

How not to SQLite

Ulf Tigerstedt

Delegation DB was slow

- Probably me suggested ditching BerkeleyDB for SQLite 2 years ago

Delegation DB was slow

- Probably me suggested ditching BerkeleyDB for SQLite 2 years ago
- Problems eased a bit, but not much.

Delegation DB was slow

- Probably me suggested ditching BerkeleyDB for SQLite 2 years ago
- Problems eased a bit, but not much.
- Migration in production was a problem, since all jobs had to be drained

Nosql used to be the hype, not doing BDB in SQL

- ```
CREATE TABLE IF NOT EXISTS
rec(id, owner, uid, meta,
UNIQUE(id, owner), UNIQUE(uid))
```

## Nosql used to be the hype, not doing BDB in SQL

- `CREATE TABLE IF NOT EXISTS  
rec(id, owner, uid, meta,  
UNIQUE(id, owner), UNIQUE(uid))`
- No types, so everything will be chars

## Nosql used to be the hype, not doing BDB in SQL

- `CREATE TABLE IF NOT EXISTS  
rec(id, owner, uid, meta,  
UNIQUE(id, owner), UNIQUE(uid))`
- No types, so everything will be chars
- No documentation

id, uid

- `uid = rand_uid64().substr(4)`



id, uid

- `uid = rand_uid64().substr(4)`
- This gives no guarantee of uniqueness, so there are loops implemented to care for this

id, uid

- `uid = rand_uid64().substr(4)`
- This gives no guarantee of uniqueness, so there are loops implemented to care for this
- It also eats randomness

## id, uid

- `uid = rand_uid64().substr(4)`
- This gives no guarantee of uniqueness, so there are loops implemented to care for this
- It also eats randomness
- Should be an `AUTOINCREMENT`

## Suggestion

- Make the schema sane

## Suggestion

- Make the schema sane
- Make a migration script

## Suggestion

- Make the schema sane
- Make a migration script
- Do the changes to use the new schema in the code