



WLCG IPv6 deployment strategy

Alastair Dewhurst
Andrea Sciabà
on behalf of the HEPiX IPv6 WG

07/07/2016



Introduction

- IPv6 requirements of the experiments have been discussed at the last GDB
- A document has been prepared and almost finalised
 - To be submitted at the July MB
- A summary is presented here

Motivations

- Exhaustion of IPv4 address space
 - A real problem at some smaller sites
- Some cloud providers may provide only (or charge less for) IPv6-only resources
- Opportunistic resources may be IPv6-only

Goals

- Eventually, replace IPv4 with IPv6 (for all Internet)!
- Allow WLCG sites to deploy IPv6-only CPUs
- Provide an upgrade path for sites to plan
 - Ideally, minimise the need of dual-stack with respect to IPv6-only, given the constraints
- Consequently, make all VO and Grid services contacted by WNs dual stack
 - Storage
 - Central services

Is the middleware ready?

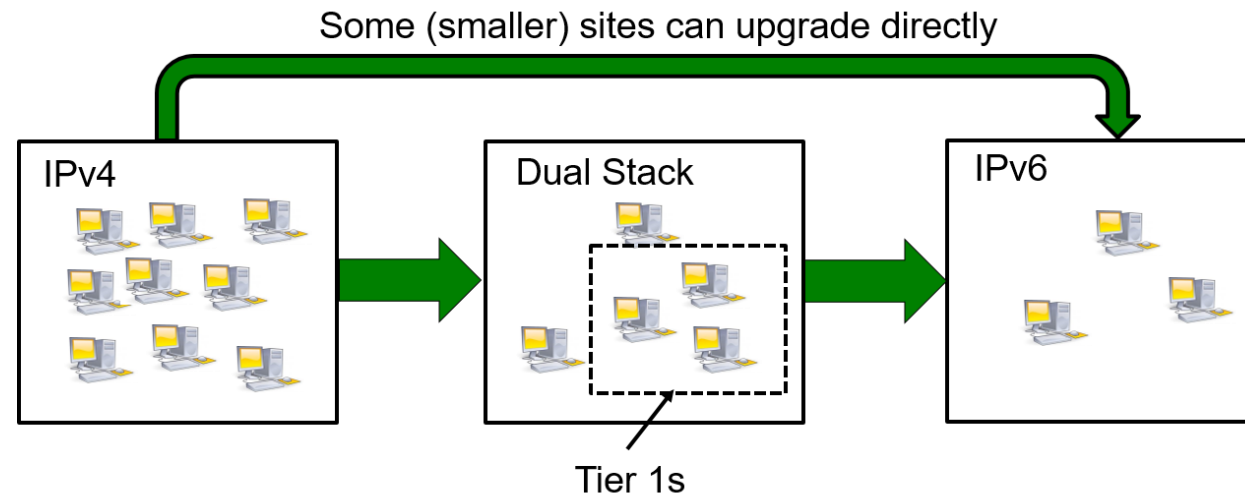
- In short, yes
- HTCondor fully supports IPv6
- CEs do
 - CREAM, ARC-CE, HTCondor-CE
 - But still some problems with CREAM client, critical for LHCb and ALICE
- xrootd 4, GridFTP, HTTP(s) do
 - Xrootd storage and redirection, GridFTP-only storage
- Most storage systems (dCache, DPM, StoRM, ...) do
 - 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

What does this mean for a site?

- IPv6-only WNs if needed
 - Dual-stack would not address the address exhaustion issue!
- IPv6-enabled CEs
 - IPv4 irrelevant if VO workload management supports IPv6
- Dual-stack storage
 - To support all CPU resources doing remote access and transfers to/from all sites
- IPv6-enabled Squid caches
 - IPv4 desirable if not all Stratum-1's support IPv6

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 all local services dual stack
 - In some cases (see batch system) may not be possible, e.g. only IPv4 is supported
 - 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% availability by April 2017
 - At least 10 Gb/s and 95% availability by April 2018


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



IPv6 peering + dual stack per sonar

work in progress

no response

ALICE status

- ALICE central services are running in dual-stack since 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
- All sites must run xrootd4
- Current situation:
 - 1/3 of sites still use xrootd3
 - 5% of SEs run in dual-stack

ATLAS status

- ATLAS WM and DM tools (PanDA, Rucio, ...) are based on HTTP and are tested on IPv6
- WNs interact via HTTP with PanDA, Rucio, CVMFS and Frontier
- Current situation and plans:
 - Pilot factories are dual-stack
 - Complete migration of services to dual stack by next April
 - Production PanDA server nodes
 - Rucio Authentication and Production nodes
 - Frontier servers at CERN, IN2P3, RAL and TRIUMF
 - ARC Control Tower

CMS status

- WM is based on glideinWMS and HTCondor, both validated on IPv6
- Federated storage via xrootd
- DM is based on PhEDEx, which uses the Oracle client for communication with central service
 - WNs do not interact with it
- Other central services are hosted on the cmsweb cluster, which works in dual-stack
- Current situation and plans:
 - glideinWMS frontends and factories are dual-stack
 - Only ~11 sites support IPv6
 - No problems seen in monitoring and computing operations
 - CMS has to enable dual-stack on all collectors and schedd's by April 2017
 - Sites have to enable dual-stack on their storage by the end of Run2
 - Sites are advised to keep IPv4 as a backup until the end of Run2

LHCb status

- WM/DM is based on DIRAC, which supports IPv6
- WNs need to access DIRAC central services, storage and optionally one of six VO boxes at Tier-1's
- Current situation and plans:
 - Only one Tier-1 and one Tier-2D storage system support LHCb in dual-stack configuration
 - Central services are being moved to dual-stack machines

Proposed 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

Executive 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**
- The VO infrastructure (e.g. central services provided by VOs) **must** provide an **equal** quality of service to both IPv4 and IPv6 resources
- Any site wishing to deploy IPv6-only CPU should contact the HEPiX IPv6 working group to discuss detailed plans
- By **April 2018** it should be possible to deploy IPv6-only CPU with relative ease
- By the **end of Run II** enough sites should have upgraded their **storage** to dual stack to allow almost complete data availability over IPv6