


The ADC Operations Story

Andrej & Ale

ADC Commissioning & Integration & Operations

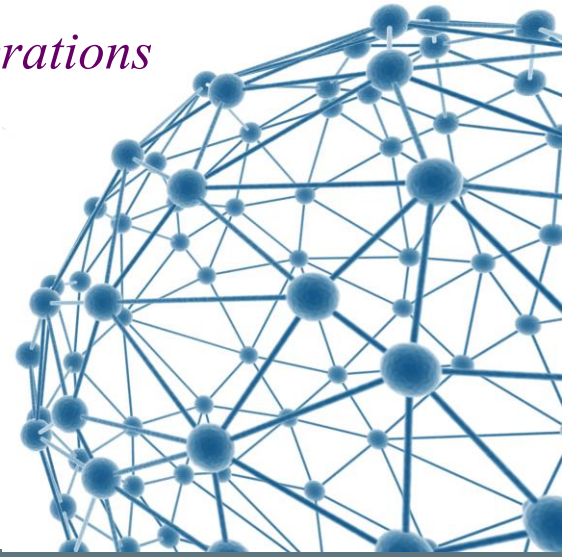

Activities

- "Standard"
  - But not so many things are standard:
    - MCORE pile reconstruction
    - Derivation Framework
- Commissioning
  - Prodsys2
  - Rucio
  - Now also Event Index
  - …and more…!

CPU Resource usage

Not good in the past months: many sites often not fully exploited

- New workflows are not optimal for the "old" ProdSys1+PandaBrokering, some needing complex job requirements
  - HighMem: not clear how much, not only 2 vs 6
  - MultiCore vs single core
  - Job duration
- Often not enough workload
- Many issues due to commissioning of other services (e.g. Rucio)
- ADC relies on the ATLAS SW: too often happened in the past months that the workflows we run were not deeply tested
  - To be addressed with ATLAS Offline

Disk Resource usage

- Too many primaries on Tier1
  - Tier1s and good Tier2s spaces should be considered on the same level of reliability. Tasks could be even assigned directly to Tier2s.
    - Good Tier2s reliability metrics are under constructions, good status
- Deletion:
  - too much time to wait for green lights, transient data cleaning manual
- New workflows changing the way in which we use the present spacetokens:
  - Merging in Tier2s is great, but we have to dimension the spaces accordingly
    - Goal is to get rid anyway of Proddisk but we need more time (Rucio migration, and better darkdata handling, see next slides)

- Near Future will solve (hopefully) most of these issues:
  - ProdSys2 will manage transient data
  - new Data Placement model: lifetime!

Job failures

- Need clear separation between SW and Site failures
  - Automate retries
- SW failures
  - still too hidden, difficult to debug
- Site failures
  - today not clearly disentangled
- Some failures ends up in darkdata
  - We need to expose these data and clean up!
- Proposal:
  - improve the HC jobs to expose better the pure site failures.
  - more HC to validate “new” workflows

Workflows latency optimization

- Network
  - Direct access vs transfer
  - Reduce startup and completion latency
    - Today DQ2+SS+FTS+Callbacks putting quite a lot of overhead: looking forward to Rucio
    - Remote I/O may speed-up for some workflows
  - MCP optimization: use network topology for job assignment
- Job retries: can be done in different sites (for non-sw issues)
- Log files:
  - Kept in Tier2s for now
  - Object store in the coming months

Opportunistic resources

- Panda brokering today can be affected by opportunistic resources not working properly:
  - E.g. too many transferring in a cloud can led to non assignments
  - We need to organize the opportunistic resources in a way that won't affect the production sites
    - First priority is to have pledged resources fully busy!

- Cloud resources usability:
  - Clouds should give to us "flexibility": we didn't see this in the September MultiCore pileup reconstruction campaign
    - More in one of the next sessions

Data loss recovery

- One of the feature of ProdSys2 is to be able to automatically recover the data losses
  - This should solve us most of the issues (biggest problems are for the unfinished/non-replicated tasks)
- What do we do for the “old” data?
  - If possible ProdSys2 can redo what’s needed
  - We should rely on multiple replication for critical data (for workflows completely different in ProdSys1 – ProdSys2)

Monitoring

- Also Monitoring tools are changing/improving
  - BigPanda is now THE PandaMonitoring
    - People getting used to it
    - More help needed in the near future
    - “campaign” monitoring, e.g. the reprocessing new monitoring: we would need something similar in near future
  - DDM dashboard(s)
    - Rucio now publishing into the DDM dashboard
    - FTS activity in a different dashboard
    - Xrootd in another different dashboard

ADC Shifts

- ADCoS Senior under revision
  - ADCoS Expert to be revisioned too
- DAST under heavy load but very useful
- AMOD to be reviewed
- Plan under discussion between ADCOps and ADC Shifts Coords

- we don’t have today people on duty over weekends

“New” frameworks: Reminder!

- ProdSys1 has to go
  - ProdSys2 first priority is to implement full ProdSys1 functionalities
  - Almost there!
- DQ2 has to go
  - Rucio first priority is to implement the same DQ2 functionalities

DQ2 → Rucio migration

- More details in the next sessions

- Soon or later (sooner than later) we have to JUMP: mid/third week of November is the target date
  - And then be agile in fixing issues on the fly

Agile deployment

- Various teams have developed their frameworks in the past year(s)
  - Often quite “on their own”
- Integration is ongoing since various months:
  - Many things/issues/details coming to surface only now:
    - It’s obvious but better to be clear: we are all ADC

… ongoing…

- Many of the new features will come with Rucio and ProdSys2
  - Cannot even test these before full integration
- can’t do all the commissioning/changes at the same time: we need to prioritize!
  - There are things that we if don’t do *now* we will keep them for the next N years!

N.B. won't be so easy…!