DE LA RECHERCHE À L'INDUSTRIE

RobinHood Policy Engine

30th of May, 2023

Yoann VALERI, yoann.valeri@cea.fr

Context – Hot & Cold

Information about data temperature and storage technologies hierarchy

- Why is there hot and cold data?

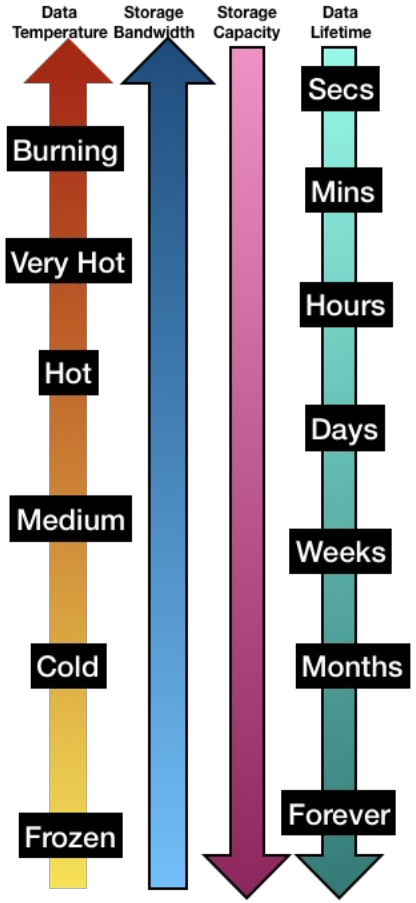

- Storage hierarchy diagram

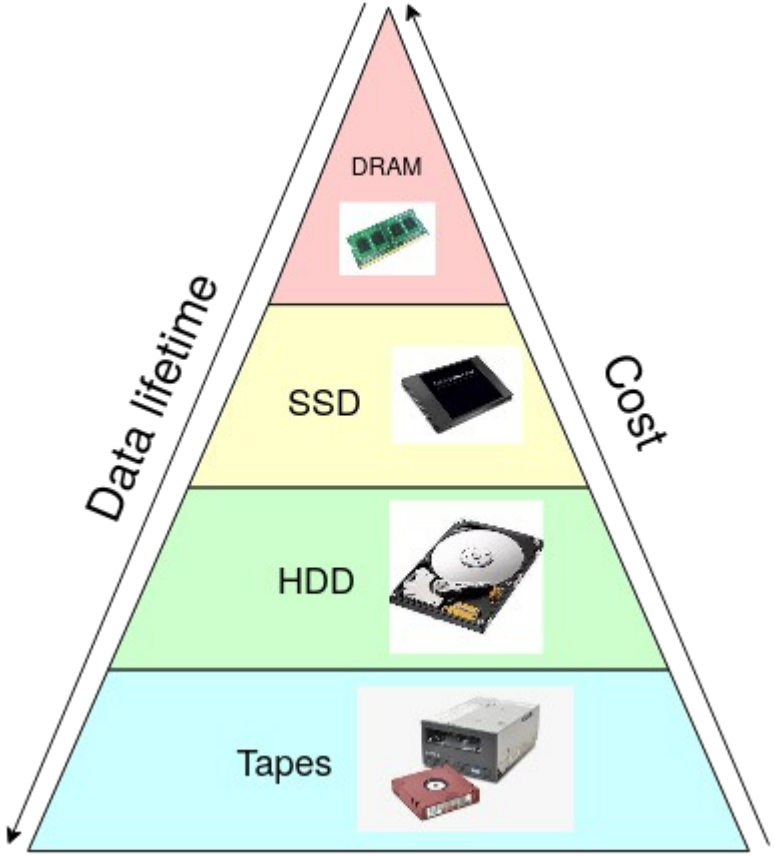

Context – Lustre HSM system

- The prototype installed at TGCC will use a Lustre/HSM setup with Phobos to manage long-term storage

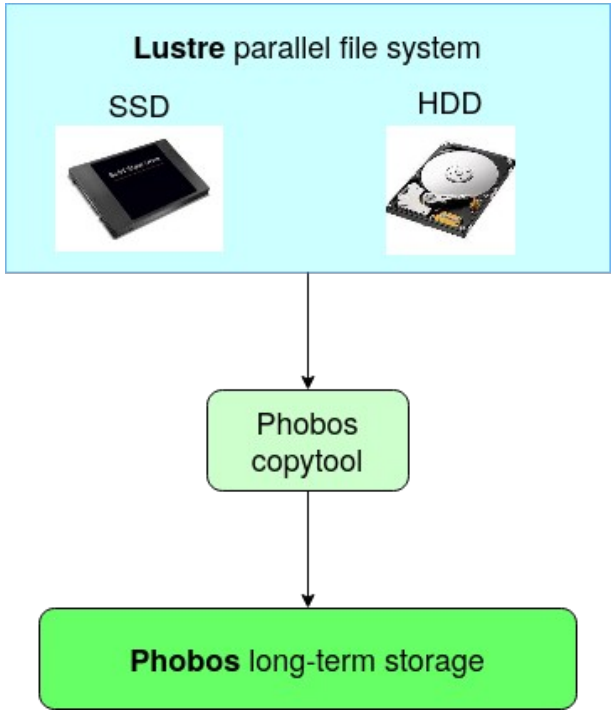

- Need for data policy migration engine → for instance : “migrate all files that haven’t been accessed in a year”
- Automate migration between tiers

Context – Lustre HSM system

- The prototype installed at TGCC will use a Lustre/HSM setup with Phobos to manage long-term storage

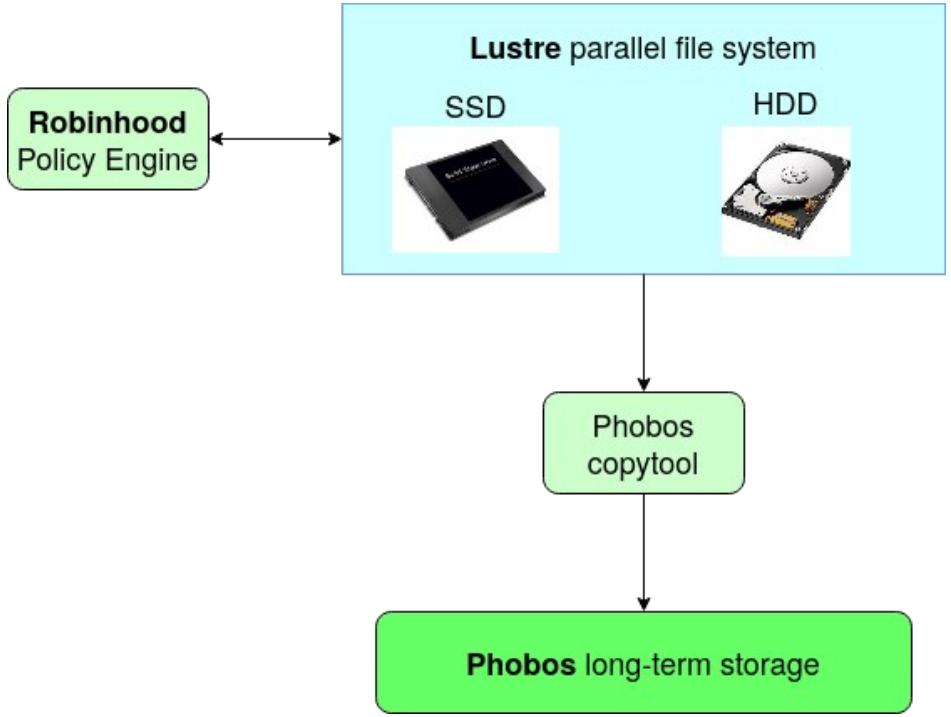

- Need for data policy migration engine → for instance : “migrate all files that haven’t been accessed in a year”
- Automate migration between tiers → Robinhood

Metadata fetching

To find the right files to migrate, we need to examine all the file system metadata

- **First solution: stat them all?**
  - **Filesystem traversal is too time consuming to do regularly**
  - **Impose a heavier load on the filesystem**
- **Our solution: mirror the metadata in a database**
  - **Iterate through each file in the filesystem**
  - **Copy the metadata of each file to the mirror**
  - **Update entries without scanning the filesytem**

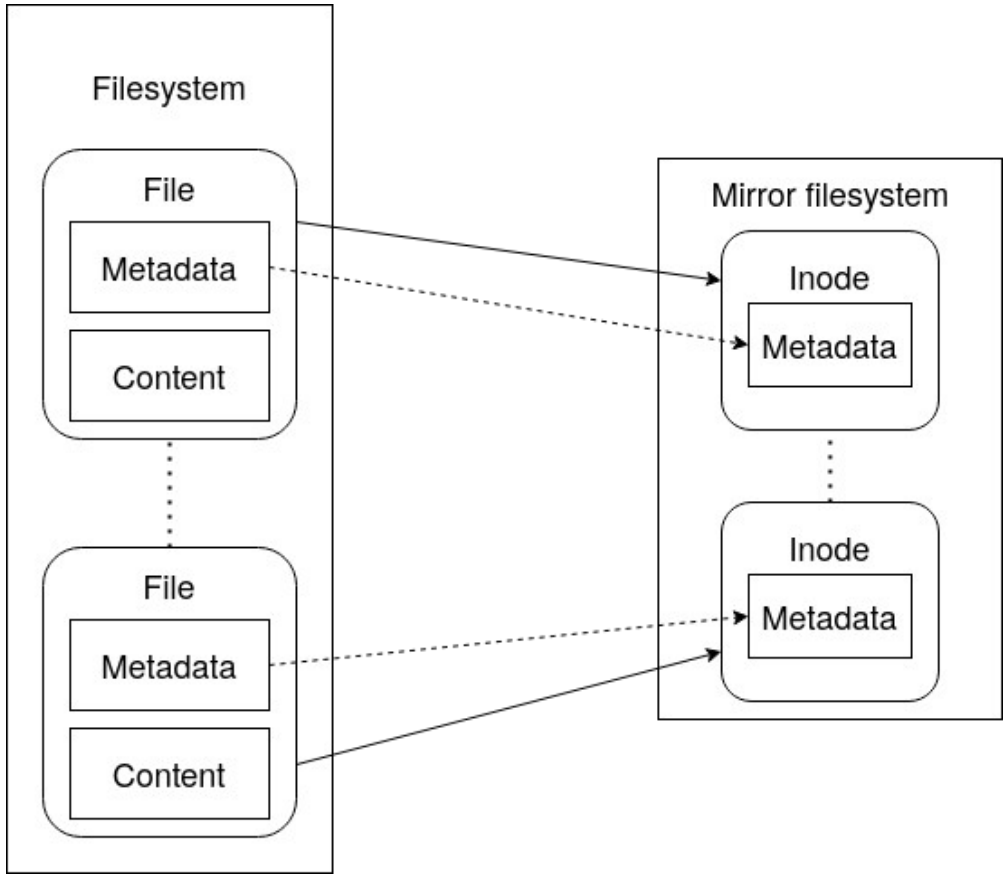

Main goals of RobinHood

1- Quickly synchronize the filesystem metadata with the mirror system

2- Easily query the mirror system

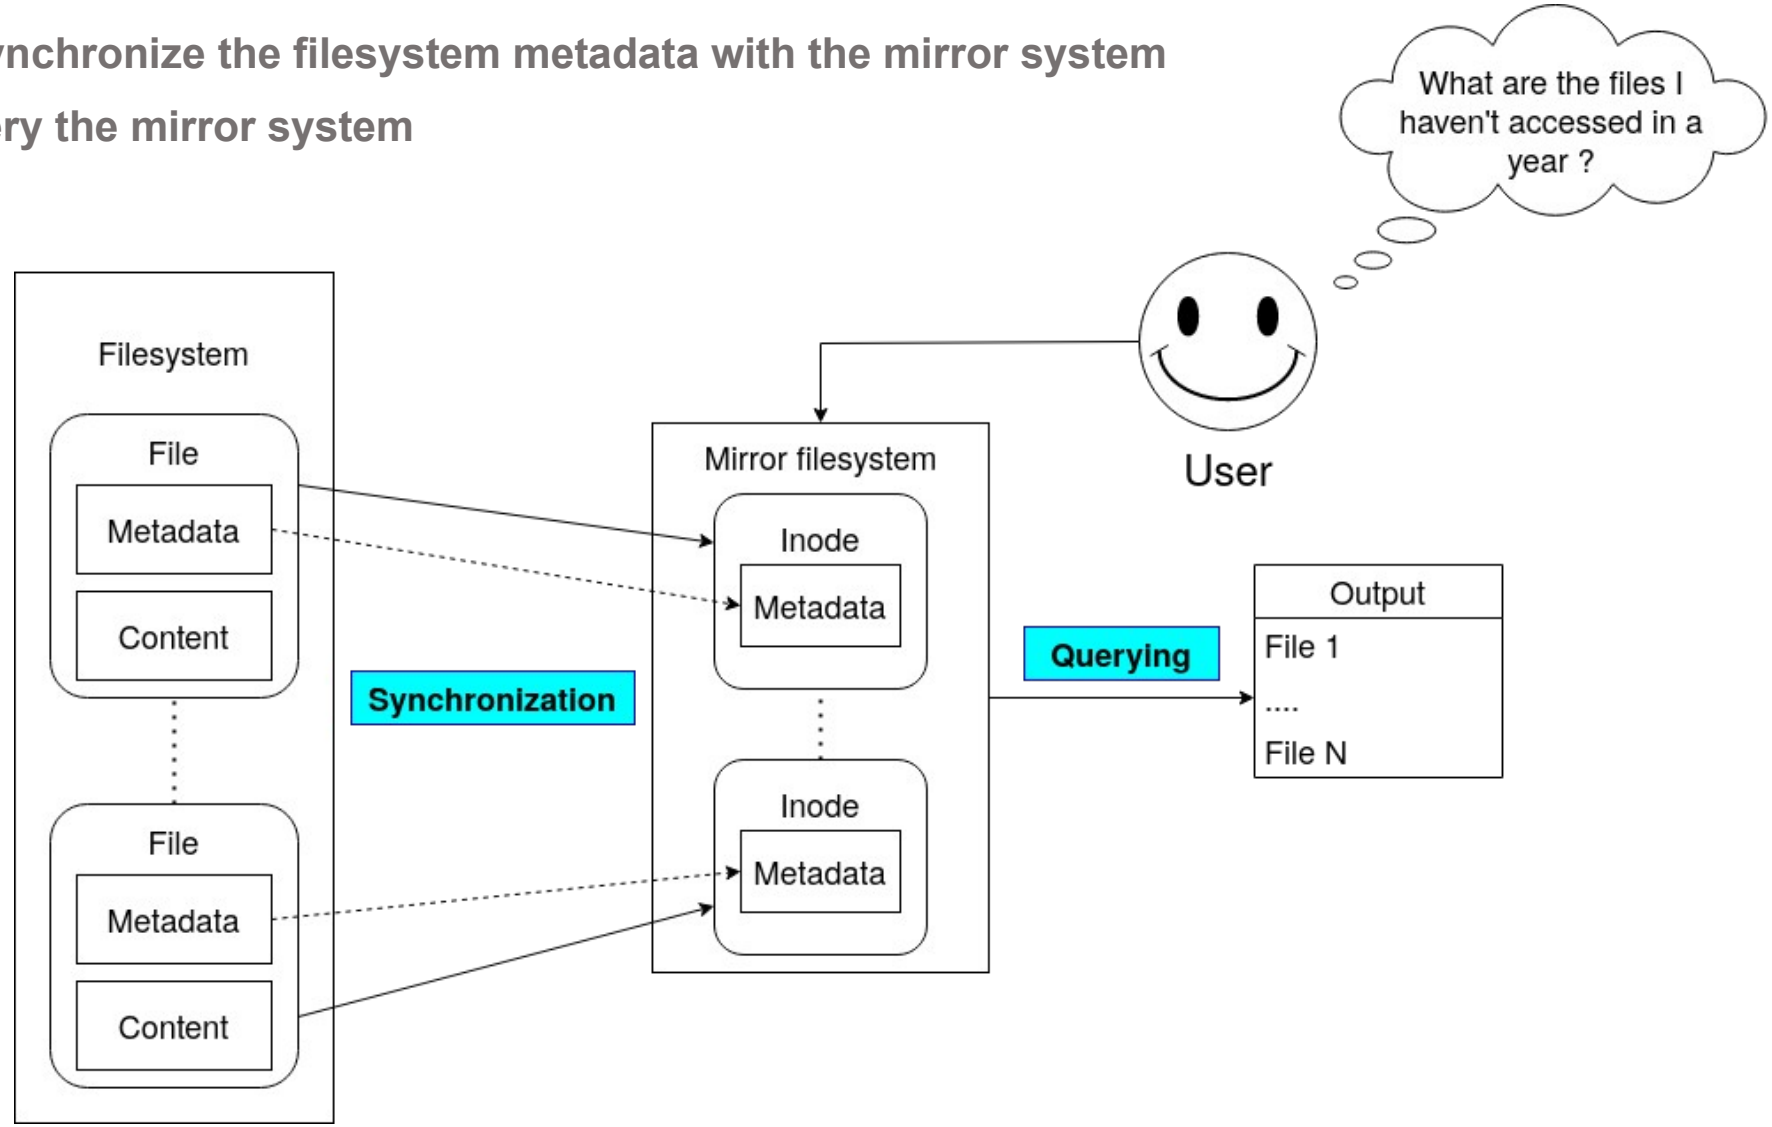

Backends & librobinhood

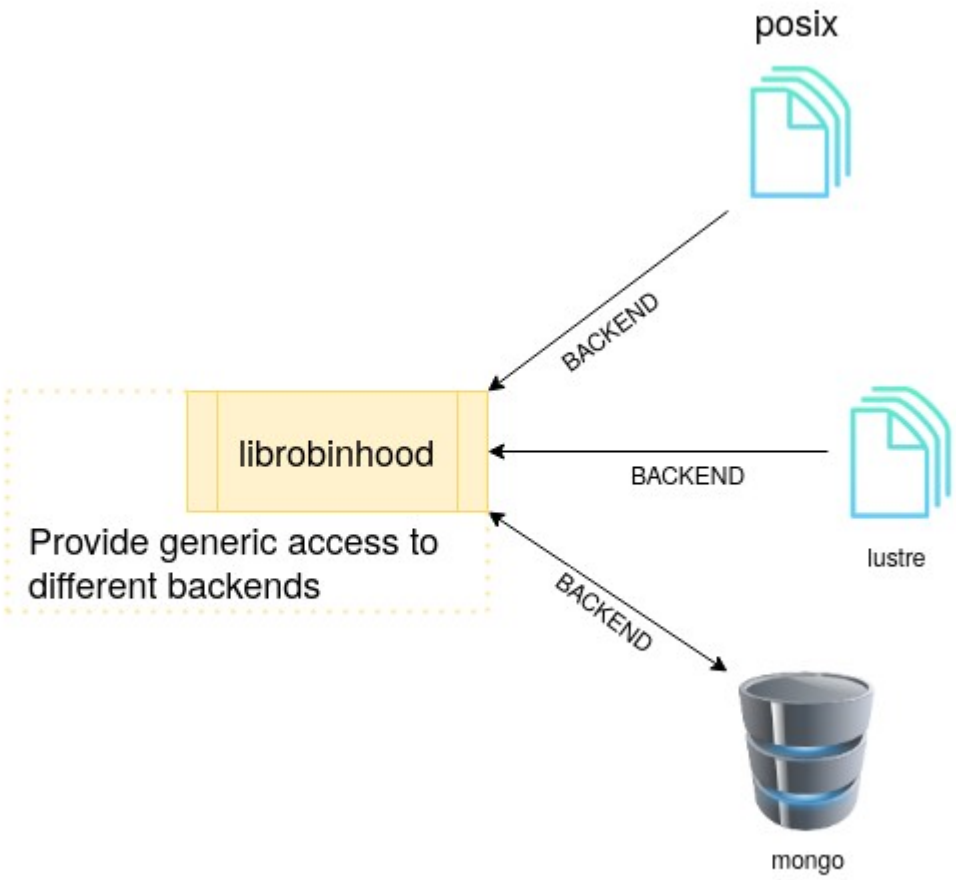

rbh-sync

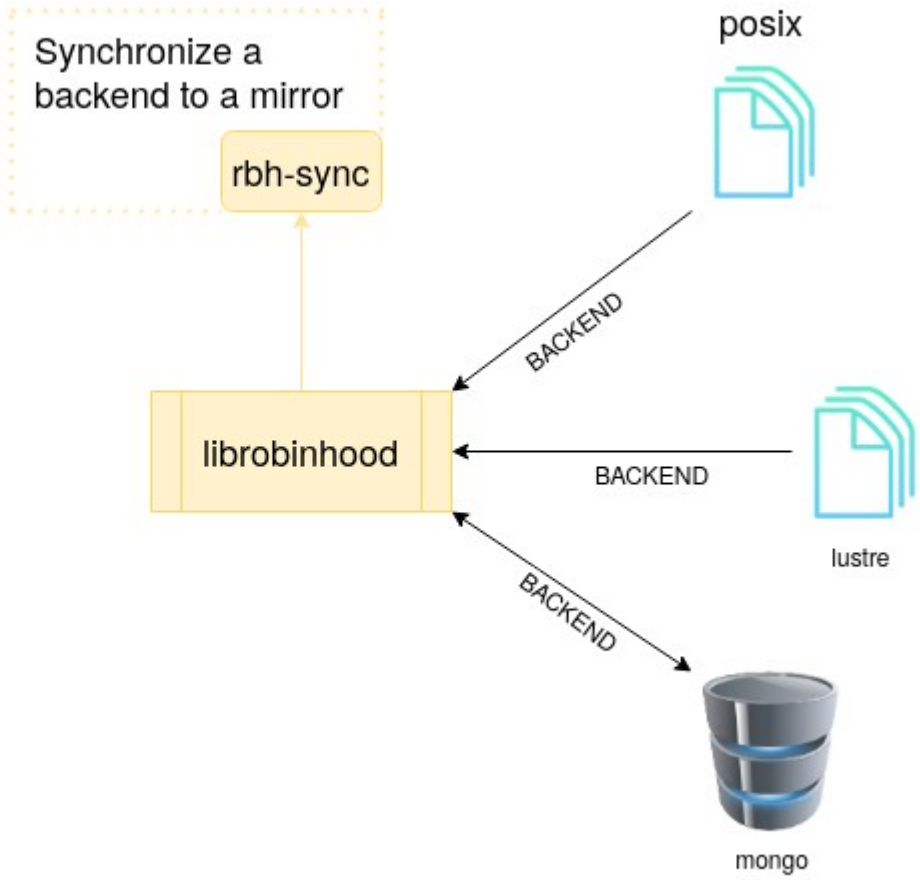
Synchronize a backend to a mirror
rbh-sync
posix
BACKEND
librobinhood
BACKEND
lustre
BACKEND
mongo

rbh-fsevents

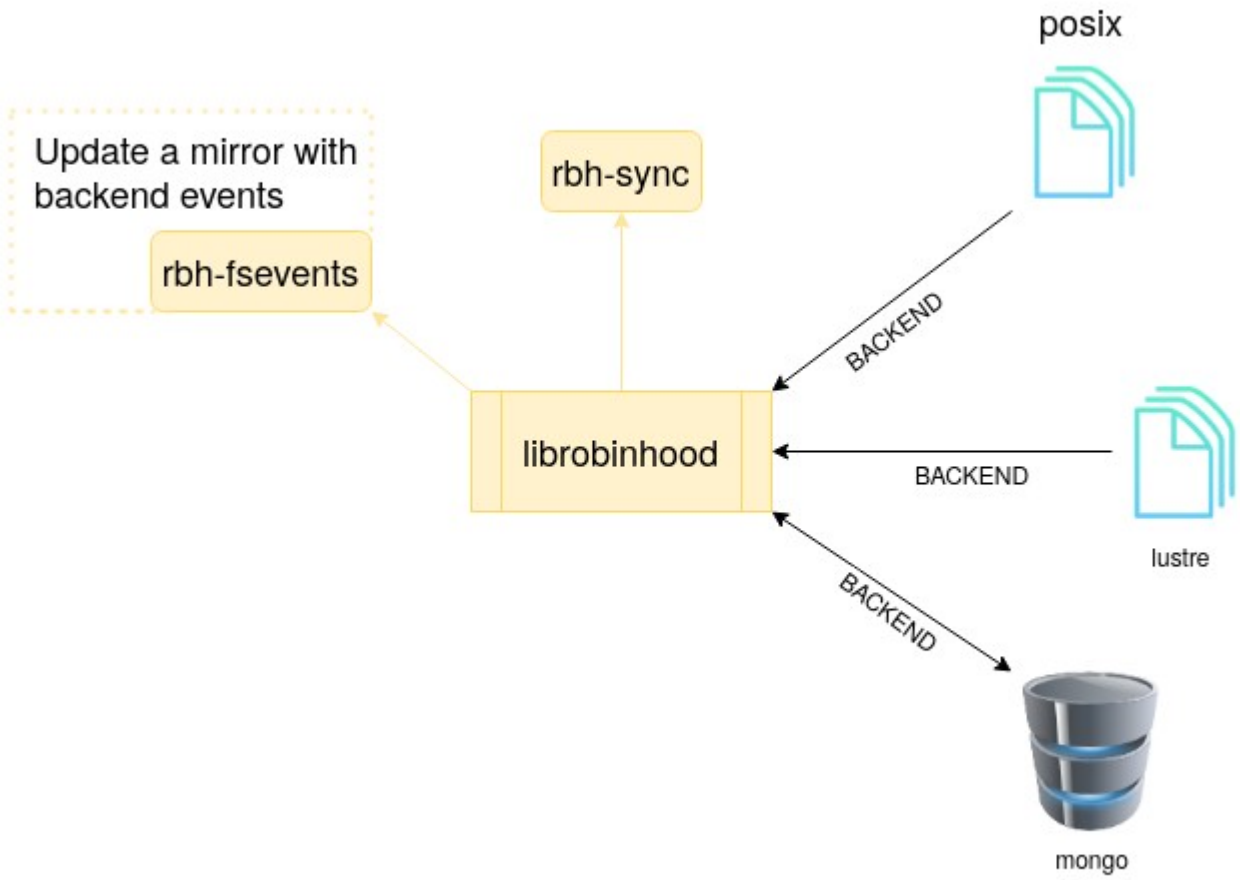

rbh-find

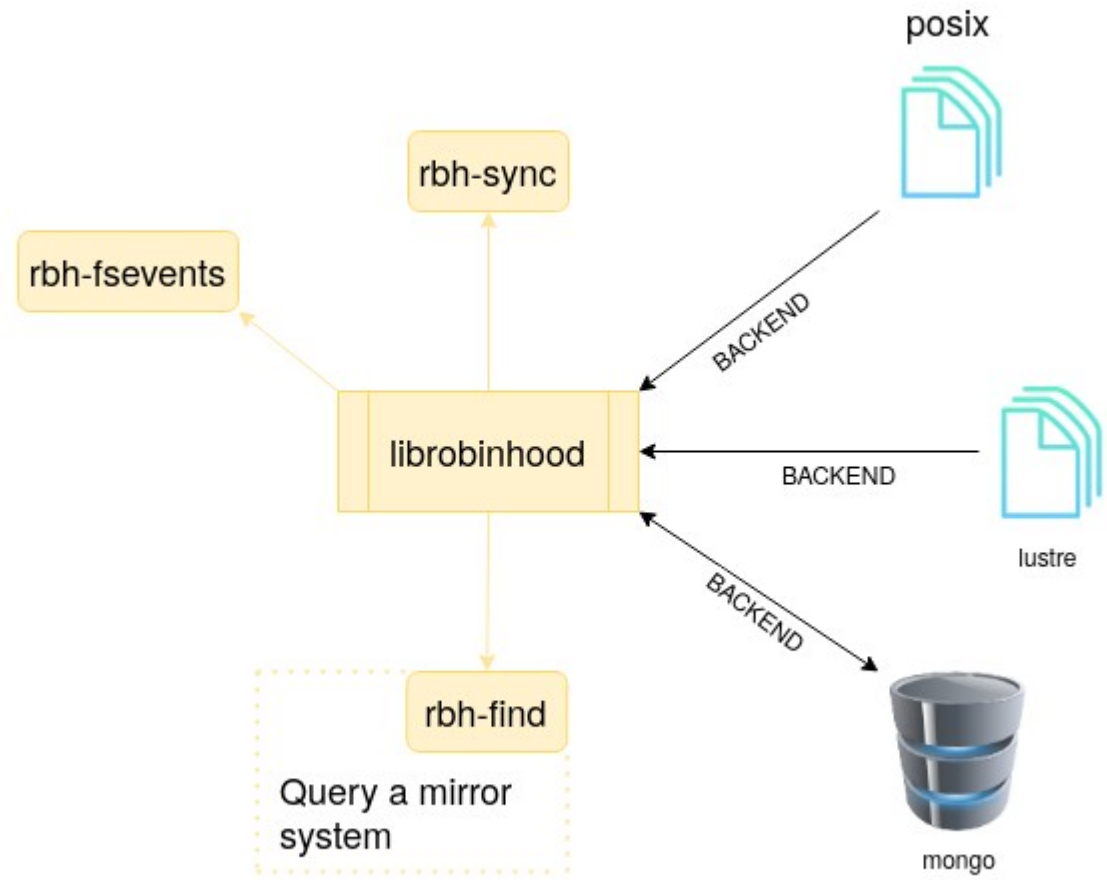
posix
rbh-sync
rbh-fsevents
BACKEND
librobinhood
BACKEND
lustre
BACKEND
rbh-find
Query a mirror system
mongo

rbh-xfind

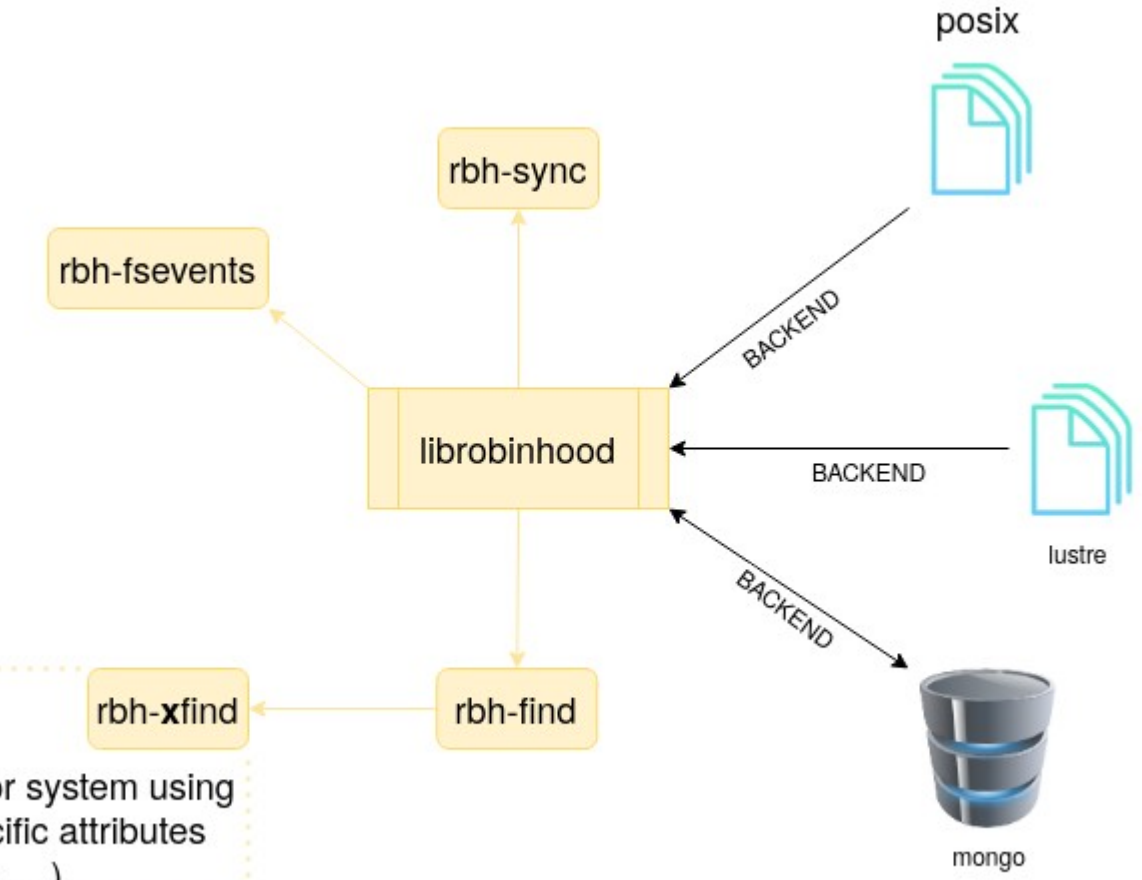

Query a mirror system using backend specific attributes (lustre, io-sea, ...)

rbh-report

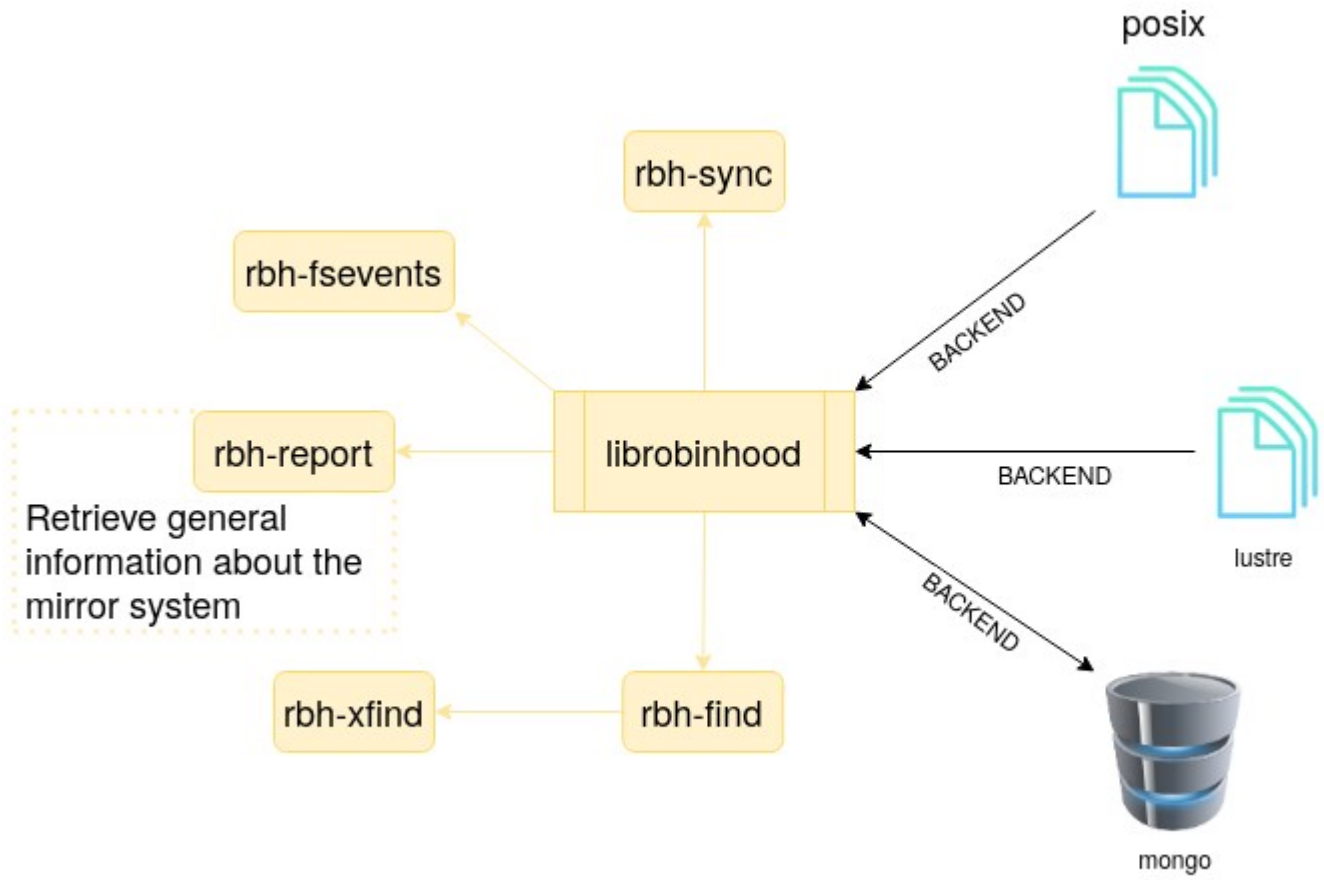
posix
rbh-sync
rbh-fsevents
BACKEND
rbh-report
librobinhood
BACKEND
lustre
Retrieve general information about the mirror system
BACKEND
rbh-xfind
rbh-find
mongo

Benefits

- Fast
  - Minimal impact on the real filesystem
  - Easily scalable with the Mongo backend while remaining efficient
- Generic
  - All tools interact with any backend through librobinhood
  - Easy to add backends
  - Easy to create tools
  - Filter a backend using a database semantic with multiple criteria
    - `rbh-find rbh:mongo:test_db -type f -size +3G -mtime +150 -name “*.txt” -ls`
- Tool suite
  - Single tool for specific actions managed by its own CLI
  - Tools can be overridden for specific options

Example of use – rbh-sync and rbh-find

```
#ls -color test_dir
/big
/huge
/small

#rbh-sync rbh:posix:test_dir rbh:mongo:test
#rbh-find rbh:mongo:test -type f -sort name
/big
/huge
/small

#rbh-find rbh:mongo:test -type f -sort size
/small
/big
/huge
```

```
#rbh-find rbh:mongo:test -type f -ls
144...292      4 -rw -r--r-- 1 root root        17 Sep 13 14:00 /small
144...293   1024 -rw -r--r-- 1 root root   1048576 Sep 13 14:10 /big
144...294 284368 -rw -r--r-- 1 root root 291188736 Sep 13 14:11 /huge

#rbh-find rbh:mongo:test -type f -rsort size -ls
144...294 284368 -rw -r--r-- 1 root root 291188736 Sep 13 14:11 /huge
144...293   1024 -rw -r--r-- 1 root root   1048576 Sep 13 14:10 /big
144...292      4 -rw -r--r-- 1 root root        17 Sep 13 14:00 /small

# rbh-find rbh:mongo:test -type f -sort size -ls
144...292      4 -rw -r--r-- 1 root root        17 Sep 13 14:00 /small
144...293   1024 -rw -r--r-- 1 root root   1048576 Sep 13 14:10 /big
144...294 284368 -rw -r--r-- 1 root root 291188736 Sep 13 14:11 /huge

# rbh-find rbh:mongo:test -type f -sort name -ls
144...293   1024 -rw -r--r-- 1 root root   1048576 Sep 13 14:10 /big
144...294 284368 -rw -r--r-- 1 root root 291188736 Sep 13 14:11 /huge
144...292      4 -rw -r--r-- 1 root root        17 Sep 13 14:00 /small
```

Conclusion – Plans for the future of RobinHood

- **Integration within IO-SEA**
  - **Task 4.2**
    - **creation of an IO-SEA backend for** librobinhood
    - rbh-find-iosea **tool for query**

- **Deduplication of source events in** rbh-fsevents
- rbh-report **tool for retrieving general information about a mirror**
- **Ongoing development**
  - **Repositories available at** https://www.github.com/cea-hpc/**{librobinhood, rbh-sync, …}**
  - **Example of patches currently in review:**
    - **https://review.gerrithub.io/c/cea-hpc/rbh-fsevents/+/552101 → management of the Lustre MIGRATE events in** rbh-fsevents
    - **https://review.gerrithub.io/c/cea-hpc/librobinhood/+/552485 → definition of a spec file and packaging methods in** librobinhood

RobinHood Policy Engine

Do you have any question ?