



Grid status

ALICE Offline week
July 20, 2017

Maarten Litmaath
CERN-IT

v1.0

Central services



- Mostly stable, a few incidents & maintenance
- In particular:
 - Apr 14-15: multiple MySQL crashes due to corruptions
 - Finally attributed to a broken expansion slot cache
 - Apr 19: upgrade of the database server running the file catalogue
 - New RAID controller and a set of disks were added to expand storage capacity and improve IO performance

Site changes




- CERN
 - Further growth of the HTCondor cluster
 - LSF resources still large
 - 40k jobs are becoming the new norm
- New “site” GRIF_IPNO_HTC
 - Batch resources moved from Torque to HTCondor
- Subatech_CCIPL ramp-up
 - 400 cores since late June
- Torino – opportunistic use of HPC cluster
 - Up to 1200 job slots since late June!
- LBNL
 - New proof-of-concept cluster for future migration

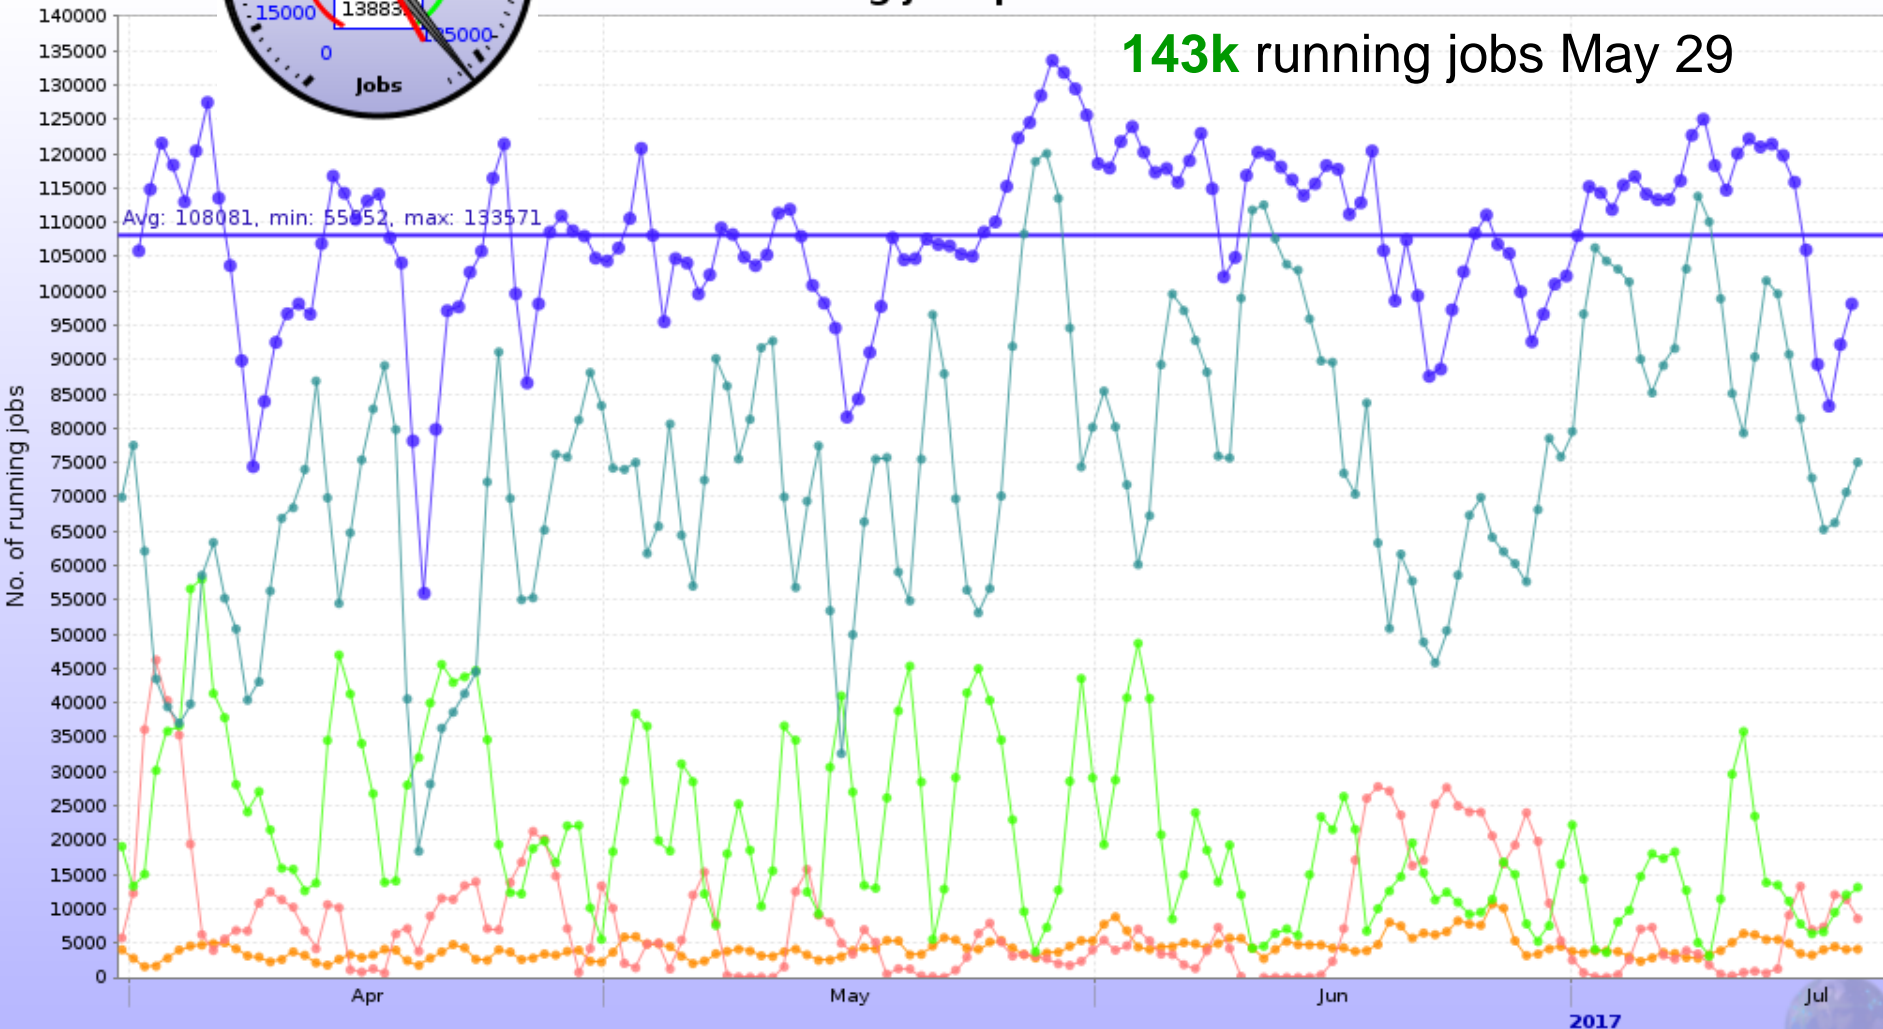
Very high activity



Running jobs per user



143k running jobs May 29



— aliproduct (MC productions) — alitrain (organized analysis) — alidaq (RAW data) — users — SUM



Taking advantage of opportunistic resources, in time for SQM (July 10-15)

Storage at CERN and beyond



- In general: **growth** and **reliability** – thanks!
- CASTOR
 - Mostly stable, a few incidents
- EOS
 - Ditto
- New EOS instance at Hiroshima

Issues at sites or with jobs



- CERN: multiple incidents affecting the new HTCondor CEs
 - Currently the gateways to 70% of the resources at CERN
 - CE configurations were further improved for robustness and scalability
 - More CEs were added as well
 - Some bugs were fixed
 - One major issue remains: CEs run out of memory due to proxy re-delegations and crash as a result
 - The devs have been asked to look into that
 - Workarounds are being considered

Middleware



- **CentOS/EL7** is steadily becoming more important
 - Various service types available in UMD-4 and/or EPEL-7
 - CREAM, EMI-UI and EMI-WN delayed to this month...
 - 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 run 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
 - New features only go into UMD-4

Containers (1)



- An isolation paradigm much lighter than VMs
- A new tool to launch containers is gaining momentum in our community and beyond:
Singularity
 - Provide desired 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
 - The WLCG Container WG will try to steer these activities across sites and the 4 experiments

Containers (2)




- They can also be used for easy deployment of pre-packaged services
 - Development hosts
 - Build and validation services
 - WLCG VOBOX
 - ...
- A Docker container for a WLCG VOBOX was set up by [Maxim Storetvedt](#) (Western Norway Univ. of Applied Sciences)
 - See his talk later in this [session](#)
- It was then configured for a new test “site” called [Nemesis](#) 😊
 - It has been working successfully alongside the other 3 “sites” submitting jobs to the HTCondor resources – see next page
- The plan is to replace the VMs of the other “sites” with such container instances
 - Better performance, more reliability

CERN HTCondor “sites”



Running Jobs



SAM



- New Availability / Reliability profile based on selected **MonALISA** metrics in use since 1 year
- 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 production on May 21

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)
 - 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>