

# Rucio “internals” discussion

---

[Martin Barisits](#)

on behalf of the Rucio team



# Agenda

---

1. Release plan
2. Roadmap and (selected) key developments for 2020
3. Selected discussion topics
4. More discussion



# Release plan

---

- Three feature releases in 2019
  - **1.19** “Fantastic Donkey” **February 2019** ✓
  - **1.20 LTS** “Wonder Donkey” **June 2019** ✓
  - **1.21** “Donkeys of the Galaxy” **October 2019** ✓
- Feature releases in 2020
  - **1.22** “Green Donkey” **March 2020**
  - **1.23 LTS** “The Incredible Donkey” **June 2020**
  - **1.24** “Aquadonkey” **November 2020**
- Release candidates for all feature releases for functional validation
- [Long Term Support release policy](#) works well
  - 2 LTS patch releases produced for now
- Do we need more feature releases?



# Roadmap: Metadata

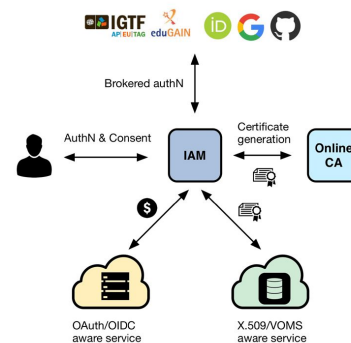
---

- Rucio metadata component will be refactored to address multiple, internal & external, backends via one metadata interface
- Backend examples:
  - DB-dids-columns: Hard defined columns in dids table
  - DB-dids-json: Generic metadata column in dids table
  - MongoDB: External mongoDB to store metadata
  - ...
- Metadata interface will decide based which backend to address for storing and querying metadata (policy driven)
- In development now, delivery Q1/Q2 2020



# Roadmap: AAI

- Token-based authentication (OIDC)
  - Extends authentication to OpenID Connect
  - Integration of IAM workflows to Rucio
  - Development done with funding from XDC project (H2020)
  - Authentication to Rucio
  - Authentication from Rucio to Storage/FTS
  - Different modes supported:
    - Orchestrate all transfers under admin token
    - Orchestrate transfers under individual user tokens





# Roadmap: AAI

---

- Rucio 1.22 includes [OIDC/OAuth2.0 AAI](#) features (See Jaroslav's presentations)
- 3 different CLI workflows were tested
- REST API accepts tokens from trusted IdPs (idp-secrets need to be deployed on Rucio server)
- WebUI in 1.22 supports OIDC login
- Testing! (And more testing!)
  - 1.22.0.dev3 deployed on WLCG DOMA Rucio instance
  - Accepts tokens from WLCG and XDC IAM
  - Need to test full chains (User → Rucio → FTS → Storage; Rucio → Storage)
  - TPC Stress test
- Requirements by other communities? (DUNE?)



# Roadmap: MultiVO

---

- See Eli's presentation
- Multi-VO mode will enable Rucio to host multiple VOs within one Rucio instance
- Designed in a way which is fully transparent to Rucio core
  - → Rucio core stays largely unchanged to avoid unwanted regressions
- ~8k+ lines of code already changed
  - Only few, minor regressions
- Additional changes coming in Q1/Q2 2020 which should bring the Multi-VO mode into a production-ready state



# Roadmap: QoS Replica Management in Rucio

- Replica management in Rucio is based on replication rules
  - Put 1 copy of `file.001` on a Rucio Storage Element (RSE) in `country=uk&type=disk`
  - `country` and `type` are RSE specific attributes
  - Rule engine finds eligible RSEs and picks one based on a set of criteria
- `type=disk` already is an expression of QoS, but it is only used in limiting the set of eligible RSEs for the replication rule

## Integration of Storage QoS in Rucio:

- QoS will become explicit part of the replication rule
- Rule concept very well suited to express VO QoS policies (cf. [whitepaper](#))
- Already some QoS-like prototypes orchestrated with Rucio: [ATLAS Data Carousel](#), [MAS](#)






# Roadmap: QoS with replication rules

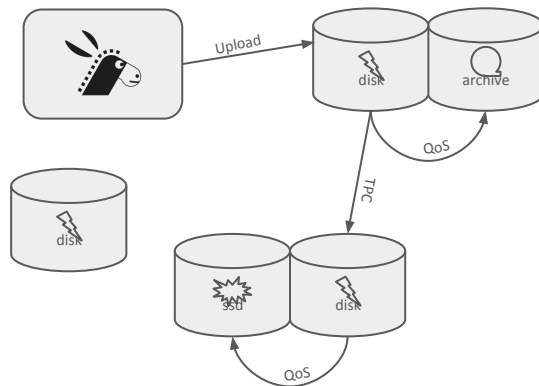
Rule today: 1 copy of data on `country=de&type=disk`

Rule tomorrow: 1 copy of data on `country=de` QoS Policy COMPUTATION

## Today: TPC between RSEs




## Tomorrow: TPC and Storage internal QoS transitions



## Conceptual decision:

- Separate RSEs conceptually into
- 1 RSE with multiple QoS zones





# Roadmap: Scalability improvements

---

- Reaper 2.0 already addressed scalability concerns for the deletion
  - Fully scalable due to dynamic partitioning algorithm
  - Already used in production
- Rule daemons are fully scalable, but limited in terms of rule size
  - Current rule processing is done atomically, which hits memory limits due to the way we deal with database objects
  - Benchmarked and profiled memory usage of daemon and identified several improvements
  - Full Scalability of rule daemons will be implemented within 2020



# Roadmap: Other improvements

---

- Stronger support for community policies
  - Path algorithms, schemas, policies
- More features for Subscriptions
  - Programmable (function-based) destination selection
  - More flexibility in expressing computing model/policy in subscriptions (See above)
- Monitoring
  - Some monitoring improvements to better track actions and workflows impacting data



# Python Support

---

- Clients
  - Python 2.6 dropped with Rucio **1.20 LTS** release line
  - Support py2.7, 3.6, 3.7, 3.8
  - Py2.7 EOL reached, but still need to support it (for now) due to massive use in our community
- Server (and daemons)
  - Support py2.7, 3.6, 3.7, (3.8)
  - As the server environment is fully under the operators control
  - **Proposal:** Drop py2.7 support with Rucio **1.24** release (November 2020)



# Documentation

- Partially discussed in Tuesdays Operations session
- More and more communities are moving Rucio into production (or are very close to)
- It is becoming very clear that we need to “invest” in user/operator facing documentation; Nick’s categories of documentation:

Howtos	Tutorials	Concepts	API
<ul style="list-style-type: none"><li>- Many questions &amp; answers on Slack → <b>Nothing is persisted</b></li><li>- <b>Operators recipes/manual</b> only in peoples head</li></ul>	<ul style="list-style-type: none"><li>- Significant improvements with dev containers</li><li>- Will add minikube “ops” tutorial which was made for the workshop</li><li>- <b>User tutorials</b> are not in the best state</li></ul>	<b>Mostly OK</b> (Also covered by publications)	<b>OK</b> (With slight improvements)



# Documentation

- Partially discussed in Tuesdays Operations session
- More and more communities are moving Rucio into production (or are very close to)
- It is becoming very clear that we need to “invest” in user/operator facing documentation; Nick’s categories of documentation:

Howtos	Tutorials	Concepts	API
<ul style="list-style-type: none"><li>- Many questions &amp; answers on Slack → <b>Nothing is persisted</b></li><li>- <b>Operators recipes/manual</b> only in peoples head</li></ul>	<ul style="list-style-type: none"><li>- Significant improvements with dev containers</li><li>- Will add minikube “ops” tutorial which was made for the workshop</li></ul>	<p>Mostly OK (Also covered by publications)</p>	<p>OK (With slight improvements)</p>

- Need to find the right tools!
  - Stackoverflow-like platform for questions?
  - Move all doc to github pages for easy contribution
- Need people to spend time on this

**Tutorials** are not in the best



# Development model

---

- Current development model requires two Pull Requests for any patch/enhancement
  - One PR against master and one against the next feature development branch
- Different workflow compared to most other open-source projects
- Difficult for newcomers to get used to
  - We have helper-tools to take care of it, but newcomers often just want to use the github web interface
- Will change to a different model, not requiring the two PRs
  - More work for the release manager, but more comfortable over all
- ETA: Probably Q2 2020, possible with 1.23 LTS release



# Databases

---

- We know Rucio scales well to Multi-Billion rows sized databases running on Oracle
- Little experience how it scales with MySQL/MariaDB and PostgreSQL
  - Volunteers for evaluation?
- Some concerns due to single-point-of-failure of the DB
  - CERN experience is that our database operates very well and very stable
    - Downtimes are extremely rare
    - Internal failover with instances inside cluster
  - If stronger redundancy/failover is needed, this should be done at a database/infrastructure level
    - Very hard to include in the business logic of Rucio
    - Usually very dependent on the specific infrastructure setup of the community



**Other topics?**



# Please contact us!

---

- Unfortunately we didn't have the Tutorial, Meet the Team and Coding Camp sessions for individual discussions
- Please join our [Slack](#) for discussions or contact [me](#) to schedule individual discussions



# More information

---

Website



<http://rucio.cern.ch>

Documentation



<https://rucio.readthedocs.io>

Repository



<https://github.com/rucio/>

Images



<https://hub.docker.com/r/rucio/>

Online support



<https://rucio.slack.com/messages/#support/>

Developer contact



[rucio-dev@cern.ch](mailto:rucio-dev@cern.ch)

Publications



<https://rucio.cern.ch/publications.html>

Twitter



<https://twitter.com/RucioData>