




Rucio Commissioning

ADC Operations



Rucio Commissioning

- Won't go into the details of the commissioning
 - ToDo list and priorities didn't change respect to 6 weeks ago
 - <https://indico.cern.ch/event/276501/session/1/contribution/3/material/slides/1.pdf>
 - <https://indico.cern.ch/event/276501/session/17/contribution/58/material/slides/1.pdf>
 - But a lot of work has been done in the meanwhile!



AleDiGGi
Log Out

AtlasComputing

ATLAS Collaboration

ATLAS TWiki
ATLAS Protected
ATLAS Computing
Public Results

Create a LeftBar for this page

- Index
- Changes
- Notifications

TWiki > AtlasComputing Web > AtlasComputing > AtlasDistributedComputing > ADCCommissioningMilestones (10 Oct 2014, SimoneCampana)

ADCCommissioningMilestones

- + [Rucio Commissioning Milestones](#)
- + [Prodsys-2 Commissioning Milestones](#)

Rucio Commissioning Milestones

- 03/10/2014 Running single PanDA job, reading input from Rucio and storing output in Rucio. **DONE**
- 10/10/2014: same as above, but at scale O(10k) jobs. **DONE**
- 17/10/2014: running JEDI tasks (of various types) with input and output in Rucio.
- 21/10/2014: deployment of DQ2 client v2.6
- 24/10/2014: running Prodsys-2 tasks with input and output in Rucio.
- 24/10/2014: achieve horizontal scalability of Rucio services
- 31/10/2014: Full chain stress test successfully concluded (1M files/day for 1 week)
- 31/10/2014: full support for tape data
- 03/11/2014: Start migration of datasets from DQ2 to Rucio
- 24/11/2014: Rucio in production
- 15/12/2014: decommission DQ2 service

Prodsys-2 Commissioning Milestones

- 24/10/2014: at least 10% of MC production in Prodsys-2
- 01/11/2014: legacy group production running in Prodsys-2 only
- 03/11: 40% at least 40% of MC production in Prodsys-2
- 17/11/2014: Reprocessing validated in Prodsys-2
- 17/11/2014: Derivation Framework in production from Prodsys-2 point of view
- 24/11/2014: HLT reprocessing commissioned in Prodsys-2 (to be used in M7)
- 01/12/2014: at least 100% of MC production in Prodsys-2
- 15/12/2014: decommissioning Prodsys-1 service



Rucio Commissioning Milestones Progress

Summary of the milestones

03/10/2014 Running single PanDA job, reading input from Rucio and storing output in Rucio. **DONE.**

10/10/2014: same as above, but at scale $O(10k)$ jobs. - Status: see tasks below. (Need to understand Holding, Failed subscriptions - new task will be injected today 14Oct, failure rate should be fully understand and comparable with normal tasks)

17/10/2014: running JEDI tasks (of various types) with input and output in Rucio. - Status: evgen, simul, merge, recon **DONE**

21/10/2014: deployment of DQ2 client v2.6 - Status: Ready, more testing useful (JIRA blocking tickets: None; Minor issues: [RUCIO-724](#)) **DONE**

24/10/2014: running Prodsys-2 tasks with input and output in Rucio.

24/10/2014: achieve horizontal scalability of Rucio services

31/10/2014: Full chain stress test successfully concluded (1M files/day for 1 week). **DONE**
(02/10 - 10/10)

././.....: Simple check for stuck rules (WebUI or CLI or DB qry)

31/10/2014: full support for tape data - Status: write to tape, stage and deletion of single file works. Biggest issue is split of disk/tape by conveyor and telling FTS to bringonline ([RUCIO-672](#))

03/11/2014: Start migration of datasets from DQ2 to Rucio

24/11/2014: Rucio in production

ProdSys2 task with input/output in Rucio

- Full chain is ongoing
 - Not all task finished (evgen and simul OK)
 - not a real issue, more an error while defining the chain.
 - Still assigned manually
- Not spot up to now real issues related to Rucio

Horizontal scalability

- Present situation:
 - possible to scale horizontally with “fixed” configuration
 - **Judge:** rucio-daemon-prod-01, singleprocess, multithread
 - **Conveyor:** rucio-daemon-prod-02,06,07,08, multiprocess, multithread
 - **Hermes:** rucio-daemon-prod-02, singleprocess, singlethread
 - **Undertaker:** rucio-daemon-prod-03, singleprocess, multithread
 - **Kronos:** rucio-daemon-prod-03, singleprocess, singlethread
 - **Automatix:** rucio-daemon-prod-04, singleprocess, multithread
 - **Transmogrifier:** rucio-daemon-prod-04, singleprocess, singlethread
 - **Necromancer:** rucio-daemon-prod-04, singleprocess, singlethread
 - **Reaper:** rucio-daemon-prod-05, multiprocess, multithread
 - Trying to increase the available VMs now
- “Automatic” scalability:
 - The database change for automatic configuration of scalability of daemons is in production. The patch to support this is implemented, not in production yet.

Rucio Tape management: stagein

- Blocker issues resolved
 - Retries with timestamp
 - Bringonline parameters passed to FTS3
- Sorting out now details of RSEs for tapes
 - One RSE for “staging area”
 - Mostly internal, between judge and conveyor

Rucio *now*

Current Aim – Decommission DQ2

- DQ2 infrastructure should be decommissioned by end of October (end of Quattor support)
 - This is rather optimistic
 - End of the year is a better estimate
- For DQ2 decommissioning we need to concentrate on getting DQ2 functionality in Rucio
- Many things work but not at the scale we need
- The major things currently missing for minimal DQ2 functionality
 - ongoing → – handling tape data (only minimally tested and fragile)
 - done → – collectors running (propagation of ATLAS info into Rucio)
 - done → – multiple FTS support (being able to copy files with more than one FTS)
 - done → – source selection (being able to specify specific source sites)
 - done → – transfer retries/stuck rule repairing (handling failures better)
 - – Stability and scalability of transfer/deletion daemons
- ATLAS Distributed Computing is complicated and all the small details add up to a lot of work
 - Some of these things are not just developers' responsibility, but on ops too

Ongoing: spot inconsistencies, under fixing

done

tested

Other issues

- Things we can live with if they are done soon after DQ2 decommissioning

- Still to be done – fine-grained permissions on DDM endpoints
 - Need group/country spaces with permission based on voms attributes
- Ongoing → – DaTRI replacement
 - Current DaTRI can use new DQ2 clients until Rucio is ready (rucio UI?)
- done → – blacklisting/downtime handling
 - We can handle this manually for now
- To be validated – accounting
 - Rucio produces a dump similar to DQ2 which can be consumed by the current accounting system
- lost file recovery and monitoring
 - Special cases can be handled manually
- ongoing → – rucio cli
 - Not so urgent due to new DQ2 client
- Still to be done – monitoring history of files/datasets
 - Very useful for DQ2 ops, but can live without it for a short time
- Ongoing → – subscription monitoring (similar to dr checker)
 - Critical for T0 export, should be there by then
- See rucio-ui – better transfer logic (choose disk source over tape, more than one file per FTS job, all the complex stuff in DQ2 SS)
 - Fine tweaks which can be applied after experience

Recovery done,
monitoring ongoing

Patched, not deployed

ongoing

DQ2 – Rucio migration

November

- Jump from DQ2 (infrastructure, i.e. CentralCatalog, SiteServices, Deletion agent) to Rucio (as dataset catalog plus all infrastructure)
 - Migration of datasets in background:
 - Forecast is of 1 week of needed migration time
 - Minimize the time in which DQ2 and Rucio handle together production data
 - The “hot” (used by Panda) datasets can be affected (despite the DQ2 2.6 “transparent” clients), hopefully for a short period (order of one day)

Backup

Horizontal scalability from Vincent



© Not Invented Here .com

Barnes • Southworth • Zugale

