

Discussed on 2022-08-04;
Comments included

Rucio 1.30 release roadmap

“The Donkeynator”

[Martin Barisits](#)

on behalf of the Rucio team




Release plan

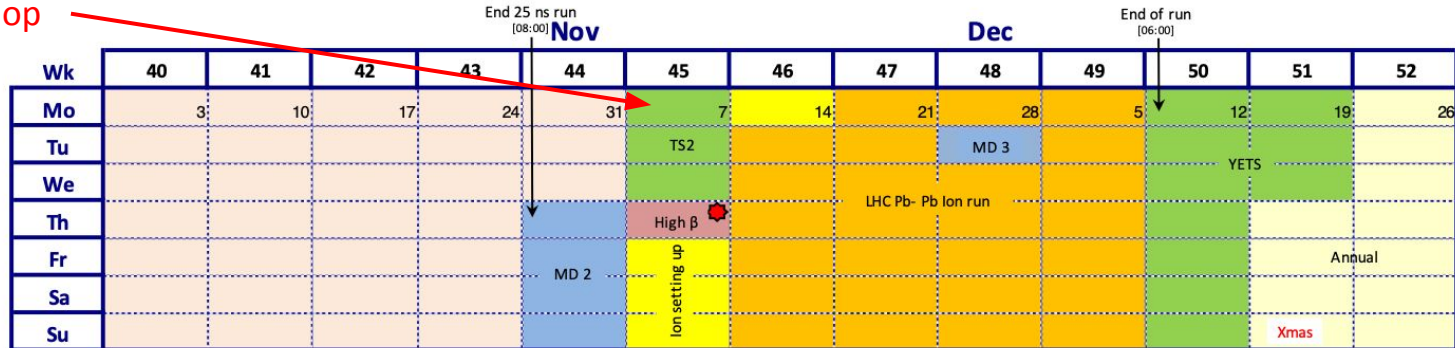
- Three feature releases in 2022
 - ~~1.28~~ ~~“Teenage Mutant Ninja Donkeys”~~ ~~March 2022~~
 - ~~1.29 LTS~~ ~~“Into the Donkeyverse”~~ ~~July 2022~~
 - 1.30 “The Donkeynator” November 2022
- LTS support period
 - ~~1.20 will be supported until 2021-06~~
 - ~~1.23 will be supported until 2022-07~~
 - 1.26 will be supported until **2023-07**
 - 1.29 will be supported until **2024-07**



Release date



Rucio workshop





Release retrospective outcome 1/2

- 1.29 release was one of our largest releases
 - Plan was very motivated, however, we were significantly “off-target” from the release roadmap
- Motivation and courageous goals are good, but we need to ensure that our roadmaps are realistic
 - Rucio community relies on our plans
 - Missing a feature for a release can create significant delays for others
- It is the responsibility of each component lead and developer to assess if the objectives are realistic
 - I will help with this assessment, but I cannot easily assess your external impacts
 - Operational issues and external priorities are normal and need to be part of our planning
- **Be courageous, but think about and question how much development load you have and how realistic it is to achieve the set goals in the given timeframe**



Release retrospective outcome 2/2

- We will add 1-2 checkpoints to each release sprint
- At the checkpoint we assess how much progress we made and where we should be at that point in time
 - If we identify discrepancies we discuss our options
 - Do nothing as situation is still under control
 - Add additional developers to a task
 - Re-prioritize: Drop tasks from the release to re-focus the attention to the remaining tasks
- Obviously, this does not help us in the planning, but it should make sure that the tasks with the highest priority make it into the release
- Collected data could be helpful to guide our decision making for future release roadmaps
- **For 1.30: Checkpoint 1: September 15**

Authentication [4 → 3]

@massnig, @rizart



Priorities

- ~~OIDC+SSO optimization and documentation~~ [#5437](#)

Tasks

-

Ongoing long term tasks

- New token workflows with storage-specific authorization [#4791](#)

Ideas

-

Clients [9 → 8]

@massnig, @rcarpa, @joeldierkes



Priorities

-

Tasks

-

Ongoing long term tasks

-

Ideas

-

Consistency checks [3 → 4]

@faluchet, @ivmfna1



Priorities

- Further integration of CMS-style consistency checking [#5438](#)
 - Changes included; Need integration in toolkit now

Tasks

-

Ongoing long term tasks

- Auditor upgrade [#3437](#)

Ideas

-

Core & Internals [27 → 23]

@bari12, @massnig, @gabrielefronze



Priorities

- Document and cleanup discrepancies between surl and lfn2pfn algorithms [#5129](#)
- Reduce to <20 issues
- ~~• Calculation of container length and events [#5422](#)~~
- ~~• Add session to API layer [#5439](#)~~
- Introduce heartbeat handler to all daemons [#5443](#)
 - 4 daemons missing!

Tasks

-

Ongoing long term tasks

-

Ideas

-

Data Pre-Placement [0 → 0]

@bari12



Priorities

-

Tasks

-

Ongoing long term tasks

-

Ideas

- Re-write C3PO with a plugin-approach

Database [9 → 6]

@massnig, @bari12



Priorities

- ~~Further integration of temporary tables to other major workflows~~ [#5441](#)
 - ~~This needs sub-tasks!~~
- Exchange of function-based indices with normal indices [#5440](#)

Tasks

- ~~Add missing query hints to code~~
 - ~~Somebody from CMS might look into this~~

Ongoing long term tasks

- ~~Optimize most common workflows~~ [#4793](#)
- Migration of all queries to new SQLAlchemy syntax (1.4 -> 2.0) [#5247](#)

Ideas

-

Dataset deletion [2 → 3]

@bari12



Priorities

- Allow undertaker to delete file
dids [#5154](#)

Tasks

-

Ongoing long term tasks

-

Ideas

-

Deletion [5 → 4]

@cserf



Priorities

- ~~Re-evaluate greedy = 1TB constant~~ [#5358](#)
- ~~Re-factor/cleanup deletion workflow~~ [#5442](#)
- Foreign key error [#5733](#)

Tasks

- Improve globus deletion workflows [#4575](#)

Ongoing long term tasks

-

Ideas

-

DIRAC [1 → 1]

@cserf



Priorities

-

Tasks

-

Ongoing long term tasks

-

Ideas

-

Docker & Kubernetes [0+3+4 → 2+4+4]

@mlasnig, @rcarpa, @ericvaandering



Priorities

-

Tasks

- Automatic scaling of kubernetes daemons
- Helm-chart testing [#h/95](#)

Ongoing long term tasks

-

Ideas

-

Documentation [9+4 → 12+3]

@bari12, @joeldierkes



Priorities

- ~~Improve the Client API documentation page #5444~~
- Create a good and clean Rucio installation (Server+daemons) page #5445
- Remove old documentation and forward to the new one #5446

Tasks

- Some re-organization of documentation needed
 - Most basic workflows, installation etc. need to be explained
- Some pages do not use verbose as they should
-

Ongoing long term tasks

-

Ideas

-

Functional tests [0→ 2]

@cserf



Priorities

-

Tasks

- Add documentation

Ideas

-

Ongoing long term tasks

-

Lifetime model [1 → 1]

@dchristidis



Priorities

-

Tasks

-

Ongoing long term tasks

-

Ideas

- Introduce Follower mode
 - ATLAS: Should at least advertise it

Messaging [4 → 4]

@massnig



Priorities

- ~~Hermes2 replacement of Hermes1~~ [#4610](#)

Tasks

- Replace hermes1 with hermes2

Ongoing long term tasks

-

Ideas

-

Metadata [2 → 5]

@robbarnsley



Priorities

- Merge `list_dids` and `list_dids_extended` [#5448](#)

Tasks

- Update and re-design `core.meta` module [#5224](#)
- Finalize external postgres extension [#5629](#)

Ideas

-

Ongoing long term tasks

-

Monitoring & Logging [3 → 4]

@maany, @mlasnig, @rcarpa



Priorities

- Further reduce logging [#5449](#)
 - Create overview of logs per component
- Rework prometheus/statsd metric names [#5804](#)

Tasks

- Prometheus client groundwork has been done
 - Needs more metrics and work

Ongoing long term tasks

-

Ideas

-

MultiVO [1 → 1]

@Thysk



Priorities



Tasks



Ongoing long term tasks



Ideas



Policies [3 → 3]

@jamesp-epcc



Priorities

-

Tasks

-

Ongoing long term tasks

- Switching permissions/schema/policy to policy packages [#4798](#)

Ideas

-

Probes & Alarms [1+7 → 1+4]

@dchristidis, @ericvaandering



Priorities

- ~~Dead code in the probes (And ATLAS specific code) to be cleaned [#p/83](#)~~
- Contributor guide [#p/84](#)

Tasks

- Move semi-required probes (e.g. obsolete calculation) to abacus daemon
- Convert CMS probes to SQLAlchemy

Ongoing long term tasks

-

Ideas

- ~~Long Term: Moving away from Nagios probes~~
- Moving probes to prometheus

Protocols [7 → 5]

@massnig, @rcarpa



Priorities

- Create a documentation page about integrating S3-type storage [#5450](#)
- Add support for S3 virtual hosted buckets [#5451](#)
- ~~Remove unused boto protocol~~ [#5346](#)

Tasks

- Some updates needed for xrootd implementation

Ongoing long term tasks

-

Ideas

- Rucio without rsemanager.py

Rebalancing [0 → 0]

@cserf



Priorities



Tasks



Ongoing long term tasks



Ideas



Recovery [6 → 2]

@cserf



Priorities

-

Tasks

-

Ongoing long term tasks

-

Ideas

- Archive support for necromancer
 - TPC transfer for archive contents?
 - Student project

Release management [6 → 4]

@bari12



Priorities

- Add aliases for all daemons to the release [#5131](#)
- Compile a clear list where the daemon names need to be changed [#5132](#)
 - Helm-charts
 - Monitoring
 - ...
- ~~Cleanup old tickets~~
- ~~Remove author lists in header of each module~~
- ~~Python 2.7 deprecation~~ [#5435](#)

Tasks

-

Ongoing long term tasks

- Rename daemons to understandable names [#4795](#)

Ideas

- Dependency testing of target platform

Replicas [6 → 7]

@massnig



Priorities

-

Tasks

- Add distance based sorting [#4020](#)
- –random functionality [#5418](#)

Ongoing long term tasks

-

Ideas

-

REST & API [7 → 9]

@bziemons, @bari12



Priorities

-

Tasks

-

Ongoing long term tasks

- Versioning for the REST API
[#4796](#)

Ideas

- Look into pagination for (heavy) streaming APIs
- Start thinking about API 2.0

Rules [16 → 16]

@bari12



Priorities

- Fix consistency of rule grouping dependent on order of ATTACH [#5251](#)
- ~~Flag transfer errors coming from tape sources~~ [#5410](#)
- Bug with rule consistency on CONTAINERS [#5447](#)

Tasks

- Reevaluate rule option: if the RSE expression doesn't point to the same RSEs anymore we should be able to reevaluate the rule. Would help a lot for decommissioning sites [#536](#)

Ongoing long term tasks

- Introduction of QoS to Replication Rules [#3419](#)

Ideas

- Possibility of a rule dry-run

Subscriptions [1 → 5]

@cserf



Priorities

-

Tasks

- Include the creation of delayed rules in subscriptions

Ideas

-

Ongoing long term tasks

-

Testing [9 → 18]

@maany, @joeldierkes



Priorities

- Write a developers testing guideline/manual [#5452](#)
- Add python 3.10 testing to CI [#5145](#)
- Re-factor the different test workflows to a common one
- Switch test to pytest

Tasks

- Improving usability of CI tests locally [#4047](#)
- More tests to factories
- ~~Make use of VO-Tests~~

Ongoing long term tasks

- Mypi static code analysis [#4455](#)

Ideas

-

Traces [2 → 5]

@mlasnig, @maany



Priorities

-

Tasks

-

Ideas

-

Ongoing long term tasks

- Trace schema enforcement

Transfers [12 → 9]

@rcarpa, @ericvaandering



Priorities

- Re-work/cleanup distance table [#5454](#)
- Re-factor conveyor-throttler

Tasks

- Remove full_mode from conveyor-receiver [#5453](#)

Ongoing long term tasks

-

Ideas

-

WebUI [13+11 → 2+29]

@maany, @cserf



Priorities

- ~~Finalize authentication workflows in WebUI 2.0 #5135~~
- ~~Remove DB dependencies completely from WebUI 2.0 #4319~~
- Release a beta webui 2.0 beta
 - X509, userpass, oidc auth
 - High priority views
 - User dashboard

Tasks

-

Ongoing long term tasks

- WebUI 2.0

Ideas

-