



ALICE

A Large Ion Collider Experiment

European Organisation for Nuclear Research



AliEn status report

Miguel Martinez Pedreira





Infrastructure

- New catalogue machine
 - SSDs → performance significantly boosted
 - Soon similar move on TaskQueue

- Updated MySQL version based on system (Costin)
 - Very few changes to [j]AliEn
 - Changes on some behaviours
 - Timeouts, config items, backup method ignored, even readline usage...

- Updated LDAP version based on system
 - Replication mechanism also updated
 - From slurpd to syncrepl
 - Quite some changes in config



New | | Fixed

- TLS communications
 - Authen's configured and running (rest ASAP)
 - After patch in AliEn
- Sending batch job parameters to MonaLisa
 - Helpful for site admins to debug jobs, linking queueId with batchId
- Pending: using latest as default AliEn version
 - apparently wrong env set
 - With libgfortran as red herring? under debug...
 - Some ideas in this sense: deleting the libraries folder (x86_64...) coming from build (is not there in classic AliEn installation); use alien-installer;
 - system gfortran should be used anyway
 - Then, will need to be debugged on the real setup...
 - As said in T1/T2 workshop:
 - New OpenSSL, Apache, MySQL, libs, ...

+ New | | Fixed

- Perl defuncts on WNs
 - Fixed, verified by sites
 - And about security: no more vulnerabilities reports means no vulnerabilities [found]...?
- Some DB queries optimized
 - Job & Catalogue quotas
- Stopping jobs with missing packages
 - Comparison AliEn vs CVMFS
 - See later slide about PackMan
- Brokers issue (~mid June)
 - 2 problems at once: job with wrong (long, repeated) set of packages led to discover length of packages column for JOBAGENT was not enough in some cases
 - painful debug session, brokers missfunctioning, slow+bad matching!

+ New | | Fixed

- Job TTLs required to be $\leq 24h$
- Big traces (from catalogue creation failure) default off
- ClusterMonitor startup
 - Recreation of httpd/conf automatically
 - Renaming old to httpd/conf.<timestamp>
 - Can be avoided touching '\$ALIEN_HOME/httpd/keep'
 - Checking if previously running, proper kill
 - Storing reliable PID to stop/start
 - Was failing in some cases...



Optimizers

- Job optimizers probing+forking now works
 - DB handlers were in bad state after forking and reusing
 - First attempt: just create a new one per child
 - Not working smoothly
 - Second attempt: clone parent handler + flags
 - InactiveDestroy
 - No problem on forking and getting proper results since then!

- OCDB→CVMFS sync
 - Being done catching the files under OCDB folders
 - Failures in CVMFS caused unsync because of faulty publications
 - Using ssh script: will CVMFS provide an API ? Or a way to detect failures
 - We can parse output of the script even, though is not so nice
 - Any 'fail' or 'code' always coming back when failed that we can use ?
 - Or, if there is absolutely no way from CVMFS side, just adding more work for AliEn/Catalogue side ? E.g.: the process tries to stat the files published after ~2h and see if they are there...
 - What if not there ?



Optimizers

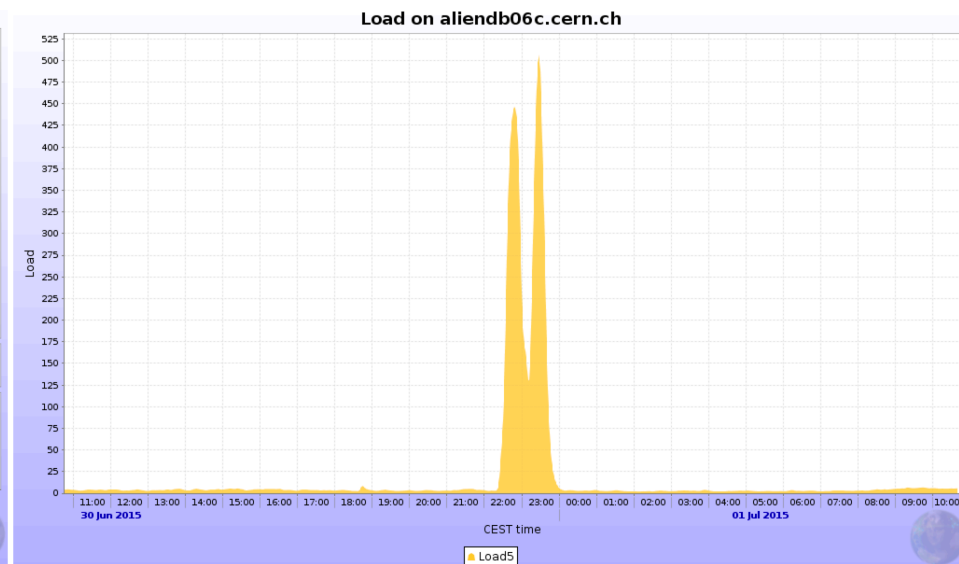
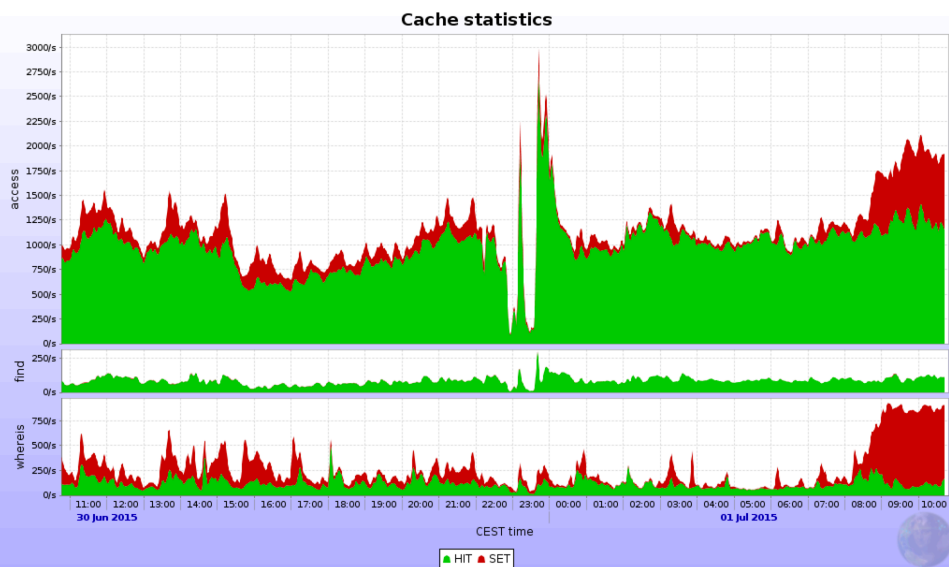
- We need more monitoring for the Optimizers
 - Right now, checked through the parent → not enough!
 - A possibility is to use PIDs of the processes and subprocesses and establish status and thresholds checks
 - Depending the type of optimizer
 - ML offers something that fits ? -> ApMon parameters + notifications from ML
 - Or AliEn itself, cronjob, or similar
 - jAliEn re-implementation

+ PackMan + Packages

- PackMan issues last week (ALICE week) and ongoing
 - Optimizer doing registration in the DB dies unexpectedly
 - Specifically doing find over /alice/packages
 - New API for publication coincides with start of problems
 - But no relation found so far...
 - In any case, the process should not die
 - Extra debug plus DB handler creation put in place as first approach
- In general, what do we want?
 - We have different list of packages in CVMFS and PackMan, they are not the same (the used ones I understand yes)...
 - Which has preference?
 - E.g.: ML doesn't allow submission if not visible in CVMFS. AliEn allows submission if existing in PackMan (all packages). Puts to error by an optimizer with a custom PackMan (CVMFS right now) after comparing availability of packages.
 - PackMan needed to keep flexibility (be able to register and run anything anywhere, as was done before)

+ Catalogue

- Catalogue load spikes (recent)
 - Coming from OCDB object finds from jobs
 - The TextCache usage practically stops
 - Fixed by splitting latest OCDB table



- Also lately jobs thrashing with too many queries (event loop?)
 - We need more capacity [and more control? E.g. call hz]



Operation in GSI

- Xrootd proxy working ok so far...
 - Few fixes/hacks
 - Going up to 1K jobs for now

Services status

AliEn: v2-19.276

ClusterMonitor: **OK**

PackMan: n/a

CE: **OK**

CE info:

Max running jobs: 1000

Max queued jobs: 50

Proxies status AliEn proxy: **OK** (1 day, 23:58)

Delegated proxy: n/a (n/a)

Proxy server: n/a (n/a)

Proxy of the machine: n/a (n/a)

Current jobs status

Assigned: 0

Running: **964**

Saving: 7

Accounting

(last 24h)

Success jobs: **2467** (profile)

Error jobs: **221** + **4** expired

kSI2k units: **2356** / pledged

Site averages

(last 24h)

Active nodes: 308.9

Average kSI2k/node: 2.456

-
- SGE interface being completed+improved
 - Missing methods
 - To report proper status of the batch system
 - Running in production partition now

+ EOS

- ORNL + Kolkata new sites
 - ORNL is functional (used within AliEn) but being customized
 - Kolkata had initial setup as test, still pending
 - Will have MGMs and FSTs together
- New eos-deploy
 - Can setup any combination of MGMs and FSTs
- New wiki with installation and configuration steps, testing and other notes
 - http://alien2.cern.ch/index.php?option=com_content&view=article&id=140&Itemid=144

+ Questions / Comments

- Record of 82K+ running jobs 😊
- We still miss a real test environment
 - Even more important now with more jAliEn dev
 - [partially] mirroring DBs + LDAP is an option
- Questions?

+ jAliEn (+Pavlo's talk)

■ JobAgent

- Implementation up to preparation for execution (30-40%?)
- Payload will run in a new process
 - Easier and more precise to control than threads
 - More isolation
- A JBox will run per JobAgent
 - Bidirectional communication

■ Packages

- PackMan basic interface + CVMFS implemented

■ Broker (matching)

- Job matching + site probing for submission ready

