



# LHCb plans for 2008-09 shutdown

Philippe Charpentier (CERN, LHCb)





## Core Software and applications

- Prepare Gaudi for 2009 data taking
  - Intensive cleaning of Gaudi internals (with ATLAS)
    - ☆ Workshop last week, ready in January
  - Move to latest version of ROOT (new schema evolution)
  - Work on reducing size of ESD (rDST and DST)
    - ☆ First attempts: reduction by 30%
- Applications
  - Certification of simulation with latest generators and Geant4
    - ☆ Ongoing (Geant4 9.1 patch 2)
    - ☆ Small simulation productions for validation
  - Alignment and calibration framework
    - ☆ Used to some extent with the few LHC-beam data (TED runs)
      - \* Reminder: cosmics are interesting but only for detector commissioning, not alignment (vertical detectors)
    - ☆ Simulation ongoing with displaced detectors (2008 survey)

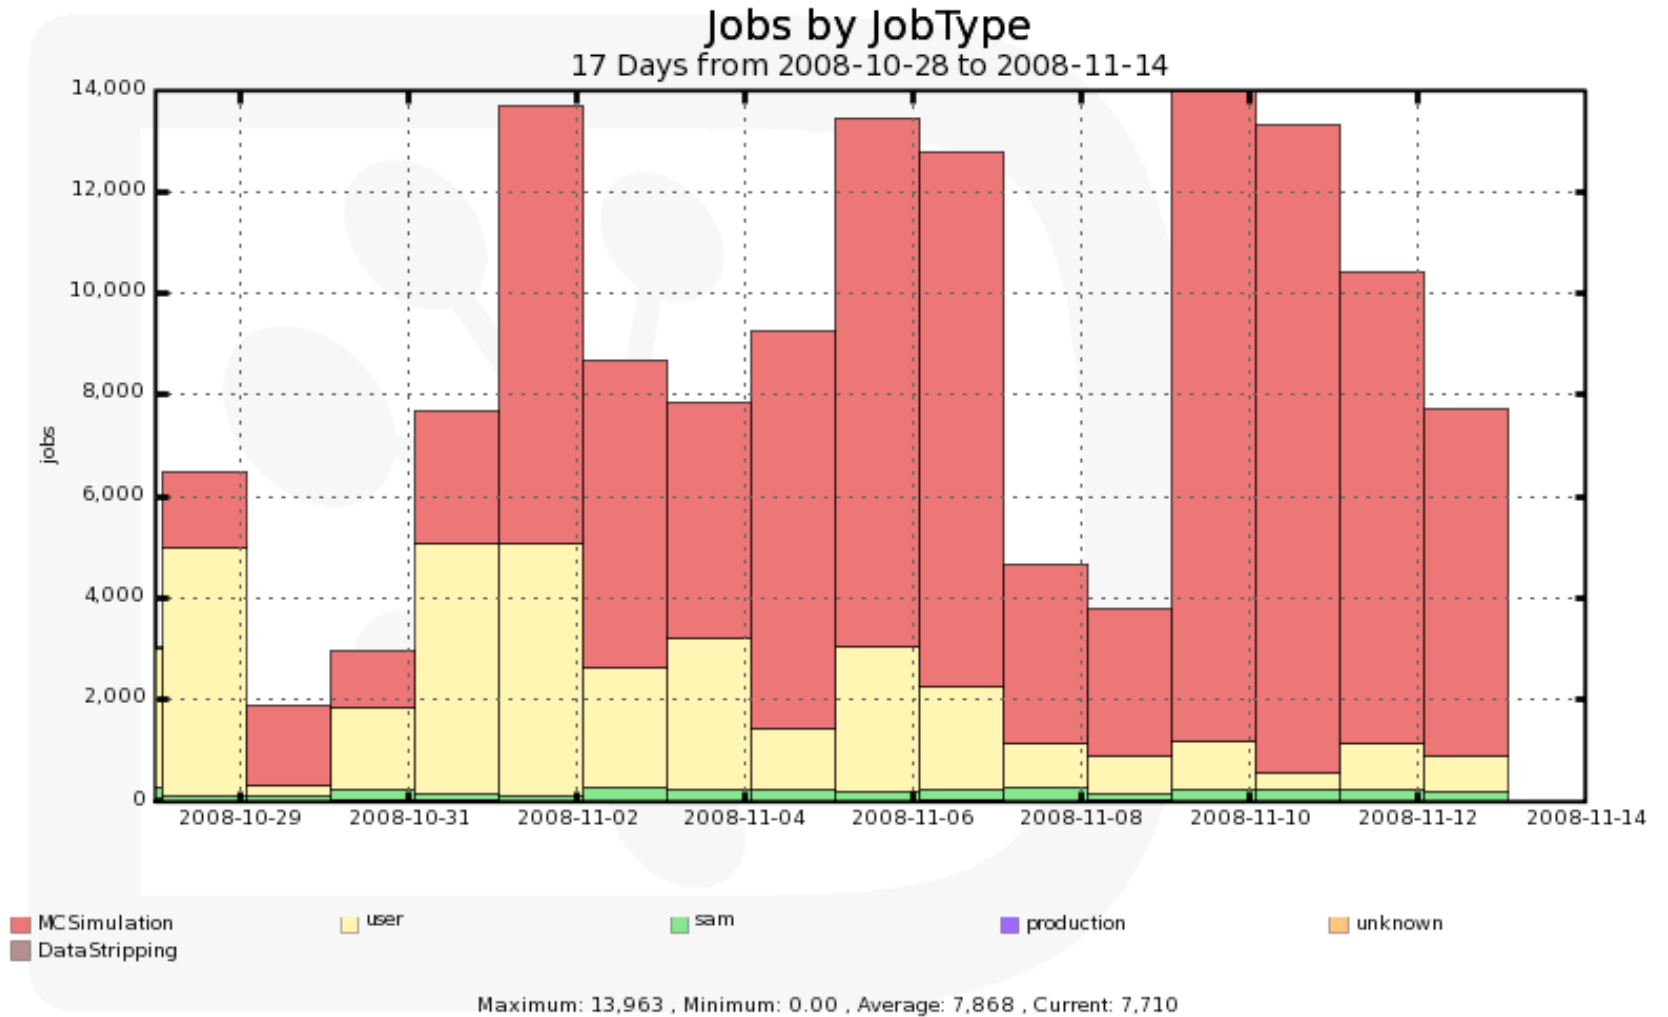


- **Reminder: this is a complete rewrite of DIRAC**
  - **First released version was in Spring**
- **Used for production activities since July**
  - **Simulation workflows (complex DAG with no input)**
  - **Reconstruction workflows (single application with input)**
  - **Stripping workflows (complex DAG with input)**
    - ☆ **First for continuation of "DC06" productions**
      - \* **Needed for physics studies and comparisons**
    - ☆ **DC06 is now ended (still one small remaining production to go)**
    - ☆ **Concentrate on support for most recent versions of applications**
- **Just started analysis on DIRAC3**
  - **Baseline is still DIRAC2 (pilot submission through LCG-RB)**
    - ☆ **Decommissioned end of 2008**
  - **Ganga now has a DIRAC3 backend**
  - **Users started to migrate to DIRAC3**
    - ☆ **Still using private pilots**
    - ☆ **Move to generic pilots as soon as role=pilot is deployed on Tier1s**



# DIRAC3 jobs

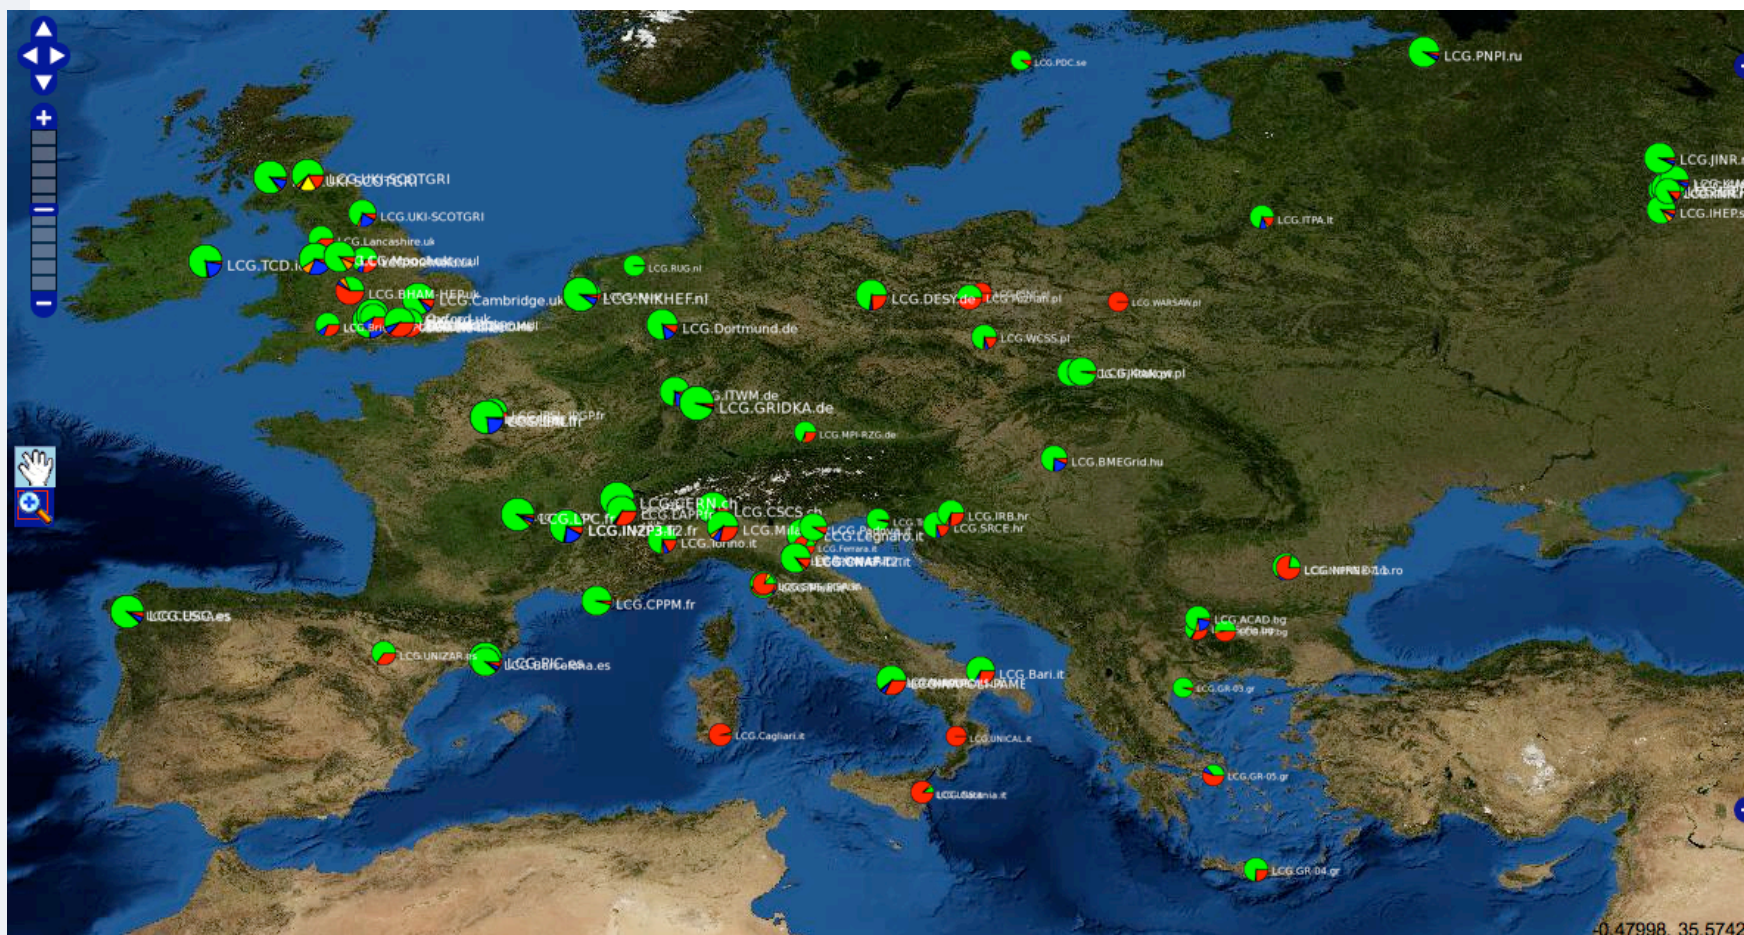
LI-HCB SHUTDOWN ACTIVITIES





# DIRAC3 portal

## LI-HCB SHUTDOWN ACTIVITIES





- Production manager
  - Schedules production work, sets up and checks workflows, reports to LHCb operations
- Computing shifters
  - Computing Operations shifter (pool of ~12 shifters)
    - ☆ Covers 14h/day, 7 days / week
    - ☆ Computing Control room (2-R-014)
  - Data Quality shifter
    - ☆ Covers 8h/day, 7 days / week
  - Both are in the LHCb Computing Control room (2-R-014)
- Daily DQ and Operations meetings
  - Week days
- Grid Expert on-call
  - On duty for a week
  - Runs the operations meetings
- Grid Team (~6 FTEs)
  - Shared responsibilities (WMS, DMS, SAM, Bookkeeping...)



- **Simulation for physics studies**
  - **More realistic detector**
    - ☆ Using 2008 geometry, tuned detector response
  - **Benchmark physics channels**
    - ☆ Statistics being defined by physics groups
    - ☆ ~100-200 Mevts (i.e. 200-400 kjobs)
- **Analysis**
  - DC06 MC-data analysis moved to DIRAC3
  - 2008 simulation (DC09?) MC-data analysis
  - Physics reach studies (toy MC, long fitting processes)
    - ☆ Running at all sites (no input data)
- **Full Experiment System Test (FEST'09)**
  - See next slides

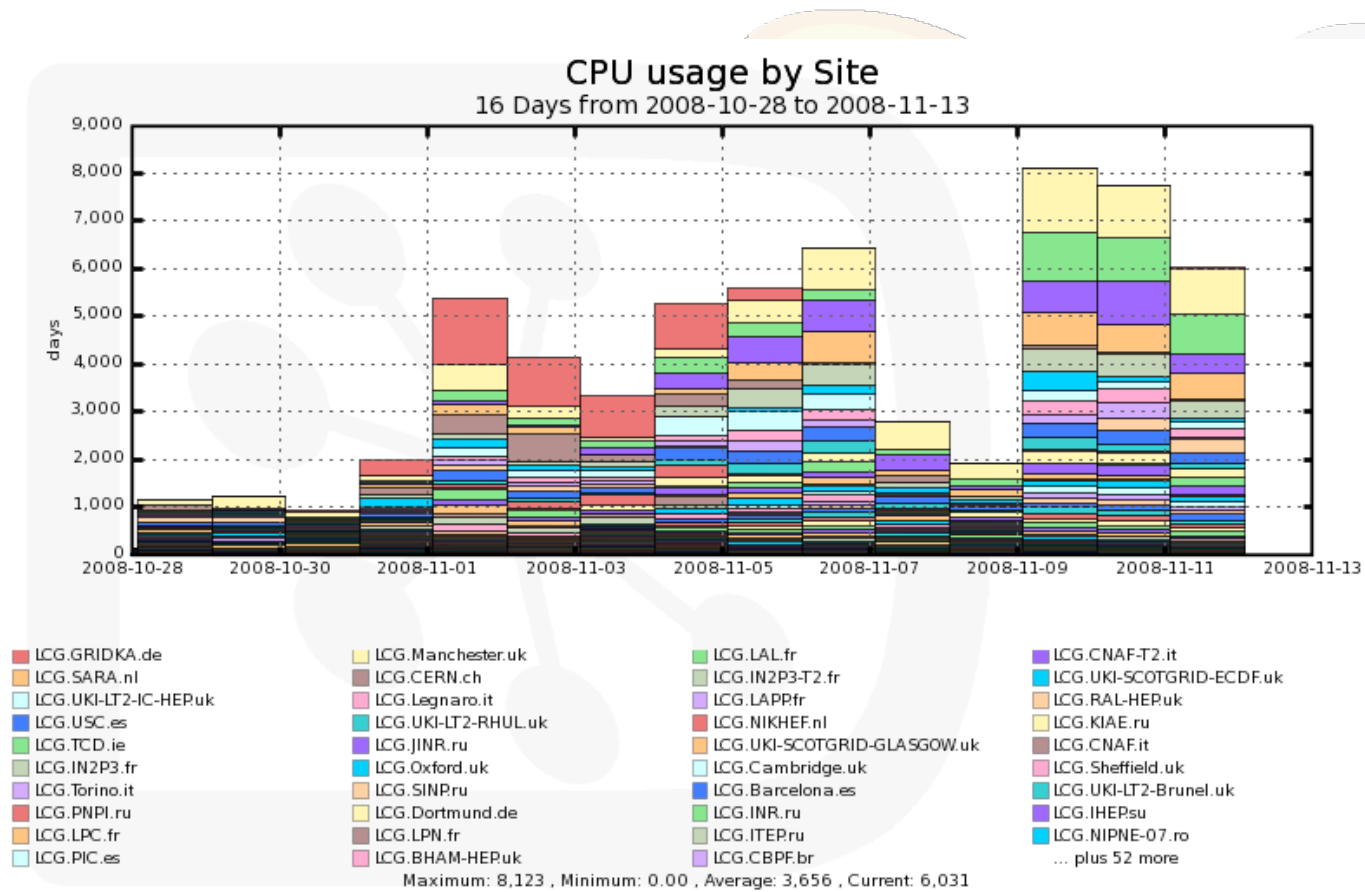


- **Aim**
  - **Replace the non-existing 2008 beam data with MC**
  - **Points to be tested**
    - ☆ **LO (Hardware trigger) strategy**
      - \* Emulated in software
    - ☆ **HLT strategy**
      - \* First data (loose trigger)
      - \* High lumi data (b-physics trigger)
    - ☆ **Online detector monitoring**
      - \* Based on event selection from HLT e.g. J/Psi events
      - \* Automatic detector problems detection
    - ☆ **Data streaming**
      - \* Physics stream (all triggers) and calibration stream (subset of triggers, typically 5 Hz)
    - ☆ **Alignment and calibration loop**
      - \* Trigger re-alignment
      - \* Run alignment processes
      - \* Validate new alignment (based on calibration stream)





- Simulate 100 Mevts (minimum bias)
  - Only store raw data format (streamed files)
  - Completed in 6 days (central part of the plot below)






- File merging (into 2-3 GB files)
  - Small files at CERN on a TOD1 space
  - Merging being done at CERN
- Online developments
  - Event injector
    - ☆ Read MC files
    - ☆ Emulate LO trigger (conditions may be varied)
    - ☆ Creates multi-event packets (MEP as front-end does)
    - ☆ Send MEP to an HLT farm node
  - Event injector control system
    - ☆ Emulation of the standard Run Control
    - ☆ Simulates a regular run, but using event injector as source
  - Multiple streams
    - ☆ Using HLT classification as criterion
      - \* Was not needed for 2008 run, hence was delayed
  - Readiness
    - ☆ Tests in December, operational in January



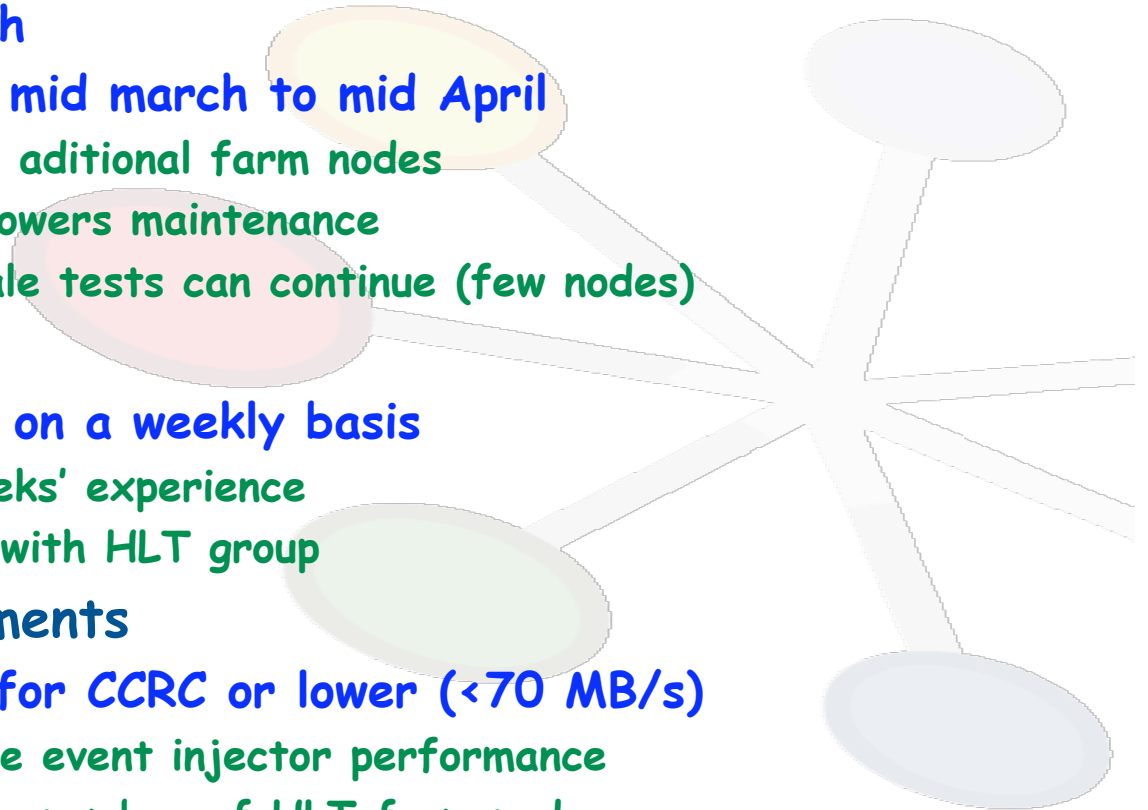
- **Description of the activity**
  - Define FEST conditions
  - Start a run
    - ☆ Events are injected and follow the normal path
  - Files export to Tier0 and distribution to Tier1s
    - ☆ CCRC'08 repetition
  - Automatic reconstruction jobs at CERN and Tier1s
- **Short test periods**
  - Typically a week
  - Depending on results, take a week interval for fixing problems
- **Vary conditions**
  - LO parameters
  - Event rates
  - HLT parameters
  - Trigger calibration and alignment loop





## Timeline for FEST'09

- First full integration tests
  - January
- FEST'09 runs
  - February, March
  - Necessary stop mid march to mid April
    - ☆ Installation of additional farm nodes
    - ☆ Pit 8 cooling towers maintenance
    - ☆ Very small scale tests can continue (few nodes)
- Run scheduling
  - Will be decided on a weekly basis
    - ☆ Using past weeks' experience
    - ☆ In agreement with HLT group
- Resource requirements
  - Same order as for CCRC or lower (<70 MB/s)
    - ☆ Depends on the event injector performance
    - ☆ Depends on the number of HLT farm nodes
  - Small CPU needs for reconstruction





## Additional activities on FEST'09

- Calibration and alignment loop
  - Performed on CERN CAF
    - ☆ Dedicated LSF batch queue on dedicated machines (currently 10 slots)
  - Mostly new calibration validation
    - ☆ Using low rate calibration stream
    - ☆ Short jobs for fast response (10 kevt/s, i.e. 10,000 s)
  - Update ConditionsDB
    - ☆ Check streaming synchronisation
  - Restart reconstruction
- Exercise physics analysis for first data
  - Concentrate on early physics
    - ☆ Validate MC generators (cross-sections of various channels)
  - Use small statistics for extracting first physics results
    - ☆ Typically 100 Mevt/s
      - \* Possibly re-use same events twice
    - ☆ Inclusive differential cross sections



## Additional shutdown activities

- Depends on the LHC machine schedule
- Larger scale FEST'09 in May
  - Full HLT farm
- LHC TED run data
  - May?
  - Using particles from dumping the SPS beam inside the injection tunnel (TED stopper)
  - 2 to 5 muons per  $\text{cm}^2$
  - Allows detector alignment (LHCb's cosmic runs!)
- LHCb prepares for being able to take data in May



- Currently user jobs use private pilots
  - Production jobs use generic pilots with role=production (Production Manager's credentials)
- We would like to move to generic pilots for all jobs
  - VOMS role=pilot
    - ☆ See VO card, to be installed at all sites
  - Allows to apply VO policy inside DIRAC
    - ☆ Relative priorities production / analysis (even with multiple groups)
    - ☆ Keep site management simple
      - \* VOMS roles: lcgadmin, pilot (used for submission)
      - \* VOMS roles: production, user (used by payload)
  - Test of generic pilots
    - ☆ Accepted in June, postponed waiting for deployment of role=pilot
      - \* Also waiting for ganga's DIRAC3 backend (include analysis jobs)
    - ☆ Will switch to generic pilots as soon as enough sites are OK
    - ☆ WARNING: unknown roles should be understood as "user", not as "sgm" as in many sites



- CPU normalisation
  - Need to set a max time limit in jobs
    - ☆ Benefit from short queues
  - Which unit?
  - Queues time limit **MUST** be expressed in a standard unit
    - ☆ WLCGUnit? E.g. 1000 kSI2k?
  - Jobs used CPU **MUST** be expressed in the same unit
    - ☆ OS returns real seconds
    - ☆ Need conversion factor per WN
      - \* Why not an environment variable? Parameter in /proc/cpu?
      - \* Should be identical to that used by the batch system
  - Normalisation needed for "time left" utility
- SAM jobs - Site reliability
  - Reliable results for storage
  - Still problems with elaborate CE test sensors
    - ☆ Added yesterday a few simple sensors in order to make results conservative
    - ☆ Apologies for the long-standing problem (lack of manpower...)



- No reason to change drastically the overall amount of data
  - Real data
    - ☆ Split year in two parts:
      - \*  $0.5 \cdot 10^6$  s at low lumi
      - \*  $4 \cdot 10^6$  s at nominal lumi ( $2 \cdot 10^{32}$ )
    - ☆ Trigger rate constant: 2 kHz
  - Simulation:  $2 \cdot 10^9$  events
- New assumptions for (re-)processing and analysis
  - Lack of experience from 2008 non-data
    - ☆ More re-processings in a first instance in 2009
    - ☆ More calibration checks (done at Tier0)
  - Envision more analysis at CERN with first data
    - ☆ Increase from 25% (TDR) to 50% (first part) and 35% (second part)
    - ☆ Include SW development and testing (LXBATCH)
  - Almost same events sizes and processing times as in TDR
    - ☆ Small adjustments due to more precise estimates
      - \* Some up, some down...
- If we don't use all resources, it's just it's not time yet...



- **New resource requirements document**
  - **Being prepared**
    - ☆ **Below are very preliminary numbers**
  - **Take into account existing data**
    - ☆ **To be scrapped at some point when new data comes**
- **First hints**
  - **Tier2s: no changes (CPU increase by 10%)**
  - **Tier1s:**
    - ☆ **CPU increase 20%**
    - ☆ **TxD1 increase 10% - should add T1D0 caches**
  - **CERN (Tier0 + CAF + LXBATCH + Analysis)**
    - ☆ **CPU increase by a factor 3: 1 MSI2k to 3 MSI2k**
    - ☆ **TxD1 modest increase - should include T1D0 caches**
- **Important note**
  - **CPU requirements expressed in MSI2k.years (i.e. integrated)**
  - **Sites should allow time variations!**
    - ☆ **3 MSI2k.year # 3 MSI2k maximum share**



- Due to lack of 2008 real data
  - Use MC data for further testing the whole system
  - FEST'09
    - ☆ Full system test including HLT, online monitoring, streaming etc...
    - ☆ Short periods (one week) from January to March
    - ☆ More FEST in May if no beam
- 2009 real data
  - Use TED runs (as soon as SPS is running)
    - ☆ Allows completion of calibration and alignment
  - First pilot run
    - ☆ Equivalent to 2008 run
    - ☆ Useful for detector understanding / tuning
    - ☆ First physics results expected even with  $10^8$  events
  - Nominal (LHCb) luminosity run
    - ☆ Assume  $4 \cdot 10^6$  seconds
- Resource needs are being re-evaluated
  - Minor changes except for CERN CPU (3 MSI2k)