

PIC Tier-1

Batch System Evolution

Pre-GDB @ CERN – 12/05/2015

C. Acosta, R. Cruz, J. Flix

PIC BS current situation

Torque 2.5.13 + Maui 3.3-4:

- handles Grid submission + local user submissions
- Fairshare / priorities (users/SAM) in place – accounting through APEL pretty good
- MCFLOAT installed: dynamic farm partitioning to manage multicore jobs
- HighMem queue for ATLAS: 256-cores available and shared with the farm

No instabilities/scalability issues observed [currently: 68.5 kHS06 – 5400 slots]

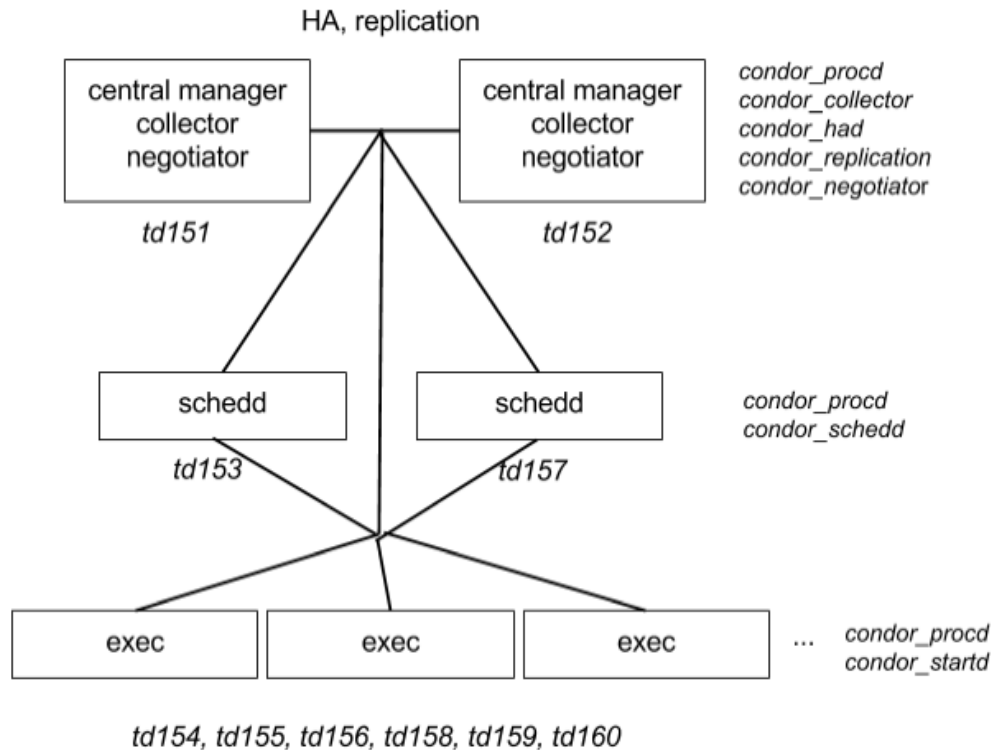
Concerns on (future) **scalability** and **security updates** of the product

- No (official) support since ~3 years → almost a defunct product...

Alternative: HTCondor

- Open-source, community product → years of good support, responsive team, clear path
- Used by many WLCG sites and growing – RAL smooth migration seen as a booster
- Scaling seems not an issue – it allows for opportunistic expansion into a cloud/HPC
- WLCG exps. have expertise in using it, but local submitters in PIC to be re-educated
- Deployment plan: small portion for Grid submission (WLCG) → migration of the Tier-1/Tier-2 services → small portion for local job submissions → **users re-education**
- moving all of the resources to HTCondor

Testing of HTCondor in PIC



- Minimal **test instance deployed**
- Some (first) functional tests done:
 - local submissions from schedds
 - SCRATCH → /home/execute
 - Partitionable slots tests (# of cpus)
 - HA on collector/negotiator tested
 - CE's will act as schedds
 - Accounting groups (quotas) ok
 - Condor TAGS ok
 - Ganglia Monitoring ok

ARC-CE (*in progress*): installed & schedd – *it is being configured*

Next functional tests to come: Grid authentication on ARC-CE, remote job submissions, accounting, user priorities, fair-share, cgroups tests, ...

Agreed on regular meetings w/CERN. Other migrating sites welcome