



IPv6 in production: its deployment and usage in WLCG

David Kelsey

STFC UK Research and Innovation

At CHEP18, Sofia, 9 July 2018

Thanks to all colleagues in the HEPiX IPv6 working group

Active in HEPiX IPv6 Working Group 2018

- M Babik (CERN), M Bly (RAL), T Chown (Jisc), J Chudoba (Prague), C Condurache (RAL), A Dewhurst (RAL/ATLAS), T Finnern (DESY), C Grigoras (CERN/ALICE), B Hoefft (KIT), D P Kelsey (RAL), F Lopez Munoz (PIC), E Martelli (CERN), R Nandakumar (RAL/LHCb), K Ohrenberg (DESY), F Prelz (INFN), D Rand (Imperial), A Sciaba (CERN/CMS)
- many more in the past, and others join from time to time
 - *apologies to any I have missed*
- And thanks also to WLCG operations, WLCG sites, LHC experiments, networking teams, monitoring groups, ...

IPv6 WG – History @ CHEP

- CHEP13: *“WLCG and IPv6 – the HEPiX IPv6 working group”*
 - Initial work on a testbed (lots of work needed on storage services)
- CHEP15: *“The production deployment of IPv6 on WLCG”*
 - First dual-stack services in production
- CHEP16: *“Deployment of IPv6-only CPU resources at WLCG sites”*
 - Plans to support IPv6-only CPU
- CHEP18: *“IPv6 in production: its deployment and usage in WLCG”*
 - Full deployment of dual-stack services well underway

Outline

- WLCG Management Board approved timelines
- WLCG dual-stack service endpoint tracking
- Tier-1/Tier-2 status
- IPv6 traffic on LHCOPN & FTS transfers over IPv6
- Monitoring
- Work in progress
- Summary

WLCG deployment plan: timeline



- By April 1st 2017
 - — Sites can provide IPv6-only CPUs if necessary
 - — Tier-1's must provide dual-stack storage access with sufficient performance and reliability
 - At least in a testbed setup
 - — Stratum-1 service at CERN must be dual-stack
 - — A dedicated ETF infrastructure to test IPv6 services must be available
 - — ATLAS and CMS must deploy all services interacting with WNs in dual-stack
 - — All the above, without disrupting normal WLCG operations
- By April 1st 2018
 - — Tier-1's must provide dual-stack storage access in production with increased performance and reliability
 - — Tier-1's must upgrade their Stratum-1 and FTS to dual-stack
 - — The official ETF infrastructure must be migrated to dual-stack
 - — GOCDB, OIM, GGUS, BDII should be dual-stack
- By end of Run2
 - — A large number of sites will have migrated their storage to IPv6
 - — The recommendation to keep IPv4 as a backup will be dropped

WLCG services status (dual-stack)




results from the LHC Expts VO feeds

(http://orsone.mi.infn.it/~prelz/ipv6_vofeed/) (~27%)



Fraction of endpoints listed in the CERN central BDII (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). (http://orsone.mi.infn.it/~prelz/ipv6_bdii/).

CERN Non-LHCOPN traffic – showing dual-stack EOS turn-on in Jan 2018



[Graphic help ?](#)

<https://netstat.cern.ch/monitoring/network-statistics/ext/?q=IPv6&p=EXT&mn=Internet&t=Yearly>

Tier 1 and Tier 2 status

Deployment of IPv6 and dual-stack storage

Tier-1 Status



- All 14 Tier 1 sites are now peering with LHCOPN over IPv6
- Dual-stack storage is gradually deployed
- Transferring (FTS) over IPv6: IN2P3, INFN, JINR, KIT (some), NDGF, Nikhef, RAL, RRC-KI, SARA, pic
- Transferring (FTS) over IPv4: ASGC, BNL, KISTI, KIT (some), RRC-KI, FNAL, TRIUMF
- FTS servers still on IPv4: FNAL, BNL
- CAUTION – we are not able to monitor all data protocols, so apologies to any TIER 1 wrongly shown as red
- Eg. TRIUMF says “data transfer is primarily using ipv6 now”

Sites IPv6 connectivity

Name	Type	LHCOPN IPv6 peering	LHCONE IPv6 peering	LHCONE IPv6 peers	NREN IPv6 peers	IPv6 LAN	dualstack perfSONAR	dualstack storage percentage by 1st April 2017	dualstack storage percentage by 31st May 2017	dualstack storage percentage by 31st July 2017
CH-CERN	Tier0	Yes	Yes	GEANT, ESnet, CERNlight	GEANT, SWITCH, RENATER, SURFnet, NORDUnet, ASGCnet, KREOnet, Internet2, CANARIE	Yes	Yes	5	15	15
TRIUMF	Tier1	Yes	Yes	Canarie	BCNET	Yes	Yes	0		
ES-PIC	Tier1	Yes	Yes	RedIRIS	RedIRIS	Yes	Yes	100	100	100
FNAL	Tier1	Yes	Yes	ESnet	ESnet	Yes	Yes	0	0	

<https://hepiv6.web.cern.ch/sites-connectivity>

And table continues ...

Tier 1 – some ongoing issues (dCache/gridftp)

At DE-KIT: for CMS IPv6 dCache (dual-homed servers)


- FTS is deployed dual-stack and for the dCache doors a special configuration is deployed with a dual-stack setup
- When gridftp used in dual-homed dcache-pools and a protocol transition from ipv6 to ipv4 is necessary, then the transfer is broken
 - Actually appears to be a problem with gridftp
- DE-KIT moving to use of dual-stack instead of dual-homed

BNL: “IPv6 has been disabled for BNL instance of FTS due to issues observed between dCache sites and plain gridftp sites for the production transfers”

Tier-2s: GGUS tickets to all 115 Tier-2 sites (not OSG)


https://twiki.cern.ch/twiki/bin/view/LCG/Wlcvlcv6#WLCG_Tier_2_IPv6_deployment_stat

Tier-2 IPv6 deployment status [07-07-2018]



WLCG Tier 2 IPv6 Deployment Status (excl OSG)

Tier-2 IPv6 deployment status [05-07-2018]



32% of Tier 2 IPv6 deployments complete – confirmed by experiments

LHCb Disk Storage – Dual Stack

Preliminary estimates (Raja Nandakumar) from storage accounting

- Tier 1 11.1 out of 18.7 PB (~60%) is dual-stack
- Tier 2 0.8 out of 2.8 PB (~30%)
- Overall 11.9 out of 21.5 PB (~55%)

Other experiments have not yet calculated these fractions

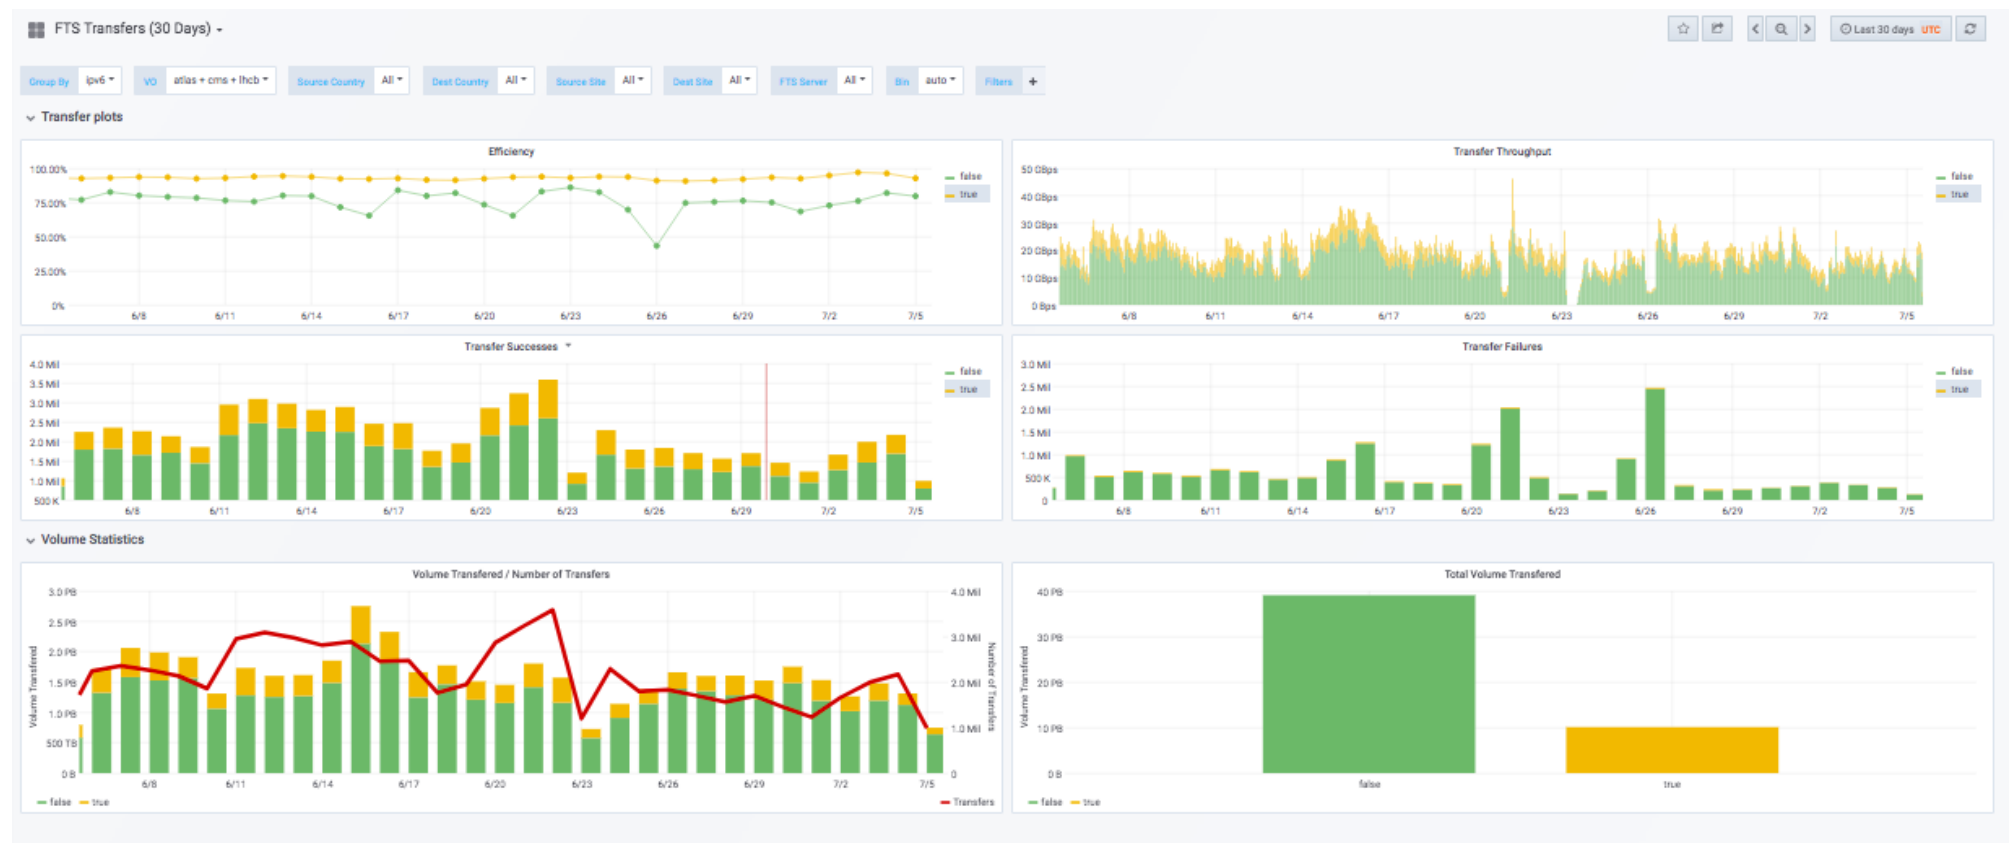
FTS – transfers over IPv6
(xRootD transfers not yet instrumented
to track IPv6 vs IPv4)

FTS transfer monitoring

Approximately 20% of data transferred in the last 30 days went over IPv6

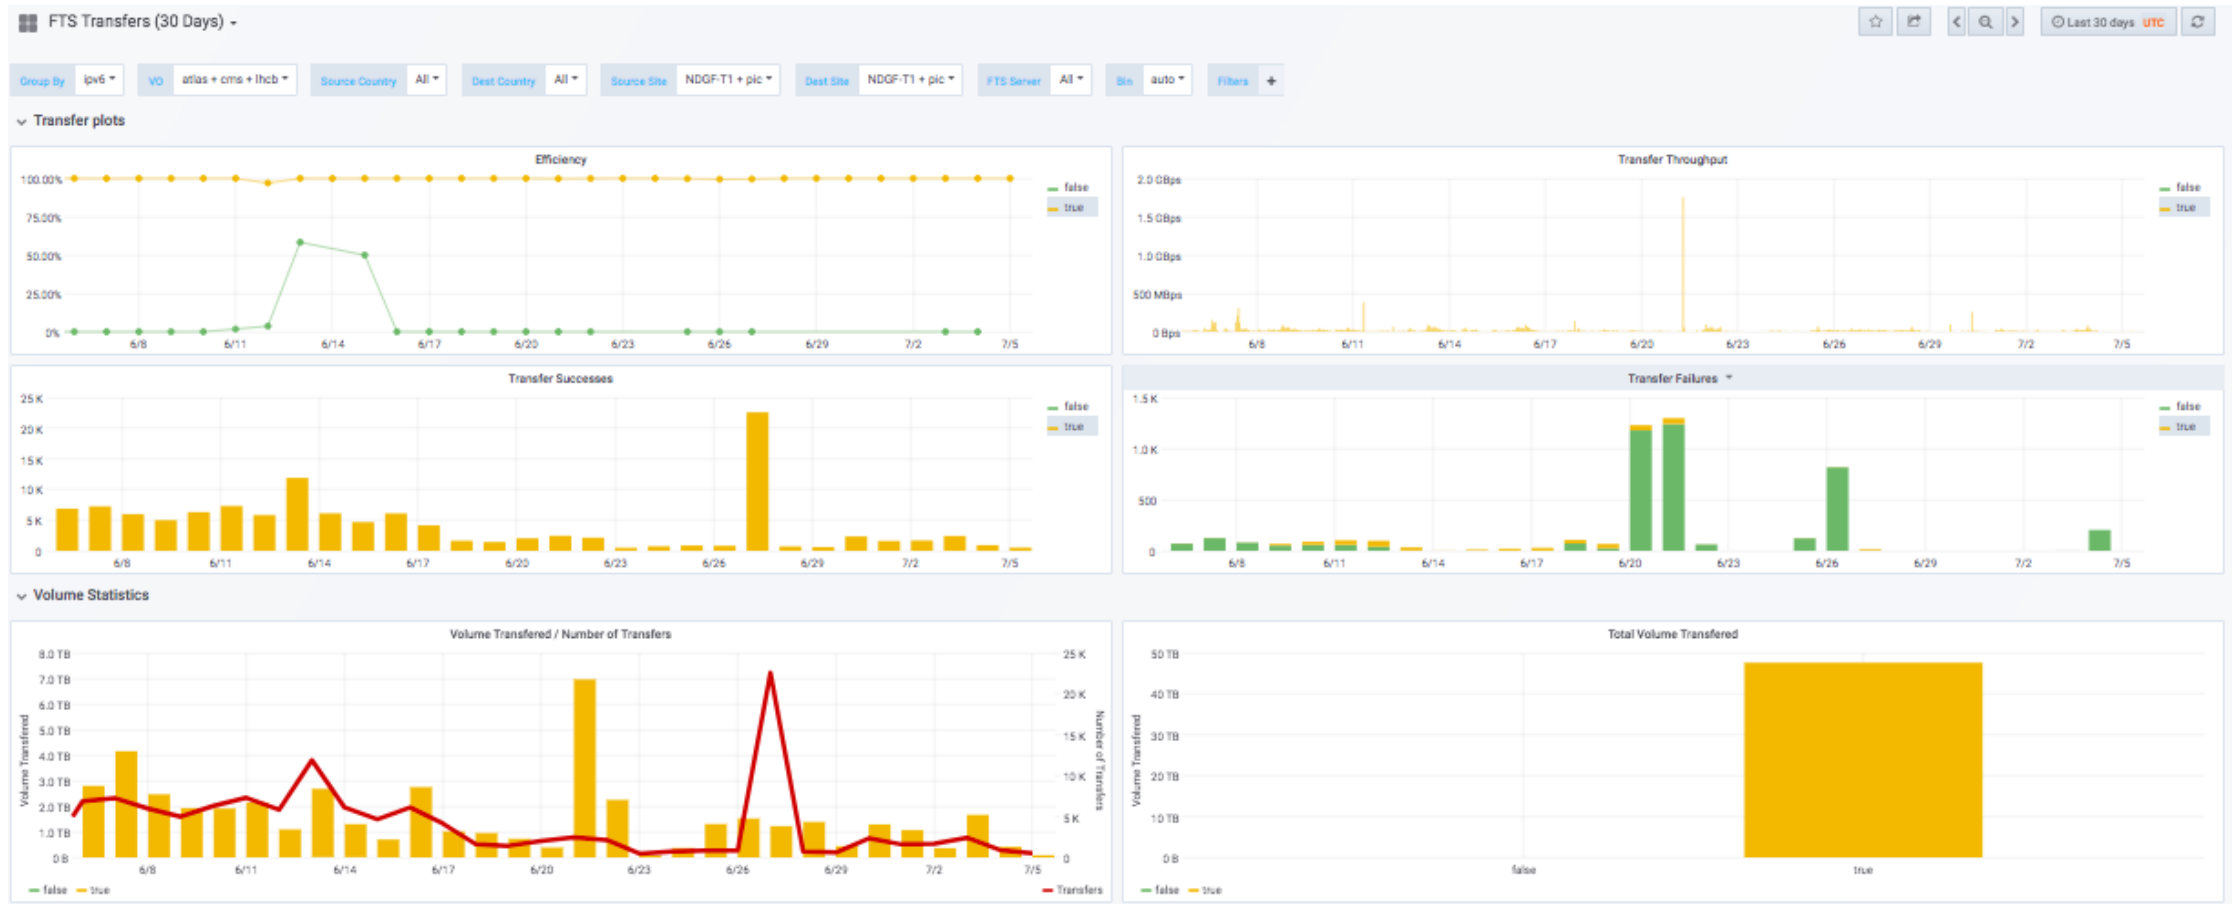
Transfers over IPv6 appear to be more efficient - interesting!

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



FTS Data Transfers (between 2 Dual-Stack Tier-1s)

“NDGF-T1” to/from “pic” (~100% IPv6)



IPv6 traffic – CERN LHCOPN

- Two routers at CERN (LHCOPN Tier 1 to CERN & LHCONE access links)
- For period 7-Jun-2018 to 6-Jul-2018
- Processed all network flow information (Edoardo Martelli)
 - data transited by CERN
- IPv4 ~26.7 PB (82.5%)
- IPv6 ~5.7 PB (17.5%)
 - This seems low – not yet investigated source/destination pairs
- Note: average total bandwidth was
 - IPv4 87.3 Gbps
 - IPv6 18.1 Gbps

Monitoring

perfSONAR & IPv6 - news

Marian Babik & Duncan Rand

- GEANT deployed IPv6 perfSONAR [instances](#) on LHCONE
- Grafana [dashboards](#) updated to version 5
 - IPv6 dashboard was introduced
- Dual-stack meshes have been re-configured
- Replacement of old dual-stack meshes – now underway
 - Create dual-stack LHCOPN with both IPv4 and IPv6 for all tests (Done)
 - Change all current experiment meshes to contain IPv6 throughput and tracepath
 - Create a dedicated IPv6/IPv4 latency mesh - only to be used for debugging specific cases


- WLCG LHCOPN, USATLAS USCMS are now also testing over IPv4 & IPv6 where possible
- LHC VO meshes still testing over IPv4 – will be addressed soon



OPN Mesh Config - IPv6 TCP BWCTL Test Between OPN Bandwidth Hosts

■ Throughput >= 900Mbps
 ■ Throughput < 900Mbps
 ■ Throughput <= 500Mbps
 ■ Unable to retrieve data
 ■ Check has not yet run


! Found



OPN Mesh Config - IPv6 Traceroute Test Between OPN Bandwidth Hosts

■ Number of Paths is <= 1
 ■ Number of Paths is >= 1
 ■ Number of Paths is >= 2
 ■ Unable to retrieve data
 ■ Check has not yet run

✓ No



<http://psmad.grid.iu.edu/maddash-webui/index.cgi?dashboard=OPN%20Mesh%20Config>

Experiment Test Framework (ETF) news

- ETF IPv6 instance
 - Switched to IPv6-only in February 2018 to test for possible issues
 - (https://etf-ipv6.cern.ch/etf/check_mk/)
 - Myproxy is still IPv4 and is on a critical path as unable to get credentials otherwise
- 2 Experiment instances running (others can request)
 - CMS and LHCb IPv6 instances were created
 - https://etf-cms-ipv6.cern.ch/etf/check_mk/
 - https://etf-lhcb-ipv6.cern.ch/etf/check_mk/
- For both LHCb and CMS have started to publish results to SAM3 (QA)
 - Able to aggregate IPv6 results and compute IPv6-only profiles
 - sample profile already exists in SAM3 QA

Other “work in progress”

- pic Tier 1 – investigating why transfers made over IPv4 when both ends are IPv6?
 - Good news – successfully runs over IPv4 (no problems seen!)
- ATLAS and CMS are testing IPv6-only worker nodes
 - ATLAS: SiGNET (Slovenia) – production work on IPv6-only WN – works well
 - CMS: Lincoln, Nebraska – lots of fun with config, sys admin, s/w compatibility

IPv6 deployment on WLCG - Summary



- WLCG to support use of IPv6-only CPU resources by end of Run 2
- **Good steady progress towards this goal!**
- Tier-1s should already have production storage accessible over IPv6
 - > 50% are currently transferring data over IPv6 (perhaps more)
- Tier-2s asked to be ready by the end of 2018 (~32% done)
- ~20% of FTS transfers today over IPv6
 - No monitoring yet for xRootD
- ~45% perfSONAR hosts now reporting IPv6-enabled
- There are still some issues to fix and/or understand
- Tier 2's – please keep us informed via your ticket
 - And let us know if you need advice
- **Many thanks** to all of you who have contributed to this endeavour!

Questions?