



IPv6-only networking on WLCG

Marian Babik (CERN), Bruno Hoefft (KIT), David Kelsey (STFC UKRI),
Edoardo Martelli (CERN), Francesco Prelz (INFN),
Duncan Rand (Imperial), Andrea Sciabà (CERN) et al

CHEP2019, Adelaide, 7 Nov 2019

On behalf of all co-authors in the HEPiX IPv6 working group

Active in HEPiX IPv6 Working Group – last 12 months

- M Babik (CERN), M Bly (RAL), T Chown (Jisc), D Christidis (U Texas/ATLAS), J Chudoba (Prague), C Condurache (RAL/EGL.eu), T Finnern (DESY), C Grigoras (CERN/ALICE), B Hoeft (KIT), D P Kelsey (RAL), R Lopes (Brunel), F López Muñoz (PIC), E Martelli (CERN), A Manzi (CERN), R Nandakumar (RAL/LHCb), K Ohrenberg (DESY), F Prelz (INFN), D Rand (Imperial), A Sciabà (CERN/CMS)
- 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-0/Tier-1/LHCOPN/LHCONE status
 - Tier-2 status
- Data transfers
- Monitoring
- IPv6-only networking
- Summary

Network and pS at Tier-1s – IPv6 status

Bruno Hoeft

- TW-ASGC Perfsonar server is now running
- RRC-KI-T1 is connected to with IPv6 to LHC[OPN/ONE], the Perfsonar server not running

improved


| Tier-1 | LHCOPN | LHCONE | IPv6 Perfsonar |
|---------------------|--------|--------|----------------|
| CA-TRIUMF | OK | OK | LHC[OPN/ONE] |
| CH-CERN (Tier-0) | OK | OK | LHC[OPN/ONE] |
| DE-KIT | OK | OK | LHC[OPN/ONE] |
| ES-PIC | OK | OK | LHC[OPN/ONE] |
| FR-CCIN2P3 | OK | OK | LHC[OPN/ONE] |
| IT-INFN-CNAF | OK | OK | LHC[OPN/ONE] |
| KR-KISTI | OK | OK | LHC[OPN/ONE] |
| NGDF | OK | OK | LHC[OPN/ONE] |
| NL-T1 - NIKHEF | OK | OK | -- |
| NL-T1 - Sara-Matrix | OK | OK | LHC[OPN/ONE] |
| RRC-KI-T1 | OK | OK | LHC[OPN] |
| RRCC-JINR-T1 | OK | OK | LHC[OPN/ONE] |
| TW-ASGC | OK | OK | -- |
| UK-T1-RAL | OK | OK | LHC[OPN] |
| US-T1-BNL | OK | OK | LHC[OPN/ONE] |
| US-T1-FNAL | OK | OK | LHC[OPN/ONE] |

IPv6 traffic on LHCOPN & LHCONE at CERN


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

Edoardo Martelli

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



IPv6 on FTS and at Tier-1s

- FTS servers at BNL, CERN and RAL work in dual stack, while FNAL has still problems
- GridFTP transfers happen also via IPv6 at most Tier-1s
- IPv6 transfers do not happen at
 - RRC-KI
 - FNAL FTS IPv6 transfers are still problematic, but should be fixed soon
- **Fraction of Tier-1 disk storage on IPv6:**
 - ALICE: 78%
 - ATLAS: 96%
 - CMS: 100%
 - LHCb: 94%
 - **All VOs: 96%**

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

Andrea Sciaba


- WLCG set a target for end 2018 for deploying IPv6 on storage systems (and perfSONAR)
- The deployment campaign was launched in November 2017
- Steady progress ([status](#))
 - **About 70%** of T2 sites have storage on dual stack



Tier-2 status (cont'd)


| Experiment | Fraction of T2 storage accessible via IPv6 |
|------------|--|
| ALICE | 85% |
| ATLAS | 59% |
| CMS | 89% |
| LHCb | 75% |
| Overall | 73% |

Status vs. time



FTS transfer monitoring - last 30 days


Approximately **57%** of data transferred via FTS in the last 30 days went over IPv6



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

% of FTS traffic over IPv6 - last 2 years

WLCG FTS IPv6 traffic over last 2 years



Data points:
Reports at HEPiX
meetings and
CHEP18/CHEP19

Experiments Test Framework (ETF)

Marian Babik



News (IPv6)

- Experiment instances are ready
 - ATLAS, CMS and LHCb now have IPv6-only instances
 - **SAM tests now running from both IPv4-only and IPv6-only instances**
 - **Experiments can now create SAM profiles that mix IPv4 and IPv6 results**
 - Ready to generate WLCG reports with IPv4/IPv6 results

perfSONAR IPv6 Mesh


Marian Babik & Duncan Rand

OPN Mesh Config - OPN IPv6 Bandwidth - Throughput



! Found a total of 4 problems involving 3 hosts in the grid

[LINK](#) to dashboard



- Re-configured meshes
 - IPv6 pS tests now included in all meshes
- New monitoring is in place which reports test “efficiency” wrt IPv4 and IPv6
 - Efficiency = number of working destinations / total destinations

IPv6-only networking

- Our main use case - ready for use of (opportunistic) IPv6-only CPU
- **BUT there are other drivers for IPv6-only:**
 - a) new CERN machine room and lack of public IPv4 addresses
 - Use only IPv6 addresses for external public networking?
 - b) multiONE (different communities using LHCONE)
 - multiple overlay networks
 - different addresses for each community
 - sites likely have lack of sufficient IPv4 address space

IPv6-only networking (2)

- Running a dual-stack IPv4/IPv6 infrastructure is **complex**
- Large companies (e.g. Facebook, EE/BT) use IPv6-only internally
 - Then use tools like NAT64/DNS64/464XLAT for legacy world
- CERN EOS infrastructure also uses IPv6-only internally

When/how do we **simplify** and move to IPv6-only in WLCG?
IPv6 working group Phase 3

- The fraction of data transfers on IPv6 is getting much larger (>50%)
- When the amount of IPv4 traffic on LHCOPN is close to zero
 - Turn off IPv4 entirely on LHCOPN?
 - simplify routing tables and tracking problems is easier
- MultiONE/LHCONE may also be using IPv6-only

IPv6-only networking on WLCG?

- Fixing dual-stack endpoints that prefer to use IPv4 and not IPv6
- Turning off IPv4 at a Site is clearly their own decision
 - they may have many other needs (not WLCG) for IPv4
- We need to include experience of sites already doing IPv6-only CPU
 - UKI-LT2-Brunel - successful IPv6-only cluster for LHCb, ATLAS, CMS & LSST
 - Also SiGNET, T2_US_Nebraska, UKI-T2-QMUL, ...
 - More testing is essential
- Transition tools such as NAT64 can be used once core is IPv6-only
- WLCG may need a date for “end of support” for IPv4-only clients
 - e.g. start of LHC Run4?

Summary



- WLCG is ready to support use of IPv6-only CPU resources
 - **Good steady progress towards this goal!**
- Tier-1s should all have production storage accessible over IPv6
 - 96% of Tier-1 storage is available via IPv6
- Tier-2s 70% sites done
 - 73% of Tier-2 storage is dual-stack
- ~57% of FTS transfers today over IPv6
- **~50% LHCOPN+LHCONE traffic observed at CERN is IPv6**
- WG Phase 3 - we are planning for move to IPv6-only services
- ***message to new research communities - build on IPv6 from start***



Questions?