The ALICE Offline Environment - Status and Perspectives

Federico Carminati

on behalf of the ALICE Core Offline Team

26/03/2009

Outline

- AliRoot
  - Simulation
  - Conditions data
  - Reconstruction
  - Visualization
  - Alignment
  - Analysis
- The computing model
  - Resources
  - Data taking scenario
- Summary

Simulation

- Geometry
  - Geometry “as built”
  - Extensive automatic internal consistency checks
  - Account of survey data and alignment
- Generators: possibility to include new ones in a transparent way
- Particle transport: possibility to use in production Geant3, Fluka and Geant4 thanks to the Virtual MC
- Digitization and raw data: detector specific, fully aware of the data taking conditions
- Ongoing improvements in the CPU and memory consumption
- In general: this is the most stable part of AliRoot

Conditions data - Shuttle

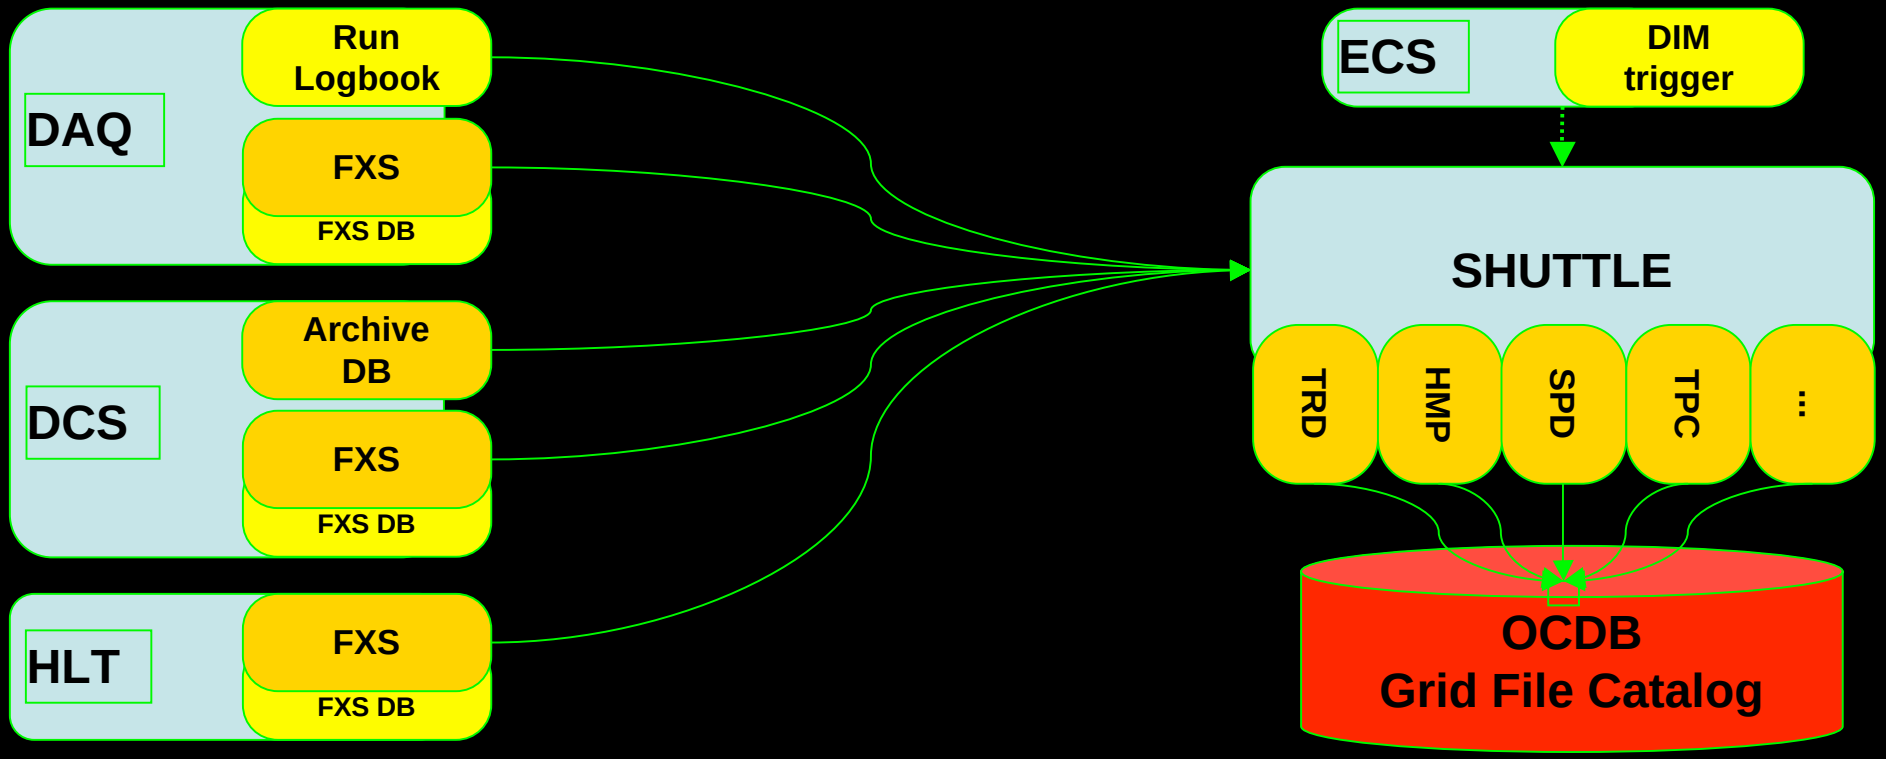

No alternative system to extract data (especially online calibration results) between data-taking and first reconstruction pass!

Conditions data - Shuttle (1)

- Shuttle (subsystem DBs to Grid conditions data publisher) system is in operation since 2 years
  - In production regime for the whole 2008
- Detector algorithms (DAs) within Shuttle have evolved significantly, ready for standard data taking
- High stability of primary sources of conditions data: DCS, DAQ DBs and configuration servers
- Toward the end of last cosmics data taking period (August) – all pieces, including DAs fully operational

Conditions data – Shuttle (2)

- Major efforts concentrated on adding more conditions data
  - Critical LHC parameters
  - New detector's and control hardware
- Conditions data access is the area with least problems on the Grid
  - Both in terms of publication and client access for processing and analysis

Reconstruction

- New developments to meet the requirements that came during the cosmic data taking
  - Prompt online reconstruction.
  - Parallel PROOF based offline reconstruction.
- Further improvements in the algorithms
- Optimization of the CPU and memory consumption.

Prompt

- Recent development
- Very useful for high-level QA and debugging
- Integrated in the AliEVE event display
- Full Offline code sampling events directly from DAQ memory

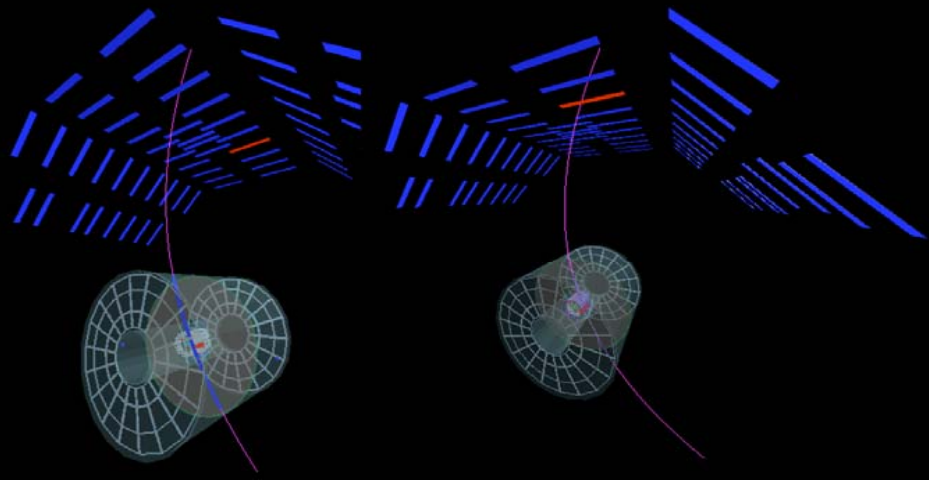

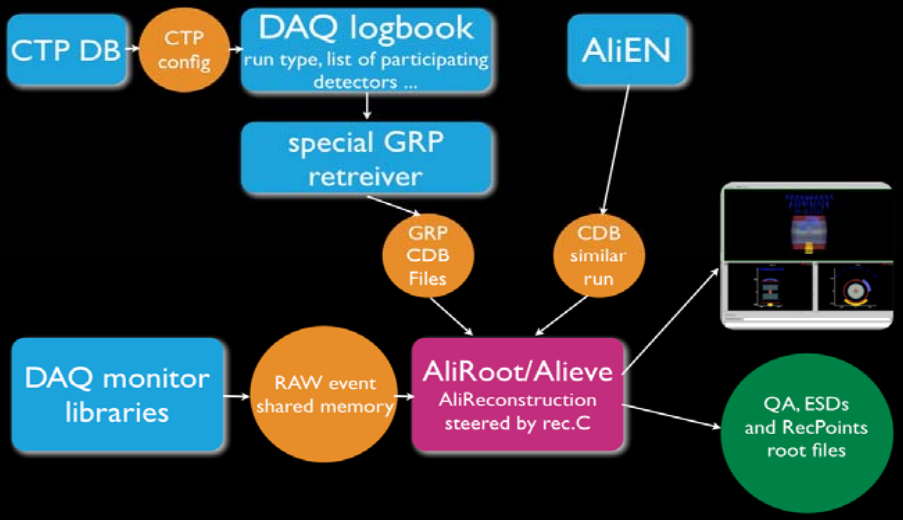

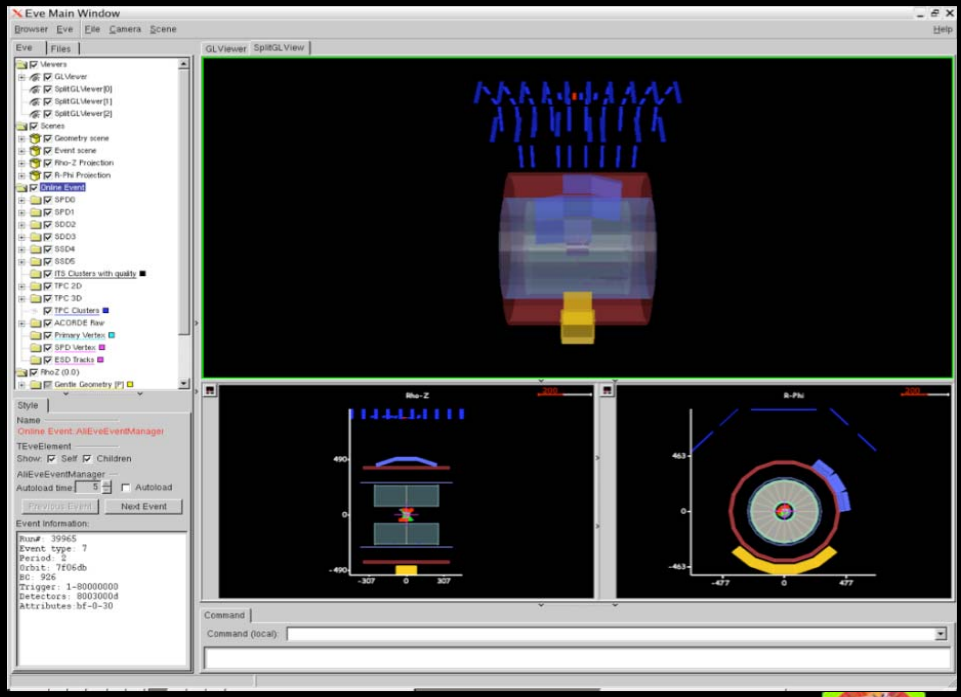

Parallel Reconstruction of Raw Data

- Needed for fast feedback from reconstruction
  - Understand ALICE detector and reconstruction software
  - Debug, tune and optimize reconstruction code
- Based on PROOF (TSelector)
  - Runs on Proof clusters (CAF, GSI AF)
- Transparent
  - User does not notice a difference w.r.t to running locally
- Minimal data flow between components:
  - Common (conditions and options) data accessed once from the client machine
  - Workers access raw-data events directly from AliEn (via xrootd)
- Minimal I/O on the workers
- Fully operational, provides ~30-fold speed-up in the processing rate on current CAF

Offline reconstruction

- Detector reconstruction parameters
  - Several beam/multiplicity/luminosity conditions
  - Taken into account on event-by-event basis
- Quasi-online reconstruction status
  - All runs from 2008 cosmics data processed
    - Emphasis on ‘First physics’ detectors
    - Selected runs already re-processed as ‘Pass 2’ and ‘Pass 3’
- Re-processing of all cosmics data – general ‘Pass 2’
  - After completion of alignment and calibration studies by detectors

Offline reconstruction (2)

- Development of quasi-online processing framework
  - Further refinement of Online QA
  - Speed up the launch of reconstruction jobs to assure ‘hot copy’ of the RAW data
  - January 2009 – detector code readiness review and new set of milestones adapted to the run plan
- The middleware and fabric are fully tested for ‘pass 1’ (T0) RAW data processing
  - To a lesser extent at T1s – limited replication of RAW to save tapes

AliEve – ALICE event-display

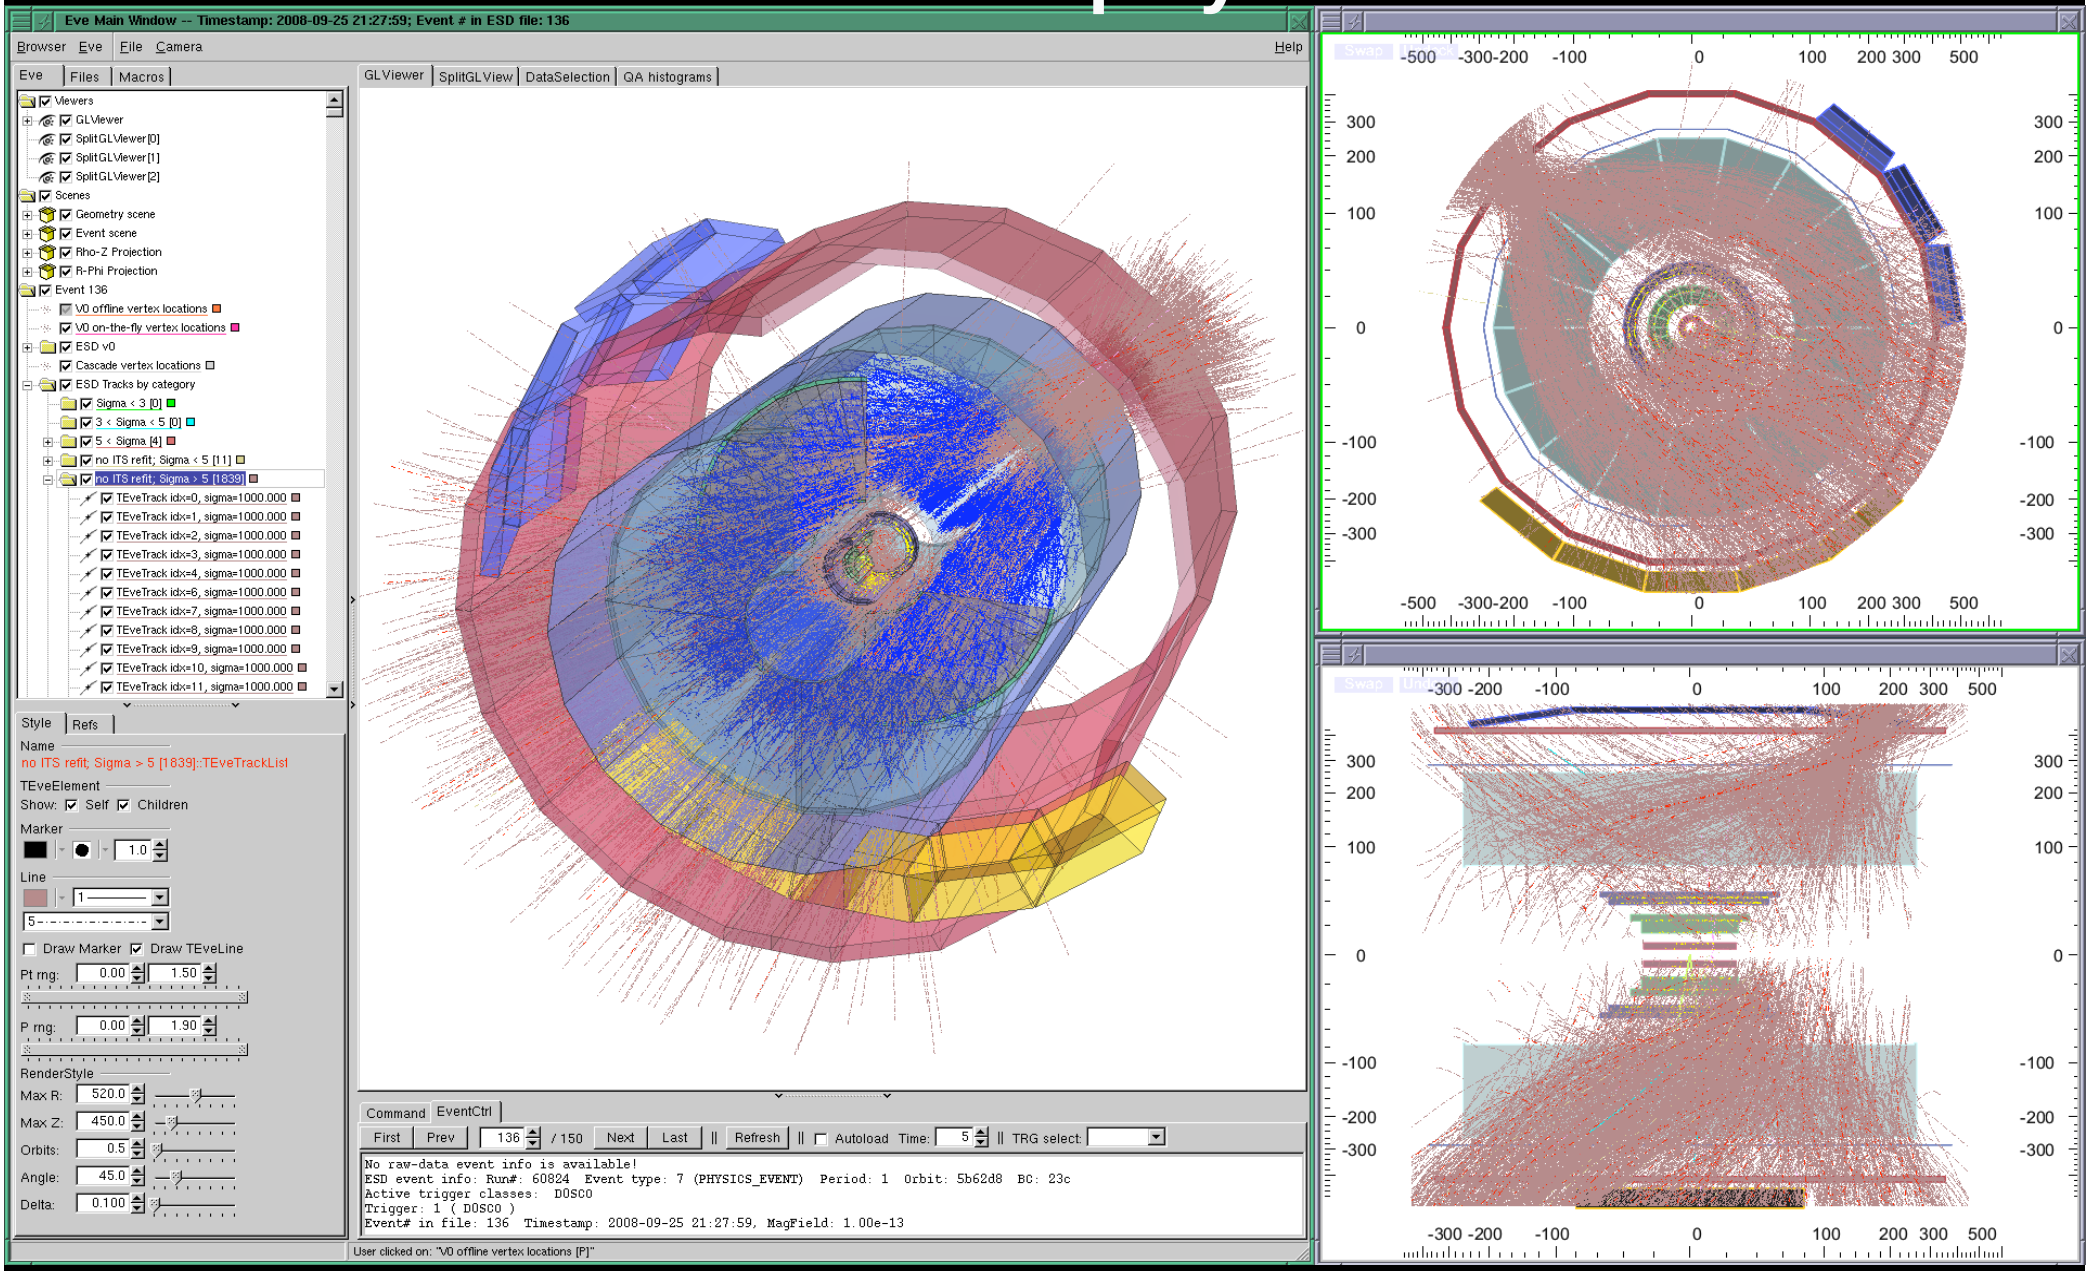

cosmic shower with ~2000 reconstructed tracks

Cosmic track in the MUON arm: 25/03/09

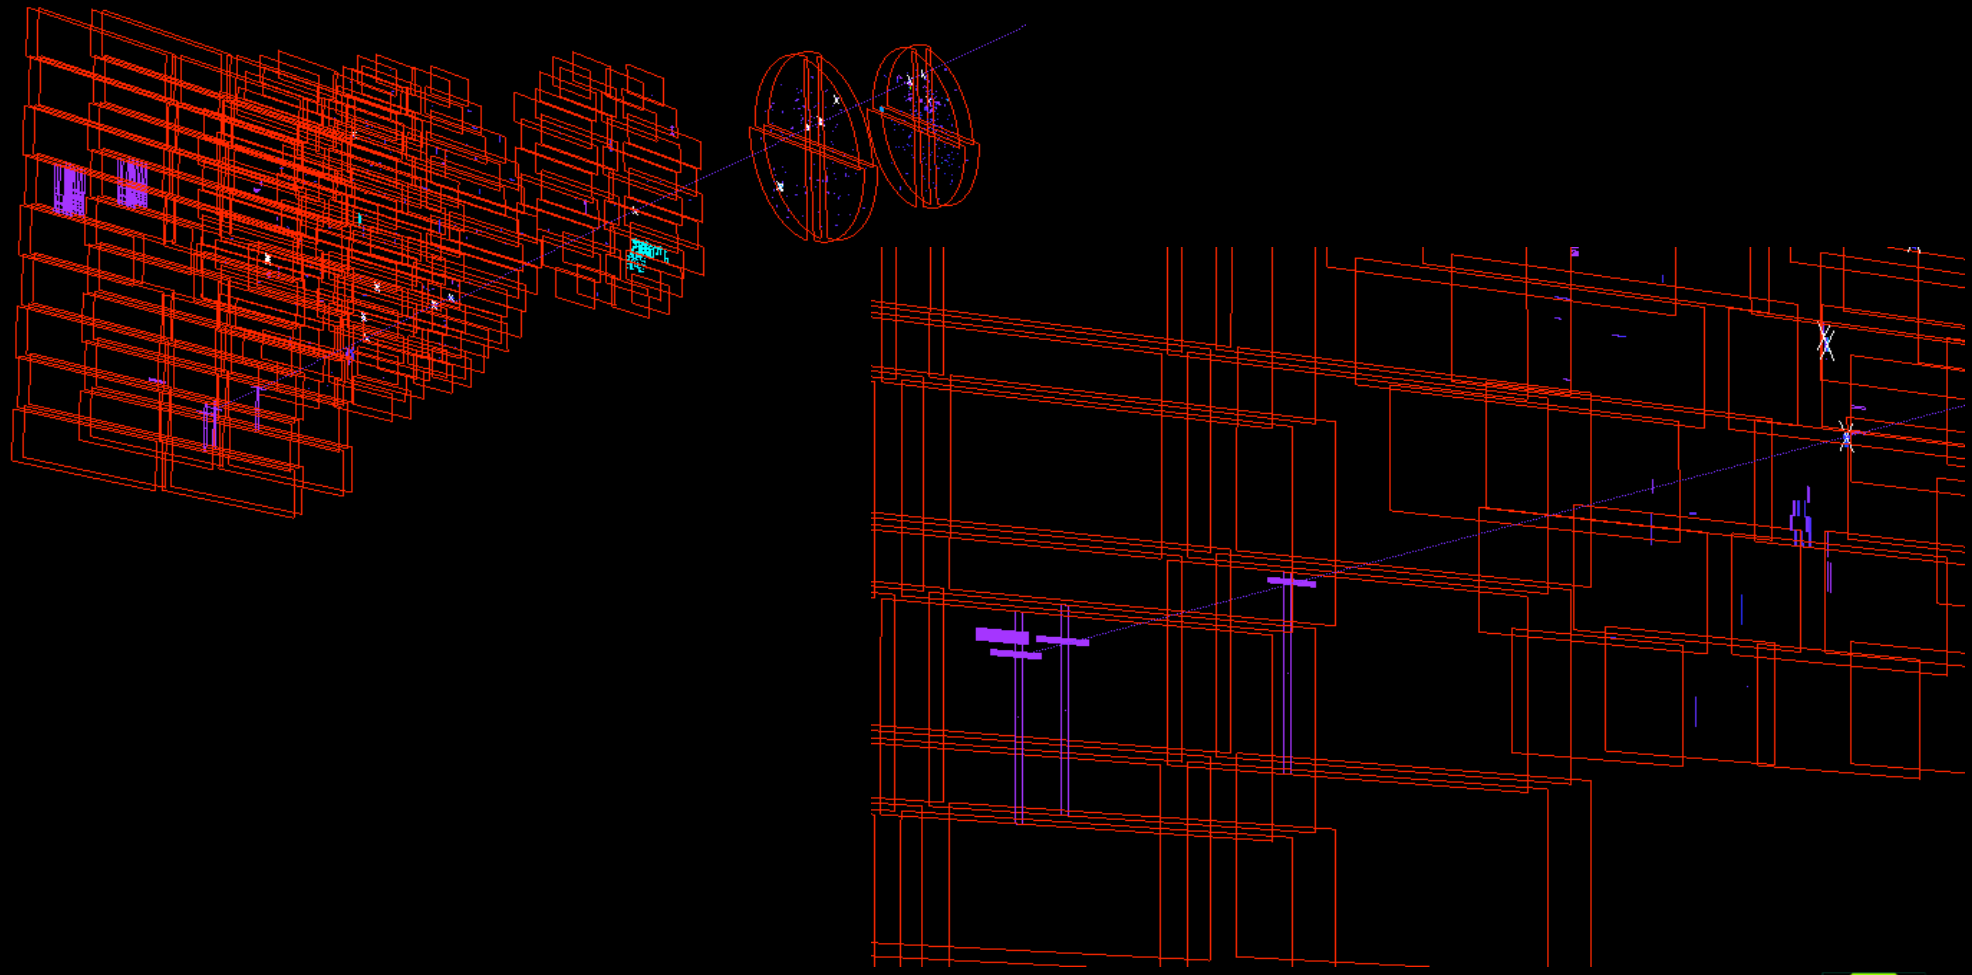

V0 Visualization

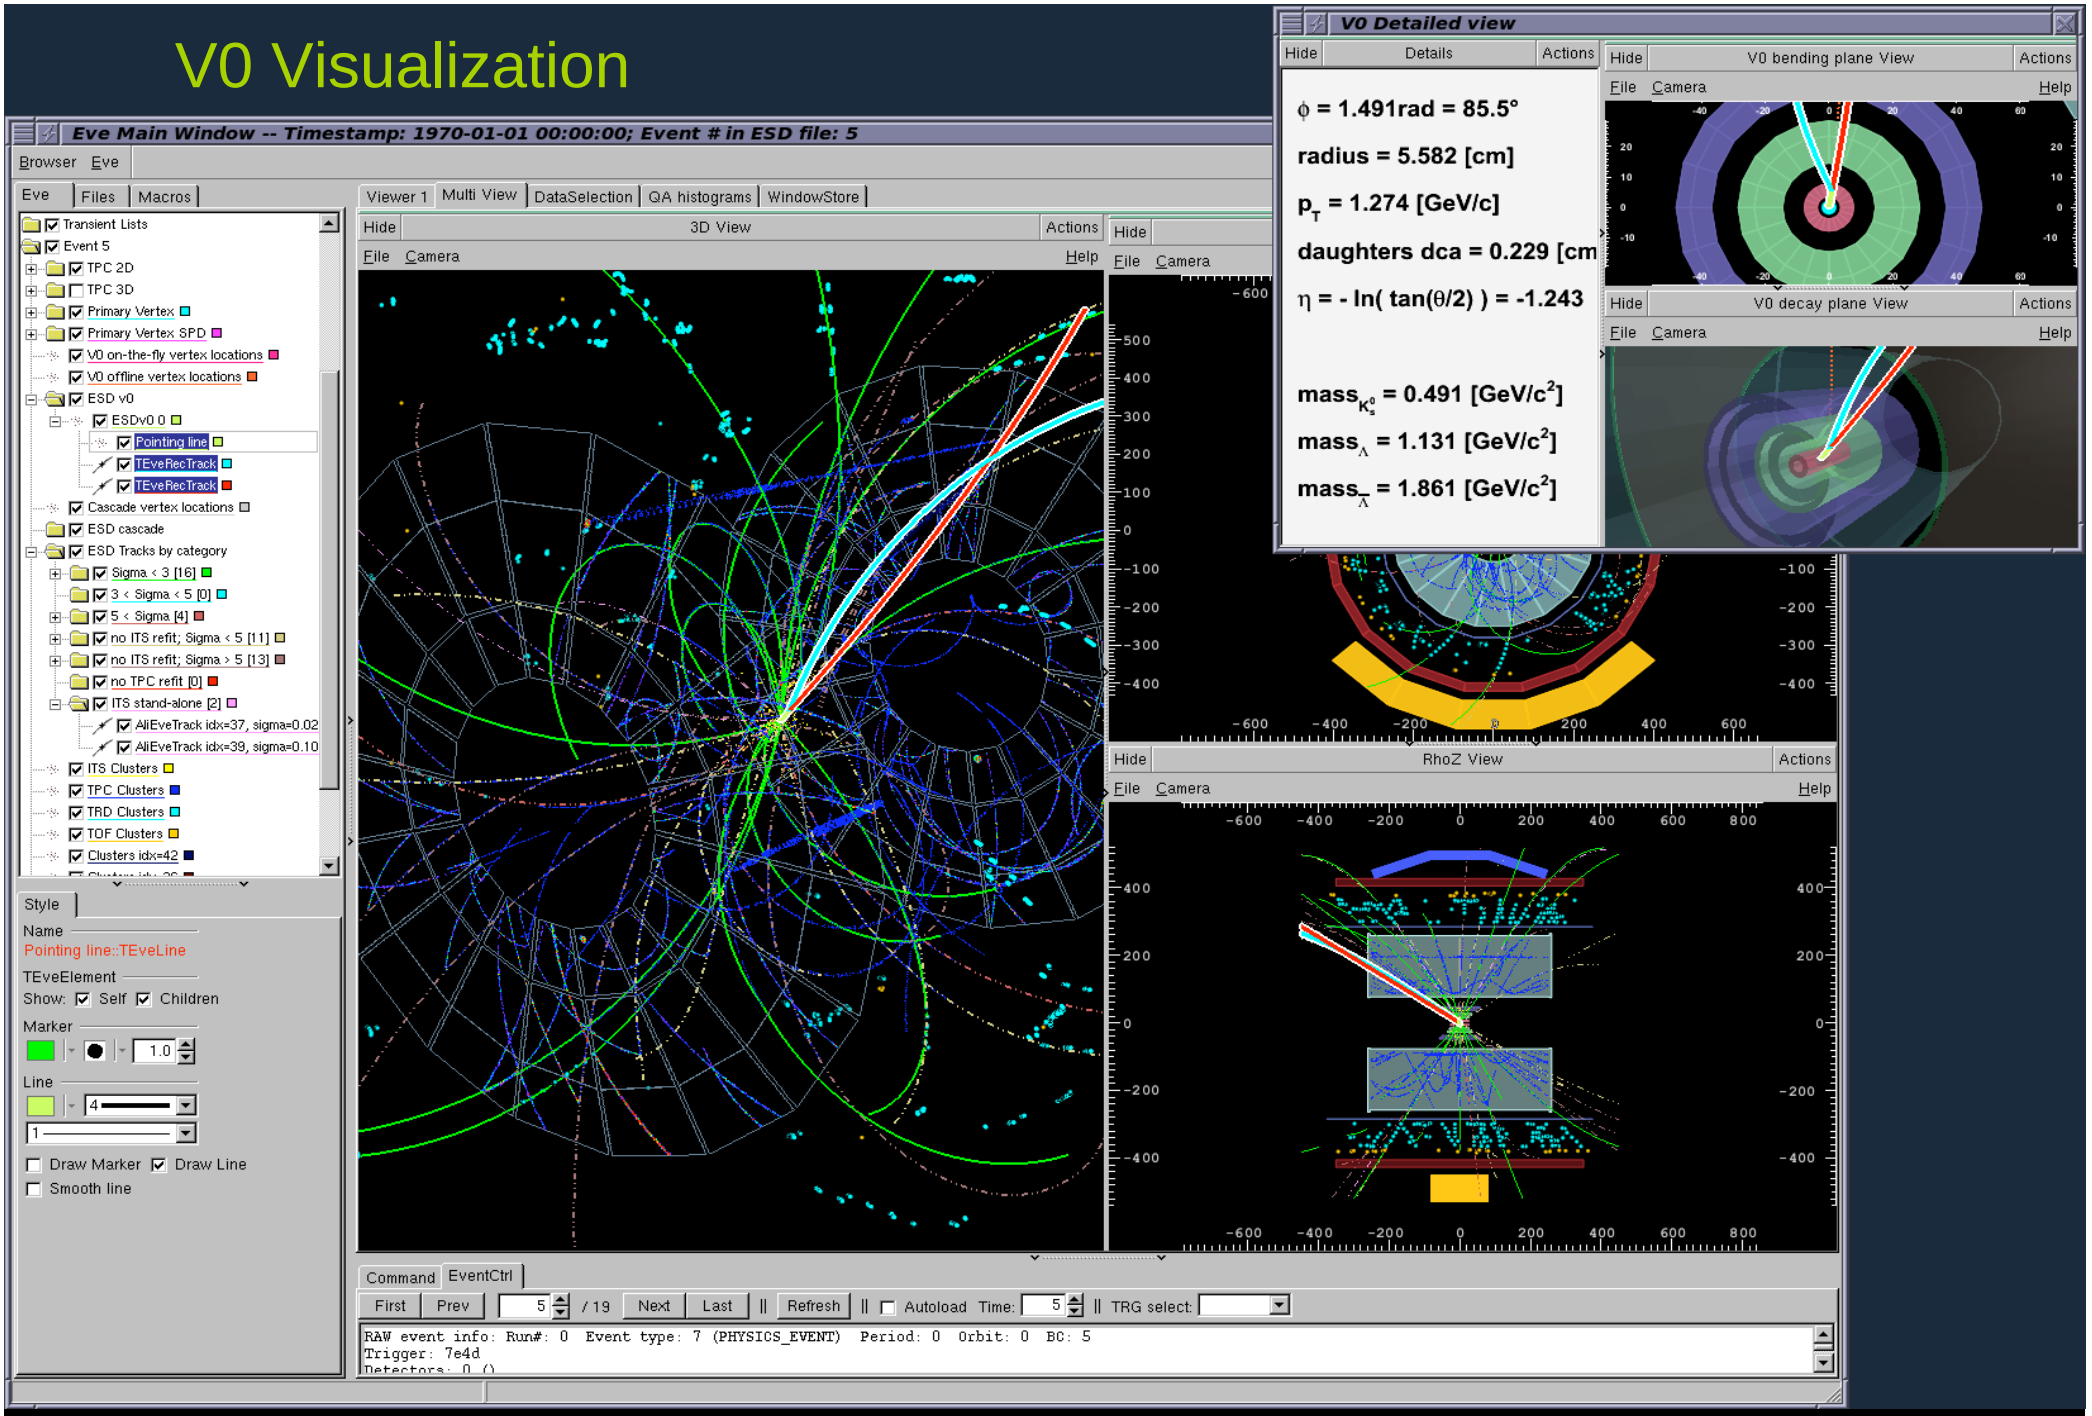

Alignment

Two approaches:
- Millepede/Millepede2
- Iterative (Rieman fit)

Millepede VS Iterative:
track-to-track Δxy at y = 0 (SPD only)

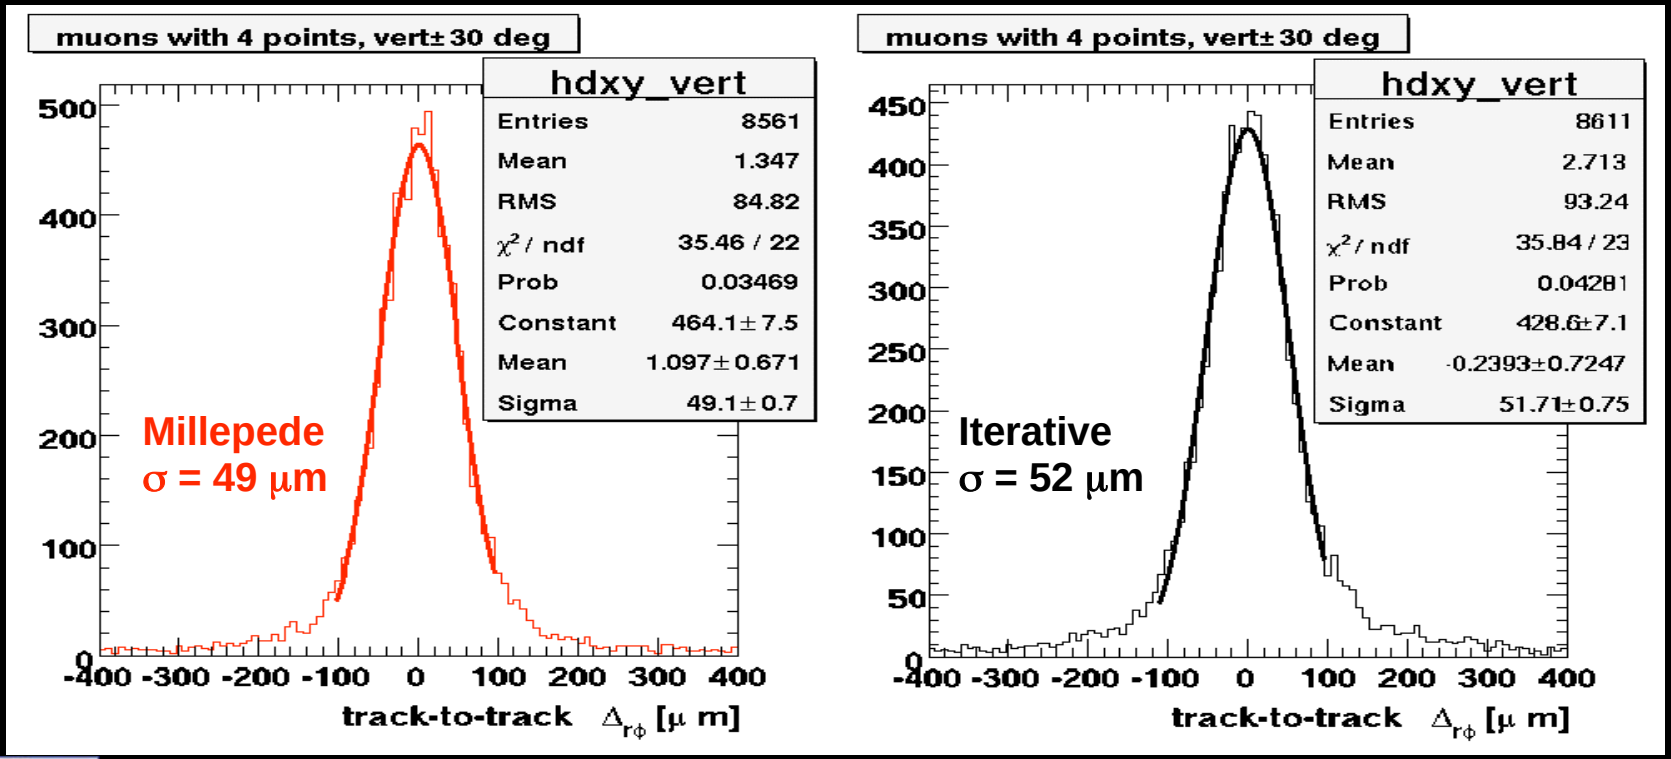

Analysis train

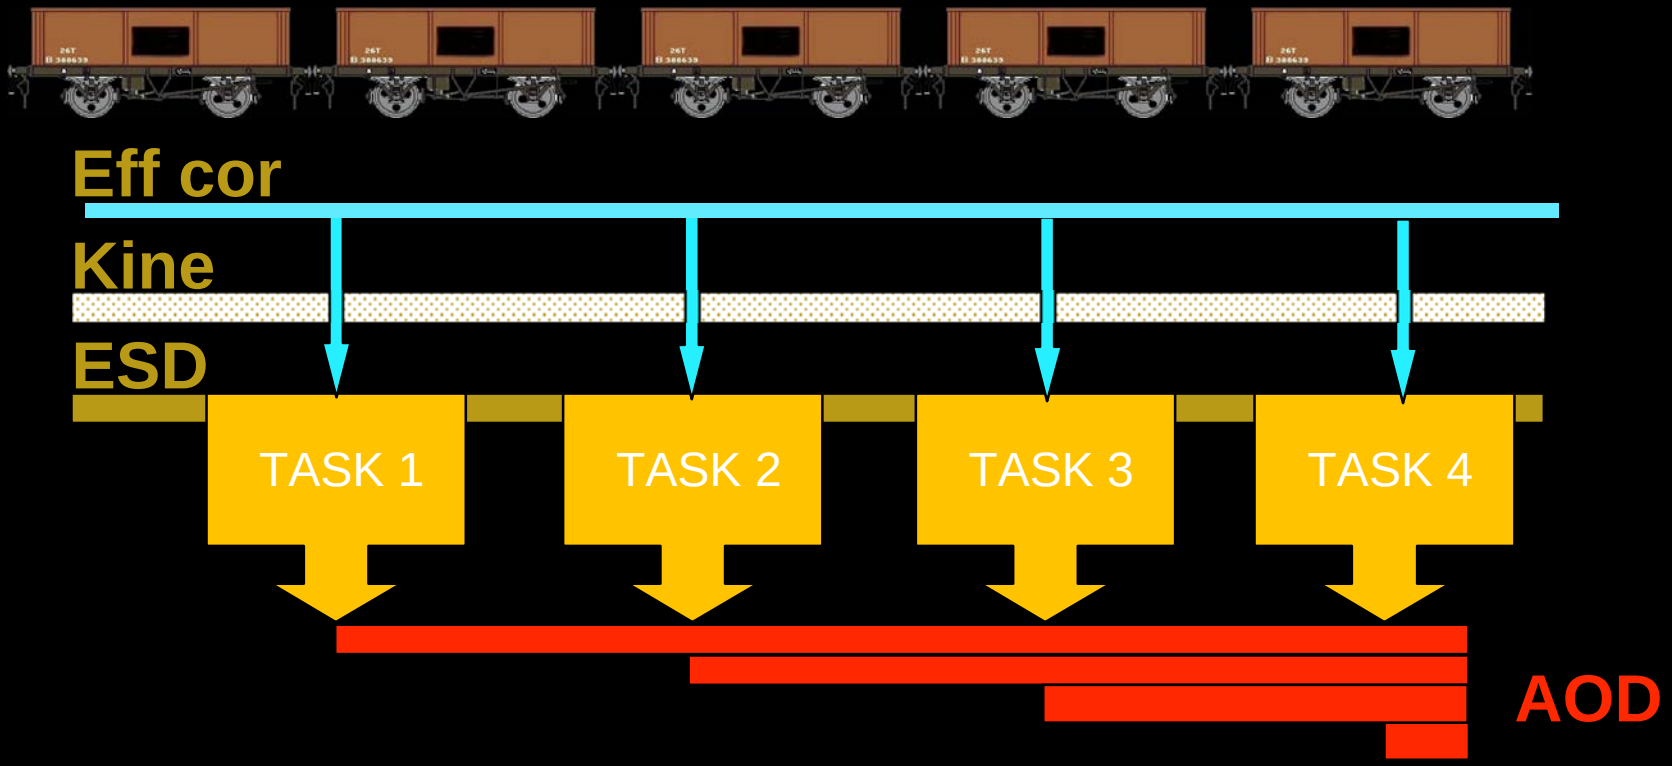

- AOD production will be organized in a 'train' of tasks
  - To maximize efficiency of full dataset processing
  - To optimize CPU/IO
  - Using the analysis framework

Analysis train: a transparent approach

PROOF SETUP

TProof::Open("user@lxb6046")
gProof->UploadPackage("pack.par")
gProof->EnablePackage("pack")
....

StartAnalysis("proof")

+

MY MACHINE

MyAnalysis.C

StartAnalysis("local")

+

AliEn SETUP

CREATE + CONFIGURE GRID PLUGIN

StartAnalysis("grid")

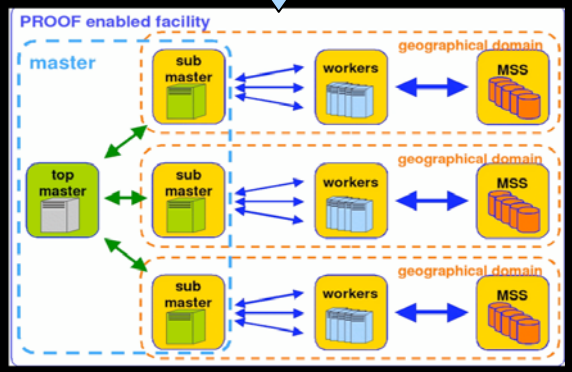

MyResults.root

AliEn 2 @GRID

Analysis train: experience so far

- The framework was developed during the last 2 years and fully adopted by ALICE users
  - Mostly integration efforts, a lot of feedback from users
  - Framework became very stable in all modes
- Very good CAF experience, stability still suffers for GRID analysis jobs
  - 5-10 concurrent CAF users daily
- Simplified procedure to include existing analysis modules in a train and run it in AliEn
  - Self-configured cars (wagons) improve efficiency
  - Light analysis module libraries will be migrated much more frequent in GRID than our offline software

Resource overview

Missing ~40% of financial resources

Parameter	Now	CTDR	Ratio
pp RAW	1.0MB	0.2MB	5*
Pb RAW	35MB	13.8MB	2.5
ESD pp	0.04MB	0.04MB	1.0
ESD Pb	6.3MB	3.0MB	2.1
AOD pp	5kB	16kB	0.3
AOD Pb	1.3MB	0.34MB	3.8
Reco pp	6.8s	6.5s	1.0
Reco Pb	800s	810s	1.0

No Root compression yet

Outside CERN!

* was 22!!

		2008		2009		2010		2011		2012	
		T1	T2	T1	T2	T1	T2	T1	T2	T1	T2
CPU	Requested	7.2	4.6	11.2	17.4	23.56	25.11	31.41	33.48	41.88	44.63
	Missing	-23%	84%	-5%	-26%	-42%	-28%	-45%	-30%	-49%	-38%
Disk	Requested	2,151	1,217	9,363.7	9,950.0	7,973	10,256	10,630	13,674	14,173	18,232
	Missing	11%	41%	-58%	-65%	-15%	-53%	-9%	-54%	-17%	-58%
MSS	Requested	2,431		11,704.9		20,788		29,932		39,076	
	Missing	23%		-44%		-52%		-55%		-54%	

Resources

- There is a serious deficit in the Computing Resources pledged to ALICE
- We have considered alternative scenarios where we fit within the available resources / investments
- A reduction in the MC will have adverse effects
- A reduction in the number of reconstruction passes is very risky and may lead to reduced quality of physics

Data taking scenario

- Cosmics
  - Resume data taking in July 2009, ~300TB of RAW
- p+p runs
  - Running at maximum DAQ bandwidth
    - Few days @ 0.9 GeV (October 2009)
    - 11 months @ 10 TeV
  - Machine parameters at P2 - optimum data taking conditions for ALICE
  - Computing resources **must be sufficient** for quasi online processing
  - Address the ALICE genuine p+p physics program and provide baseline measurements for AA

Data taking scenario (2)

- A+A run
  - Fall 2010 - a standard period of Pb+Pb running
  - Computing resources **must be sufficient** to process these data within 4 months after data taking (as foreseen in the Computing Model)
  - Results to be presented at QM@Annecy (the LHC QM) in Spring 2011
- Monte Carlo
  - 2009-2010 are standard years for Monte Carlo production

Summary

- Good progress in the development of the ALICE offline software
  - Mature simulation code
  - Stable, fully operational software for conditions data
  - Improved reconstruction, possibility for fast feedback
  - Powerful visualization based on EVE
  - Operational alignment
- Successful processing of the cosmic data in 2007-2008
- Computing model verified with cosmic and simulated data
- Possible shortage of resources in 2009-2010
- Ready for the first LHC collisions!

Backup

Computing resources

- Computing Resources (new requirements following the LHC scheduled announced after Chamonix)
  - 10 months of continuous pp running with an average data rate 3 times larger than the one in a standard year of data taking → 50% more data that impacts mainly storage but less CPU needs
  - 1 month of PbPb running equivalent to a standard year of data taking
  - Reduction of Monte Carlo for PbPb
  - The availability of resources, in particular for the PbPb data, remains a worrisome issue
  - No major new contributions anticipated

Table 2.: CPU requirements for 2009-2010 and comparison with previous requirements

	TO new	CAF	T1	T2	T0	CAF	T1	T2	T0	CAF	T1	T2
	new requirements (MSI2K)				old requirements (KSI2K)				variation (%)			
2009Q1	7,9	2,6	8,0	8,1								
2009Q2	7,9	2,6	8,0	8,1	9,1	2,6	19,9	14,3	-11 %	1 %	-55 %	-37 %
2009Q3	7,9	2,6	8,0	8,1								
2009Q4	8,1	2,6	10,7	9,0								
2010Q1	8,4	2,6	10,7	9,0								
2010Q2	8,4	2,6	10,7	9,0	9,1	2,6	23,6	25,1	0 %	0 %	9 %	-19 %
2010Q3	8,5	2,6	10,7	9,0								
2010Q4	9,1	2,6	25,6	20,2								

Table 3.: Disk requirements for 2009-2010 and comparison with previous requirements

	CERN	T1	T2	CERN	T1	T2	CERN	T1	T2
	new requirements (PB)			old requirements (PB)			variation (%)		
2009Q1	1,7	2,4	1,7						
2009Q2	1,9	3,0	2,6	2,5	9,9	9,6	-4 %	-56 %	-54 %
2009Q3	2,2	3,6	3,5						
2009Q4	2,4	4,3	4,4						
2010Q1	2,6	4,9	5,3						
2010Q2	2,9	5,5	6,2	4,2	9,9	10,3	8 %	-0 %	21 %
2010Q3	3,1	6,1	7,0						
2010Q4	4,5	9,9	12,4						

Table 4.: Custodial Storage (integrated) requirements for 2009-2010 and comparison with previous requirements

	CERN	T1	CERN	T1	Tape	T1
	new requirements (PB)		old requirements (PB)		variation (%)	
2009Q1	3,3	2,4				
2009Q2	3,4	3,6	7,7	10,6	-52 %	-44 %
2009Q3	3,6	4,7				
2009Q4	3,7	5,9				
2010Q1	4,1	7,0				
2010Q2	4,6	8,2	8,1	19,7	-18 %	-41 %
2010Q3	5,0	9,3				
2010Q4	6,7	11,6				

Requirements vs pledges

		2009		2010	
		T1	T2	T1	T2
CPU	Requested	10,7	9,0	25,6	20,2
	Missing	-4%	43%	-41%	-0%
Disk	Requested	4 263,54	4 380,9	9 869	12 365
	Missing	-9%	-6%	-31%	-52%
MSS	Requested	5 887,01		11 648	
	Missing	5%		-12%	