

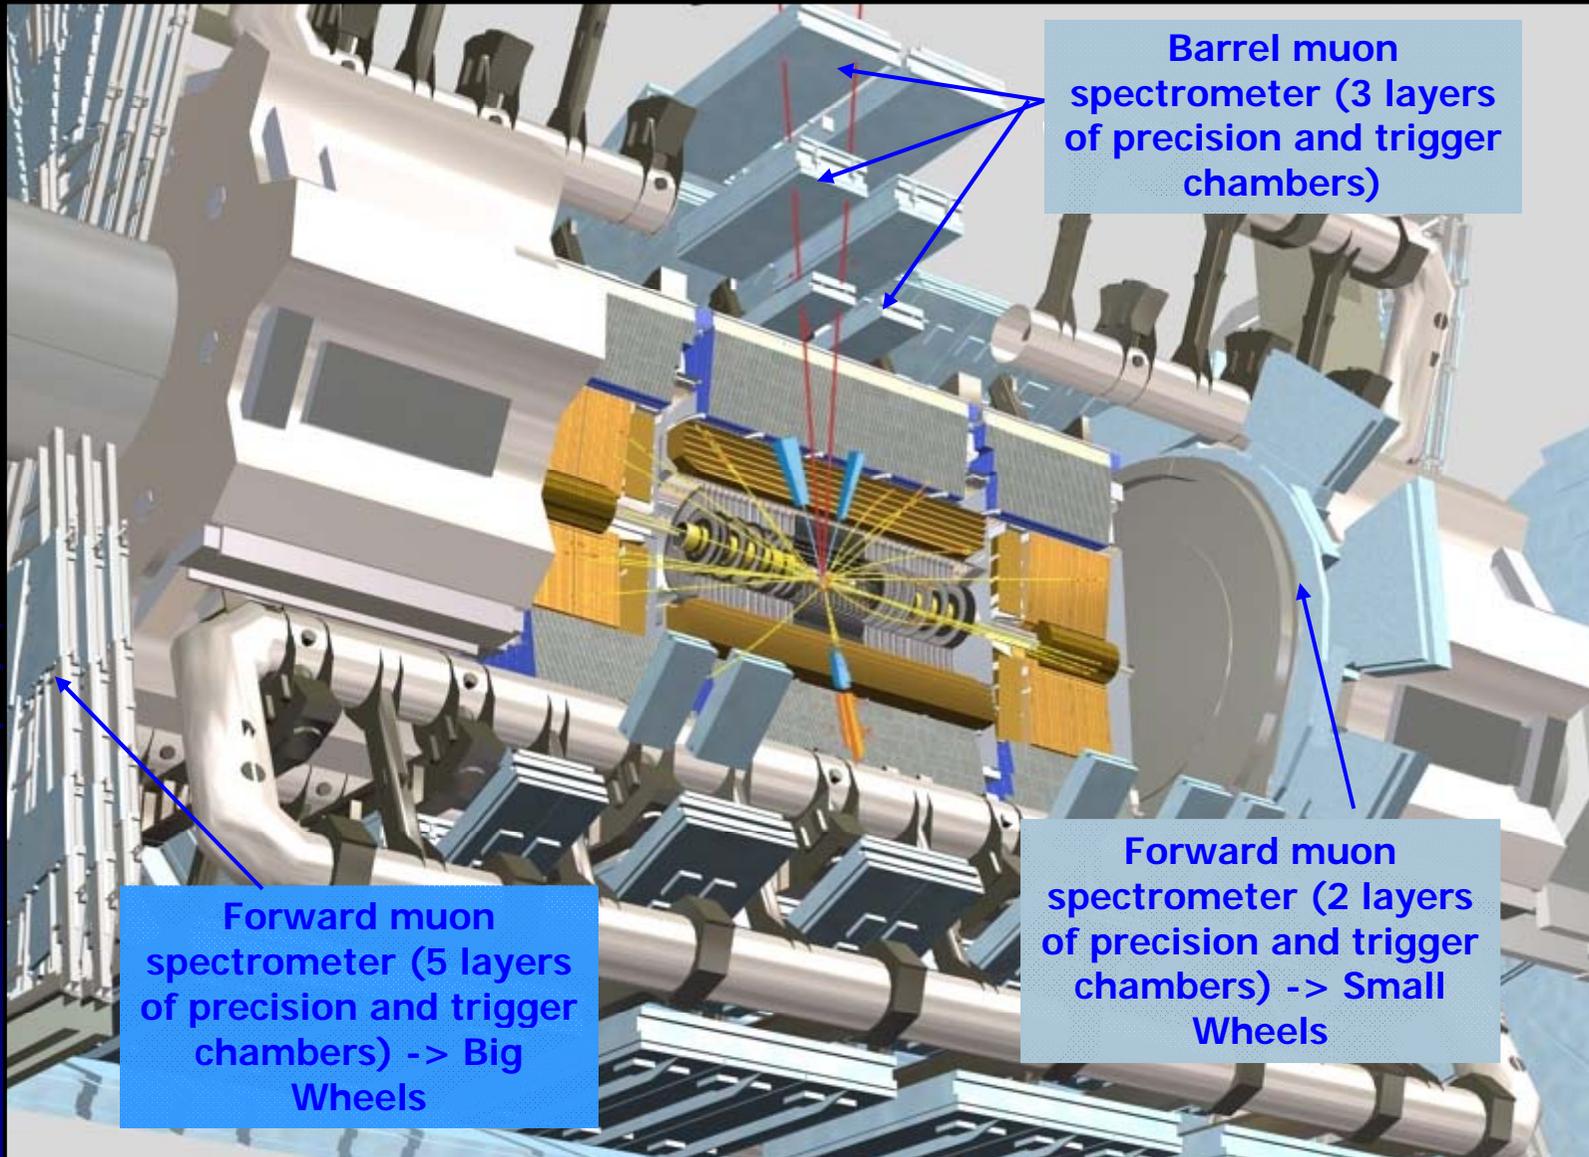
The ATLAS Conditions Database Model for the Muon Spectrometer

*Monica Verducci
INFN Roma1
3rd November 2008
ACAT08 Erice (Italy)*

Outline

- *Introduction @ the ATLAS Muon Spectrometer*
- *Muon Spectrometer Data Flow: Trigger and Streams*
- *Muon Conditions Database Storage*
- *Software Infrastructure*
- *Applications and Production test*

The ATLAS Muon Spectrometer

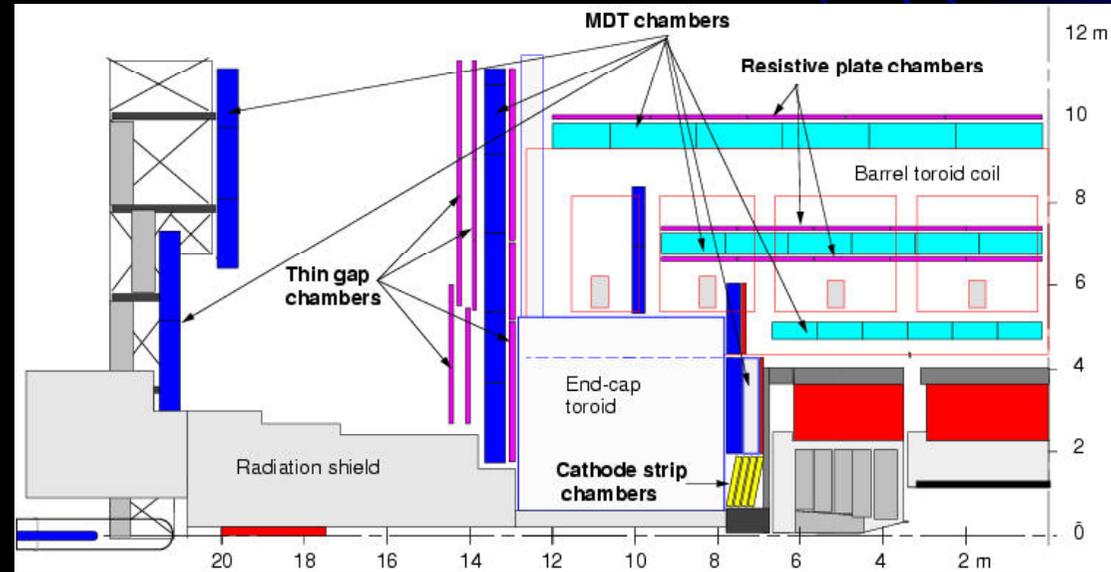
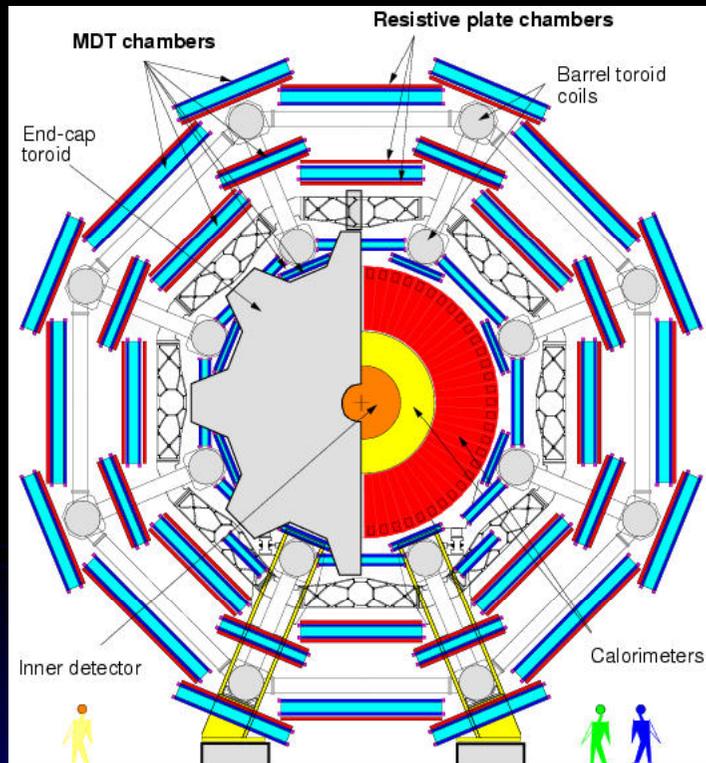


Barrel muon spectrometer (3 layers of precision and trigger chambers)

Forward muon spectrometer (5 layers of precision and trigger chambers) -> Big Wheels

Forward muon spectrometer (2 layers of precision and trigger chambers) -> Small Wheels

The Muon Spectrometer Layout



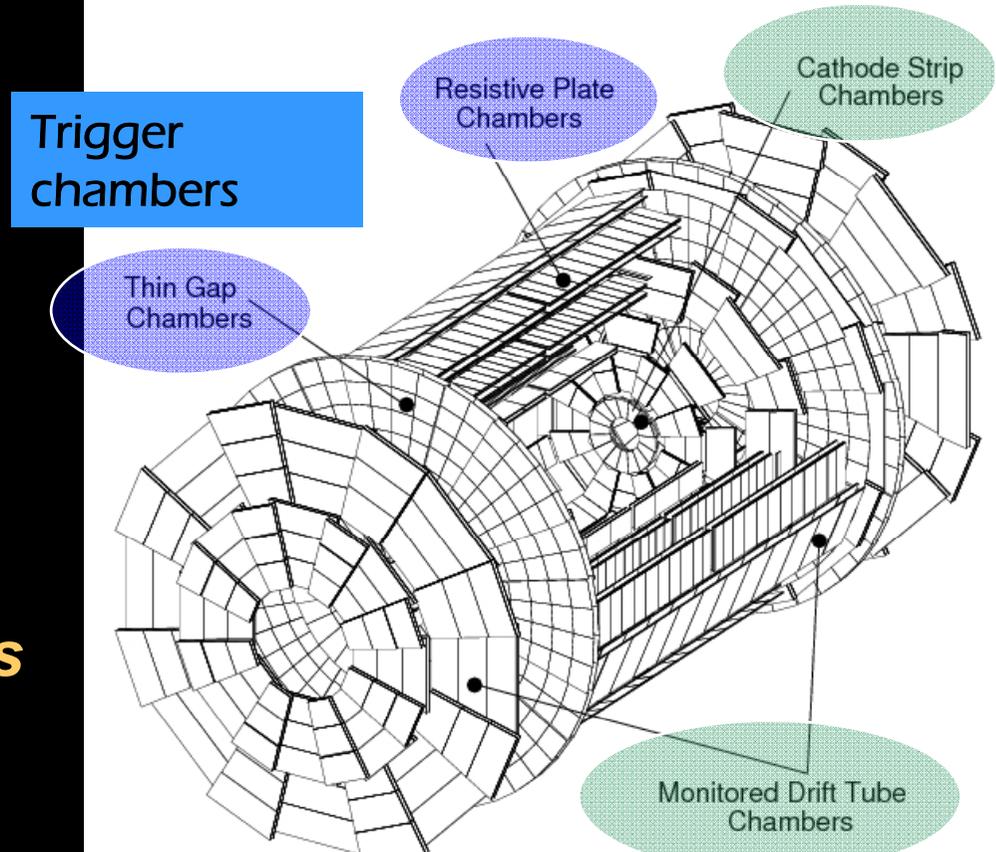
Three toroidal magnets create a magnetic field with:

- **Barrel:** $\int Bdl = 2 - 6 \text{ Tm}$
- **Endcaps:** $\int Bdl = 4 - 8 \text{ Tm}$

- **RPC & TGC:** Trigger the detector and measure the muons in the xy and Rz planes with an accuracy of several mm.
- **CSC:** Measure the muons in Rz with $\sim 80 \mu\text{m}$ accuracy and in xy with several mm. Cover $2 < |\eta| < 2.7$
- **MDT:** Measure the muons in Rz with $\sim 80 \mu\text{m}$ accuracy. Cover $|\eta| < 2$

Muon Spectrometer granularity

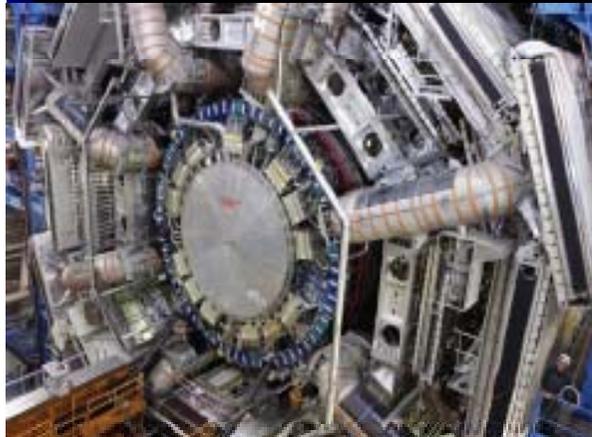
- **Cathode-Strip Chambers (CSC)** : 32 chambers, 31k channels
- **Monitored Drift Tube (MDT)**: 1108 chambers, 339k channels
- **Thin Gap Chambers (TGC)**: 3588 chambers, 359k channels
- **Resistive Plate Chambers (RPC)**: 560 chambers, 359k channels



Need a good resolution in the timing, pT and position measure to achieve the physics goals! Extremely fine checks of all the parts of each subdetector! Huge amount of information...

Precision chambers

ATLAS Event Flow



TRIGGER

Event Selection

Output Streams

Detector Parameters

Configuration DB

subset

Non event data

Conditions DB

ATHENA

Offline reconstruction

Primary
Pathological
Express
Align
Calibration &

10^9 events/s => 1GHz
1 event ~ 1MB (~PB/s)

Hierarchical trigger system
~MB/sec
~PB/year raw data

Muon Calibration Stream

- *Muons in the relevant trigger regions are extracted from the second level trigger (LVL2) at a rate of ~1 KHz*
- *Data are streamlined to 3 Calibration Centres*
 - *Ann Arbor, Munich, Rome (from to Naples for RPCs)*
 - *100 CPUs each*
- *The stream is useful also for alignment with tracks and trigger efficiency studies, as well as for the other Muon detectors (trigger and precision chambers)*
- *~1 day latency for the full chain*
 - *From data extraction, to calibration computation at the Centres, to writing the calibration constants in the Conditions DB at CERN*
- *Need to carefully design the data flow and the DB architecture*

The MUON “Non Event Data”

A typical ATLAS “Non-Event data” could be a:

- *Calibration and Alignment data (from express and calibration streams for a total data rate of about 32MB/s, dominated by the inclusive high pt leptons (13% EF bandwidth= 20Hz of 1.6MB events). RAW Data -> 450 TB/year. More streams are now subsumed into the express stream)*
- *PVSS Oracle Archive, i.e. the archive for the DCS « slow controls » data, and DAQ via OKS DB.*
- *Detector configuration and connectivity data, specific subdetector data*

Mainly used for:

- *Diagnostic by detector experts*
- *Geometry, DCS*
- *Sub-Detector hardware and software*
- *Data defining the configuration of the TDAQ/DCS/subdetector hardware and software to be used for the following run*
- *Calibrations and Alignment*
- *Event Reconstruction and analysis*
- *Conditions data*

Muon Conditions Data

- *Calibration for the precision chambers*
- *Alignment from sensors and from tracks*
- *Efficiency flags for the trigger chambers*
- *Data Quality flags (dead / noisy channels) and final status for the monitoring*
- *Temperature map, B field map*
- *DCS information (HV,LV,gas...)*
- *DAQ run information (chamber initialized)*
- *SubDetector Configuration parameters (cabling map, commissioning flags...)*

Calibration Stream
& offline algo



Analysis algorithms



Hardware Sensor



OKS2COOL and
PVSS2COOL



Constructor parameters



Storage of the 'non-event' data

- *There are different Database storage solution to deal the different hardware and software subdetector work point.*
 1. *Hardware Configuration DB*
 - *Oracle private DB, architecture and maintenance under detector's experts*
 2. *Calibration & Alignment DB*
 - *Oracle private DBs, one for the MDT Calibration (replicated in three centers) and one for the Alignment sensors.*
 3. *Data Quality DB*
 - *Cool INTR server for the development*
 4. *Condition DB*
 - *Contains a subset and less granularity information*
 - *Cool Production DB*

Conditions DataBase

- *The Conditions data are non-event data that could:*
 - *Vary with time*
 - *May exist in different versions*
 - *Data coming from both offline and online*
- *The Conditions DB is mainly accessed by the ATLAS offline reconstruction framework (ATHENA)*
- *Conditions Databases are distributed world-wide (for scalability)*
 - *accessed by an “unlimited” number of computers on the Grid: simulations jobs, reconstruction jobs, analysis jobs,...*
- *Within ATLAS, the master conditions database is at CERN and using Oracle replica mechanism will be available in all Tier-1 centers*
- *The technology used in the Conditions DB is an LCG product: COOL (COnditions Objects for LHC) implemented using CORAL*

Cool Interface for Conditions Database

- *The interface provided by COOL:*
 - *LCG RelationalAccessLayer software which allows database applications to be written independently of the underlying database technology (Oracle, in MySQL or in SQLite).*
- *COOL provides a C++ API, and an underlying database schema to support the data model.*
 - *Once a COOL database has been created and populated, it is possible for users to interact with the database directly, using lower-level database tools*
- *COOL implements an interval of validity database*
 - *Database schema optimized for IOV retrieval & look-up*
 - *objects stored or referenced in COOL have an associated start and end time between which they are valid.*
 - *times are specified either as run/event, or as absolute timestamps in agreement with the meta-data stored.*

- *COOL data are stored in folders (tables)*
 - *Database = set of folders*
 - *Within each folder, several objects of the same type are stored, each with their own interval of validity range*
- *COOL folders can be*
 - **SingleVersion:** *only one object can be valid at any given time value*
 - *DCS data, where the folder simply records the values as they change with time*
 - **MultiVersion:** *several objects can be valid for the same time, distinguished by different tags*
 - *calibration data, where several valid calibration sets may exist for the same range of runs (different processing pass or calibration algorithm)*

<i>Since</i> (Time)	<i>Until</i> (Time)	<i>ChannelId</i> (Integer)	<i>Payload</i> (Data)	<i>Tag</i> (String)

MuonSpectrometer Examples

1. **DCS**: Temperature or HV values depends on the loV and are relative simple and small → Inline Payload
2. **Calibration Data and Alignment**: Parameters with a high granularity, more parts can give the same loV → Reference or CLOB Payload

1.

Since	Until	Ch Id	Payload		Tag
<i>Evt10</i> <i>Run1</i>	<i>Evt20</i> <i>Run10</i>	1	<i>HV</i>	<i>Temp</i>	<i>DCS</i>

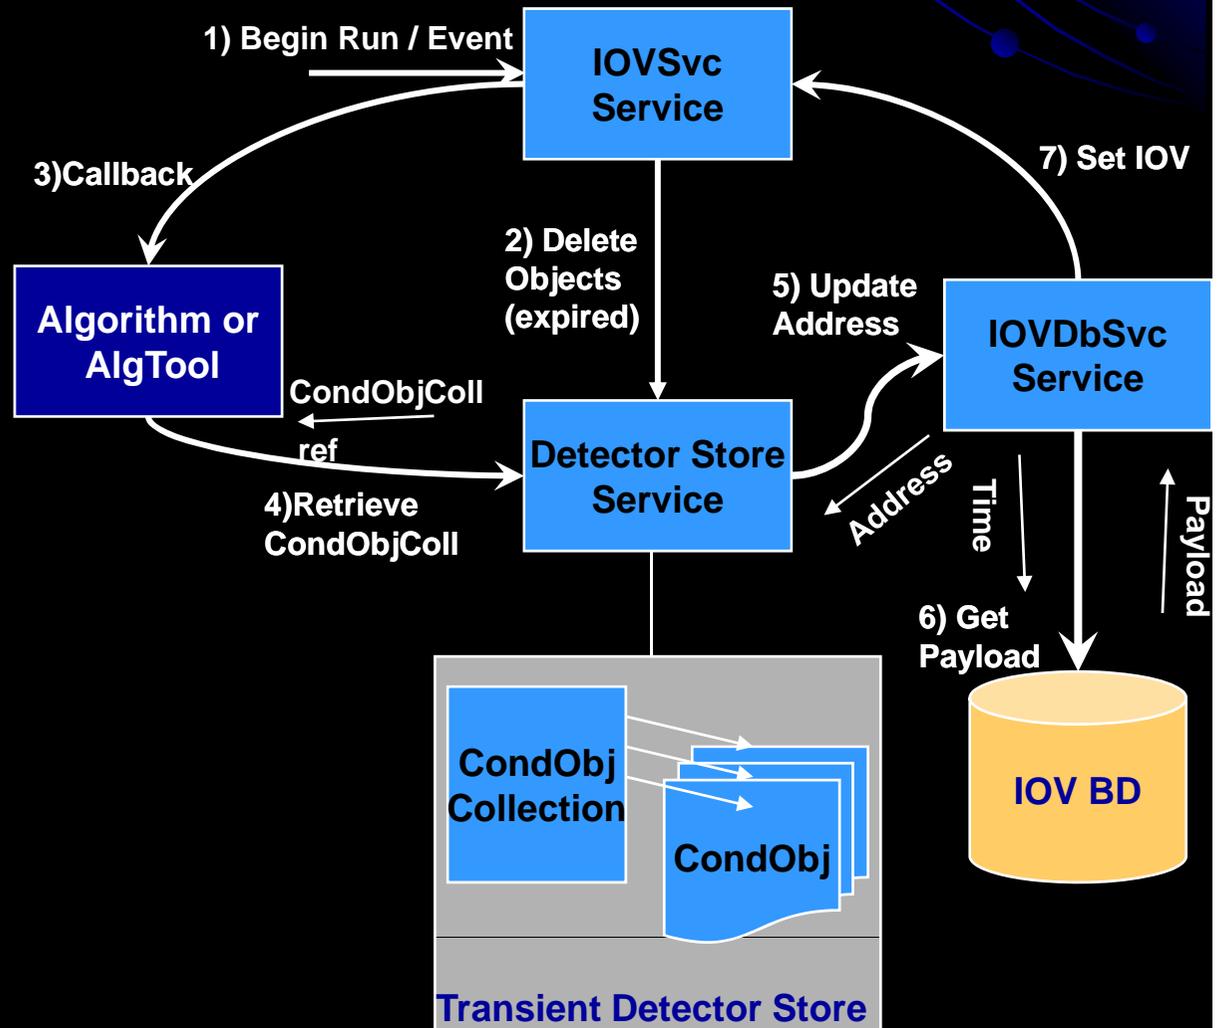
2.

Since	Until	Ch Id	Payload		Tag
<i>Evt10</i> <i>Run1</i>	<i>Evt20</i> <i>Run10</i>	1	<i>T0</i> <i>CLOB</i>		<i>Cosmics</i> <i>M4</i>

Access by the Reconstruction

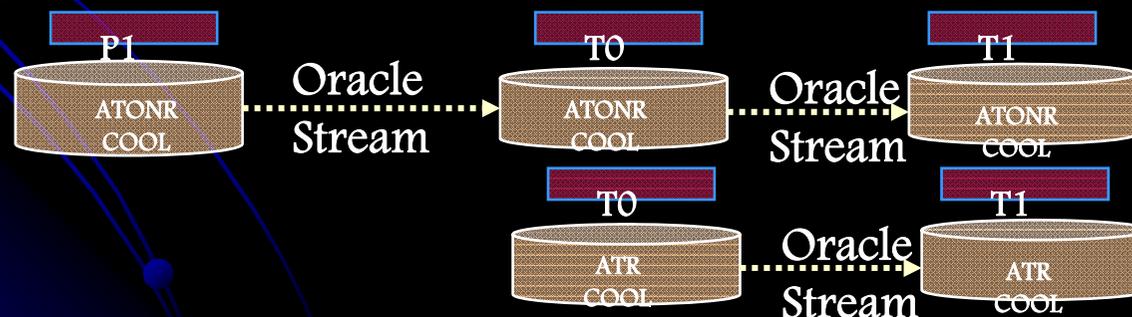
Access to COOL from Athena is done via the Athena **IOVDbSvc** (provides an interface between conditions data objects in the Athena transient detector store (TDS) and the conditions database itself).

Reading event data, the IOVDbSvc ensures that the correct Cond Data Obj are always loaded into the Athena TDS for the event currently being analyzed.



Conditions DB Deployment in ATLAS

- The available servers dedicated to the storing the CondData are: **ATONR, ATR, ATR (old cool version)** the replica schema identifies the different servers.
- Each subdetector has an its own schema object: **ATLAS_COOLONL_XXX** (with reader, writer and owner account in atonr) and **ATLAS_COOLOFL_XXX** (with reader, writer and owner account in atr) .
- The folders structure is defined by the subdetector experts, included the format of data and the table structure and it can be different for each system.
- Different instances exist due to the superposition of the run number (COMP200, OFLP200...)
- Hierarchical tagging and global tagging to improve the granularity and flexibility of the data



Commissioning and Tests

- Tests of all the chain, transfer of data (streaming), access to the data in reconstruction job have been tested
- The cosmics data have been stored successfully (in particular alignment and calibration info)
- The Muon data replica and access have been tested inside the overall ATLAS test with some dummy data:
 - The production schema for ATLAS have been replicated from the online RAC ATONR to ATR and then on to the active Tier-1 sites, more than 1900 runs, 80 GB of COOL data replicated to Tier-1s, ~1.9 GB/day by oracle streams.
 - Tests on the access by ATHENA and on the replica/transfer data between Tier1 and Tier0 have been done, good performance @ Tier1 (~200 jobs in parallel).

Schema	#folders	#chan	Chan payload	N/run	Total GB
MDT	1+1	1174	CLOB: 3kB+4.5kB	0.1	13.0

Conclusions

- *The Access and architecture for most of the Muon Conditions Data have been extensively tested.*
- *The Software Commissioning has provided a good opportunity to tune the structure DB and the reconstruction interface.*
- *The ATLAS overall Conditions Data stress test did not outline any particular problem in the muon data access.*

Backup

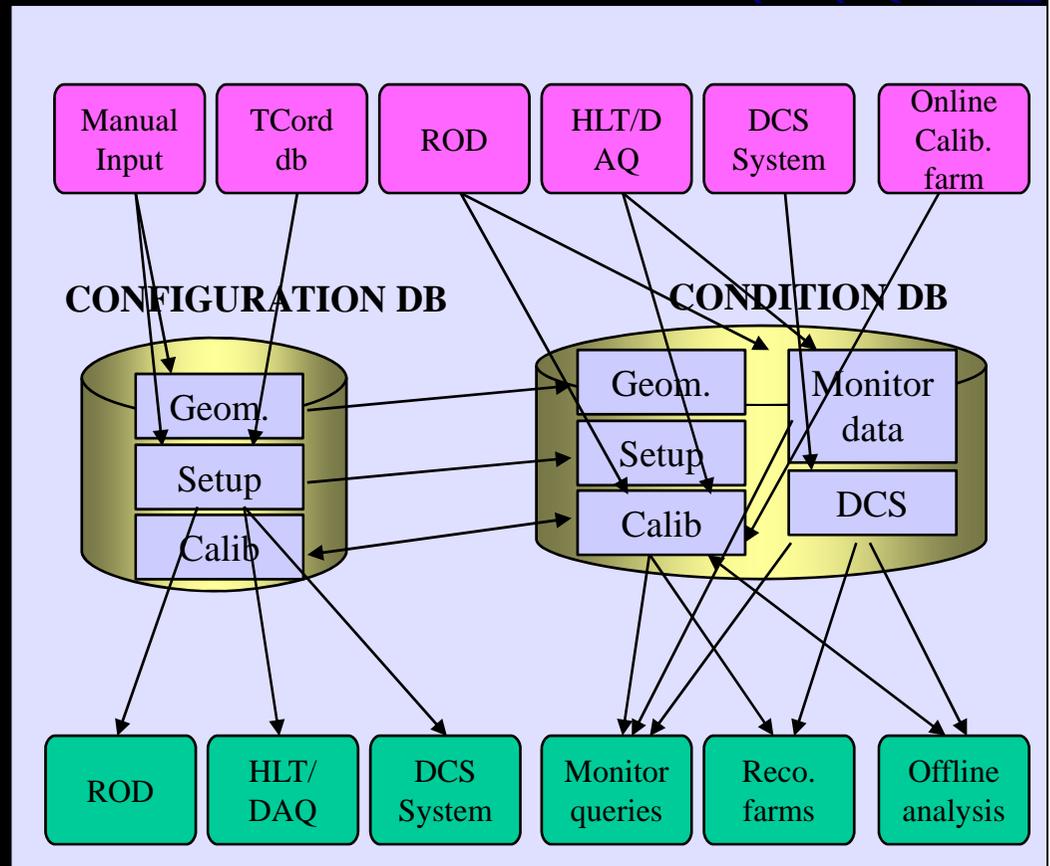
Conditions Data Schemas

- *Different CondDB servers exist, providing different use and parallelizing the accesses*
 - *Dedicated to the online ATONR is replicated outside point1, at Tier0 and Tier1*
 - *ATLR is dedicated to the offline reconstruction and access (not available at point1)*
- *The Schema objects are defined by the subdetector names ATLAS_COOLONL_XXX (atonr) and ATLAS_COOLOFL_XXX (atlr), three different users are defined for the W,R,O privileges.*
- *Different instances exist due to the superposition of the run number (COMP200, OFLP200...)*
- *Hierarchical tagging and global tagging to improve the granularity and flexibility of the data*

Data Storage Capability

The “non-event data” are stored in:

- **Configuration DB**
 - Data needed at the start of the run to configure (i.e. TDAQ, DCS & subdetector hardware)
- **Conditions DB**
 - Data needed to describe the “condition” of the event (Event Reconstruction and analysis, Diagnostic, Calibration and Alignment)



ATLAS Database Storage Requirements: ~16 TB in 2009 at CERN (plus 1TB for catalog, bookkeeping system, monitoring,...)

Some numbers: CondDB

- *ATLAS daily reconstruction and/or analysis job rates will be in the range from 100k to 1M jobs/day*
- *For each of ten Tier-1 centers that corresponds to the Conditions DB access rates of 400- 4000 jobs/hour*
- *Each reconstruction job will read 10-100MB of data*
- *Atlas requests to Tier-1s is a 3-node RAC cluster dedicated to the experiment.*
- *Expected rate of data flow to Tier-1s is between 1-2 GB/day*

ATLAS Workload

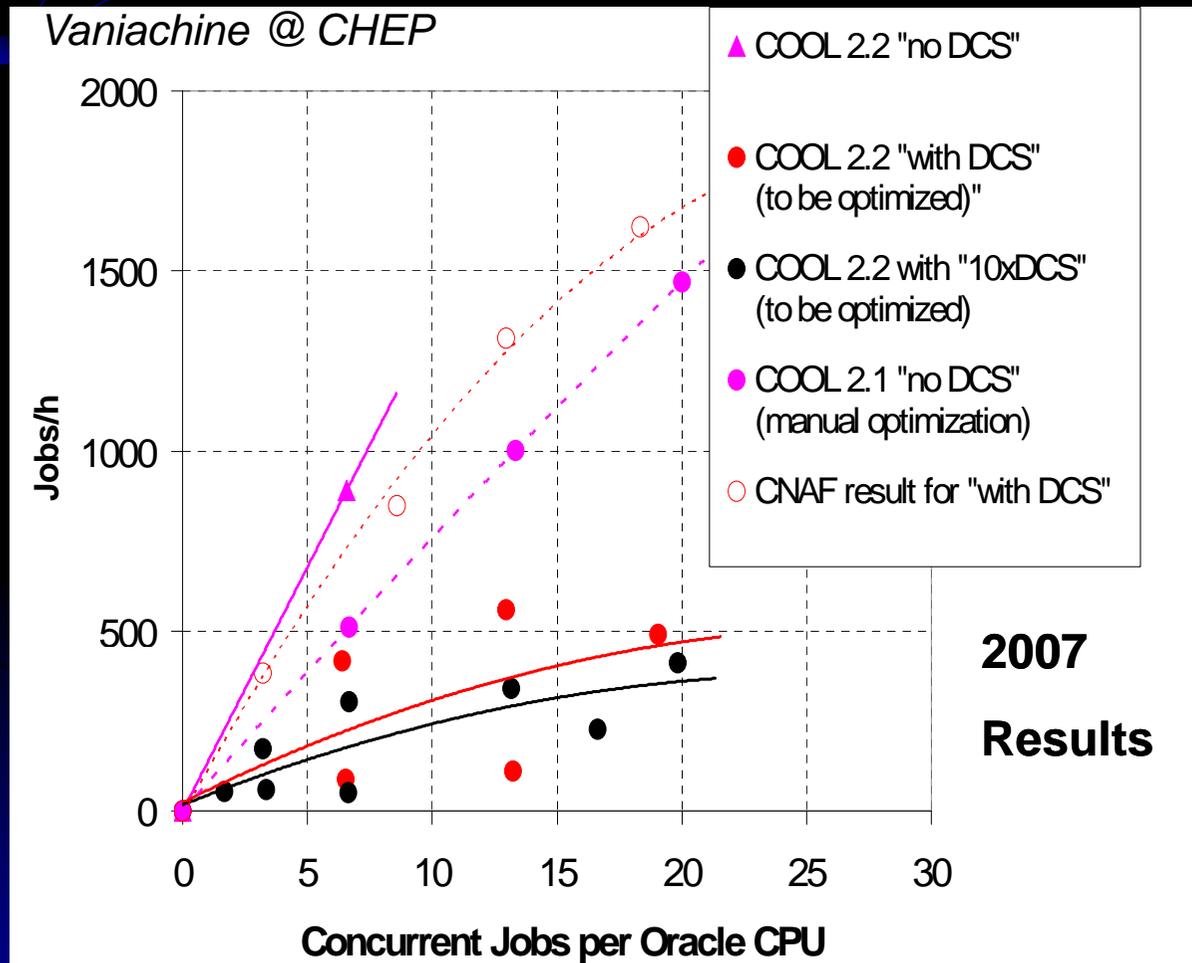
- Regular conditions data updates on Online RAC, testing propagation to Offline RAC and further to ten Tier-1s
 - Test workload using multiple COOL schemas, mixed amount/type of data:
 - This data is split over 5 database schemas: ATLAS_COOLONL_ID, ATLAS_COOLONL_CALO, ATLAS_COOLONL_MUON, ATLAS_COOLONL_GLOBAL and ATLAS_COOLONL_TDAQ, in each case using the database instance name **OFLP200**.

Schema	#folders	#chan	Chan payload	N/run	Total GB
INDET	2	32	160 char	1	0.16
CALO	17	32	160 char	1	1.3
MDT	1+1	1174	CLOB: 3kB+4.5kB	0.1	13.0
GLOBAL	1	50	3 x float	6	0.19
TDAQ/DCS	10+5	200+1000	25 x float	12	59.6
TRIGGER	1	1000	25 x float	12	5.4

- Data in some sense a 'best guess' ATLAS conditions DB load (dominated by DCS)

Presented by Hawkings and Vaniachine @ CHEP

Workload scalability results



In ATLAS we expect 400 to 4,000 jobs/hour for each Tier1

For 1/10th of the Tier1 capacities that corresponds to the rates of 200 to 2,000 jobs/hour

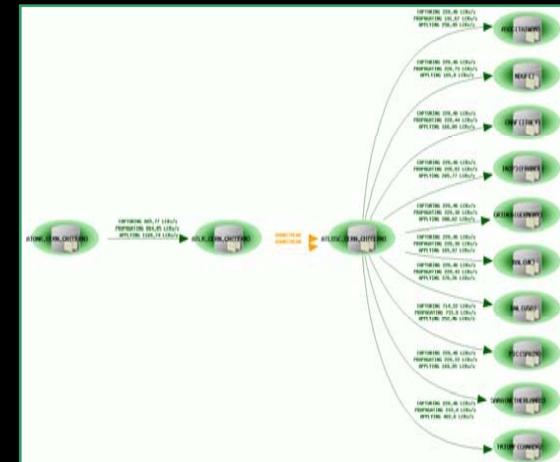
Good Results!

For the Access by Athena we have obtained 1000 seconds per job events at ATLR due to the DCS and TDAQ schema access!

Oracle stream Replication Test

- *Production phase started on April, 1st 2007, with 6 destinations*
- *We have now the **10** destinations actively receiving data*
- *Since April, 13th more than **60GB** of COOL test data have been replicated*
- *A cron job runs twice per hour adding one run's worth of data, roughly 20 MB per run, which amounts to 1GB/day volume. Tests have been successful in increasing the volume to **2GB/day** over several days. Problems arise mainly with memory and CPU issues on the replicating machine.*
- *When a Tier-1 has a failure, procedures are in place to isolate the site and make it « catch up » with the others.*
- *These procedures were used several times during these tests and where successful*
- *There was a formal recovery exercise on June, 13th, on the **3D DBA Workshop**, which involved most of the Tier-1s DBAs.*

Presented by F.Viegas @ CHEP07



ATLAS Data Volume @LHC

Bunch-crossing frequency: 40 MHz
~ 20 collisions p-p per bunch crossing

10^9 events/s => 1GHz
1 event ~ 1MB (~PB/s)

Hierarchical trigger system
~MB/sec
~PB/year raw data

Reduction of the event via 3 Levels of Trigger
into 4 different output streams (200Hz, 320 MB/s):

- **Primary stream** (5 streams based on trigger info: e,m,jet)
- **Calibration and Alignment Stream** (10%)
- **Express Line Stream** (Rapid processing of events also included in the Primary Stream 30 MB/s, 10%)
- **Pathological events** (events not accepted by EF)

