

Monte Carlo production for the CLIC CDR

Our experience

S. Poss, C.B. Lam, J. Strube, C. Grefe

CERN

September 22, 2011

Table of contents

- Introduction
- Production framework
- CLIC productions
- Other considerations
- Conclusion

Introduction

- 6 benchmark analysis:
 - $e^+e^- \rightarrow h\nu_e\bar{\nu}_e$,
 - $e^+e^- \rightarrow H^+H^-$, $e^+e^- \rightarrow H^0A$,
 - $e^+e^- \rightarrow \tilde{q}_R\tilde{\bar{q}}_R$,
 - $e^+e^- \rightarrow \tilde{\ell}\tilde{\bar{\ell}}$ ($\ell = e, \mu$),
 - $e^+e^- \rightarrow \tilde{\chi}^\pm_i\tilde{\chi}^\mp_j$, $e^+e^- \rightarrow \tilde{\chi}^0_i\tilde{\chi}^0_j$,
 - $e^+e^- \rightarrow t\bar{t}$ (500 GeV).
- Plus all the backgrounds

Total number of events generated in the last 6 months: $17 \cdot 10^6$, events processed: $\sim 100 \cdot 10^6$.

Part I

Framework

Production framework

DIRAC has a Transformation System:

- Create a workflow object (XML representation of a job), and let the system create the jobs for you
- Productions are identified unique ID
- Automatic resubmission of failed jobs: minor monitoring needed

It **saved us a huge amount of time and man power.**

Production framework

The DIRAC transformation system comes with a useful set of API commands: get properties, change them, update inputs, etc.

Full support for the ILD and SiD software frameworks, plus support for WHIZARD and PYTHIA code. Additionally we have support for generator level cuts (StdhepCut tool from Lars Weuste), Overlay input, etc.

Defining steps

- Due to uncertain CPU usage, it was decided to split all steps in different jobs.
- Implies specific productions for:
 - Generation: WHIZARD/PYTHIA.
 - Simulation: Mokka/SLIC.
 - Reconstruction: Marlin/LCSIM-SLICPandora-LCSIM.
- Limitation: few number of events per step (10 – 50), so small files to carry around.
- Linking of productions is straightforward, more on that later: **data driven production management**.

All files produced are stored in the DIRAC replica and Metadata catalog

Replica catalog

- Use of LFC is somewhat **limited in performance sometimes** (more true for the LHC though)
- **DIRAC provides a catalog** to overcome this
- Has also the possibility to **write and read in the LFC via the same interface**.

In the futur (coming month), when adding a file in the DIRAC FC, it will also end up in the ILC LFC. It would be nice that the files be always added in both catalogs.

Metadata catalog

- ILC stores the metadata encoded in the file names.
- Somehow restricts the number of meta info that one can add to a given file once the file is created.
- DIRAC FC comes with a metadata catalog, that sets the **metadata at directory level** (e.g. production ID, software packages, cross section, polarisation, etc.) and at File level (number of events, etc.).
- Flexible as one can add (and remove) metadata on the fly.
- Searchable: find EvtType=tt Datatype=SIM
- Possibility to set ACLs.
- Also comes with convenient API commands.

Web interface is under development, should be available for tests in the coming weeks.

Production Monitoring

Monitoring done with web site:

- Overview of production statuses: % of the files processed, failure rate
- Access to jobs from a given production in one click
- Monitoring and accounting plots also available

Part II

CLIC Production

Defining a production

1. Define the input data (if any) via a metadata query (e.g. meta['ProdID']=186)
2. Define the application and versions you want to use
3. Define the output files (if any)

Data driven procedure: if new files are added to the catalog that correspond to the input data query, jobs are automatically created.

Generating Events

Use WHIZARD for most samples:

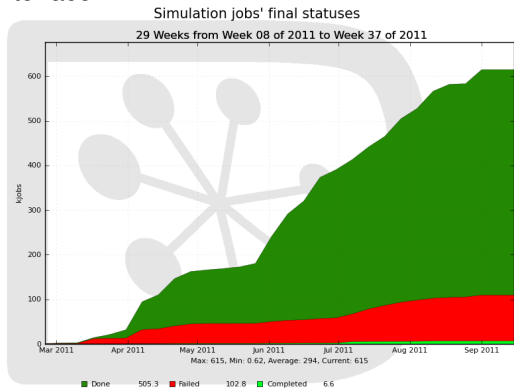
- Handling of **beam spectrum and ISR included**
- Adding a **new process** is straight forward
- Does the correct computation of diagrams and interference:
accurate cross section and kinematics estimation
- Limitation: **final states have no width**, cannot do $t\bar{t}$, W^+W^- or Z^0Z^0 .

Use PYTHIA for those 3 samples:

- Final states have the correct width
- Not so accurate estimation of cross section or kinematics, but simple enough to be very similar to those produced by WHIZARD
- Has **beam spectrum interface via CALYPSO**, ISR is PYTHIA default
- Limitation: **no flexibility at all**, adding a channel cannot be done easily

Simulation

Define productions for both detector concepts, but very similar interface.



Generated on 2011-09-22 06:57:37 UTC

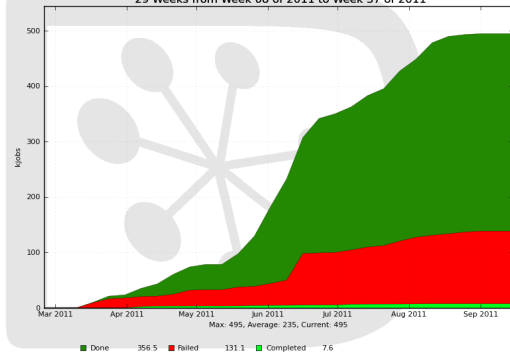
Reconstruction

As for the simulation: dedicated productions are created, automatically using files produced by the simulation step.

Without Overlay:

Reconstruction (no overlay) jobs' final statuses

29 Weeks from Week 08 of 2011 to Week 37 of 2011

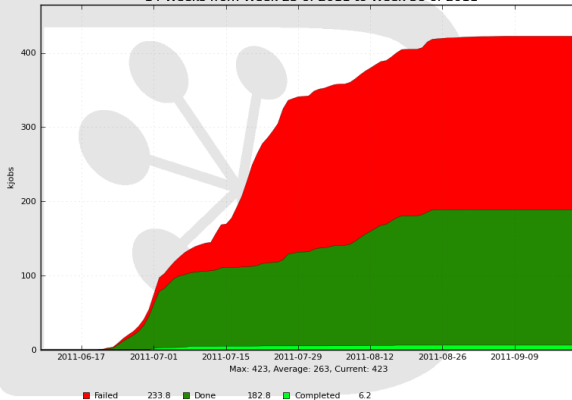


Generated on 2011-09-22 06:57:04 UTC

Reconstruction with overlay

Reconstruction (with overlay) jobs' final statuses

14 Weeks from Week 23 of 2011 to Week 38 of 2011



Generated on 2011-09-22 06:58:25 UTC

Overlay handling (1)

CLIC detector benchmark require to reconstruct the events with overlaid $\gamma\gamma \rightarrow had.$ 60 BX with 3.2 interactions per bunch Xings are considered for every signal event.

Simulated background files contain 100 events each (more would be better but CPU is a constrain). For a 10 signal events file, one need ~ 2000 bkg events, or 20 files per job.

Overlay handling (2)

This did not work: storages (CERN in particular, but also IN2P3) do not cope with such a load (2000 jobs trying to access 20 random files each represent a huge number of queries).

Solution: **merge the files randomly once** to reduce the number of files needed: merge them by packs of 200.

Works but a bit tedious: requires to have all the files locally once. Plus one needs a **very large number of files** to have enough combinaisons.

Needed solution: **use random access** (LCIO v1.51 has it, but not LCSIM). Then only one big file is needed on every site with direct access. No more transfers are needed.

Part III

Prospects

Prospects

Testing running in the US sites: when it runs, we'll need to see for disk space here and there. CERN disk requests will be made: for the moment, only 40Tb are available on disk (as much as needed on TAPE though).

Will finish implementing dedicated service that will know the available processes, to easy users' lives. For the moment, a text file holds all the processes available and the corresponding WHIZARD version.

Will push to get the File Catalog web browser.

Part IV

Conclusion

Conclusion

- Produced in 6 month \approx 100 million events in 9 months
- Used $2 \cdot 10^{10}$ KSI2K seconds of CPU in \sim 20 sites
- Data produced correspond to \sim 130 TB
- Most problems were due to storage access, in particular for the overlay files
- ILC DBD could benefit from DIRAC usage

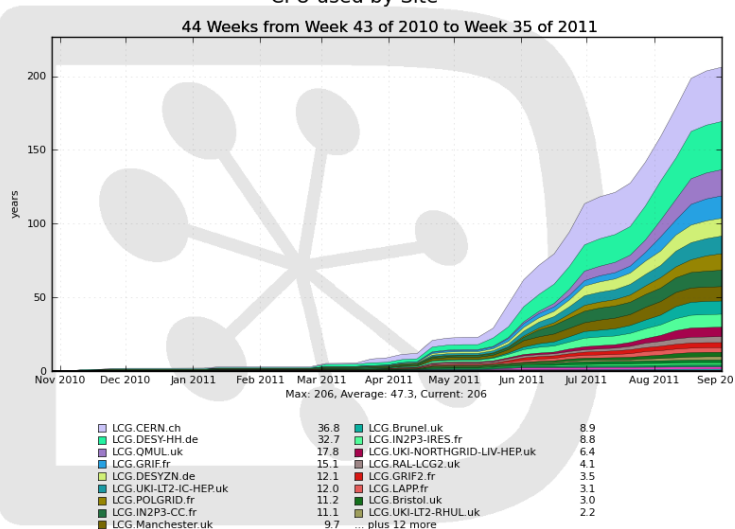
At least, the data produced should be stored in both catalogs.

Part VI

Backup slides

Resources used: CPU

CPU used by Site



Generated on 2011-09-08 13:38:05 UTC

Resources used: Storage

Site	Production (TB)	User (TB)
CERN	128	20
IN2P3	4	9
RAL	4	28
KEK	0.02	0
IMPERIAL	1.6	0
TAU	$4 \cdot 10^{-4}$	0

Sites that can be used in addition: DESY, Bristol, BONN, RALPP

Technical aspects

Dropped pilot mode as some sites did not support it, using private pilot mode

Steering files are installed on the sites like applications (dependency relation), so no need to pass them in the input sandbox of the job.

Software installed in the Shared areas of the sites when/where possible. Does not use SAM framework (yet) as DIRAC takes care of the software installation and removal.

Services availability

100 million events correspond to ~ 4 million jobs. On average 1000 concurrently running.

File catalog is the most expensive service in terms of CPU and simultaneous queries to DB. Had to duplicate service (not DB).

It would be useful to replicate the DB and other services to have more instances available (more VO boxes needed).

Same problem with JobManager with high load.

File names conventions

- Base path is /ilc/prod to separate from /ilc/user data.
- Path holds most relevant info (machine, energy, process, detector, data type (gen, SIM, REC, DST), prodID)
- File name is process_prodID_jobID.ext (ext=stdhep,slcio)

Data validation

When simulating data with 2 frameworks, it's necessary to validate it before running. For that purpose **TOMATO** was designed: convert the stdhep to slcio (stdhepjob supported by DIRAC), and run TOMATO (Marlin processor) to create histograms of "significant" distributions.

"Significant" is specific to the different analysis, so every working group needs to provide the variables to be plotted.

Was used for the $t\bar{t}$ analysis.