

# PIC Tier-1 IPv6 deployment

WLCG Ops Coord. meeting – CERN – 07/07/2016

J. Flix (PIC/CIEMAT) [remote]

for the PIC team

# IPv6 history at PIC Tier-1

- We started IPv6 tests in early **2013**
  - Main FW non IPv6 compatible. Handmade ip6tables FW
- **Feb 2013**: Joined the gridftp testbed of HEPiX IPv6 Working Group
- **Jun 2013**: Enabled dual-stack on production perfSONAR boxes
- **Aug 2013**: Testing dCache dual-stack compatibility
- **May 2015**: New main FW in production fully IPv6 compatible
- **January 2016**: Testing our non IPv6 compatible batch system (Torque/Maui) with dual-stack CREAM-CE and WN
- **March/April 2016**: Production dCache headnodes (SRM, gridftp, gsidcap | dcap, xrootd) and ATLAS,CMS , LHCb and ATLAS Tier-2 pools
- **May 2016**: Testing IPv6 in HTCondor and HTCondor-CE
- **May 2016**: CMS local PhEDEx nodes and testing CMS xrootd redirectors
- **June 2016**: WN tests with IPv6 and LHCb VOBOX moved to dual-stack

# IPv6 history at PIC Tier-1

- We started IPv6 tests in early **2013**
  - Main FW non IPv6 compatible. Handmade ip6tables FW
- **Feb 2013**: Joined the gridftp testbed of HEPiX IPv6 Working Group
- **Jun 2013**: Enabled dual-stack on production perfSOM nodes
- **Aug 2013**: Testing dCache dual-stack compatibility
- **May 2015**: New main FW in production for dual-stack compatible
- **January 2016**: Testing our non IPv6 compatible batch system (Torque/Maui) with dual-stack CREAM-CE and HTCondor
- **March/April 2016**: Production dual-stack dCache headnodes (SRM, gridftp, gsidcap|dcap, xrootd) and ATLAS LHCb and ATLAS Tier-2 pools
- **May 2016**: Production IPv6 in HTCondor and HTCondor-CE
- **May 2016**: Production IPv6 on CMS local PhEDEx nodes and testing CMS xrootd redirectors
- **May 2016**: WN tests with IPv6 and LHCb VOBOX moved to dual-stack

All the work in collaboration with the HEPiX IPv6 WG and the VOs

# Farm IPv6 tests 1/3

- Currently, private IPv4 deployed on WN. Routing forced to FW (NAT) and then to GPN
  - New strategy for IPv6 deployment → Public IPv6 address reachable via LHCOPN/ONE and GPN
- CREAM-CE (ce01.pic.es) in dual-stack, Torque/Maui not IPv6 compatible and one WN in dual-stack with 24 slots
  - Enabled for ATLAS, CMS, LHCb and ipv6.hepix.org
  - Jobs are submitted to **ce01.pic.es:8443/cream-pbs-ipv6** (special queue)
- New squid server, IPv6 compliant, installed for these tests  
<http://squid01-test.pic.es/>
  - CVMFS populated using this squid, then 'frozen' and manipulated, so jobs sent to this WN use this squid server for conditions data


# Farm IPv6 tests 2/3

- SAM tests (and production jobs) for **LHCb** working fine
  - Indeed LHCb has asked for an IPv6-only WN (see later)
- SAM tests for **ATLAS** OK after the PandaQueue was properly configured
- SAM tests for **CMS** OK, except the xrootd-access test that **fails**
  - When using IPv6, this test is failing atm... No problem are observed with IPv4
  - Debugging with the help of CMS AAA Federation (Marian Zvada) and XRootD experts

ce01.pic.es [Link to data](#)

### Test history ce01.pic.es using ATLAS GENERAL


720 hours from 2016-06-07 12:00 to 2016-07-07 12:21



ce01.pic.es [Link to data](#)

### Test history ce01.pic.es using CMS CRITICAL FULL

720 hours from 2016-06-07 12:00 to 2016-07-07 12:21



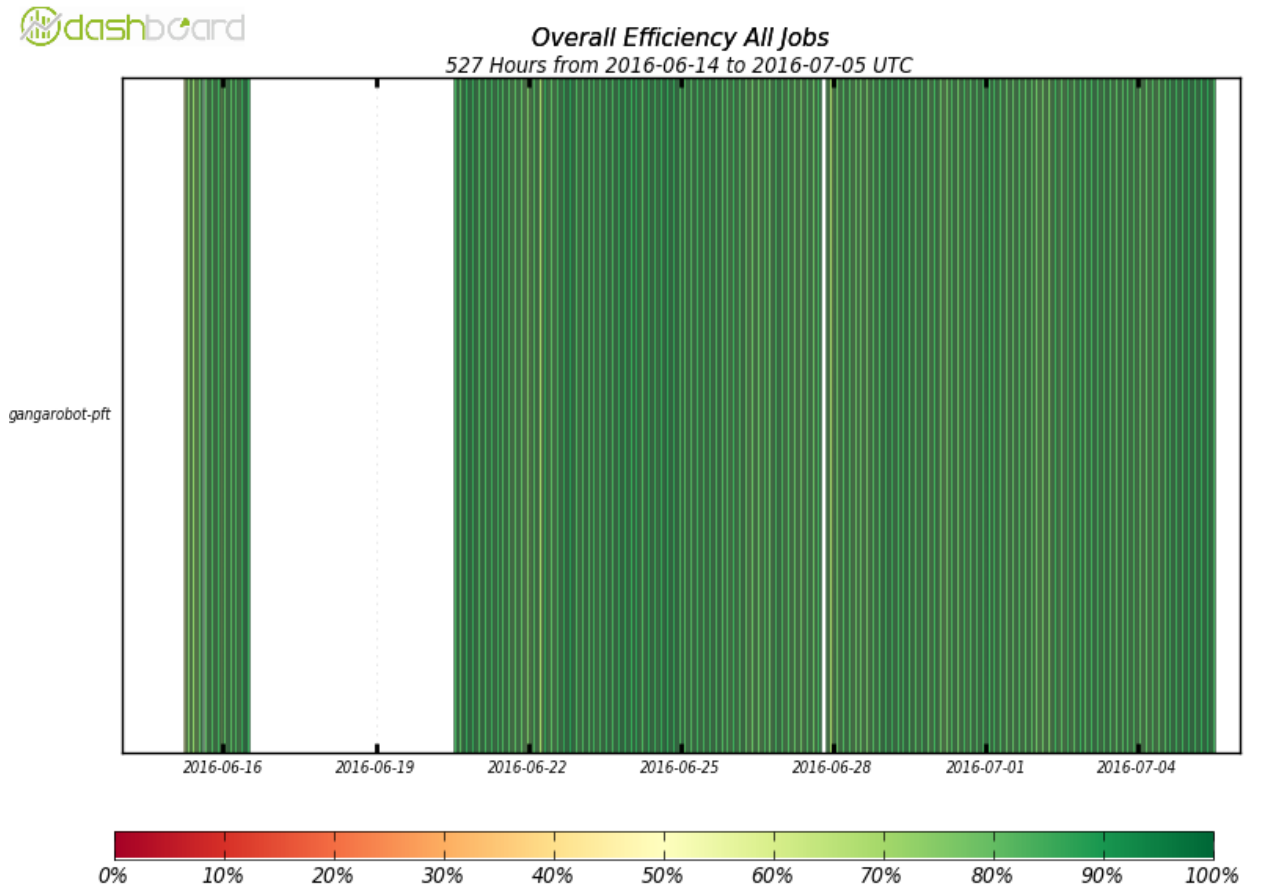

← TFC from this WN differs from Git

← xRootD CMS problems

# Farm IPv6 tests 3/3

- Setting up **HammerCloud** jobs that would routinely run at the test IPv6 environment
  - ATLAS and CMS are now executing HC jobs in this dual-stack node (since ~1 month)
    - CMS has commissioned their pilot factory under IPv6 by means of these activities
      - For CMS it was noticed that a pilot would take up to an hour to start pulling work. This was caused by tens of configuration files being retrieved from the factory to the WN via *wget*, each having to timeout before IPv4 being successfully tried. As this was corrected, the pilot running in the dual stack WN could contact the factory via IPv6 and start pulling work normally
  - PIC has offered the first IPv6 environment for HammerClouds
    - which has an added value for real production activities commissioning under IPv6

# HC - ATLAS

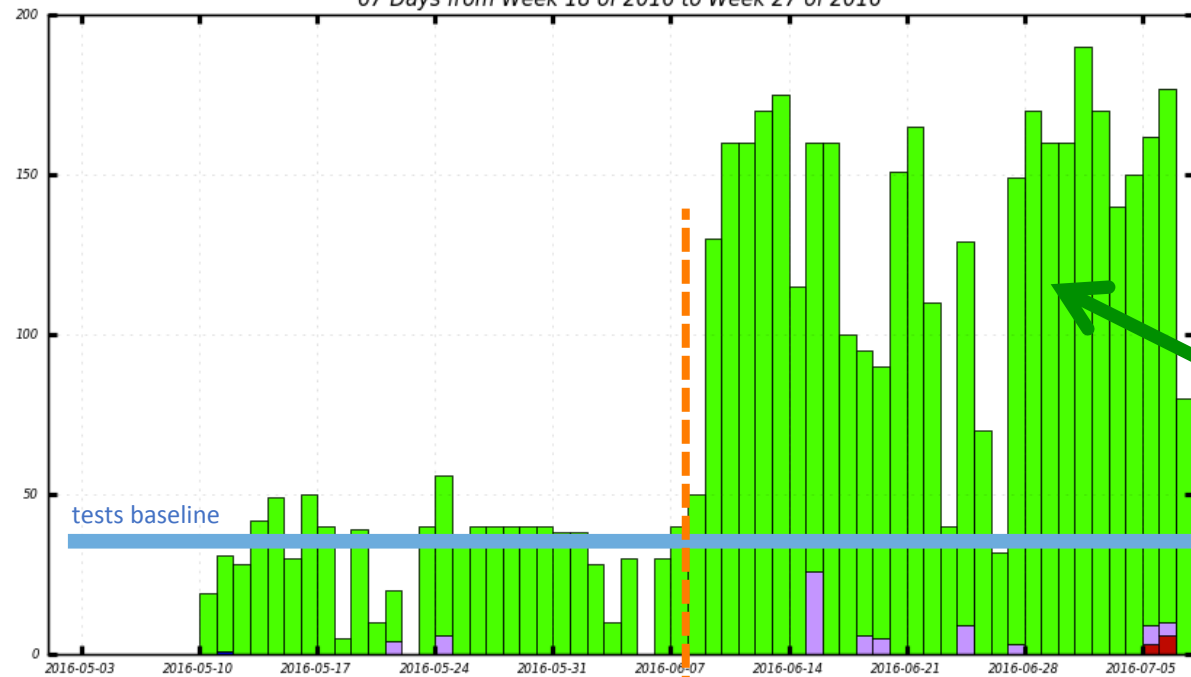




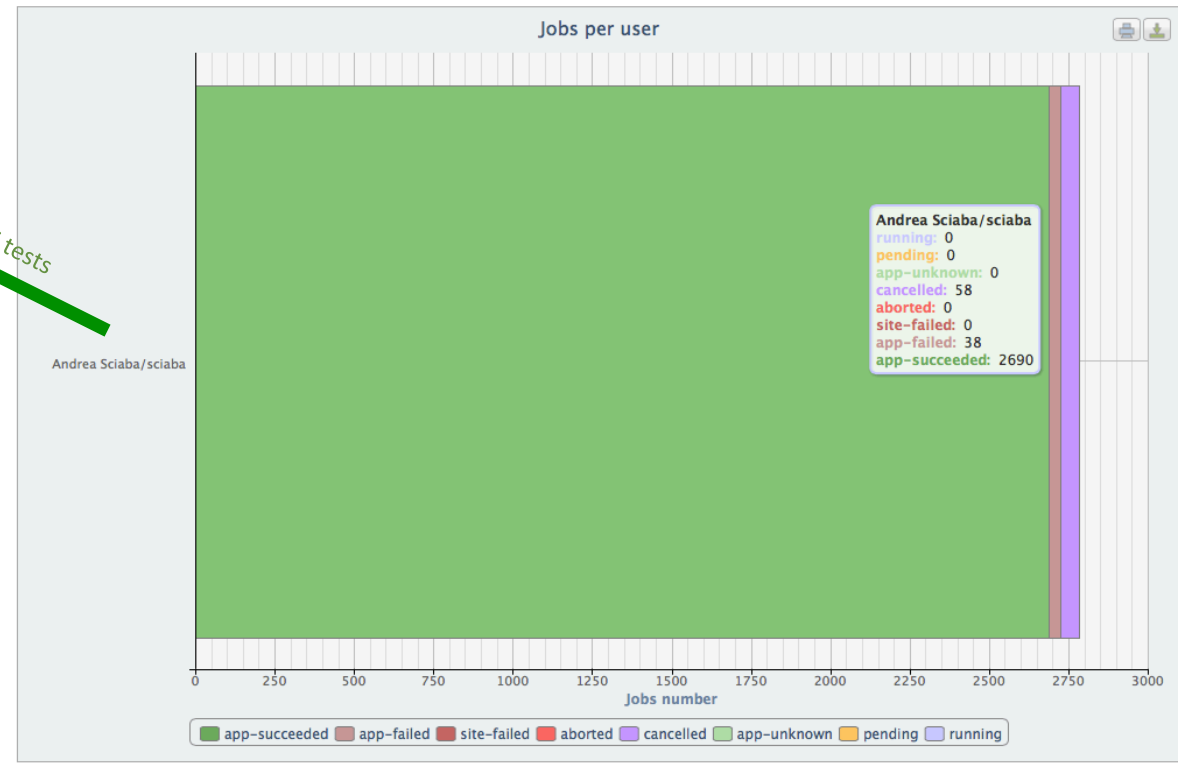
# HC – CMS



Number of Successful, Failed and/or Aborted Jobs  
67 Days from Week 18 of 2016 to Week 27 of 2016

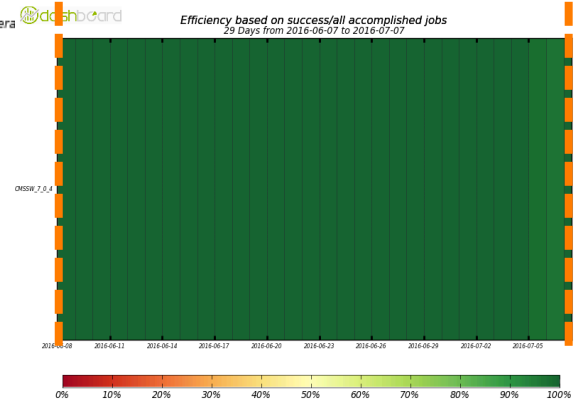


HC IPv6 tests



■ Number of Successful Jobs  
■ Number of Cancelled Jobs  
■ Number of Failed and/or Aborted Jobs  
■ Number of Unknown-Status Jobs

Maximum: 190.00 , Minimum: 0.00 , Average: 45.00  
 Efficiency based on success/all accomplished jobs  
 29 Days from 2016-06-07 to 2016-07-07



# Data federations

- **ATLAS:**

- The XRootD services are now operative (again, see today's note) and will be soon moved to dual-stack

- **CMS:**

- One of the local xrootd redirectors has been moved to dual stack ([xrootd02-cmst1.pic.es](http://xrootd02-cmst1.pic.es)), while keeping the other one IPv4 to cross check any effects ([xrootd01-cmst1.pic.es](http://xrootd01-cmst1.pic.es))
- the xrootd door [xrootd-cmst1-door.pic.es](http://xrootd-cmst1-door.pic.es), serving data to AAA, was moved to dual stack by mid-May 2016

# HTCondor tests (IPv6-only)

- **HTCondor 8.5.5** dual-stack testbed in place:
  - 2 Central Managers in High Availability
  - 1 ARC-CE (legacy, not dual-stack) + 2 HTCondor-CEs (thanks Iain B. Steers!)
  - 2 local schedds
  - 11 WNs available (132 cores) → 10 WNs dual-stack and 1 WN IPv6-only
- **LHCb** interested to start soon tests in the IPv6-only WN in this setup

# Next steps

- Debug the XRootD CMS problem
- Open HTCondor testbed for ATLAS, CMS and LHCb
  - Run SAM and HammerCloud tests
  - IPv6-only WN is provided, since HTCondor is IPv6 compliant, LHCb is interested to test it
- Enstore tests in IPv6
  - Working in close contact with the developers for testing this service
- These tests are particularly interesting for integrating all of the missing IPv6 elements and see the impact on real jobs