




# Rucio at BELLE II

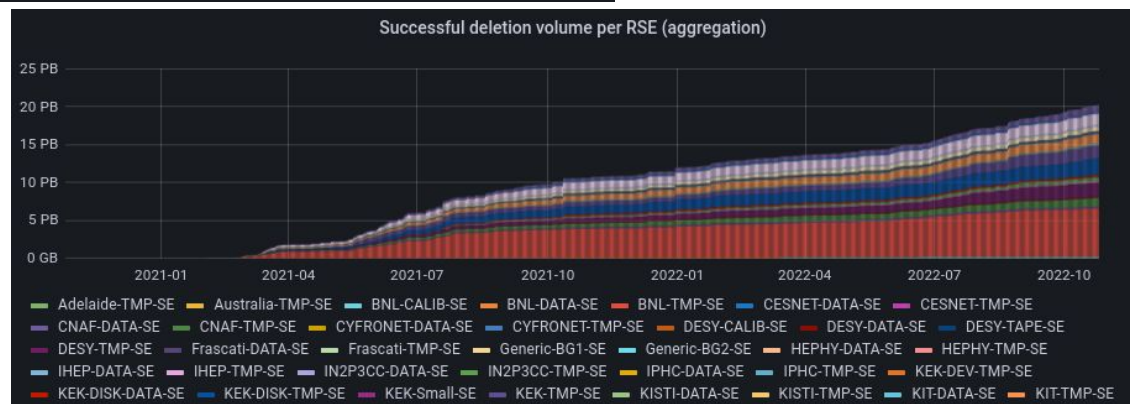
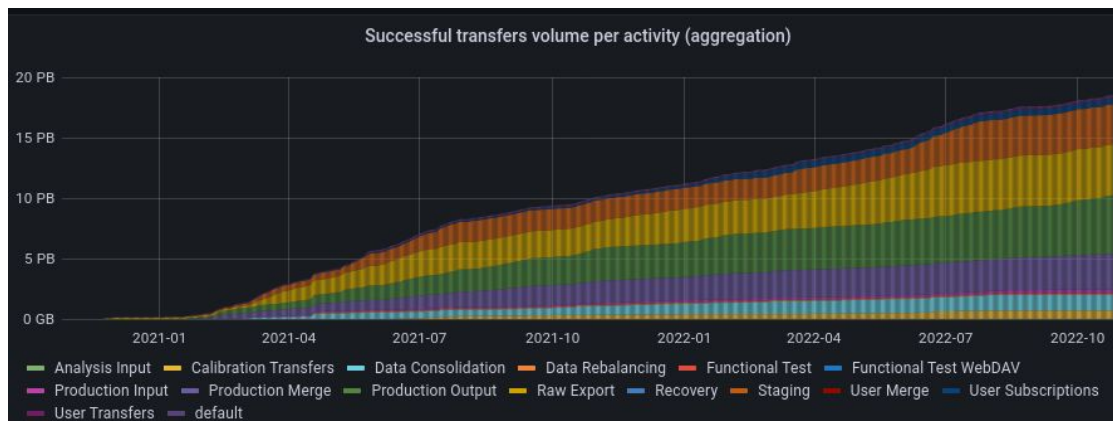
Cédric Serfon (BNL) on behalf of DDM team

November 11th, 2022



# Introduction

- Belle II now uses Rucio in production for 1.5 years
- No major issues. It fits for Belle II purpose



# Overview

- Running on DB on Postgres with 2 frontend servers + 2 daemons (VMs) + 1 WebUI + 2 tracers
- We are currently still using 1.28.7 since the version of base DIRAC we use (v7r2) is not totally python3 compatible
  - Plan to migrate beginning of next year to v7r3py3
  - Once it is done, plan to move to 1.30.x
- Monitoring stack on InfluxDB, elastic, Grafana

# Main features used

- We use all the “main” daemons :
  - Judge
  - Conveyor (not using preparer yet)
  - Transmogrifier
  - Reaper
  - Undertaker
  - Hermes2
- In particular all the rules other than the ones created by DIRAC are generated by subscriptions
- More daemons under test or used on demand :
  - Kronos
  - Recovery

# Belle II Rucio unit-tests

- In the past, we suffered from some regression due to a change not taking into account Belle II convention
- In addition some methods specific for Belle II were introduced (e.g. `dirac add_file`)
- For these reasons, votest test suite was introduced
  - It includes specific Belle II tests
  - In addition, most of the other tests that only worked with ATLAS convention also work with B2 convention
  - We also spotted some methods that can only work for ATLAS convention (see later)

Summary

Jobs

✓ setup

✓ test (atlas, centos7, postgres14, ...)

✓ test (belleii, centos7, postgres14, ...)

Run details

Usage

Workflow file

test (belleii, centos7, postgres14, 3.6, votest, votest-belleii)

succeeded 18 hours ago in 12m 28s

Build Images and Run Tests

```
4818 XFAIL test_auditor.py::test_auditor_check_survives_failures_and_queues_failed_rses
4819 reason: [NOTRUN] loops endlessly sometimes
4820 XFAIL test_rse.py::test_get_rse_attribute[True]
4821 FIXME: Calling functions which change the rse attribute should invalidate the cache.
4822 XFAIL test_authentication.py::test_saml_success
4823 The WebUI isn't linked to CERN SSO yet so this needs to be fixed once it is linked
4824 XFAIL test_authentication.py::test_saml_fail
4825 The WebUI isn't linked to CERN SSO yet so this needs to be fixed once it is linked
4826 XPASS test_judge_evaluator.py::TestJudgeEvaluator::test_judge_dataset_grouping_all This is a test for a known bug. See issue 5251
4827 = 507 passed, 299 skipped, 4 xfailed, 1 xpassed, 70 warnings in 265.79s (0:04:25) =
```



# Metadata

- 2 kind of metadata :
  - Accounting metadata : Metadata that are used for accounting
  - Generic metadata : All the rest
- Accounting metadata were imported to Rucio to generate accounting summaries that are injected to influxDB
- Generic metadata : We started some tests to evaluate if we can use Rucio to store them. More information in dedicated talk

# Staging

- Staging from data on TAPE is still a manual operation
- Working on a new agent part of BelleDIRAC to automatically stage datasets :
  - A top level “directory” i.e. container is provided by the production system
  - Container is resolved into a list of datasets
  - Staging requests are sent (rules on the datasets) taking into account size of buffer and of the destination RSEs
  - If volume to stage higher than buffer size, throttling applied
- This is kind of daemon could also be useful to some other collaborations. Maybe something can be implemented in Rucio ?

# Use of Rucio features in analysis tools

- Rucio CLI is not used by the end-users
- But Rucio API and Rucio features are being used :
  - download\_client now used by the end-user download tool. Some minor change was needed to handle “/” in the file names
  - Implementation of a new feature to set an automatic lifetime on the rules created by the Rucio File Catalog in belledirac
    - Based on scope, e.g. all user\* scopes get a 30 days lifetime
    - Lifetime of the rule is extended whenever a new file is added
  - Start testing global quota
  - Introduction of specific RSEs only usable by local or country accounts (a la LOCALGROUPDISK)



# Policy package

- Policy package implemented for Belle II containing :
  - Schema
  - Belle II tuned permissions
  - extract\_scope, lfn2pfn algorithm
- Comments :
  - generic policy maybe too generic/restrictive. Difficult to use as is
  - Some “tricks” implemented (e.g. possibility to enable access based on country) would be great to share

# Recent problem founds

- Reaper memory leak. Seems to be related to gfal deletion with gsiftp
- Problem with conveyor in release 1.28.x (multi-source failed jobs not properly handled)
  - Fix in 1.29.x
- Some WebUI views based on dumps are expected to have these dumps in CET (not even UTC)
- Identification of some methods only working for ATLAS.  
Non-exhaustive list :
  - REST API for creating DID sample
  - Method to extract scope, name for PFN
  - Redirector not working with “/”

# Things to improve

- How to check why a rule is stuck ?
- Possibility to force a stuck rule reevaluation immediately
- <https://github.com/rucio/rucio/issues/536> ;)
- Possibility to have archived tables into another database with API transparently using the proper DB

# Misc.

- WebDAV TPC : We recently moved ~70% of our TPC transfers to WebDAV
- Getting rid of voms :
  - IAM test instance available. Test are being performed. A few open questions (e.g. how to handle different robot certificate that share the same email)
  - Token : We need Dirac v8r0 to start having a look at time (early next year ?)
- More in Silvio's talk at WLCG meeting

# Conclusion

- Belle II is happily using Rucio for the last 1.5 year
- Test of some new functionalities ongoing : generic metadata, popularity
- New challenges ahead (e.g. tokens)