

Minos and I:

My adventures with keepup,
roundup and batch processing

Marek Zielinski
Rochester/Fermilab

My story with Minos batch

- I started in May 2014, as Guest Scientist in SCD, taking over some duties from Rashid
 - He nicely (though briefly) explained some things about roundup, then told me: “We did not even touch on where the real work is. You’ll learn from experience.”
- Art and Adam were of tremendous help
- I learned about detectors, scripts, bluearc dirs, dcache dirs, lists, logs, error codes, spill, cosmic, atmos, sntp, mrnt, historical artifacts...
- I helped to run and monitor daily keepup, roundup, top-up, catch-up, MC reco, re-reco...
- When things work, they run in cron. When they don’t, one has to do detective work.
 - Errors are not too frequent, but can be quite unique
 - Need to learn how to respond
 - Experience comes with time

Batch by examples

- A repeating error: beam db not OK
 - Cron keepup will not process such subrun
 - Standard action: resubmit subrun by hand after some time and hope it will work now
 - Better: ask nicely Robert or Donatella about status
 - Perhaps even better: learn how to check beam database, resubmit when OK
- More rarely: fpe
 - Standard action: resubmit without opt
 - Not so fast! Last week we suddenly had hundreds of ND subruns failing with fpe. Experts found a culprit in calibration tables.
- Conclusion: be alert, check everything, question everything, notify experts, learn.

Batch in steps (1)

- Keepup runs in cron
 - Any error fixing or reprocessing requires resub by hand
 - also creating input list and modifying status lists by hand
 - know proper flags for special processing
 - Modify crontab file for special situations (eg. no beam for a period of time, a prolonged detector or db problem)
 - Watch cron emails, status lists
 - good/bad/zap/nospill runs
 - Watch files in staging dirs
 - esp. nearcat, farcat, neardet, fardet
 - files older than a few days need investigation
 - Check for jobs that died on grid nodes
 - fetching node logs currently requires stealing identity of the submitter

Batch in steps (2)

- Roundup runs in a loop
 - stop/restart by hand for maintenance days
 - start by hand after crashes/system problems
 - if roundup does not start, check STOP flags, then logs
 - recently creating dirs in /pnfs fails with srm errors and roundup stops
 - quick solution: create dirs by hand; then fix the problem
 - know how to force roundup for incomplete runs
 - watch ROUNDUP/WRITE staging dir
 - old files indicate problems with copying roundup output to dcache or with samweb registration
- Similar comments apply to big production passes and MC reco processing
 - but using different flags and/or scripts

Living in interesting times

- All was quite fine till the fall shutdown.
- Then we lost slf5. Howie's scripts were not slf6 or grid capable. It was a bit scary.
- Before end of shutdown:
 - Adam created KGA (KeyGear and Ambrosia)
 - They run with jobsub_client and on FermiGrid
 - I converted keepup scripts to samweb, and Art later converted roundup
 - We upgraded keepup to elm4
 - Experimentally discovering various things that need to be changed when going to a new code release...

All's well that ends well

- Transitioning from old to modern scripts went better than we had courage to expect
 - From slf5 only to slf6 capable
 - From sam to samweb
 - From condor_glidein to jobsub_client
 - From Minos cluster to the grid
 - From heavy use of bluearc to mostly dcache
- And expanded and more coherent params
- And improved ability to diagnose job problems

Passing the torch

- Since last May I have been running the offline meeting
 - Great venue to learn about status and current issues
- Paola B and Jeny T are attending the offline meetings since the beginning of the year
 - Learning, asking questions
 - Helping to resolve cert/proxy issues
- Anna M joined us in Feb (and seems to enjoy it 😊)
- Minerva rep (Harold YR) is involved with ND issues and Minos meetings – it helped communication w/ Minerva
- We had several Minos batch training events
 - A full length (1.5 hr) tutorial movie is available at affordable price (BYOS, it's 510 MB)
 - “Marek's distilled notes for minos batch work” (aka cheat-sheet) are wildly popular
 - Q&A sessions and emails

Keep in mind

- IMHO, the most important for the production processing crew is to:
 - Stay in touch with Minos colleagues
 - Always know the current status of the experiment
 - Question all unexpected outcomes of processing, investigate, alert experts
- This is not a push-button operation
- Experience is everything, comes with time

How about future?

- cvmfs transition is ongoing
- KGA: (almost) OSG ready, let's do it
- Roundup – replace it also?
 - Use FTS with concatenation?
- Analysis scripts can/should follow the path to the bleeding edge of FIFE
 - samweb, jobsub_client, outside grid, pure dcache
 - No more bluearc (ok, ok), cpn, srm...
- Plug Minos scripts into FIFE batch processing tools when available
 - Automated re-try of failed runs (with specific error codes) would be a nice feature
 - Web monitoring of job status/progress

The last seven bullets

- It was interesting experience...
- ...a bit challenging at times...
- ...during a big transition for Minos batch...
- ...which, in the end, was very successful
- The Minos team is fantastic
- It has been a blast!
- Thanks, and keep enjoying chocolates, guitar, the best song ever written about keepup, the Schreckenbergerer, yoyos, hyperspheres...