
ARC-HTCondor incompatibility

David Cameron
On behalf of ARC team
WLCG operations coordination, 3.3.21

Timeline

- 2020-12-16: Issue of GDPR violation raised to EGI SVG
 - ARC publishing DNs in clear text through it's LDAP infosystem interface
- 2020-12-17: Patch to schema released for sites to deploy
 - This patch was subsequently retracted after finding it caused issues with WLCG sites
- 2021-02-11: ARC 6.10 released which hashes DNs
 - Found to be buggy and retracted
- 2021-02-15: ARC 6.10.1 released with bugs fixed
- 2021-02-18: Sites running ARC 6.10.1 report problems with jobs submitted by HTCondor
- 2021-02-24: Agree with WLCG to update pilot factories with patched HTCondor binary rather than undo GDPR fixes
 - For ATLAS this was done on 2020-03-02
- 2021-03-09: HTCondor official release with fixes to work with old and new ARC versions

Details of HTCondor incompatibility

- HTCondor submits jobs through the ARC's GridFTP interface
- It periodically queries all jobs belonging to the owner from ARC's LDAP interface
 - LDAP query with filter `nordugrid-job-globalowner="/DC=CH/DC=cern/..."`
- Thus job submission works but it cannot find the status of the jobs

```
027 (21124897.000.000) 02/17 07:49:07 Job submitted to grid resource
GridResource: nordugrid ce01.gla.scotgrid.ac.uk
GridJobId: nordugrid xxx.yyy
```

...

```
029 (21124897.000.000) 02/17 08:04:11 The job's remote status is unknown
```

- The fix is to change the LDAP filter to an OR of the original DN and the hashed DN
 - Thus maintaining compatibility with old and new ARC versions

Notes

- ARC's web service interface also exposes the DN in the GLUE2 schema, but the https communication is authenticated and encrypted
- ARC clients do not use **nordugrid-job-globalowner**
 - Submission systems such as ARC Control Tower and DIRAC are not affected
- The medium-term plan is for HTCondor to switch away from legacy LDAP/GridFTP to ARC's new REST interface
- In WLCG all jobs are submitted by robot proxies
 - No personal information has ever been exposed so no impact from GDPR issue

Lessons learned

- More testing!
 - It is impossible to know what every user does but HTCondor is a major client and so should be tested when “API” changes
- HTCondor submission tests have been added to ARC’s CI/CD testbed
 - An ATLAS Panda queue has been created to submit HammerCloud jobs from HTCondor to CEs running the latest ARC nightly release
- People from different communities can fix problems together with calm polite discussions :)