Update from the HEPiX IPv6 working group

David Kelsey (STFC, UK Research and Innovation)
on behalf of the HEPiX IPv6 Working Group

WLCG GDB
14 Oct 2020

On behalf of all members of the HEPiX IPv6 working group

Recently active in the HEPiX IPv6 Working Group

- M Babik (CERN), M Bly (RAL), T Chown (Jisc), D Christidis (U Texas/ATLAS), J Chudoba (Prague), C Condurache (EGI.eu), C Grigoras (CERN/ALICE), B Hoeft (KIT), D P Kelsey (RAL), E Martelli (CERN), S McKee (U Mich), R Nandakumar (RAL/LHCb), K Ohrenberg (DESY), F Prelz (INFN), D Rand (Imperial), A Sciabà (CERN/CMS), D Stockland (Imperial)
- Many more in the past and others join from time to time
- *and thanks also to WLCG operations, WLCG sites, LHC experiments, networking teams, monitoring groups, storage developers…*

Outline

- WLCG transition to dual-stack IPv4/IPv6 storage
  - Tier-1/Tier-2 status
- FTS data transfers – fraction over IPv6
- Monitoring (ETF and perfSONAR)
- IPv6-only networking
  - Why? How?
- Summary

For Working Group meetings see

https://indico.cern.ch/category/3538/

Meetings of working group (2020):
16-17 Jan (CERN – F2F)
3 June (half-day virtual meeting)
29-30 Sep (two half-days)

One-hour update meetings:
10 Mar, 30 Apr, 2 Jul, 3 Sep

Upcoming virtual meetings:
19-20 Jan 2021 (2 half-days - assumed virtual)
One-hour updates: 22 Oct, 26 Nov 2020

IPv6 and Tier-1 storage

- FTS servers at BNL, CERN and RAL work in dual stack, while Fermilab is still IPv4-only
  - But Fermilab T1 storage is dual-stack
- GridFTP transfers happen also via IPv6 at most Tier-1s
- IPv6 transfers do not happen at
  - RRC-KI
  - Fermilab FTS IPv6 transfers are still not happening

- Fraction of Tier-1 disk storage on IPv6 *(not recently updated)*
  - ALICE: 78%
  - ATLAS: 96%
  - CMS: 100%
  - LHCb: 94%
  - All VOs: 96%

IPv6 traffic on LHCOPN & LHCONE at CERN

TWiki > LHCOPN Web > LHCOPNEv4v6Traffic (2020-09-08, EdoardoMARTELLI)

LHCOPN and LHCONE IPv4 and IPv6 traffic volumes seen at CERN Tier0

IPv4 and IPv6 traffic volumes month by month

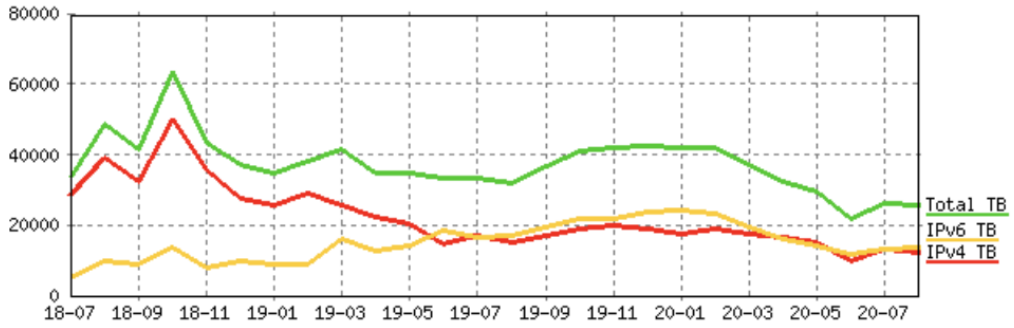

Percentage of IPv6 traffic over the total

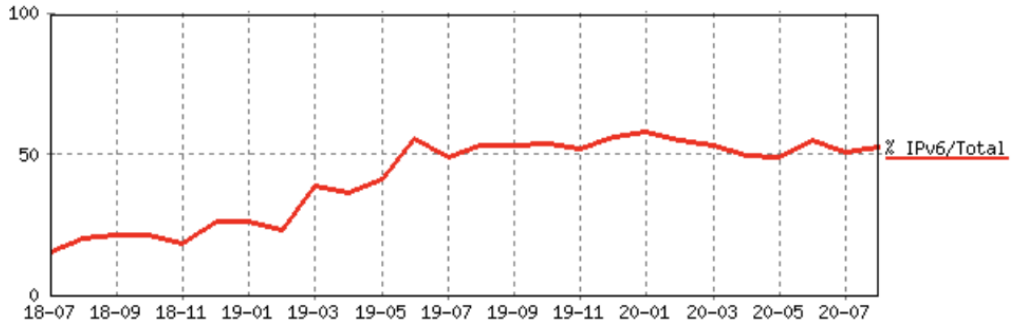

IPv6 traffic on LHCOPN & LHCONE as seen at CERN

- ~50% of all traffic is IPv6
- From June 2019 onwards

LINK to these plots

Tier-2s: GGUS tickets to all Tier-2 sites

- The deployment campaign was launched in November 2017
- Steady progress (status)
  - About 76% of Tier-2 sites have storage on dual stack

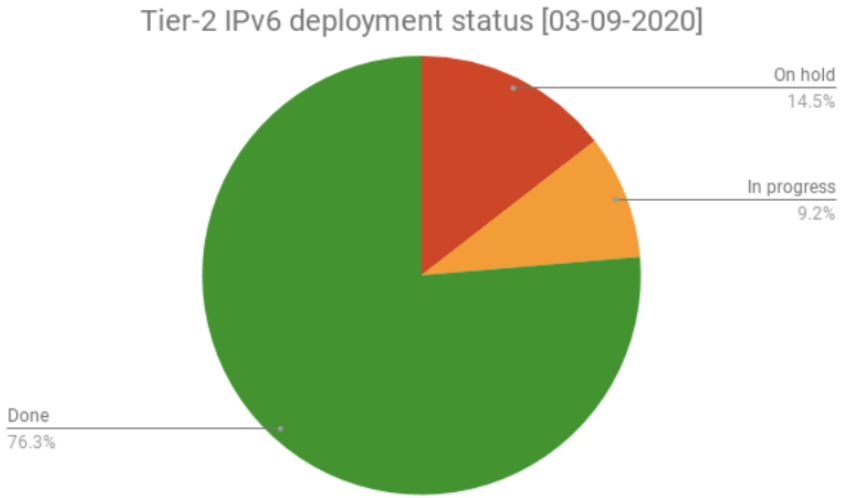

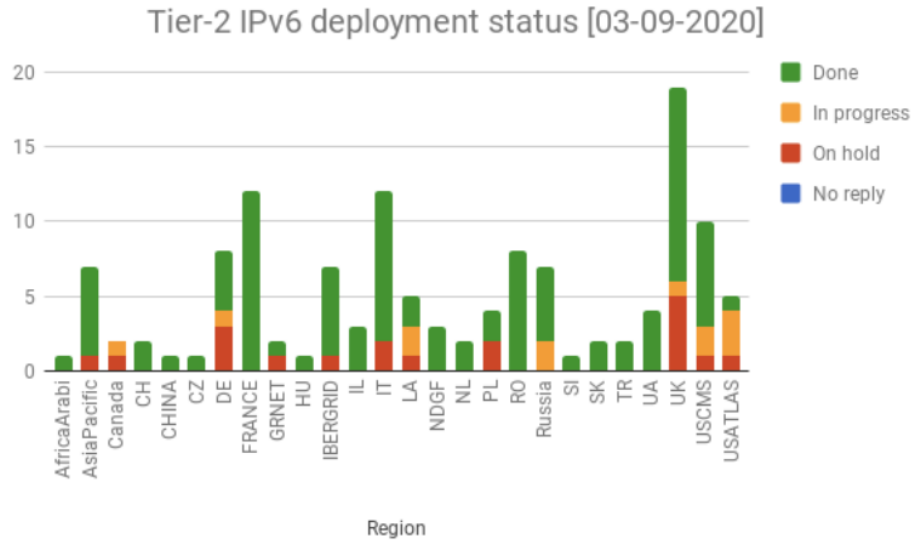

Tier-2 status (cont'd)

Experiment	Fraction of T2 storage accessible via IPv6
ALICE	86%
ATLAS	62%
CMS	93%
LHCb	75%
Overall	76%

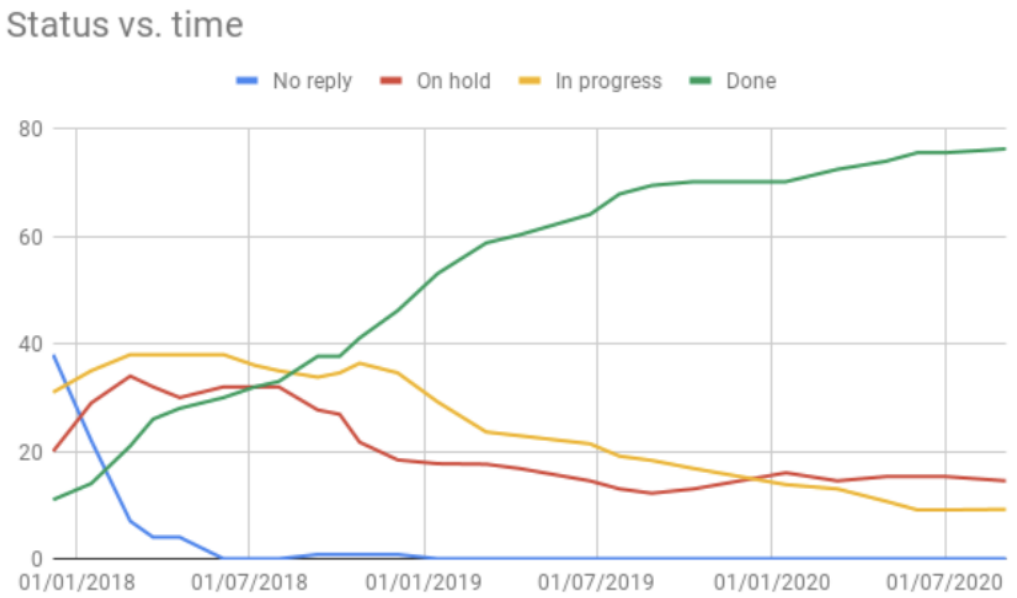

Data Transfers - FTS transfer monitoring - August 2020

Approximately 59% of data transferred via FTS in Aug 2020 went over IPv6

https://monit-grafana.cern.ch/

% of FTS traffic over IPv6 - last >2 years

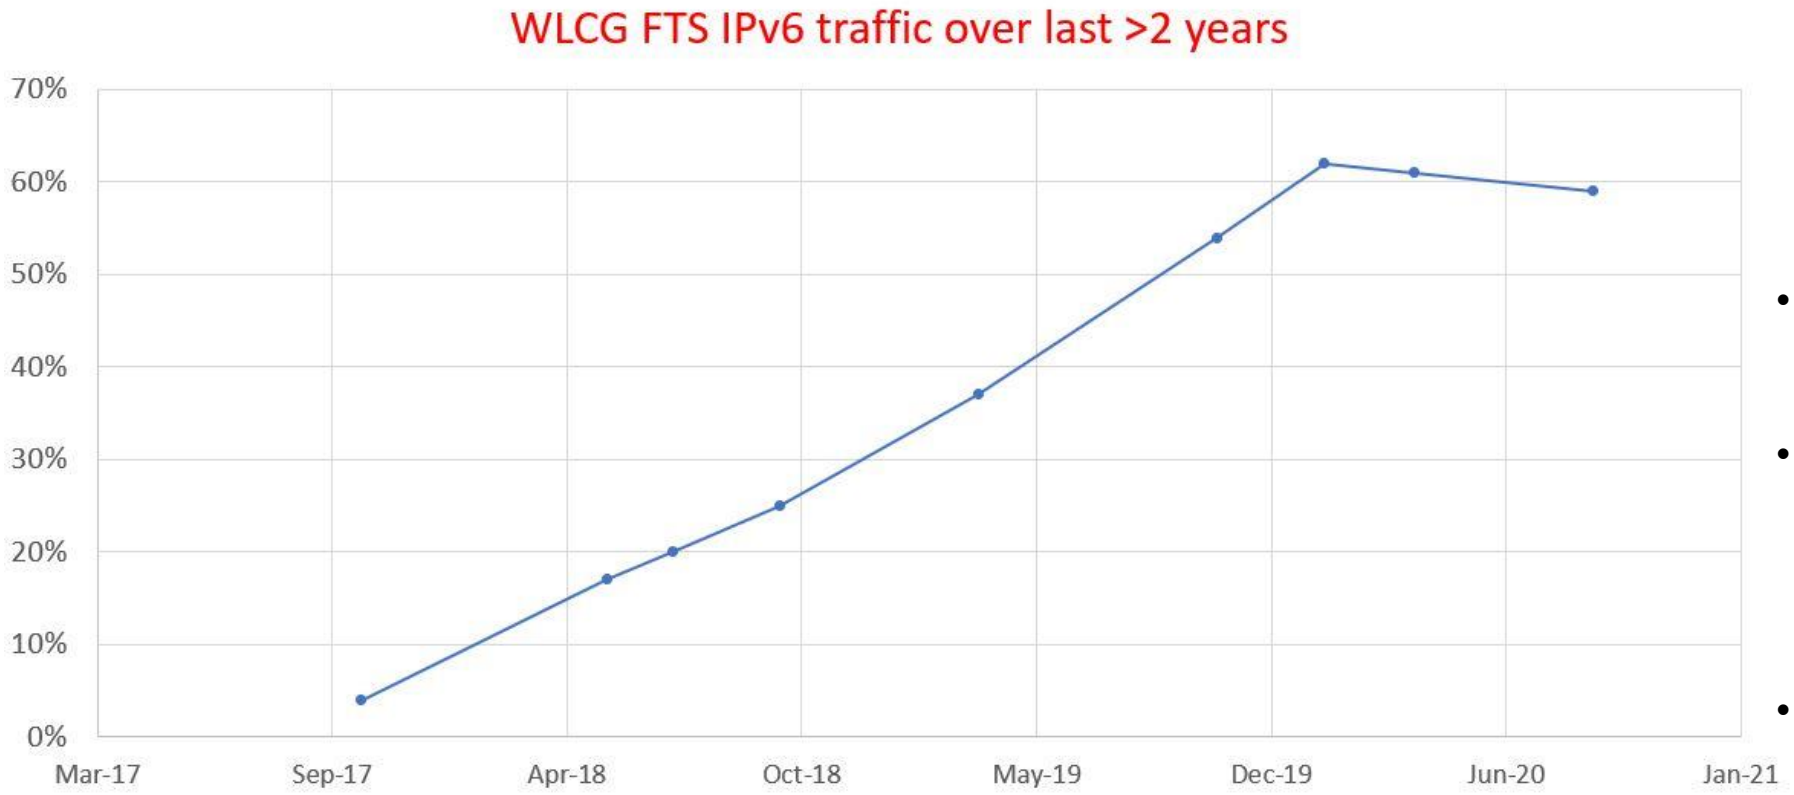

- IPv6 %traffic no longer increasing during 2020
- Since April 2020, total traffic reduced and mix of CMS to ATLAS traffic changed
- all contribute to small downward trend in 2020

Monitoring & IPv6

- Important tools for tracking/troubleshooting IPv6
- See presentations by Marian Babik at IPv6 WG 30 Sep 20

- Experiment Test Framework ATLAS, CMS and LHCb now have IPv6-only ETF instances
  - running both IPv4-only and IPv6-only instances
  - Able to generate WLCG reports with IPv4/IPv6 results

- perfSONAR
  - IPv6 measurements included in all meshes

IPv6-only networking

- Our main use case for dual-stack storage was and still is:
  - Be ready for use of (opportunistic) IPv6-only CPU
- BUT there is a new requirement for data transfers over IPv6
  - Talk by Shawn McKee on 29th Sep 20 (IPv6 working group meeting)
- Research Networking Technical Working Group (RNTWG)
  - The ability to understand WAN network flows is too limited
    - Preparing for HL-LHC
  - Need new methods to mark and monitor network use
  - label traffic at the packet level to indicate experiment and activity
    - better understand use of the network – data in flight, not just end points
  - Packet Marking sub-group started work
  - Decided to focus on IPv6 marking (and backport to IPv4 if possible)

IPv6-only networking (2)

Important driver for IPv6-only

- Running a dual-stack IPv4/IPv6 infrastructure is complex
- Dual-stack everywhere – not desirable as the “end of transition”
- Large companies (e.g. Facebook, EE/BT) use IPv6-only internally
  - Then use tools like NAT64/DNS64/464XLAT for legacy world
- CERN EOS infrastructure uses IPv6-only internally

- Our proposed plan is to simplify and move to IPv6-only in the majority of WLCG services and clients
  - Ongoing support for IPv4-only clients where needed (via use of NAT64?)

Work to achieve IPv6-only networking on WLCG

- Fix dual-stack endpoints that prefer to use IPv4 rather than IPv6
- More testing of IPv6-only clusters
- Encourage more sites to use dual-stack CPU
  - Many sites have successfully moved to dual-stack worker nodes
  - IPv6 CPU will more naturally transfer data over IPv6

- WLCG may need to agree a date for “end of full support” of IPv4-only clients
  - e.g. before the start of LHC Run4?
  - Transition tools such as NAT64/DNS64 can be used once core is IPv6-only

Summary

- WLCG is ready to support use of IPv6-only CPU resources
  - Tier-1: 96% of storage is available via IPv6
  - Tier-2: 76% sites & storage are now dual-stack
- ~60% of FTS transfers today over IPv6
- ~50% LHCOPN+LHCONE traffic observed at CERN is IPv6
- But growth currently “on hold”
- IPv6 WG is working on the move to IPv6-only services
  - Much simpler to manage
  - IPv6 data transfers to help netflow packet marking
- ***Many thanks to all who contribute to this work – volunteers to join the working group are always welcome!***

Questions?