



# WLCG IPv6 deployment strategy

Alastair Dewhurst  
Andrea Sciabà  
on behalf of the HEPiX IPv6 WG

19/07/2016, WLCG MB

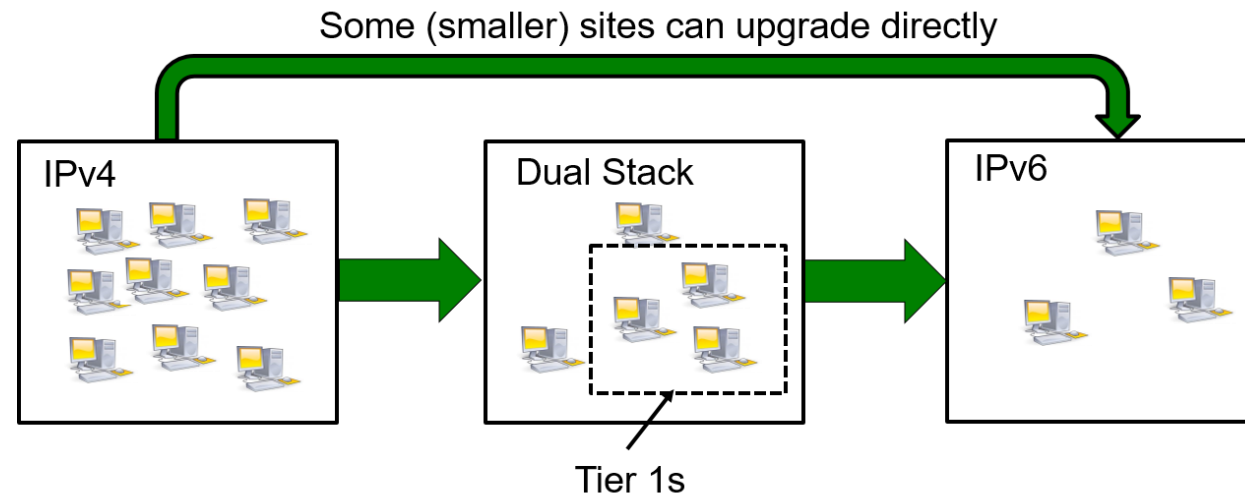


# Introduction

- IPv6 requirements of the experiments have been discussed at the last GDB.
- A document has been submitted to the MB for **approval**.
  - A summary is presented here.
- The main aims:
  - To provide a viable migration path for sites needing to switch to IPv6.
  - To allow sites to make long term planning decisions regarding their network setup.
  - To allow VOs to make use of IPv6-only CPU resources should they become available in future.

# Typical migration path for a site

- Deploy IPv6 on the network infrastructure
- Deploy a dual-stack perfSONAR instance
- Make storage dual-stack
  - Applies to all sites, for the sake of allowing others to remotely access data via IPv6
- Make local services dual stack
  - Eventually it should not be necessary, e.g. IPv4 could be decommissioned
- Make all WNs IPv6-only
  - Allow for a “grace period” during which IPv4 is kept as backup
  - Use IPv4 private addresses if needed (e.g. by batch system)



# Role of Tier-0/1 sites

- Extremely important as they provide access to, and distribute, a lot of data
- They run several central Grid and VO services
  - FTS, Frontier, Stratum-0/1, VOMS, MyProxy, WM/DM services, ETF, Dashboard monitoring...
- Therefore, critical to make services dual-stack
  - Even if availability of IPv4 addresses is not an issue at the site
- Storage performance and reliability should not depend on the IP protocol version used. Proposed targets:
  - At least 1Gb/s and 90% reliability by April 1st 2017
  - At least 10 Gb/s and 95% reliability by April 1st 2018

# ATLAS, CMS and LHCb

- ATLAS, CMS and LHCb are all aiming to have their critical central services dual stack by April 1st 2017.
- The VOs do not rely on a fully federated storage model.
  - Can run some workflows on IPv6-only resources as long as there is one dual stack storage.
- The more sites that provide dual stack storage the easier it becomes to integrate IPv6-only CPU resources into their computing model.
- All VOs urge caution from sites when migrating as we are in a data taking period.

# ALICE

- ALICE central services have been dual stack for more than a year.
- Storage is **fully** federated
  - Any site can access data from any site
- To support IPv6-only resources, **all data** must be available on some IPv6-enabled storage
- ALICE will not be able to support IPv6-only CPU resources by April 1st 2017.
- ALICE can support IPv6-only CPU resources as soon as enough sites have upgraded their storage to dual stack.
  - Goal of the end of Run II.

# Summary

- Sites **can** provide IPv6-only CPU resources from **April 2017** onwards **if necessary**
- Sites **can** provide IPv6-only interfaces to their CPU resources, **if necessary**
- Any site wishing to deploy IPv6-only CPU resources should contact the HEPiX IPv6 working group to discuss detailed plans.
- From April 1st 2017 central services **must** be accessible via IPv6 and by April 1st 2018 an **equal** quality of service via both IPv4 and IPv6 is required.
- All sites are encouraged to upgrade their storage to dual stack. CERN and the Tier 1s will be **required** to provide dual stack access to their storage from April 1st 2017.
- By the **end of Run II** enough sites should have upgraded their **storage** to dual stack to allow almost complete data availability over IPv6
  - ALICE aim to be able to support IPv6-only CPU resources by the end of Run II.

# Backup






# Is the middleware ready?

- In short, yes
- HTCondor fully supports IPv6
- CEs do
  - CREAM, ARC-CE, HTCondor-CE
    - But still some problems with CREAM client, critical for LHCb and ALICE
- xrootd 4, GridFTP, HTTP(s) do
  - Xrootd storage and redirection, GridFTP-only storage
- Most storage systems (dCache, DPM, StoRM, ...) do
  - CASTOR does not (and never will) but possible to provide a “gateway” service via xrootd/GridFTP if needed
  - EOS does, in versions using xrootd4
- All central Grid services do
  - MyProxy, VOMS, FTS, CVMFS, BDII, Frontier, ...
  - Although often not yet deployed in dual-stack

# Proposed timeline

- **By April 1<sup>st</sup> 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 1<sup>st</sup> 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
  - GOCD, 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

# Tier-1 status and plans



## IPv6 tier-1 site readiness ticket

status at Tuesday 07. June



Ticket-ID	Type	VO	Site	Priority	Resp. Unit	Status	Last Update	Subject	
121896		none	BNL-ATLAS	very urgent	OSG(Prod)	assigned	2016-06-06	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	KO
121895		none	USCMS-FNAL-WC1	very urgent	OSG(Prod)	solved	2016-06-03	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121894		none	RAL-LCG2	top priority	NGI_UK ▶ assigned	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	ok
121893		none	Taiwan-LCG2	top priority	ROC_AsiaPacific	in progress	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121892		none	SARA-MATRIX	top priority	NGI_NL	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121891		none	RRC-KI-T1	top priority	ROC_Russia ▶ assigned	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121890		none	NIKHEF-ELPROD	top priority	NGI_NL	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121889		none	NDGF-T1	top priority	NGI_NDGF	solved	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121888		none	KR-KISTI-GSDC-01	top priority	ROC_AsiaPacific	waiting for reply	2016-06-02	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121887		none	JINR-T1	top priority	ROC_Russia	in progress	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	?
121886		none	INFN-T1	top priority	NGI_IT	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121885		none	IN2P3-CC	top priority	NGI_FRANCE	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121884		none	pic	top priority	NGI_IBERGRID	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121883		none	FZK-LCG2	top priority	NGI_DE	solved	2016-06-01	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121882		none	CERN-PROD	top priority	ROC_CERN	in progress	2016-06-05	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	✓
121881		none	TRIUMF-LCG2	top priority	ROC_Canada	solved	2016-06-07	Tier-1 LHCOPN IPv6 Peering, incl. dualst...	ok

