



# IPv6-only networking – an update from the HEPiX IPv6 working group

Martin Bly (STFC UKRI), David Kelsey (STFC UKRI),  
Andrea Sciabà (CERN)

HEPiX, Nikhef, 16 Oct 2019

# On behalf of all colleagues 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/EGI.eu), T Finnern (DESY), C Grigoras (CERN/ALICE), B Hoeft (KIT), D P Kelsey (RAL), 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



*Since HEPiX in San Diego (25-30 March 2019)*

- *Two F2F meetings at CERN (2/3 May and 17/18 Sep 2019)*
  - *plus 3 meetings by Vidyo*
  - *Chasing site deployment and Data monitoring/dashboards to show IPv4 vs IPv6*
  - *Starting plan for Phase 3 - move to IPv6-only networking*
- WLCG dual-stack service endpoint tracking
- Tier-0/Tier-1/LHCOPN/LHCONE status
- Tier-2 status
- Data transfers
- Monitoring
- IPv6-only networking
- Summary

# WLCG services status (dual-stack) *all services, not just storage*



results from the LHC Experiment VO feeds  
([https://orsone.mi.infn.it/~prelz/ipv6\\_vofeed/](https://orsone.mi.infn.it/~prelz/ipv6_vofeed/)) (~44%)



Francesco Prelz

Fraction of endpoints listed in the CERN central BDII ([lcg-bdii.cern.ch](https://lcg-bdii.cern.ch)) where the DNS returns a dual-stack IPv6-IPv4 (A+AAAA) resolution (green line) or an IPv6-only resolution (blue line). ([https://orsone.mi.infn.it/~prelz/ipv6\\_bdii/](https://orsone.mi.infn.it/~prelz/ipv6_bdii/)).

# Tier-0/Tier-1/LHCOPN/LHCONE status

# Network and pS at Tier-1's









Bruno Hoefft

- All sites connected to LHCONE
- TW-ASGC Perfsonar server currently down
- RRC-KI-T1 is connected to with IPv6 to LHC[OPN/ONE], the perfsonar server is running from beginning of this week (Oct. 14)

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 to selected T1's



Link	spring 2019		fall 2019	
	IPv4	IPv6	IPv4	IPv6
KIT→CERN	62%	38%	25%	75%
CERN→KIT	61%	39%	22%	78%
IN2P3→CERN	25%	75%	38%	62%
CERN→IN2P3	36%	64%	23%	77%
NL-T1→CERN	26%	74%	28%	72%
CERN→NL-T1	19%	81%	12%	88%
NDGF→CERN	79%	21%	18%	82%
CERN→NDGF	19%	81%	20%	80%


[LINK](#) to these plots

# IPv6 traffic on LHCOPN & LHCONE at CERN

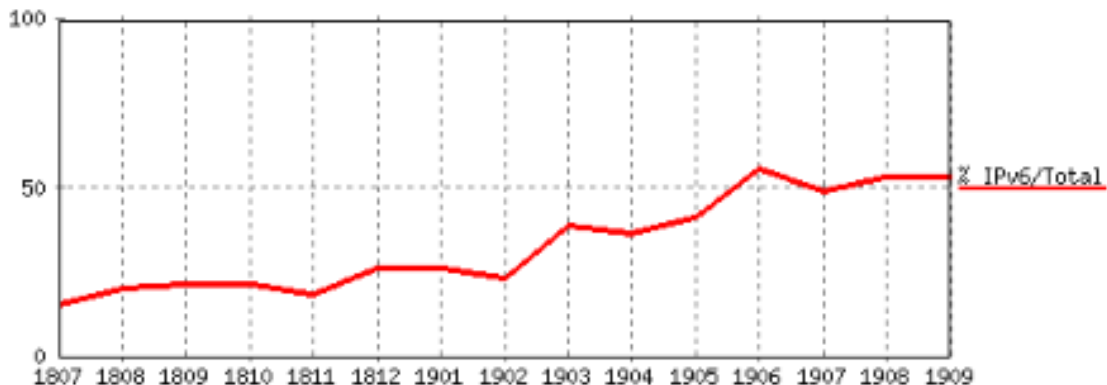
## LHCOPN and LHCONE IPv4 and IPv6 traffic volumes seen at CERN Tier0

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-1's

- FTS servers at BNL, CERN and RAL work in dual stack, while FNAL has still problems
- GridFTP transfers happen also via IPv6 at
  - IN2P3, JINR, NDGF, RAL, SARA-MATRIX, NIKHEF, CNAF, ASGC, PIC, TRIUMF, KIT, KISTI, FNAL
- 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
  - GGUS tickets sent to all non-US sites
  - US sites are tracked via the experiments
  - Sites made aware of the WLCG plans and asked to report plans and give updates
- 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




# Tier-2 status: reasons for delays

- 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
  - **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

Reason of delay [11-10-2019]



Time scale for deployment [10-10-2019]



# T2 deployment observations

- Tentative extrapolation to the future based on site replies
  - Within weeks: 80%
  - Within months: 94%
  - However, the estimates of the time to completion are found to be optimistic
- Most sites are responsive and provide detailed information
  - For some however regular pinging is a must...
- Several sites must wait for their campus infrastructure to become IPv6-ready
- It is evident that IPv6 is being deployed or will soon be deployed at the vast majority of the remaining sites and that the WLCG request is taken seriously

# Data transfers over IPv6 (FTS)

(Note: xRootD and HTTP transfers not yet instrumented to track IPv6 vs IPv4, but then only a small fraction of FTS transfers use xRootD or HTTP)

# FTS transfer monitoring - last 30 days


Approximately 54% 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  
also CHEP18





# Monitoring

# ETF News

Marian Babik

- Experiments instances ready
  - ATLAS, CMS and LHCb have now 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
- Other news
  - ETF has been migrated to Docker, all production instances now run in containers
  - Kubernetes support is currently being tested in pre-production
  - CMS interested to use ETF to monitor volunteer resources (including VAC-model in which jobs are pulled from ETF HTCondor pool)


# perfSONAR news

Marian Babik & Duncan Rand

- **Campaign to update all perfSONARs to CC7 and 4.2 ongoing**
  - Contacting sites directly (will issue tickets if necessary)
  - SL6 unsupported since Q4 2018
- CERN new perfSONAR servers in production
  - 2 servers with 40 Gbps/100 Gbps
- GÉANT is refreshing perfSONAR hardware at core points of presence
  - including four LHCONE interconnects at Frankfurt, Geneva, London and Paris
- Good progress with maddash
  - Solved most pS problems at Tier-1s
  - Found some bugs, already fixed by developers
- **NSF-funded IRIS-HEP and SAND**
  - Both have tasks/objectives related to perfSONAR

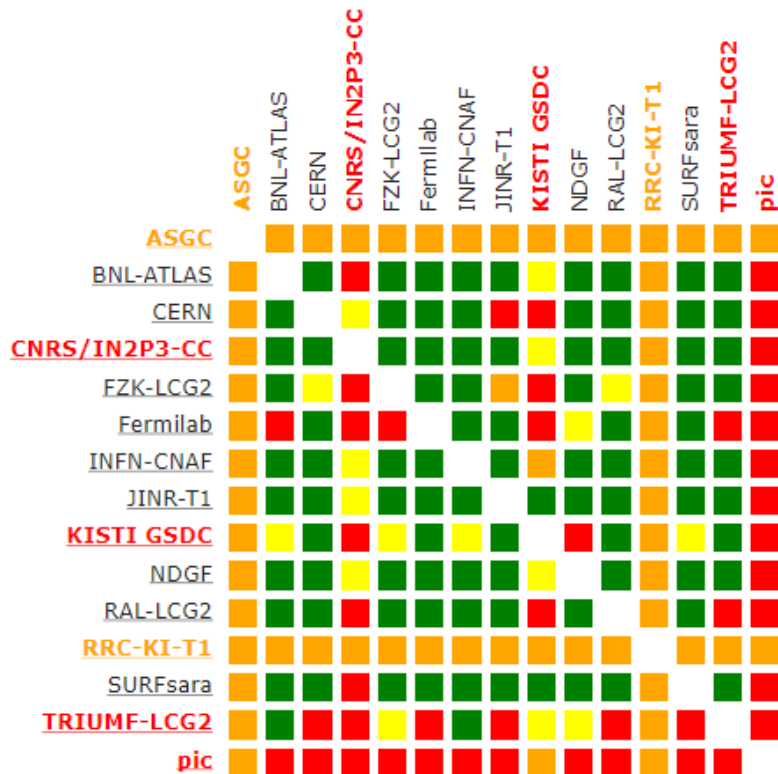
# perfSONAR IPv6 Mesh

OPN Mesh Config - OPN IPv6 Bandwidth - Throughput



! Found a total of 7 problems involving 6 hosts in the grid

[LINK](#) to dashboard



- Re-configured as was agreed
  - 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
- New monitoring is in place which reports test “efficiency” wrt IPv4 and IPv6
  - Efficiency = number of working destinations / total destinations
  - For each perfSONAR node and each mesh, reports missing results in the perfSONAR’s local measurement archive ([link](#))

# IPv6-only networking

- Main use case is to be ready for use of opportunistic IPv6-only CPU
- BUT there are other drivers
  - new CERN machine room likely to not have enough public IPv4 addresses
    - only IPv6 addresses for external public networking
  - multiONE (several different communities using LHCONE)
    - needs multiple overlay networks with different addresses per community
    - sites likely to not have sufficient IPv4 address space

# IPv6-only networking (2)

- Running a dual-stack IPv4/IPv6 infrastructure is complex
- Many large companies (FaceBook, EE, ...) use IPv6-only internally
  - CERN EOS infrastructure 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
  - could consider turning off IPv4 entirely on LHCOPN
  - simplifies routing tables and tracking problems is easier
- MultiONE/LHCONE may also be using IPv6-only

# IPv6-only networking on WLCG?

- Turning off IPv4 at a Site is clearly their own decision
  - they may have many other needs (not WLCG) for IPv4
- We need to examine experience of sites already doing IPv6-only
  - Brunel, QMUL, Slovenia, Nebraska, ...
- Continue study dual-stack endpoints that choose IPv4 and not IPv6
- 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?
- **message to new research communities - build on IPv6 from start**

# 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
- 54% of FTS transfers today over IPv6
  - No monitoring yet for FTS use of xRootD or HTTP (but that is small)
- **>50% LHCOPN+LHCONE traffic observed at CERN is IPv6**
- Phase 3 - we are planning for move to IPv6-only networking





# Questions?