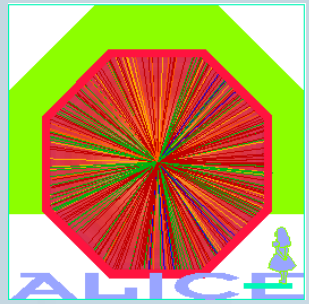

Databases in ALICE

L.Betev

LCG Database Deployment and Persistency Workshop

Geneva, October 17, 2005

Outline

- Transversal field and function of ALICE databases
- Examples of database applications
- Conditions DB and data collection models
- Medium term plans
- Conclusions

Transversal field of the ALICE DBs

- Detector Control System (DCS):
  - Application scope: configuration of systems and devices (modules and channels), front-end configuration (busses, thresholds); Archiving of monitored detectors and devices parameters
  - Size: millions of records, Tera bytes
- High Level Trigger (HLT):
  - Application scope: mini-DST like TAG/ESD database for physics studies and offline event selection
  - Size: up to 10^9 events and 30TB per year
- Detector Construction DB (DCDB):
  - Application scope: use by individual sub-detector groups and integration, repository and flow management for modules, components and their test data, cables, racks
  - Size: millions of records, Tera bytes
- Experimental Control System (ECS):
  - Application scope: inclusion/exclusion of sub-detectors to a partition
  - Size: small number of small records

Transversal field of the ALICE DBs

- Data Acquisition (DAQ):
  - Application scope: parameter repository and resources assignment to DAQ tasks: configurations (current and stored), run parameters (current and stored)
  - Size: possibly large number of small records
- Trigger:
  - Application scope: repository for trigger classes (input to Central Trigger Processor), definition of trigger masks
  - Size: large number of small records
- Offline Conditions DB (CDB) – coupled to the Grid:
  - Application scope: non-event data for offline data reconstruction and analysis
  - Size: millions of files, Tera bytes
- Distributed file catalogue, metadata and TAG DBs:
  - Application scope: raw data, production files, software catalogue and event tags
  - Size: hundreds of millions of records, Giga bytes

Database applications - DCDB

- Contains detector properties gathered during constructions time
- Each detector group has one instance of DCDB with (detector) specific dictionary (PostgreSQL)
- Read only copy at CERN (Oracle)
- Content
  - Description of detector components
  - Measurements of response
  - Location of detector components
  - Geometry
- Usage & lifetime
  - Input for calibration and geometry database
  - Lifetime of ALICE

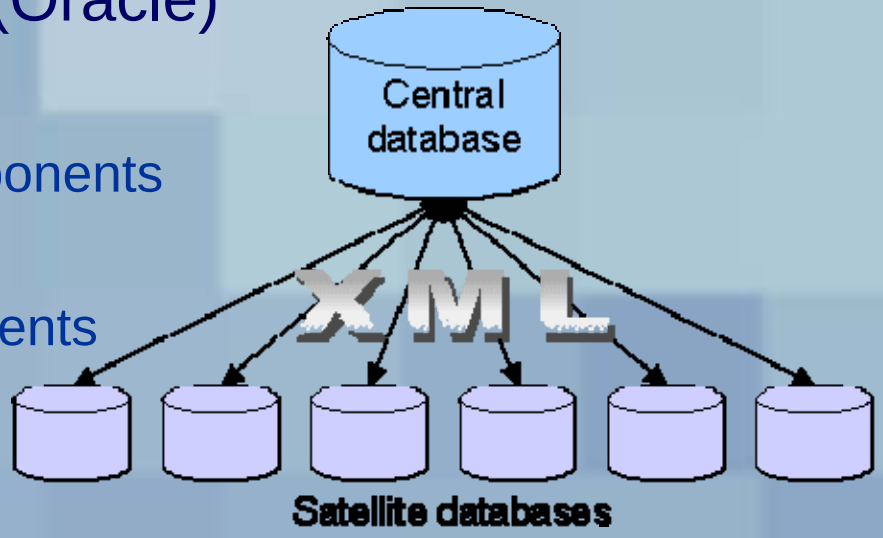

Database applications – DAQ and ECS logbook

- Contains messages stored by DATE infoLogger, statistics extracted from logbook stream for each run

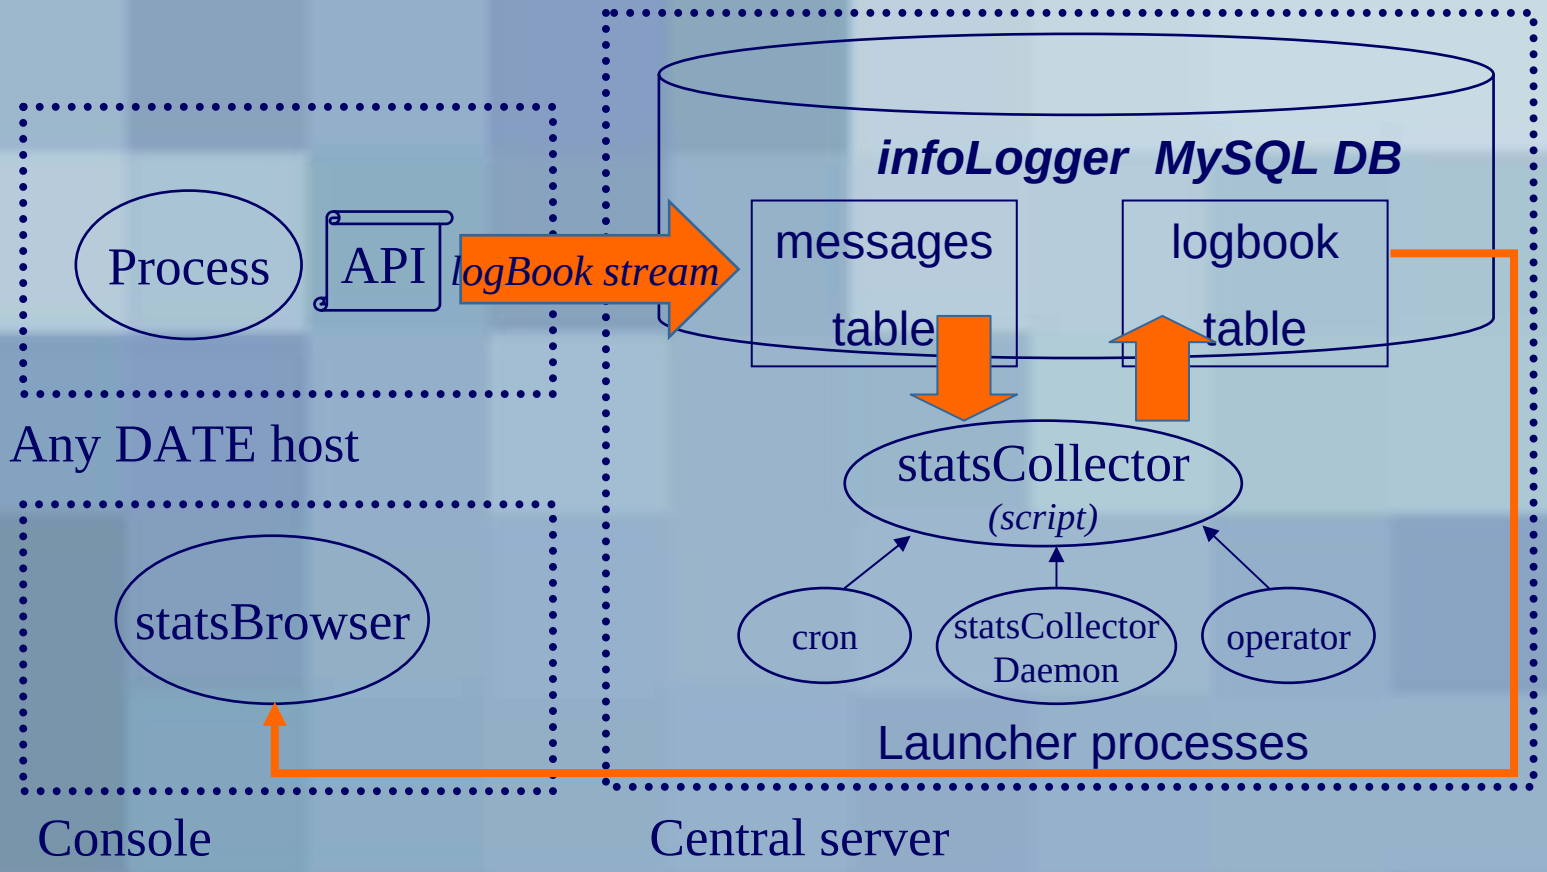

Databases relations and connectivity

- Examples from previous slides:
  - DCDB – “quasi” distributed system, central repository is read-only, updates are infrequent
  - DAQ – completely closed central system
- ***The biggest challenges remain the Grid file catalogue and offline conditions databases:***
  - Used in a heterogeneous and often ‘hostile’ computing environment (the Grid).
  - Contains data from wide variety of sources

Offline conditions DB relations

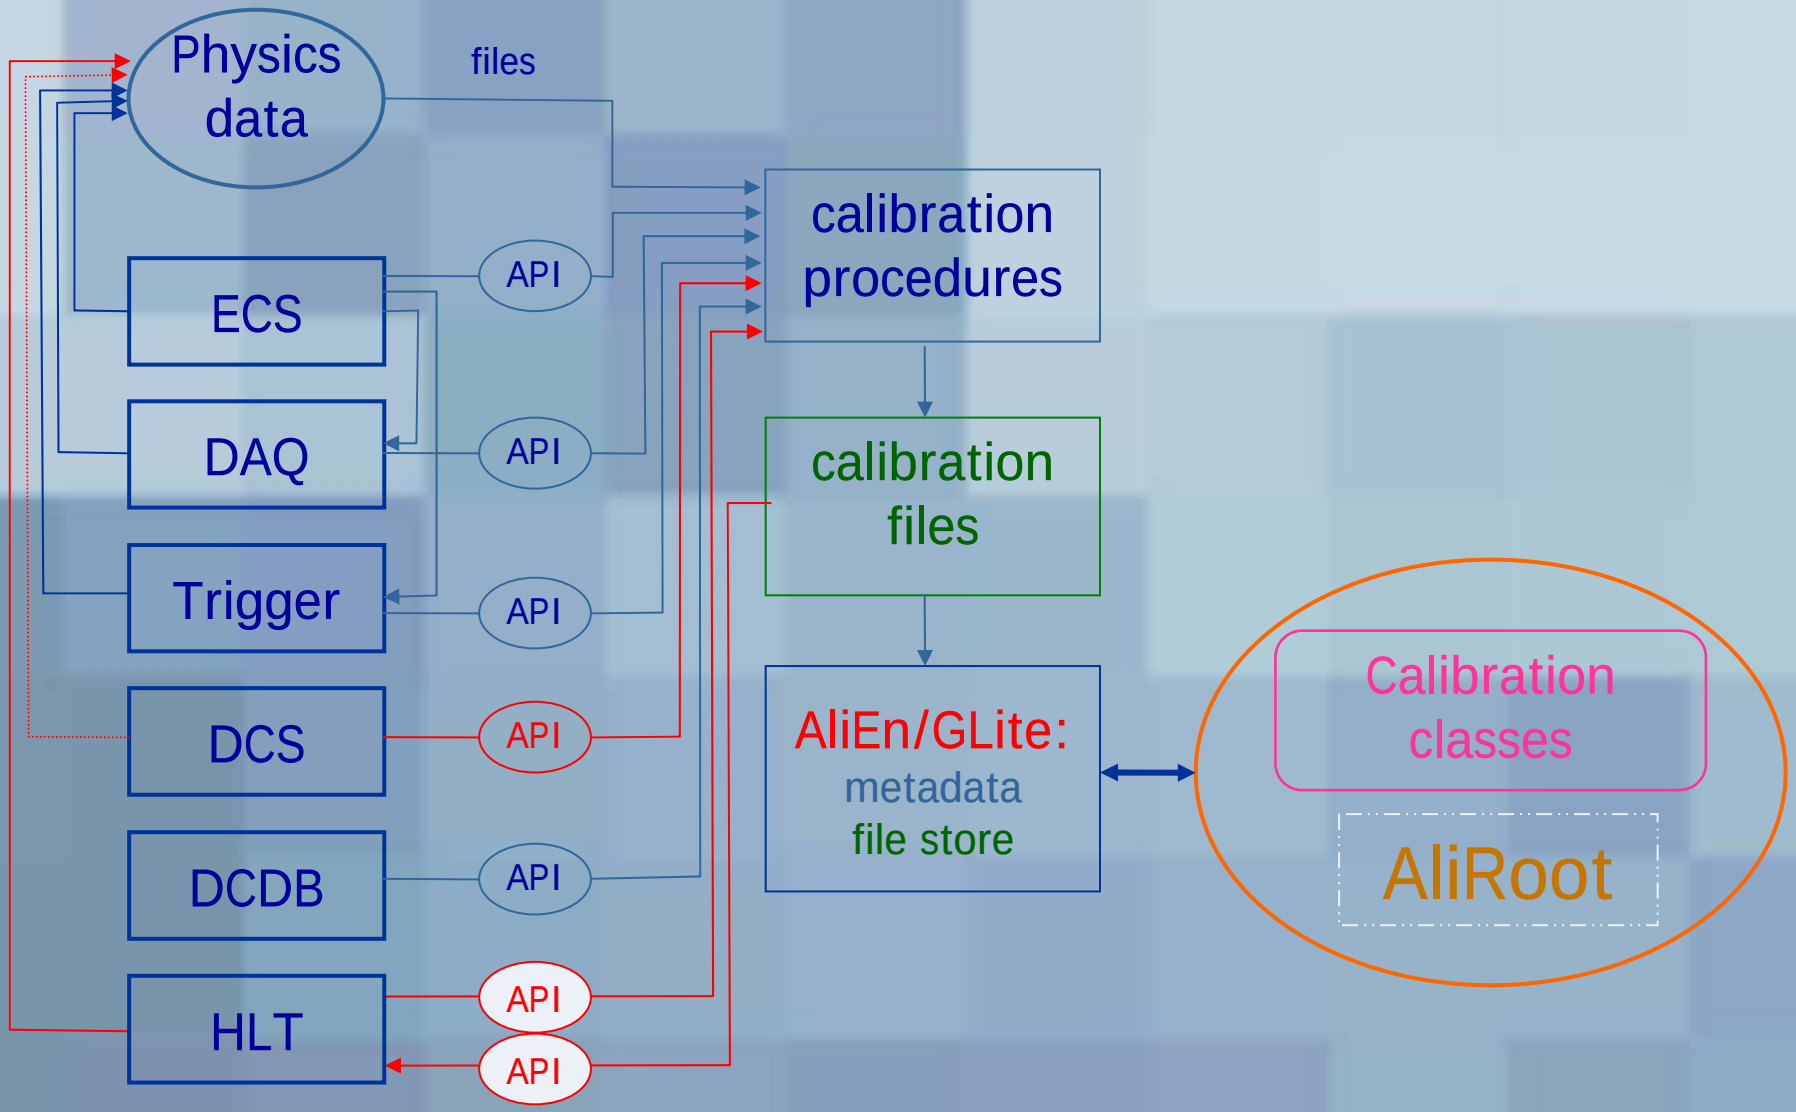

Considerations for Conditions DB

- Sources for conditions DB and relations to other DBs are already quite complicated:
  - All databases potentially containing conditions information are “closed”, i.e. only accessible at CERN
  - It would be difficult to provide access methods to all DBs from the production and analysis code
- ALICE uses ROOT as offline framework base:
  - Naturally defines the technology choice for object store – all conditions data are stored as ***root files***
  - Additional condition - these files are read-only

Considerations for Conditions DB (2)

- Conditions data should be accessible in a Grid distributed production and analysis environment
  - The root files are registered in the ***Grid Distributed File Catalogue***:
  - ***No need*** for distributed DBMS in a traditional sense and with all accompanying problems (access, replication, authentication) – a very big plus
  - Assures worldwide access to all files and associated tags
- Drawbacks:
  - Replication of information
  - Depends almost entirely on the Grid file catalogue functionality

Conditions DB – content and access

- Update frequency:
  - The condition DB will have rather infrequent updates (max: ***once per run***)
  - The objects themselves can have rich internal structure with finer granularity, depending on the detector calibration / alignment needs
- Access framework:
  - Simple user interface and data identification (strings)
  - Automatically get the right valid object
  - Modification of the objects is handled through versioning
  - Same interface for different types of data sources (local file, Grid)

Conditions DB – content and access (2)

- Data organisation - simple filesystem-like storage:
  - Calibration/DETECTOR/FirstRun-LastRun.version

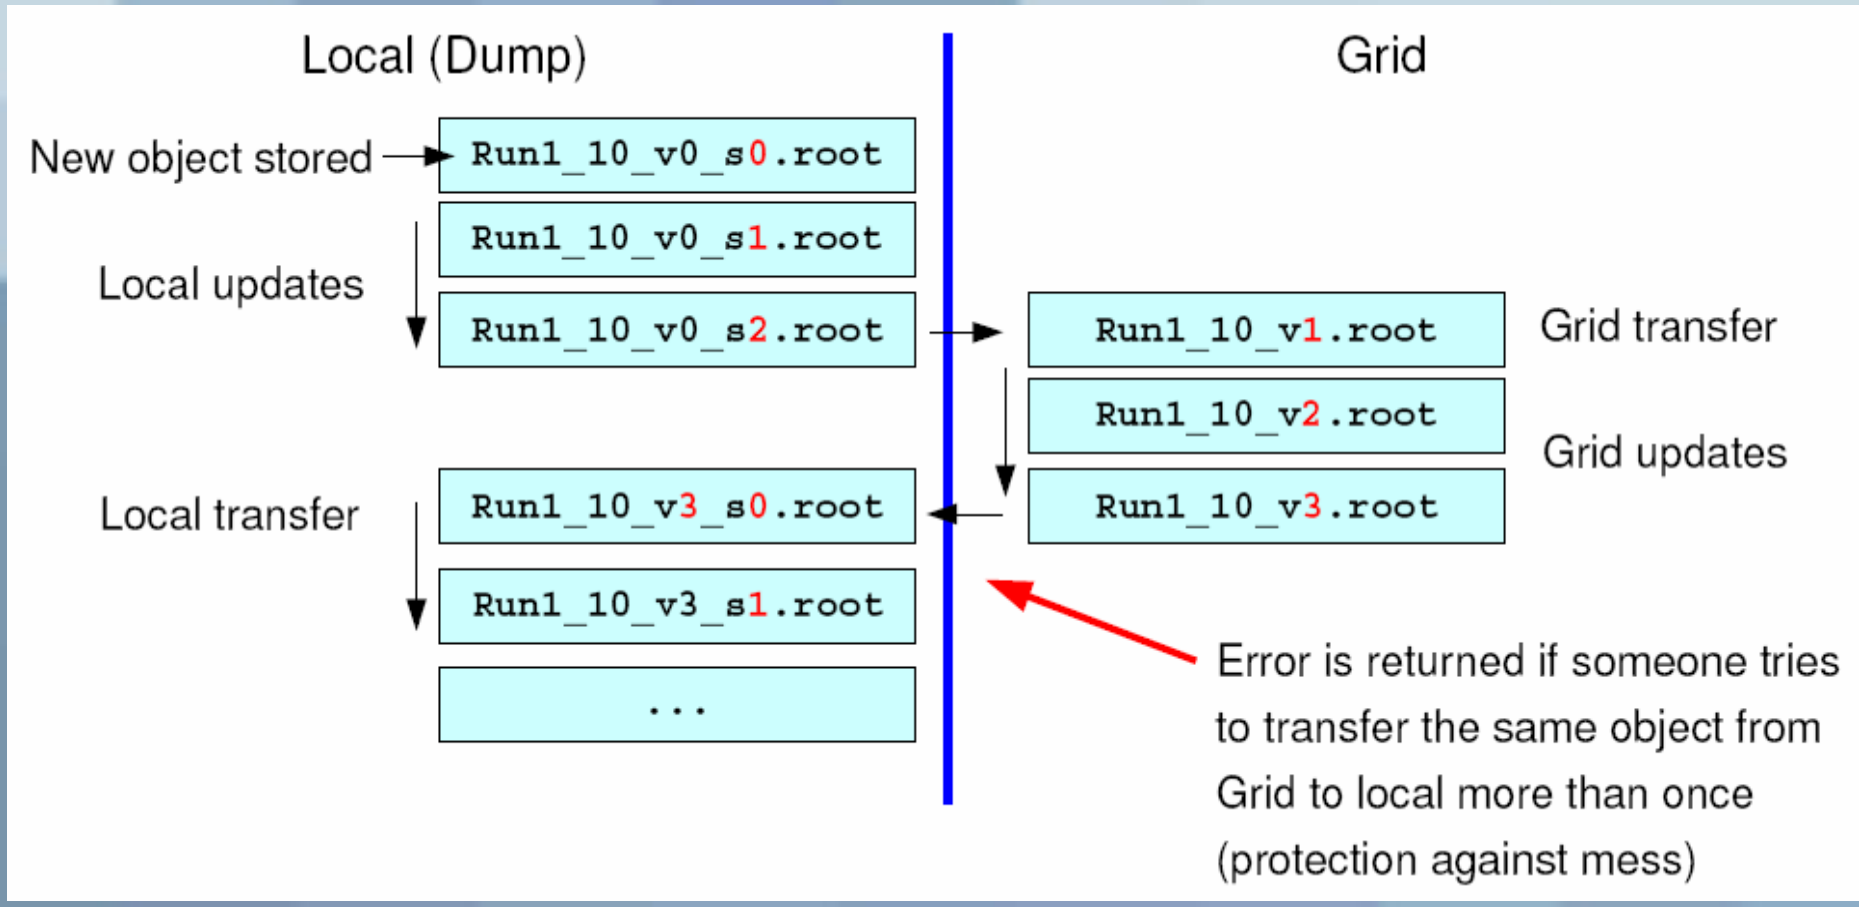

Conditions DB – information gathering

- Automated system – Shuttle:
  - Collecting conditions data from external DB at predetermined time intervals (for example at the end of every run)
  - Processing collected data using detector-specific preprocessor
  - Storing conditions data in the Conditions DB

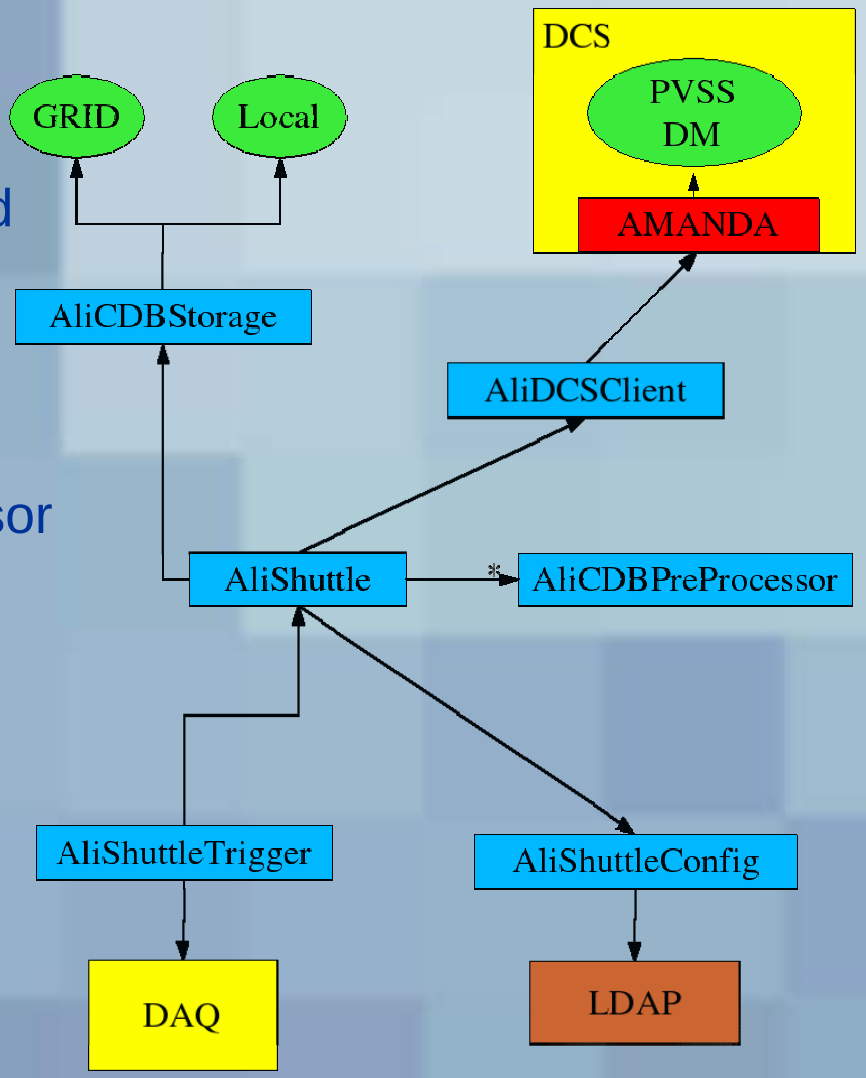

Conditions DB – information gathering (2)

- Similar systems can be used to process conditions data from other sources:
  - Only new interfaces needed
  - The processing of the data is done in AliRoot (no change of framework)
  - Allows detector groups to do analysis and reduction of the raw data volume before storing it in the Conditions DB

Medium term plans

- Grid file catalogue is used routinely in ALICE since several years in the scope of the physics data challenges (PDC04, PDC05, SC3)
- ROOT Grid access classes are now mature
- AliRoot Conditions DB access framework is complete
- Beginning of 2006 – Test of ALICE TPC with operational DCS and DAQ:
  - Test of both DBs and in addition the offline Conditions DB
- Beginning of 2006 – Physics Data Challenge ’06:
  - Final test of the entire ALICE offline framework on the Grid, including test of Conditions framework

Conclusions

- ALICE has a rich field of databases used in the various groups for wide variety of tasks
- The development of the majority of DBs is well advanced:
  - Some of them are already in production since several years
- The biggest challenge remaining is to gather the information from all sources and make it available for Grid data reconstruction and analysis:
  - In this context, ALICE has decided to re-use the already existing Grid file catalogue technology
  - And enrich the AliRoot framework with tools allowing connectivity to all other DBs currently in existence