





DE LA RECHERCHE À L'INDUSTRIE

RobinHood-v4 progress report

What is RobinHood?

RobinHood provides efficient and easy to use means to replicate and query any filesystem's metadata.

Filesystem

search: $O(n)$ 



Database


search: $O(1)$ (with an index)

What's new?


Extended attributes (xattrs)

Custom metadata attributes






key=0x1234...




●
key={
 .some = ["struc", "tured"],
 .data = 3306,
}

- string-indexed mappings
- sequences
- scalars (sized integers, strings, binary blobs)





●
layout=[{...}, {...}, ...]










layout=[{...}, {...}, ...]


TODO

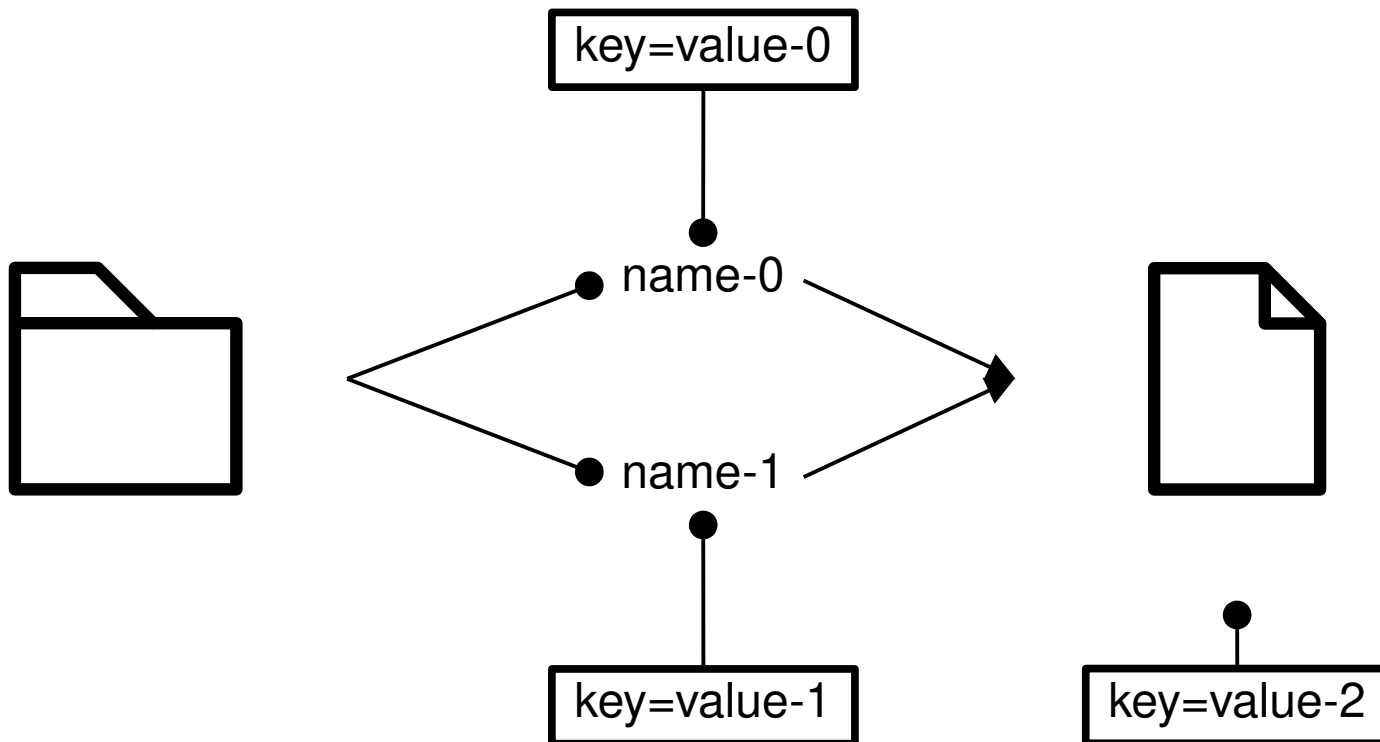


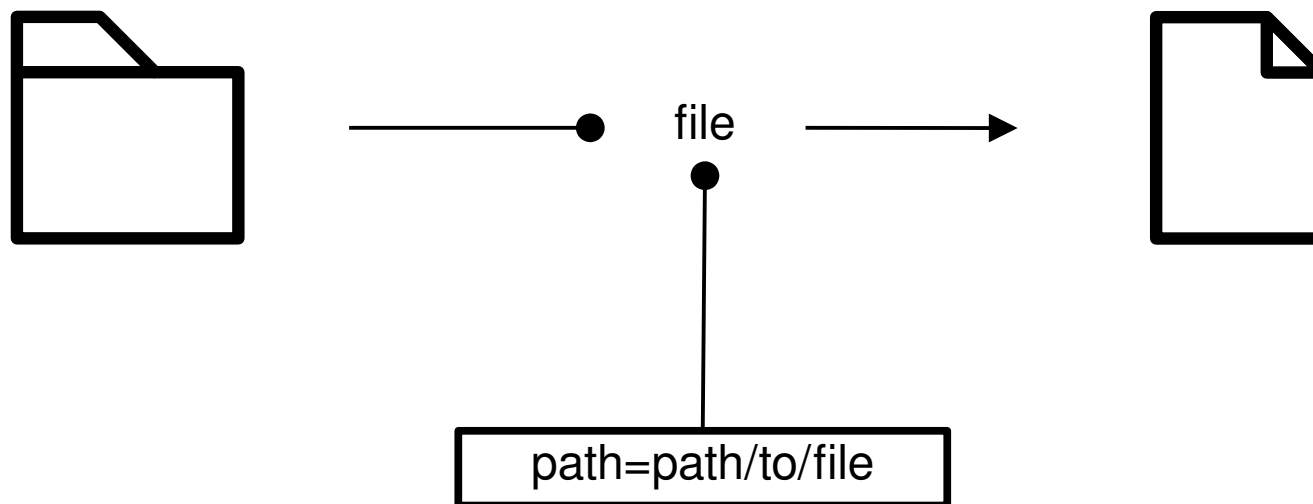








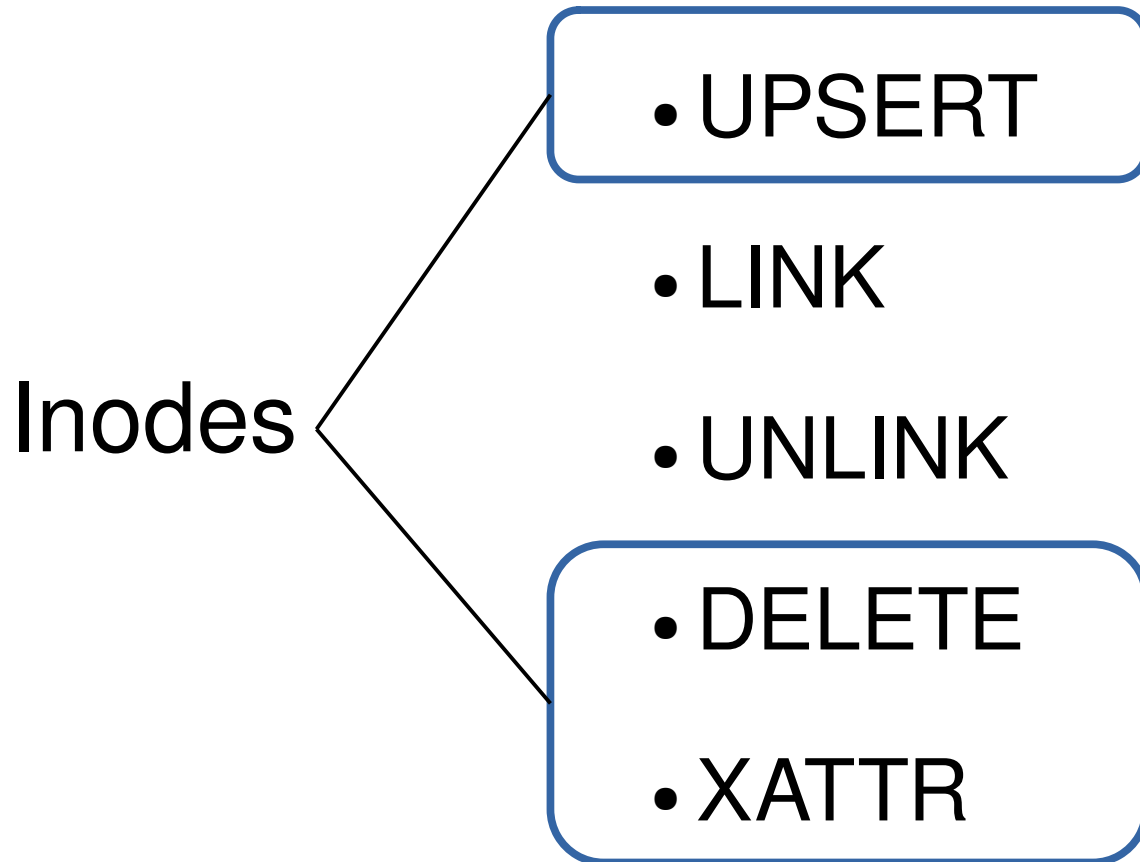


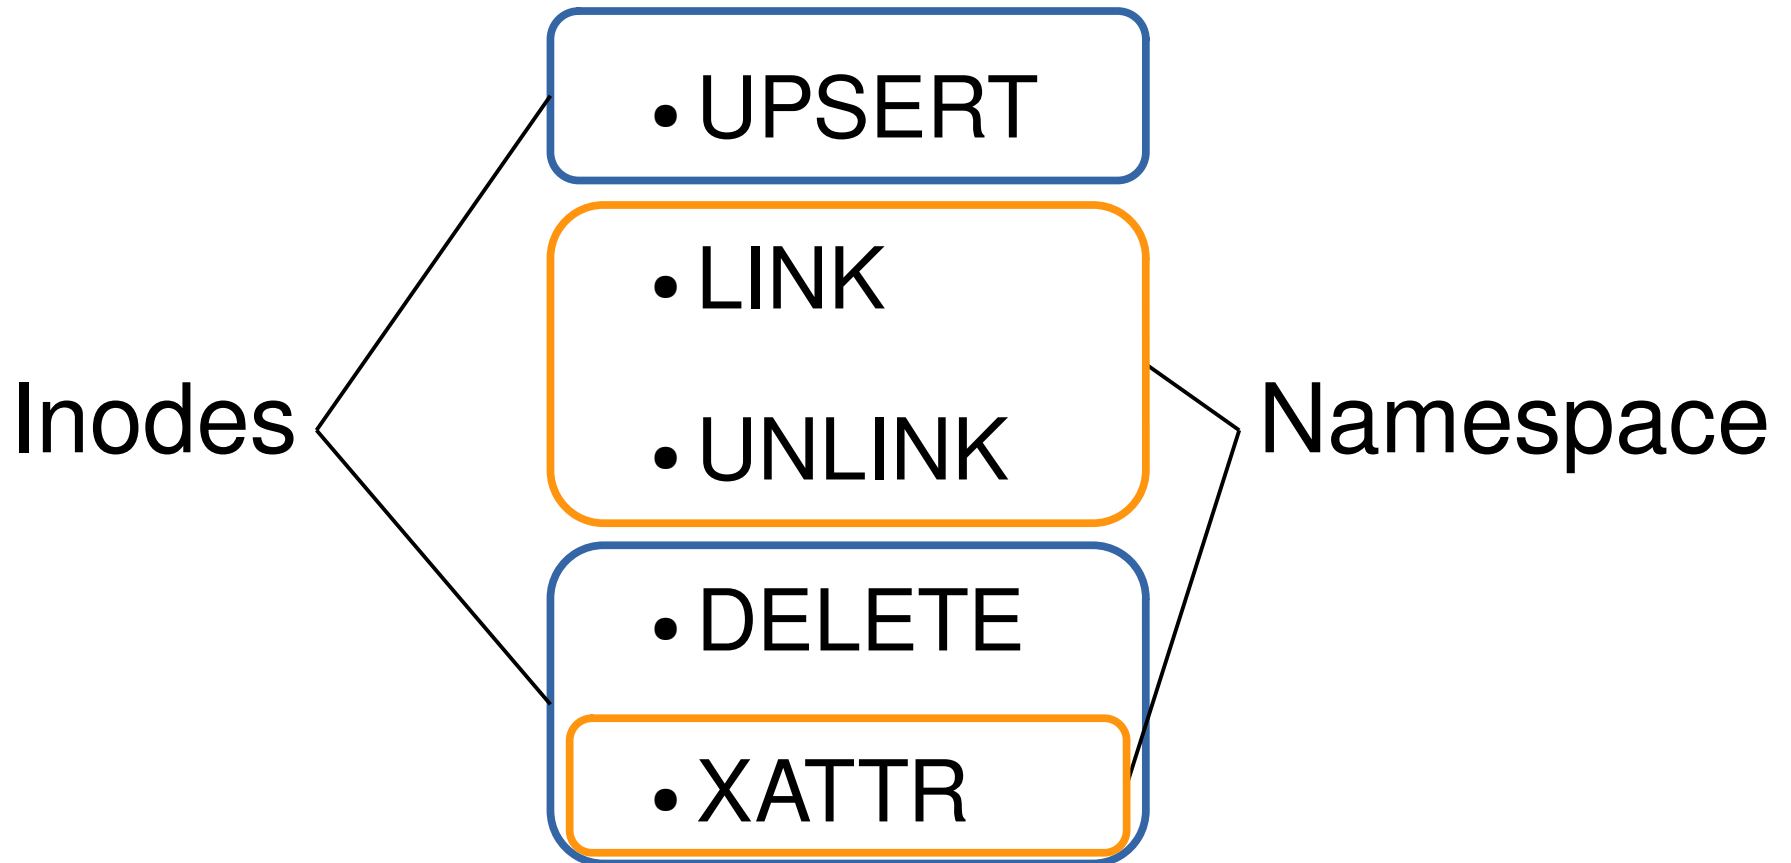


FSEvents


Metadata update records

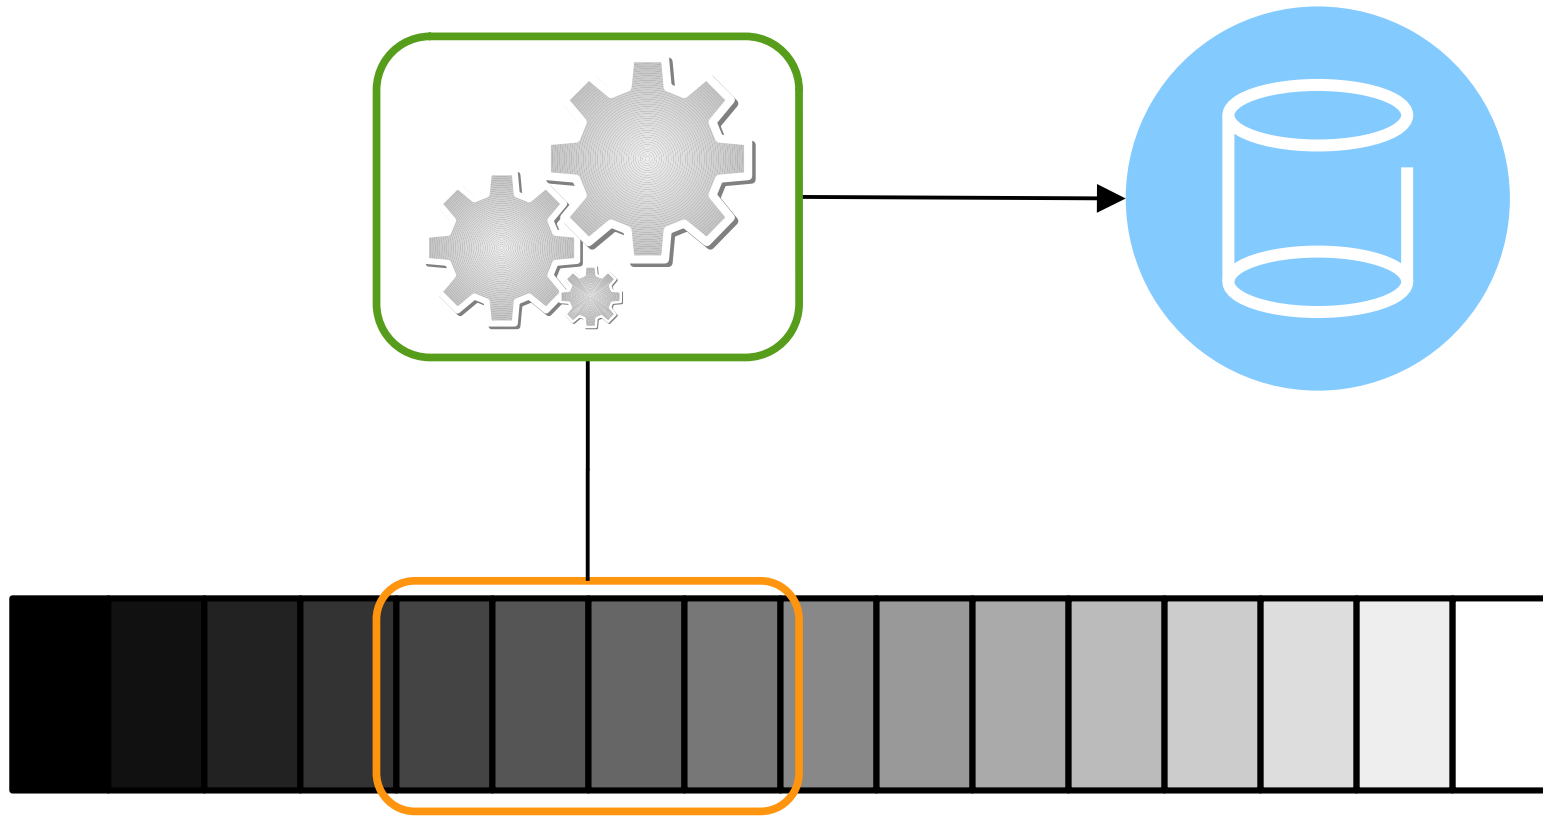
- UPSERT
- LINK
- UNLINK
- DELETE
- XATTR

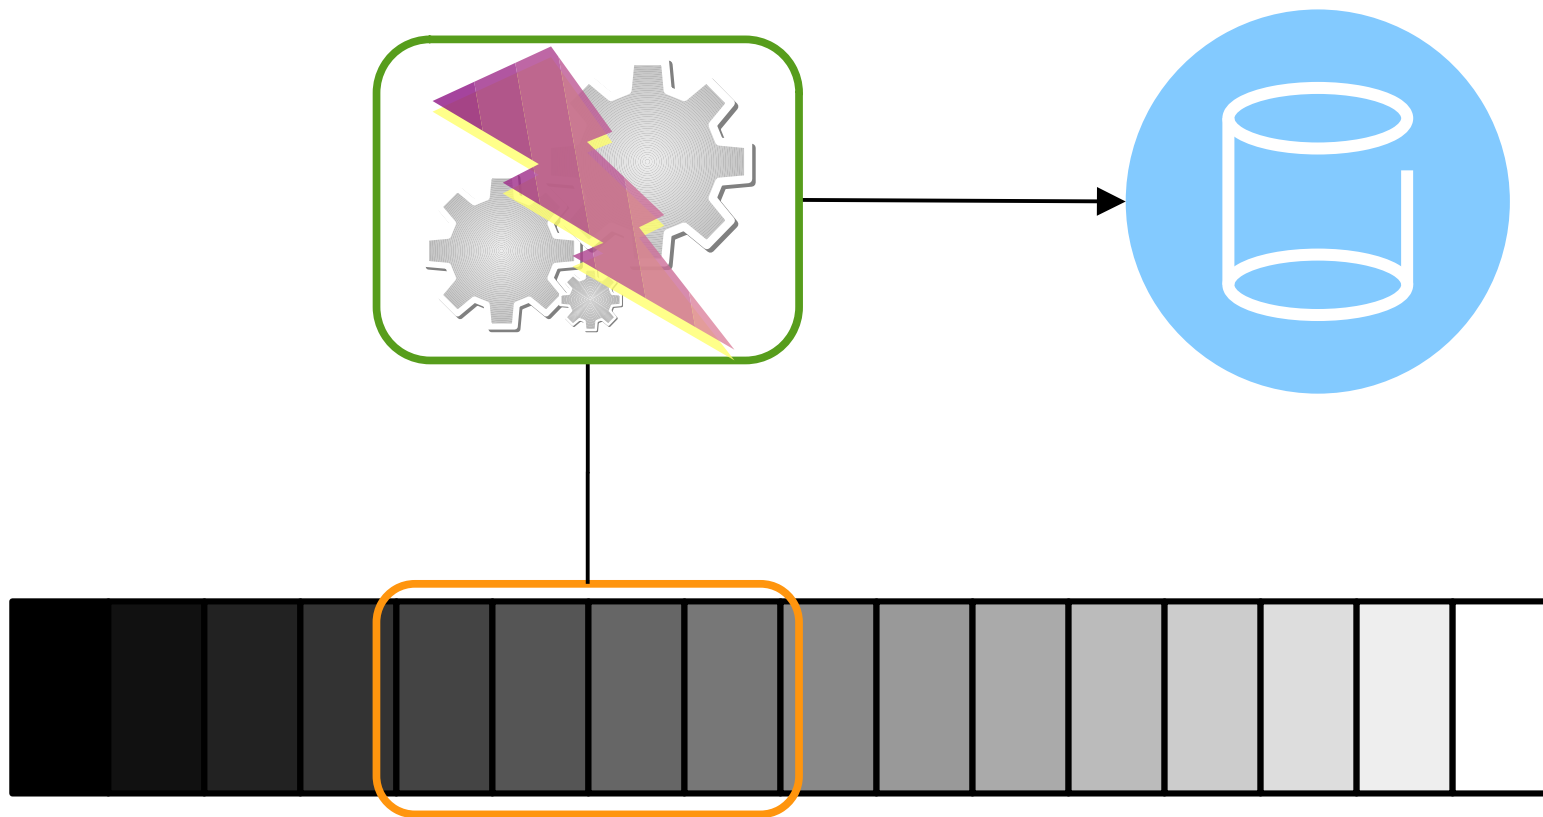






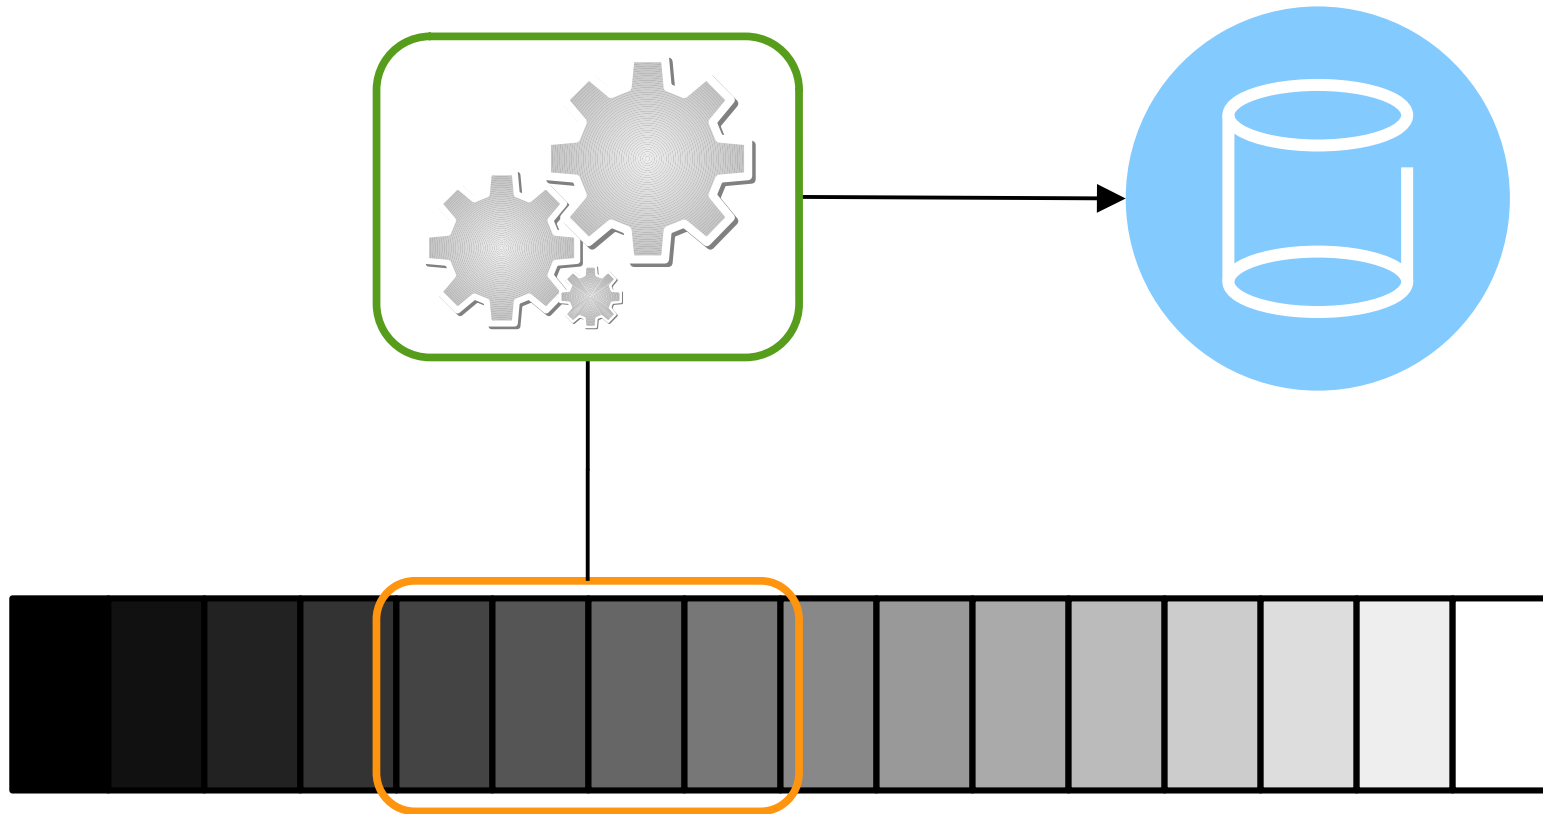
FSEvents are replayable

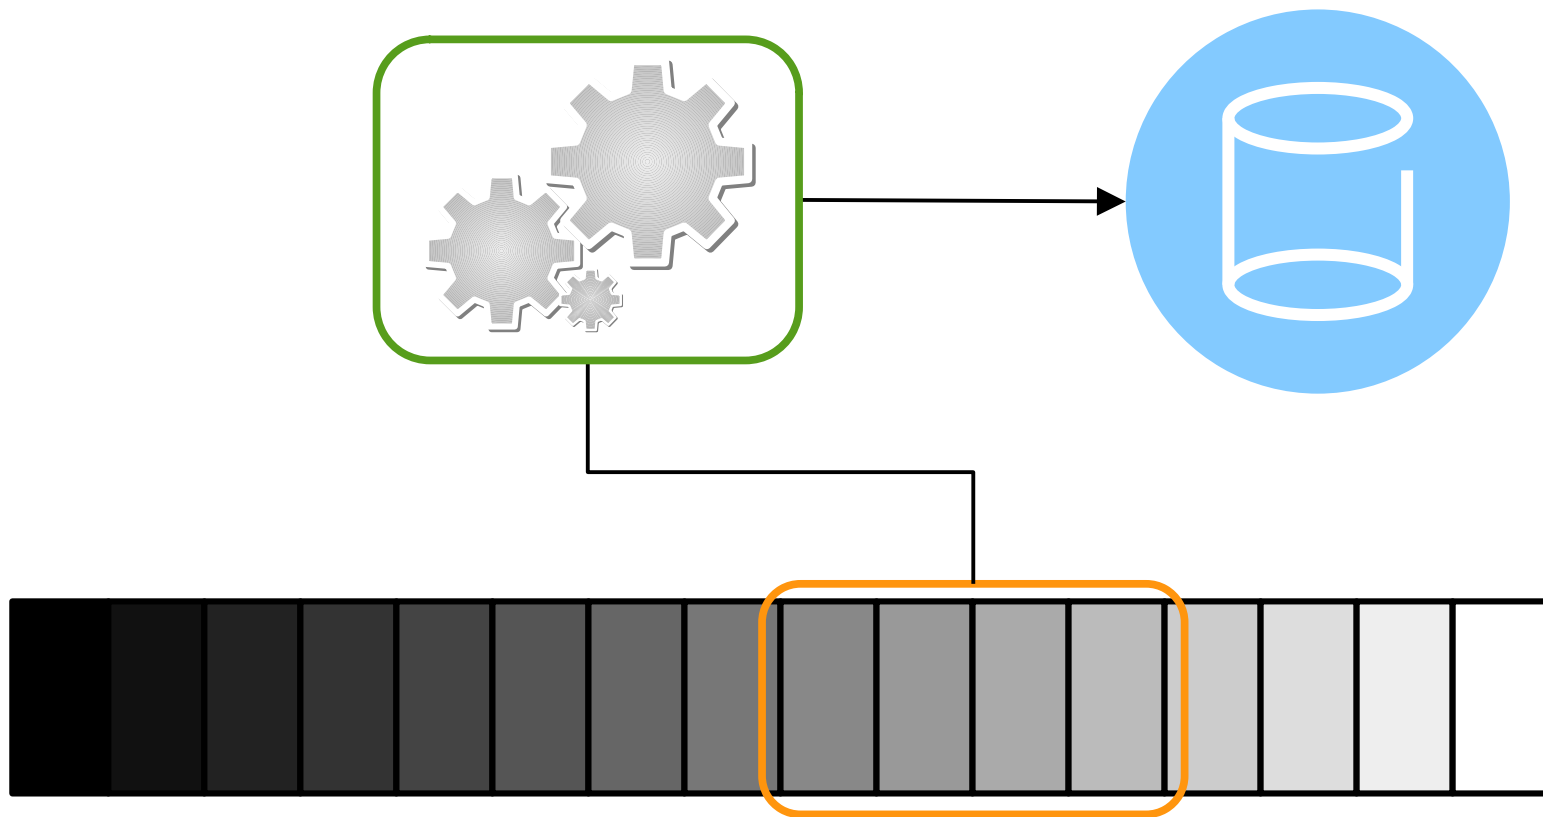









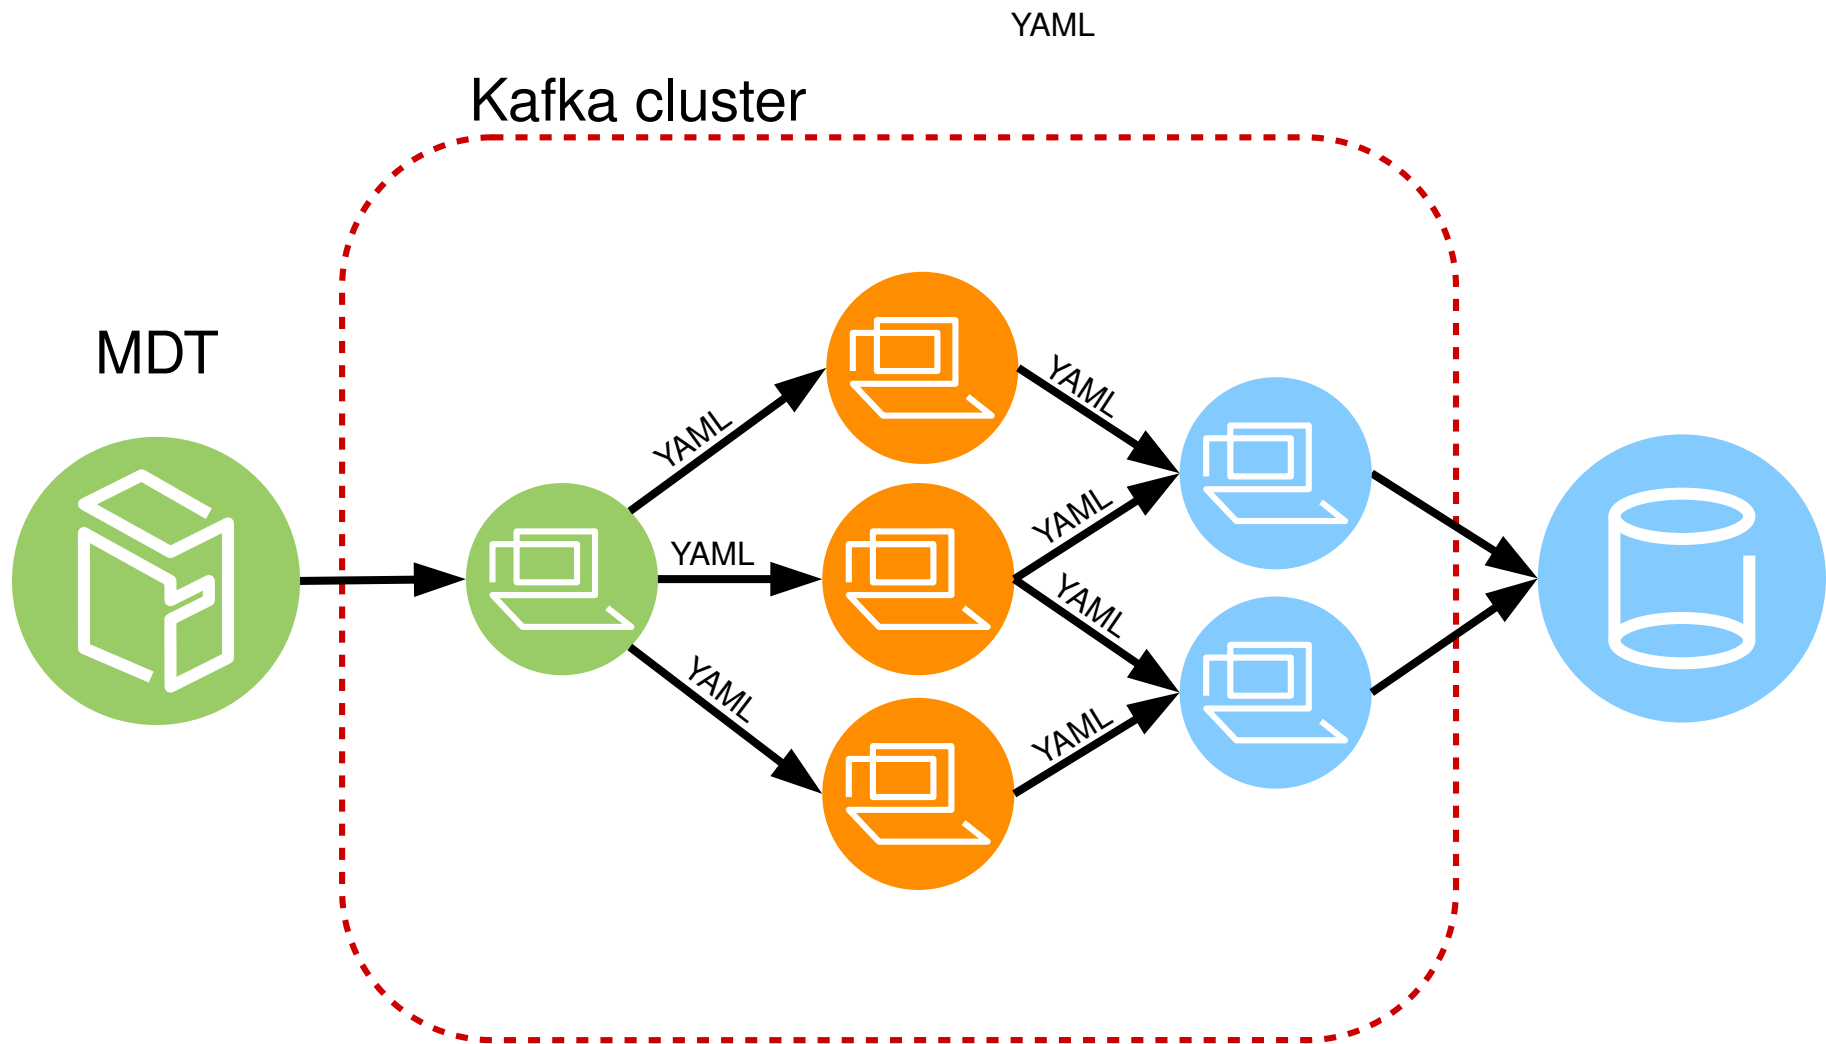




FSEvents



YAMML




<https://github.com/cea-hpc/miniyaml>

What's next?


Lustre changelog support

Lustre changelog records



FSEvents

FSEvent deduplication mechanism



Python binding

```
1 from datetime import datetime, timedelta
2
3 from robinhood.backends import MongoBackend
4 from robinhood.filter import FilterField
5
6 ninety_days_ago = datetime.now() - timedelta(days=90)
7 filter_ = FilterField.STATX.ATIME >= ninety_days_ago.timestamp()
8
9 for fsentry in Backend("my-filesystem").filter(filter_):
10     fsentry.unlink()
```

Demo time!



Thank you