Karlsruhe Institute of Technology

IPv6 @ LHCOPN + LHCONE

Bruno Hoeft / KIT (HEPiX-WG)

STEINBUCH CENTRE FOR COMPUTING - SCC

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-1: deployment of dual-stack on production storage, CVMFS and FTS by April 2018

Deadline not achived

Deadline not achived

Network and pS @ Tier-1‘s

- All sites connected do LHCOPN/ONE except RAL
- RRC-KI-T1 don’t have IPv6 on perfSONAR

Tier-1			
	LHCOPN IPv6 Peering	LHCONE IPv6 Peering	dual stack Perfsonar
ASGC	✓	✓	LHC[OPN/ONE]
BNL	✓	✓	LHC[OPN/ONE]
CH-CERN	✓	✓	LHC[OPN/ONE]
DE-KIT	✓	✓	LHC[OPN/ONE]
FNAL	✓	✓	LHC[OPN/ONE]
FR-CCIN2P3	✓	✓	LHC[OPN/ONE]
IT-INFN-CNAF	✓	✓	LHC[OPN/ONE]
NGDF	✓	✓	LHC[OPN/ONE]
ES-PIC	✓	✓	LHC[OPN/ONE]
KISTI	✓	✓	LHC[OPN/ONE]
NL-T1	✓	✓	LHC[OPN/ONE]
RAL	✓		LHC[OPN]
RRC-KI-T1	✓	✓	??
JINR	✓	✓	LHC[OPN/ONE]
Triumf	✓	✓	LHC[OPN/ONE]

IPv6 and FTS @ Tier-1‘s

- All FTS servers are dual stack, except the FTS server at FNAL
  - The IPv6 fraction of transfers could be larger than it is currently
- GridFTP transfers happen also via IPv6 at
  - IN2P3, JINR, NDGF, RAL, SARA-MATRIX, NIKHEF, CNAF, ASGC, PIC, Triumf, BNL, KIT, KISTI
    - DE-KIT : Redeployment of the storage network (move from a dual-homed setup (which does not work well at dCache with GridFTP) to a dual-stack setup realised
    - TRIUMF : just finished migration of all dCache nodes to dual-stack servers
    - KISTI : not visible at FTS monitor → Alice XROOT transfers only
- They do not at
  - FNAL and RRC-KI
- Application dual-stack readyness
  - falling back ipv6 to IPv4 does not work @ globus
    → ARC-CE IPv6 disabled until all Workernodes dual-stack

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

Period	IPv4 TB	IPv6 TB	Total TB	% IPv6/Total	Note
1807	28464	5202	33666	15.45	
1808	39139	9666	48805	19.81	estimation, some data missing
1809	32562	8714	41276	21.11	estimation, some data missing
1810	50016	13686	63702	21.48	estimation, some data missing
1811	35470	7940	43410	18.29	
1812	27478	9587	37065	25.87	
1901	25501	8991	34492	26.07	
1902	29217	8833	38050	23.21	
1903	25540	15934	41474	38.42	
1904	22205	12731	34936	36.44	

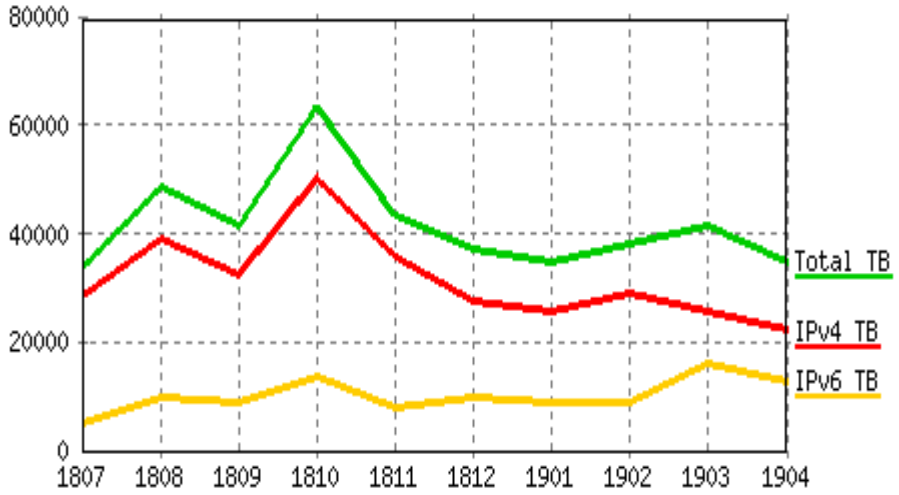

LHCOPN and LHCONE
IPv4 and IPv6 traffic volumes seen at CERN Tier0

Period	IPv4 TB	IPv6 TB	Total TB	% IPv6/Total	Note
1807	28464	5202	33666	15.45	
1808	39139	9666	48805	19.81	estimation, some data missing
1809	32562	8714	41276	21.11	estimation, some data missing
1810	50016	13686	63702	21.48	estimation, some data missing
1811	35470	7940	43410	18.29	
1812	27478	9587	37065	25.87	
1901	25501	8991	34492	26.07	
1902	29217	8833	38050	23.21	
1903	25540	15934	41474	38.42	
1904	22205	12731	34936	36.44	

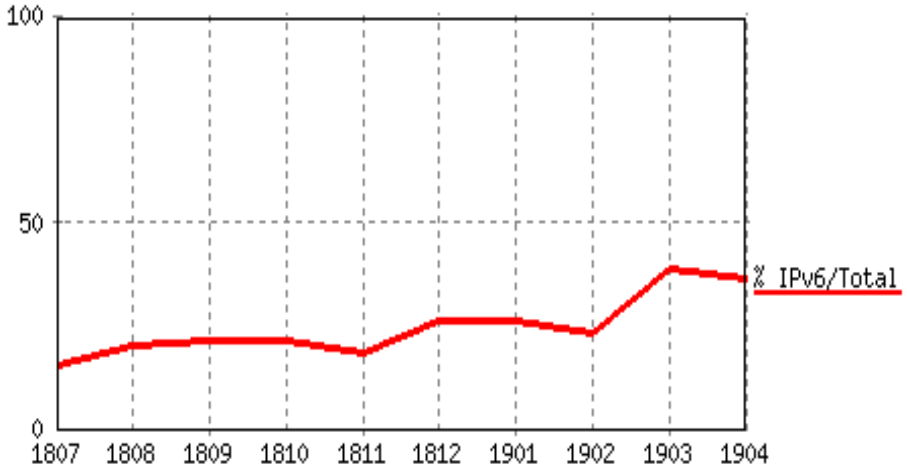

@DE-KIT – after Alice, Atlas, CMS, LHCB and Belle-2 dual-stack enabled

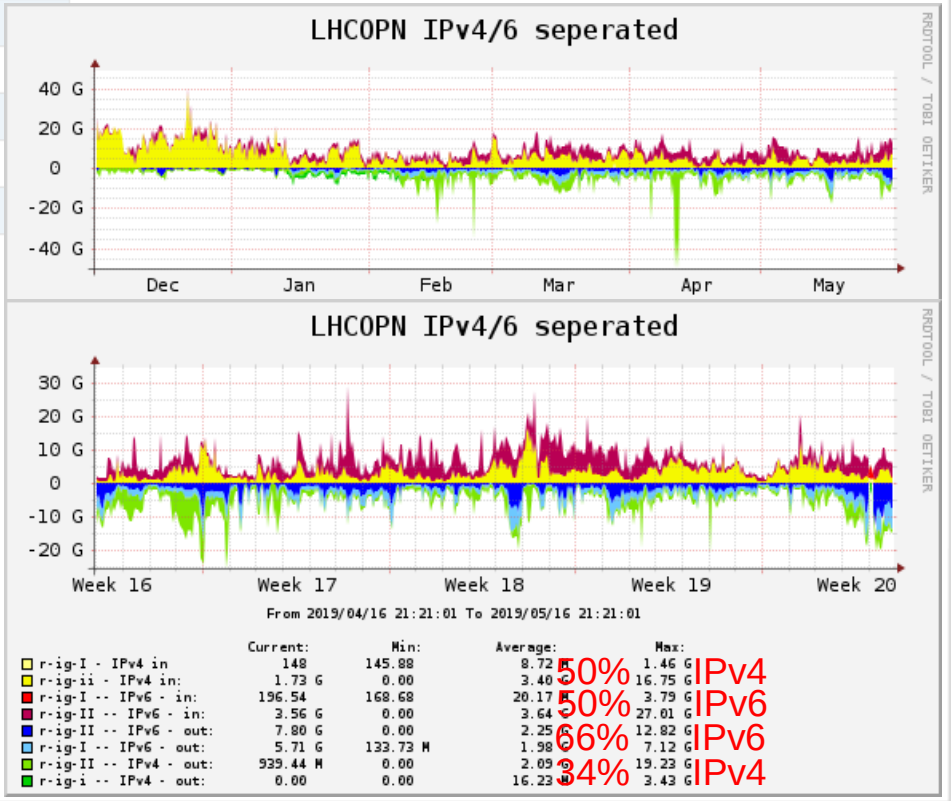

IPv6 Tier-2 sites

- The deployment campaign was launched in November 2017
  - Still ongoing -- Steady progress (status)
  - About 60% of T2 sites have storage on dual stack

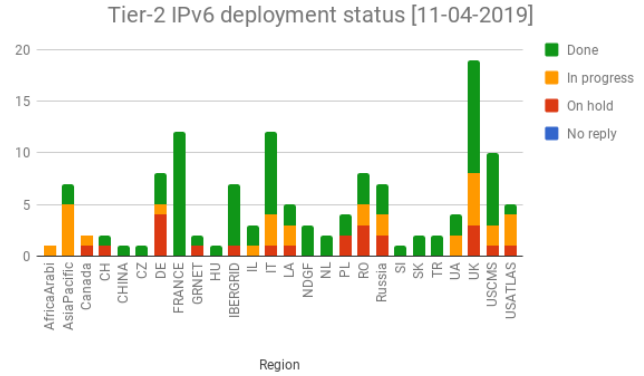

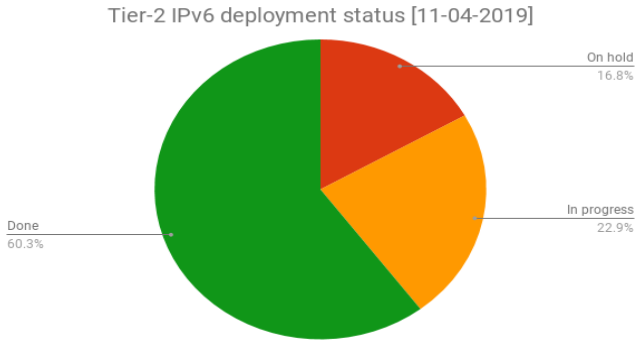

Experiment status

VO	T2 storage on IPv6 (%)	
ALICE	72	
ATLAS	49	
CMS	83	
LHCb	63	
WLCG	65	← everage status

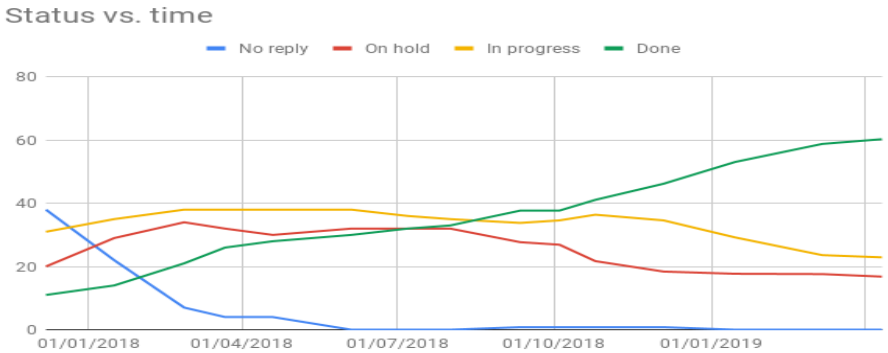

T2 deployment observations

- USATLAS and USCMS sites are now also tracked
  - Not via GGUS but via the experiments
- Regions differ greatly with respect to their status
  - NDGF, IBERGRID, France, USCMS lead the pack
    - No change
- Very few sites (<5) say they won't meet the deadline
    - 40% of tier-2 are still not ready yet
- 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
    - Still the case → country programms investing to enable IPv6 at R&E institutions

perfSONAR News

- 4.1.6 is the latest version (4.2.0 soon - preemptive scheduling & gridftp)
- Campaign to update all perfSONARs to CC7 and 4.1 on-going
  - Thanks to Duncan most of the T1s updated and work fine
  - UK and FR meshes also in very good shape
  - Still few issues to resolve (RRC-KI has no IPv6, TRIUMF waiting for hardware, NDGF lat. off)
- 100 Gbps perfSONARs
  - SARA, CSCS already have 100Gbps; CERN is now 40Gbps (100Gbps to be connected), BNL has 80Gbps (plans 100Gbps), KIT has 100Gbps in testing
- Dedicated IPv6 Meshes integrated e.g.: in LHCOPN/ONE

…

9

Conclusion

- IPv6 peering of Tier-1 sites in place
- last Tier-1 sites are getting IPv6 „production“ ready (Q1 2019)
  - except FNAL and RRC-KI
- Several Tier-2 site on a good track – 40% did not match the timeline (end of 2018)

- IPv6 Minitoring
  - url of NRENs → collect them at the HEPiX-WG page
    - Still no central NREN URLs IPv4/6 throughput graphs
    - Tier-x approaching their NREN individual -- seperating IPv4/6 traffic

Dualstack deployment includes additional complexity

- First investigations: identifying areas for IPv6 only deployment (LHCOPN?)

Karlsruher Institut für Technologie

Questions
Suggestions
Discussion