



# Grid status

ALICE Offline week  
Nov 3, 2016

Maarten Litmaath  
CERN-IT

v1.0

# Central services



Machine	Machine status			Machine type						Mem	
	Online	Uptime	Load	Kernel	Machine model	CPU	CPUs	MHz	Total ▲	Used	
8. db6c	✓	48d 9:46	3.84	4.4.0-36-...	ProLiant DL380 Gen9	Xeon E5-2667 v4 3.20GHz	32	1298	1.476 TB	162.9 GB	
7. db6b	✓	422d 11:02	1.68	3.19.0-26...	ProLiant DL380 Gen9	Xeon E5-2687W v3 3.10GHz	40	1200	755.8 GB	145.7 GB	
9. db6d	✓	48d 4:42	0.5	4.4.0-36-...	ProLiant DL380 Gen9	Xeon E5-2687W v3 3.10GHz	40	1200	755.8 GB	293.4 GB	
11. db6f	✓	9d 11:42	0.17	4.4.0-45-...	ProLiant DL380 Gen9	Xeon E5-2687W v4 3.00GHz	48	1200	755.8 GB	157.9 GB	
12. db6g	✓	9d 11:37	0.28	4.4.0-45-...	ProLiant DL380 Gen9	Xeon E5-2697A v4 2.60GHz	64	1199	755.8 GB	154.5 GB	
15. db9	✓	387d 3:45	0.04	3.13.0-65...	ProLiant DL380p Gen8	Xeon E5-2690 2.90GHz	32	1200	377.9 GB	26.77 GB	
29. alienvm1	✓	281d 10:12	0.39	3.13.0-76...	ProLiant DL380p Gen8	Xeon E5-2690 v2 3.00GHz	40	1200	377.9 GB	35.39 GB	
13. db7	✓	747d 22:31	2.37	3.13.0-37...	ProLiant DL380p Gen8	Xeon E5-2690 2.90GHz	32	1200	377.9 GB	11.08 GB	
37. pcaliweb02	✓	698d 10:34	0.1	3.13.0-40...	ProLiant DL380p Gen8	Xeon E5-2690 2.90GHz	32	1200	377.9 GB	6.951 GB	
44. alientest02	✓	6d 10:19	0.44	4.4.0-45-...	ProLiant DL380p Gen8	Xeon E5-2690 v2 3.00GHz	40	1632	377.9 GB	61.47 GB	
41. pcalimonitor4	✓	8d 9:41	2.13	4.4.0-45-...	ProLiant DL380p Gen8	Xeon E5-2690 2.90GHz	32	1205	377.9 GB	10.87 GB	
39. pcalimonitor2	✓	8d 11:41	1.45	4.4.0-45-...	ProLiant DL380p Gen8	Xeon E5-2690 2.90GHz	32	1226	377.9 GB	27.27 GB	
38. pcalimonitor	✓	8d 9:39	29.87	4.4.0-45-...	ProLiant DL380p Gen8	Xeon E5-2690 v2 3.00GHz	40	1848	377.9 GB	91.81 GB	
10. db6e	✓	9d 11:43	0.64	4.4.0-45-...	ProLiant DL380 G7	Xeon X5680 3.33GHz	24	1634	283.4 GB	120.7 GB	

- Mostly stable, a few incidents
  - DB locks, connection leaks, MySQL hangs
  - API server slowdown due to catalog bug (fixed)
  - One broken gcc/alienv update
  - Maintenance not always transparent
- Unavailability Sep 14-15 due to big network intervention
  - All user and grid activity had to be stopped
  - In parallel the File Catalog was moved to a new, more powerful machine!

# New sites

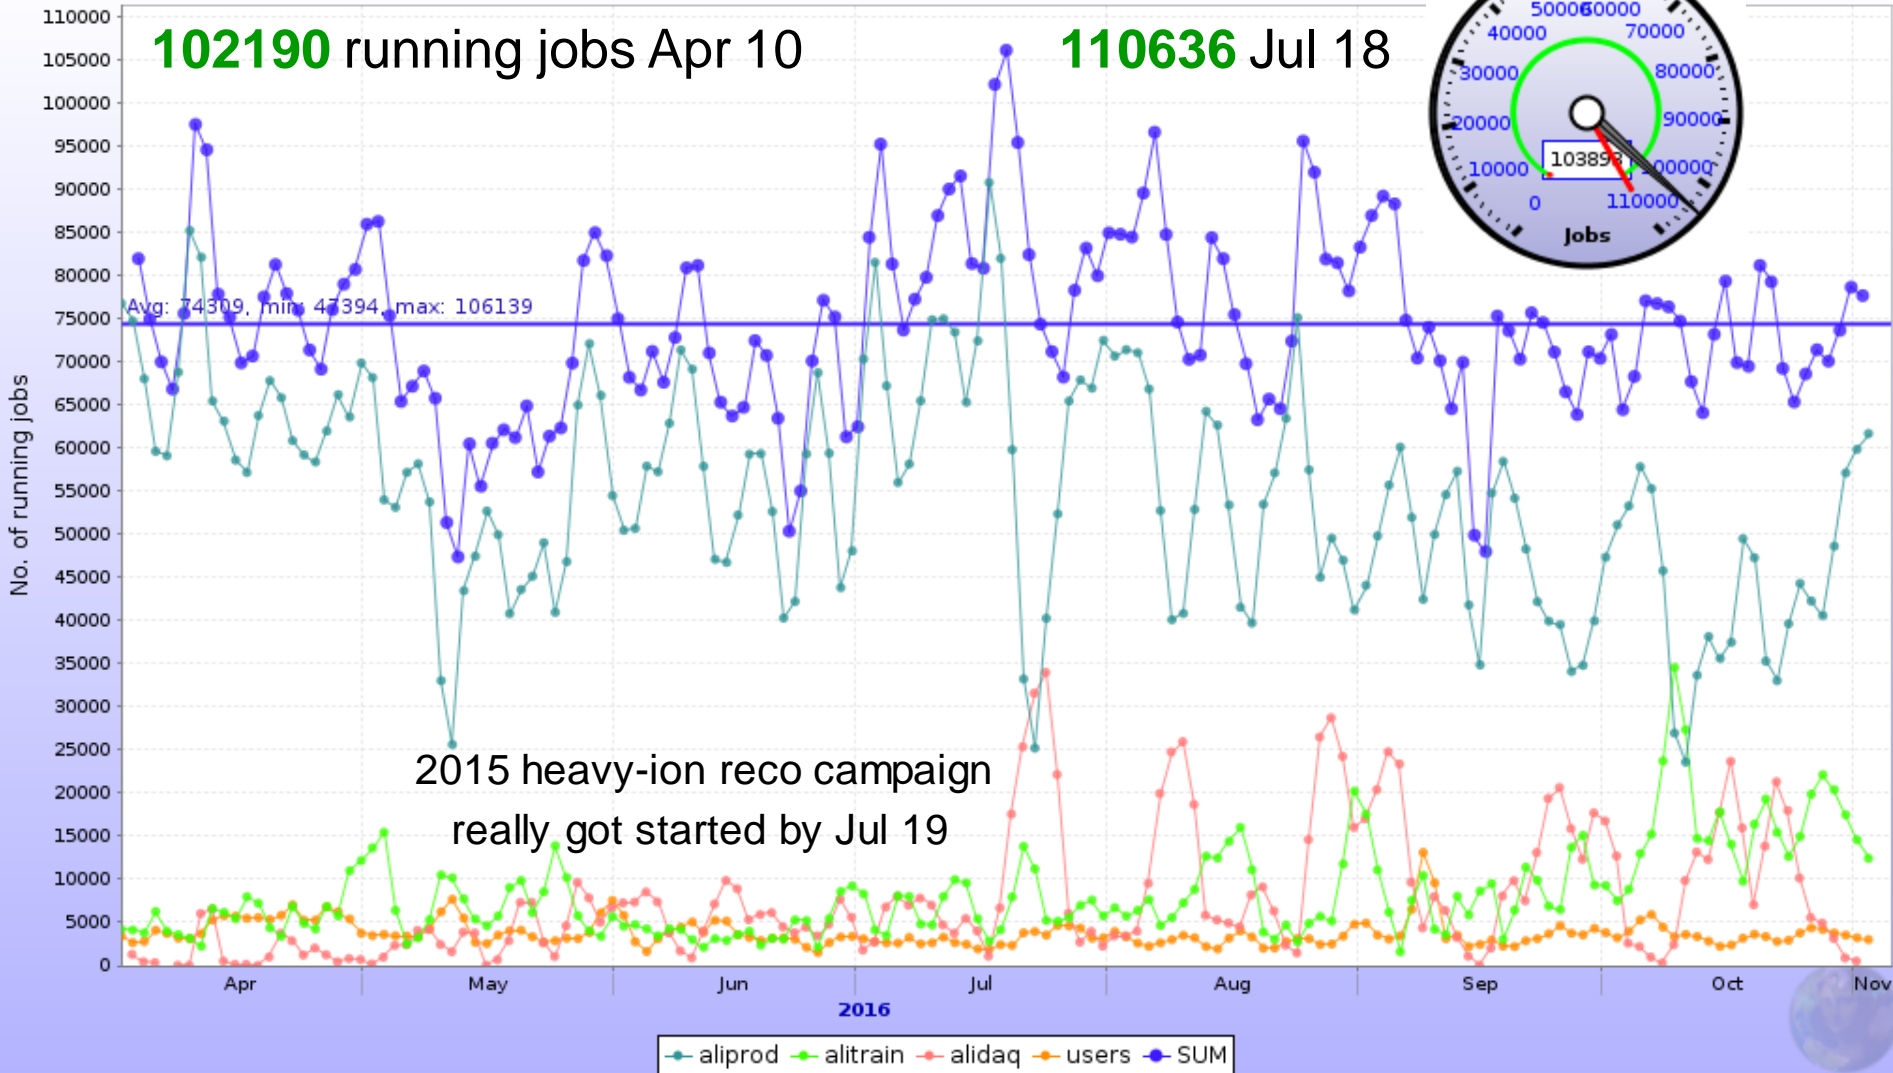


- **Vienna**: up to 1500 job slots!
- **UPB** (Bucharest): in progress
- New CERN “sites” for redundancy and load sharing, submitting to the steadily increasing HTCondor resources (currently 30% vs. 70% LSF)
  - CERN-CORONA
  - CERN-MIRAGE
- **Altaria**: CERN extension into T-Systems cloud
  - Details further in this presentation

# High activity



Running jobs per user



# Storage at CERN and beyond



- **CASTOR**
  - Mostly stable, a few incidents
  - Disk pool setup for the proton-ion run and beyond agreed with CASTOR team
- **EOS**
  - Mostly stable, a few incidents
  - New head nodes with 512 GB memory were rolled in transparently
- **EOS crashes at CERN and other sites on Oct 13**
  - Clients unexpectedly used signed URLs instead of encrypted XML tokens
  - The switch was due to one DB table being temporarily unavailable plus a wrong default (fixed)
  - The EOS devs have been asked to support the new scheme and avoid that unexpected requests crash the service

# Issues at sites or with jobs



- CERN: multiple incidents affecting CREAM CEs
  - Currently the gateways to 70% of the resources at CERN
  - One time it was our fault: the sandbox partitions filled up with debug logs (fixed)
- KISTI: CVMFS Squid services incident
  - “Random” job failures for 5 days
- Asian sites: CVMFS issues involving the Stratum-1 at ASGC
  - Mitigated and being investigated for a real fix
- CERN-AURORA upgrade to broken AliEn version
  - Non-trivial upgrades will be announced better
- ...
- Occasional high error rates due to jobs themselves
  - Badly behaved code, e.g. using way too much memory
  - Output from previous attempts not cleaned up
  - ...

# Altaria



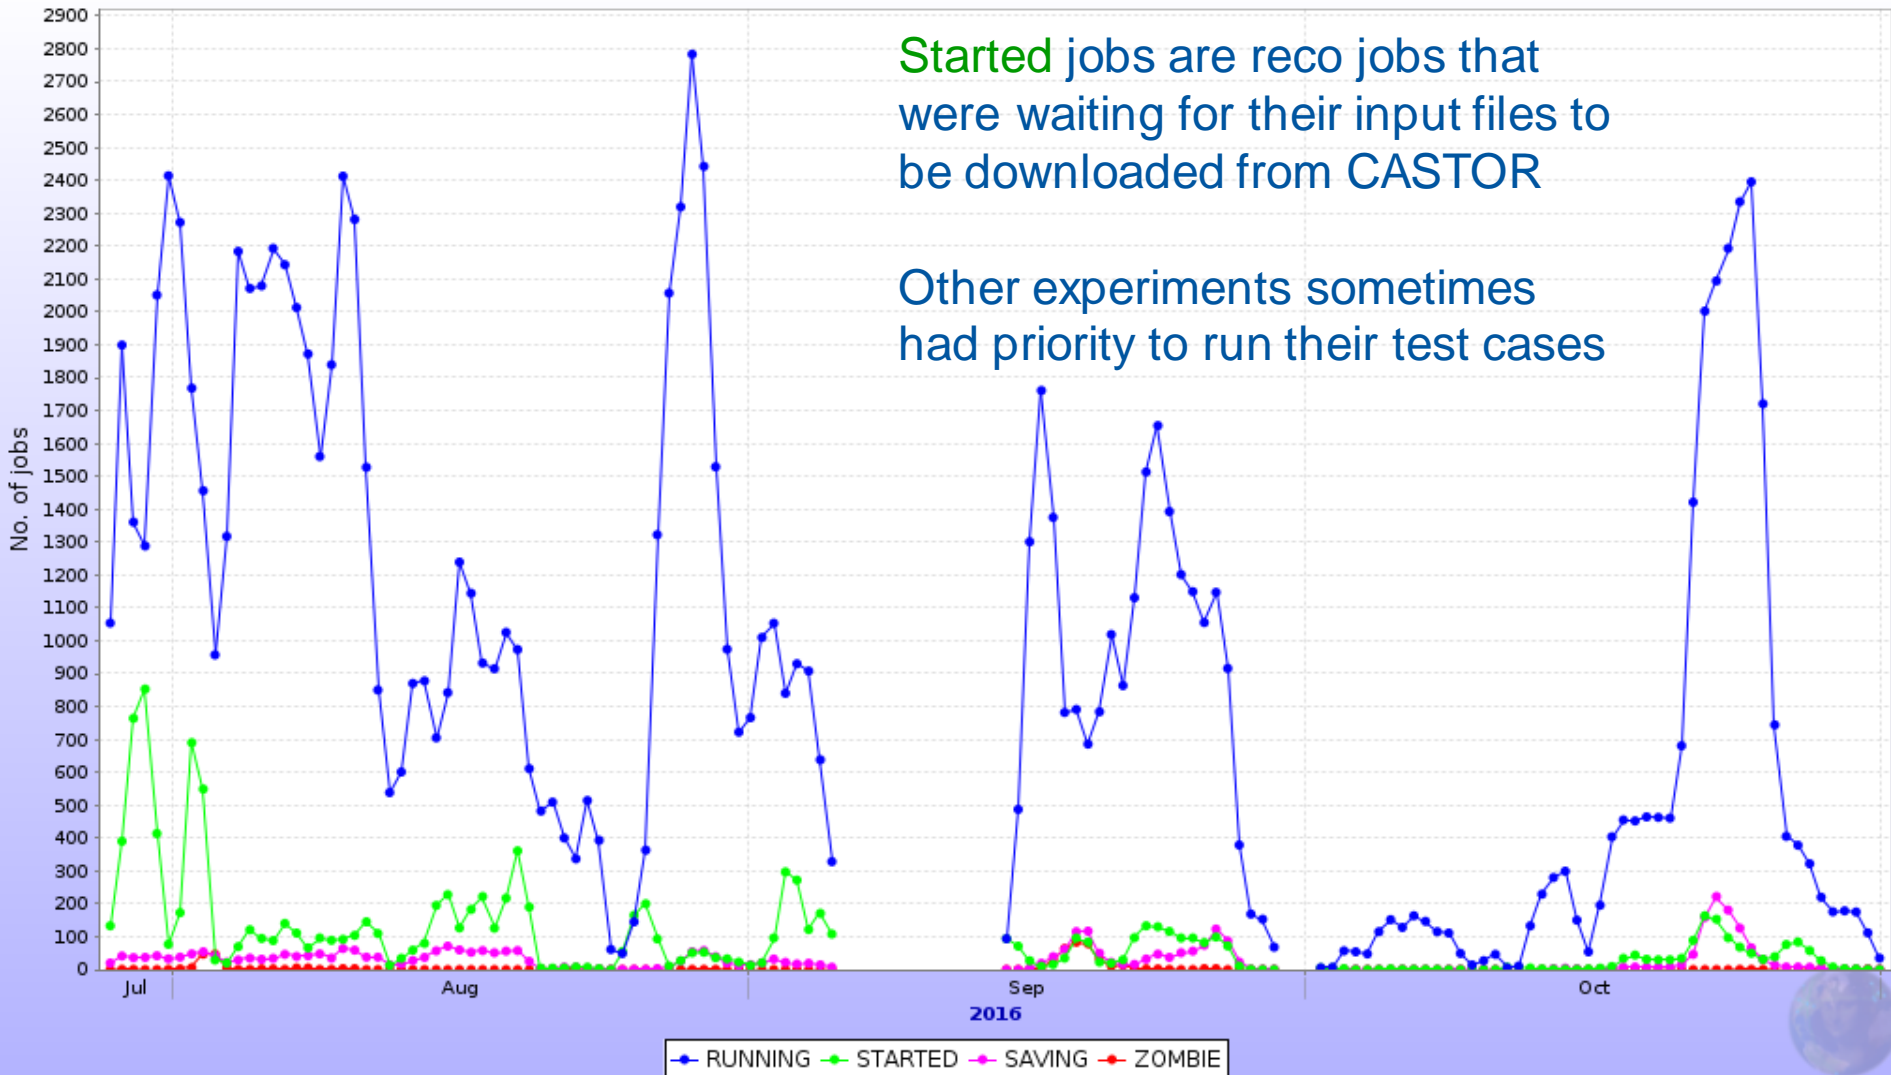
- A virtual site to drive a temporary extension of CERN resources located in a cloud at T-Systems, Germany
  - A common project to test how cloud resources can be procured and used efficiently
- Up to 3500 cores were available, shared by the 4 experiments
- ALICE were the first to start using them as of July 28
- The project ended on Oct 31
- The resources were used for reco, MC and user analysis
  - Trains were excluded to reduce the load on the 10-Gbit link to CERN
- The exercise has helped the IT cloud team gain useful insights
  - And allowed 630k ALICE jobs to complete successfully!
  - The error types and rates were compatible with those of the CERN “sites”



## Active jobs in Altaria

Started jobs are reco jobs that were waiting for their input files to be downloaded from CASTOR

Other experiments sometimes had priority to run their test cases





# Middleware



- **SL5** has been deprecated for MW as of May
- **CentOS/EL7** is becoming more important
  - Some services already available
  - CREAM expected by Dec
    - VOBOX will follow
- **SL6** still the default, but the experiments are preparing for physical worker nodes running CentOS/EL7
  - Containers (or VMs) could still provide SL6
  - ALICE jobs have been tested OK on CentOS/EL7

# SAM



- New Availability / Reliability profile based on selected **MonALISA** metrics in use since May
- So far no big issues were reported
  - And it allowed NDGF-T1 to appear at last !
- **Reminder: SE test failures will reduce the A / R!**
- Corrections can always be applied as needed
- Test job submission to the HTCondor CE still to be added

# Xrootd reminder



- Sites should continue upgrades to Xrootd  $\geq 4.1$ 
  - Half of the sites have done that already, thanks!
  - Required for IPv6 support
- Communication via LCG Task Force list as usual for expert advice
- SL6 hosts can have xrootd for ALICE installed 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>