



ALICE

A Large Ion Collider Experiment

European Organisation for Nuclear Research



# AliEn status report

Miguel Martinez Pedreira



# + APIs

- dcache issue
  - having lfn-like pfn
    - `root://srm.ndgf.org:1094//alice/cern.ch/user/a/alitrain/PWGJE/Jets_PbPb_2011/104_20140609-0951/lego_train.C`
    - `root://srm.ndgf.org:1094//alice/disk/14/41166/ea0fce6a-e98a-11e3-abef-c7fc858f3c77`
  - thought to be on new api only, because of new envelope creation
    - but found also on addMirror commands on original user api
  - didn't find anything wrong in envelope creation
    - so debugging the call itself -> alien\_cp
    - difference on treatment on 1<sup>st</sup> and 2<sup>nd</sup> replicas
      - URL != TURL
      - addMirror with wrong parameter
      - either modify alien\_cp or patch the mirroring (xurl)
- G tables maintenance as usual (+2.X M/day)

# + APIs

- Close SE in env causing wrong SE selection
  - leaving `alien_CLOSE_SE` unset
  - autodiscovery works OK now
- Command fixes
  - `ps` after long-time broken
  - `masterjob`
  - trying to go over more commands
    - feedback appreciated 😊
- Merged version (see later) looks stable in `api08`
  - `access` and `tables` used correctly compared to `api03`
  - should we put it there also ?
    - easy to revert...

# + IPv6

- AliEn (v2-21 branch) built with PERL 5.18.2
  - using the buildsystem is not trivial
  - some packages had to be updated or replaced
- Some changes needed in the code
  - deprecated operations (e.g. defined)
- Even if new Perl is IPv6 ready...
  - Socket libraries have to be added (IO::Socket::[IP | INET6])
  - Checking network-based packages to see if they are ready
    - Some use INET6, others IP, some try both, others none...
- Good news, all services running in local VO
- Still To-Do
  - make sure network packages are updated and using ipv6 -> rebuild
  - test with IPv6 stack only to make sure it works
  - xrootd 4

# + AliEn code unification

- Merged version
  - all repos/scripts unified
  - svn repo added
  - running on all kind of services, stable
  - tested also running local site + basic jobs
- Pending...
  - put it in CVMFS
    - when done and working, only one version used everywhere
  - new installation being fixed
    - db tables / structure differences, initial values, services creation...
    - rebuild needed
      - MySQL, lots of modules, perl update, new modules needed, etc...

# + ERROR\_SPLIT issue

- Some jobs failing to SPLIT for apparently no reason
- Found issue when updating the status in the DB
  - InnoDB Deadlock
  - Maybe increased due to high amount of jobs lately?
- Two solutions added
  - Fix on the DB interface (transaction retry)
  - Enabled flag in MySQL that avoids? this
- This flag brought another issue
  - DB binlog growing too fast (JDLs mainly, see later)
    - using cache in some places to avoid calls
  - Finally flag disabled, the first fix seems to be doing the work



# ZIP 64

- Archives > 4GB truncated currently
- Added support to zip bigger things
- Using `Archive::Zip::SimpleZip`
  - adds the needed features
  - but keeps a similar use and options like the previous zip library
- Unzip with `IO::Uncompress::Unzip`
  - again, few changes to code thanks to similar interface
- Tested successfully
  - manual archiving from the prompt
  - job output archives
- Thanks to M. Tahawi (Summer Student)



# OCDB -> CVMFS

- Until now, manual process
- All registered files in AliEn checked
  - Catch if is an OCDB object
  - then processed by an optimizer
    - creates folders hierarchy like in CVMFS, perms...
    - uploads them to CVMFS
- OCDB catalogue in sync without problems since this tool is used
- Permissions issue
  - have to be properly set before upload
- Checks number of times a file was tried to be uploaded and when, filesystem issues, logs, etc...



# + JobAgent

- Bug in JA time-to-live logic
  - the time left was miscalculated and making JAs stop much earlier
    - and also implying that longer jobs had less chances to be picked up
- Jobs with environment variables set on JDL
  - and not cleaned up between executions
  - first try: saving/restoring the full env of the jobagent was causing troubles
  - unset only the variables set
- New feature, used ?
  - JA control process checks for 'traceLog' files
    - and put messages contained on those files in the job trace
    - Asked to add this in order to debug jobs (to find problematic stage)
- Spy
  - first sending command to node
    - was causing a die thus proxied call to CM wouldn't execute
  - not on all versions anyway...



# What next

- CVMFS package listing simplification
  - Site JDLs
  - But specially broker matching
    - needs to be smaller, maybe just a flag
  
- Job JDL DB simplification
  - JDL tables are big due to the text they contain
  - misused in some cases
    - origJdl vs resultsJdl
    - Masterjob vs Subjobs
  
- DB (Catalogue) optimization
  - load of recent days brings up again the issue
  - idea of using slave to load the balance (for read)
  - why not making the whole DB transparently HA, LB ?
  - also queries done from AliEn can be improved
  
- JSON
  - again rebuild issue
  - code should be fast

# + What next: jAliEn

- Currently most core operations implemented
  - used actively by Monalisa
    - PFNs, GUIDs, LFNs, jobs...
- TJAlien plugin for ROOT in good shape
  - developments during the summer
  - main commands implemented
- JobAgent from PERL to Java
  - so we can have the full chain, even if some parts depend on perl
  - and see how the plugin works
- Other central parts missing (Optimizers, Broker, Manager...)
- Dealing with packages, dependencies, easier
  - just needs updated JVM





# Questions

- Record of 69586 running jobs 😊
- Questions?