



# Operations

Maarten Litmaath  
CERN-IT-SDC

ALICE Offline week  
Nov 27, 2015

# Central service support



- Keep them running OK
  - Occasional high loads on DB, API servers or network
    - E.g. due to data replication fallback via our hosts
  - Air conditioning incident
- Improve the SW
  - For example, v2-19-276 has become the default
    - Requiring RFC proxies on the VOBOX
  - See talk by Miguel

# Site support



- Alert and/or assist sites having issues concerning their VOBOX, CE, WN, or SE
  - Noticed in MonALISA and/or reported by users/admins
- Help new sites get started
  - Bye-bye and thanks to LLNL, welcome ORNL !
  - Cibirong and Bandung are back
  - Wuhan also back, but suffering poor connectivity
  - Tsukuba in progress
  - HPC → see talk of Andrey K + Pavlo, and talk of Markus
  - Clouds → see page 9
- Assist sites with changes
  - SE additions, upgrades, migrations, cleanups
    - See talks of Miguel, Andrey Z, Luciano
  - Transition to HTCondor CE
    - Thanks to Pavlo

# User support

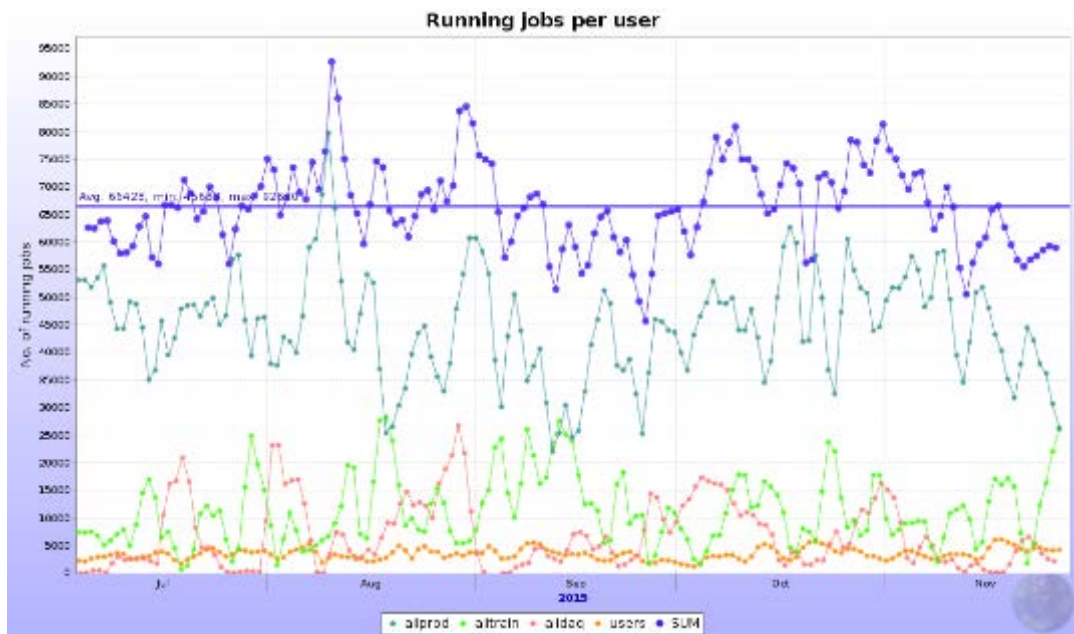


- “Cannot get token!”
  - Often due to issue with user machine or network
- Mistakes in JDL
  - Bad TTL, inconsistent SW versions, ...
  - AliEn now **catches** some of these
- Segfaults / memory leaks not seen in a quick test
  - Uninitialized variables, undeclared transients, double deletes, insufficient cleanup, ...
- Problem with installing AliEn, AliRoot, ...
  - Dario’s **exact recipes** usually are the answer
  - Giulio also got into the game 😊
- Many **experts** on the **analysis task force** and **operations** lists help other users → **thanks!**


# High activity



- 97k concurrent jobs briefly reached on Aug 10
- 90k+ reached a few more times
- Taking advantage of increased + opportunistic resources
  - CERN lifted 15k LSF cap → usually 20k+ jobs since months
- Large analysis component in preparation for QM 2015



66k average Jul 3 – Nov 26



# CASTOR



- Various incidents led to a reorganization for HI data taking:
  - 1 “fast” pool for **DAQ + writing to tape**
  - 1 “slow” pool for **reco, replication** to T1 sites, **recalls** from tape
  - New files in pool 1 **automatically** get replicated to pool 2
  - **Many disk servers** that were prepared for retirement were temporarily **added**
  - **Enough space** to keep all 2015 HI data on disk until early next year
  - **Thanks** to the CASTOR and Fabric teams!

# Heavy ion reconstruction



- Important changes in the code and workflow have been implemented to **reduce the memory usage**
  - Tested with 2011 heavy ion reference data
  - If all goes well, for this year's heavy ion data the reco jobs will only need ~2.5 GB RAM
- To be on the safe side, **special arrangements** were made with the sites that will receive HI raw data
  - **CNAF, FZK, KISTI** and **SARA** have set up dedicated high memory queues
  - At **CERN** the jobs can request 2 cores and hence have twice the memory
  - All setups were tested with normal jobs
  - **Thanks** to the sites for their good support !
- **Normal queues will be used wherever possible, depending on the actual memory needs**

# Middleware



- **SL5** will be deprecated for MW next spring
  - Already there have been openssl compatibility issues
    - Will only get worse
- **CentOS/EL7** more important next year
  - **SL6** remains the default
- HTCondor batch system becoming more popular
  - Now also comes with its own CE
    - Already used a lot in OSG, now also at CERN
  - The ARC CE remains an alternative



# Birds and clouds



- A new AliEn job submission module for the HTCondor CE is being developed by Pavlo
- To be used at CERN and other sites moving to such a setup
  - Some code improvements still in progress
- Also used for job submissions to a special CE in front of commercial cloud resource projects from CERN-IT
  - This way AliEn just sees another “site” with a standard CE
  - The cloud management aspects are handled behind the scenes by IT department experts
- The *Altaria* site will be used for that
  - Located on Eyjafjallajökull ! 😊



# SAM



- New Availability / Reliability profile based on selected MonALISA results remains not yet critical
- Its introduction was delayed because of issues with the Dashboard GUI that needed to be fixed
- We intend to give the new profile a try next year
- Sites will first be alerted to compare old and new profile results
  - SE test failures will reduce the A / R!
- Test job submission to the HTCondor CE still to be added

# Xrootd



- Sites need to start planning upgrades to Xrootd  $\geq 4.1$ 
  - Stable and well supported
  - Required for IPv6 support
    - Will become more important next year
- T1 sites are particularly concerned
  - They receive raw data exported from T0
  - Latest CASTOR incompatible with old xrd3cp
    - Gateway hosts are being used as a temporary workaround
- Communication via LCG Task Force list as usual for expert advice
- SL6 hosts now can have xrootd for ALICE installed through rpms!
  - <http://linuxsoft.cern.ch/wlcg/>
  - Thanks to Adrian Sevcenco!

# Networks



- Successful workshop on network evolution in Asia Sep 22-24 at **KISTI**
- All ALICE sites in Asia were represented
- **LHCONE peering** was agreed between network infrastructure providers **TEIN**, **KREONET2** and **ASGC**

# Sites mostly OK – thanks!



## However...

- VOBOX or WN **issues**
  - CVMFS problem, CE not ready for jobs, myproxy running low, ...
  - Absence of “system” library
    - HEP\_OSlibs rpm helps avoid that
- Files **unavailable** due to SE problem
- Admins please check site issues page
  - <http://alimonitor.cern.ch/siteinfo/issues.jsp>
- **Subscribe to relevant notifications**
  - <http://alimonitor.cern.ch/xml.jsp>