



ALICE Monitoring

Costin.Grigoras@cern.ch



ALICE

Summary

- Monitoring infrastructure
- Information sources
- Collecting and storing
- Presentation
- Feedback to production system, dashboard and SAM



Monitoring infrastructure

- All components instrumented with ApMon
- Distributed MonALISA collectors, one per site in each VoBox
 - also doing data aggregation
- Central collector for long term data archival and presentation
 - and specialized collectors for sites and other interested parties



Information sources

- Jobs: ~real time monitoring of job resource usage and of host health (CPU and Wall time and specs, memory, traffic, disk IO ...)
- Services: VoBoxes and Central Services in detail
 - AliEn services, box health, inter-site connectivity tests
 - also controlling the DNS balancing



ALICE

Information sources (2)

- Storage: all xrootd data nodes, all xrdcp/xrd3cp transfers, health, occupancy
- Networking: tracepath and available bw per stream between all sites, kernel buffer sizes, IPv6 readiness
- Other data: production accounting, I/O eff. of sites vs storages, SE functional status, experiment data taking and electronic logbook ...



ALICE

Collecting and storing

- Centrally collecting as much as possible aggregated data only (per site, user, masterjob, production etc)
 - details are quickly forgotten and old data is further aggregated
 - 7 years of data in some 370GB
- Same database for monitoring data, accounting, production management, experiment data taking info, detector QA info and many more
 - allowing for quick production assessment and control plus all kinds of correlations we have or haven't imagined yet



ALICE

Presentation

- A single web interface for all operations with separate sections for the interested parties:
 - site accounting and alerts for sites
 - real time view of users' jobs, following their progress, memory usage, with access to logs, Grid catalogue browser, file editor etc
 - managers set up meta production requests and follow their automatic progress
 - users can also join organized analysis trains by a few clicks
- 200K pages per day, 0.1s avg response time



Presentation (2)

My jobs | My home dir | Catalogue browser | LEGO Trains | Adminis

ALICE Repository

- ALICE Repository
- Google Map
- Shifter's dashboard
- Run Condition Table
- Production Overview
- Production info
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- Build system
- HepSpec
- Dynamic charts

close all

This page: [bookmark](#), [URL](#)

Running jobs trend

Jobs: 35053

Status of masterjob 301985348

```

301991001 : trace, 1.6 GB RSS, 4.598 GB VMS (did not run, max RSS: 1.6 GB, Virt: 4.598 GB) (resubmitted 2 times)
301991003 : trace, 443.4 MB RSS, 3.395 GB VMS (did not run, max RSS: 443.4 MB, Virt: 3.395 GB) (resubmitted 1 time)
301991005 : trace, 2.175 GB RSS, 4.251 GB VMS (did not run, max RSS: 2.175 GB, Virt: 4.251 GB) (resubmitted 2 times)

```

RUNNING (11671)

```

301985441 : trace, 729.9 MB RSS, 2.965 GB VMS (running for 2:03, 99.45% CPU @ ALICE::CNAF::LCG, max RSS: 729.9 MB, Virt: 2.965 GB)
301985442 : trace, 817.7 MB RSS, 1.914 GB VMS (running for 2:03, 89.77% CPU @ ALICE::LLNL::PBS, max RSS: 817.7 MB, Virt: 1.914 GB)
301985443 : trace, 829.9 MB RSS, 2.232 GB VMS (running for 2:02, 94.35% CPU @ ALICE::Hiroshima::Hiroshima, max RSS: 829.9 MB, Virt: 2.232 GB)
301985444 : trace, 847.4 MB RSS, 2.213 GB VMS (running for 2:02, 95.13% CPU @ ALICE::Hiroshima::Hiroshima, max RSS: 847.4 MB, Virt: 2.213 GB)
301985445 : trace, 836.4 MB RSS, 1.93 GB VMS (running for 2:02, 89.08% CPU @ ALICE::LLNL::PBS, max RSS: 836.4 MB, Virt: 1.93 GB)
301985446 : trace, 699.3 MB RSS, 2.933 GB VMS (running for 2:03, 99.31% CPU @ ALICE::CNAF::LCG, max RSS: 699.3 MB, Virt: 2.933 GB)
301985449 : trace, 713.7 MB RSS, 2.947 GB VMS (running for 2:03, 98.11% CPU @ ALICE::CNAF::LCG, max RSS: 713.7 MB, Virt: 2.947 GB)
301985450 : trace, 708.7 MB RSS, 2.803 GB VMS (running for 1:29, 125.4% CPU @ ALICE::CERN::LCG, max RSS: 708.7 MB, Virt: 2.803 GB)
301985451 : trace, 709.6 MB RSS, 2.943 GB VMS (running for 2:03, 99.35% CPU @ ALICE::CNAF::LCG, max RSS: 709.6 MB, Virt: 2.943 GB)
301985452 : trace, 726.2 MB RSS, 2.249 GB VMS (running for 1:28, 91.1% CPU @ ALICE::LBL::CONDOR, max RSS: 726.2 MB, Virt: 2.249 GB)
301985453 : trace, 877.6 MB RSS, 1.973 GB VMS (running for 2:02, 89.61% CPU @ ALICE::LLNL::PBS, max RSS: 877.6 MB, Virt: 1.973 GB)
301985454 : trace, 736.9 MB RSS, 1.9 GB VMS (running for 1:29, 42.54% CPU @ ALICE::CCIN2P3::LCGSGE2, max RSS: 736.9 MB, Virt: 1.9 GB)
301985455 : trace, 708.5 MB RSS, 2.799 GB VMS (running for 1:29, 65.75% CPU @ ALICE::CERN::LCG, max RSS: 708.5 MB, Virt: 2.799 GB)
301985456 : trace, 864.1 MB RSS, 3.492 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 864.1 MB, Virt: 3.492 GB)
301985457 : trace, 916.2 MB RSS, 2.946 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 916.2 MB, Virt: 2.946 GB)
301985458 : trace, 699 MB RSS, 1.915 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 699 MB, Virt: 1.915 GB)
301985459 : trace, 725.1 MB RSS, 2.258 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 725.1 MB, Virt: 2.258 GB)
301985460 : trace, 562.3 MB RSS, 2.203 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 562.3 MB, Virt: 2.203 GB)
301985461 : trace, 878.3 MB RSS, 2.721 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 878.3 MB, Virt: 2.721 GB)
301985462 : trace, 706.4 MB RSS, 2.721 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 706.4 MB, Virt: 2.721 GB)
301985463 : trace, 853.9 MB RSS, 2.721 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 853.9 MB, Virt: 2.721 GB)
301985464 : trace, 837 GB RSS, 2.893 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 837 GB, Virt: 2.893 GB)
301985465 : trace, 706.2 MB RSS, 2.721 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 706.2 MB, Virt: 2.721 GB)
301985466 : trace, 838.5 MB RSS, 2.721 GB VMS (running for 2:02, 99.75% CPU @ ALICE::CERN::LCG, max RSS: 838.5 MB, Virt: 2.721 GB)
301985475 : trace, 1.122 GB RSS, 3.478 GB VMS (running for 9:19, 143.2% CPU @ ALICE::CERN::LCG, max RSS: 1.122 GB, Virt: 3.478 GB)
301985476 : trace, 317.4 MB RSS, 2.706 GB VMS (running for 2:02, 0.028% CPU @ ALICE::NIHAM::PBS64, max RSS: 317.4 MB, Virt: 2.706 GB)
301985477 : trace, 715.6 MB RSS, 2.227 GB VMS (running for 1:28, 90.79% CPU @ ALICE::LBL::CONDOR, max RSS: 715.6 MB, Virt: 2.227 GB)

```

04 Jul 2013

Legend: RSS (red), Virtual (cyan)



ALICE

Feedback to production system

Monitoring data is most useful when automatic actions can be implemented based on it

- storage replica selection based on SE distance and status (both read and write)
- automatic production management (job submission and resubmission, following complex meta processing chains)
- remote services control, deploying Grid-wide software upgrades
- central services load balancing
- alerts on any event with open subscription, annotating downtime periods on the time series display



ALICE

Feedback to/from dashboard and SAM

- Some part of the monitoring information has high relevance for 'general' monitoring
 - case in point - SAM
- The reverse is also the case - instead of gathering information ourselves, external sources can provide meaningful data
 - case in point - job information from BDII
- Other cases can be identified and consolidated for more harmonious picture
 - our hope is that this TF will be able to do just that



Summary

- ALICE has a vertically integrated monitoring system
 - raw and aggregated data is open for anybody to consume
 - dedicated exports are in place for filling common monitoring databases
- Aiming to a tighter integration with the common tools in order to offer a clearer picture to the users