Requirements of the MDT Calibration and Alignment Centre

Oliver Kortner

Max-Planck-Institut für Physik

ATLAS-Besprechung, 09.10.2006

Outline

1. Calibration and aligment tasks in the muon spectrometer.
2. Data-stream dedicated to calibration and alignment.
3. Hardware and software requirements for the calibration centre.
4. Time scale.
5. Operation plan.

Calibration and Alignment Tasks

The regular calibration tasks

- Weekly synchronization of all drift-tube channels.
- Daily determination of the r-t relationship of each chamber.
- Daily determination of the spatial resolution of each chamber.

The regular alignment tasks

- Initial alignment with straight muon tracks – 1 run with no toroid field planned for the beginning of each year.
- Alignment of chambers without optical alignment sensors by means of tracks in overlaps with chambers equipped with optical sensors.
- Monitoring of the alignment with low-p_T muon tracks.

Calibration and Alignment Data-Stream

Data-stream dedicated to calibration and alignment

- Needed in order to collect sufficient statistics for calibration and alignment.
- Data in a level-2 trigger tower can be put out at a rate of 2.5 kHz.
- Composition of the data stream:
  - 2.25 kHz for calibration ($\sim \frac{1}{3}$ to each calibration centre).
  - 250 Hz for alignment with overlap tracks (100% to Munich, 0% to Rome and Michigan).
- Size of muon-tower event: $\sim$1 kB.
- Data volume for 15 h of data taking (1 run): 135 GB.
- Data will be transmitted to the calibration centres in blocks of 1-5 GB.

Data Transfer Scenarios

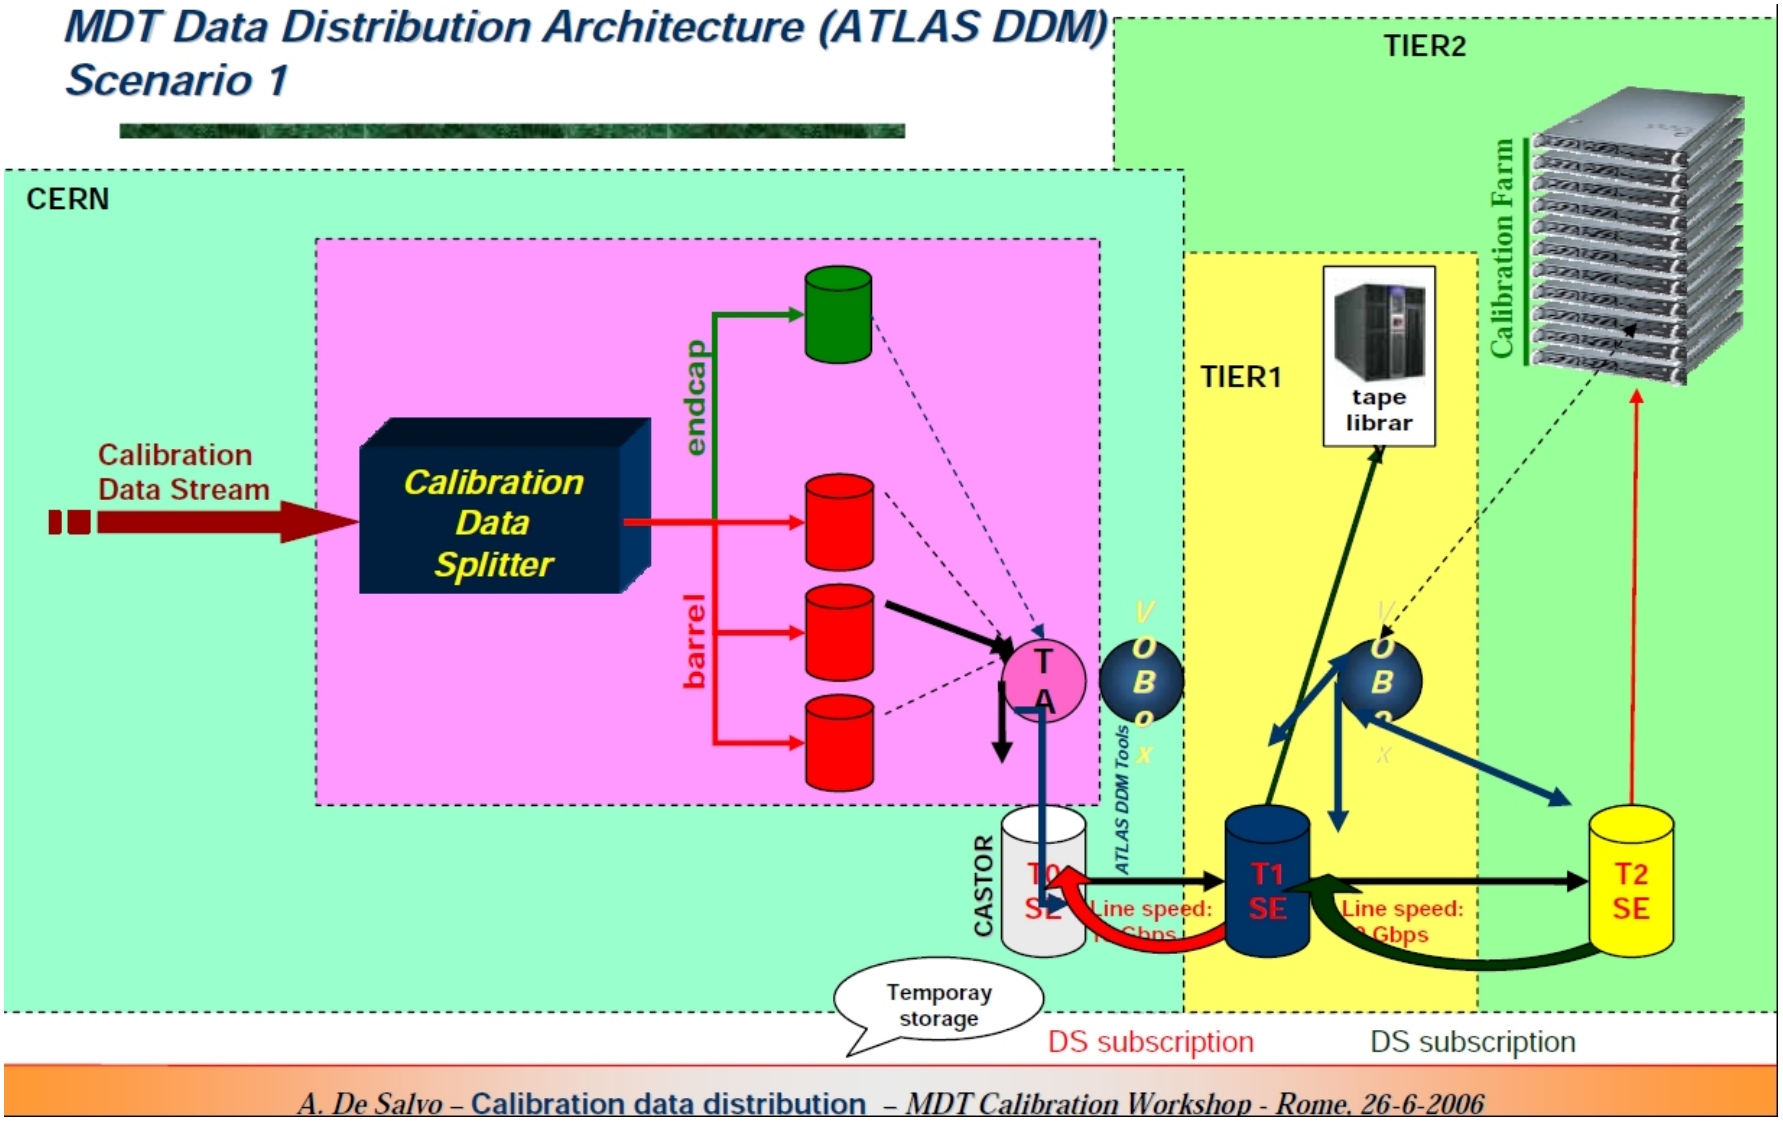

Data Transfer Scenarios

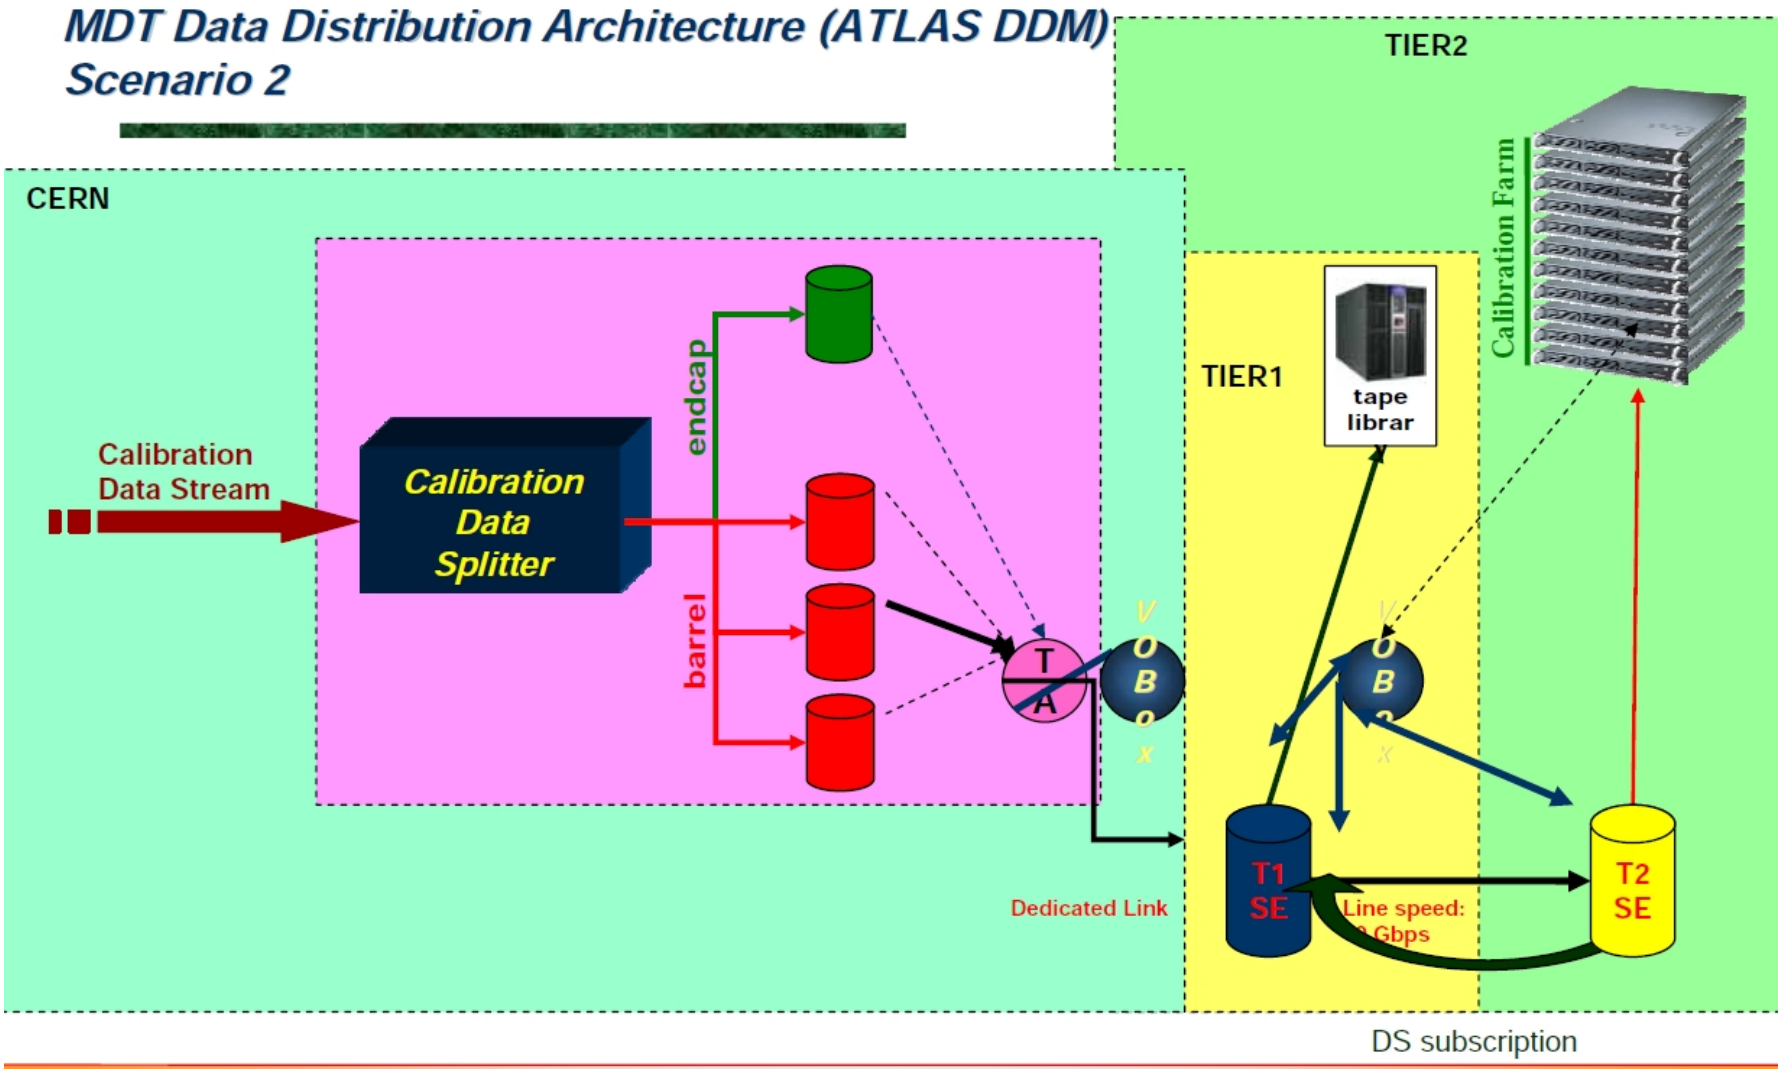

Calibration Data Flow

- The data produced at CERN is registered in the Grid Storage Elements at Tier0 (this step could be skipped in case of a dedicated link to Tier1s)
  - The data is splitted in files by the Calibration Data Splitter
  - The files are transferred to the Storage Elements as soon as they are created
  - A calibration dataset is created by the Transfer Agent for each calibration session
  - The files are registered as part of a calibration dataset in the ATLAS DDM (DQ2)
    - New files belonging to the same calibration datasets are added to the list of the dataset components as soon as they are stored and registered in the Ses
- The Transfer Agent subscribes the calibration datasets to both Tier1 and Tier2 Storage Elements
  - The subscription is performed via the DDM/DQ2 facilities
    - Interacting with the VOBOXes at Tier0 and Tier1
- The subscribed dataset are automatically transferred by the DDM
  - Tier0 → Tier1 → Tier2
  - This operation is performed in quasi-streaming mode via the FTS (File Transfer Service) channel
    - Transfer priority managed by the DDM
    - QoS assured by the centers/resource owners
  - The Tier1 will be responsible of backing up the data on tape
    - This is done for free if using CASTOR or similar systems as a target Storage Element
- The Calibration Farms in the Tier2 will be able to process the files as they will available in the local Storage Elements
  - Files belonging to an unfinished dataset may be used as soon as they are transferred
  - At the end of the calibration data processing the dataset subscription may be released from Tier0 and Tier2

Calibration Data Flow

- **Data discovery**
  - The remote calibration farms should be able to analyze the new calibration data as soon as it is ready
  - The components of a calibration dataset can be discovered via DDM
  - The dataset name should be supplied by the calibrators or there should be a clear automatic procedure to go to the next calibration dataset after the previous run will be finished
    - For example using a standard naming convention for the dataset name for each calibration
- **Data access**
  - The calibration data could be accessed directly from the local storage in the Tier2s
    - For example using GFAL (RFIO) directly from ROOT
  - The calibration farm nodes should be a separate partition with respect to the rest of the Tier2 infrastructure
    - The calibration Worker Nodes should be excluded from the Grid access but should be instrumented with all the Grid facilities to access the data
    - Dynamical farm partitioning should be possible
  - Data should be either read directly from the calibration application or by a data preparation agent, to isolate the relevant parameters for the fitting program
- **Data cleanup**
  - The calibrators should tag a calibration as done
  - After a calibration is over, the data could be in principle released (unsubscribed), thus deleted from the Tier2 storage
  - In case there is a need of reprocessing, the data could be re-subscribed from the Tier1
  - There should also be the possibility to keep the data for a longer period in the Tier2 storage
    - A "garbage collection" process should also be provided, to put a limit on the maximum amount of local Tier2 storage that can be used for calibration data in status "done" before releasing it

Calibration Data Flow

- In case one of the calibration Tier2 is temporary unable to process the data, another center should take over the responsibility of its data processing
  - The calibration data splitter could be reconfigured and restarted to redirect the data in the proper place, until the problems are fixed
    - Just changing a parameter in the configuration files
  - Should we send at the beginning the full set of the data to all the centers?
    - In this case no reconfiguration of the data splitter is needed, in case of problems
- Data throughput
  - Calibration data CERN → Tier2s: ≈100 GB/day
  - Calibration constants Tier2s → CERN (database): 50-100 MB/day

Hardware and Software Requirements

CPU requirements

- CPU power must be sufficient to process the events at least as fast as the data rate (≈ 1 kHz).
- Time to process 1 event: ≈ 0.1 ms.

$\rightarrow$ 100 CPUs needed to process 1000 events/s.

- MPI and LMU will provide 100 CPUs each in order to allow reprocessing and to avoid night shifts for calibrators.

Disk-space requirements

- The calibration centres should be able to hold the data from the stream for one year.

$\rightarrow$ 1 TB fileserver for incoming data + 5 TB fileserver for data storage needed.

- 5 TB fileserver already provided by the MDT group for this purpose.

Software Requirements

Software requirements

- LCG+DQ2 as required for Tier-2 centres.
- 1 storage element for the calibration centre.

\+ Oracle database server for the calibration database.
Only additional requirement to a Tier-2 centre.

Time Scale

- A basic version of the calibration centre must be available until end of this year to be ready for the CDC test beginning of next year.
- Basic version:
  - 30 CPUs.
  - Final file server and storage element.
  - Oracle database server.
- 80% of the CPUs must be available summer 2007 to be ready for the cosmic calibration and the first pp collisions.
- 100% of the CPUs must be available beginning of 2008 when LHC turns on collisions at 14 TeV.

Operation Plan

- We expect that Tier-2 specific tasks are taken care of by RZG or Tier-2 personal, i.e.
  - administration of the computing cluster,
  - installation of Grid tools,
  - installation of Athena,
  - administration of the Oracle database server.
- Our jobs:
  - Installation of calibration and alignment software.
  - Operation of the calibration and alignment centre (job submission, checks of calibration and alignment results, checks of correct filling of the calibration database).
  - Baseline plan: 2 shifts/day (7-15, 15-23) with 2 operators each (1 at MPI, 1 at LMU).