

PRS muon Meeting,

CERN, 7 May 2002

Status of Muon production at Legnaro

Stefano Lacaprara, INFN and Università di Padova

Summary:

- ▶ Status and problems;
- ▶ Events produced;
- ▶ Analysis farm: how to allow access to data.

Status and problems:

- ▶ Digi 10^{34} with and w/o Tk are running: most dataset are almost completed, MB_pt1 just started;
- ▶ LOT of problems!, as usual;
- ▶ Major problem with network interface on some PU servers
 - ★ Sometimes Net broken with heavy load (well below bandwidth available) $\Rightarrow$ all jobs running hangs;
 - ★ Upgrade of Net card driver, new one has a work-around to reset the Net interface in case of problem;
 - ★ hopefully solved;

► Major problem with lock-journal server:

- ★ Too many files open, $\Rightarrow$ all jobs running hangs;
- ★ moved lock-journal server to other machine;
- ★ hopefully solved;

► BUT: lot of work spent in recovering situation

- ⇓ recover federation, locks, resuming dead jobs, zombies, corpses ...;
- ⇓ Problems happened always few minutes after we left...
no CPU used during that nights (murphy's law...);
- ⇓ now running at lower rate to avoid problem (60 jobs, 120 nodes);

► Important problem: lot of crashes, exit status 139

- Unknown reason, not reproducible, $\sim 20\%$ of the jobs;
- All production sites affected by this problem, not only Legnaro, not only Muon (eg. CERN 21% jobs crashed);
- Good news: the crash is not “persistent” (as was last year), but “transient”, namely, a job resubmitted process the event where it have crashed before (sometimes crashes on some following event(s));
- enough iterations (submit-resume-submit-resume-...) will produce all the events;

- ▶ Minor (?) problem: some event ($\sim 0.1\%$) skipped during digitization
 - ↪ Events skipped because of corrupted calo hits;
 - ↪ Apparently the same problem we had last year (although less frequent), the skipping procedure was indeed developed last year trying to solve that problem;
 - ↪ Skip works, no crash, but it's a symptom of something strange (related to other crashes ??)
- ▶ All jobs with exit status not 0 and jobs resubmitted have output run not validated: must be done by hand

- ▶ Inconsistencies between no. of events apparently produces (according to log) and what is written in MetaData
 - ↪ Boss (and so RefDB) does not knows about skipped event, not serious;
 - ↪ dsDump and ds2html (which read MetaData) sometimes show collection size $<$ than events produced (according to logs) when there are resubmitted jobs (via CARF:ResumeRun cards);
 - ↪ Asked Vincenzo, investigating: a problem of dsDump (which does not count properly the no. of events produced) or a real problem (ResumeRun does not work properly?);

- ▶ Due to all these problems, we have not yet transferred all the DBs on the analysis disks;
- ▶ Transferring DBs to CERN is not available, yet, for the same reason: to achieve this we should set-up the transferring tool and start the dispatching.
- ▶ If we want to copy at CERN only a selection of a Dataset, it is possible (via cloning Digi): how much human work does it require?
- ▶ Low luminosity digitization will start as soon as the high lumi is finished.

Events produced $\mathcal{L} = 10^{34}$:

Dataset	Pythia/CMSIM	ooHit	Digi	notes
mu02_h120_WW2mu	10000	10000	10000	completed!
mu02_h140_WW2mu	10000	10000	9000	
mu02_h160_WW2mu	10000	10000	9386	
mu02_h180_WW2mu	10000	10000	9139	
mu02_h200_WW2mu	10000	10000	9000	
mu02_h150_ZZ4mu	10000	10000	9500	completed!
mu02_h130_ZZ4mu	10000	10000	10000	
mu02_h200_ZZ4mu	10000	10000	9500	
mu02_h300_ZZ4mu	10000	10000	9500	
mu02_h500_ZZ4mu	10000	10000	9000	
mu02_MB1mu_pt1	279358	156732 (56%)	11827	just started, No Tk Digi
mu02_MB1mu_pt4	404992	255732 (63%)	247246	No Tk digi
mu02_MB1mu_pt10	110513	86914 (78%)	85690	