

# AliEn status report

Miguel Martinez Pedreira



# Touching the APIs


- Bug found, not sending site info from ROOT to central side
  - was causing the sites to have network overloads and reduced efficiency
- Started to touch the ‘untouchable’ unmaintained code
  - The bug fix implied having to modify the access methods
- Found tricky code
  - hardcoded stuff for SEs
  - redundant calls
  - different use of cache and database tables in the different apis

# + Touching the APIs

- First step, adapt the code to process the site and attempt number from jobs read requests
  - also reading from the right table (SEDistance)
- In the case of APIs specially, we need somewhere to test
  - managed to create a new one in pcalice92
  - soon after, used new server to add a user api: apiserv08
- Spotted a part of the code to select SEs based on 'whereis'
  - but then almost same 'whereis' repeated
  - idea to cache them (heavy operation)
  - idea to sync it with job optimizers
    - since jobs are splitted based on the inputdata the jobs request to read
- Reordered cache usage
  - access the same
  - envelope was misused (caching the same information as access)
  - added whereis, same as in optimizer

# + Touching the APIs

- Result: we do less 'whereis' calls and use the cache better



Access cache access statistics					
Series	Last value	Min	Avg	Max	Total
1. HIT	139.8	0	189.9	1268	498451121
2. SET	160.2	0	96.4	842.1	253076986
<b>Total</b>	<b>300</b>		<b>286.3</b>		<b>751528108</b>

whereis cache access statistics					
Series	Last value	Min	Avg	Max	Total
1. HIT	151.8	0	53.98	3361	141703422
2. SET	282.7	0	61.34	780.8	161026230
<b>Total</b>	<b>434.6</b>		<b>115.3</b>		<b>302729653</b>

```

whereis : 1500000 (min: 1272, avg: 1570, max: 2868, total: 2.194 GB) : 5138099 hits
findwait : 1387 (min: 12, avg: 1511, max: 35568, total: 2 MB) : 959296 hits
jobbroker : 1 (min: 95, avg: 95, max: 95, total: 95 B) : 215 hits
access : 506829 (min: 12, avg: 2631, max: 3823, total: 1.242 GB) : 744581 hits
find : empty
envelope : empty

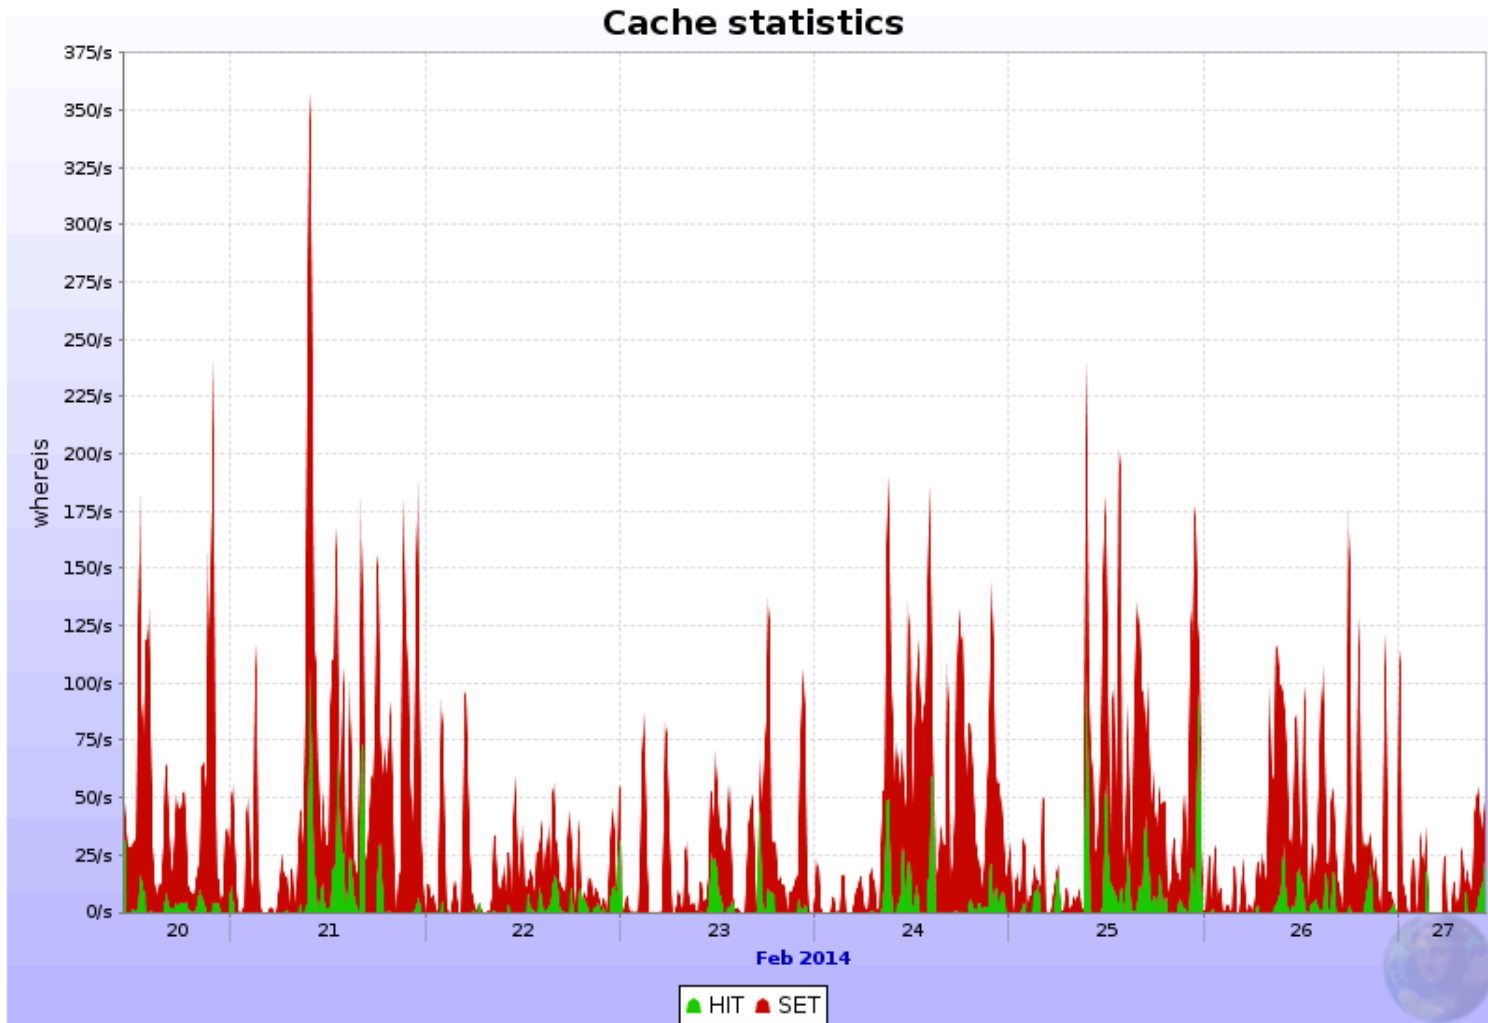
```

```

Java memory stats: 32 GB total memory, 32 GB max memory, 11.62 GB free
Java version: 1.7.0_55
Uptime: 43d 5:11

```

# + Touching the APIs



# + AliEn code unification

- Playing with the apis code raised the issue of the AliEn versions again
  - decide to start merging
  - sync of SVN and central services
    - more differences than expected...
- Initial status
  - v2-19 – CVMFS: voboxes + wns
  - CS: v2-19 + v2-20 (TQ)
  - Job APIs: v2-19.112, shared for 10 API servers
  - User API: v2-19.112 + patches, only api03

# + AliEn code unification

## ■ 1. CENTRAL – CVMFS

- replace site-side parts into the CS installation
  - JobAgent, ClusterMonitor...
    - checked all files anyway
- created alien.NEW
- put it on some production services gradually

## ■ 2. Jobs – Users APIs

- differences in access code and some manual patches from api03
- new version, alien.219\_API, put in job apis
  - small issues forming the envelope, coming from whereis result

# + AliEn code unification

- 3. APIs + CS + CVMFS
  - specially important authen, access, admin, user commands...
- Finally the one and only version! ☺ running:
  - 1 user API: apiserv08
  - Authen in db2, JobBroker, JobManager, JobInfoManager, IS in db8
  - Progressively to the rest
- What now ?
  - To be put in CVMFS
    - new version to be used/tested explicitly first?
  - SVN? (Name? now alien.FINAL)
  - Scripts, installation
  - To be fully tested...
    - But quite smooth so far
  - Differences in installations, also affect behavior
    - Fresh installation where we put the new code





# Certificates

- In the last months, several problems to access the GRID by several users
- Missing/outdated certificates in CS and/or API
- Have to add them manually
  - not updated installations
  - not clear what has to be there
    - just IGTF package?
- Automate: cronjob or tool to update some parts of the installations
  - it exists for the CVMFS one



# Other items

- SPLIT jobs not MERGING
  - JOBSTOMERGE now correctly updated
  - jobs splitting into 0 subjobs now to error
- Fix for ZOMBIEs
  - race condition between insertion-waiting and execution in the node
  - fix in a db field
- JA env cleanup between jobs
- CMreport sends more, bigger messages
- Proxy-init fix
- JA check for output size
- Option to disable catalog trace from LDAP



# Other items

## ■ dcache issue

### ■ having lfn-like pfns

- `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`

### ■ though to be on new api only, because of new envelope creation

- but found also on `addMirror` commands on original user api

### ■ under investigation

## ■ G tables maintenance

- with a 'high' estable number of jobs running -> +2M entries/day
- new table every month aprox

# + IPv6

- Starting next week
- First step: update PERL version in AliEn and see what crashes
- Update to xrootd 4.0.0 in API
- Student coming on 7<sup>th</sup> July to deal with this
- Test with IPv6 stack only to make sure it works



# Conclusion

- Still more things to do
  - JDL optimization in DB, Broker queries, improve commands...
- What to do with v2-20 and v2-21 ?
  - Catalogue conversion takes long
- HLT/Cloud incoming? (Dario)

# + [Almost] Birthday

- About to reach 400.000.000 jobs 😊



- Questions ?