



WLCG operations

ALICE T1-T2 Workshop
UPB, Bucharest
May 14-16, 2019

Maarten Litmaath
CERN-IT

v1.0

Central service main issues



- August
 - Task queue DB HW issue on a Sat afternoon
 - Fixed late evening thanks to heroic efforts by Miguel and Costin
 - Fallout from CA reorganization in AliEn
 - Some new CAs are incompatible with the old OpenSSL in AliEn
 - Several cache crashes required manual restarts
 - Due to huge query rates from badly behaved jobs
- October
 - Task queue host erratic I/O performance
 - Task queue had to be moved to another machine
- November
 - Instabilities due to raw data replication traffic
 - Cured by various improvements

Site evolution



- University of **Derby**
 - PoC cluster of 120 cores since June 19
- New site **ICM** (Warsaw) being set up
- New or enlarged EOS instances at several sites
 - The third [EOS workshop](#) was held Feb 4-5
- Growth – 2018 and 2019 pledges
 - More disk space is welcome in particular!

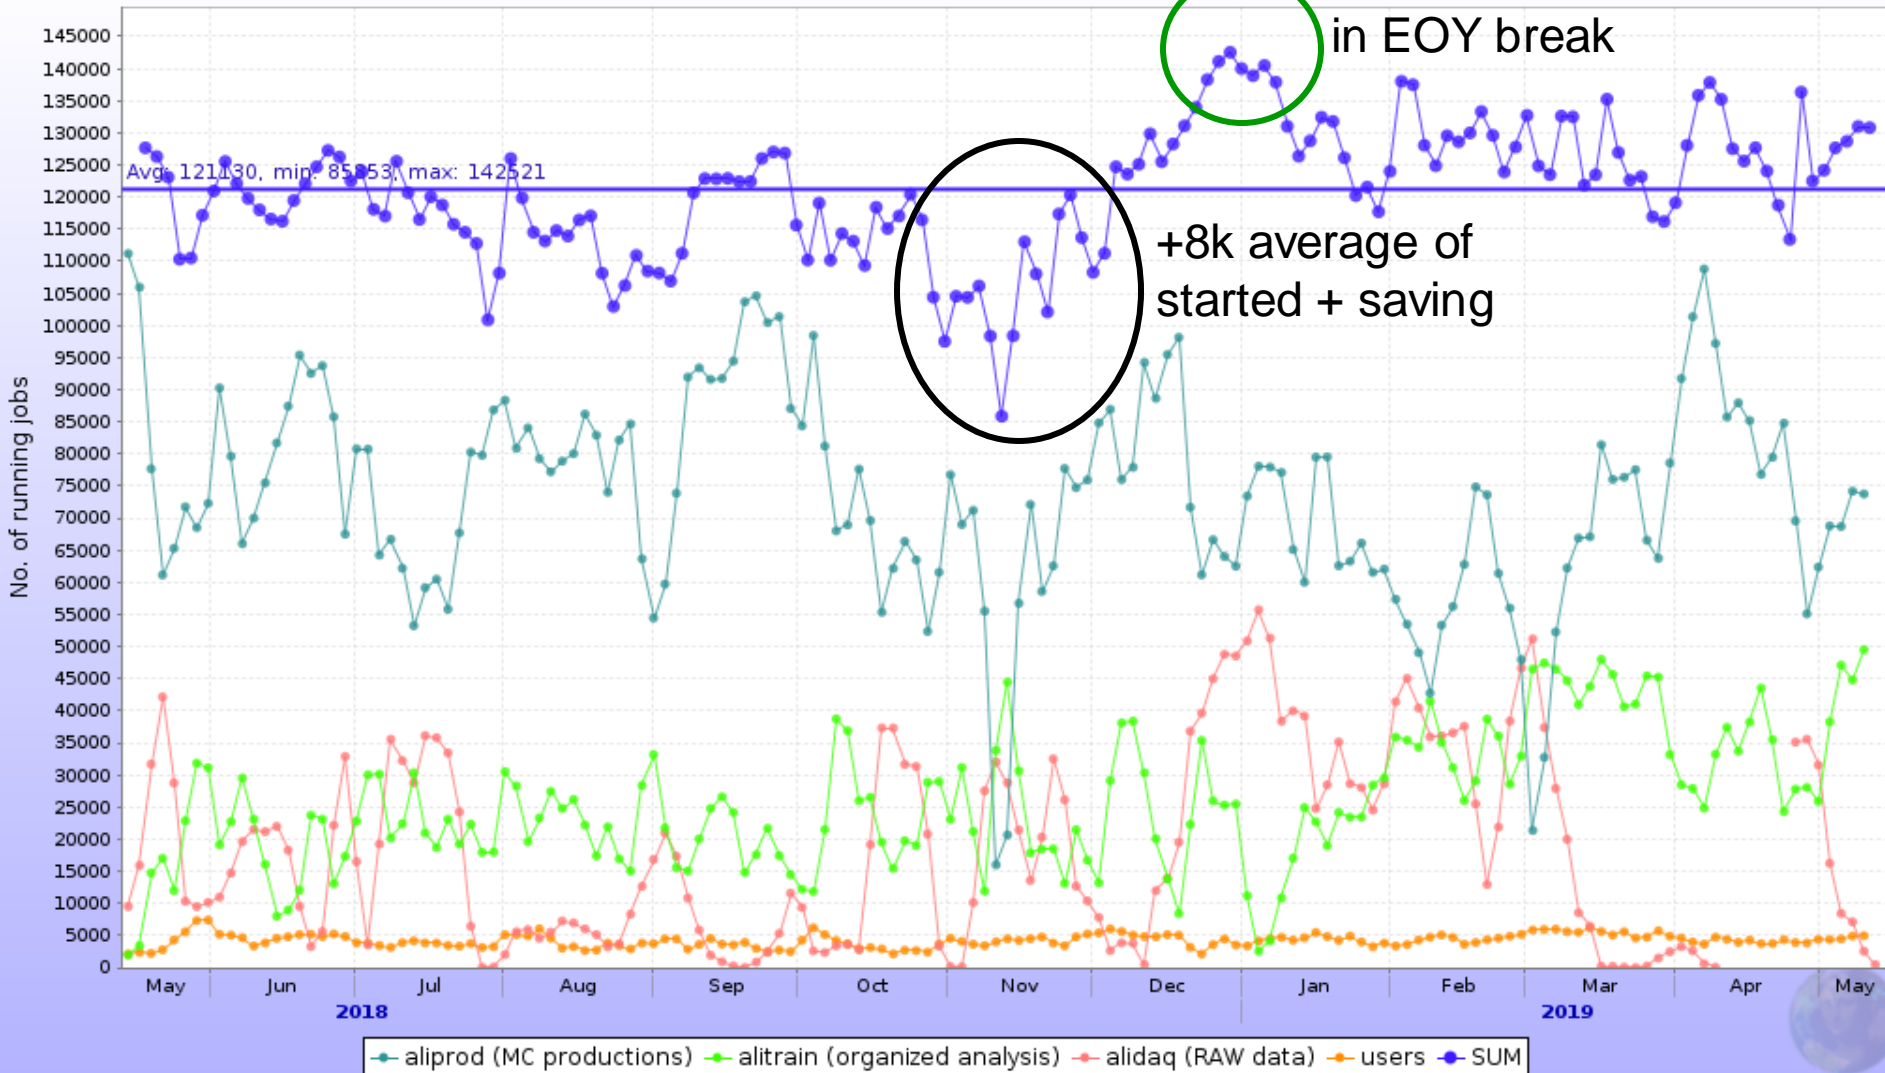
High activity



Running jobs per user

150k reached
in EOY break

+8k average of
started + saving



In time for SQM 2019, June 10-15



Storage



- Heavy-ion data taking
 - Sufficient disk space on *eosalice* to allow all data to be kept on disk for reconstruction
 - Export to the T1 sites and copy to CASTOR completed during the end-of-year break
- CERN
 - EOS name space upgraded to *QuarkDB*
 - More scalable & robust

Issues at sites or with jobs (1)



- IN2P3-CC: 110 TB lost due to RAID problem
- WN had to be rebooted everywhere to fix *L1 Terminal Fault* kernel vulnerabilities
- Occasionally high error rates due to bad jobs
 - Fallout at sites will be limited when *containers* are used
- Clermont unusable for 1 month due to huge failure rates
 - Weird errors that finally were gone unexplained
- Prague unusable for 2 months due to network issues
 - Recovered thanks to a big effort of the Prague team!

Issues at sites or with jobs (2)



- CERN: 20k “ghost jobs” discovered on March 12 and another 10k on March 20
 - Not visible in MonALISA
 - Several example jobs found stuck in *gdb*
 - During ~1.5 days the number of such jobs steadily shrunk to zero
 - Root cause was not found

Middleware



- **CentOS/EL7** has become the default
 - SL6 will be supported (critical fixes) until Dec 2020
- **CREAM CE** supported until **Dec 2020**
 - EGI Ops and WLCG Ops will provide guidelines for helping sites replace CREAM CEs with suitable choices from a set of potential alternatives
 - **ARC** – already in use at many sites since many years
 - **HTCondor** – being integrated into the UMD
 - Solutions *without a CE* – if all your customers can handle that
 - See the **Future After CREAM-CE** sessions (first, second) held on May 7 at the EGI Conference and the May 8 GDB summary
- UMD only has officially supported products
 - Products like *Maui* remain available in the Community Repository

Containers



- WLCG Containers WG intends to provide common guidelines for Singularity deployment on WN
 - Driven by CMS and ATLAS use cases
 - Should suffice also for ALICE and LHCb
- As of RHEL/CentOS 7.6 there is a `sysctl` to enable *unprivileged* mount namespaces (not yet default)
 - That allows Singularity to be deployed and used *without the need for setuid root* → avoids being affected by yet more critical vulnerabilities in that package...
- For ALICE the use of containers by jobs is linked to the deployment of **JAliEn**
 - More in Maxim's talk

EU cloud projects



- Helix Nebula Science Cloud project has finished
 - Up to ~20k cores were available to the 4 LHC experiments *plus* other participating projects until Dec
 - Virtual sites *Capella* and *Regulus* were used by ALICE to target the 2 consortia providing resources
 - 400k MC jobs delivered
 - ALICE use cases served for WLCG/HEP TCO studies
- Non-trivial to configure/debug such resources
 - Network and other issues
 - Default easy use case: MC production
- As of Jan this year the next phase has started: Open Clouds for Research Environments (OCRE)

SAM



- **Reminder: SE test failures reduce the A / R!**
- Corrections are applied as needed
 - A policy defines thresholds to reduce the number of recalculations

Xrootd reminder



- Sites should continue upgrades to **Xrootd** **>= 4.1**
 - Almost all sites have done that already, thanks!
 - Required for **IPv6** support
 - EOS → upgrade to v4 a.k.a. *Citrine*
- Communication via LCG Task Force list as usual for expert advice
- ALICE add-ons are available through **rpms**
 - <http://linuxsoft.cern.ch/wlcg/>
 - **Thanks** to Adrian Sevcenco!

Networks



- WLCG **IPv6** deployment
 - Status recently presented in Apr GDB
 - ALICE storage majority available through IPv6
 - 73% at CERN + T1 sites
 - 46% at T2 sites
 - ALICE needs *~all* storage to be dual-stack before *IPv6-only* computing can be used
- 4th Asia Tier Center Forum, held Nov 20–21
 - Hosted by **SUT** – thanks!
 - Fostering network evolution in Asia – well done!

Tips for sites – thanks!



- Possible **issues** on VOBOX, CE, WN
 - CVMFS problem, CE not ready for jobs, myproxy running low, ...
 - Absence of “system” library
 - HEP_OSlibs rpm helps avoid that
- Jobs may fail due to SE **problems**
- **Admins please check site issues page**
 - <http://alimonitor.cern.ch/siteinfo/issues.jsp>
- **Subscribe to relevant notifications**
 - <http://alimonitor.cern.ch/xml.jsp>