




Rucio and Belle II


C. Serfon on behalf of the Belle II DDM team

September 29th 2021



Introduction

- Belle II is a B-factory based at KEK (Tsukuba Japan) and an international collaboration of institutes all over the world with more than 1000 collaborators
- Phase III started in Spring 2019 and is currently in data taking mode
- Aiming at 50 ab^{-1} by 2031

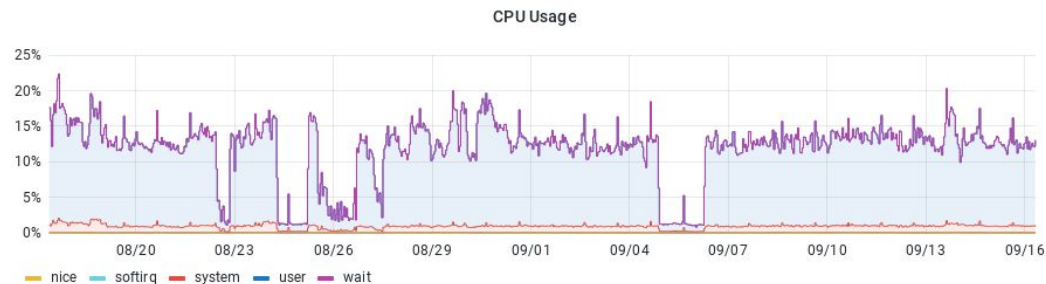


Introduction

- Lots of changes since the last Rucio workshop (March 2020). At that time, Belle II was still not using Rucio, and there was no plan yet for transition
- Since then, Belle II successfully moved to Rucio in January 2021. This migration as well as integration of Rucio with BelleDirac was presented at different meetings including vCHEP 2021.
 - [Integration of Rucio in Belle II](#)
 - [The Rucio File Catalog in DIRAC implemented for Belle II](#)
- This talk will give a bit more details of what happened after the migration, and what are the future plans

Belle II Rucio instance

- Rucio instance hosted at BNL using PostGreSQL
 - ~100M replicas/18PB registered (AFAIK, biggest PSQL instance running Rucio)



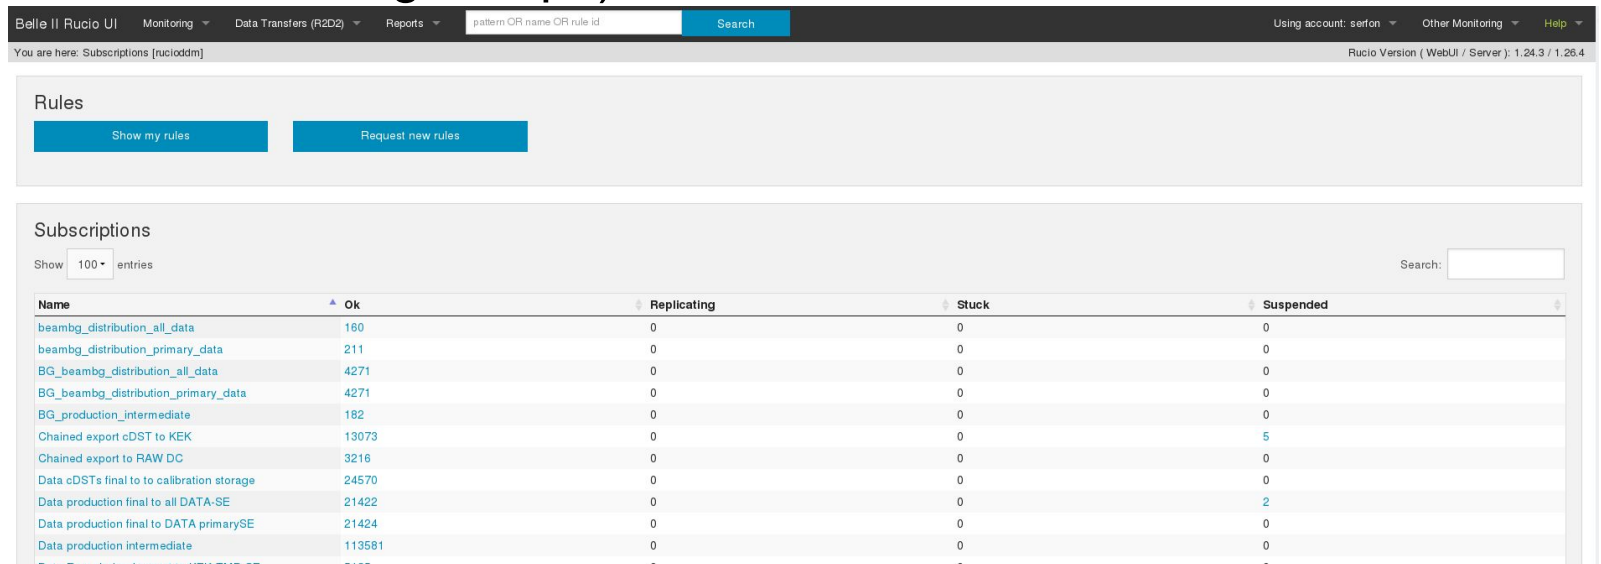
- Interaction rate 80 to 100 Hz
- Hosts :
 - 2 servers
 - 2 daemons node
 - 1 WebUI node
 - 2 tracers (introduced recently)

Developments performed for the migration

- Integration of Rucio and BelleDIRAC (more details after)
- Chained subscriptions
 - Subscription that can create a rule dependant of a first one
- Hermes2 :
 - New daemon that submits messages to different services (ActiveMQ, InfluxDB, elastic, etc.)
 - Will replace the old Hermes that is limited to ActiveMQ
- Grafana Dashboards for monitoring [transfer and deletion](#) or [site occupancy](#)
- More info on the CHEP papers

New services enabled after the migration

- WebUI :
 - Introduced mainly to monitor subscription (every rules but the one created by the Rucio File Catalog in Dirac) are generated by subscriptions
 - Had to figure out many details since lack of documentation (e.g. for views using dumps)



Belle II Rucio UI Monitoring Data Transfers (R2D2) Reports pattern OR name OR rule id Search Using account: serfon Other Monitoring Help

You are here: Subscriptions (rucioodd) Rucio Version (WebUI / Server): 1.24.3 / 1.26.4

Rules

Show my rules Request new rules

Subscriptions

Show 100 entries Search:

Name	Ok	Replicating	Stuck	Suspended
beambg_distribution_all_data	160	0	0	0
beambg_distribution_primary_data	211	0	0	0
BG_beambg_distribution_all_data	4271	0	0	0
BG_beambg_distribution_primary_data	4271	0	0	0
BG_production_intermediate	182	0	0	0
Chained export cDST to KEK	13073	0	0	5
Chained export to RAW DC	3216	0	0	0
Data cDSTs final to calibration storage	24570	0	0	0
Data production final to all DATA-SE	21422	0	0	2
Data production final to DATA primarySE	21424	0	0	0
Data production intermediate	113581	0	0	0
Data Production export to KEK TMD SE	5195	0	0	0

More services that are being deployed/developed

- Popularity :
 - Tracer enabled + Kronos. Being integrated into end-user tools
 - Lack of common tools to analyze traces (available only in ATLAS gitlab). Would be useful to have a common set of scripts
 - As for the WebUI, Thomas' departure is a big loss :(
- TAPE organized recalls :
 - Convergent evolution with what is done in ATLAS "TAPE Carousel". Service is being developed on the BelleDIRAC side
 - Maybe something can be done on the Rucio side to have a generic service that could be used by different Workflow and Workload Management Systems
- Metadata :
 - Evaluation of generic metadata in Rucio. Some read tests already performed
 - Work on implementing the possibility for DID to inherit the metadata of its parent

Interfacing Rucio and DIRAC

- Integration of Rucio into BelleDIRAC :
 - Rucio File Catalog plugin + Synchronization agents
 - Many challenges due to the differences between Rucio and DIRAC (e.g. DIRAC user/group/host vs Rucio accounts)
 - The Catalog does a bit more than a catalog since it also creates replication rules
 - Rucio and BelleDIRAC work nicely together
- The code developed for BelleDIRAC was merged into Vanilla DIRAC
 - Further changes are being done by J. Martyniak (Imperial College London) to support multiVO and rucio.cfg less setup
 - It is important that what is implemented in Vanilla DIRAC can be used by Belle II

Feedback about using Rucio

- Rucio fulfills what Belle II production system needs
- Belle II now uses Rucio for all the file transfers between SEs
- Rucio works with "datasets", nicely matches what Belle II have built with "datablocks" as the unit of data management
- Interfacing vanilla DIRAC to Rucio requires some tweaks for file-base data management
 - Rucio's feature, not allowing re-registration of deleted files, prevents rescheduled DIRAC jobs from registering output files when the file from previous failed attempt has been removed (in Belle II workflow)
 - We had to disable some features in DIRAC, not to have transfers initiated by DIRAC

Things to improve

- Operation documentation (e.g. options in configuration and what they do)
- Some part still very ATLAS biased (e.g. WebUI)
- Some topics are of interest for many communities and might not be a priority for ATLAS (e.g. metadata) :
 - Need some specific discussion/developments. Rucio Special Interest Group introduced recently can help on that
- Dedicated component responsible on component previously under Thomas' responsibility

Things to improve

- Would be good to introduce VO specific tests :
 - e.g. during our migration from 1.23.12 to 1.26.0, we suffered from a regression affecting RSE expression parsing
- Making replicas for a group of datasets is easy with a subscription, but reducing replicas, or erasing datasets require a bit of work
 - e.g. no easy way to reduce the number of replicas to reduce the number of replicas for all the datasets in a group, eg. specified by a dataset name pattern

Conclusion

- Rucio has been in use in Belle II for almost 9 months now
 - Helped in automatizing things
 - Provide new functionalities to end-users (e.g. asynchronous deletion, “collection”) that were not available before
- We’re happy that some of the developments done by Belle II (e.g. Hermes2, RFC in DIRAC) are interesting for other communities and will continue to propose new features