



## Outline

- The GSI farm
  - Cores
  - Data storage
  - People
- Calibration
- Analysis train
  - Composition
  - Schedule and performance



## GSI-wide

(1000 employees, effectively: Hades, CBM, PANDA, theory ..)

- **2500** active cores
- ALICE sponsored 1800 cores, now effectively **1440** (aging) with a special BMBF grant for GSI Tier2/Tier3
- **800 cores** blocked at all times for GRID jobs (**Tier2**, pledged)
- Cores available for **calibration and analysis**:
  - From a minimum of ~200 (sharing GSI-wide)
  - To a maximum of ~1700 (free farm)

# GSI farm: storage and network



Data storage is also GSI wide: **Lustre system**, 1 PByte

ALICE data:

	N. of events ( $10^6$ )	Data set size (TBytes)
2010 pp data, pass 2	320	80
2010 pp MC	280	75
2010 PbPb data, pass 1	24 (all)	60
2010 PbPb MC	0.2	2
<b>Total</b>		<b>217 TB</b>

Essential component: adequate network!!!



- Linux / High Performance Cluster  
(GSI wide): 4-5 people involved, **~2-3 FTE**
- ALICE group  
4-5 people deeply involved, **<1 FTE**
- Weekly computing meetings since May 2008
- ALICE users: about 60-70 people  
(calibration and analysis train)



Primary goal of the extra computing resources is the

## **DEVELOPMENT of the detector calibration procedures**

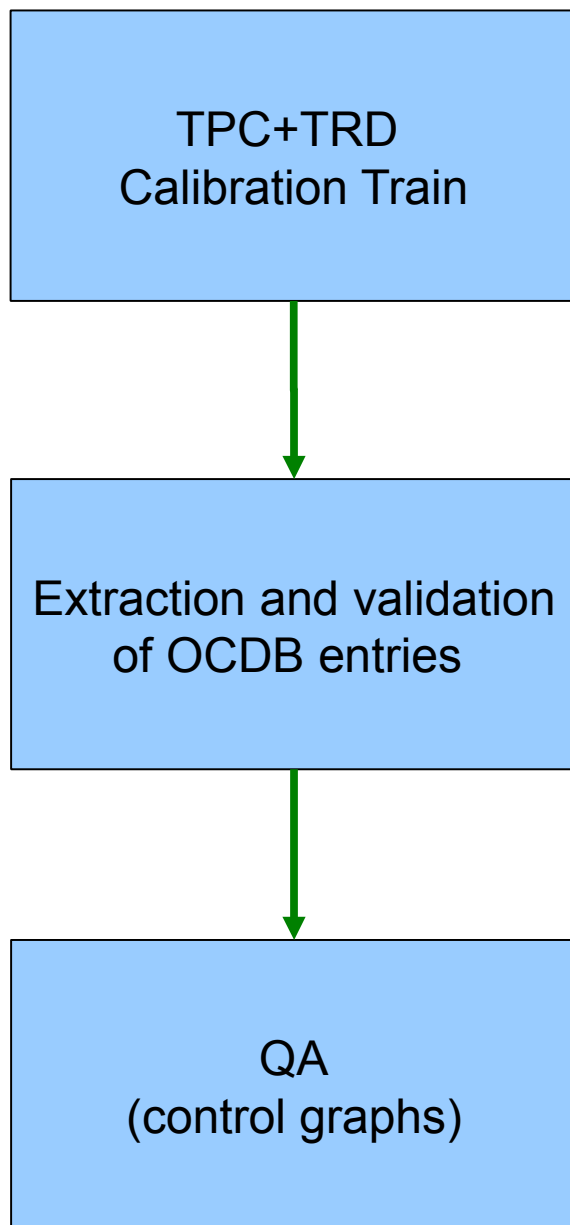
- This requires fast and frequent iterations over a subset of data (over the WHOLE data taking period)
- Iterative improvement of the calibration algorithm
- When well understood and stable, the calibration procedure can become one stable step in the chain on GRID (e.g. Pass0)

**This has been crucial for the readiness of the calibration of TPC and TRD on a relatively short time scale**

# Pass0 Complexity

- **Pass0 functionality**
  - TRD and TPC calibration (gain, v-drift, t0 ...)
  - Validation of the calibration values
  - Extraction of the OCDB entries
- **Input**
  - ESD and ESDfriends
  - OCDB entries from the previous run
- **Output**
  - Calibration objects
  - QA information

# Pass0 Development at GSI



- **Analysis flow (as many iterations as needed)**
  - Production of calibration object (TPC + TRD calibration train)
  - Creation of OCDB entries (merging, validation, extraction)
  - QA for OCDB entries (trend graphs)
- **The calibration algorithms successfully tested on LHC10b, LHC10c (10% statistics) before running on the GRID**

- ~40 directories  $\approx$  “modules” (analyses)
- Normally 20-40 tasks
- Output used by a very large community:
  - single users
  - analysis groups:
    - $\gamma$  conversion
    - identified particle spectra
    - HFE
    - $J/\psi$
    - $R_{AA}$
    - HBT
    - Check of total cross section by ALICE (!)
- All official analyses are in AliRoot and on the central trains
- Software: AliRoot revisions + local development of own task (same as if running privately with the plugin on GRID)






- 4 train flavours: pp data, pp MC, PbPb data, PbPb MC
- A data set is processed:
  - Per run
  - In jobs of e.g. 100 pp data chunks, 200 pp MC ESD files
- **RECURSIVE MERGING:**
  - Per run
  - Per LHC period
- Reporting (train performance, basic quantities)
- Test and feedback to calibration!

# Analysis train: test of calibration



Just one simple example:

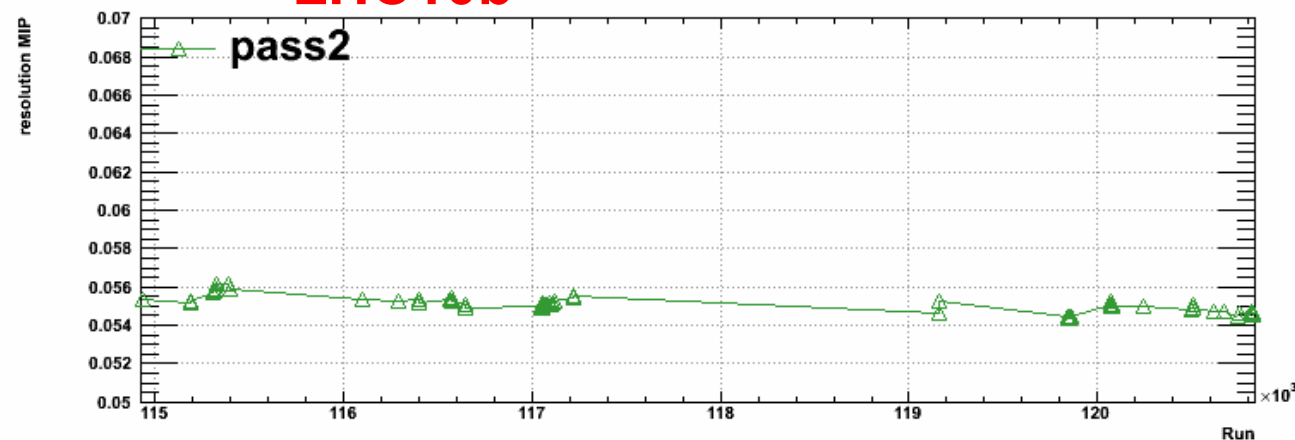
TPC dE/dx:  
MIP position



LHC10b

LHC10c

TPC dE/dx  
resolution



# Analysis train: performance



Time from job submission to the very last individual merging  
("average" train, about 800 cores)

	N. of events ( $10^6$ )	Time for the complete train	Schedule
2010 pp data, pass 2	320	Order 6 hours	Each Tuesday, Each Friday
2010 pp MC	280	Order 8 hours	Each Tuesday, Each Friday
2010 PbPb data, pass 1	24	Order 5 hours	Each Monday, Each Thursday
2010 PbPb MC	0.2	Order 2 hours	Each Monday, Each Thursday

**EACH TRAIN TYPE → 2 TIMES PER WEEK**

# Analysis train: goodies



- A fixed schedule helps people to organize their development
- With the frequent running, there is no longer the need for large scale individual analysis
- “Educational” effect: Strict monitor of resource usage by analysis!!

**Add one task at the time**

**When jumps in CPU or memory → task out (till the next train, in 3 or 4 days)**

