



Grid status

ALICE Offline week
March 30, 2017

Maarten Litmaath
CERN-IT

v1.1

Central services

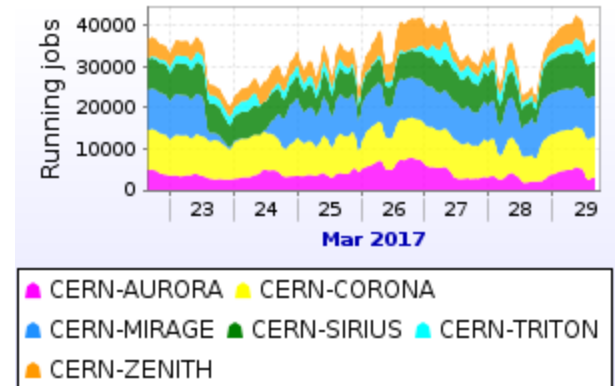


- Mostly stable, a few incidents
 - Jan 4 – yearly scheduled power cut lasted much longer than usual, way beyond the UPS battery lifetime
 - It took ~2 days to deal with all the fall-out
 - Jan 14 – AliEn CA certificates of various services expired
 - Now the CERN Grid CA is used, with automatic reminders
 - Feb 27 – Authn instabilities due to a flood of find queries hitting an ancient catalog bug
 - Remnant of a similar issue fixed last year
 - That same week:
 - Broken RAID controller slowing down the Task Queue
 - MySQL crashes due to a corruption – the master DB had to be restored from the slave
 - Week of March 13 – catalog unexpectedly became slow in the compression of various tables
 - Not clear what caused this or how it got cured exactly, but the configuration ended up better tuned as a result


Site changes



- FZK
 - Second virtual site for ARC resources
 - Original site also switched to ARC after decommissioning of the CREAM CEs
 - ARC CE performance issues under investigation
- Prague
 - Second virtual site for ARC resources
- CERN
 - Strong growth of the HTCondor cluster
 - Most LSF resources will be moved there
 - The CREAM CEs will be decommissioned
 - 40k jobs are becoming the new norm
- Altaria
 - Currently serving UK cloud sites in proof of concept
 - Details further in this presentation

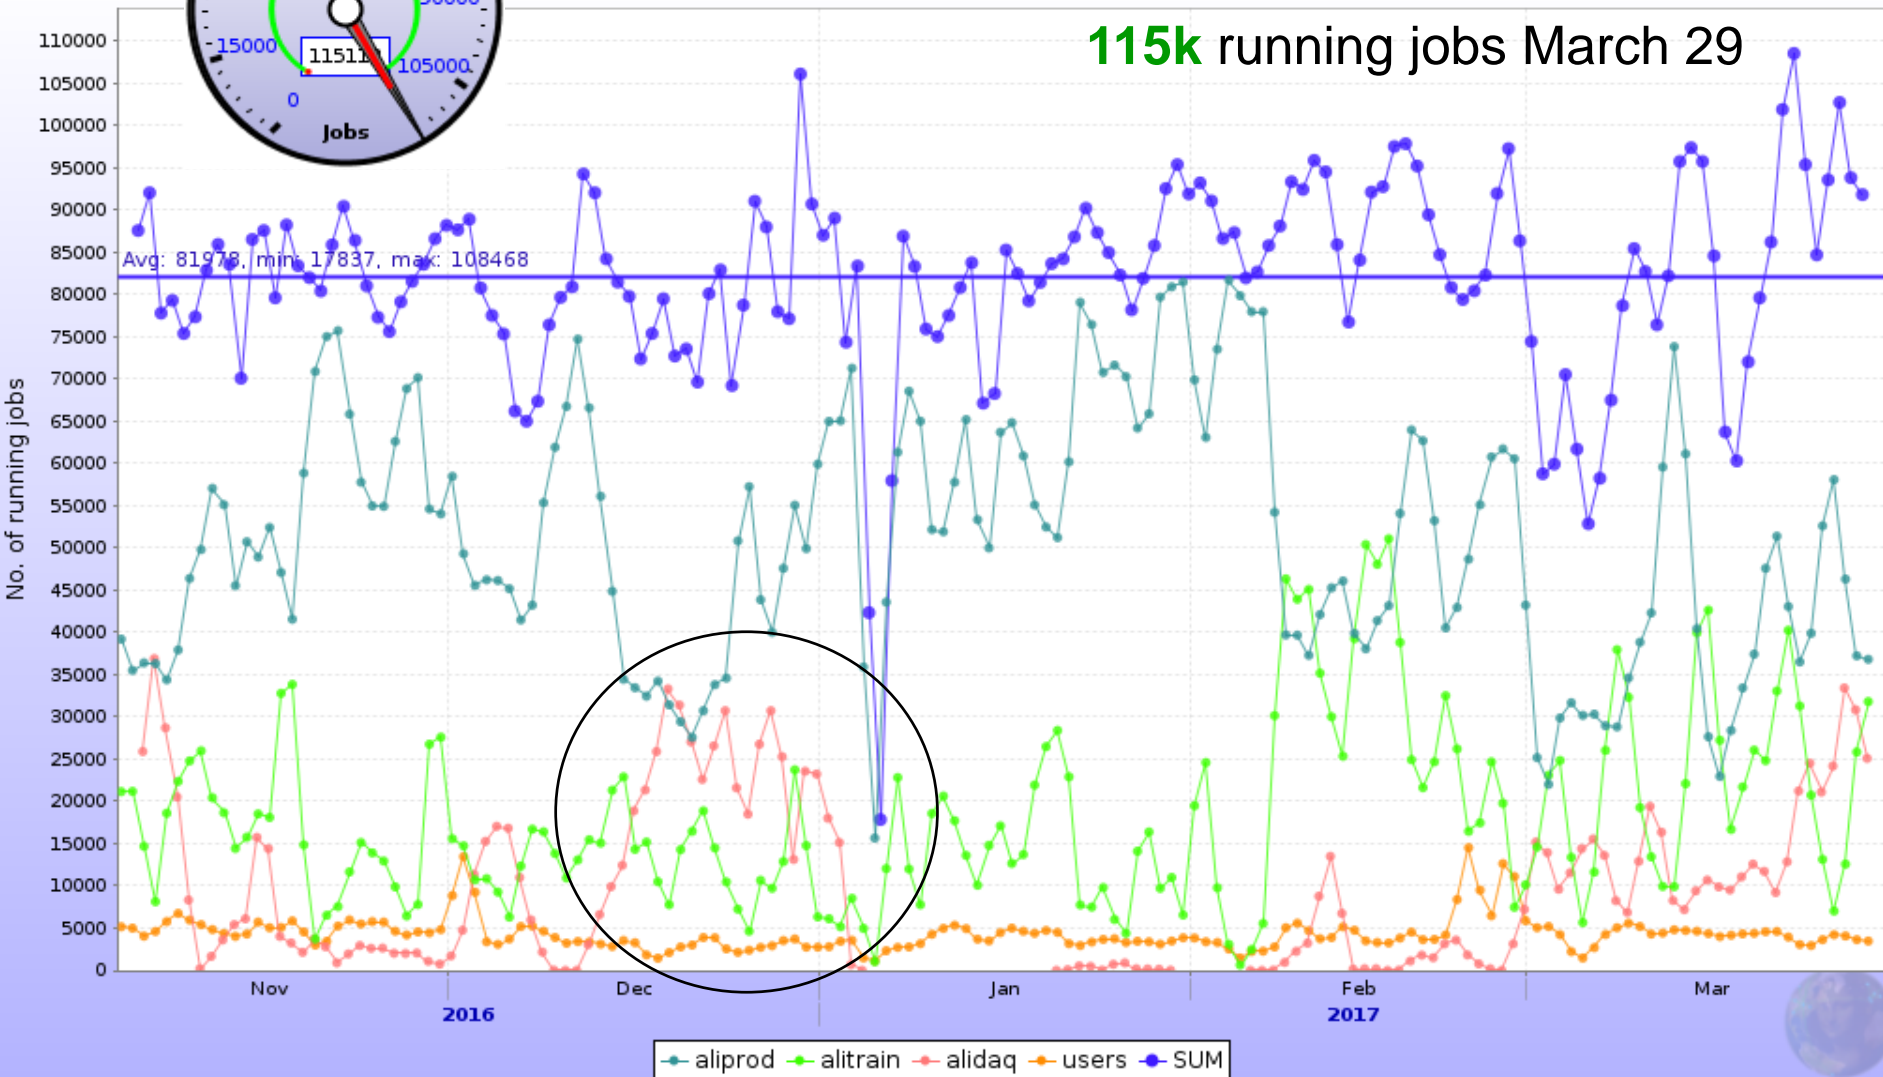


High activity



Running jobs per user

115k running jobs March 29



Most of the p-A reco done over the EOY break, in time for QM (Feb 5-11)



Storage at CERN and beyond



- CASTOR
 - Mostly stable, a few incidents
 - Recently overloaded by reco jobs at CERN *plus* T1 sites
 - Cured by limiting the concurrent transfers + request queuing
- EOS
 - Mostly stable, a few incidents
- Other sites ditto – thanks!
- Please consider **EOS** for your next SE solution
 - A first EOS workshop was held Feb 2-3
 - Check out the presentations!

Issues at sites or with jobs



- CERN: multiple incidents affecting the new HTCondor CEs
 - Currently the gateways to 50% of the resources at CERN
 - CE configurations were improved for robustness and scalability
 - Some bugs were encountered, worked around and investigated with help from the HTCondor developers
- Asian sites: more CVMFS issues involving the Stratum-1 at ASGC
 - Worked around by overriding the default configuration
 - The ASGC Stratum-1 will be removed from it
 - And by the new AliEn version v2-19.395
 - The Job Agent checks the CVMFS status of the WN to try ensure the user payload dependencies will be available
- High rates of expired jobs for a particular MC campaign
 - They were using much more memory than most jobs
 - Several sites complained
 - At CNAF we had to switch to the high-memory queue
 - With half the usual number of job slots
 - The campaign finished nonetheless and experts managed to reduce the memory consumption to normal levels for future cases

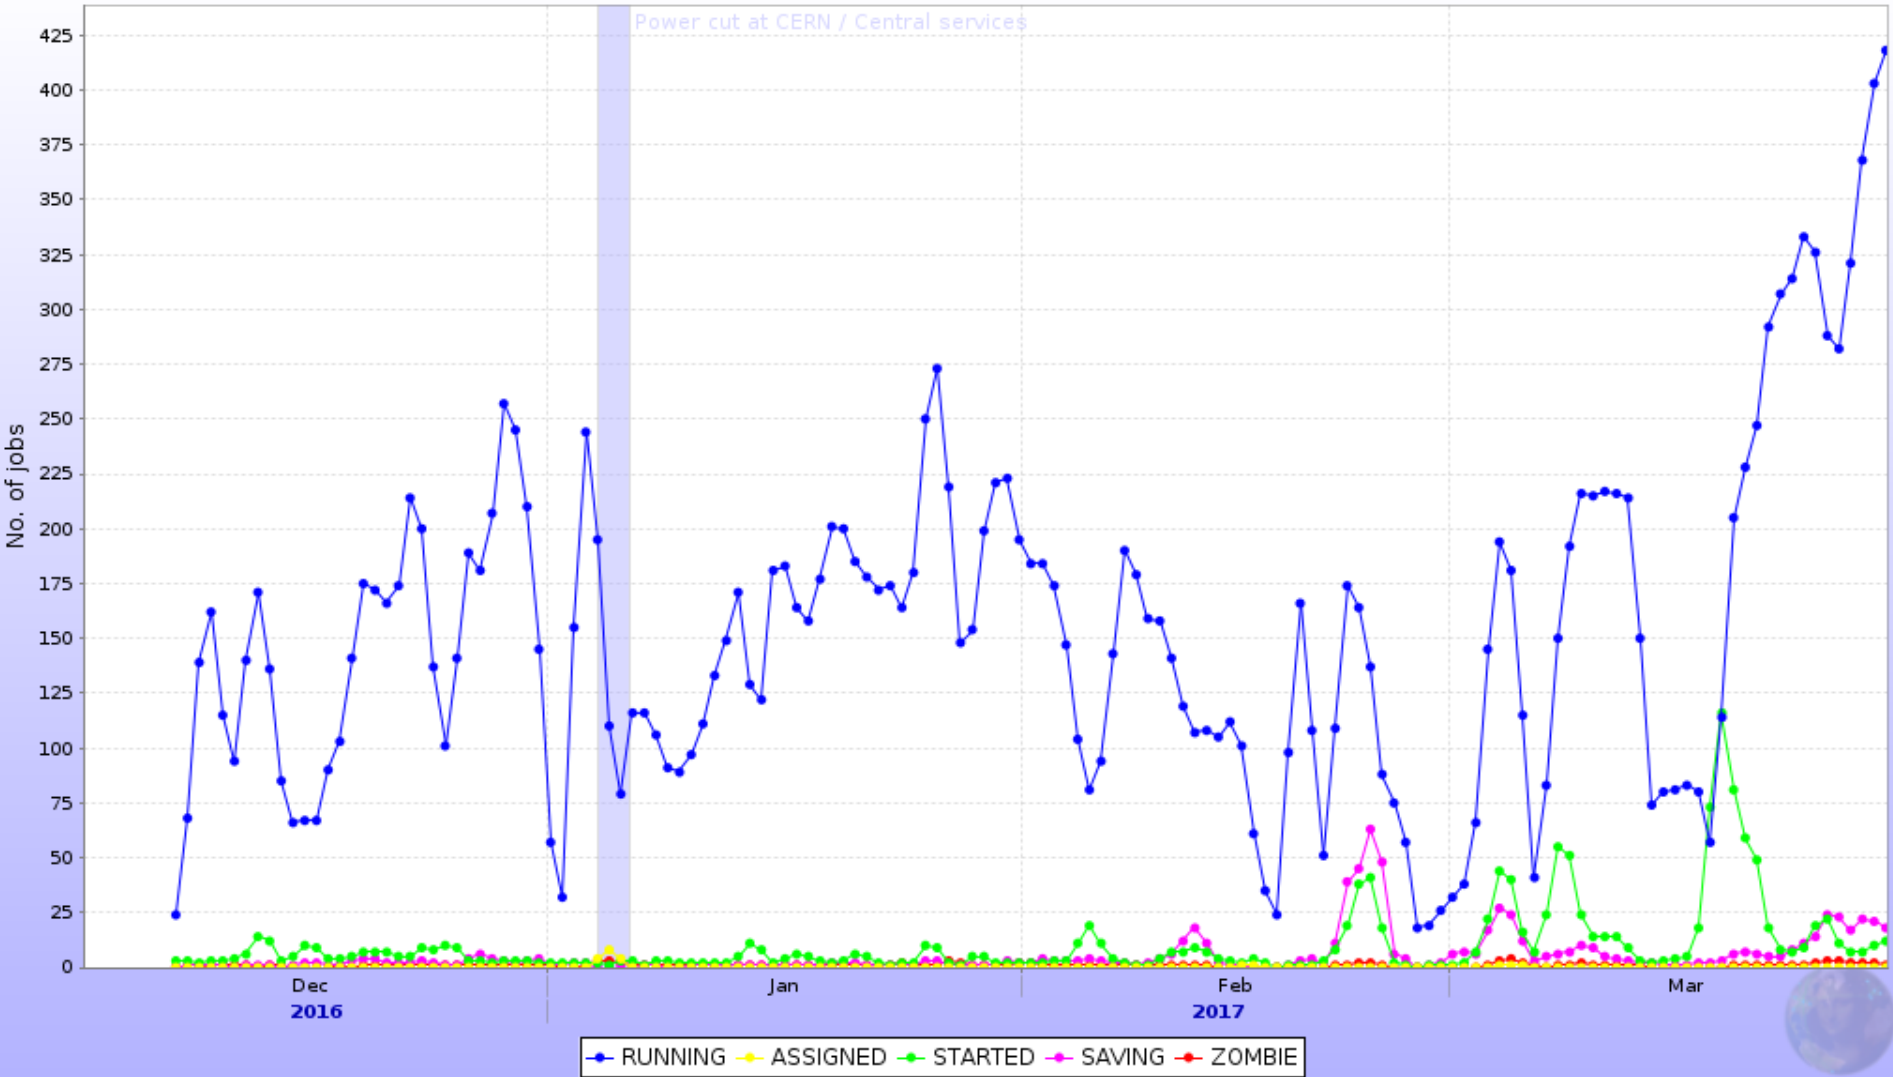
Altaria



- A virtual site to drive cloud resources
- Currently being used in a proof of concept for sites that want to provide their resources via cloud instead of traditional grid mechanisms
- In particular the UK T2 sites are moving to that model
 - Manchester (first ALICE jobs since 2008)
 - Liverpool (first ALICE jobs ever?)
 - Birmingham
 - Oxford
- Cloud VMs are configured such that they connect to an HTCondor pool at CERN to which Altaria submits its jobs
 - For monitoring and accounting it may be desirable to have an HTCondor pool per site, hosted on its own VOBOX
- The VMs are managed by the sites, AliEn just sees resources appear as if they were WN job slots
 - Managed e.g. through Vac or Vcycle



Active jobs in Altaria



Middleware



- **CentOS/EL7** is becoming more important
 - Some service types already available in UMD-4
 - CREAM, EMI-UI and EMI-WN expected by May
 - VOBOX will follow
- **SL6** still the default, but the experiments have been preparing for physical worker nodes running CentOS/EL7
 - Containers (or VMs) could still provide SL6
 - ALICE jobs have been tested OK on CentOS/EL7
- **SL6** in UMD-4 only has officially supported products
- **SL6** in UMD-3 may still be the easiest to use for certain services
 - Mind that new features only go into UMD-4

Containers



- An isolation paradigm much lighter than VMs
- A new tool to launch containers is gaining momentum in our community and beyond:
Singularity
 - US-CMS already use it at their T2 sites
 - Provide SL6 environment on CentOS/EL7 WN
 - And **isolate** each user payload from other processes → the plan is to let our Job Agents use it where available
 - No timeline yet
 - It currently needs to be *setuid root* because the kernel requires root for *mount namespace* operations
 - A future EL update (tentatively 7.4) is expected to lift that requirement

SAM



- New Availability / Reliability profile based on selected **MonALISA** metrics in use since May
- So far no big issues were reported
- **Reminder: SE test failures will reduce the A / R!**
- Corrections have been applied as needed
- Test job submission to the HTCondor CE has been added to the preprod instance
 - Production to be updated soon

Xrootd reminder



- Sites should continue upgrades to **Xrootd >= 4.1**
 - Most sites have done that already, thanks!
 - Required for **IPv6** support
- Mind that v4.6.0 has some issues (463, 465)
 - Should be fixed in v4.6.1
- Communication via LCG Task Force list as usual for expert advice
- ALICE add-ons are available through **rpms**
 - <http://linuxsoft.cern.ch/wlwg/>
 - **Thanks** to Adrian Sevcenco!

Tips for sites – thanks!



- Possible **issues** on VOBOX, CE, WN
 - CVMFS problem, CE not ready for jobs, myproxy running low, myproxy type wrong, ...
 - 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>