DE LA RECHERCHE À L'INDUSTRIE

RobinHood-v4 progress report

Commissariat à l'énergie atomique et aux énergies alternatives – www.cea.fr

What is RobinHood?

RobinHood mirrors a filesystem’s metadata, and makes it queryable

- Mirror
  - Read Replica
  - Snapshot
  - Disaster Recovery
- Query
  - Expressiveness
  - Performance
  - Efficiency

Architecture

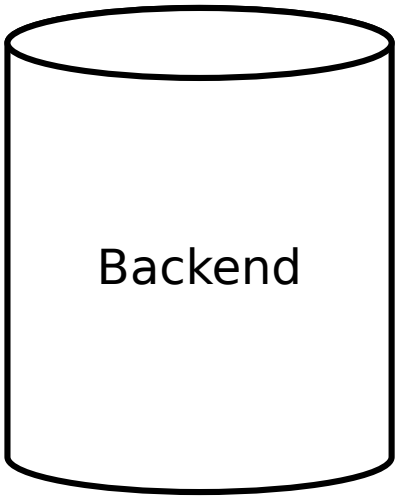

Architecture

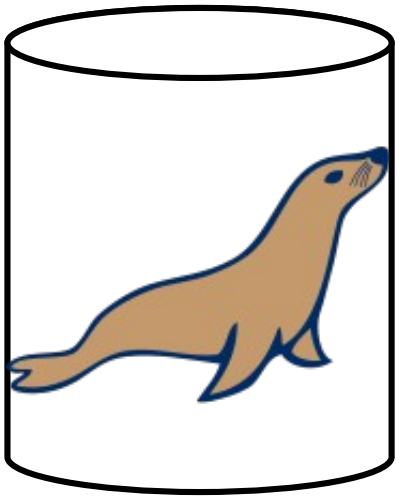

Architecture

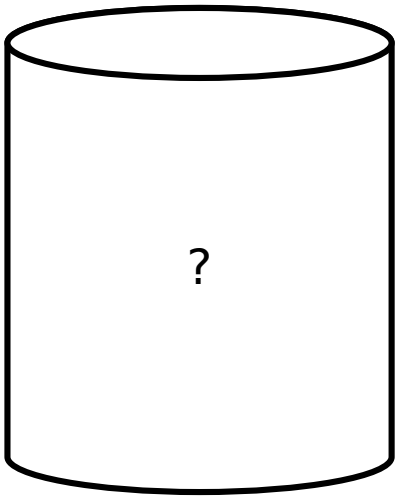

Architecture

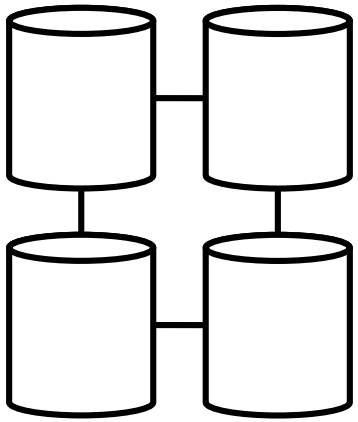

Architecture

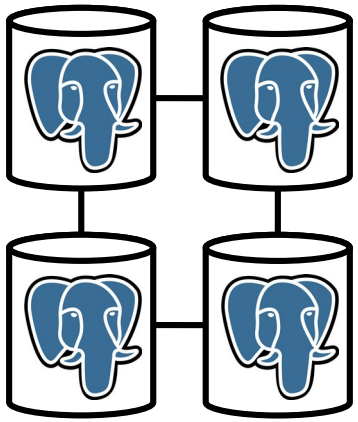

Architecture

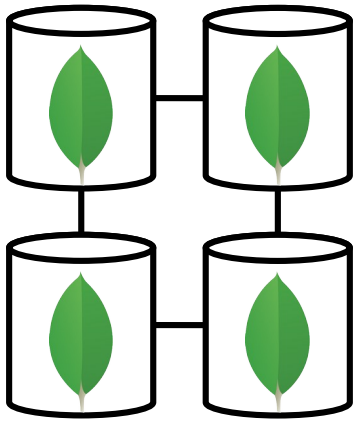

Architecture

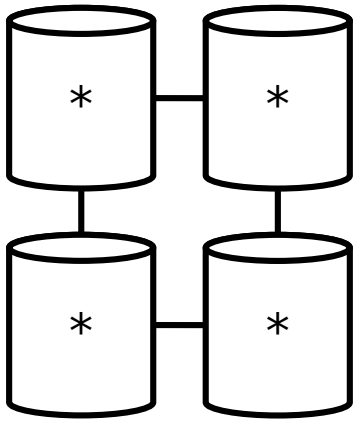

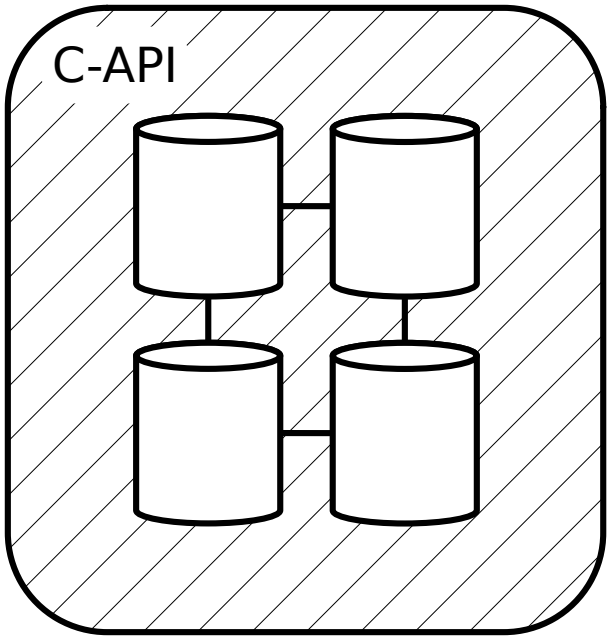
C-API

Architecture

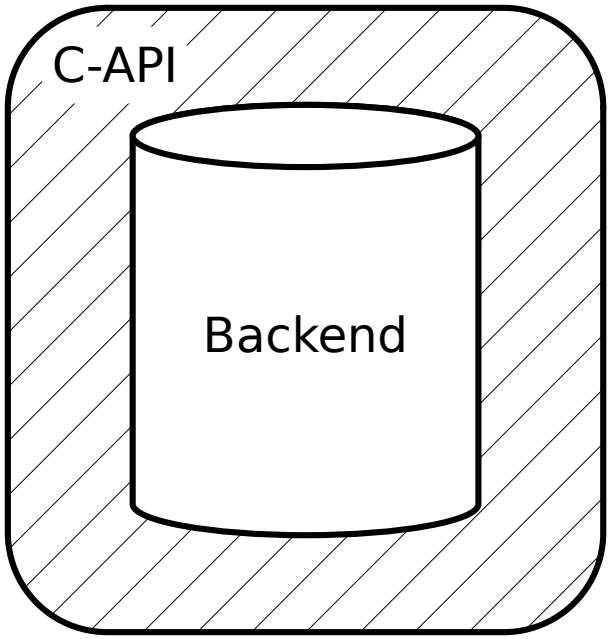

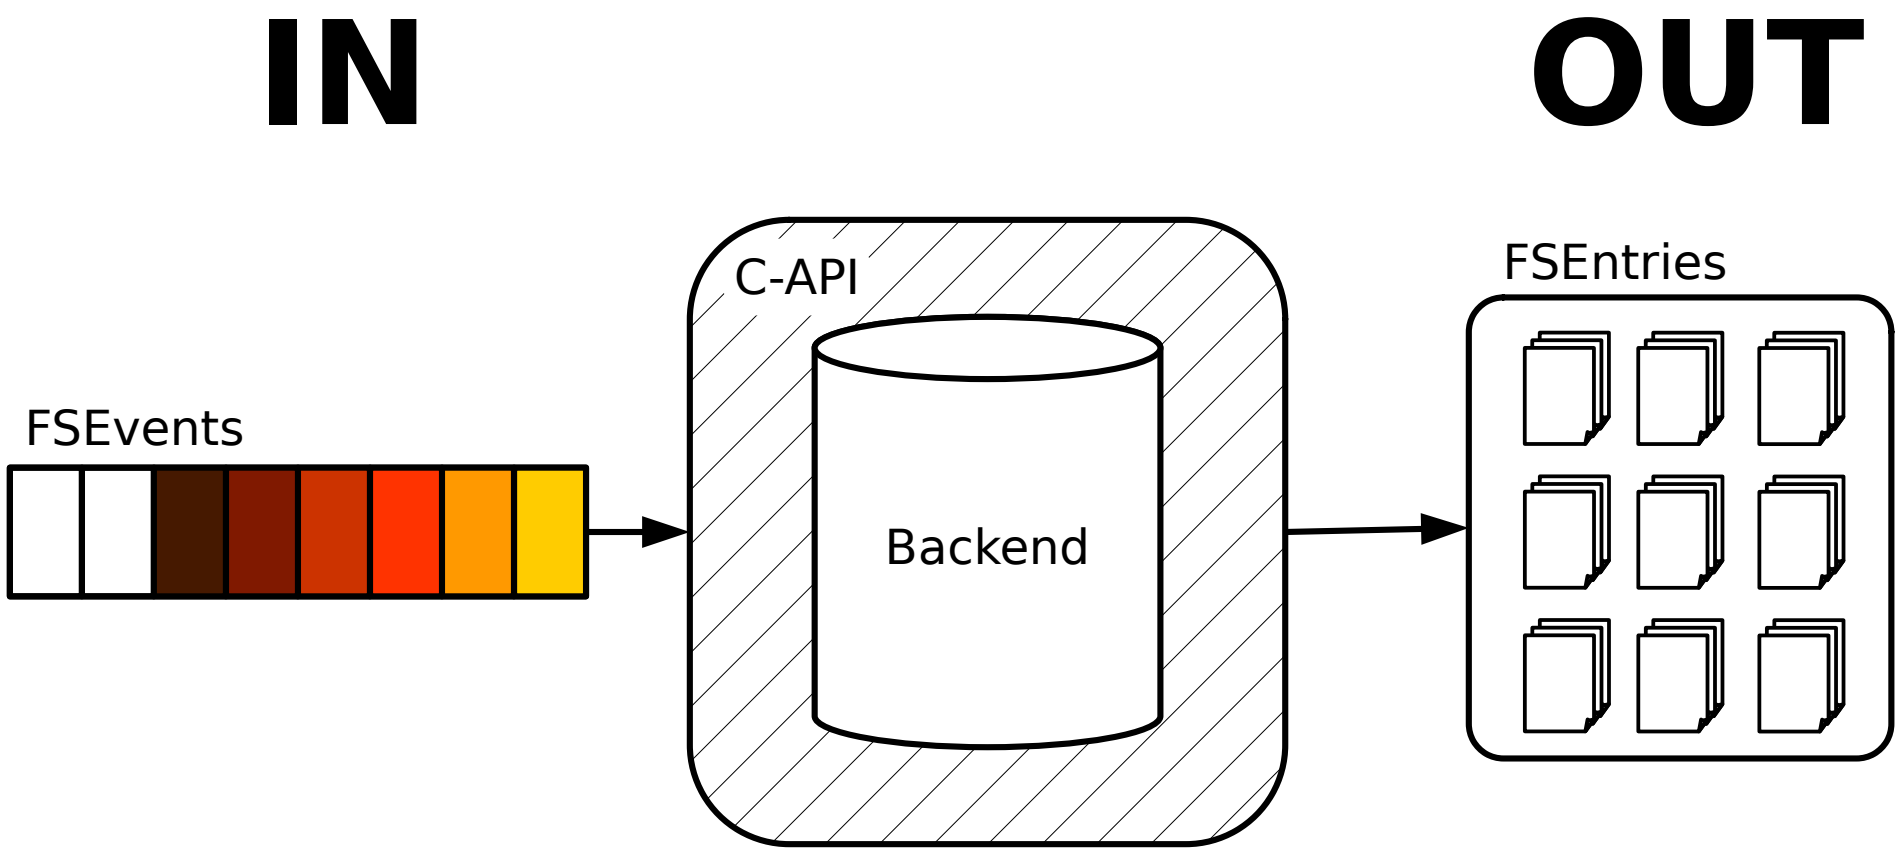
IN
OUT
C-API
FSEntries
FSEvents
Backend

Architecture

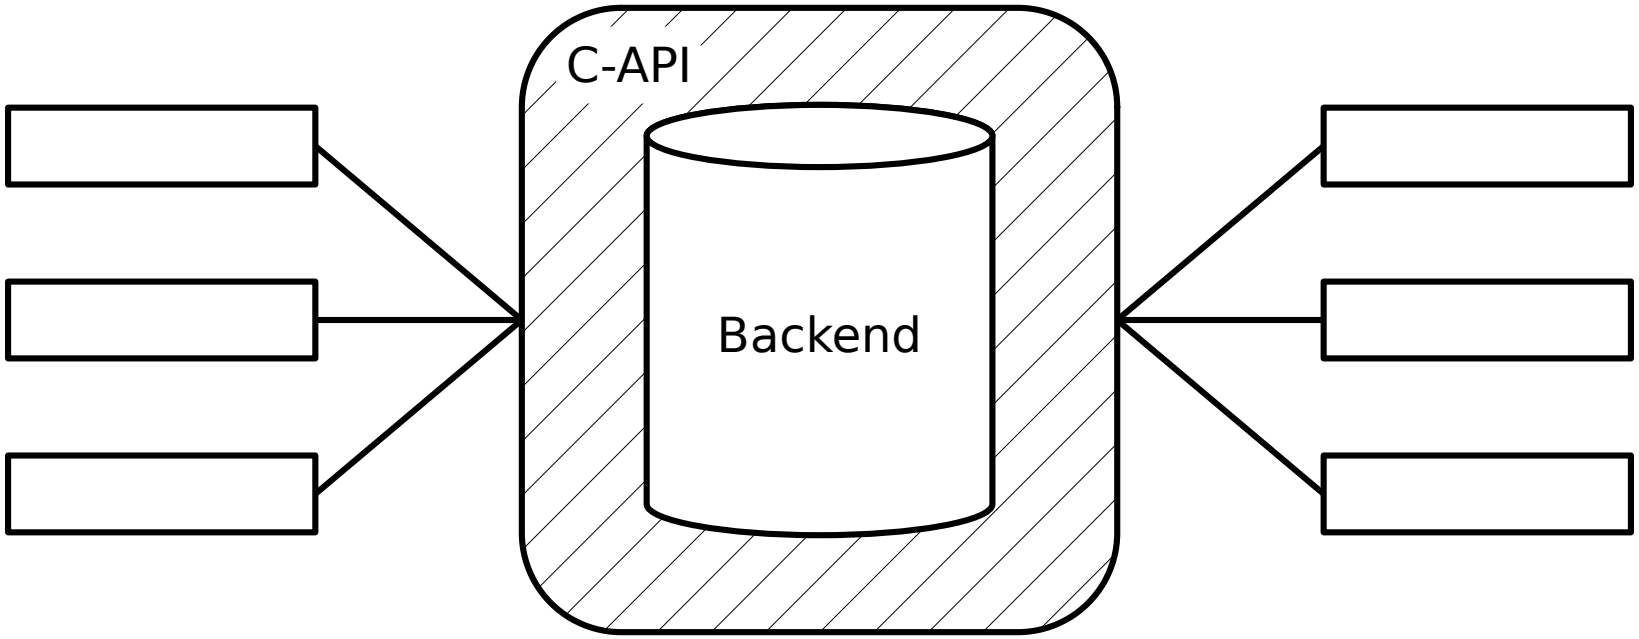

Applications

rbh-sync

rbh-sync synchronizes two RobinHood backends

rbh-sync: usage

rbh-sync SRC DEST

rbh-sync: example

```
rbh-sync rbh:posix:/mnt/scratch rbh:mongo:scratch
```

rbh-sync: URIs

- Uniform Resource Identifier
- RFC 3986:

  scheme://[userinfo@]host[:port]]path

rbh-sync: URIs

https://example.com/test

rbh-sync: URIs

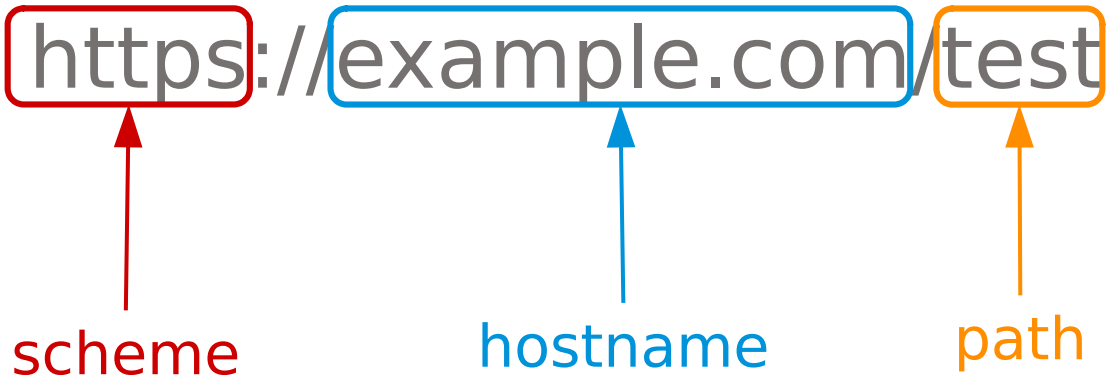

rbh-sync: URIs

rbh:mongo:scratch

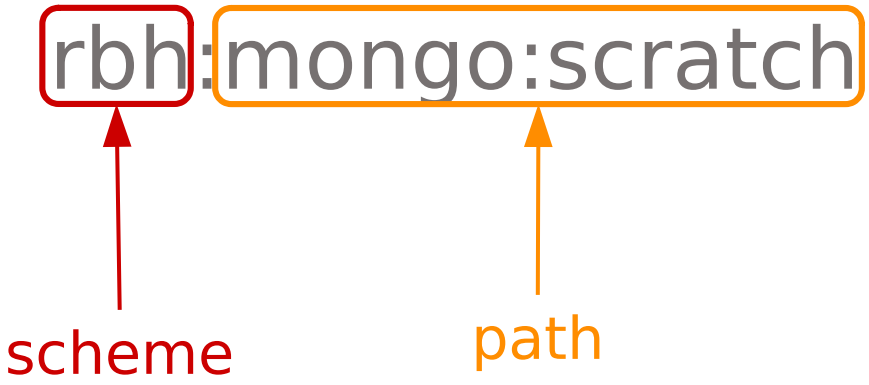
rbh:mongo:scratch
scheme
path

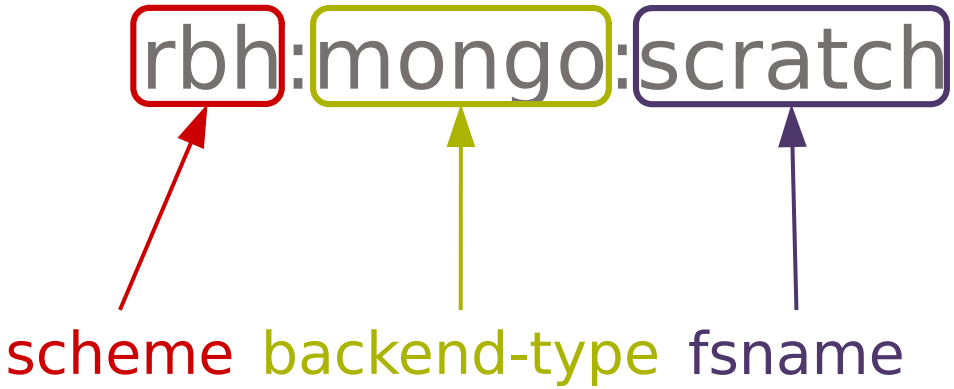
rbh:mongo:scratch
scheme
backend-type
fsname

rbh-sync: URIs

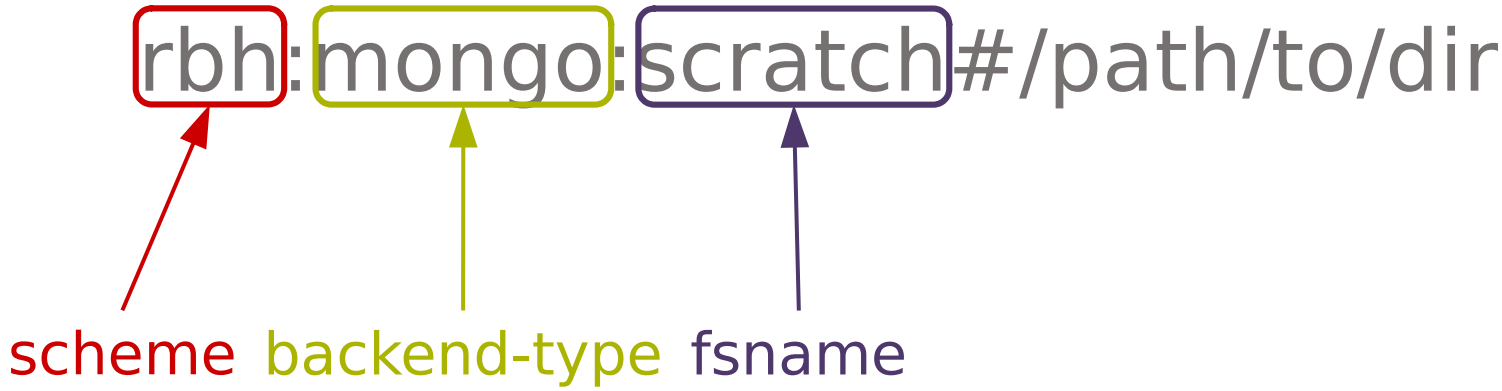

rbh-sync: URIs

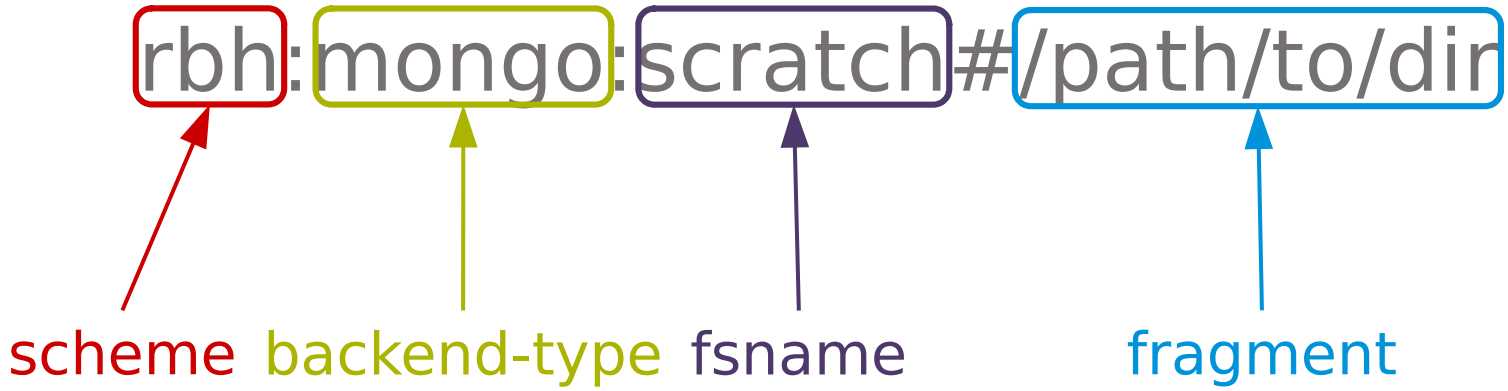

rbh-sync: URIs

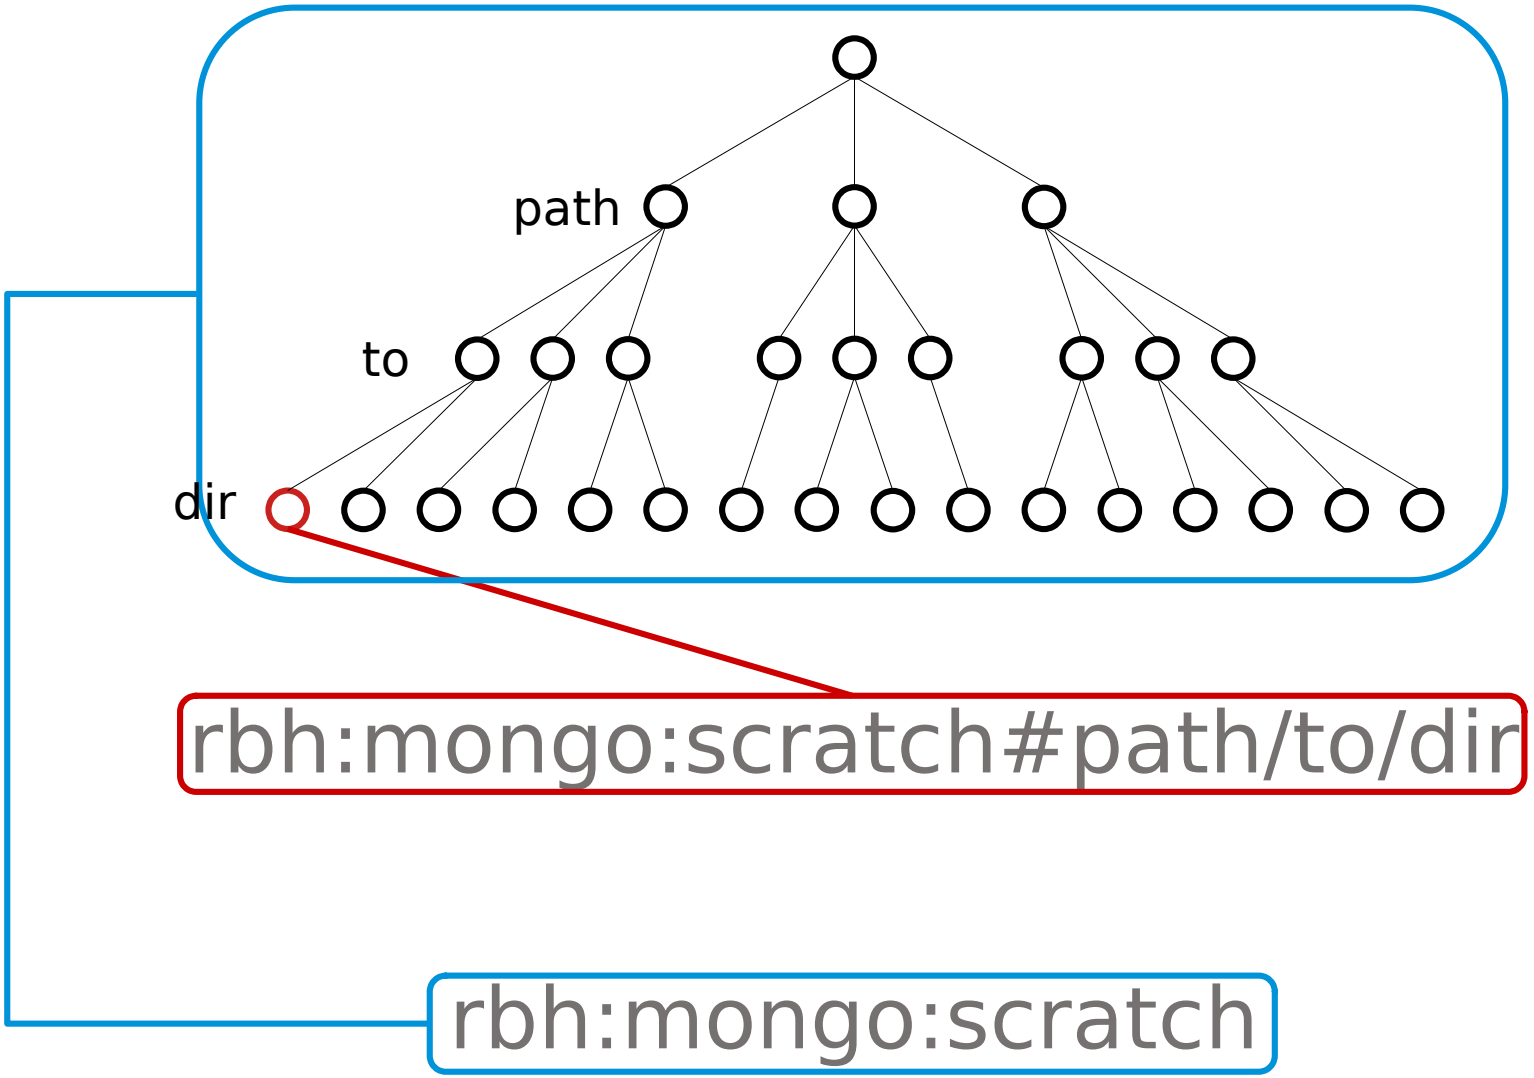

rbh-sync: another example

```
rbh-sync rbh:posix:/mnt/scratch#dir-0 \
         rbh:mongo:scratch
```

rbh-sync: another example

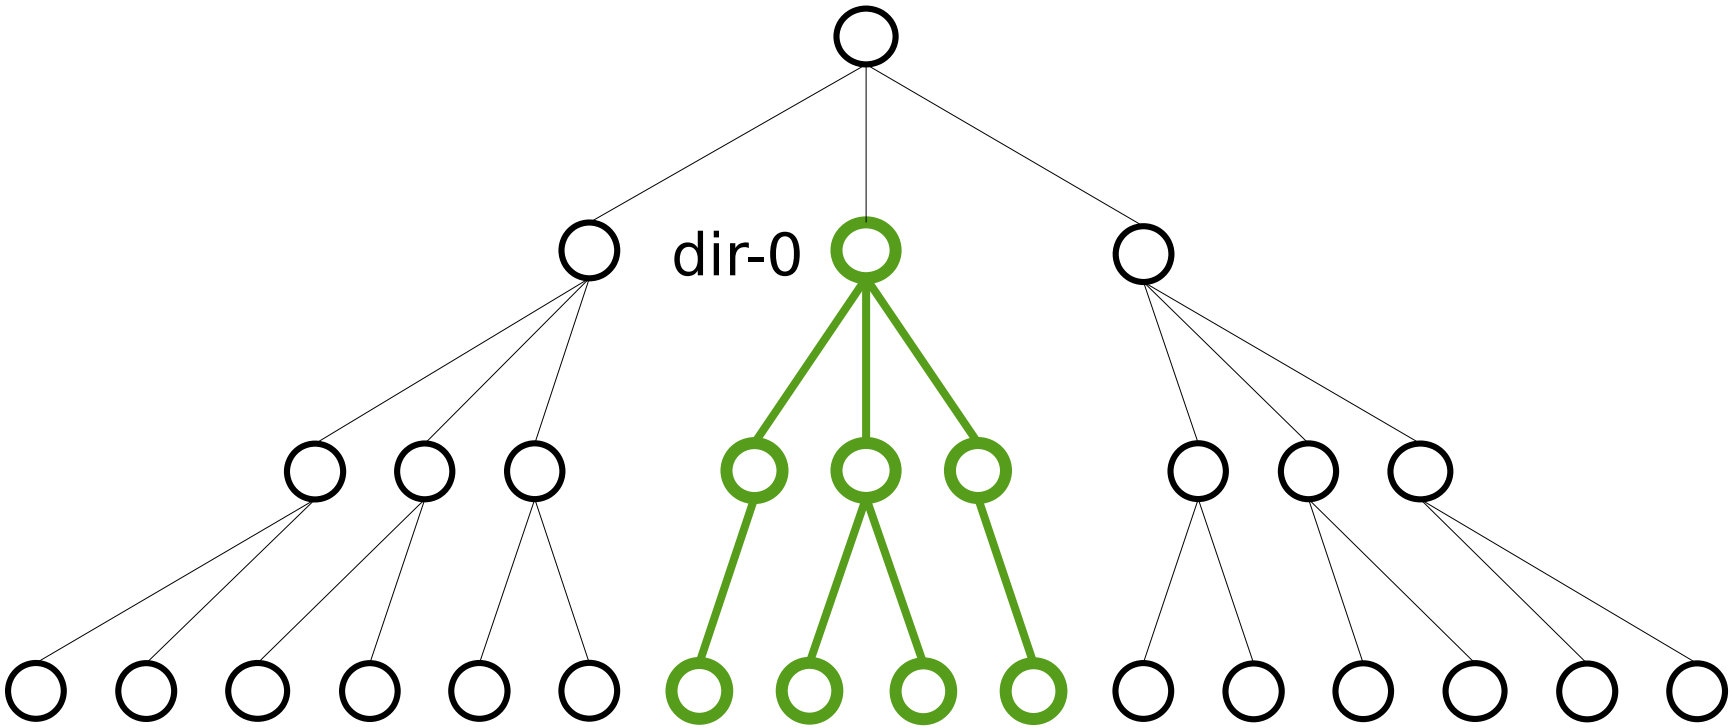

rbh-sync: parallelism

```
cd /mnt/scratch
for dir in *; do
   rbh-sync posix:/mnt/scratch#”$dir” mongo:scratch &
done
wait
```

rbh-sync: another example

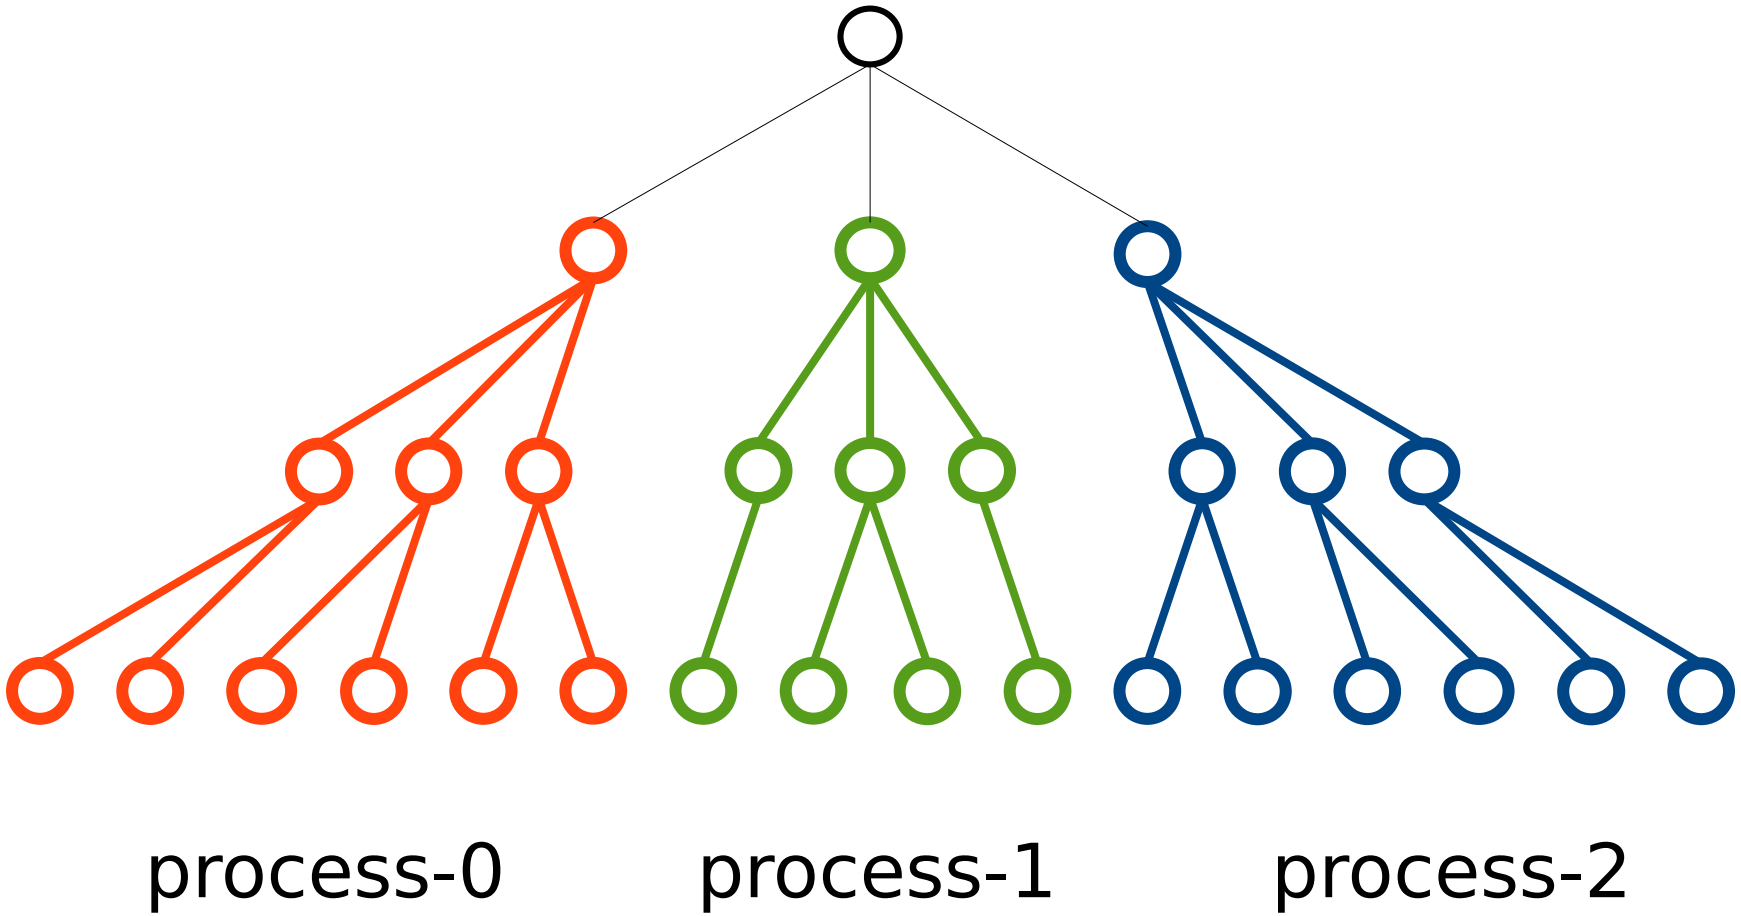

rbh-sync: performance

- 1 node:
  - 16 cores
- Lustre on RAM:
  - 1 MDT
  - 8 OSTs
  - Cold cache
- Filetree structure:
  - 4 levels of dirs
  - Branching factor of 16
  - 5th level: ~1M files of 1KiB each

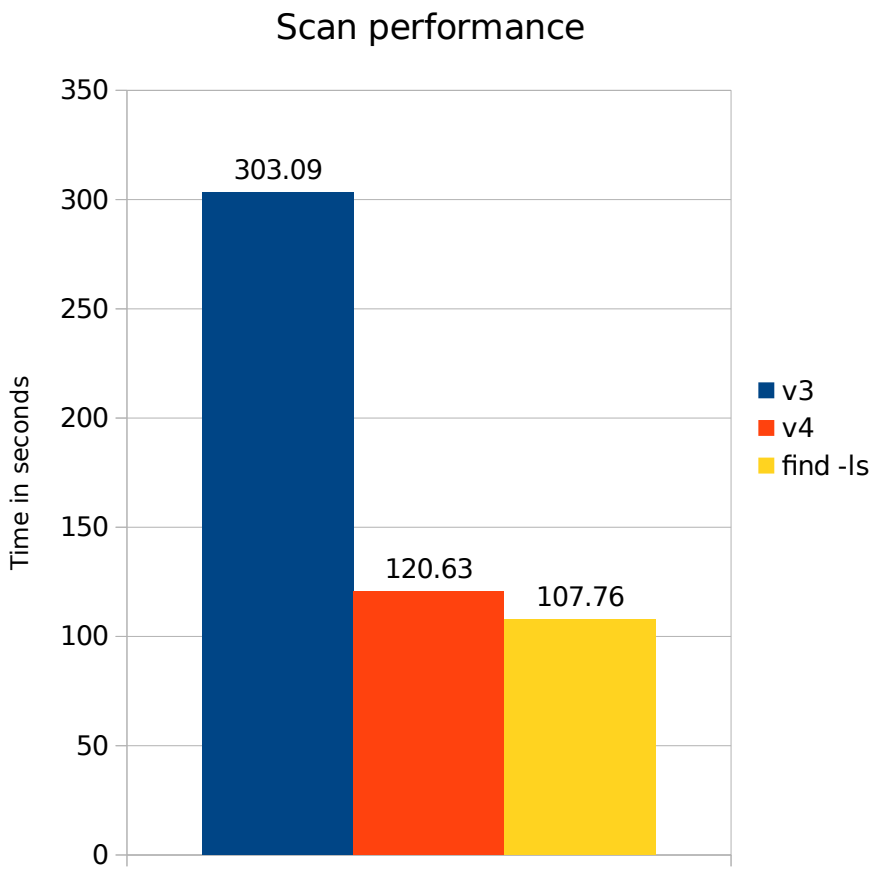

rbh-find

A clone of (gnu-)find that queries RobinHood

rbh-find: usage

rbh-find [URI ...] [[predicates] [actions] ...]

rbh-find: examples

```
rbh-find rbh:mongo:scratch -name ‘a*’ -print0
```

rbh-find: performance

- 1 node:
  - 16 cores
- Lustre on RAM:
  - 1 MDT
  - 8 OSTs
  - Cold cache
- Filetree structure:
  - 4 levels of dirs
  - Branching factor of 16
  - 5th level: ~1M files of 1KiB each

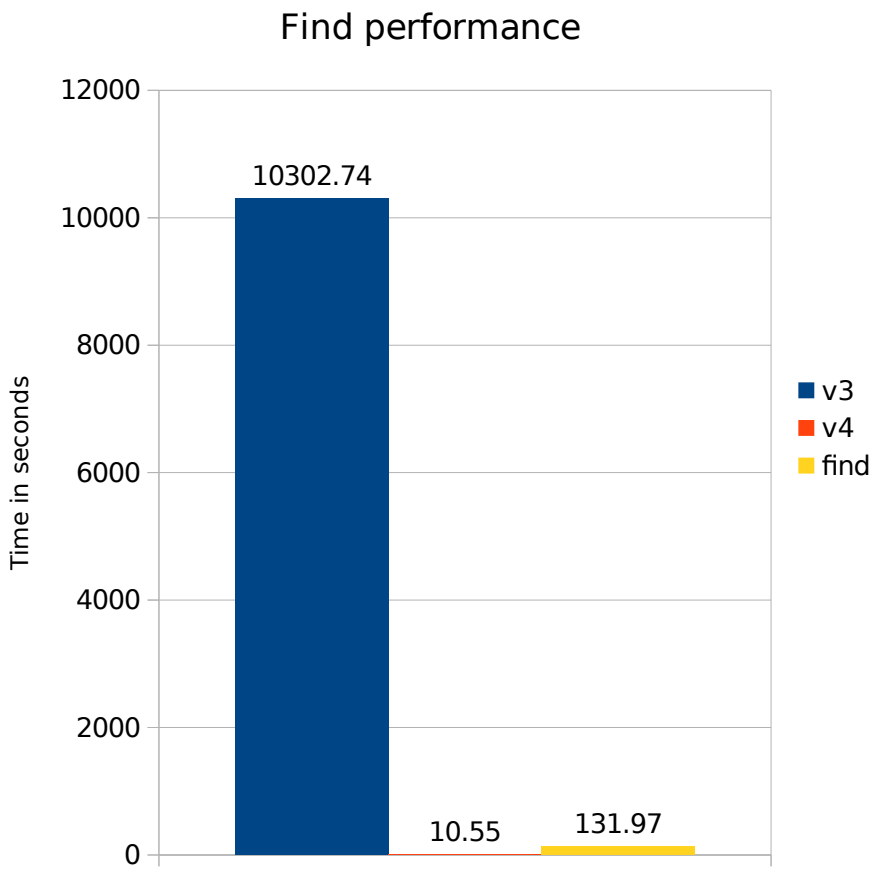

rbh-find: performance

- 1 node:
  - 16 cores
- Lustre on RAM:
  - 1 MDT
  - 8 OSTs
  - Cold cache
- Filetree structure:
  - 4 levels of dirs
  - Branching factor of 16
  - 5th level: ~1M files of 1KiB each

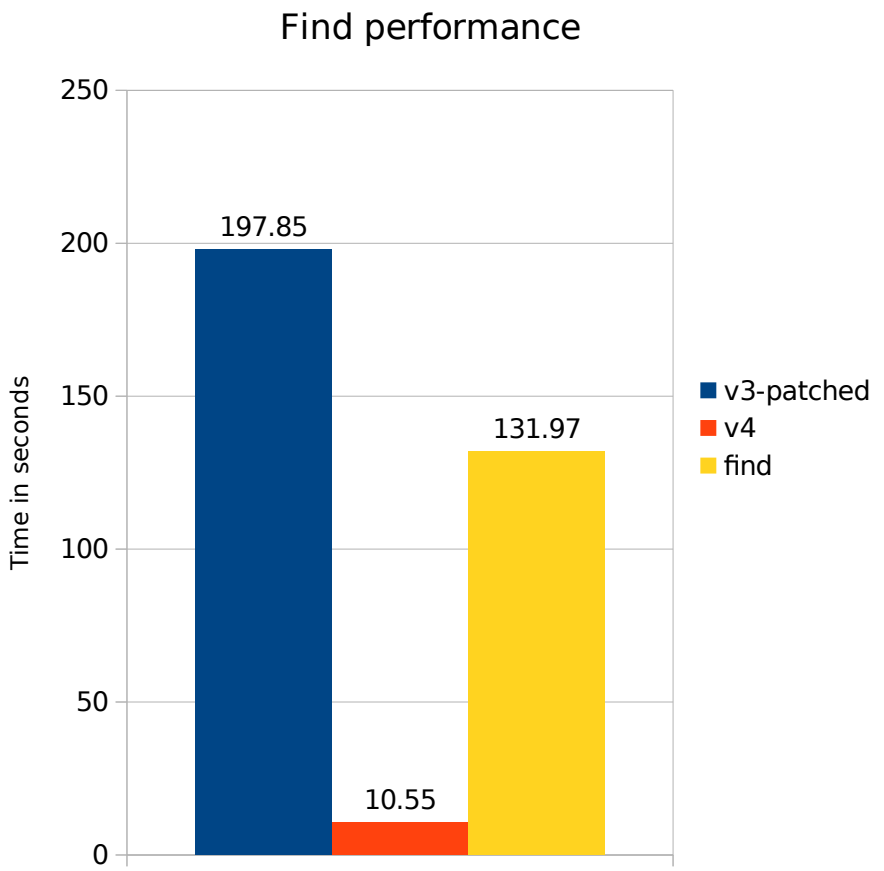

rbh-fsevent

Apply filesystem events on a RobinHood backend

rbh-fsevent: usage

TODO

Changelog processing

Before

Before

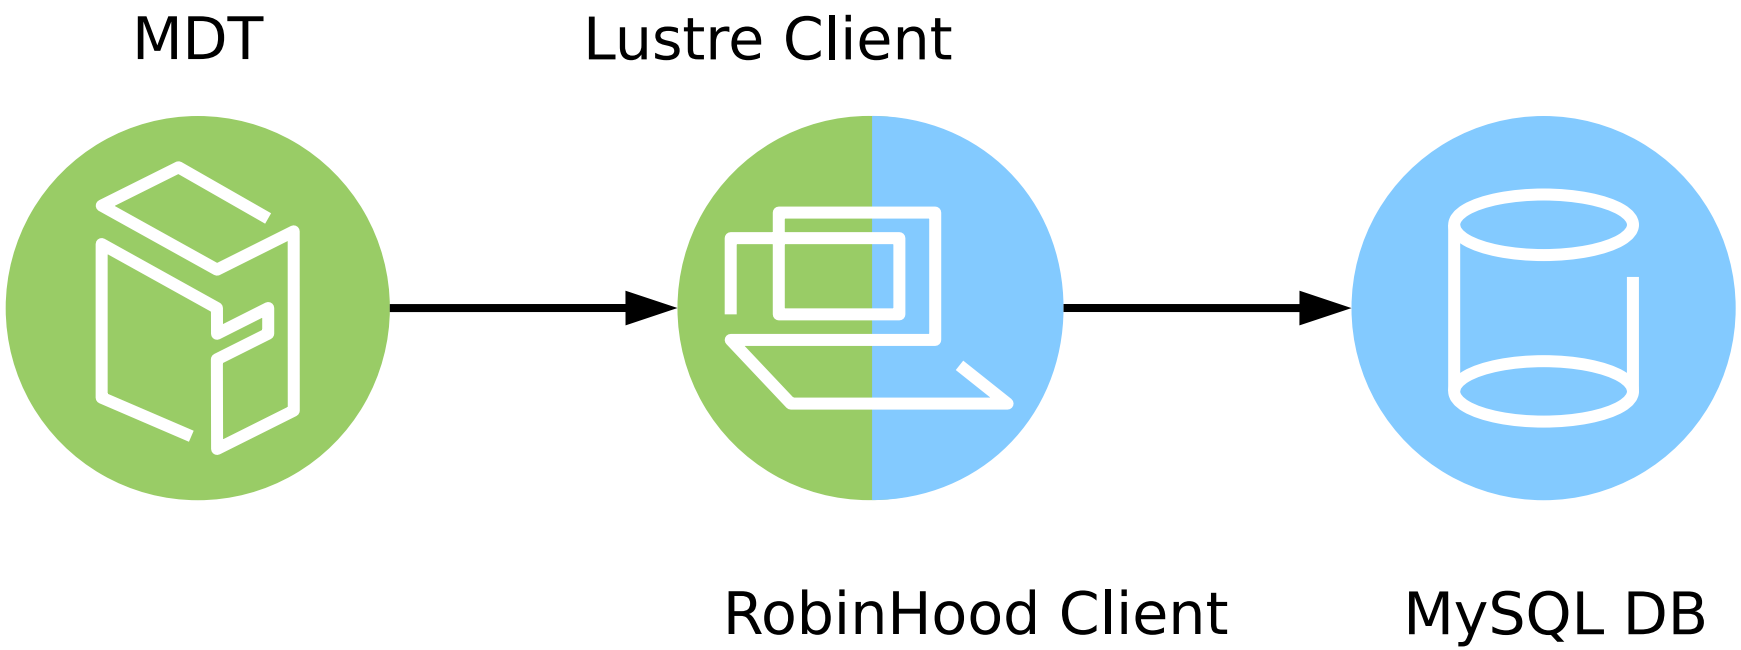

Before: overload

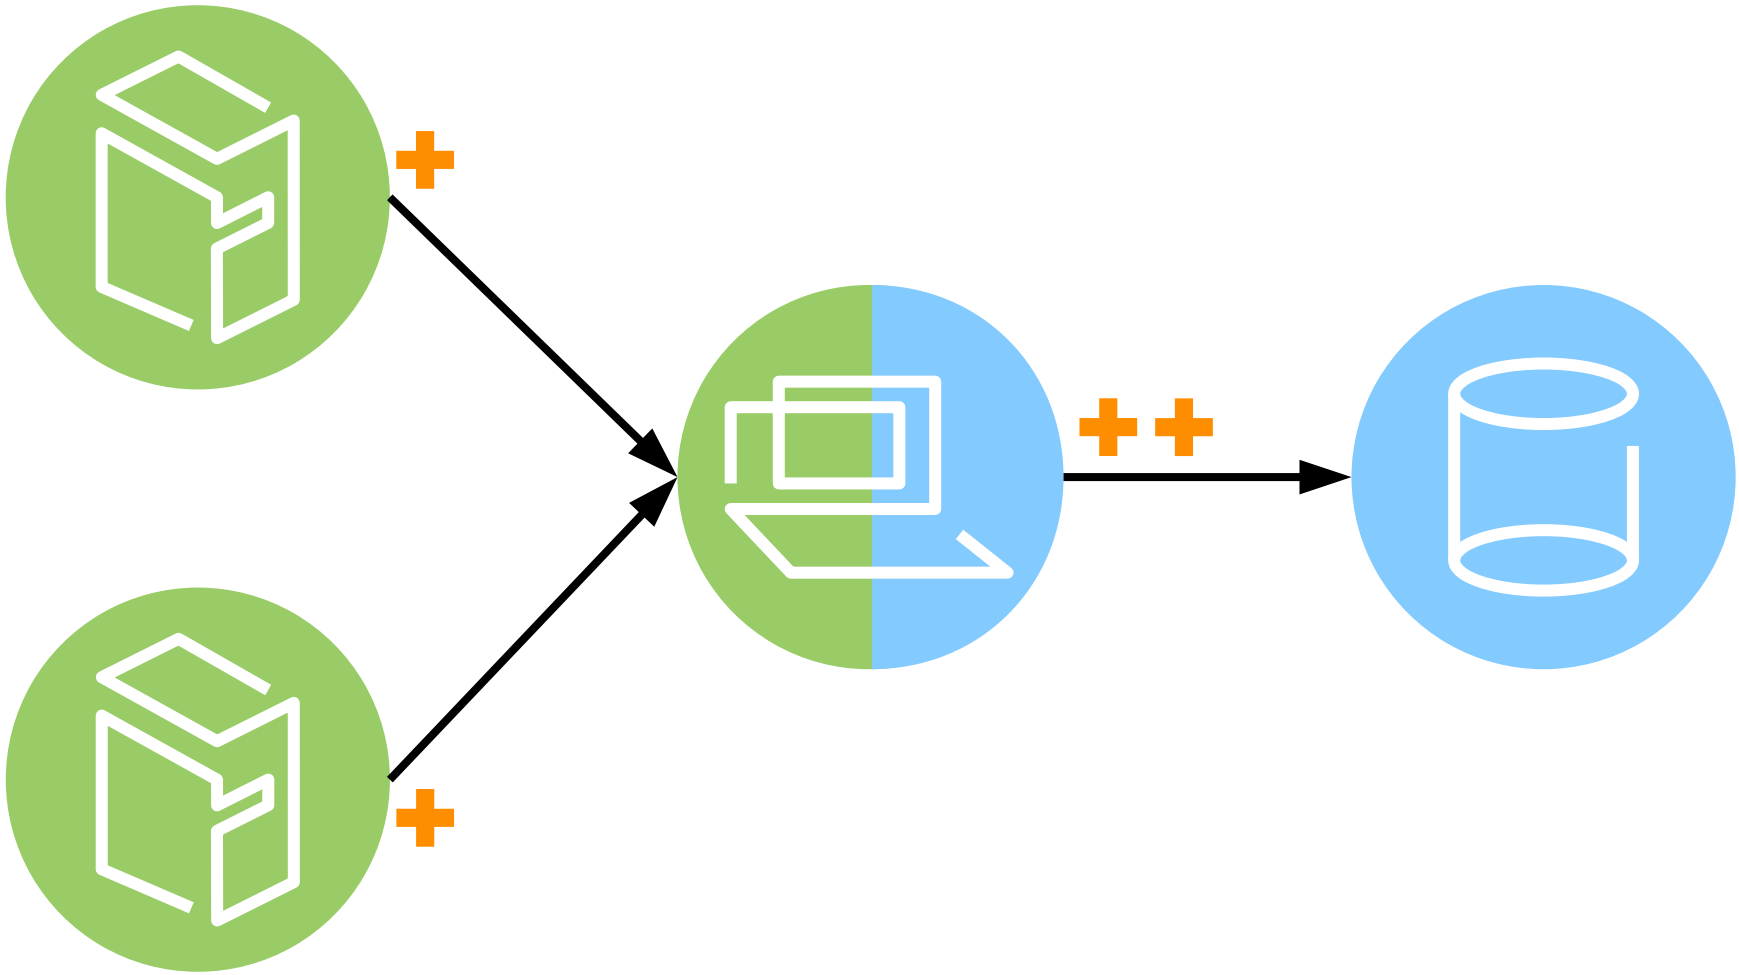

Before: imbalance

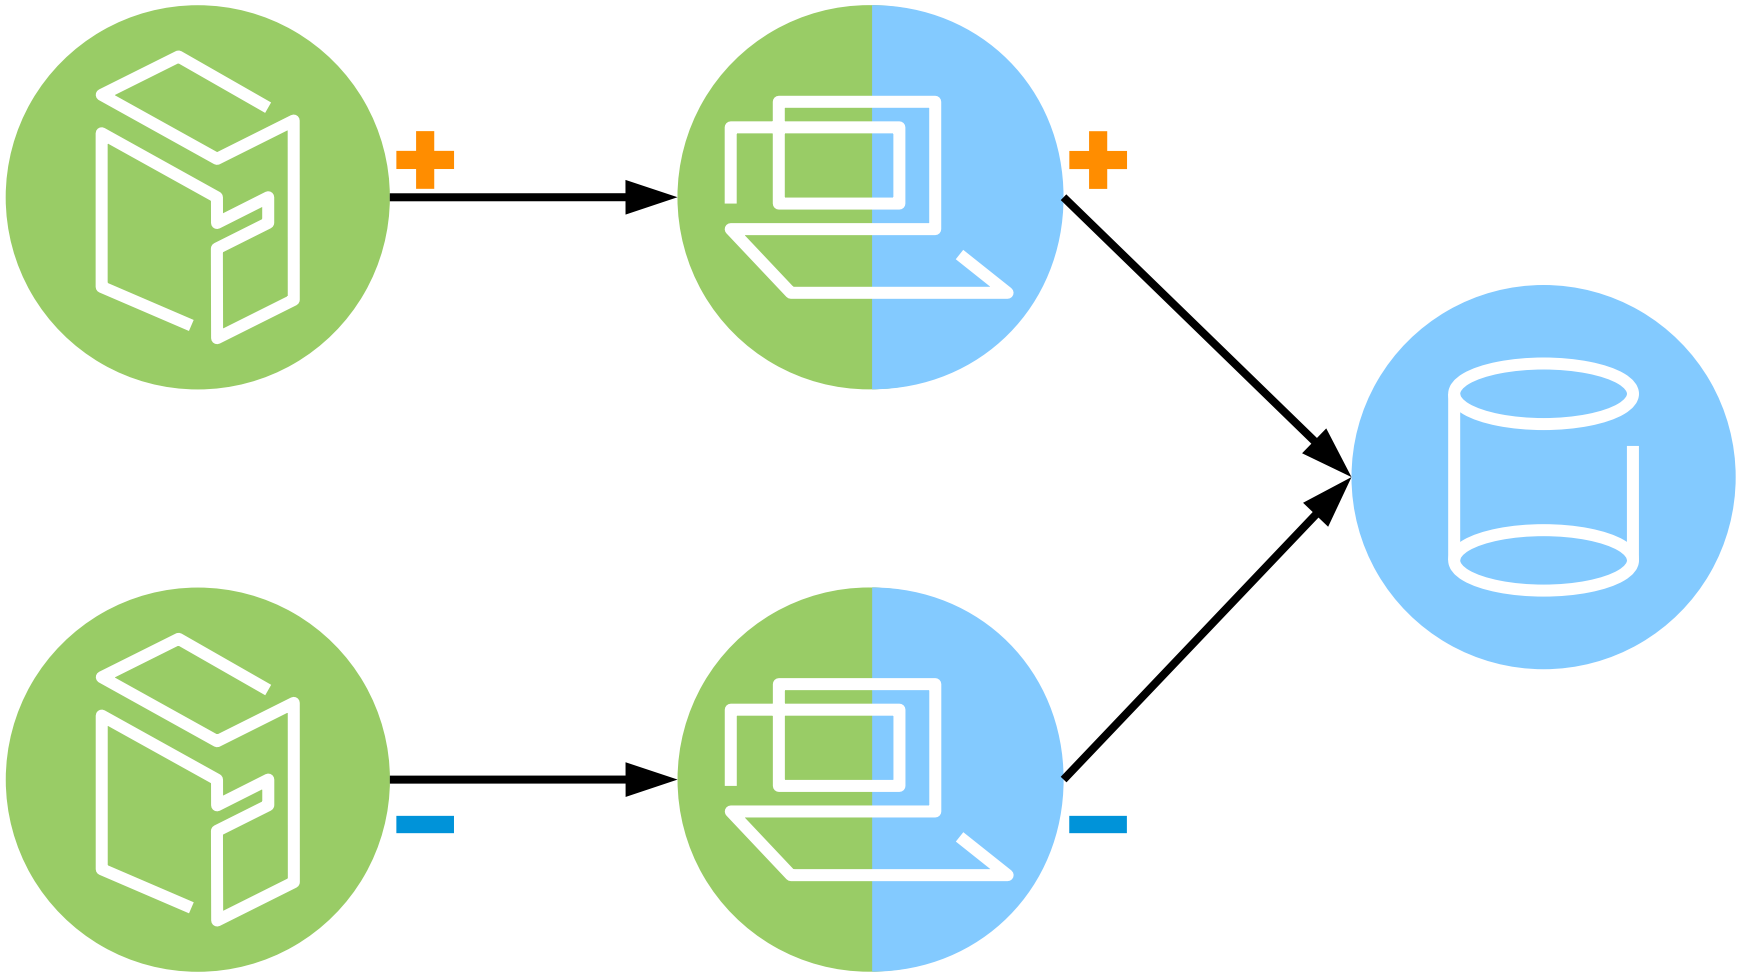

After

After: with kafka

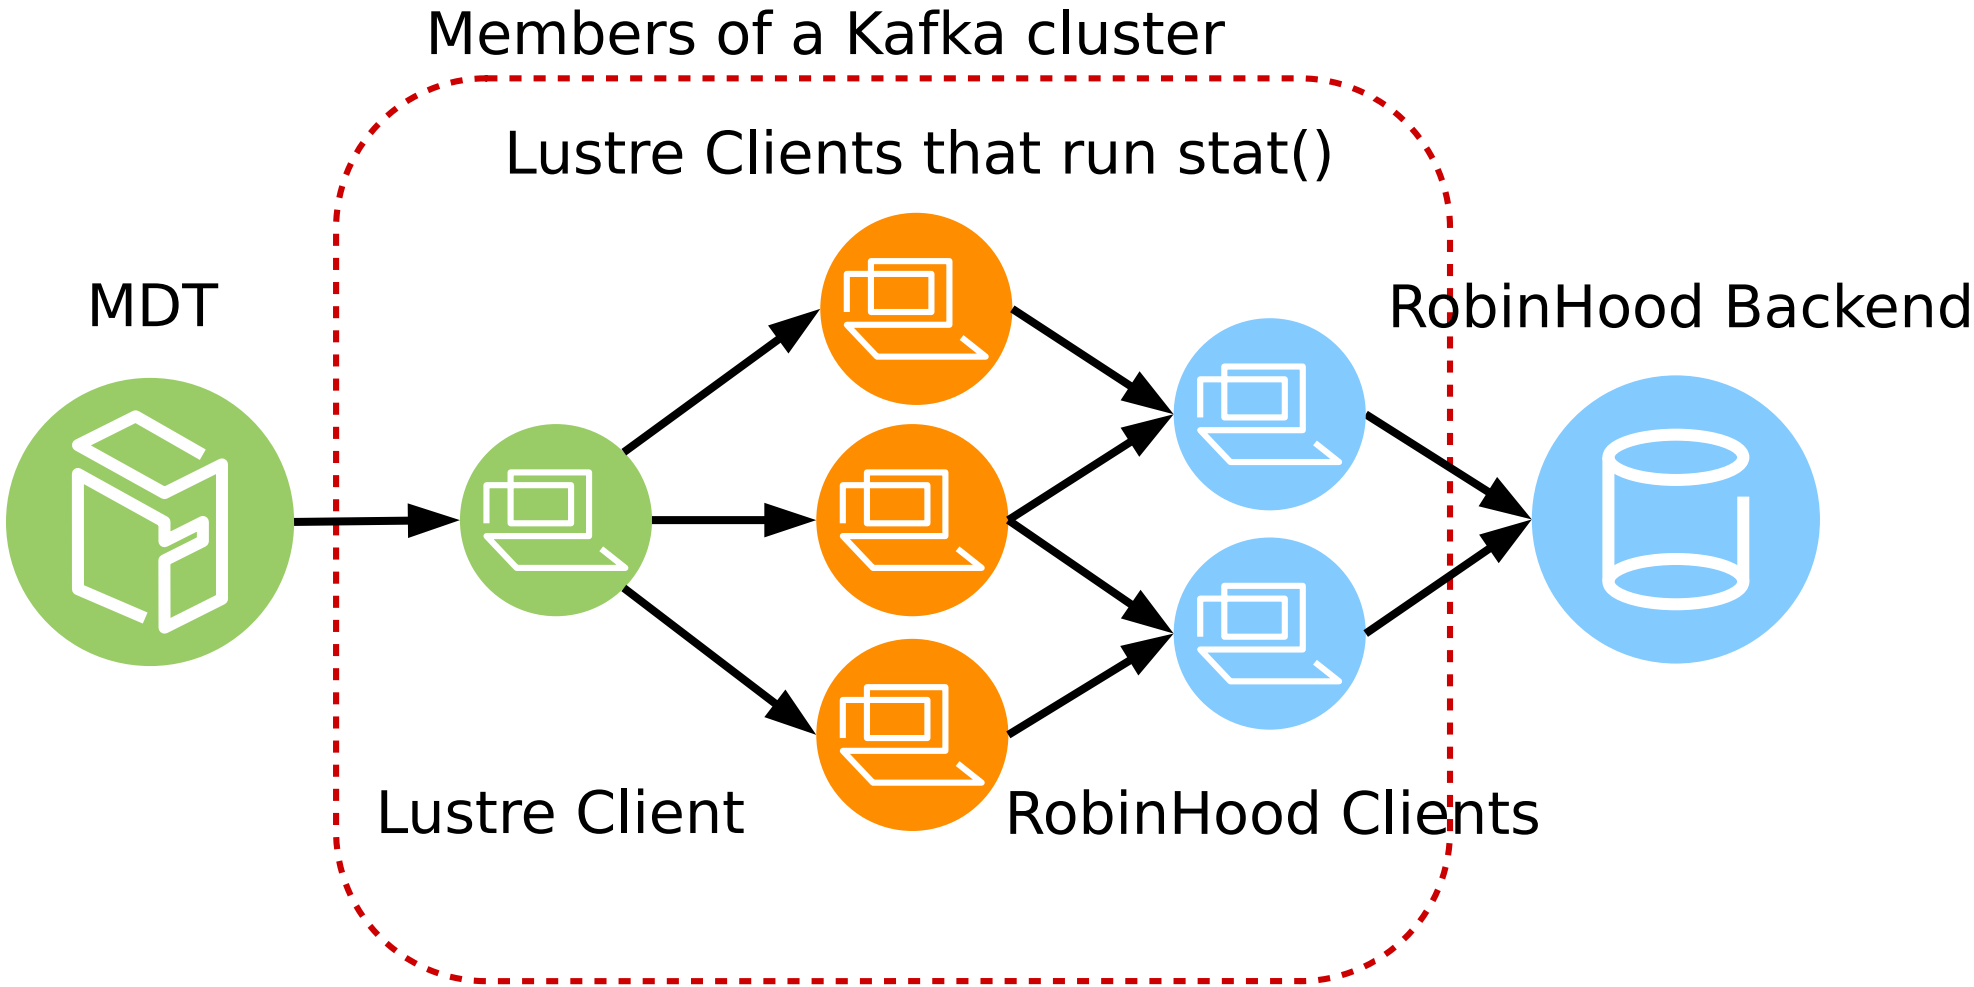

After: with kafka

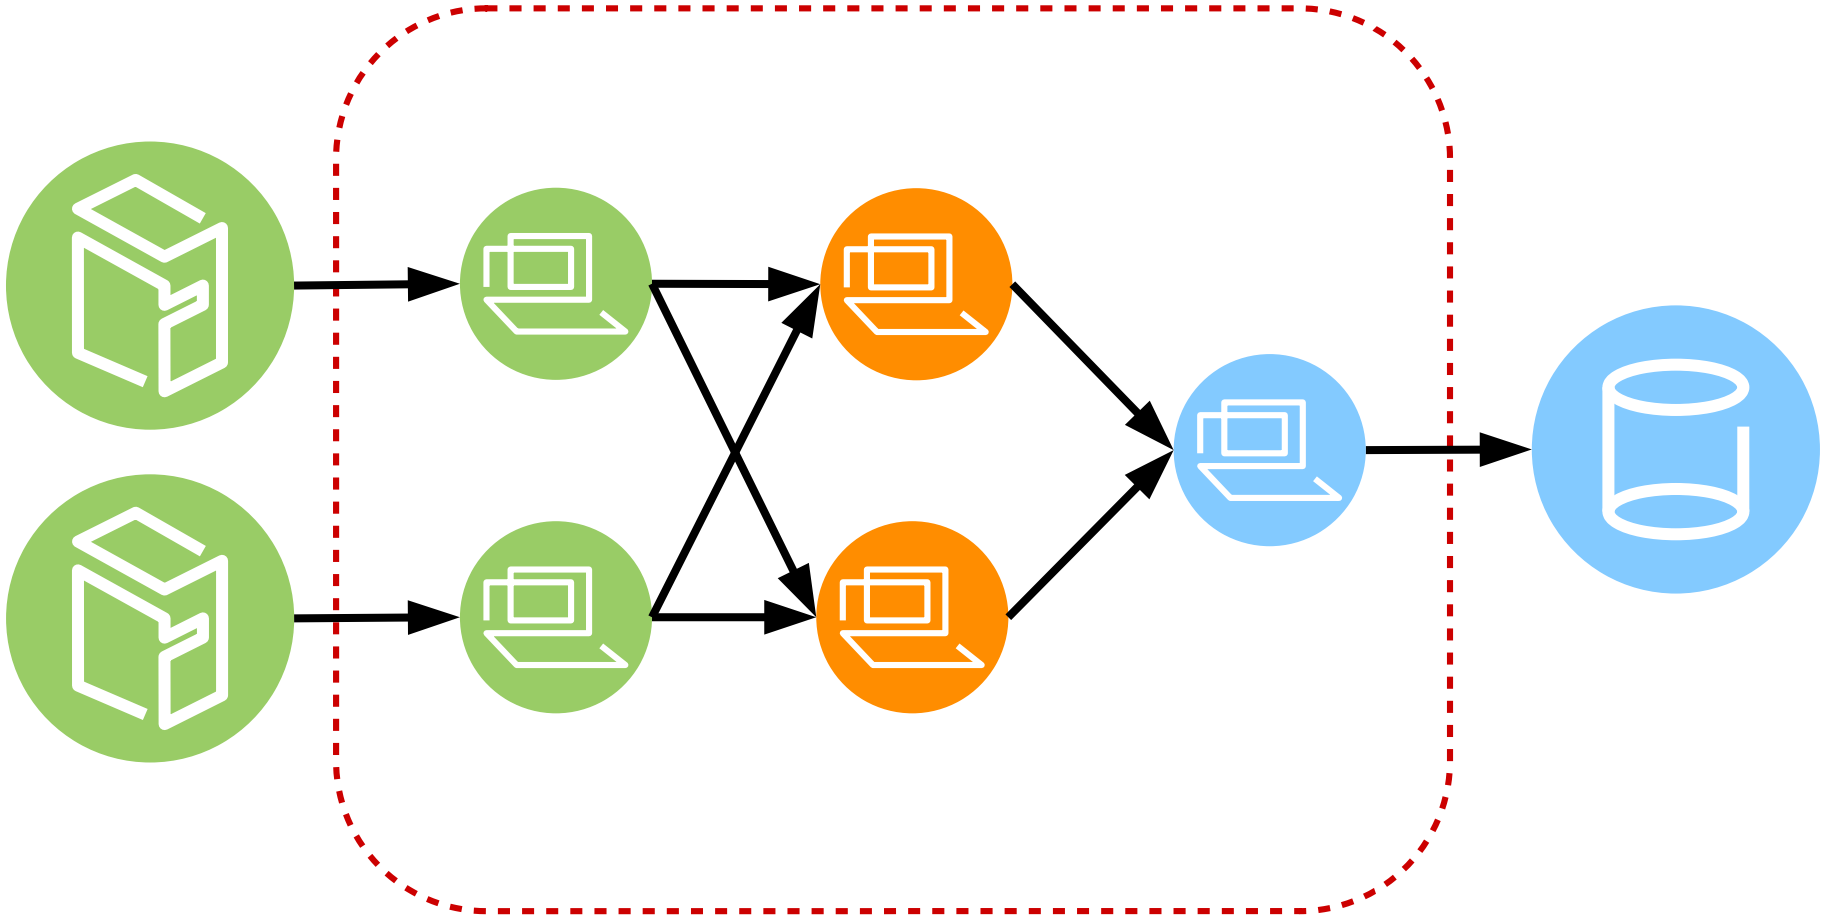

After: with kafka

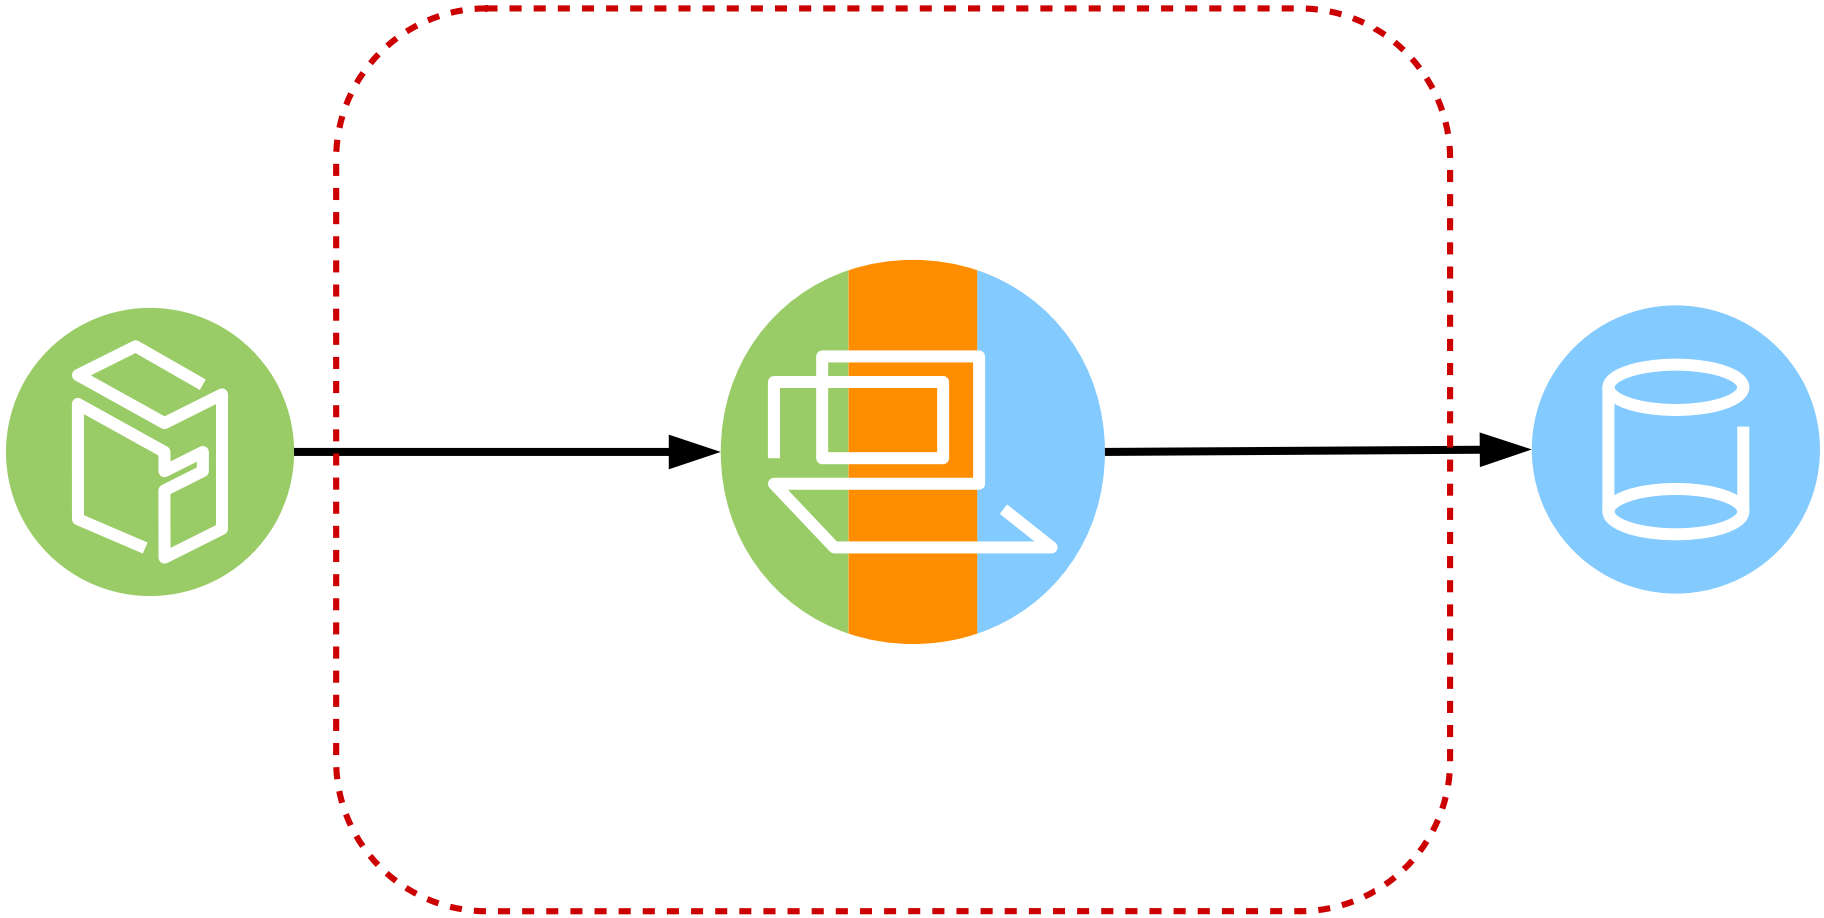

Thank you