



FROM RESEARCH TO INDUSTRY





# Robinhood v4


Scaling RBH for the future














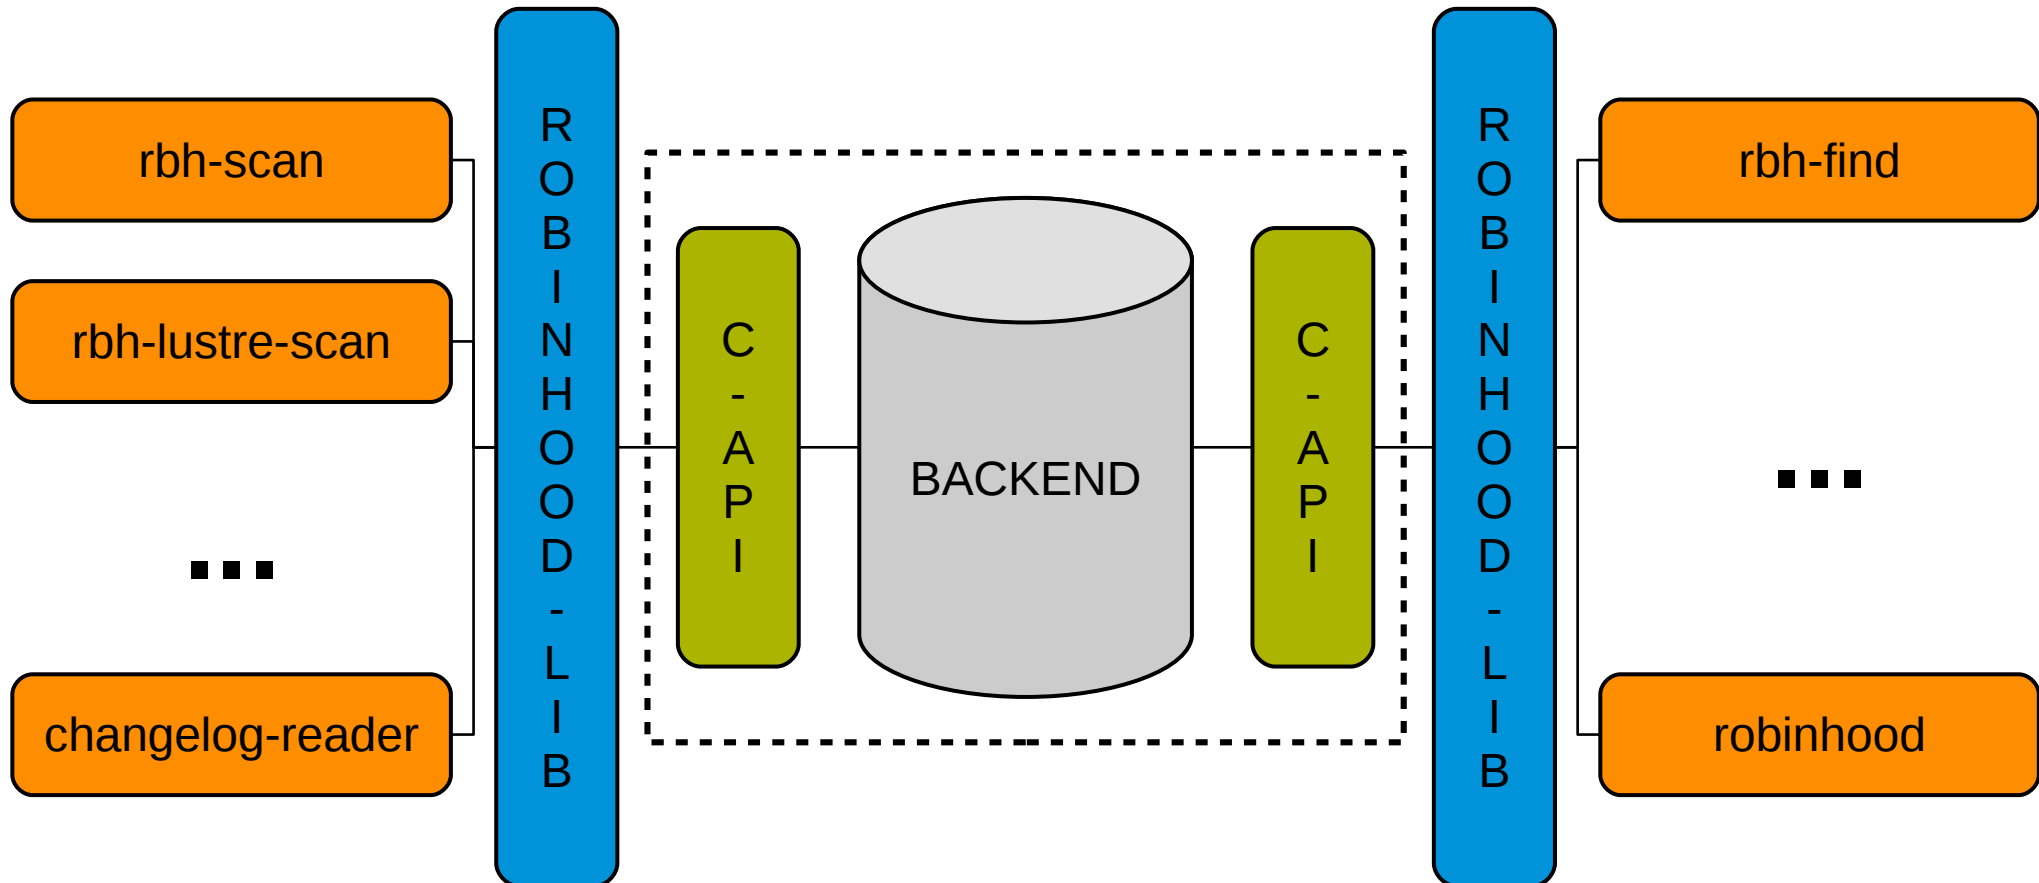


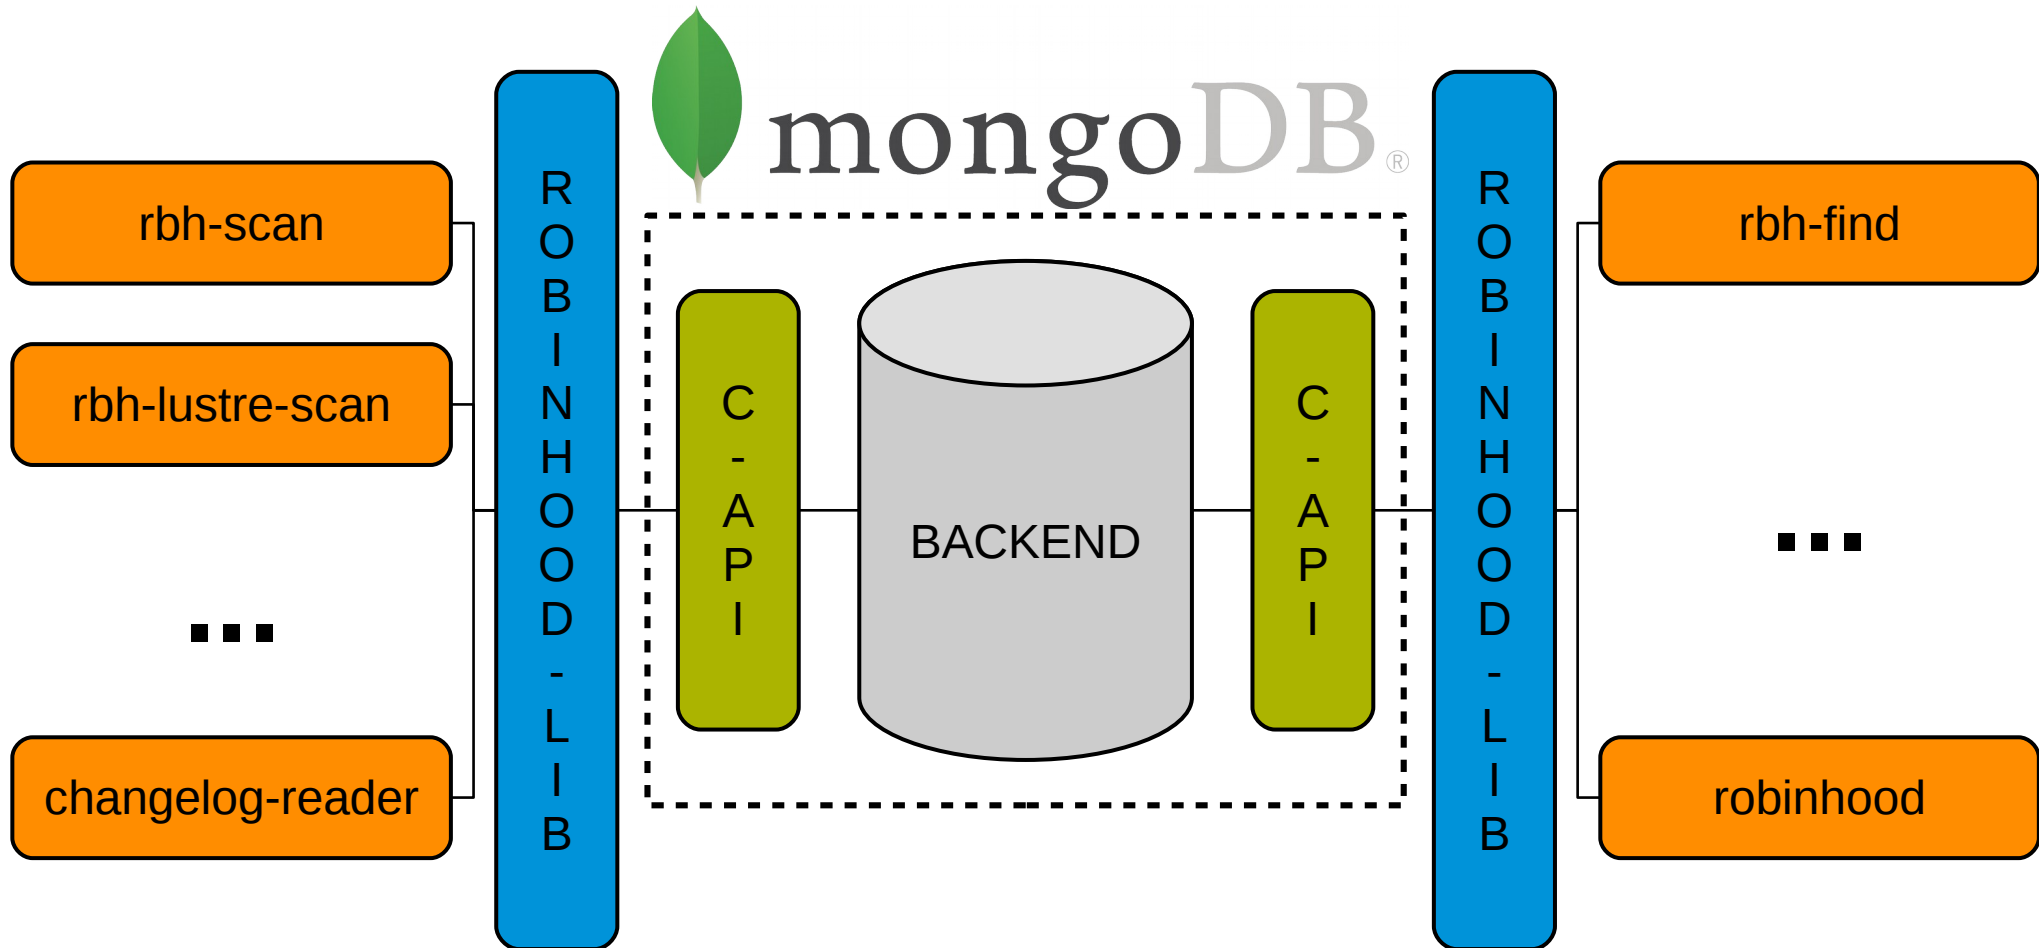


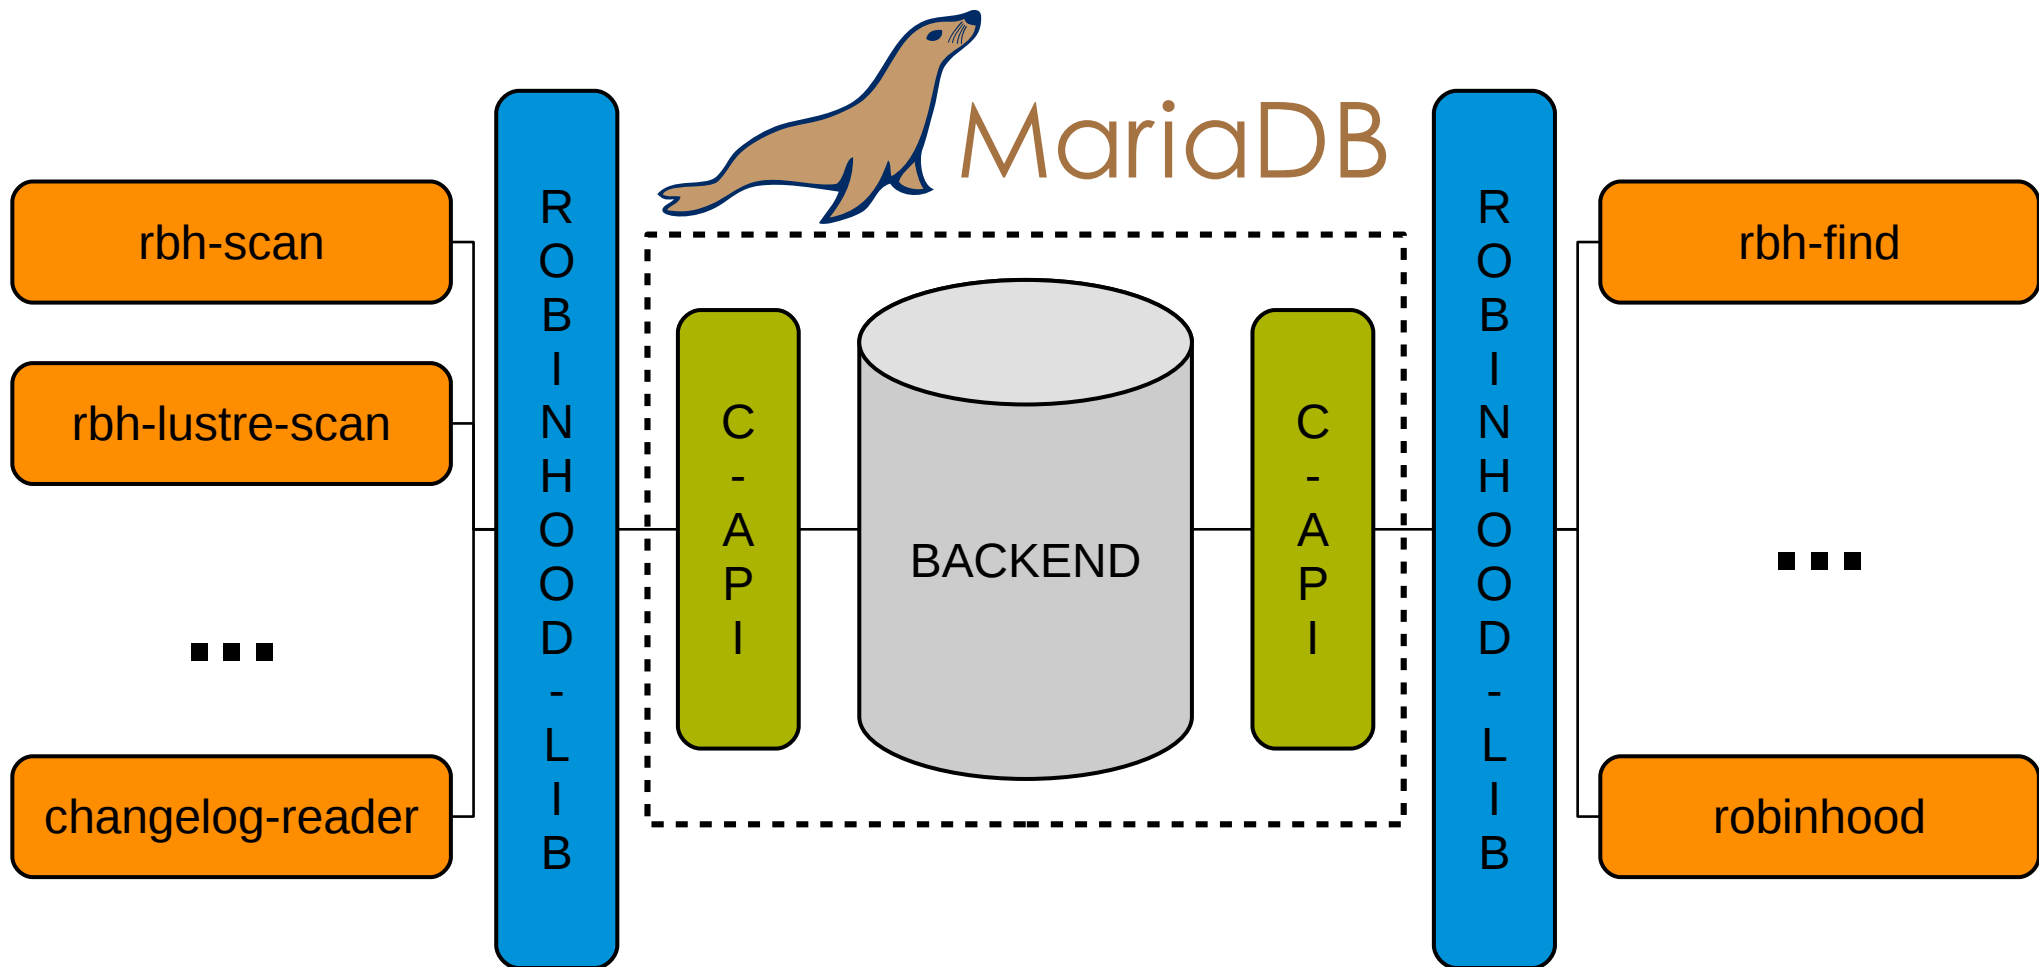


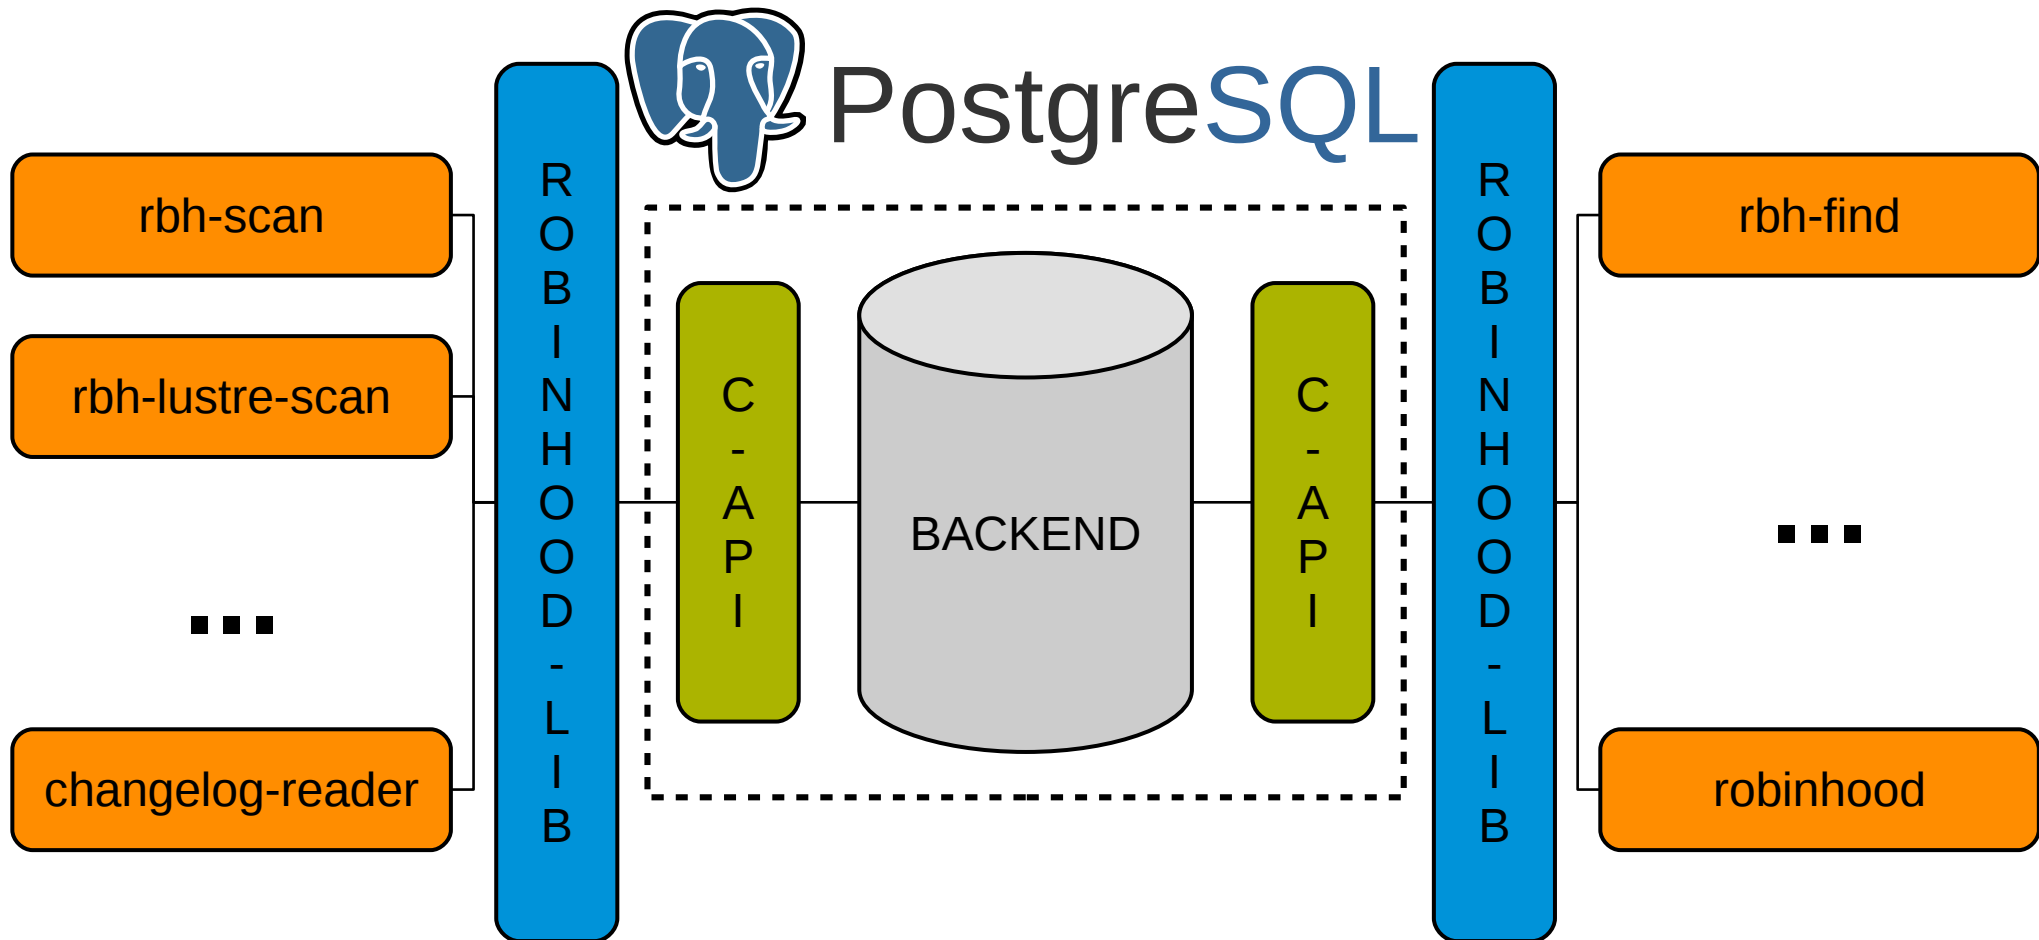
# Architecture (WIP)

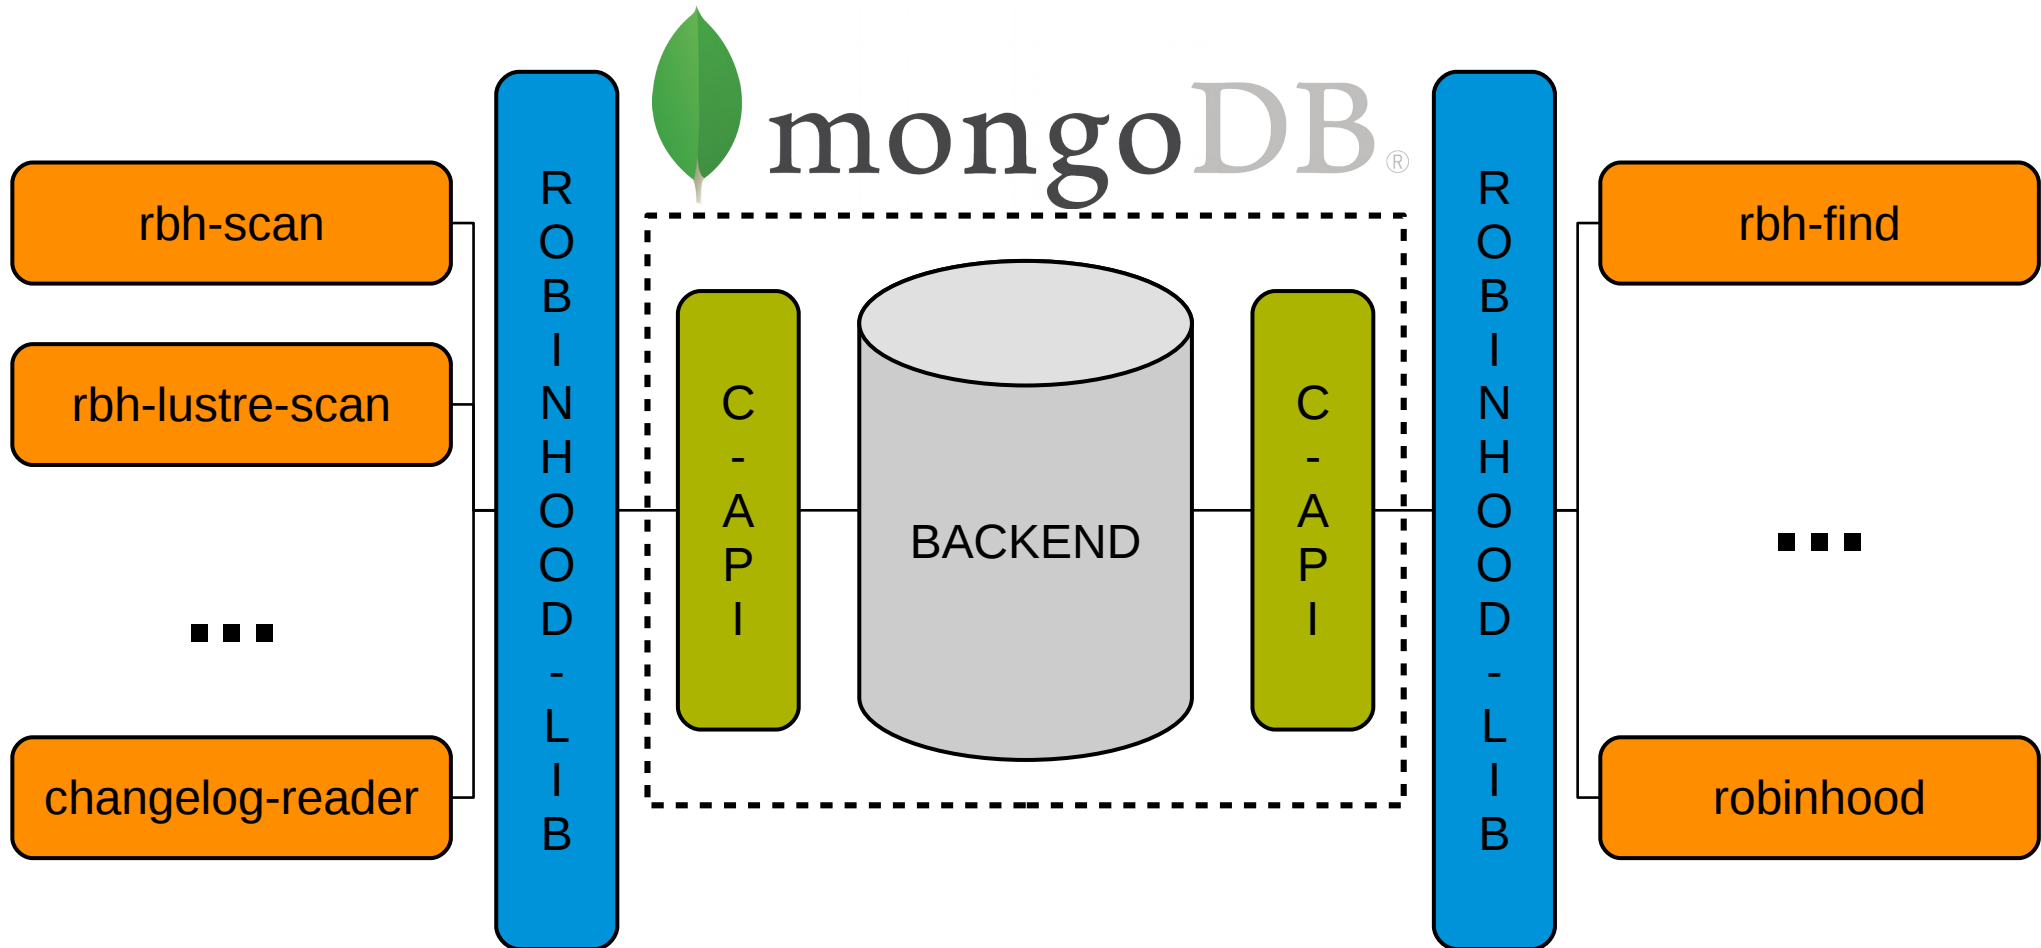












# How hard is it to implement a backend?

# How hard is it to implement a backend?

The MongoDB backend is ~600 lines of code



# How hard is it to implement a backend?

The MongoDB backend is ~600 lines of code

Most of it is boilerplate


# How hard is it to implement a backend?

The MongoDB backend is ~600 lines of code


Most of it is boilerplate \*

- Lustre filesystem
- 400k entries
- Single server (collocated mds, mgs, and oss)
- Single client (!= server)

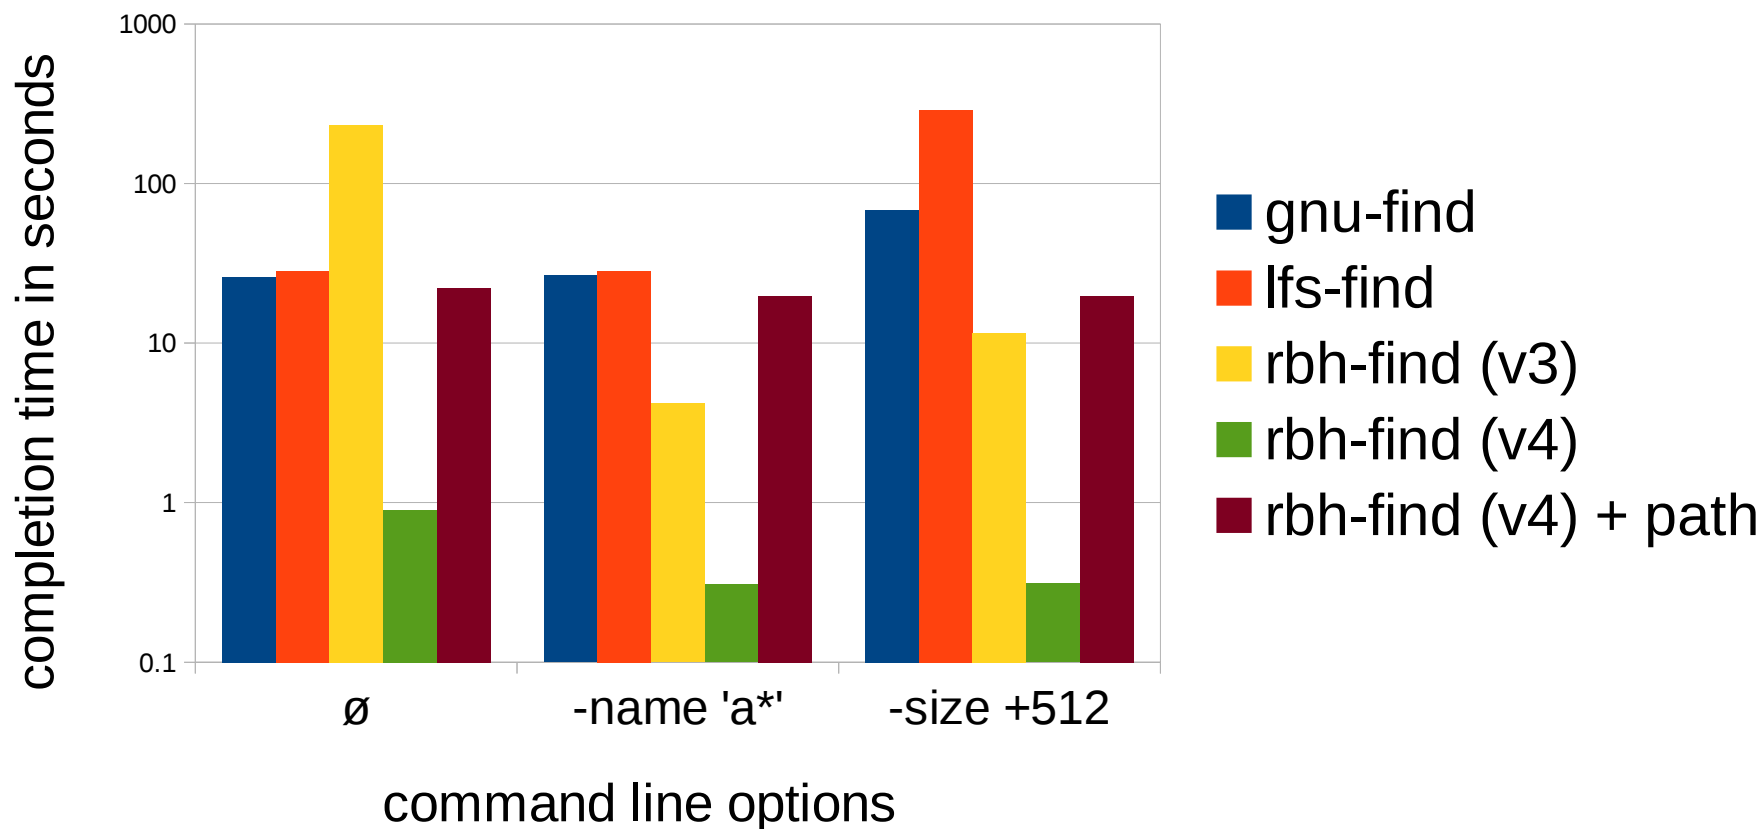
## Speed comparison of different flavors of find



Speed comparison of rbh-find (v4) and rbh-find (v4) + path



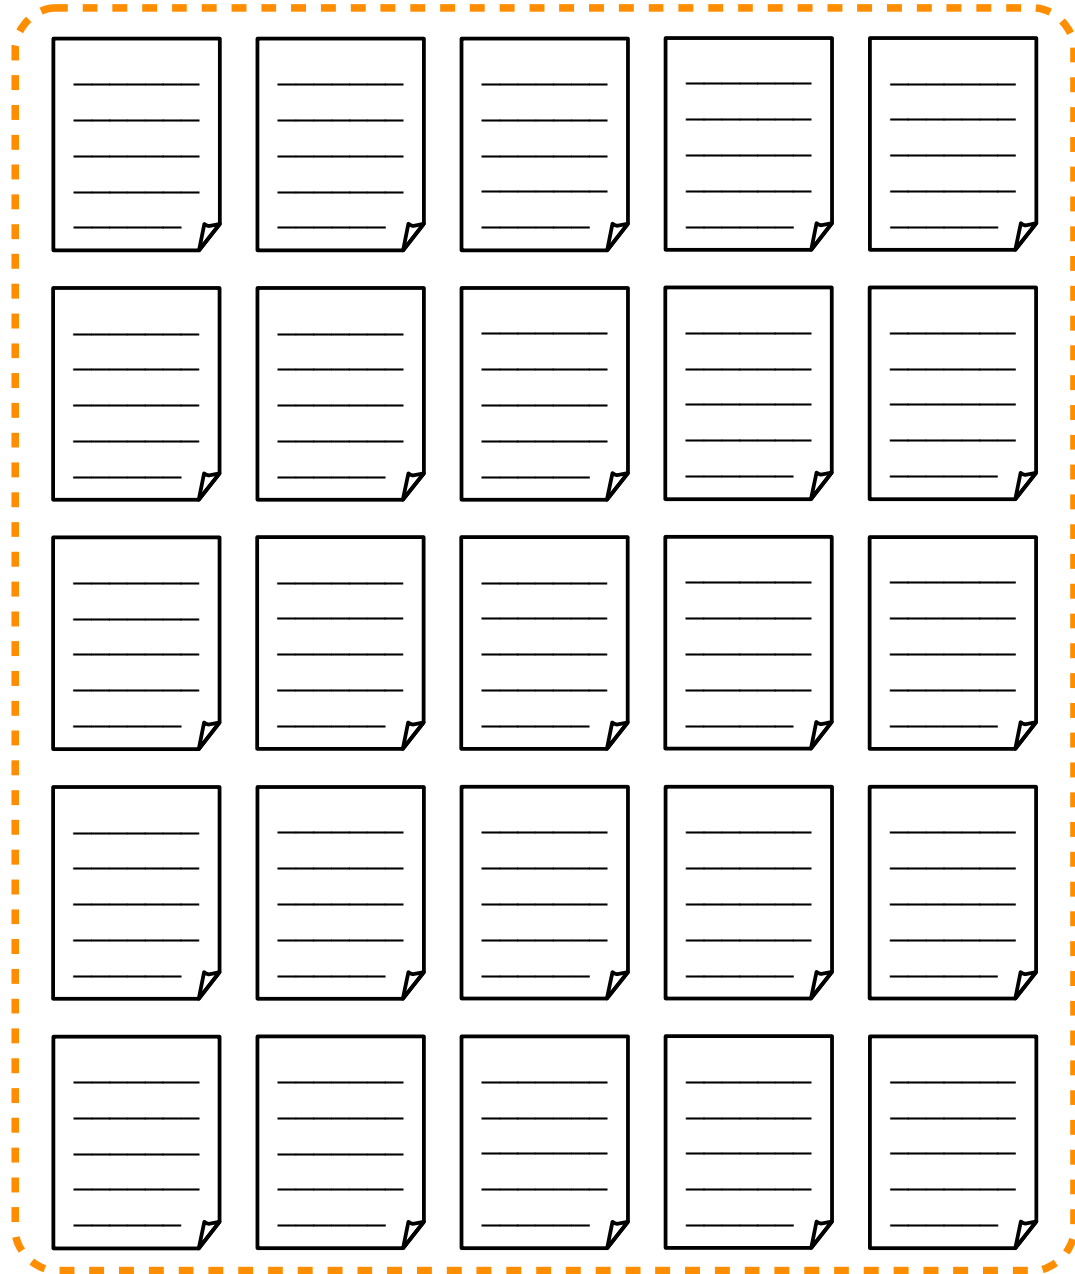
## Speed comparison of different find flavors



# How come?

<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>
<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>
<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>
<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>
<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>	<hr/> <hr/> <hr/> <hr/> <hr/>

# How come?







# How come?




# How come?




# How come?




# How come?




# How come?



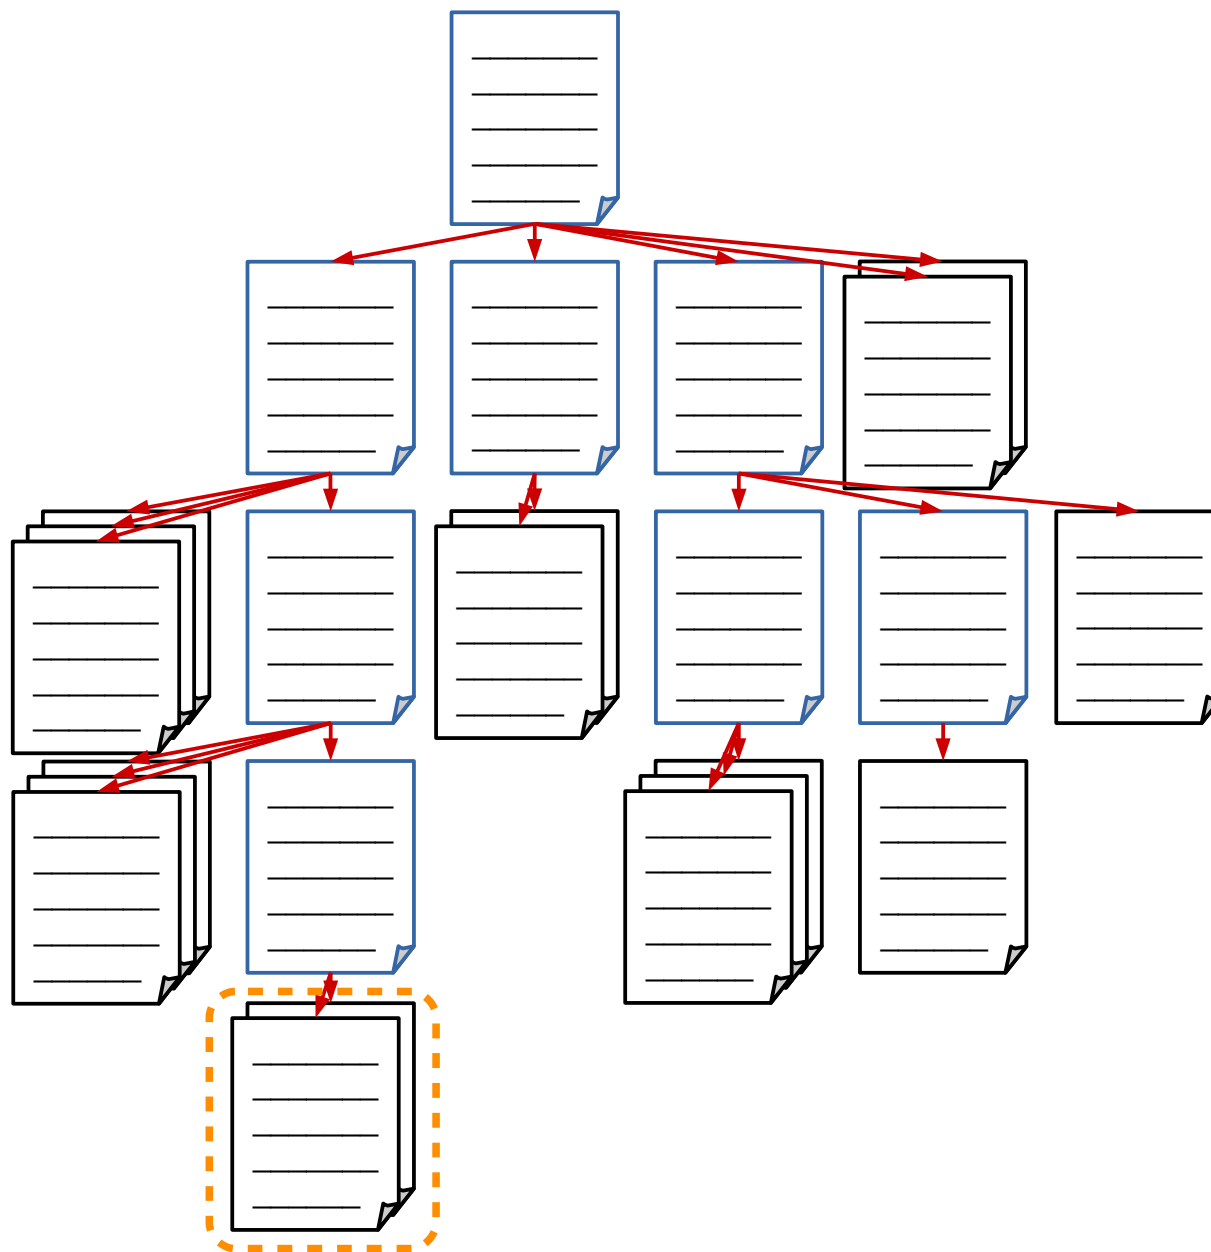
# How come?



# How come?



# How come?

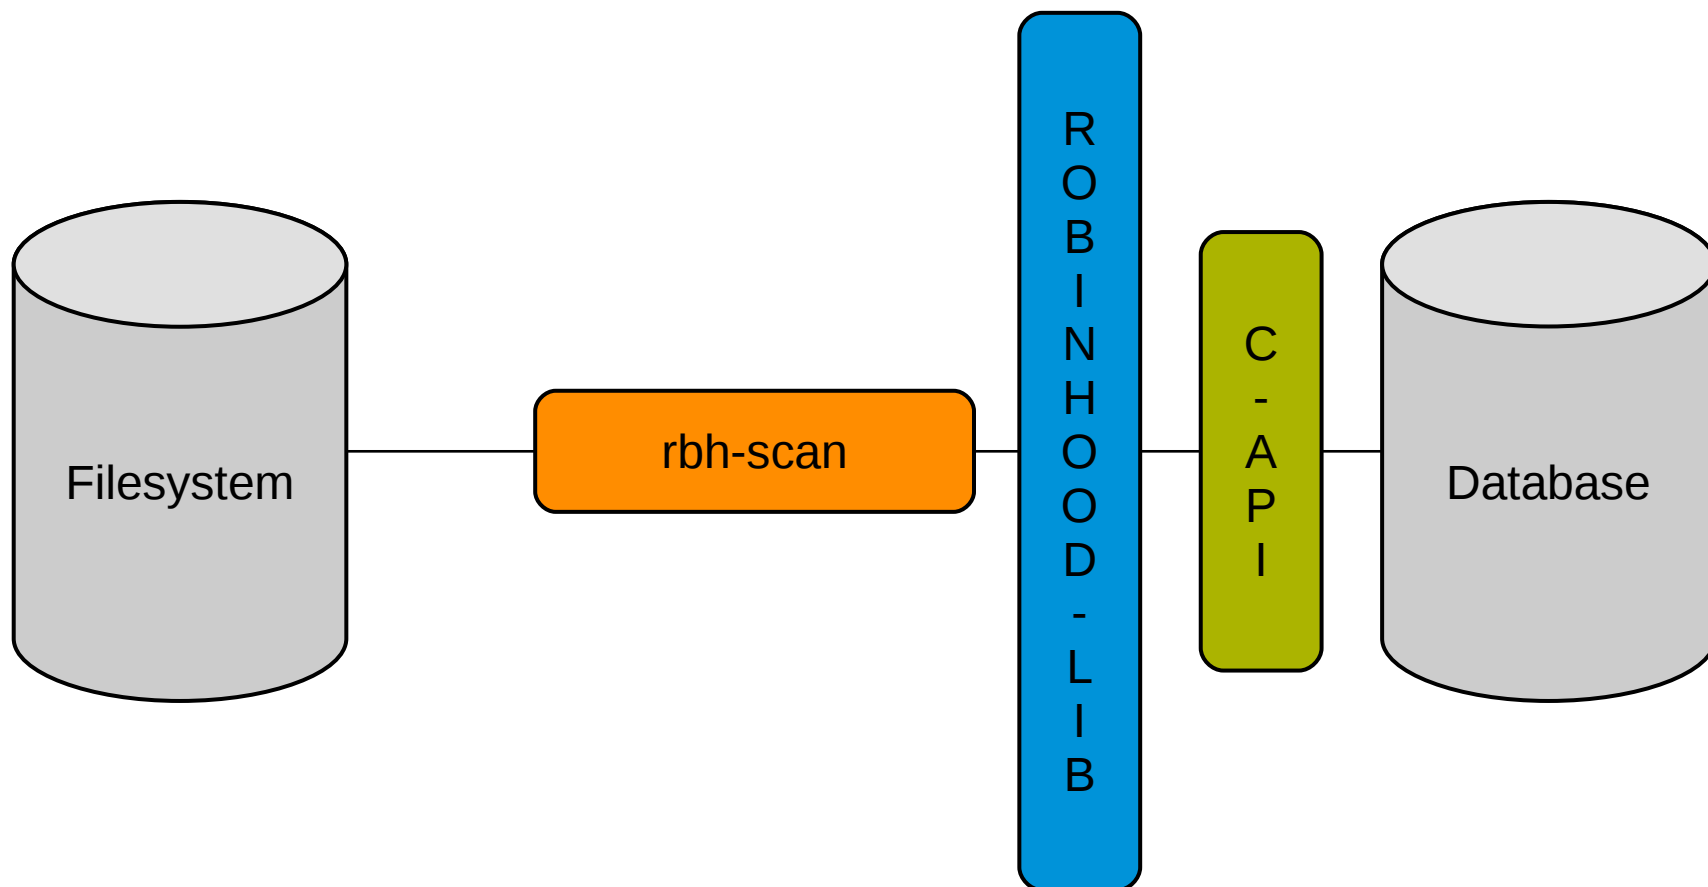




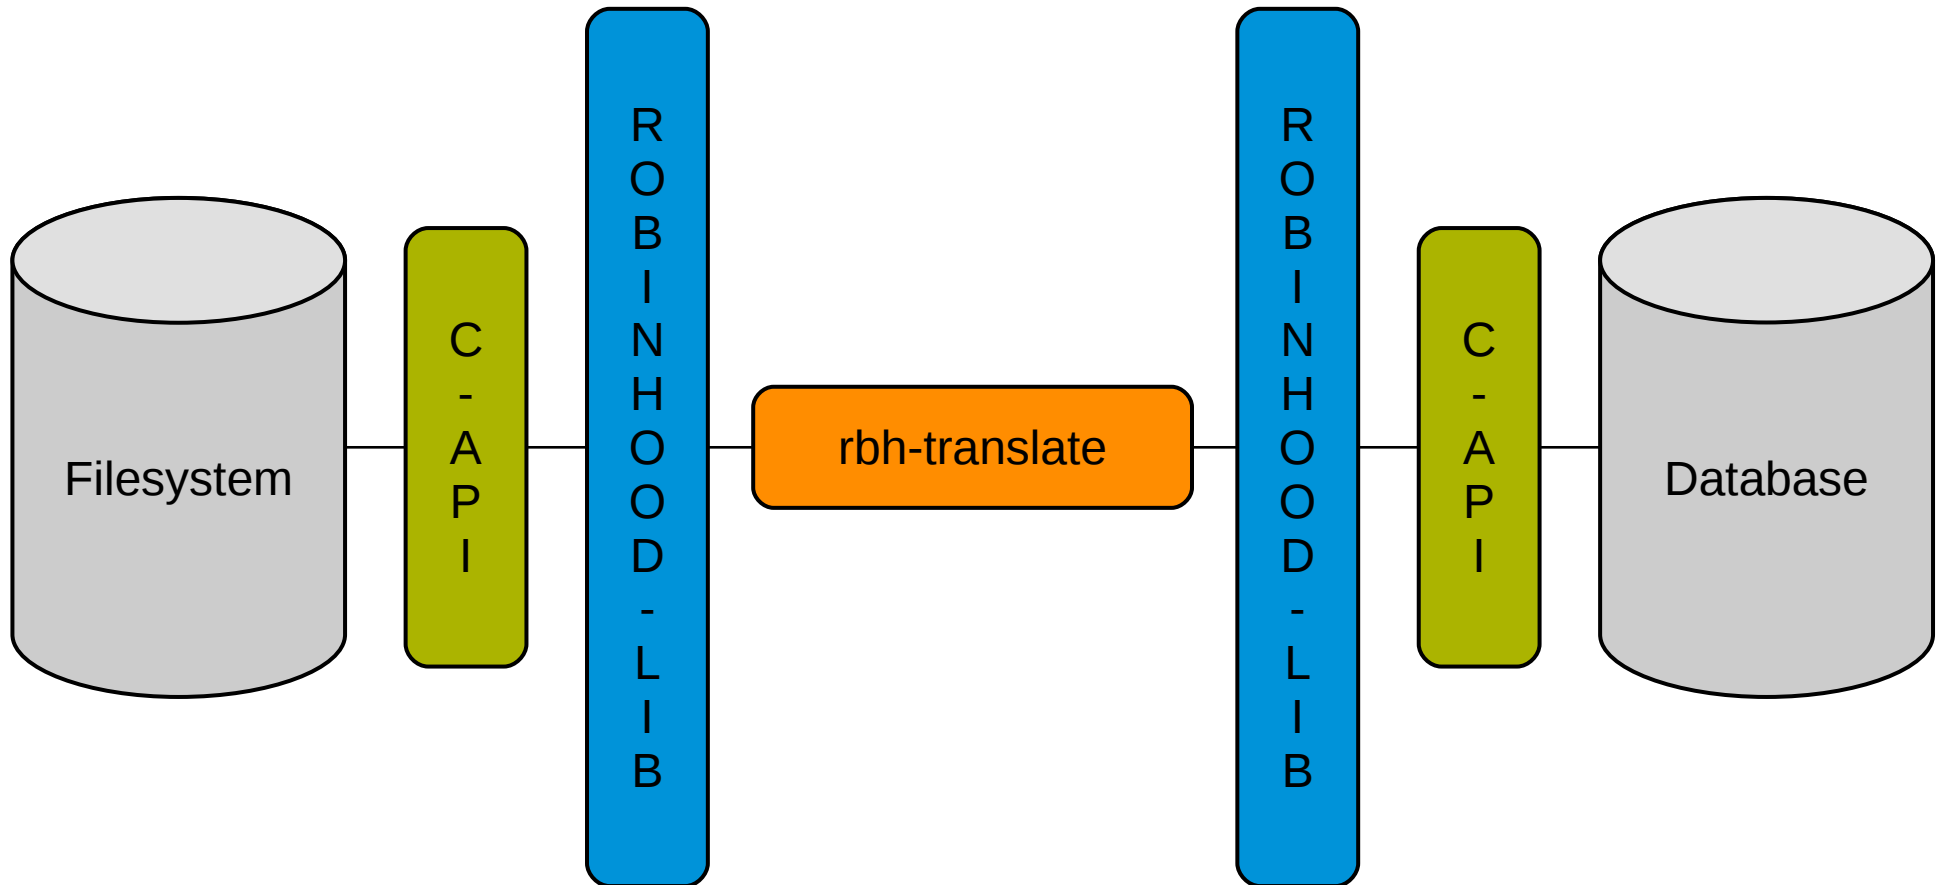
# There is still lots to do!

- changelog\_reader
- rbh-report
- web interface
- posix scanner \*

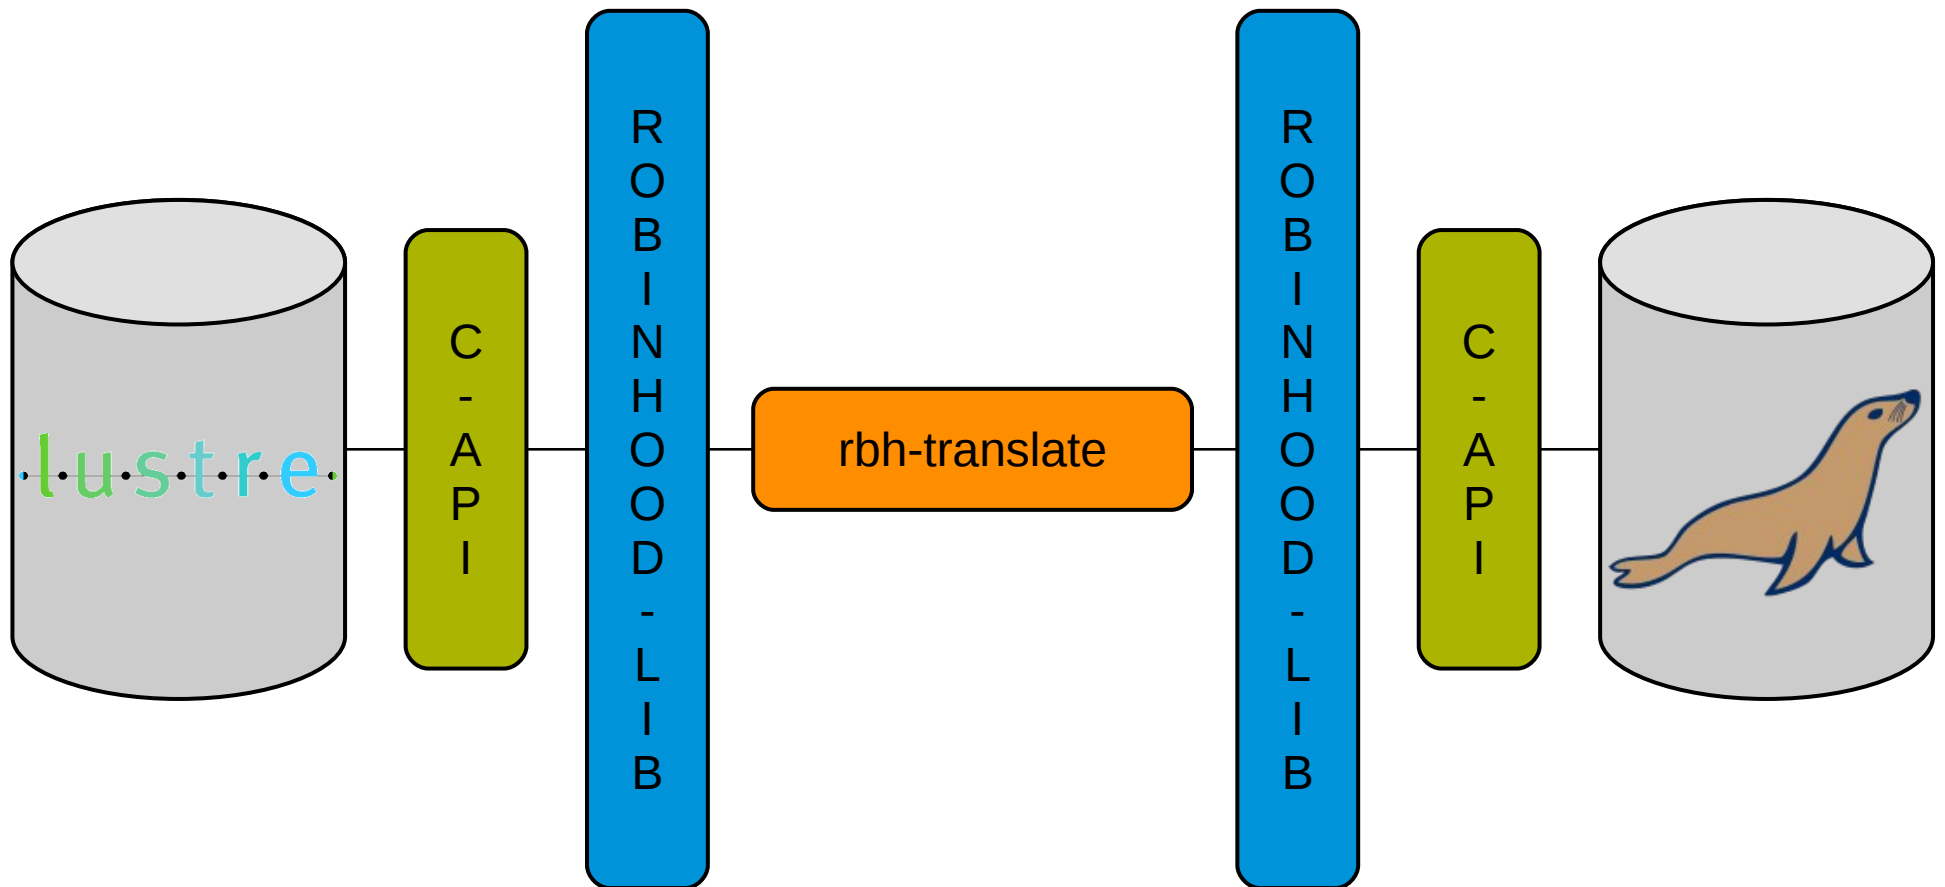
# Posix scanner or posix backend?




# Posix scanner or posix backend?



# Posix scanner or posix backend?



# Posix scanner or posix backend?





# Do we still need a policy engine?

# Do we still need a policy engine?

```
robinhood --run=purge --target=class:old
```

# Do we still need a policy engine?

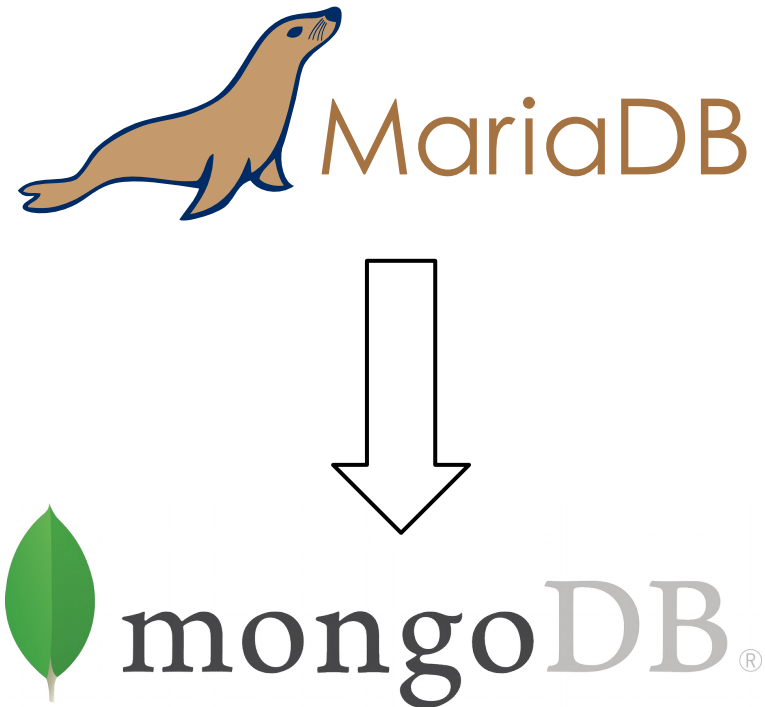
robinhood --run=purge --target=class:old

Vs

rbh-find -atime +30 -execdir rm {} \+



- Developer-friendly
- Admin-friendly
- Scalability



Thank you!

Questions?

- MongoDB's logo - unknown - fair use  
<https://www.mongodb.com/brand-resources>
- MariaDB's seal - Mike Zinner - GFDL  
<http://badges.mariadb.org/logo/Mariadb-seal-flat-browntext.svg>
- PostgreSQL's elephant - Daniel Lundin - PostgreSQL's license  
[https://wiki.postgresql.org/images/a/a4/PostgreSQL\\_logo.3colors.svg](https://wiki.postgresql.org/images/a/a4/PostgreSQL_logo.3colors.svg)
- Lustre's logo - unknown - fair use  
[https://web.archive.org/web/20010805232820im\\_/http://lustre.org/images/lustre\\_logo.gif](https://web.archive.org/web/20010805232820im_/http://lustre.org/images/lustre_logo.gif)