


IPv6 update

Andrea Sciabà
on behalf of the HEPiX IPv6 WG

WLCG GDB, 16-01-2019




IPv6 deployment on WLCG









- The main activity of the HEPiX IPv6 working group and the WLCG ops coordination IPv6 task force in the last year is the coordination and support for the deployment of IPv6 at WLCG sites
- The stated goal is to allow data on federated storage to be accessible by jobs on IPv6-only connected CPUs
- Short summary of the timeline
 - Tier-1: deployment of dual-stack on production storage, CVMFS and FTS by **April 2018**
 - Tier-2: deployment of dual-stack on production storage (and perfSonar if installed) by **end of Run2** (i.e. end of 2018)
- Last GDB update in September

IPv6 CERN traffic

- IPv6 non-LHC traffic increased by 47% (to CERN) and by 17% (to Internet) since September
- Information on IPv6 traffic on LHCOPN/ONE available only for some Tier-1's
 - See next slide

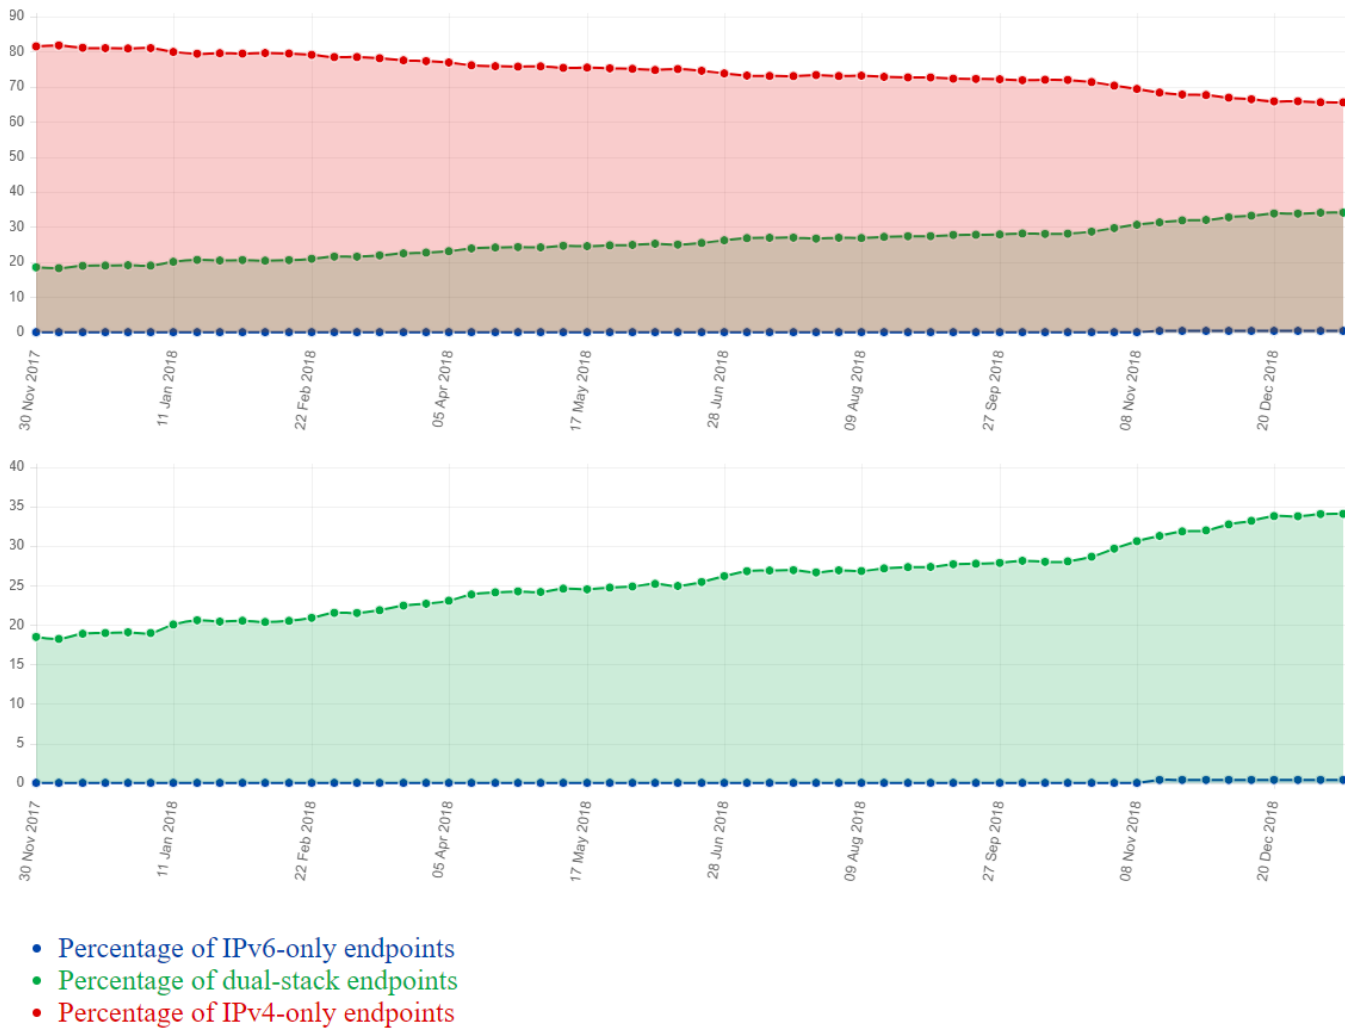


IPv6 traffic to selected T1's



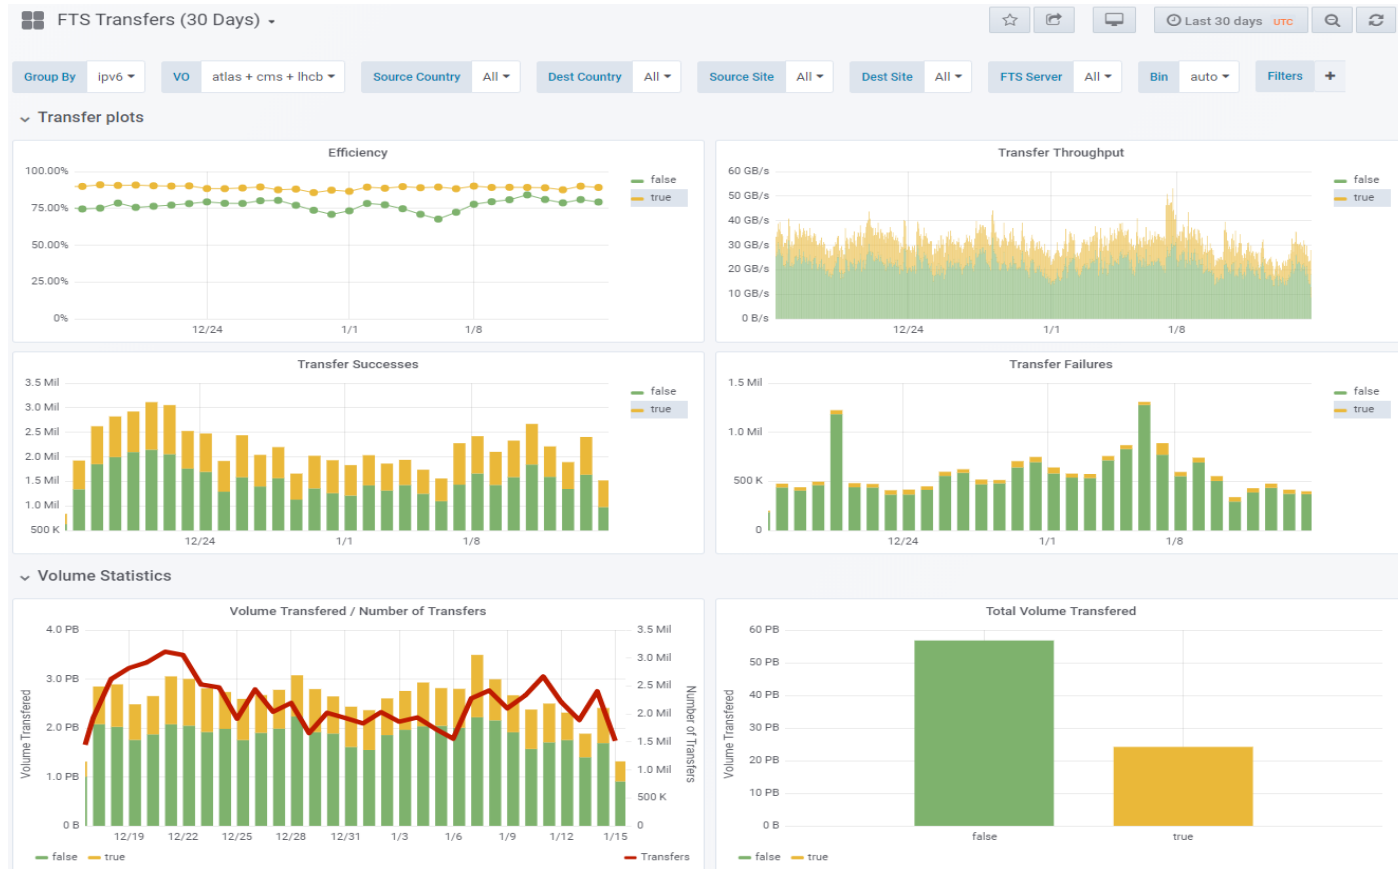
Link	IPv4	IPv6
KIT→CERN	60%	40%
CERN→KIT	64%	36%
IN2P3→CERN	58%	42%
CERN→IN2P3	64%	36%
NL-T1→CERN	31%	69%
CERN→NL-T1	48%	52%
NDGF→CERN	70%	30%
CERN→NDGF	11%	89%

LHC experiment endpoints on IPv6



- Steady increase in IPv6 addresses (+25% since September)
 - https://orsono.mi.infn.it/~prelz/ipv6_vofeed/

FTS GridFTP traffic





- IPv6 GridFTP traffic via FTS has increased further
 - From 24% of the total in September to 30% now
- Xrootd IPv6 traffic not yet monitored

Status of IPv6 at Tier-1's


- All T1's route IPv6 traffic on LHCOPN (LHCONE TBC)
- All T1's have an IPv6-enabled pS but for RRC-KI
- FTS servers at CERN and BNL work in dual stack, while at FNAL it still uses IPv4 only
 - Previously BNL FTS was defaulting to IPv4, but not any more
- IPv6 deployment completed at most T1's
 - IN2P3, CNAF, RAL, SARA-MATRIX, PIC, JINR, NDGF, NIKHEF, ASGC, BNL
- Still not 100% finished, but on a good track at
 - RAL: storage on IPv6 for ATLAS and CMS. LHCb to be done in Jan-Feb, ALICE TBD
 - KIT: storage is dual stack for all VOs but ALICE, which will be done this month
 - TRIUMF?
- Not yet at
 - FNAL: site-wise IPv6 is enabled already; deployment on storage will start in January and take 3-5 months
 - RRC-KI?

IPv6 at Tier-2 sites

- The deployment campaign was launched in November 2017
 - GGUS tickets sent to all non-US sites
 - Sites made aware of the WLCG plans and asked to report plans and give updates
- Steady progress ([status](#))
 - About 53% of T2 sites have storage on dual stack




Experiment	Fraction of T2 storage accessible via IPv6
ALICE	58%
ATLAS	43%
CMS	75%
LHCb	62%
Overall	57%




Reasons of no deployment

- Checked why sites are late in their IPv6 deployment and the expected ETA
 - **Network**: waiting for interventions on the site infrastructure or network issues to be fixed
 - **No dual stack**: infrastructure network is fine but IPv6 needs enabling on storage or perfSONAR
 - **pS/storage issues**: encountered problems that need fixing
 - **Other**: other types of problems (e.g. the site needs relocating)
 - **No info**: no meaningful information given by the site on their plans

Time scale for deployment



Reason of delay



T2 deployment observations

- Tentative extrapolation to the future
 - Within days: 60%
 - Within weeks: 75%
 - Within months: 90%
- Networking issues of interventions are by far the largest cause of delay
 - Most often, not under the control of the WLCG site staff
- Most sites are responsive and provide detailed information
 - For some however regular pinging is essential

News on IPv6 monitoring

- The IPv6 dual-stack mesh in the perfSonar dashboard to be decommissioned very soon
- IPv6 pS tests now included in all meshes
 - Throughput and traceroute tests always for both IPv4 and IPv6, latency tests for either depending on what's available on both ends
- pS monitoring now tracking the “efficiency” for both IPv4 and IPv6
 - Efficiency = number of working destinations / total destinations
 - https://psetf.aglt2.org/etf/check_mk/index.py
- CC7 campaign is ongoing, the network support team will start issuing tickets soon
- Progress with Maddash, but still some open issues

Conclusions

- IPv6 traffic and IPv6-enabled endpoints are ever increasing
- A few Tier-1 sites are lagging behind the deadline of last April
- Deployment Tier-2 is proceeding well enough
 - An acceleration was foreseen as the deadline approached and it did happen!
 - Even if far from the nominal target, an analysis of the tickets leads to the expectation that 90% of the T2's will have their storage in dual stack “within months”
 - On the bright side, CMS has already 75% of the T2 storage in dual stack
 - Differences among experiments are likely correlated to how much additional “pressure” they put on their sites, on top of the WLCG effort
- **Deadline extension?**
 - IMO, no reason to, just keep the pressure on the straggling sites, the sooner the better