

Job Processing

- Database consolidation
- Task recovery
- De-cronification
- T1 overflow primary reco
- Panda

Database consolidation

- Task_req ~ etask
 - Currently sync back and forth, e.g. status
 - T0 has no task req – data driven
 - prod is 'whim-driven'
 - Status in *requested*, pending, submitting, running, ..
- EJobExe ~ Panda (activated/archived jobs)
 - biggest table – mostly duplicate info
- EjobDef? Dynamic job definition
 - to process every file
 - some for MP(500evt), some for single core(50)?

Task Recovery

- Examples
 - repro of some events fails
 - wait for new cache
 - AOD corrupted during merge
 - lost files
 - turn of HIST to save RAM
- Create recovery task to (re)produce missing files
- Attach to original container
- What about the TAG?
 - new merged AOD – update TAG guid?
 - no. Verify AOD before upload(in merge TRF) and replicate.

De-cronification

- AKTR and jobdef has many crons
 - process bulk tasks, tasks(pending,submitting), more jobs, add attempts, abort jobs, stats
 - create output datasets, freeze them, add to containers
 - subscribe datasets, delete them
- Representation of desired state and engines to get there
- Messaging triggers between different actors
 - eg. if output datasets exist, then create jobs
 - poll for dataset, or trigger – do both
- Delay between task def, and jobs created is artificial
 - partly due to crons, and partly deliberate to allow user back-out
 - 'Submit Task' – are you sure(yes/no) – can always abort (c.f qsub)

Example: Task Request

- Task(chain) request web form(click submit)
- Sanity check – inputs,release,... accept task
- Create output datasets and containers
- Create jobs for bamboo 1st task
- Monitor loop
 - more jobs, increase maxattempt, abort jobs, notify user, jobs for next task
 - requestor can have active input
 - abort, tobefinished, more RAM
 - clickable actions for user analysis too?
- Task Done? Freeze datasets, subscribe to wherever

T1 overflow primary reco

- Expect p-p reco to keep-up at T0, for this year
- IF backlog approaches 10 day disk buffer
 - may need to use T1s
 - should prepare and test?
- T0 reco and T1 repro have differences
 - data driven prodsys task definition
- Test run was HI
 - cut short, but monitoring was criticized

Panda

- Break cloud boundaries
 - small clouds can trap big hi-prio task
 - 50k cores process high priority task
 - 10 smaller tasks better?
 - transfer to destination T1 is bottleneck
 - Merge HITS at T2 before transfer(or MP)
 - Prodsys or DDM?
 - more general if DDM does it
- Compulsory generic merging post-prod step for users

Random stuff

- Feedback or actual walltime/RAM usage from scout jobs
- Task chain monitoring