# CDF Offline Operations

Aidan Robson
Glasgow University

Run II Computing Review, 14 Sept 2005

- Data production chain
- MC production

Universal concerns:

- ❖ transparency of computing
- ❖ low operational load

# Unified computing framework

New SAM-farm went into production at end of May:
achieved unified computing framework for all aspects of CDF computing

Production
MC Production
User analysis

all use CDF Analysis Farm (CAF) technology

custom submission + Condor job management

Uniform interface “grid-like” except we choose where to send jobs
Can submit to any grid farm if a machine is set aside to act as a
CDF headnode and make it “look like a CAF”

- Already seen benefit:
  augmented SAM-farm with user CAF nodes for special processing job

Production executable packaged in “gridified” way
– contains all required libraries, can run anywhere

# Fast production turnaround

On 23 August achieved goal of

- starting processing block of data
- run-range including runs taken up until one month previously
- with final calibrations

Now proceeding in blocks of ~ 1 month

– processing begins one month after last run

– data available ~6 weeks after last run

# Stable reconstruction

Reconstruction algorithms frozen since January.

We have designated a reference platform and dataset
– with every change that potentially affects reconstruction
we validate and understand any differences

Between 24 Aug and 10 Sept:
production farm processed ~300M events
crash rate **1 in 100 million** (3 events total!)

All data taken since the last shutdown (since Dec 2004)
is processed consistently.

Recent code releases distributed under both Fermi Red Hat Linux 7.3 and Scientific Linux 3 (SL3) – we ensured executables running under OS in which they were built produce identical results

**Production farm using SL3**

# “One pass” calibration scheme

Streamlined calibration scheme required to allow 6 week processing turnaround

Three stages of calibration, different requirements on statistics/analysis

1. Central Tracker drift times / silicon pedestals, bad channels
   – totally automatic, normally available by 1 day after data-taking
2. Beamlines
   – farm runs special process continuously to generate beamfit ntuples; fitting job by cron: minimal intervention
3. Other subdetectors:
   Calorimeter, Time of Flight, EM timing (~8 systems)
   – calibration ntuples made for each subdetector
   – two-week period for calibrators to run analysis
   – calibration tables gathered and validation run on W / Z / J/□
   – results examined and formally signed off

# "One pass" (contd.)

Now going through 4th cycle of 1-pass scheme – running smoothly

6 Sept: ntuples up to 4 September sent to calibrators
19 Sept: deadline for calibrations to be in DB
22 Sept: validation samples available, processed with new calibrations
4 Oct: sign-off meeting (2 weeks for validation);
farm ready to begin processing

**Developing monitoring**

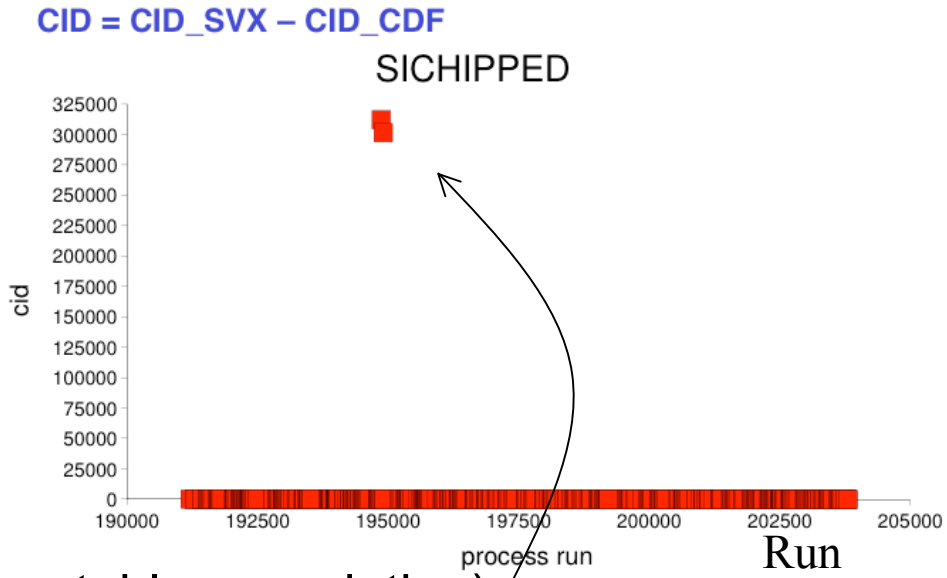


– existing monitoring checked only for presence of automatic calibrations

– developing content checks
both at ***database level***
(table ID numbers - correct run – table association)
– especially important in 1-pass scheme
and ***physics level*** (eg beam offset monitoring)

# Data Quality Monitoring (DQM)

Dedicated group (MOU-level commitment)

Enters at two stages
- shortly after data first taken and beamlines generated, to catch problems quickly
- after final processing of data, to provide good run list

# Re-reprocessing

Concern from last year's review:

"CDF should take into account that, despite their laudable plan for a single pass reconstruction, they might need further reprocessing "

Experience: We have already had to re-reprocess some of 2005 data.
Main reason: faults found in calibrations / calibration merging

- Currently on 4th cycle of 1-pass processing.
- Learned during 1st and 2nd cycles (different problems!);
  3rd cycle was problem-free.

So far spent ~1 month of farm time re-reprocessing
We believe unexpected reprocessing introduces contingency factor ~1.1

Farm has spare capacity and has capability of harnessing user CAF
– reprocessing doesn't imply need to double farm capacity

- Calibration validation maturing.
  eg DQM checked beamlines for next farms runperiod  no problems
  – new monitoring had already caught and fixed issues
    at the level of a few runs with inconsistent beamlines.

# Farm operations

- Today's capacity ~ 0.75THz corresponding to 18M ev/day
  migration to GCC compiler optimised for speed gave ~10% performance gain
- Good day of data-taking: 5M ev recorded by detector
- Farm operating efficiency (time farm operating  CPU load)
  > 90% in July/August, sustained over weeks.

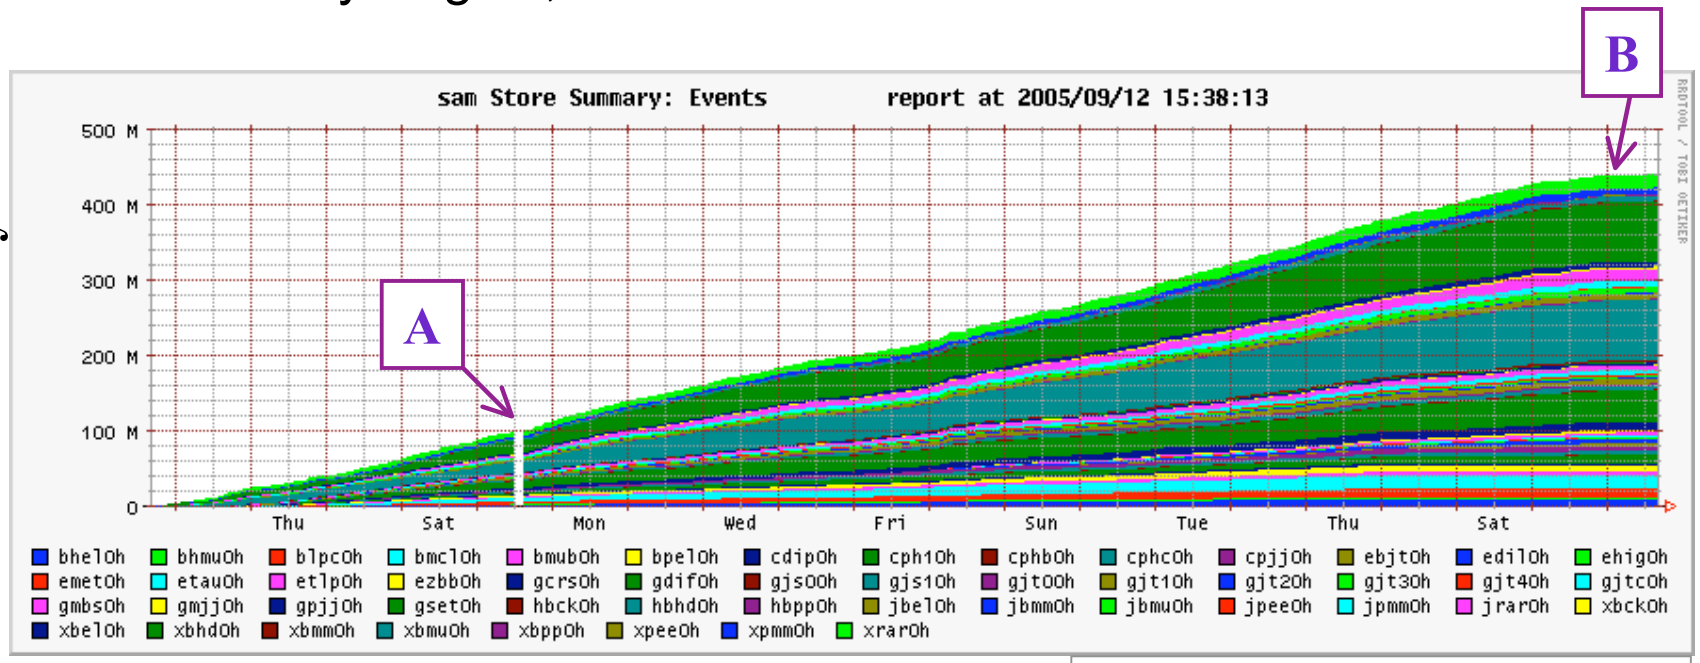


factor ~1.4 over raw data: for two streams write compressed datasets in addition to full DST

**50 datasets created from 6 raw data streams**

# Farm operations (contd.)

**A** Upload of offline luminosities
- operations job not performed optimally
- quick recovery

**B** Approaching end of processing: 3 streams completed
AND user jobs generating extreme load on database

Sample of recent miscellaneous problems leading to fall in farm CPU usage:

1. enstore robot problem
2. central CDF code machine goes down
3. file transfer requests stack up on farm output server
4. DB sees extreme load from user jobs
5. crash of server that determines tape destination for file
6. many processes hanging waiting for dCache

well-understood, fast response

more understanding needed

# Farm operations (contd.)

Attention turned to monitoring so we are able to

– triage problems

job submission? Concatenation server?
SAM? dCache? Enstore? Network?

– formulate well-defined problem-reports/requests to experts

Effort starting now to keep track of frequency of different problem sources

Currently have offline shift-person who monitors key plots and numbers
Automating and raising level; ultimate aim:

***make offline systems like any other detector subsystem monitored in control room***

# Operational load

**Farm**

**Data Handling**

**Database**

**SysAdmin**

**CAF Admin**

**Coordination**
Heads of Offline, Deputy Head
Head of Offline Operations
Offline Operations Managers
Calibration coordinator

**DQM**

**MC Production**

**Code Management**

**Calibrations**
Beamline generation
Calibration merging
Subsystem calibration
(many of them  coordinator)

development of improved monitoring reduces the load on providers of services that are needed continuously

Try to ensure each task covered by >1 person – no bottlenecks

Close interaction with ops managers; weekly reports from each group

Enlist help from physics groups - eg validation

# MC production

Run-, luminosity-dependent effects included automatically
(silicon configuration, beamline position, additional interactions)

We generate MC datasets corresponding to real runranges

- For each run generate MC events corresponding to configuration and average instantaneous luminosity
- Total # events proportional to total $\int L$
- Generate with fixed # events / unit $\int L$

– ***allows natural extension of MC runrange***; no need to restart

Result for 2002/03 data is many small files corresponding to runs – data handling system requires large files – bottleneck was manual concatenation

"CDF should invest some effort in improving the operating efficiency at the file merging stage […] Currently, manual intervention is necessary"

Problem solved – automatic concatenation now implemented
so concatenation is not a separate step

# MC production (contd.)

Remote sites set up for mc production
– select site on the command line at job submission time

Also will be able to take advantage of non-CDF-dedicated resources made available via Glide-In

MC corresponding to 2005 data will take ~2THz*months to generate
– will be generated over space of 1 or 2 months
– comparable to offsite capacity

Overall ratio of MC:Data events generated up to 2004 is ~ 1:1

# MC production (contd.)

Once MC generated:
moved to physics group buffer;
automatically read by upload server that catalogues and writes to the data handling system.

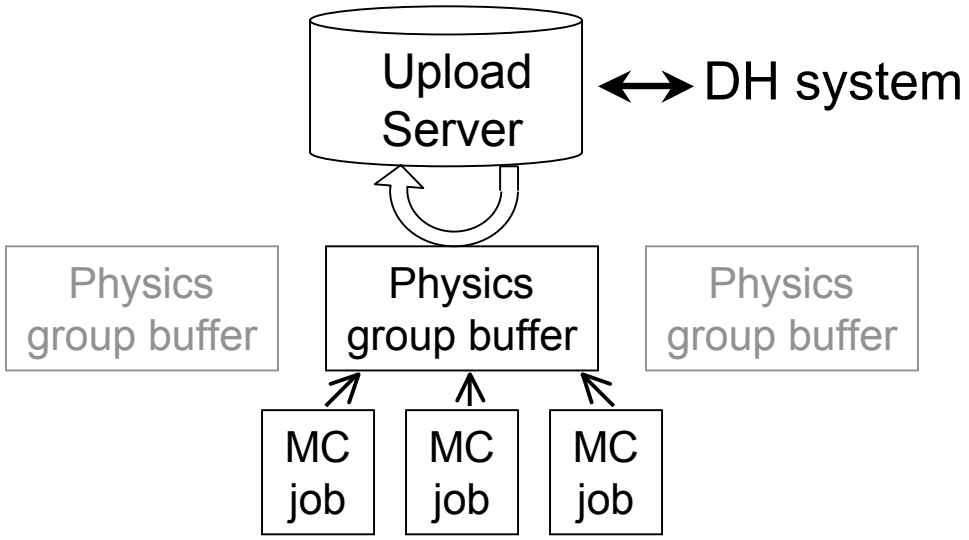


Changing role of MC production group:

2004: maintaining infrastructure and submitting jobs

2005: maintaining infrastructure and coordinating requests;
job submission moves to physics group representatives

# Ntupling task force

New centralisation of mc, data ntuples
– optimise use of CPU, disk

Analysis procedures matured:
• essentially two infrastructures used for ntuples
  – serve ~50% analyses
  – number is growing; new users choose one of the two
• can keep ntuple catalogue in DH system

**Ntuple-making becoming operational cycle**

New physics disk pool in testing for ntuple storage

# Summary

- Production cycle established and stable
  - 6-week turnaround
- MC production mature
- Ntupling developing into operations cycle; use of shared ntuples growing
- Monitoring undergoing enhancement; to be moved to the control room

# Backup

# Frontier

Frontier remote database replication fully validated and integrated in imminent code release

- draws load from Oracle DB read
  - implications for Oracle cost (licence cost/connection)
- however Frontier does not help wrt SAM database access
  - every job running SAM project writes to the database.