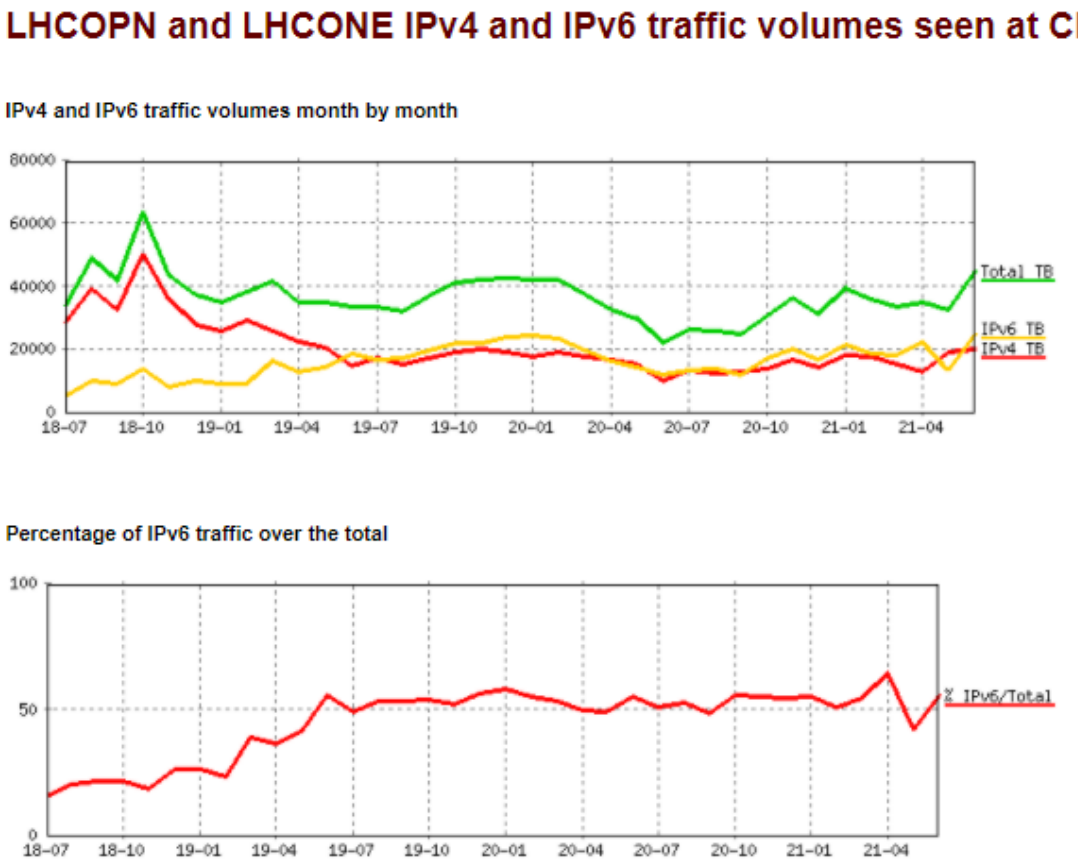

IPv6 traffic on LHCOPN/ONE as seen at CERN

- ~50% of all traffic is IPv6
- From June 2019 onwards
- Why so small? (LHCONE?)

LINK to these plots

IPv4 transfers on LHCOPN/ONE?

- Tier1s are dual-stack, but IPv4 can still be used for transfers
  - site issues - some storage end-points not yet dual-stack
  - both ends are dual-stack but configuration or request prefers IPv4
  - transfers are to/from WN's and the WNs are still IPv4-only
- encourage all sites to deploy all WNs and all services as dual-stack
- encourage all sites and all experiments to "prefer" IPv6 transfers
- Some recent examples of IPv4 traffic included networks:
  - CERN, INFN-CNAF, JINR, FR-IPBS, DESY, BCNET, NARACOM

DOE Labs and US IPv6-only directive

https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf

- They must develop plans by end of FY2021, for
- At least (20%, 50%, 80%) of IP-enabled assets on Federal networks are operating in IPv6-only environments by the end of FY (2023, 2024, 2025)

How will this affect WLCG?

- FTS transfers won't be impacted
  - all storage in WLCG MUST be at least dual stack
- IPv4-only worker nodes outside of the US will not be able to remotely read data in the USA
  - encourage deployment now of dual-stack WN's
  - Or replicate to a storage with IPv4
  - If US sites can keep IPv4 just for the storage (as an exception) the impact would be negligible

From dual-stack to IPv6-only

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

- During LS2 - planning for an IPv6-only WLCG
- To simplify operations
  - dual-stack infrastructure is the most complex!
- Large infrastructures (e.g. Facebook, EE/BT) use IPv6-only internally
- The plan - the goal we are working towards
  - IPv6-only for the majority of WLCG services and clients
  - With ongoing support for IPv4-only clients where needed
    - e..g. via use of RFC 7755/7756
      - Stateless IP/ICMP Translation for IPv6-only Data Centres
- Timetable to be defined

Preparing for an IPv6-only WLCG

- Need to test IPv6-only clusters/worker nodes
  - IPv6-only testing at CERN (Ben Jones,Luis Fernandez Alvarez)
- FTS monitoring changes will happen
  - to monitor IPv6 share on all FTS protocols
- INFN analysing sampled WN traffic to identify protocols and services using IPv4 beyond those we already know
  - Applications that store, handle or signal literal IPv4 addresses will be broken by protocol translation
  - Then work with experiments and developers to “fix” whatever is broken
- Need to test tools to allow IPv4-only clients (but who can do this?)

IPv6-only timetable?

- WLCG can already support use of IPv6-only CPU resources
- During 2021-22
  - Sites
    - deploy dual-stack WN’s and other services
    - check that configurations “prefer” IPv6
    - investigate data transfers that happen over IPv4
  - Experiments
    - check data transfer configurations - always prefer IPv6

IPv6-only timetable (2)

- First place to turn-off IPv4?
  - LHCOPN (Tier0 and all Tier1's are dual-stack)
  - When the amount of IPv4 traffic becomes very small
    - make LHCOPN IPv6-only
    - IPv4 traffic will flow via general internet links
- **WLCG IPv6-only plan**
  - MB to agree date for end of "full support" for IPv4 on WLCG Clients
  - Core services at that point can become IPv6-only
  - during LS3 and before start of Run 4?

Questions, Discussion?