

# Analysis framework - status

Andrei Gheata, Mihaela Gheata, Andreas Morsch  
ALICE offline week, 5-9 July 2010

# Parameter containers

- Allow saving configuration AND/OR post-processed data using the output container formalism
  - Already introduced last time, now few bugs fixed in the implementation
- Simple procedure allowing saving this kind of data together with the final merged outputs
  - Define an output slot in the task corresponding to this data
  - In the AddTask macro, create a container of type kParamContainer and connect to this slot
    - `mgr->CreateContainer("cMyTaskConfig", TList::Class(),  
AliAnalysisManager::kParamContainer,  
Form("%s:ConfigData", myOutputFile));`
  - In `LocalInit()` use `PostData()` for the configuration data. In `Terminate()` one can also `PostData()` for any post-processed data (like normalization or fitting)
- **Please start using this very important new feature !**
  - **Allows keeping the configuration with the results**
  - We will also do it for the physics selection task

# Alien plugin (1)

- Most development done in this area
  - The plugin became the preferred way of doing grid analysis...
- Merging via generated JDL script
  - `plugin->SetMergeViaJDL()`
  - Files for merging (JDL, merge macro, executable and validation script) are generated and copied in alien space.
  - Submitted in “terminate” mode only
  - The Terminate for the tasks get executed at the end of the merging jobs (!)
  - Currently the merging is done per master job – to be extended

# Alien plugin (2)

- `plugin->SetOverwriteMode()`
  - Will trigger overwriting input data collections AND existing output files.
  - Make sure not to have this if you need to resume a failed merging of outputs
- `plugin->SetFastReadOption()`
  - Try to reduce default xrootd timeouts to deal with the fact that currently alien does not prioritize copies which are closer or more accessible
  - Use on your own risk – you will be warned whenever you use this feature.

# Alien plugin (3)

- Improved feedback from failing steps of running via the plugin
  - Verbosity of copy/submit commands
  - Check if file copying works, printing out `alien_CLOSE_SE`
  - Better indications of what to do in case of failure
  - Automatic check if the generated dataset is empty
- “Micro-control” of plugin components
  - Enable/disable file copying in AliEn
  - Submission, merging (ON/OFF) to be implemented

# Alien plugin (4)

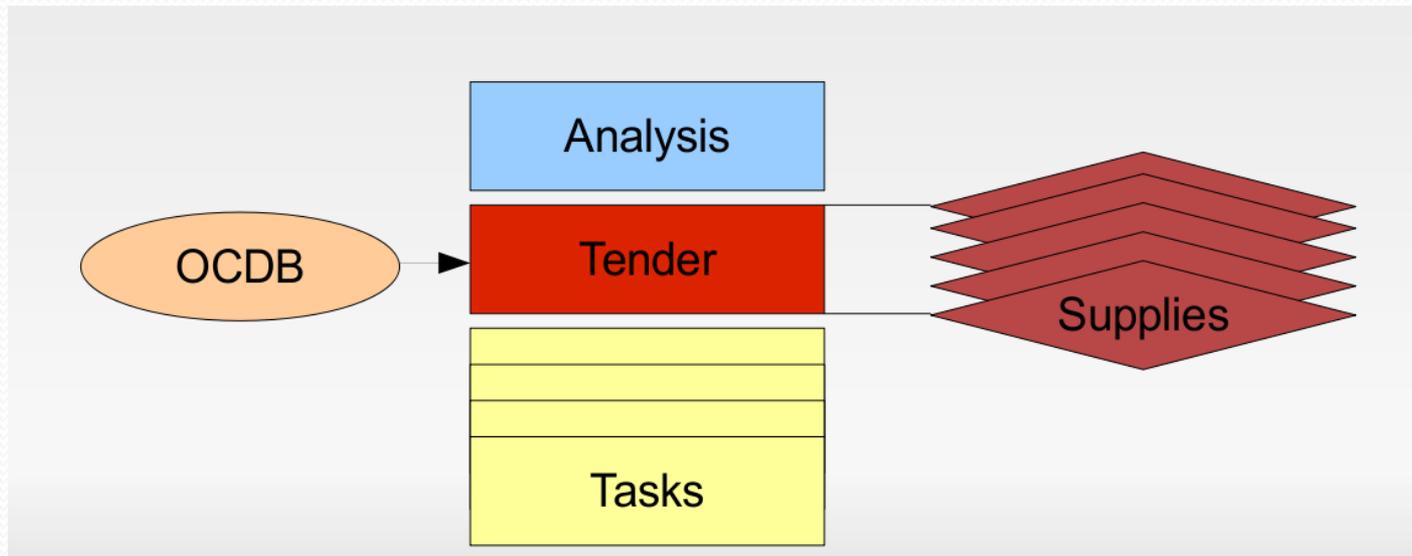
- **SetDefaultOutputs()** enabled as default
  - Do not specify your outputs by hand anymore
  - Will automatically archive all outputs in `root_archive.zip` – **SetOutputArchive produces a Fatal if default outputs are set**
- Using specific SE's via: **output@My::preferred::SE** or via: `plugin->SetPreferredSE()` not allowed anymore
  - **@disk=2** used now by default
  - Use: **SetNumberOfReplicas(max=4)**
- **stdout** and **stderr** kept only for the jobs that fail for some reason

# Plans

- Unified plugin submission both in GRID and PROOF-based analysis facilities
- Extend plugin “micro” controls to allow automatic generation and update for production jobs
- Extended list of cuts accommodating user cuts besides the standard physics selection

# Analysis tender

- Supplies accessing and applying calibration improvements w.r.t pass1 to ESD
- Tasks benefit of this change
- Event selection can be used
- AOD's can be streamed out



# Tender in production

- Used for pass1 filtering



## MonALISA Repository for ALICE



[My jobs](#) ★ | 
 [My home dir](#) ★ | 
 [Catalogue browser](#) ★ | 
 [Repository Home](#) | 
 [Administration Section](#) | 
 [ALICE Reports](#) | 
 [Events XML Feed](#) | 
 [Firefox Toolbar](#) | 
 [MonaLisa GUI](#)

### ALICE Repository

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

close all

### Running jobs trend



### PRODUCTION CYCLES

Train Details » No filter

Manage »

Production info						Jobs status				Comment
Production	Description	Status	Completion rate	Config	Results	Total	Done	Running	Waiting	
FILTER_pass2_001_LHC10c	FILTER_pass2_001_LHC10c: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10c)	Running	99%			112	111			FILTER_pass2_001_LHC10c: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10c)
FILTER_pass2_001_LHC10b	FILTER_pass2_001: data ESD + PhysSel -> AOD + Muons + Vertices	Running	98%			1296	1283	7		FILTER_pass2_001: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10b)
FILTER_pass1_001_LHC10b	FILTER_pass1_001: data ESD + PhysSel + Tender -> AOD + Muons + Vertices	Running	97%			748	727			FILTER_pass1_001: data ESD + PhysSel + Tender -> AOD + Muons + Vertices (LHC10b)
QA11_LHC10d3_Merging	QA11_LHC10d3_Merging: PWG1 QA train	Running	100%			1	1			QA11_LHC10d3_Merging: PWG1 QA train (sim)
QA11_LHC10d3	QA11_LHC10d3: PWG1 QA train	Running	30%			1773	533		1240	QA11_LHC10d3: PWG1 QA train (sim)
QA10_LHC10d_Merging	QA10_LHC10d_Merging: PWG1 QA	Running	86%			15	13			QA10_LHC10d_Merging: PWG1 QA

# Analysis trains

- PWG<sub>1</sub> QA trains running now automatically
  - Policy: A new QA train version every new AN tag
  - Memory was an issue – now under control (see next)
  - **No QA flag for the run yet as outcome of this procedure !**
- Event selected filtering trains starting
  - Producing standard AOD's (UA<sub>1</sub> jets included) + muon AOD's + Vertexing deltas
  - Mostly starting from pass2 data, but can run also on pass1 using the tender
  - **Please test your analysis on this data !**
- **PWG<sub>n</sub> physics wagons should validate their configuration used in the central analysis train for running on data !**

# Trains

- Shifter's dashboard
  - Run Condition Table
  - Production info
  - Job Information
  - SE Information
  - Services
  - Network Traffic
  - FTD Transfers
  - CAF Monitoring
  - SHUTTLE
  - Build system
  - Dynamic charts
- close all



Train Details # 170 filter

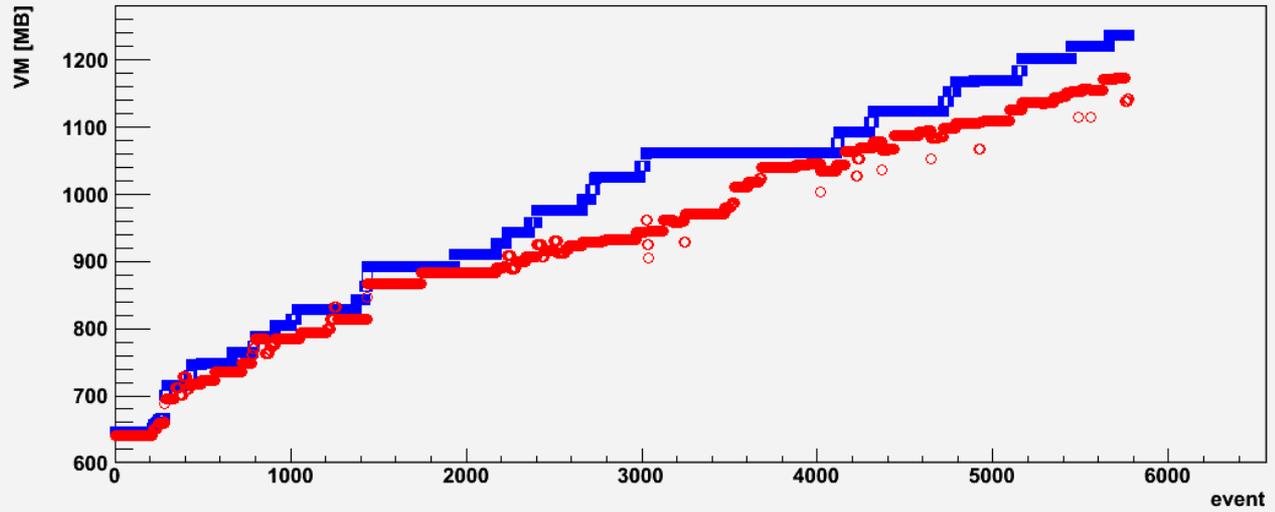
manage >

Production info						Jobs status				Comment
Production	Description	Status	Completion rate	Config	Results	Total	Done	Running	Waiting	
FILTER_pass2_001_LHC10c	FILTER_pass2_001_LHC10c: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10c)	Running	99%			112	111			FILTER_pass2_001_LHC10c: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10c)
FILTER_pass2_001_LHC10b	FILTER_pass2_001: data ESD + PhysSel -> AOD + Muons + Vertices	Running	99%			1296	1284	5		FILTER_pass2_001: data ESD + PhysSel -> AOD + Muons + Vertices (LHC10b)
FILTER_pass1_001_LHC10b	FILTER_pass1_001: data ESD + PhysSel + Tender -> AOD + Muons + Vertices	Running	97%			748	727			FILTER_pass1_001: data ESD + PhysSel + Tender -> AOD + Muons + Vertices (LHC10b)
QA11_LHC10d3_Merging	QA11_LHC10d3_Merging: PWG1 QAtrain	Running	100%			1	1			QA11_LHC10d3_Merging: PWG1 QA train (sim)
QA11_LHC10d3	QA11_LHC10d3: PWG1 QAtrain	Running	30%			1773	533		1240	QA11_LHC10d3: PWG1 QA train (sim)
QA10_LHC10d_Merging	QA10_LHC10d_Merging: PWG1 QAtrain	Running	86%			15	13			QA10_LHC10d_Merging: PWG1 QA train
QA10_LHC10d	QA10_LHC10d: PWG1 QAtrain	Running	94%			583	553	15	7	QA10_LHC10d: PWG1 QA train
QA11_LHC10d2_Merging	QA11_LHC10d2_Merging: PWG1 QAtrain	Running	44%			9	4	1	3	QA11_LHC10d2_Merging: PWG1 QA train (sim)
QA11_LHC10d2	QA11_LHC10d2: PWG1 QAtrain	Running	99%			3499	3471		17	QA11_LHC10d2: PWG1 QA train (sim)
QA11_LHC10d1_Merging	QA11_LHC10d1_Merging: PWG1 QAtrain	Running	81%			11	9			QA11_LHC10d1_Merging: PWG1 QA train (sim)
QA11_LHC10d1	QA11_LHC10d1: PWG1 QAtrain	Running	98%			1517	1487			QA11_LHC10d1: PWG1 QA train (sim)
QA10_LHC10c_Merging(pass2)	QA10_LHC10c_Merging: PWG1 QAtrain	Running	70%			131	93	2	1	QA10_LHC10c_Merging: PWG1 QA train - pass 2
QA10_LHC10c(pass2)	QA10_LHC10c: PWG1 QAtrain	Running	96%			5968	5750	77	118	QA10_LHC10c: PWG1 QA train - pass 2
QA9_LHC10d_Merging	QA9_LHC10d_Merging: PWG1 QAtrain	Completed	68%			19	13			QA9_LHC10d_Merging: PWG1 QA train
QA9_LHC10d	QA9_LHC10d: PWG1 QAtrain	Completed	99%			1079	1078			QA9_LHC10d: PWG1 QA train
QA9_LHC10b_Merging	QA9_LHC10b_Merging: PWG1 QAtrain	Completed	99%			137	136			QA9_LHC10b_Merging: PWG1 QA train
QA9_LHC10b	QA9_LHC10b: PWG1 QAtrain	Completed	99%			2339	2336			QA9_LHC10b: PWG1 QA train
QA9_LHC10c_Merging	QA9_LHC10c_Merging: PWG1 QAtrain	Completed	100%			11	11			QA9_LHC10c_Merging: PWG1 QA train
QA9_LHC10c	QA9_LHC10c: PWG1 QAtrain	Completed	100%			164	164			QA9_LHC10c: PWG1 QA train
QA8_LHC10d_Merging	QA8_LHC10d_Merging: PWG1 QAtrain (no TRD)	Completed	100%			1	1			QA8_LHC10d_Merging: PWG1 QA train (no TRD)
QA8_LHC10d	QA8_LHC10d: PWG1 QAtrain (no TRD)	Completed	100%			26	26			QA8_LHC10d: PWG1 QA train (no TRD)
QA8_LHC10b_Merging(pass2)	QA8_LHC10b_Merging: PWG1 QAtrain (no TRD)	Completed	97%			75	73			QA8_LHC10b_Merging: PWG1 QA train (no TRD) - pass 2
QA8_LHC10b(pass2)	QA8_LHC10b: PWG1 QAtrain (no TRD)	Completed	99%			839	835			QA8_LHC10b: PWG1 QA train (no TRD) - pass 2
QA8_LHC10c_Merging	QA8_LHC10c_Merging: PWG1 QAtrain (no TRD)	Completed	97%			99	97			QA8_LHC10c_Merging: PWG1 QA train (no TRD)
QA8_LHC10c	QA8_LHC10c: PWG1 QAtrain (no TRD)	Completed	98%			4840	4768			QA8_LHC10c: PWG1 QA train (no TRD)

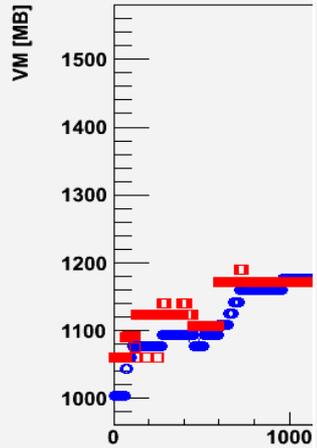
# Task buzzer

- Investigating some huge leaks of  $O(100K/\text{event})$
- Testing tcmalloc from Google performance tools
  - In the hope of fixing fragmentation problems
  - The hope was all that's left at the end...
- Found guilty tasks creating in-memory ntuples and added a couple of blockers in Savannah
  - Some more strict policy for the task output will probably need to be enforced
- Very time consuming process cannot continue like that
  - A new automatic tool is running in grid – the task buzzer...
  - What it does: memory profiling for single tasks

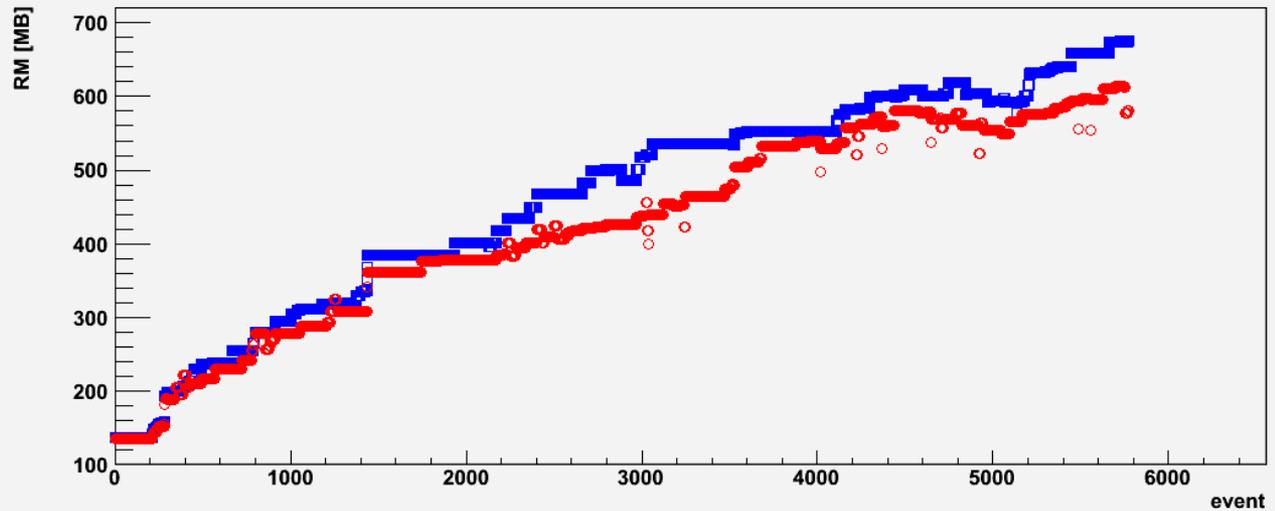
VM[MB] ITStracking tcmalloc:1328 malloc:1350



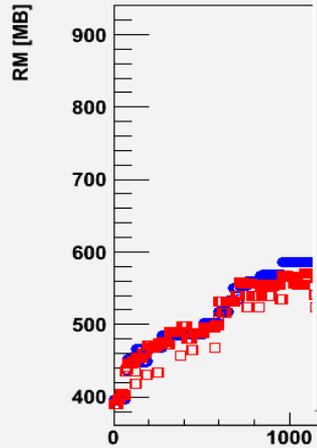
VM[MB] all tasks: tcm=21



RM[MB] ITStracking tcmalloc:1328 malloc:1350



RM[MB] all tasks: tcm=21



event

# Buzzer in action

- Run after every deployed AN tag
- Will start buzzing when detecting leaks...

ALICE Repository

CE Repository

Google Map

Shifter's dashboard

Run Condition Table

Production info

- RAW production cycles
- Analysis train**
- MC production cycles
- MC production requests

Job Information

SE Information

Services

Network Traffic

TD Transfers

CAF Monitoring

SHUTTLE

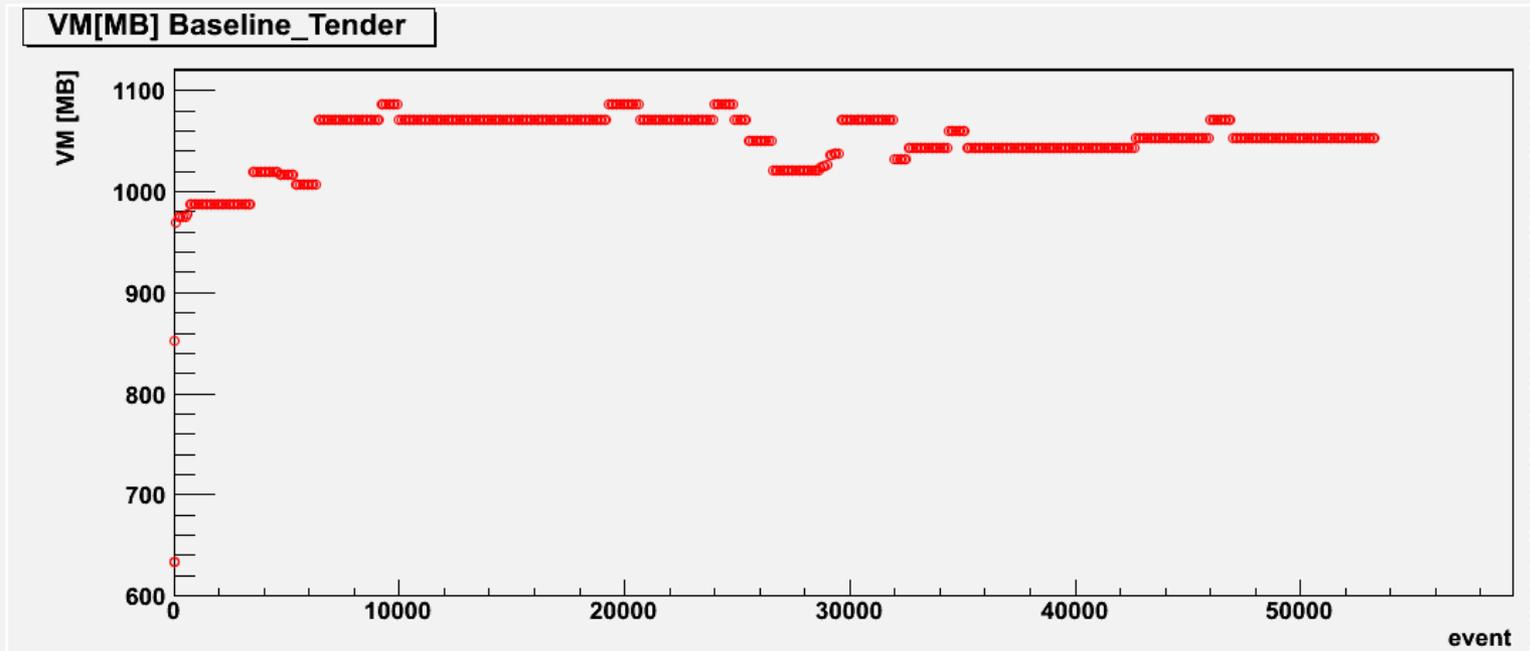
Build system

Dynamic charts

Running jobs trend



A gauge showing the number of running jobs. The scale ranges from 10500 to 21000, with major ticks at 10500, 14000, 17500, and 21000. The needle is positioned at approximately 17500.



# Documentation

- Analysis framework FAQ
  - [http://aliweb.cern.ch/Offline/AliRoot/FAQ.html#Question\\_of\\_the\\_day](http://aliweb.cern.ch/Offline/AliRoot/FAQ.html#Question_of_the_day)
  - An attempt to collect useful information in one place
  - Will hopefully grow bigger