Considerations for a CBM QoS Architecture

1. Basic considerations

These considerations build on the estimate of the CBM computing requirements for operation at SIS-100 (CBM-CN-18001), from where it becomes obvious that the costs of storage media for CBM are significant and probably a limiting factor for the experiment operation. The basic guideline for the following considerations is thus that a persistent data level and a corresponding storage and distribution strategy shall not be defined if

a. the production of these data is deterministic, i.e. is defined by another data level plus the production procedure, and
b. the production of the data consumes limited computing resources.

This means that whenever reasonably possible, preference is set for re-production of derived data over their storage. As an example, full reconstructed data (ESD-type) are not stored, since reconstruction from raw data is expected to be fast.

2. Persistent data levels

The persistent data levels for CBM offline computing are:

1. **RAW**: Selected de-contextualised (“unpacked”) data extracted from the data stream as delivered by the CBM data acquisition. Typical objects are digis (digital single-channel information), typical containers are raw events.
2. **SIM**: Full MC information of simulated data. CBM intends to provide simulation statistics equivalent to the triggered experiment data statistics. Of about 10% of these data, the full MC information will be kept.
3. **AOD**: Analysis Object Data serving as input for high-level user analysis. AOD are derived from RAW through calibration, reconstruction and skimming, or from SIM through detector response simulation, calibration, reconstruction and skimming. Different AOD types may be defined serving different physics analysis objectives. Typical objects are tracks, typical containers are reco events.
4. **PAR**: Parameter data required for the production of the AODs (calibration, reconstruction). These are needed for high-level physics analyses. Typical parameter sets comprise the experiment configuration (geometry, settings) defined at the start-up of the experimental run, the running conditions monitored and recorded during the experiment operation, and calibration parameters obtained through an analysis of RAW data. Parameters are typically managed through appropriate data bases.
5. **PHY**: Physics-level results, usually in a binned and inclusive format. These are derived from AOD and constitute the experiment results to be made public.

3. Data operations

The anticipated workflow can be outlined by the following steps:

1. Experiment data taking: RAW data are delivered to on-site storage media (RAW_HOT). CBM aims at taking data for three months per year. The estimated raw data volume is 18 PB/a.

2. RAW data are copied to long-term archive (RAW_COLD). Two copies of RAW have to be archived in two distinct data centres, one being FAIR/GSI. The archival process can already be started during experiment operation.
3. Offline-level calibration parameters are derived by analysis of a subset of RAW. This happens at the host lab, where RAW_HOT is available.
4. AOD are created from RAW and PAR.
5. AOD are transferred to the participating data centres, serving a regional CBM user community. The type of AOD may depend on the principal analysis aims of the respective community.
6. AOD are subjected to high-level physics analysis.
7. SIM data are generated in correspondence with the experiment settings, conditions and statistics. AOD production from SIM data is performed on-the-fly. 10% of the SIM data are stored even after AOD production. Simulation can be performed on any of the participating data centres.
8. AOD produced from SIM are subjected to physics analysis, complementing the analysis of the corresponding experiment data.

4. Quality-of-Service policy

The following service levels can be identified on the base of the workflow outlined above. They do not include PAR and PHY data because of their negligible size compared to experiment and simulated data.

Name	Usage	Volume	Characteristics	Example media
RAW_COLD	Long-term storage of prime experiment data.	18 PB/a	High reliability and long-term stability	Tape
RAW_HOT	Availability for calibration and production of AOD for 2 years after data taking	36 PB	Low latency	Disc
AOD	Availability for user-level physics analysis up to 5 years after data taking	18 PB	Low latency	Disc
SIM	Availability for MC-level analysis up to 3 years after production	16 PB	Low latency	Disc
AOD_SIM	Availability for user-level physics analysis up to 5 years after data taking	18 PB	Low latency	Disc