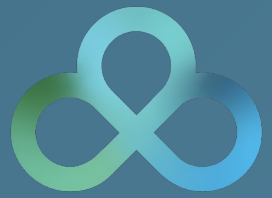

EDB
Postgres for the AI Generation

Understanding Consistency in PostgreSQL Replication

Boriss Mejías
Holistic System Software Engineer
Air Guitar Player
5th of June, 2024

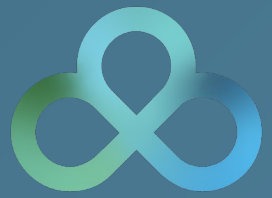

EDB
Postgres for the AI Generation

Understanding Consistency in PostgreSQL Replication

Boriss Mejías
Solutions Architect
Air Guitar Player
5th of June, 2024

Eventual Consistency

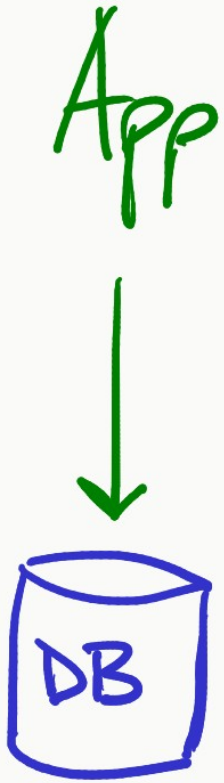
App
DB

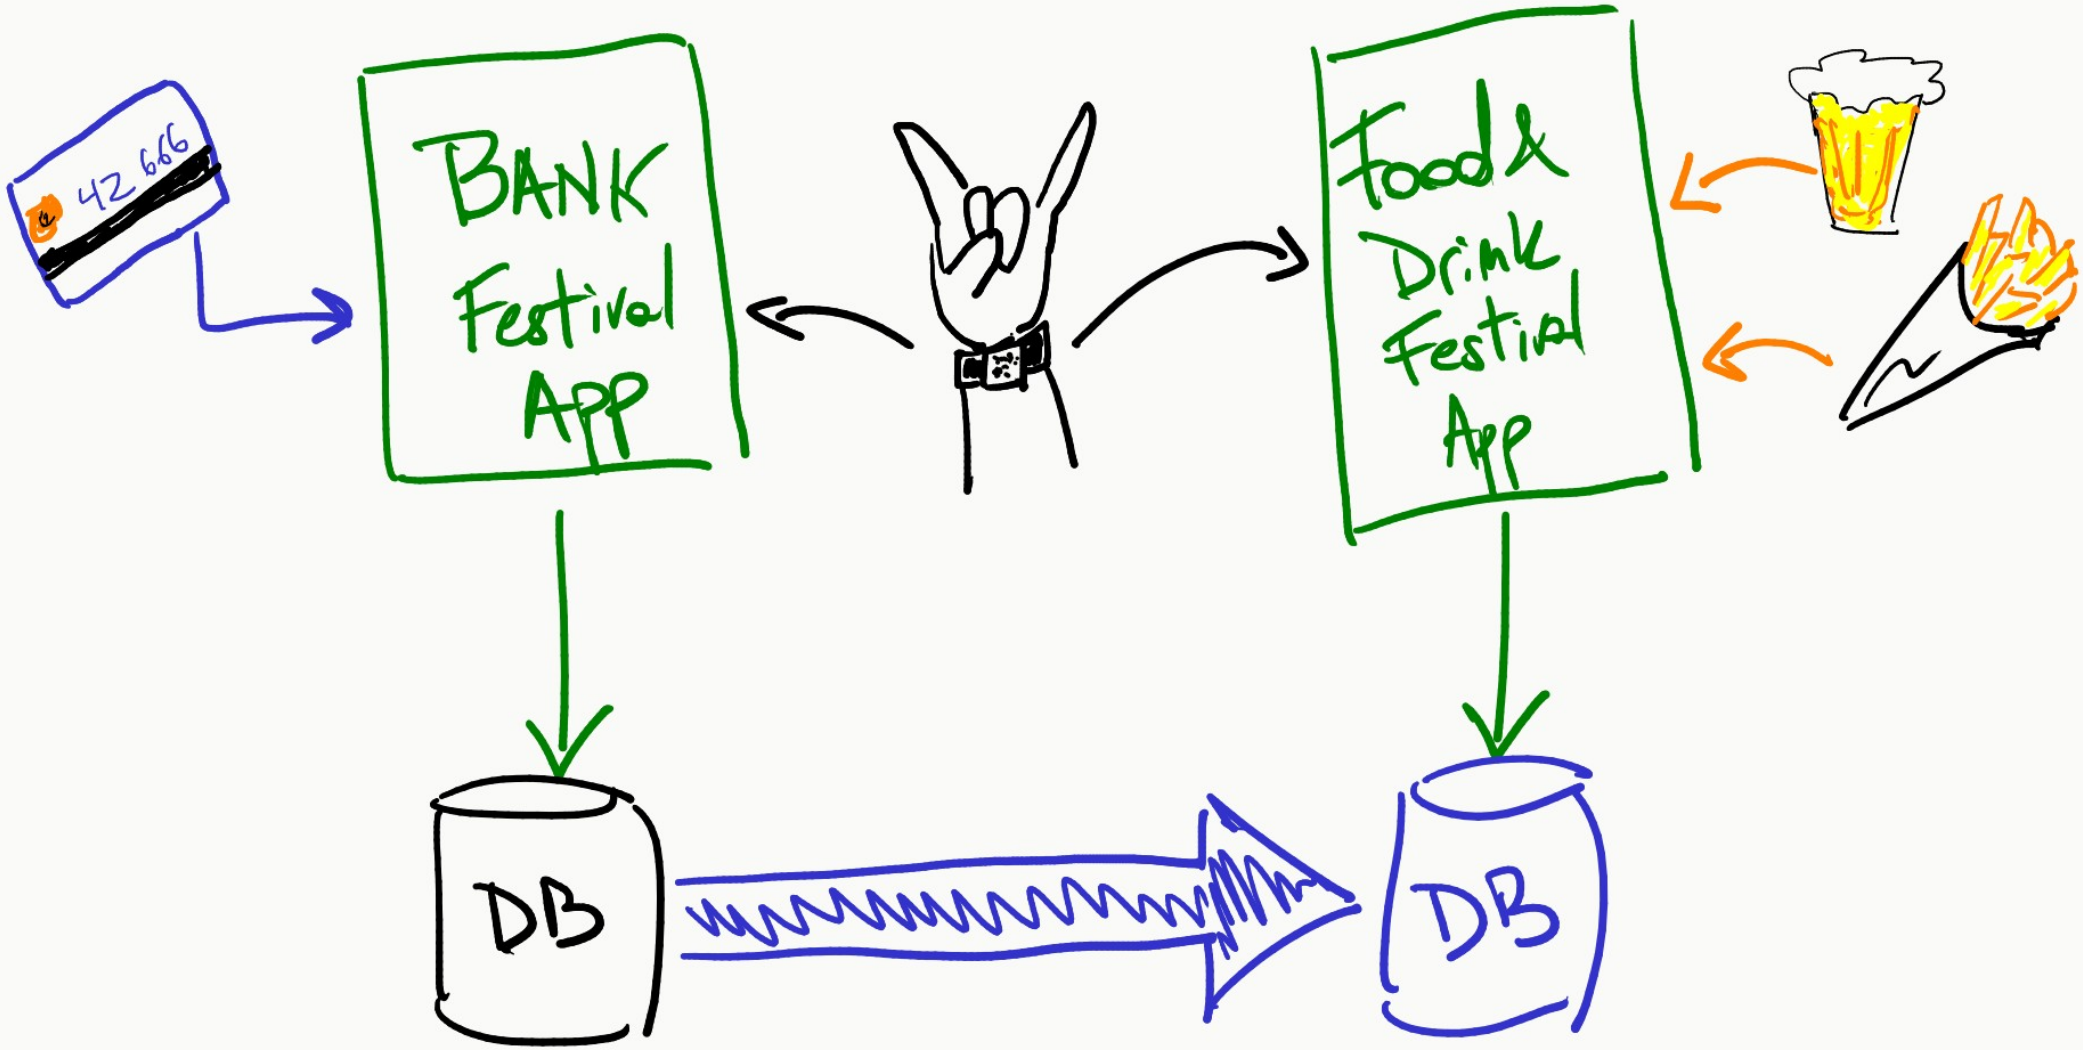
42 666
BANK
Festival
App
Food &
Drink
Festival
App
DB
DB

Consistency

Every read receives the most recent write or an error

Replication Basics

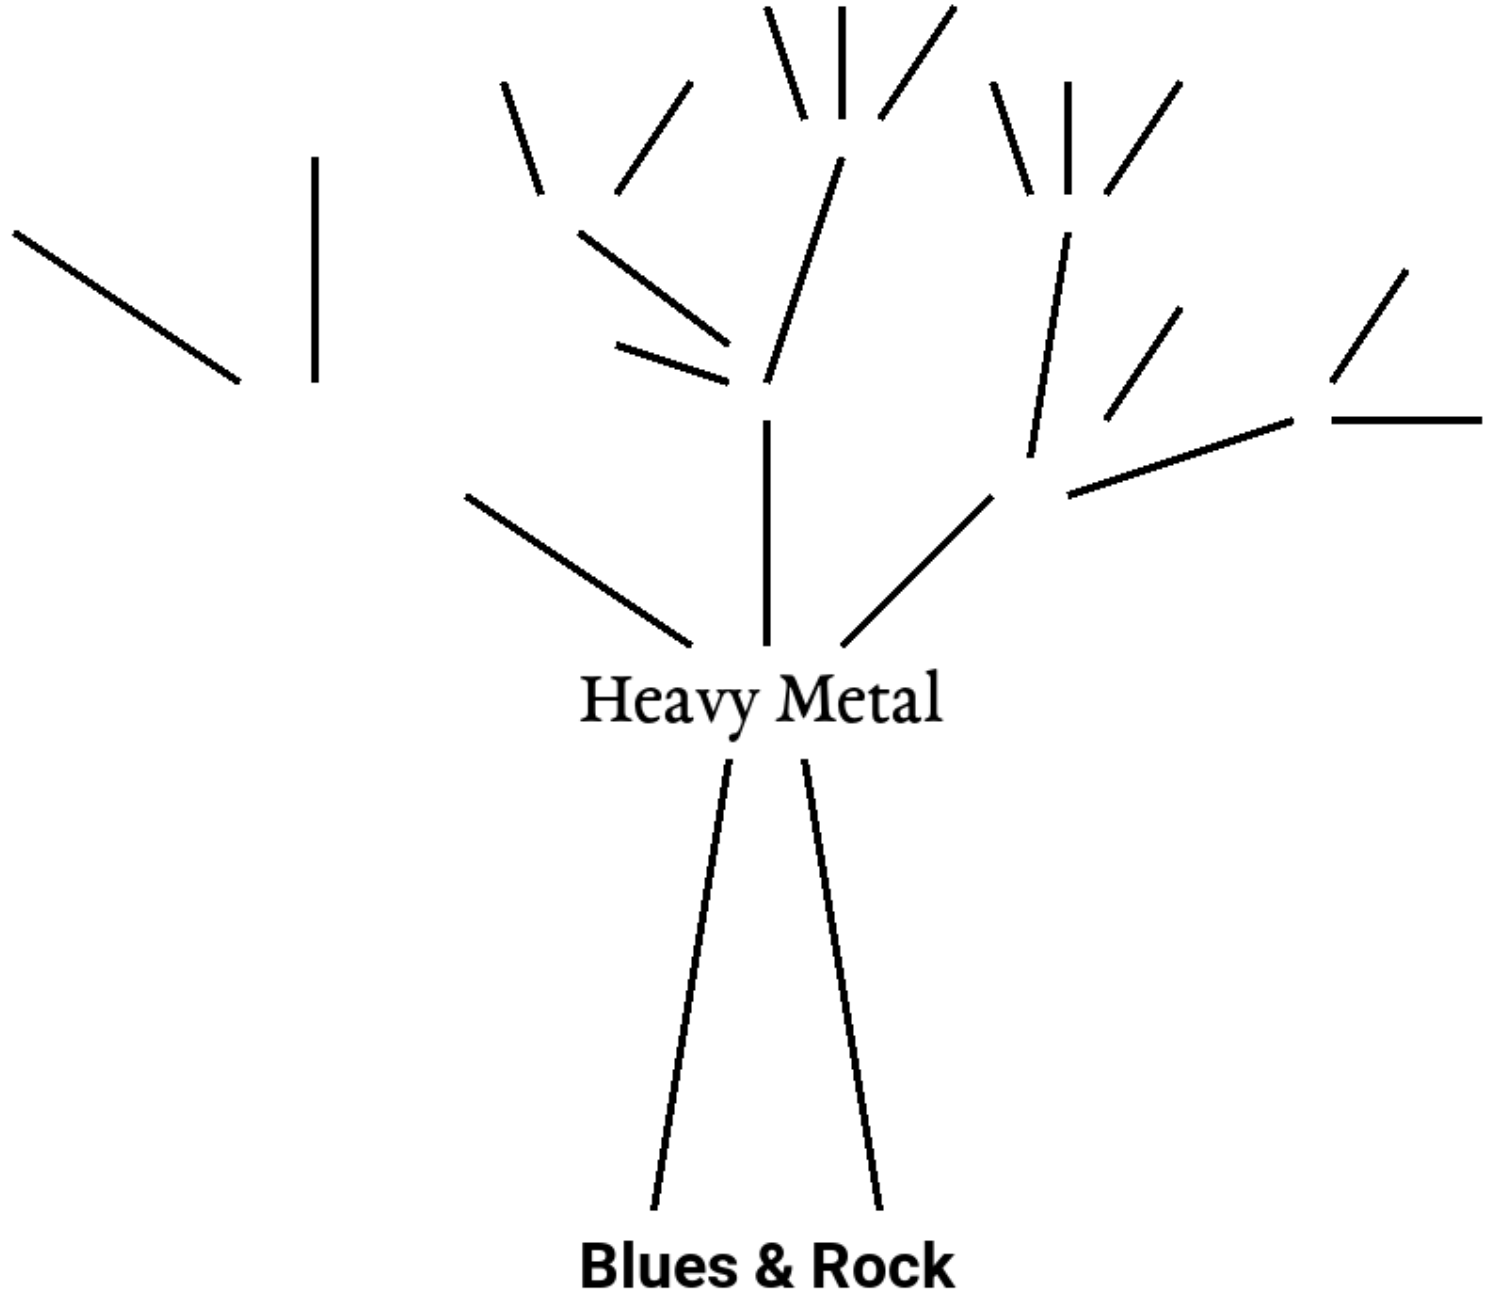
Heavy Metal
Blues & Rock

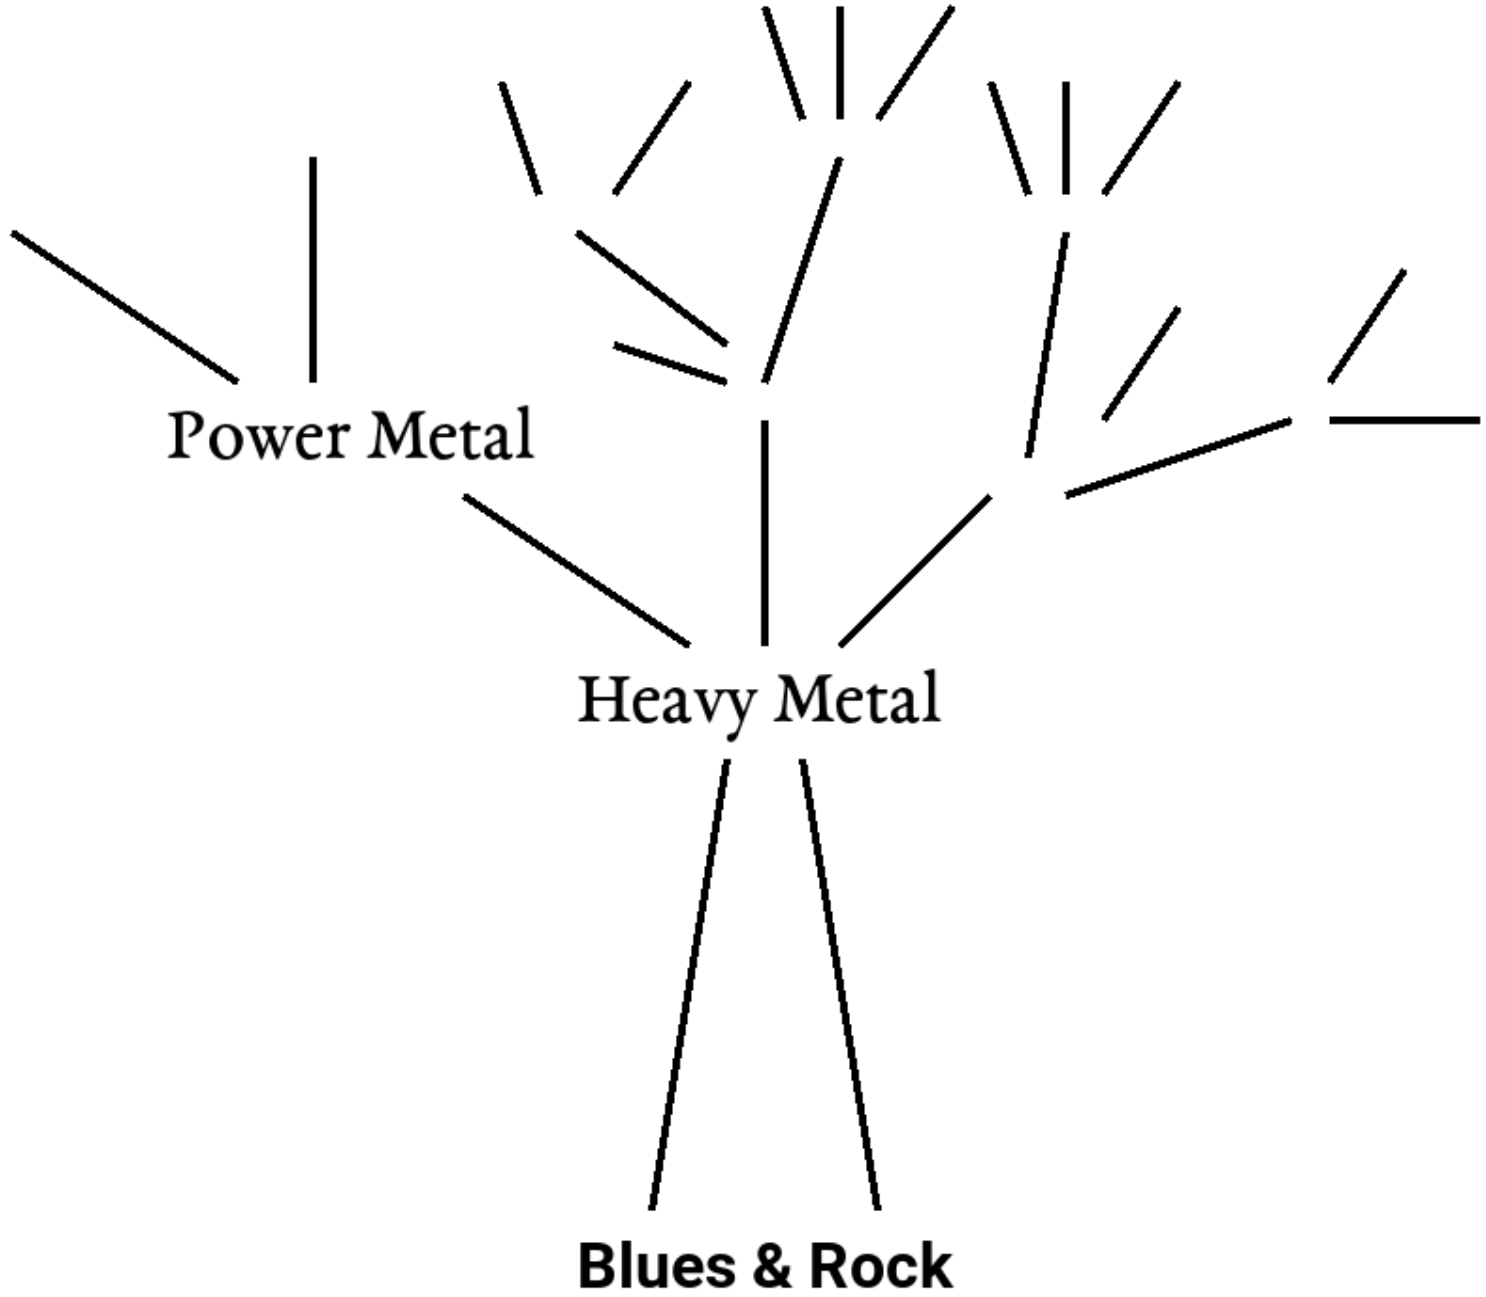
Power Metal
Heavy Metal
Blues & Rock

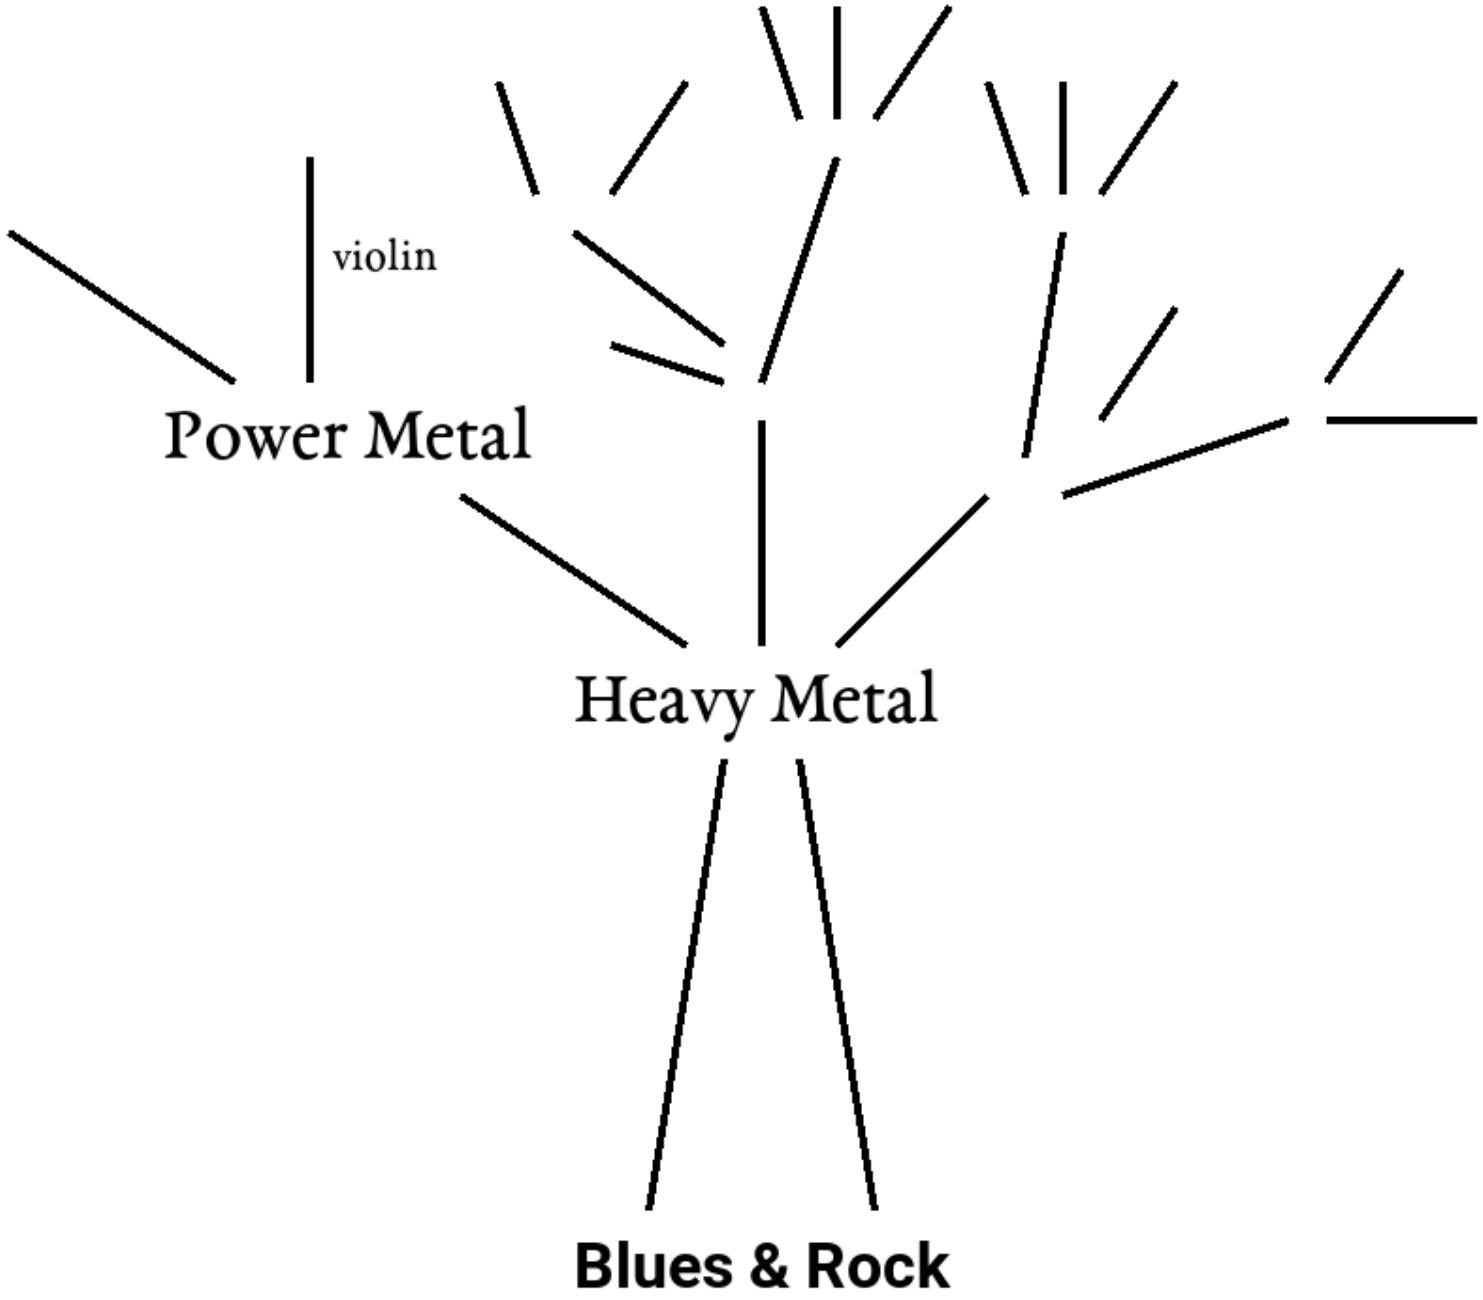
violin
Power Metal
Heavy Metal
Blues & Rock

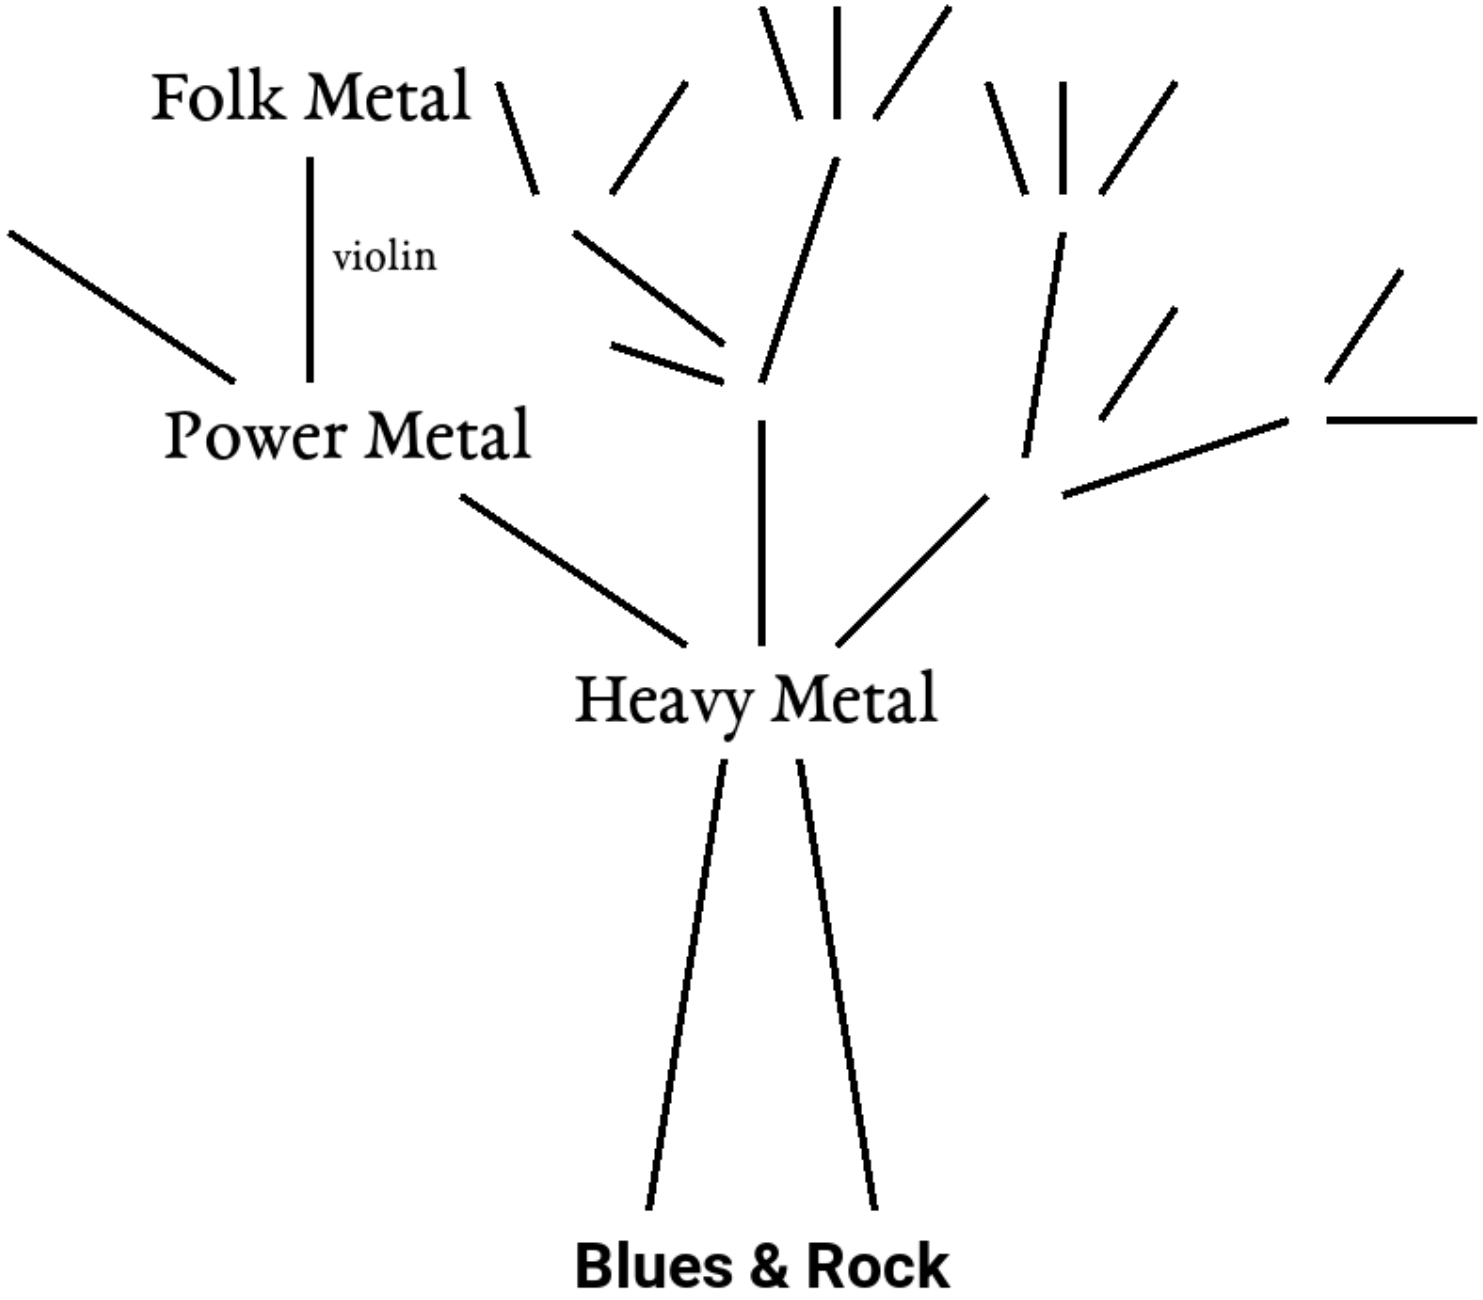
Folk Metal
violin
Power Metal
Heavy Metal
Blues & Rock

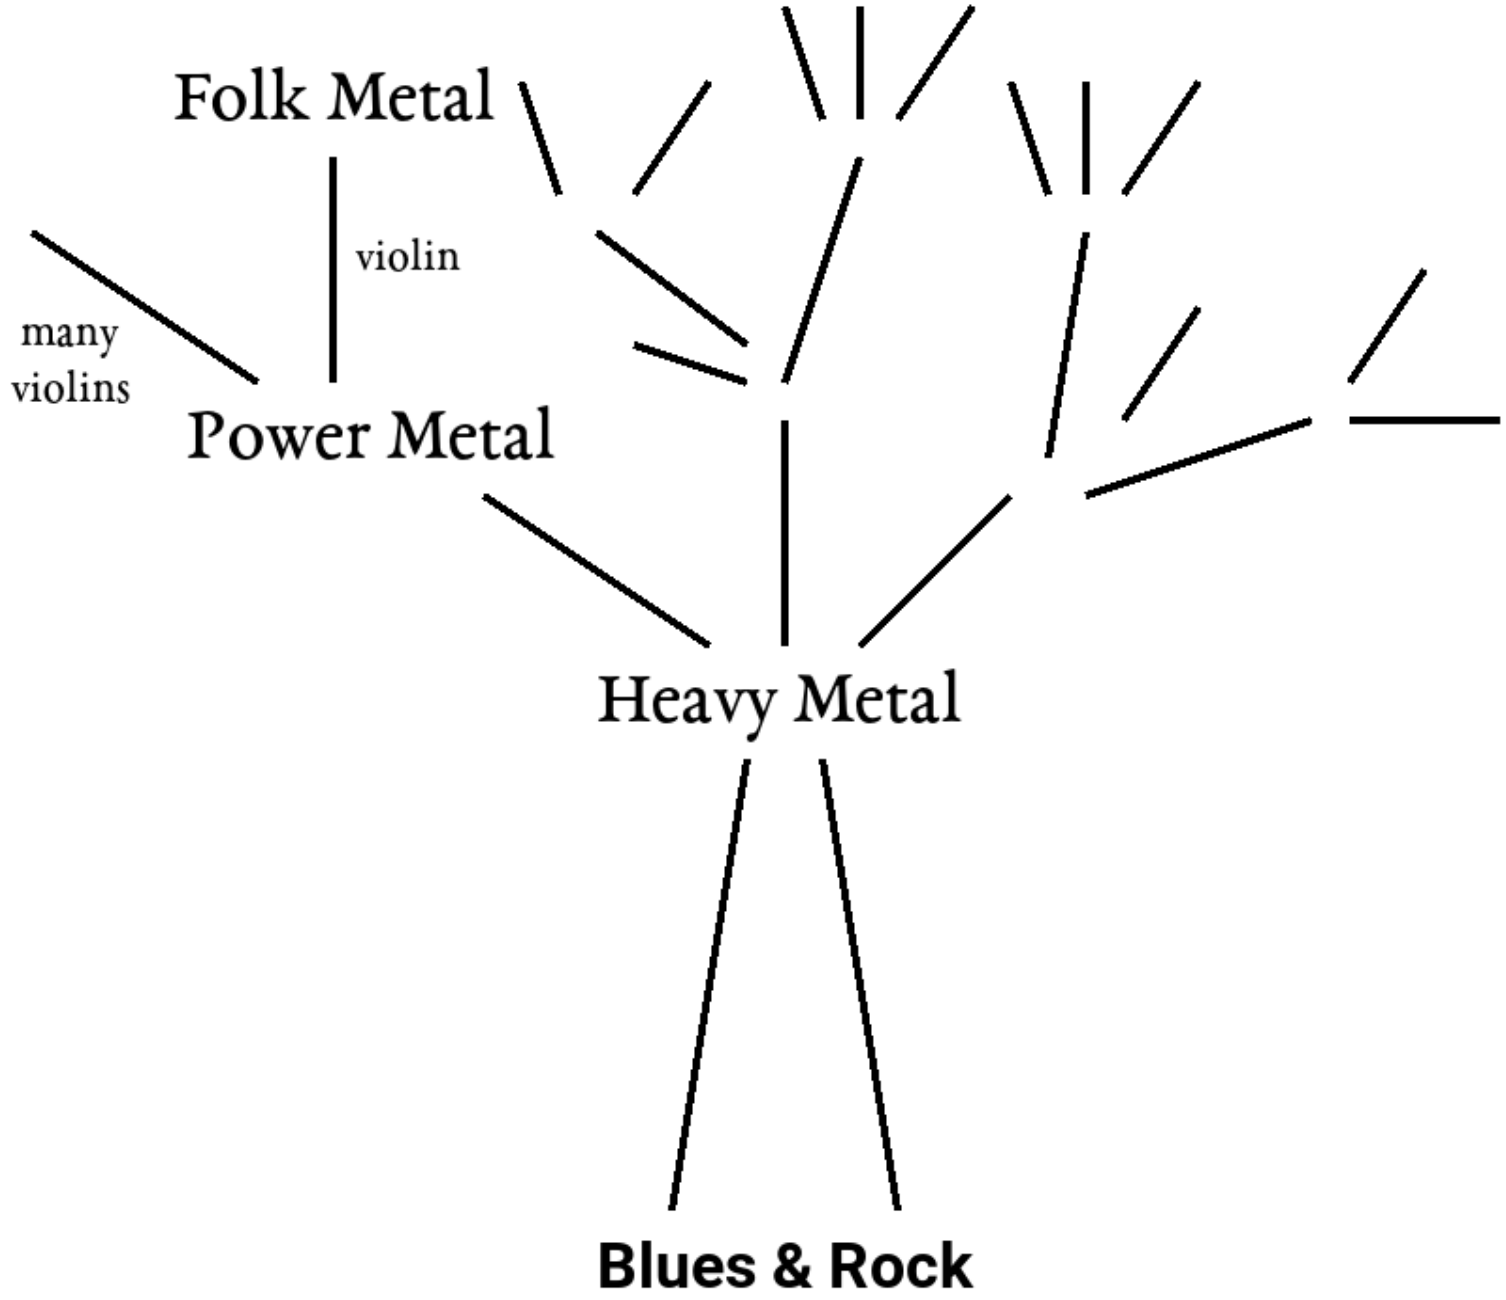
Folk Metal
violin
many
violins
Power Metal
Heavy Metal
Blues & Rock

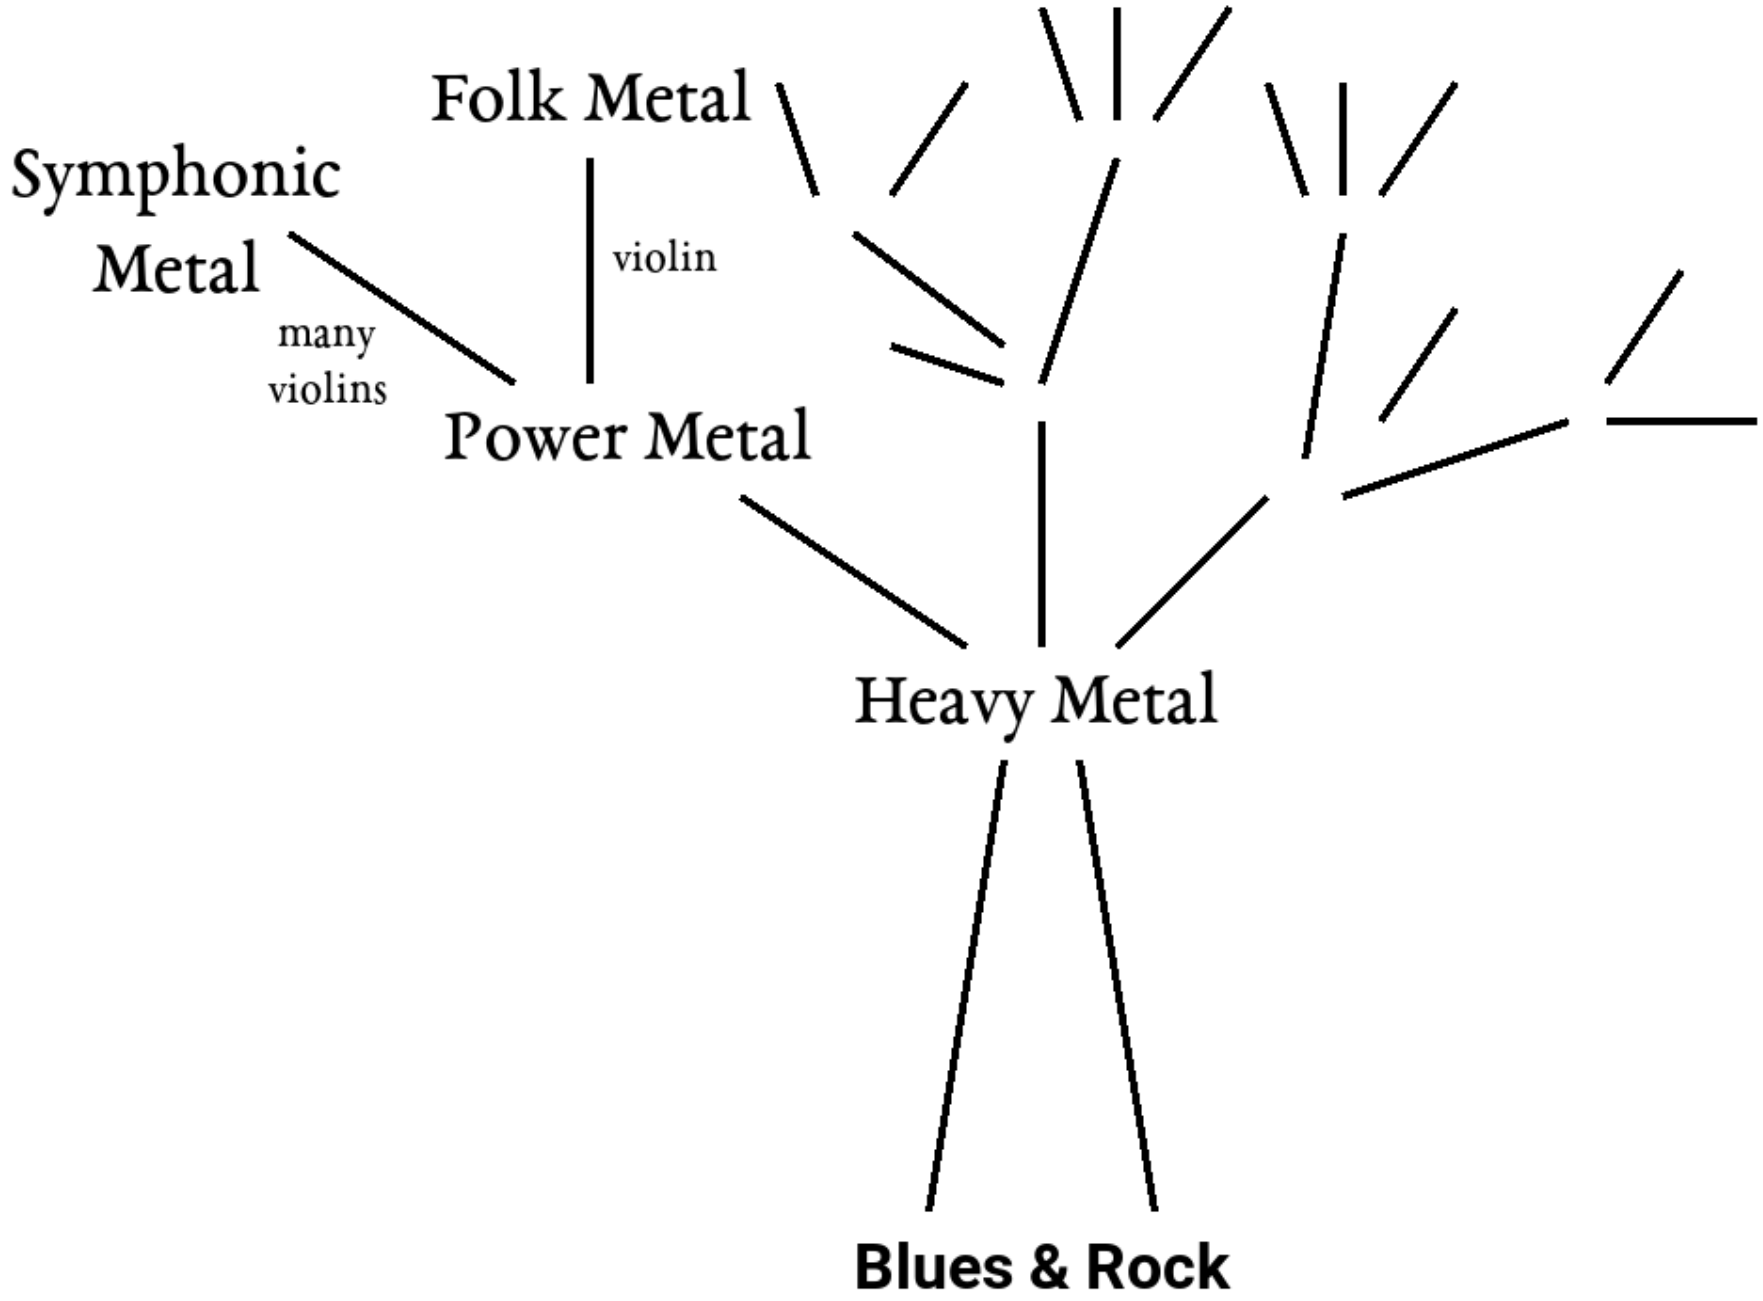
Folk Metal
Symphonic Metal
violin
many violins
Power Metal
Heavy Metal
Blues & Rock

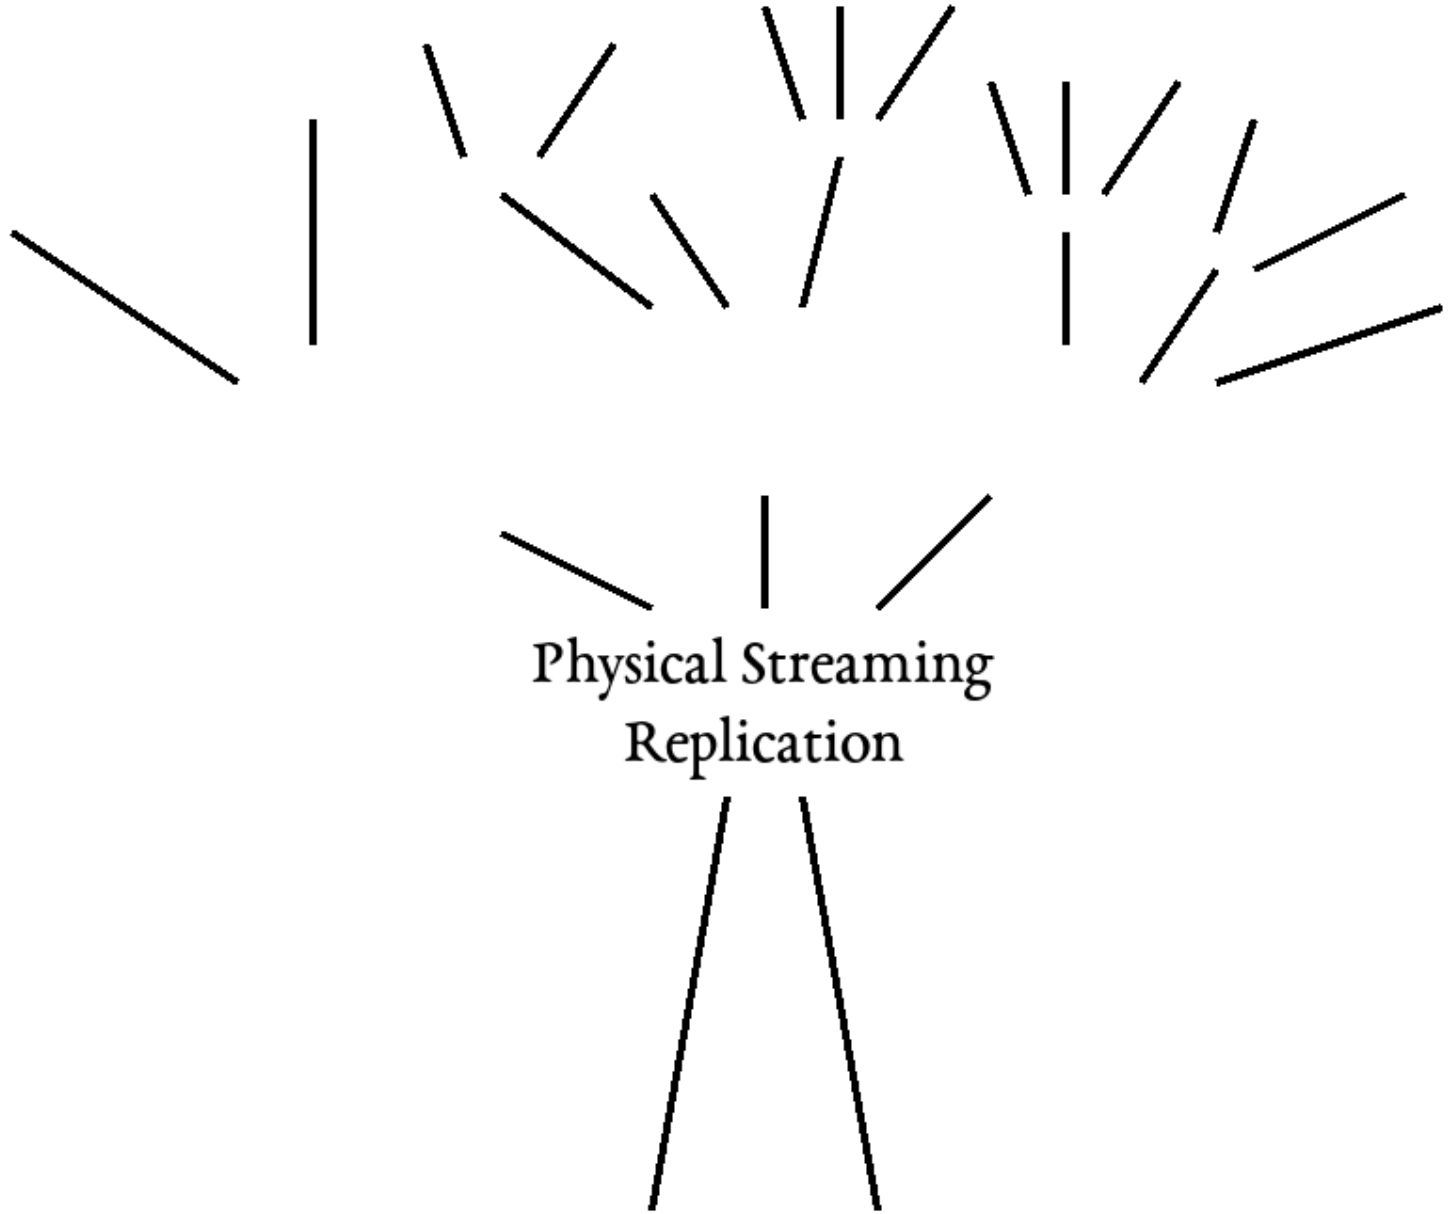
Physical Streaming
Replication
Backup & Recovery

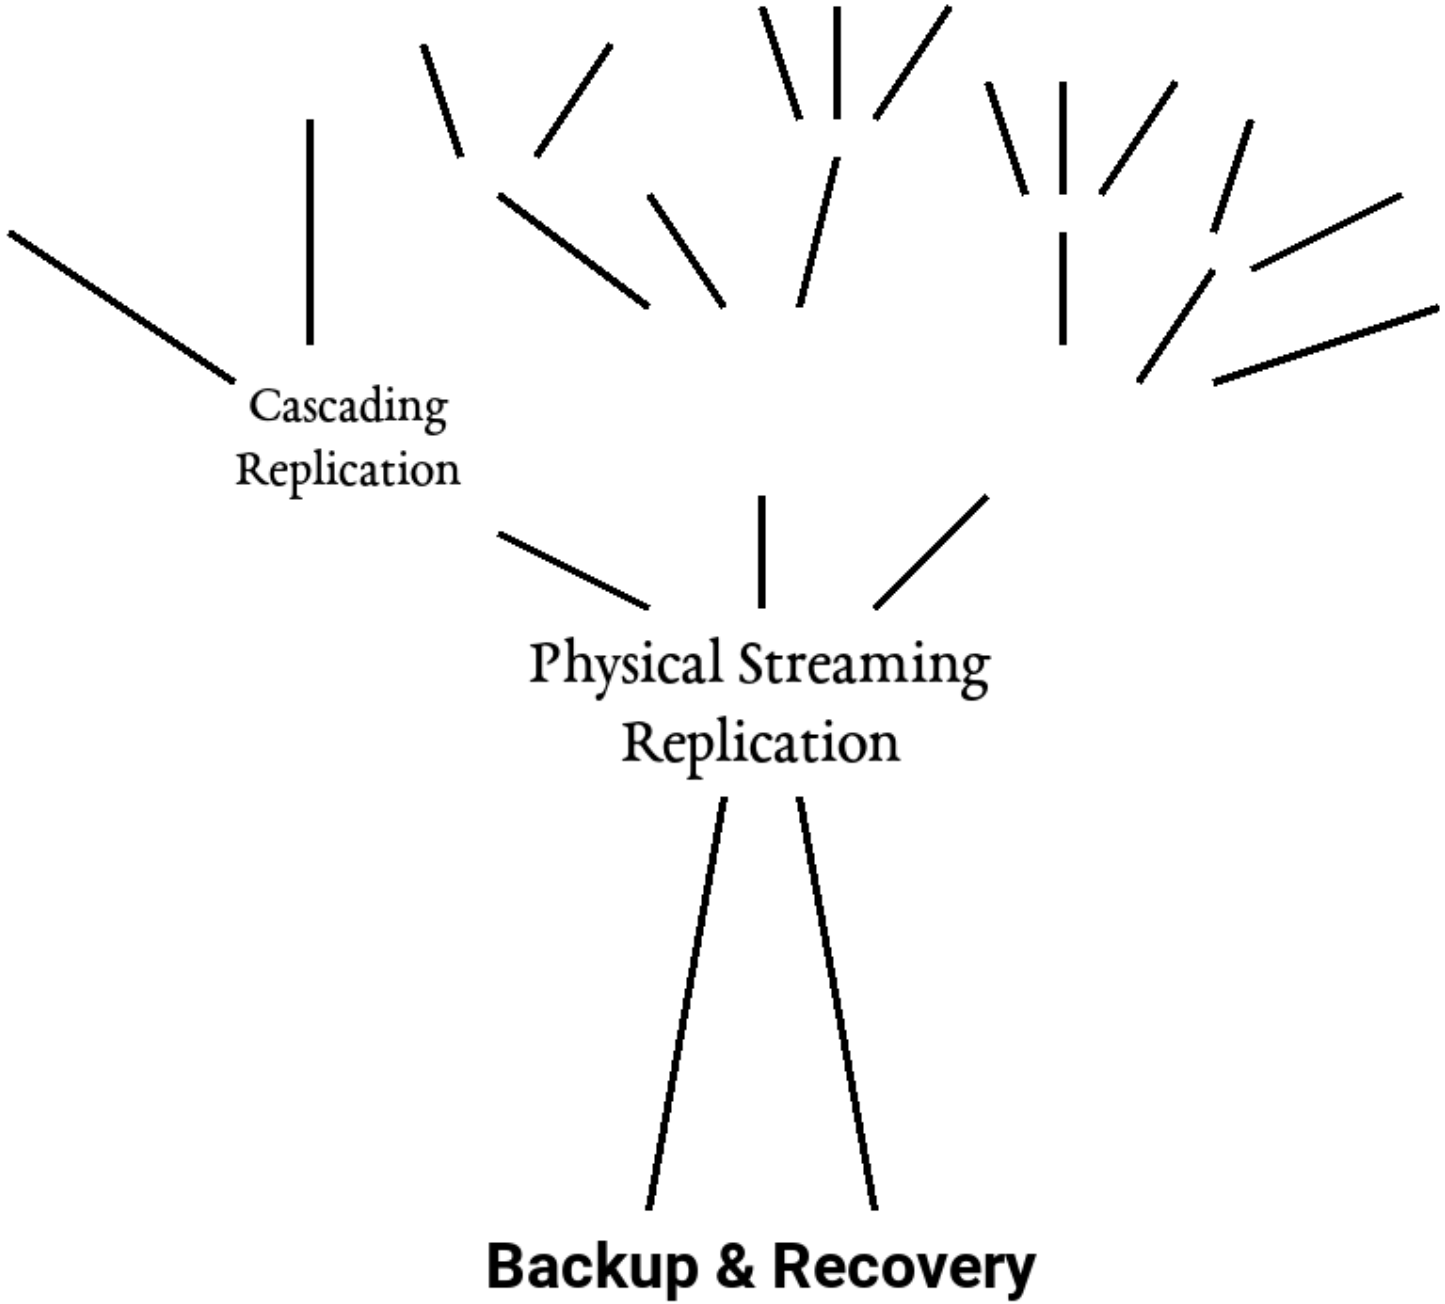
Cascading
Replication
Physical Streaming
Replication
Backup & Recovery

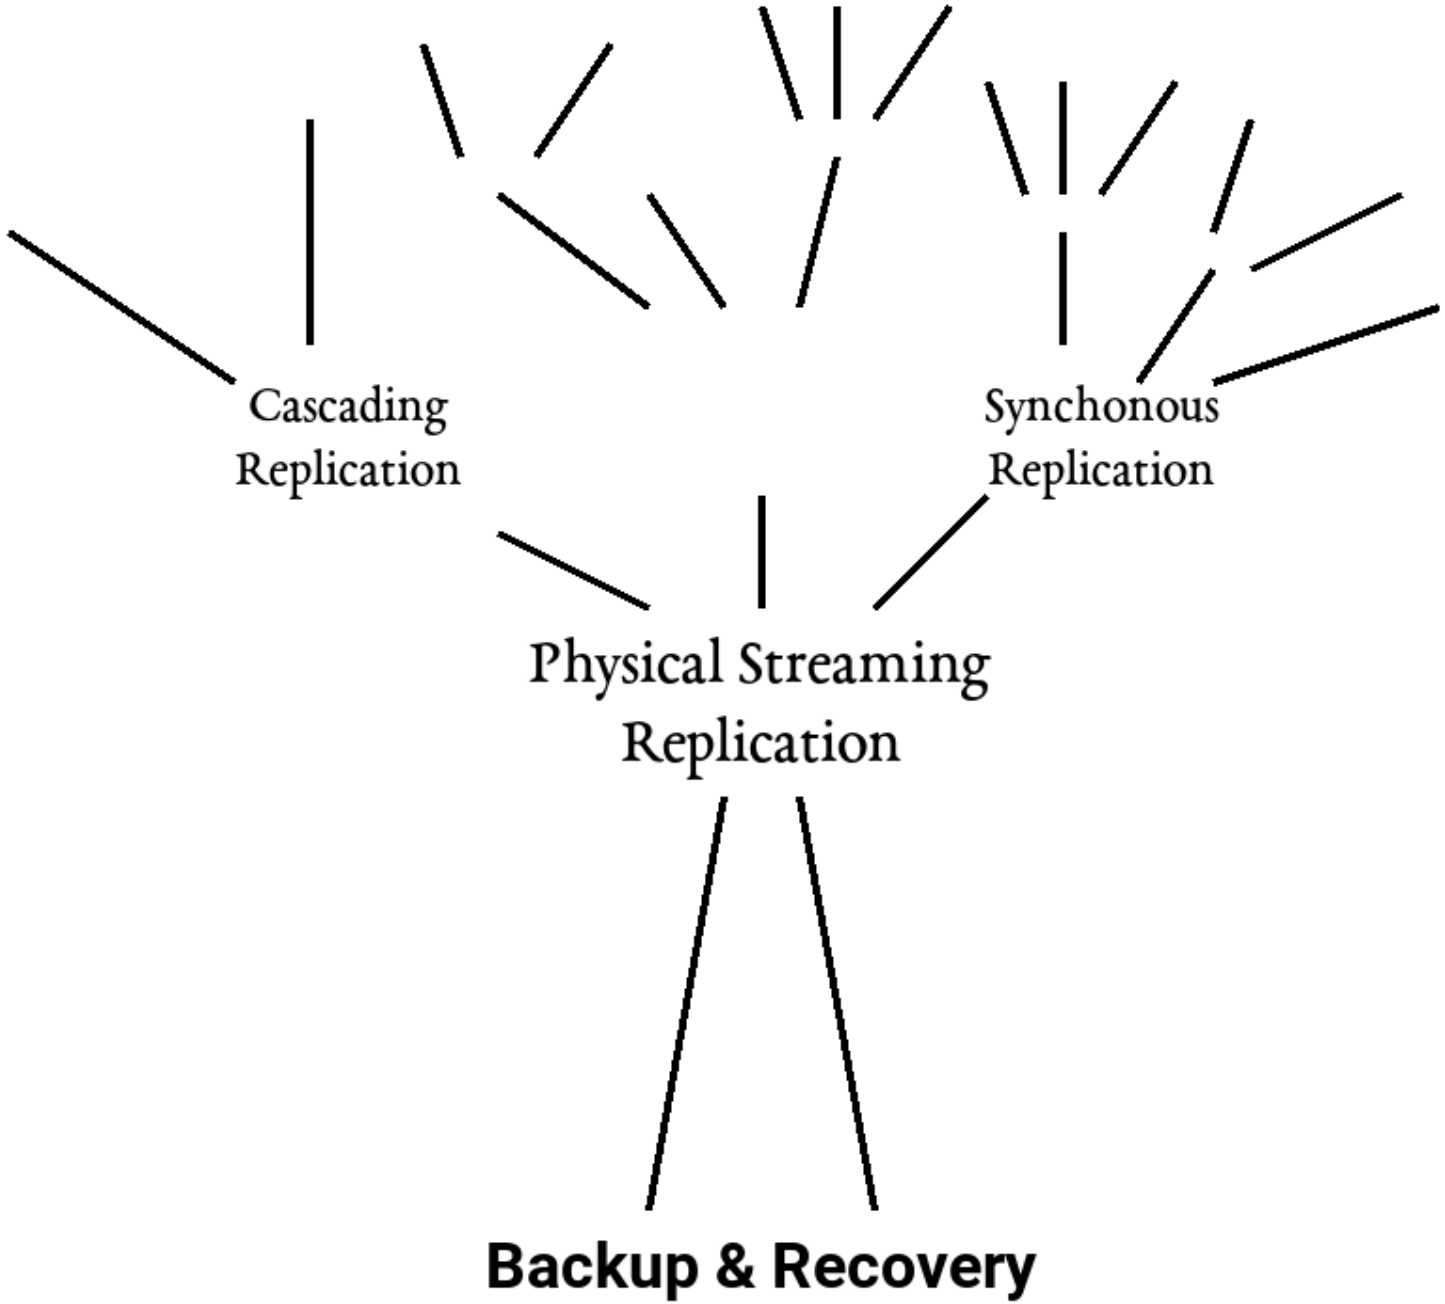
Cascading
Replication
Synchonous
Replication
Physical Streaming
Replication
Backup & Recovery

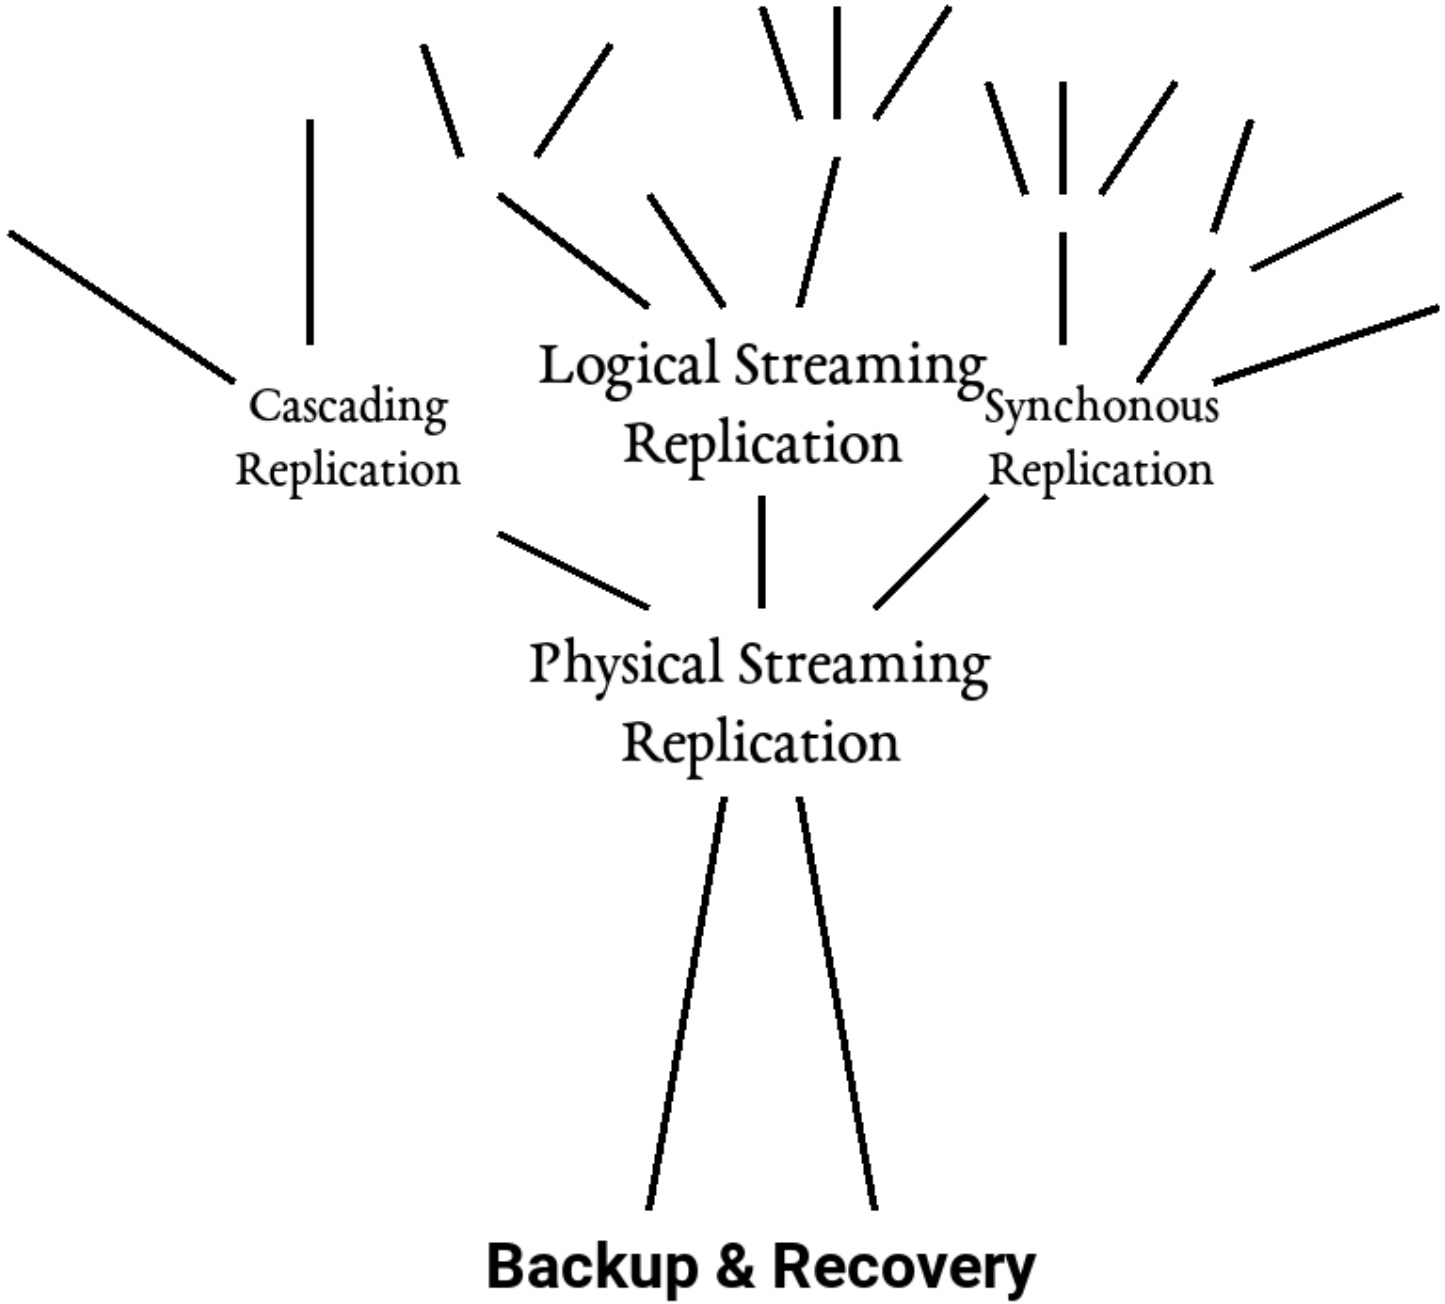
Cascading Replication
Logical Streaming Replication
Synchonous Replication
Physical Streaming Replication
Backup & Recovery

Table to store metalhead and balance

Table to store metalhead and balance

```
CREATE TABLE cashless (
     id_metalhead BIGSERIAL      PRIMARY KEY
   , balance      NUMERIC(5, 2) NOT NULL
);
```

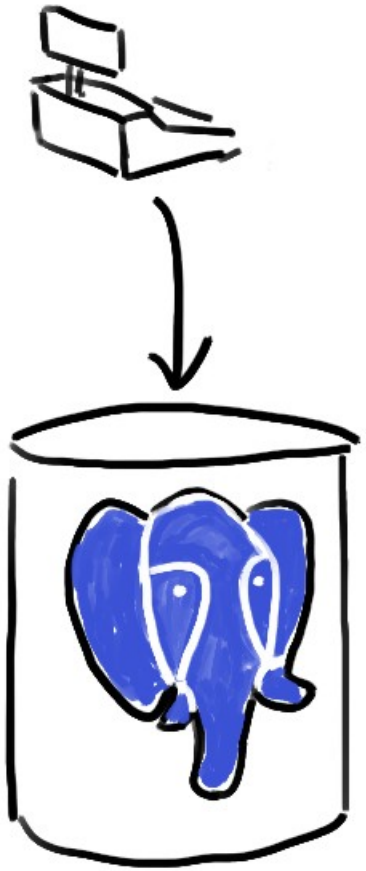

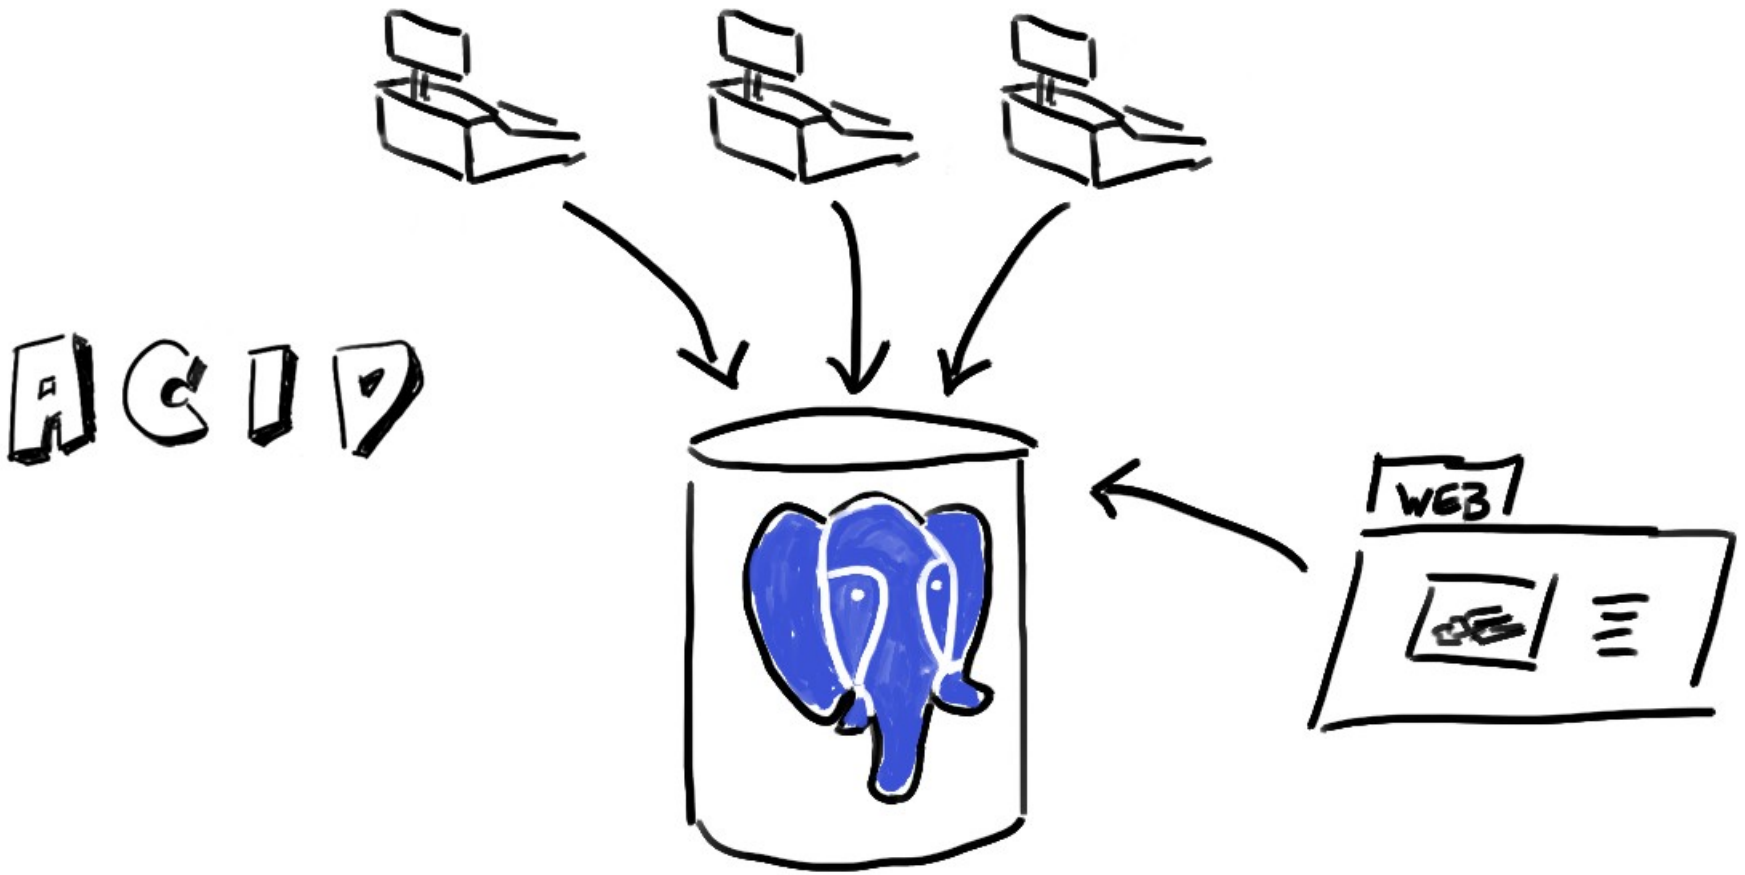

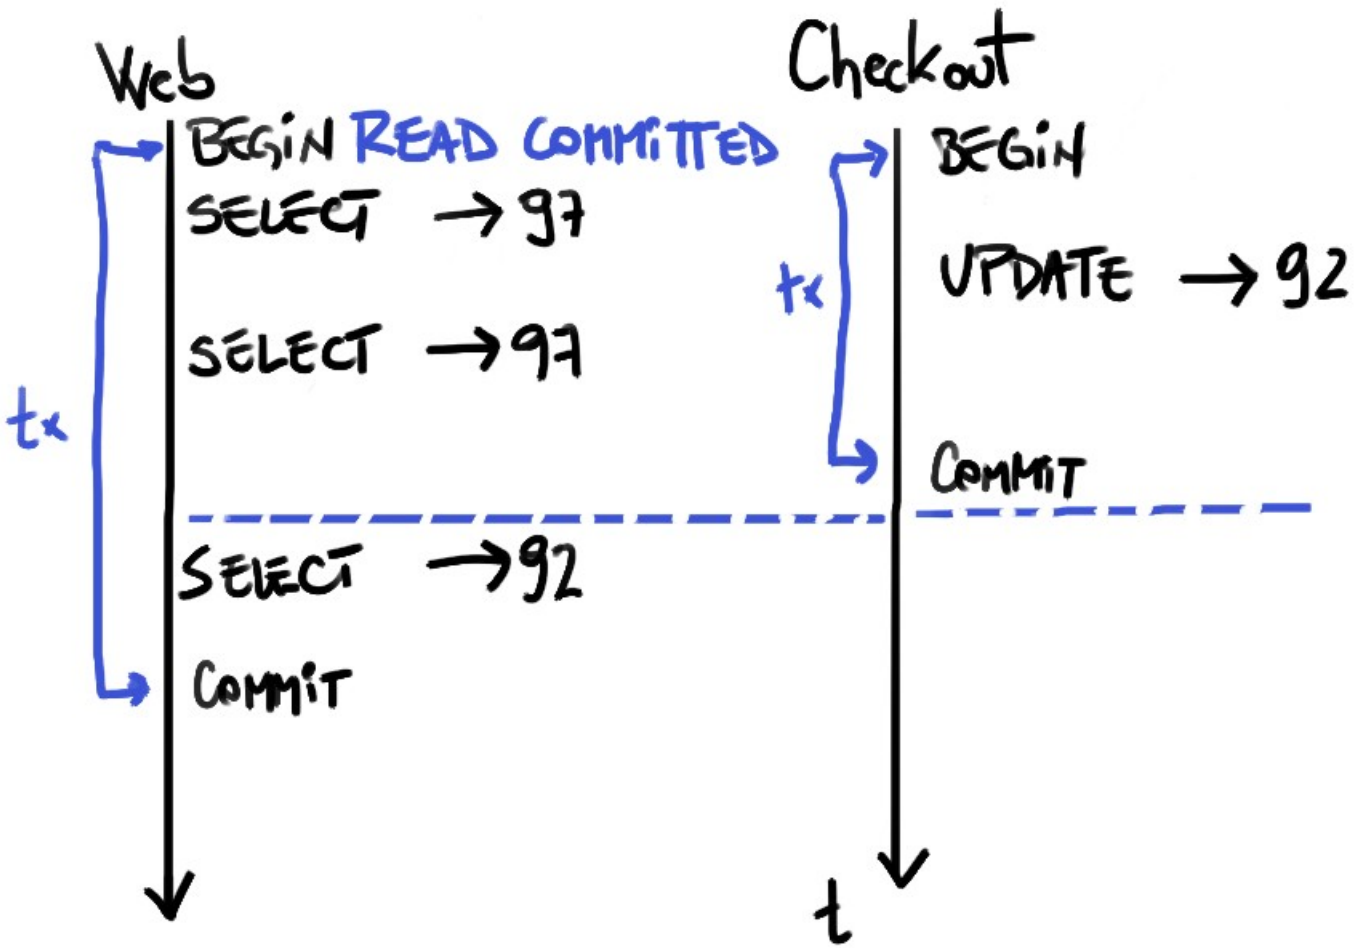

Consistency

Every read receives the most recent write or an error

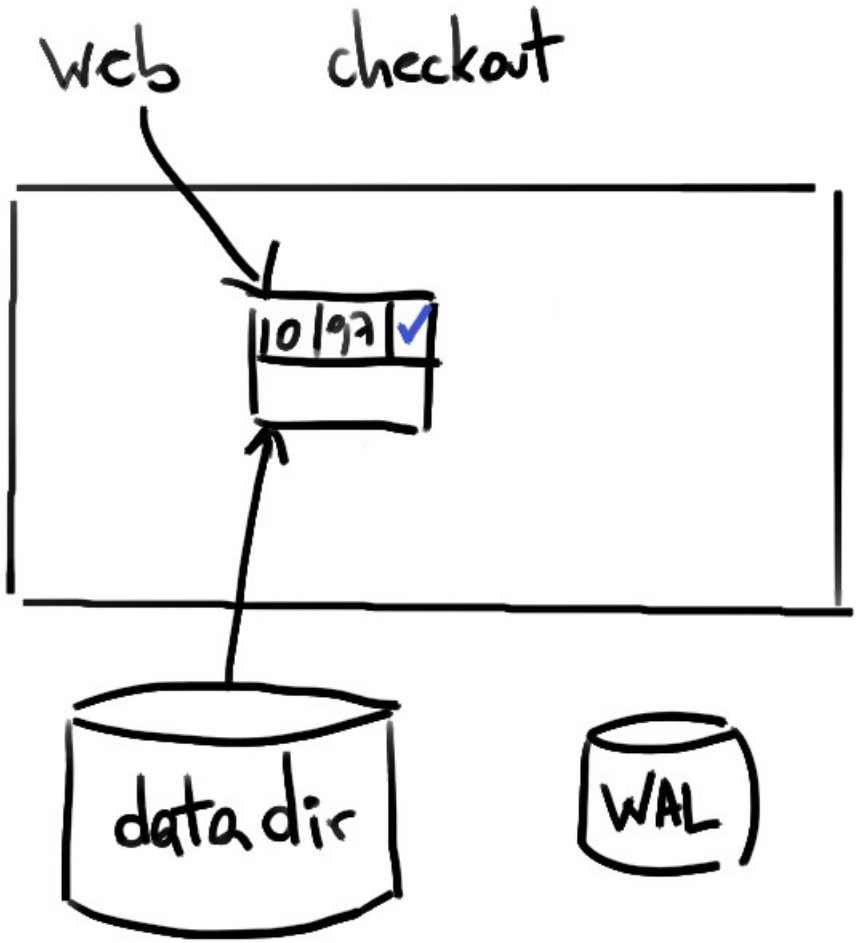

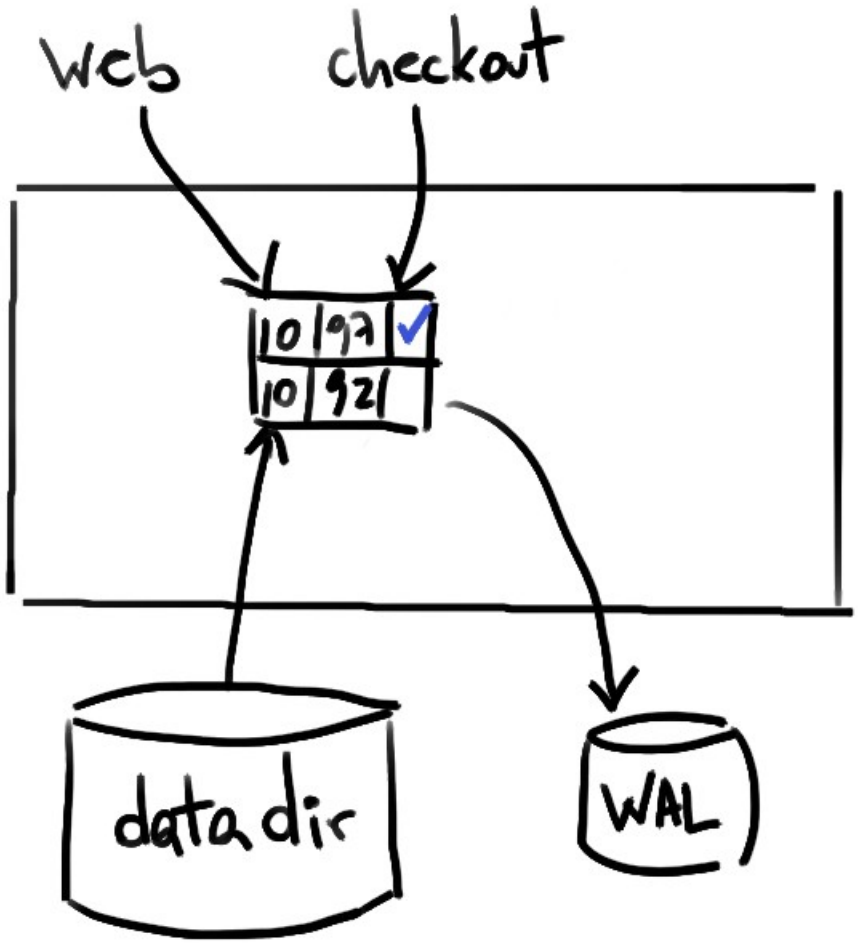

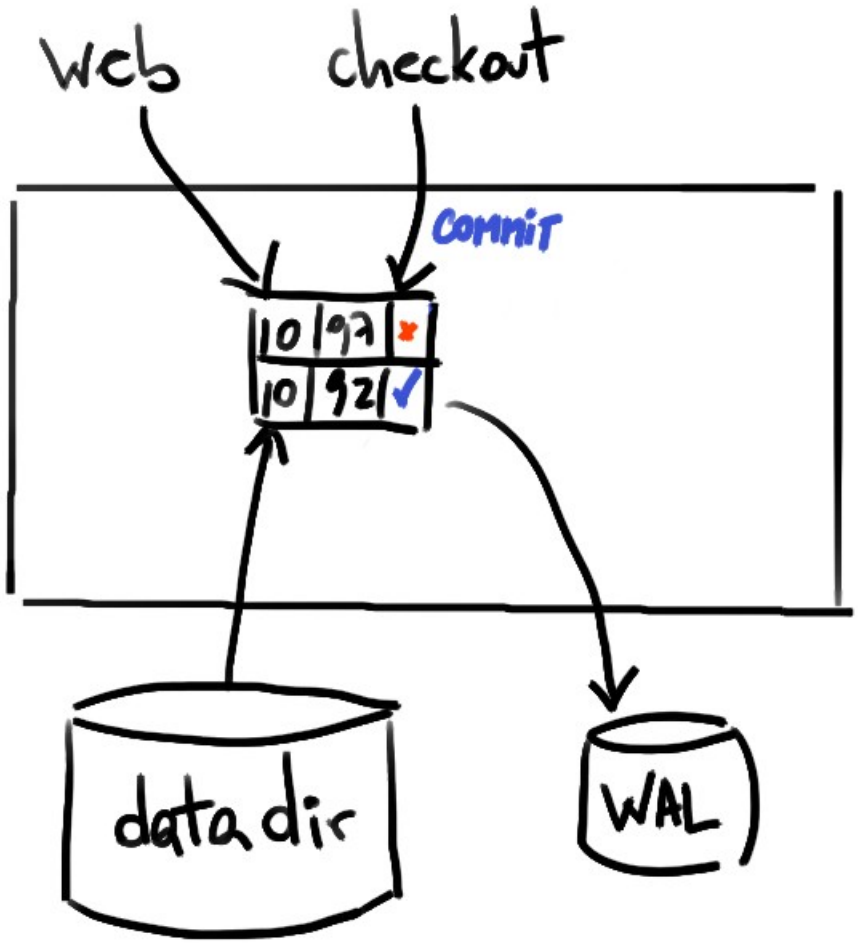

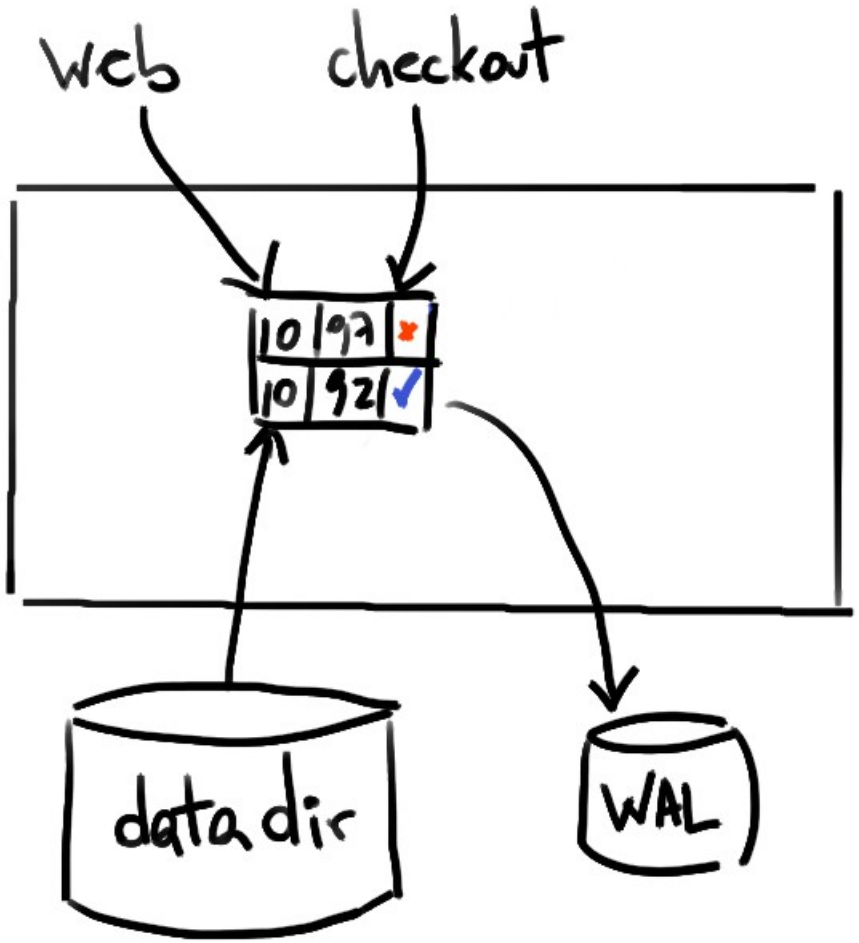

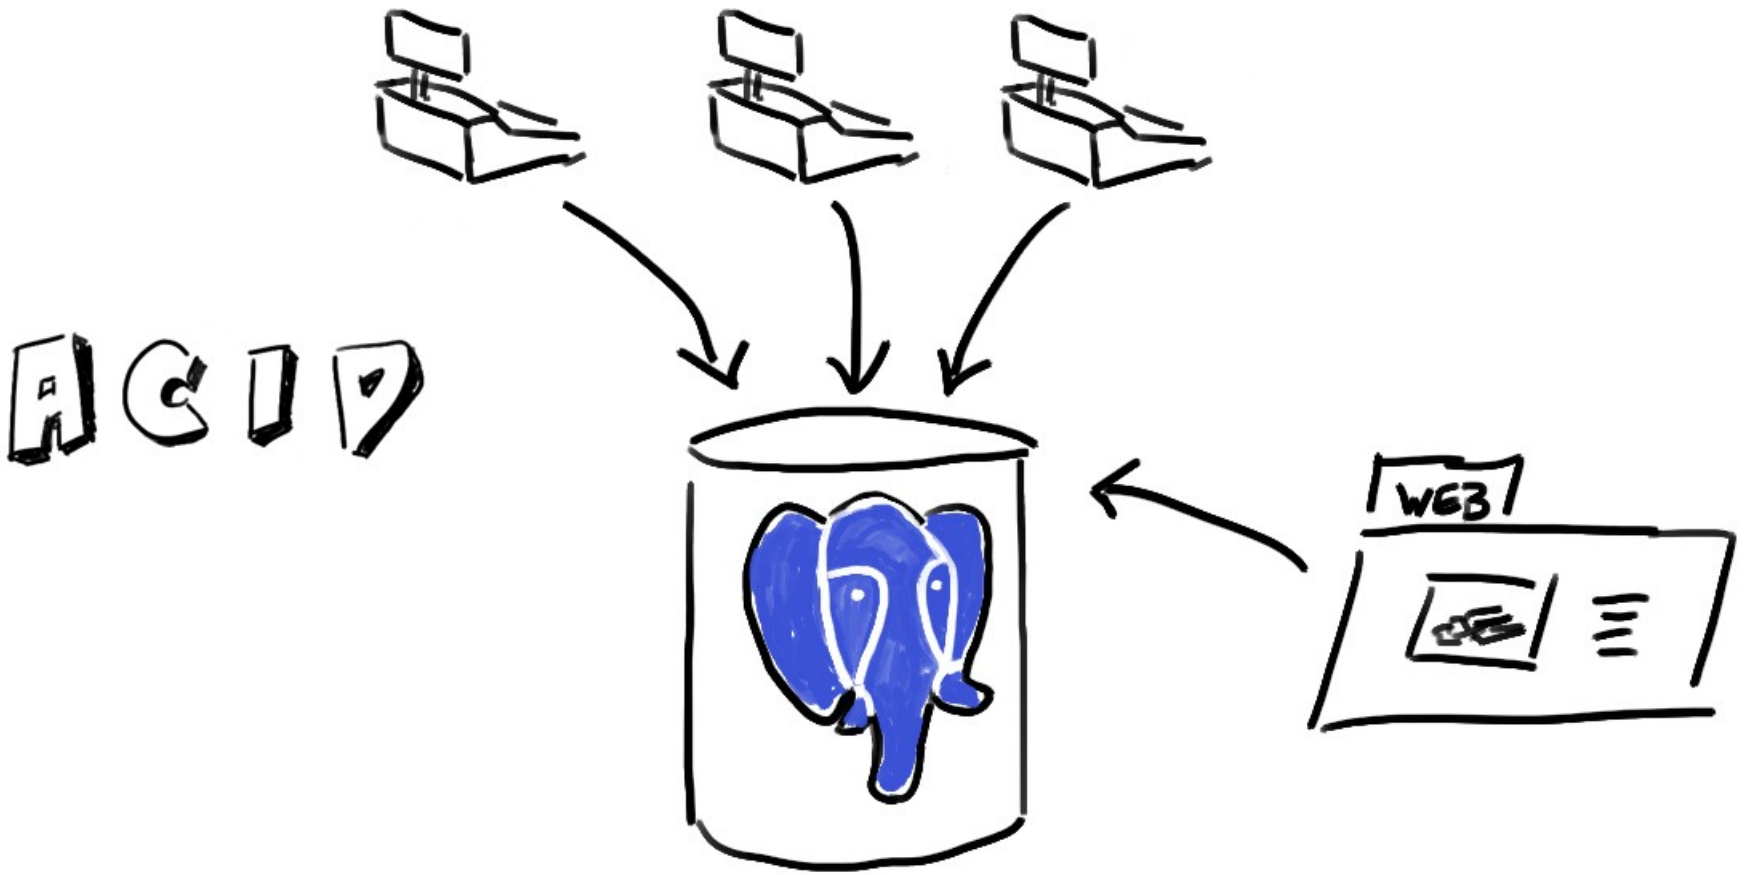

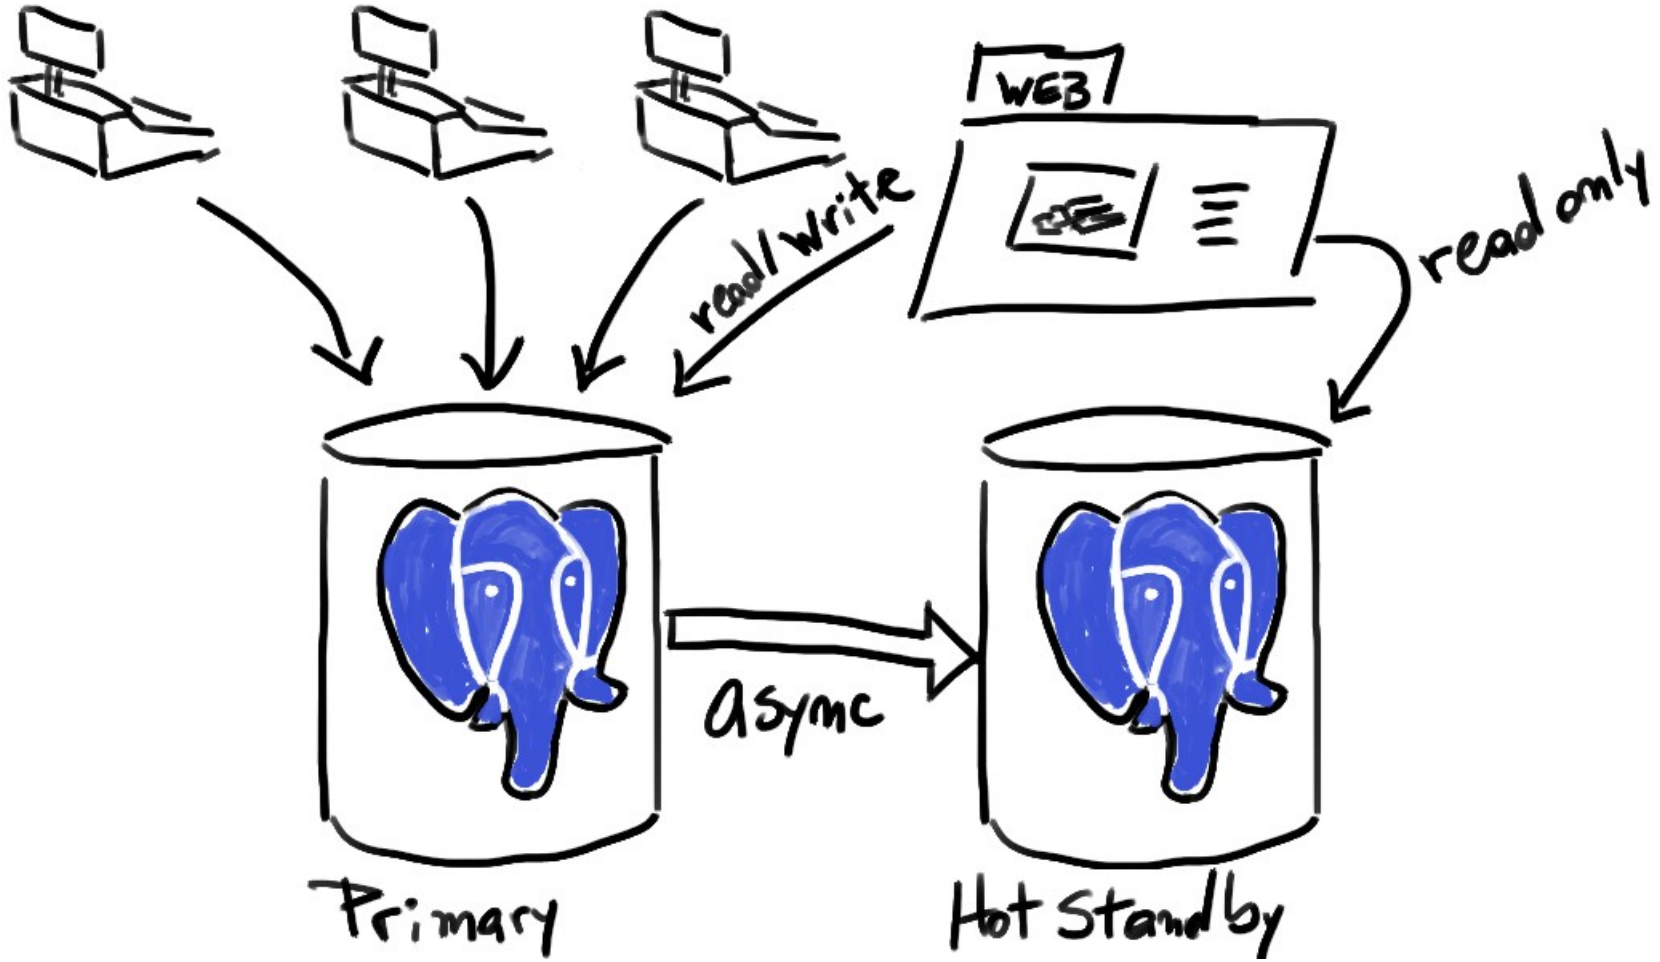

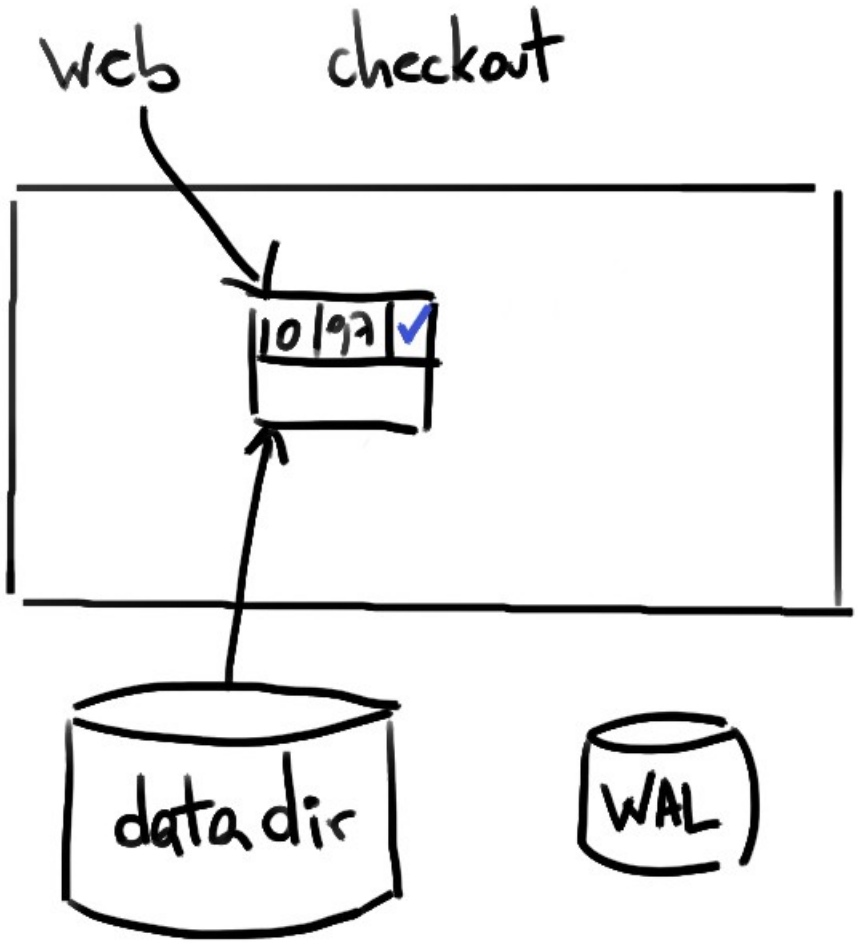

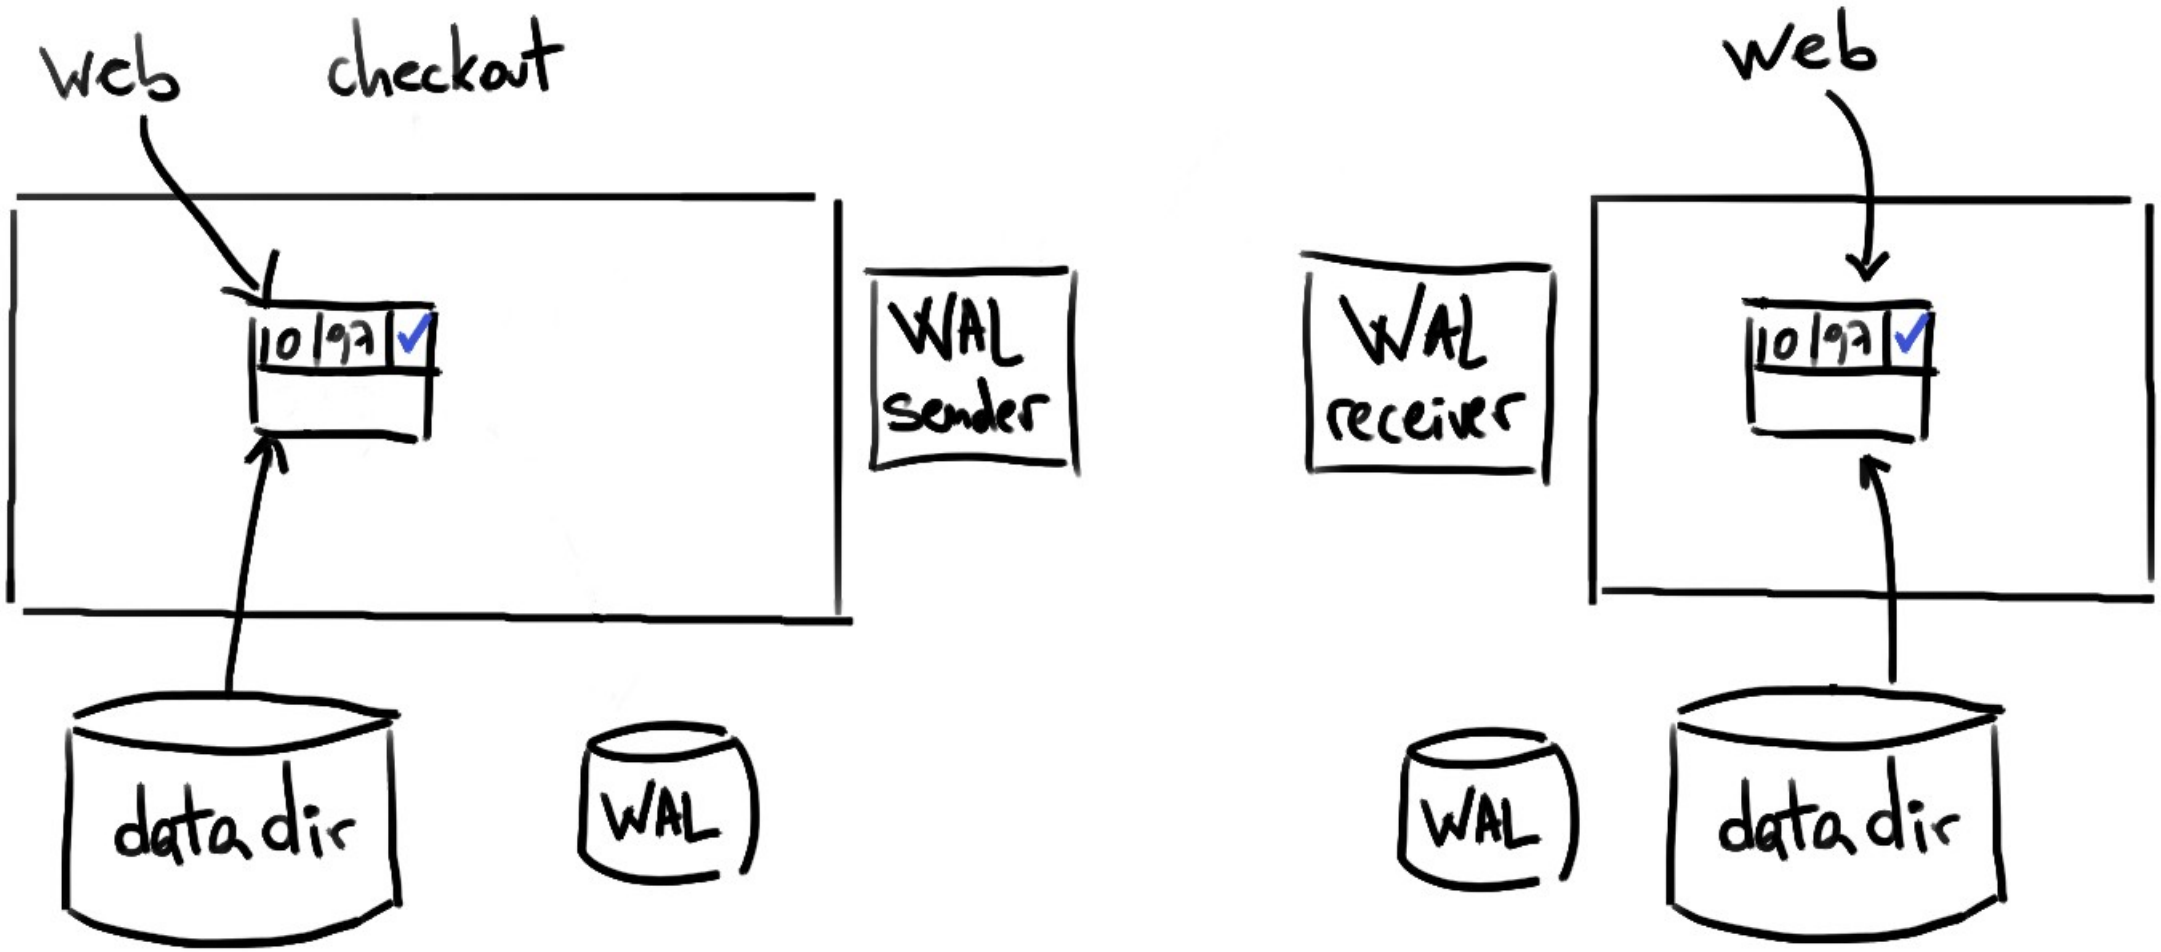

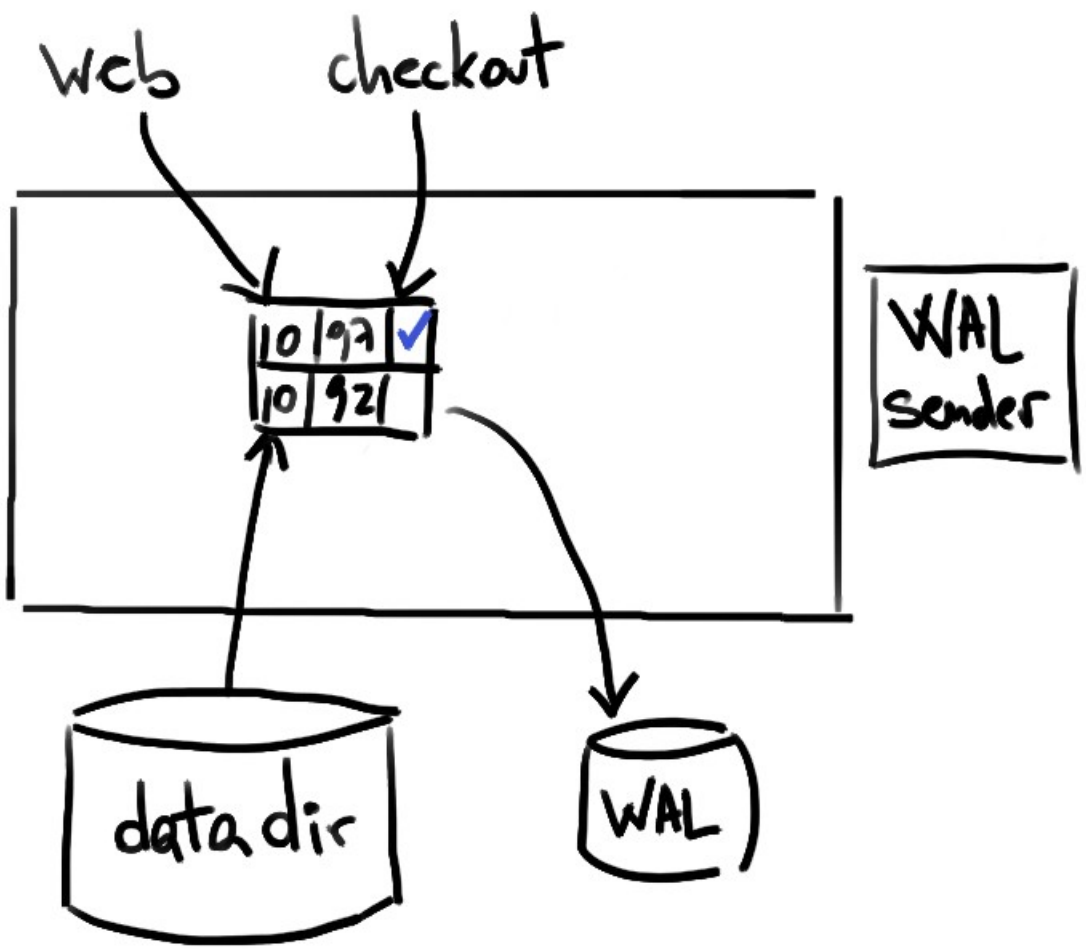

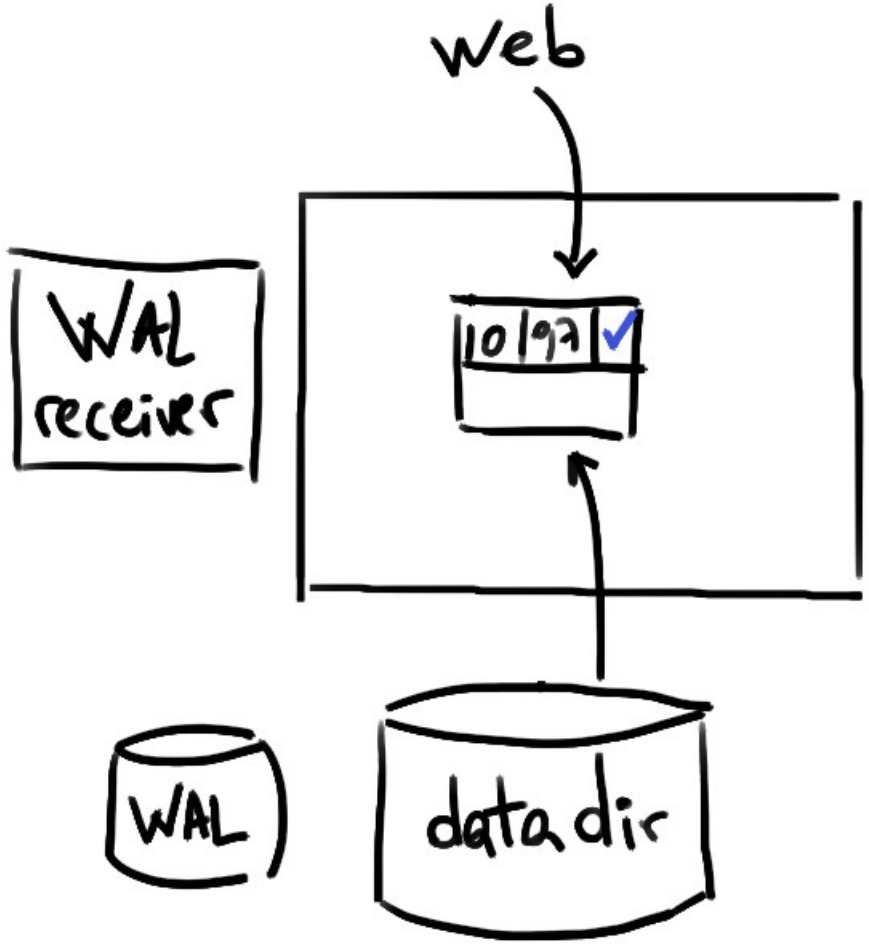

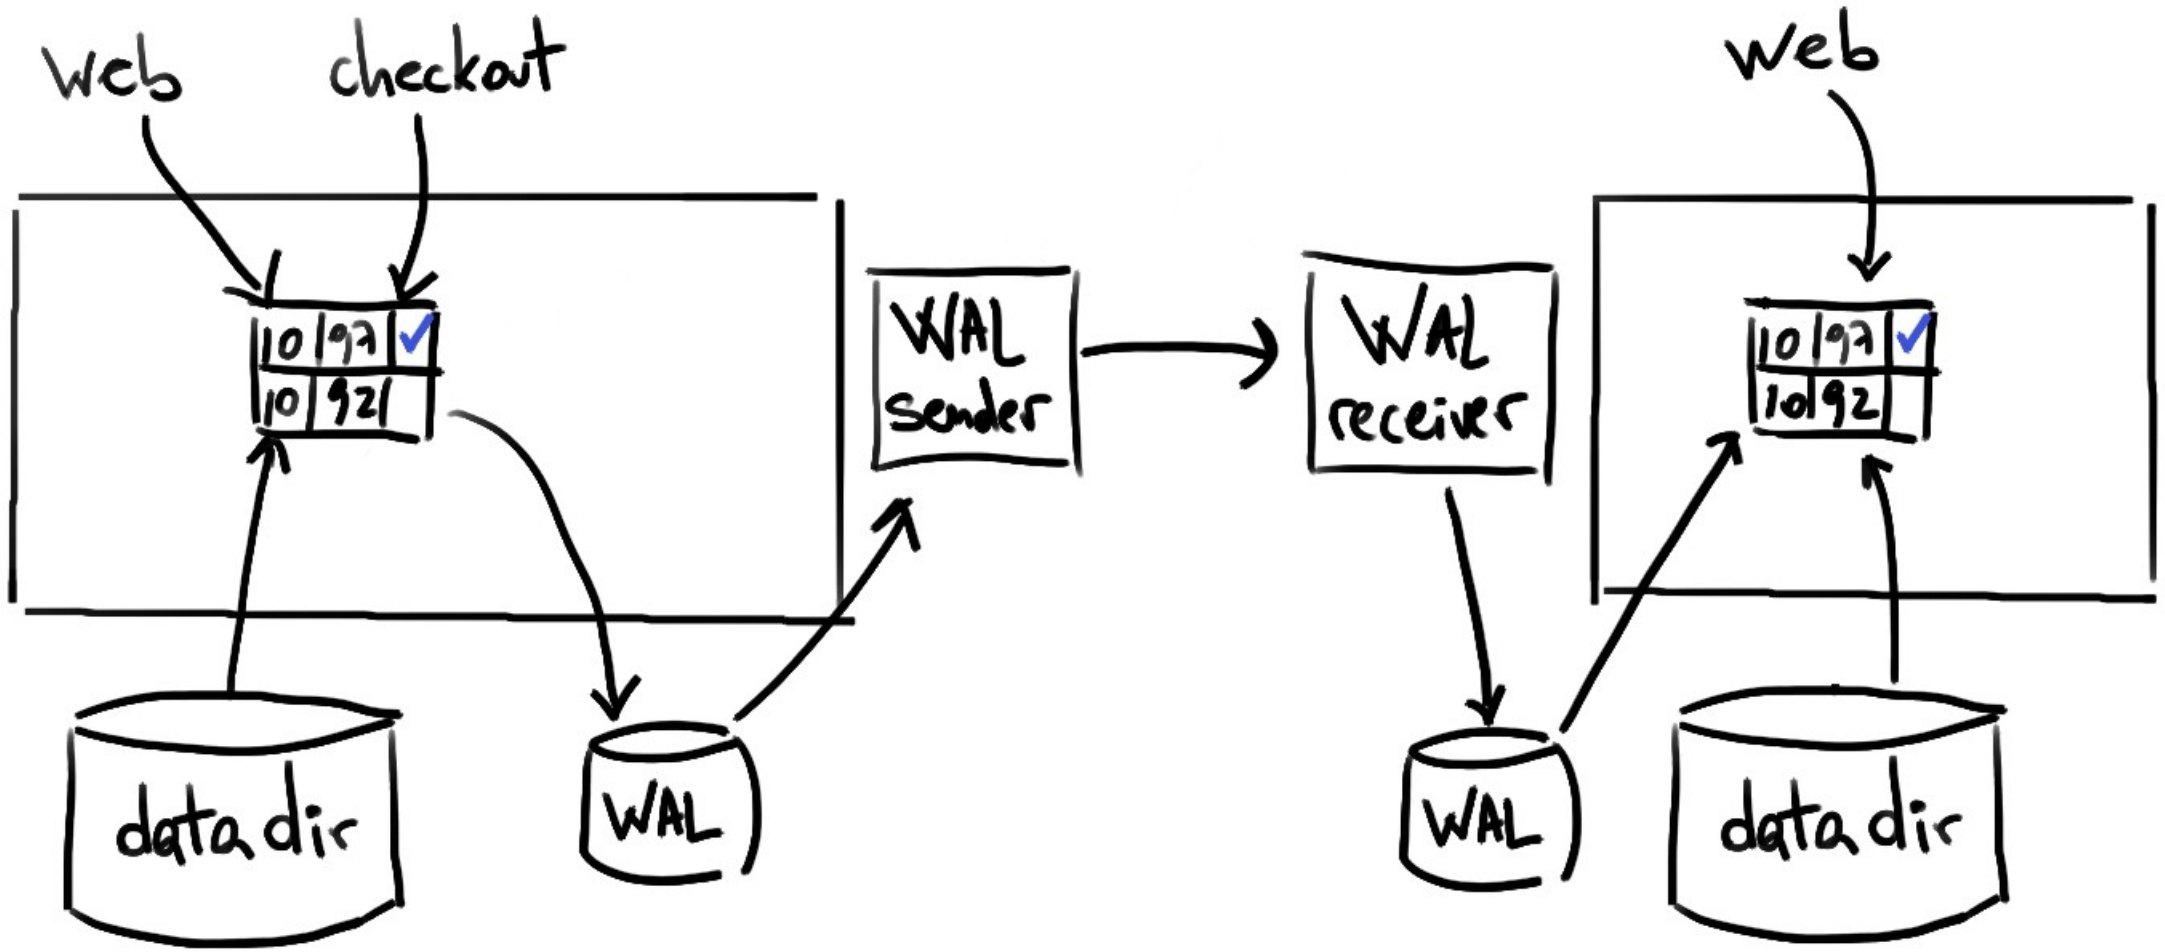

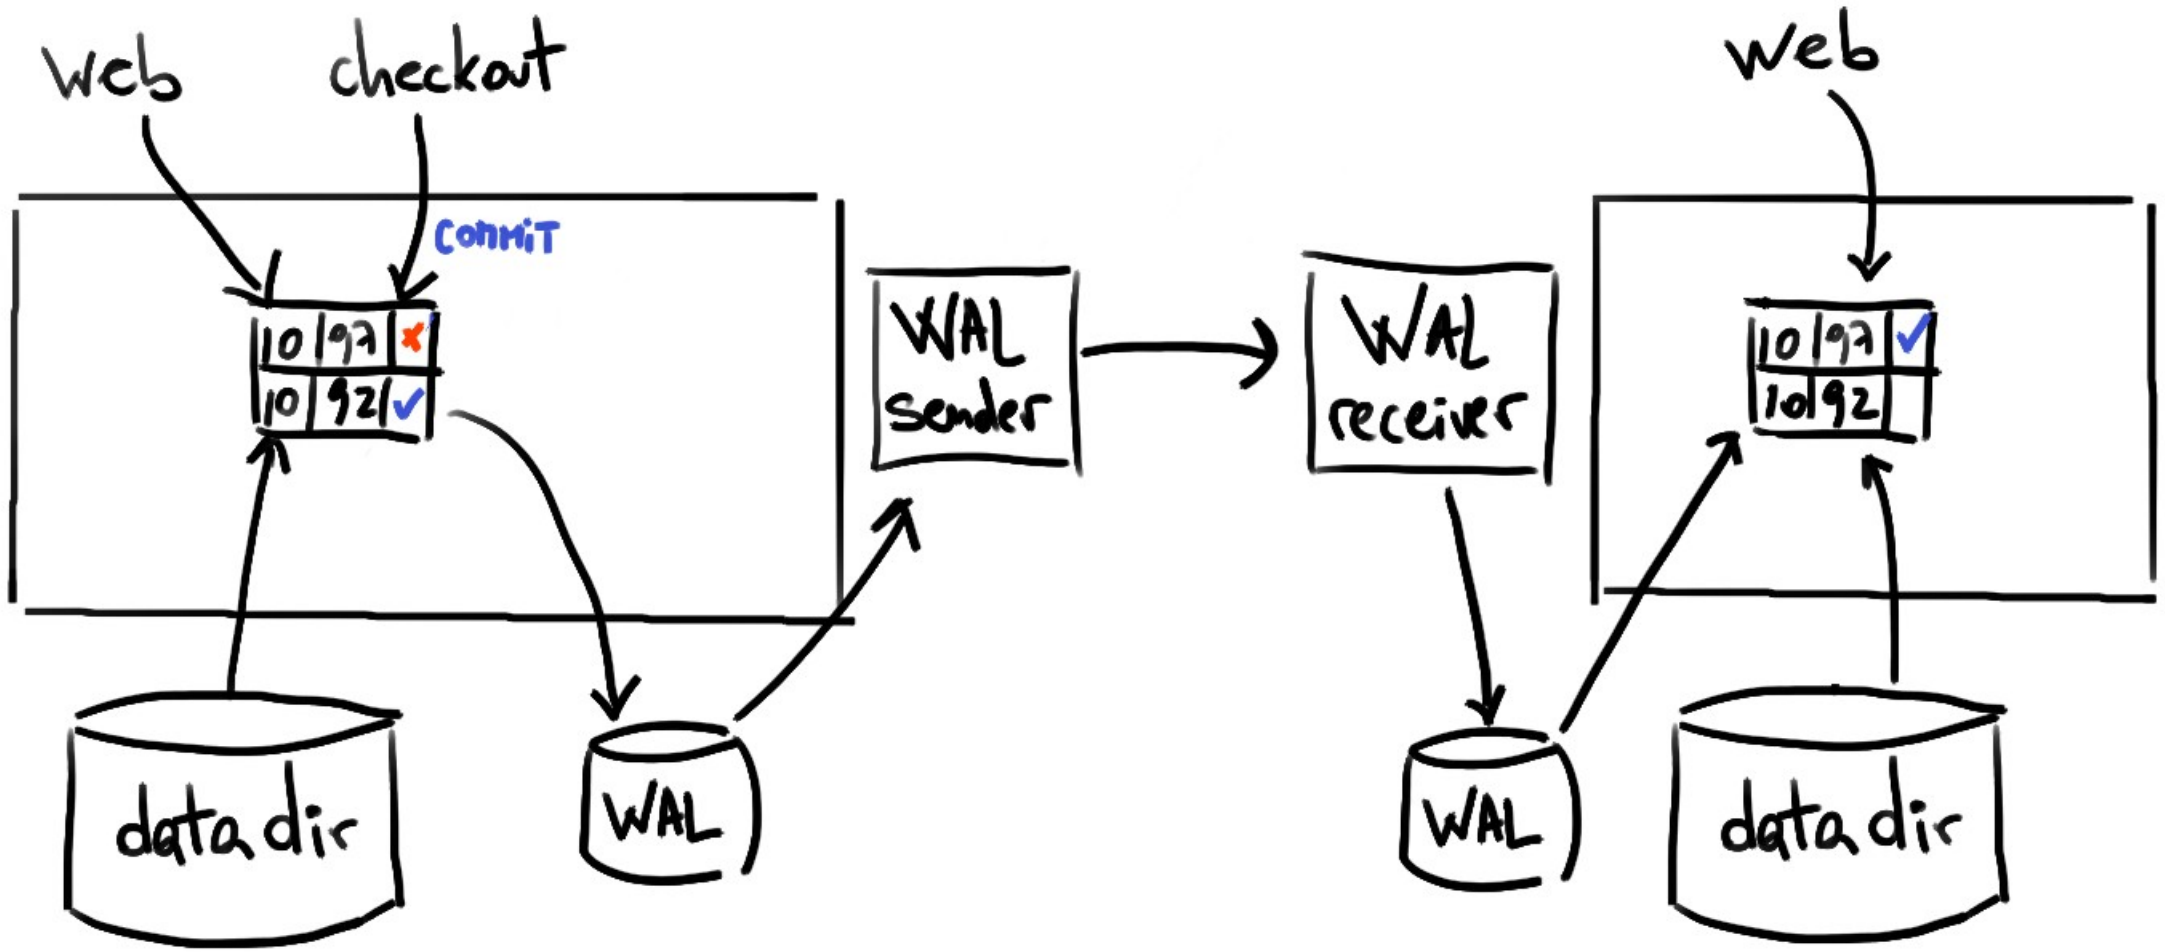

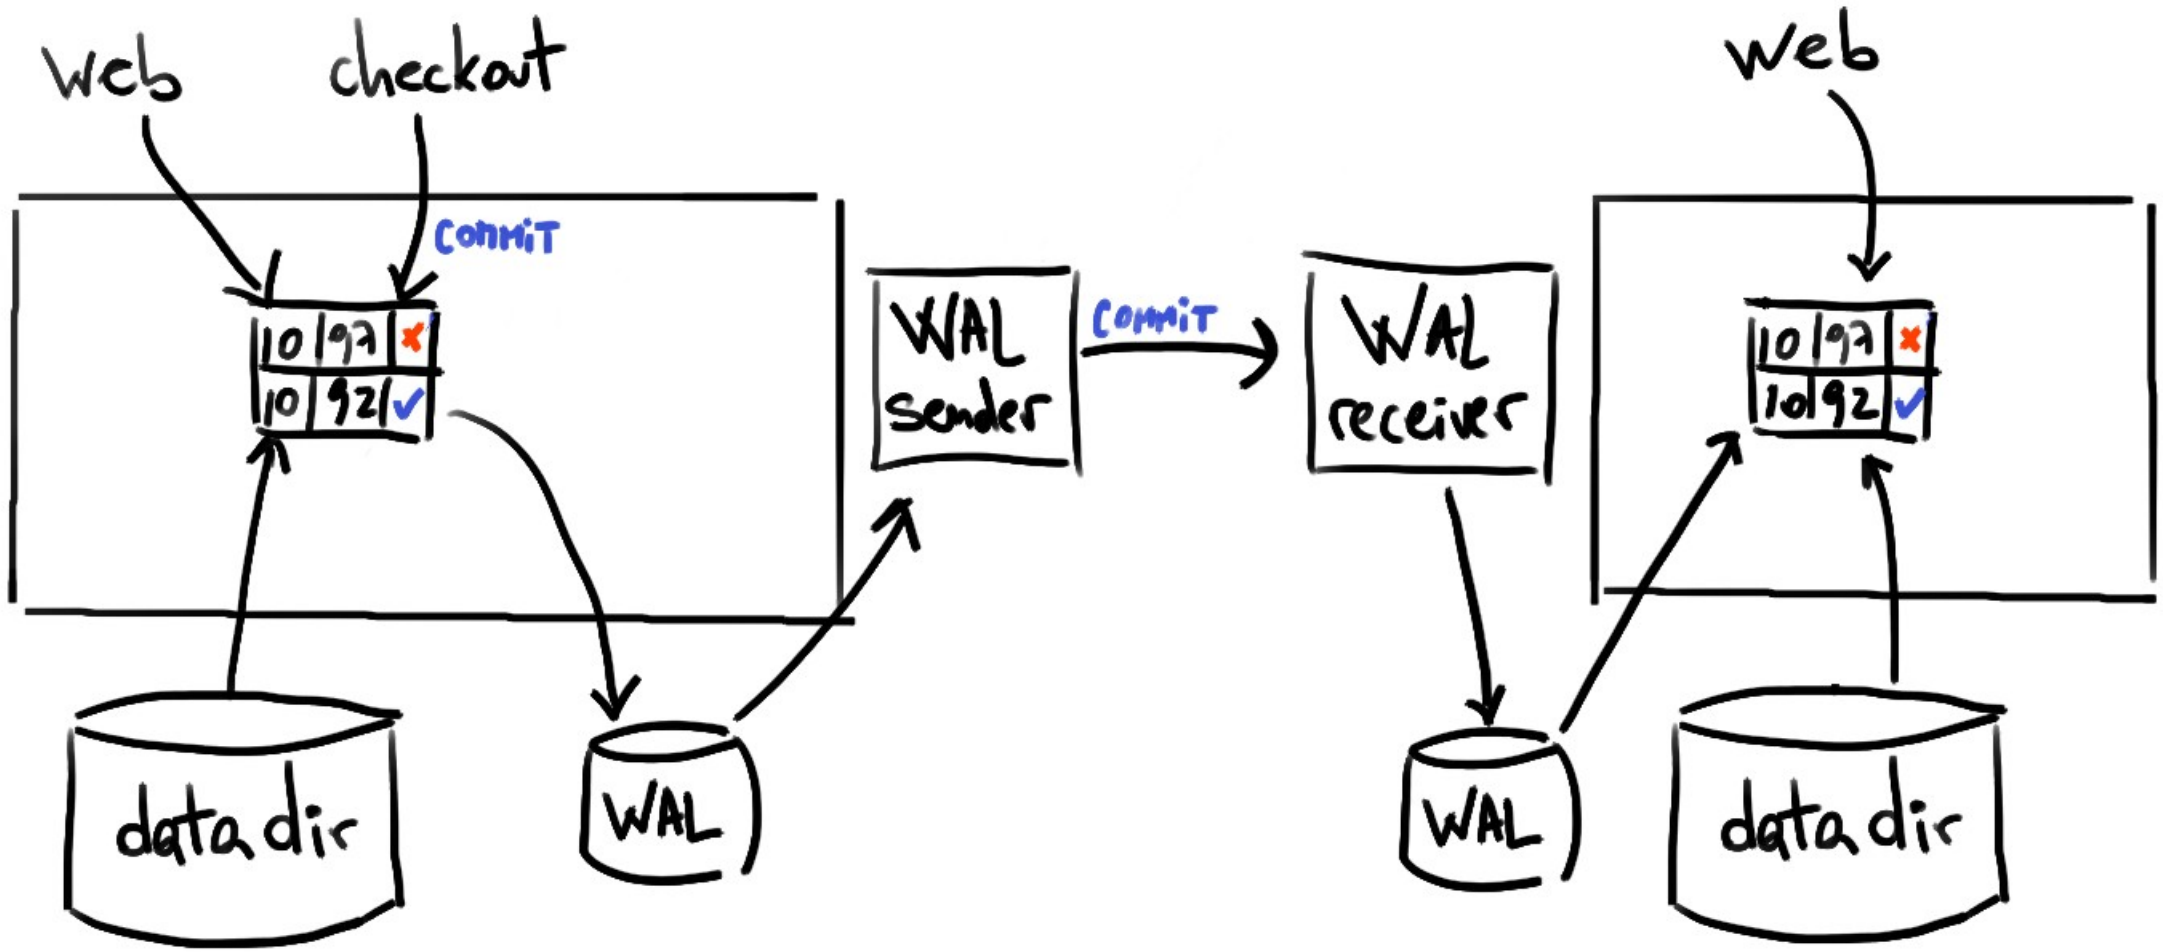

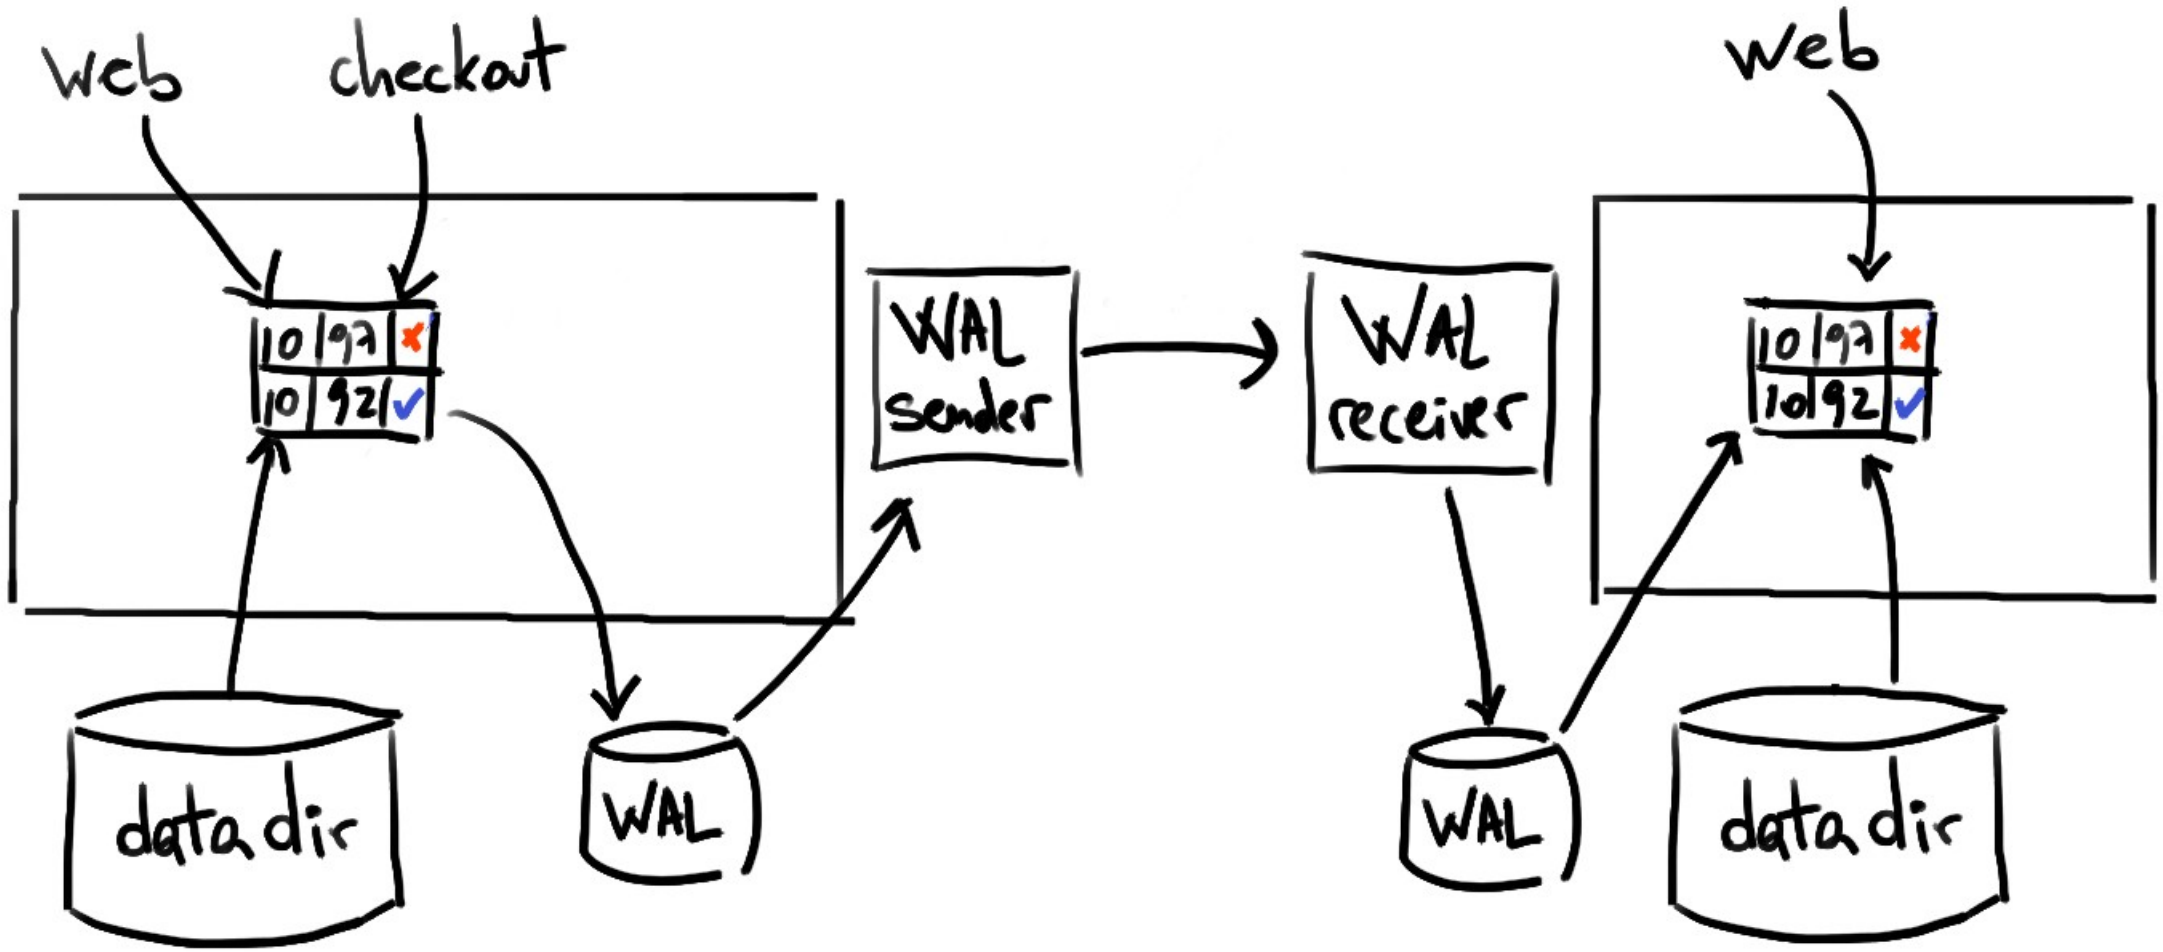

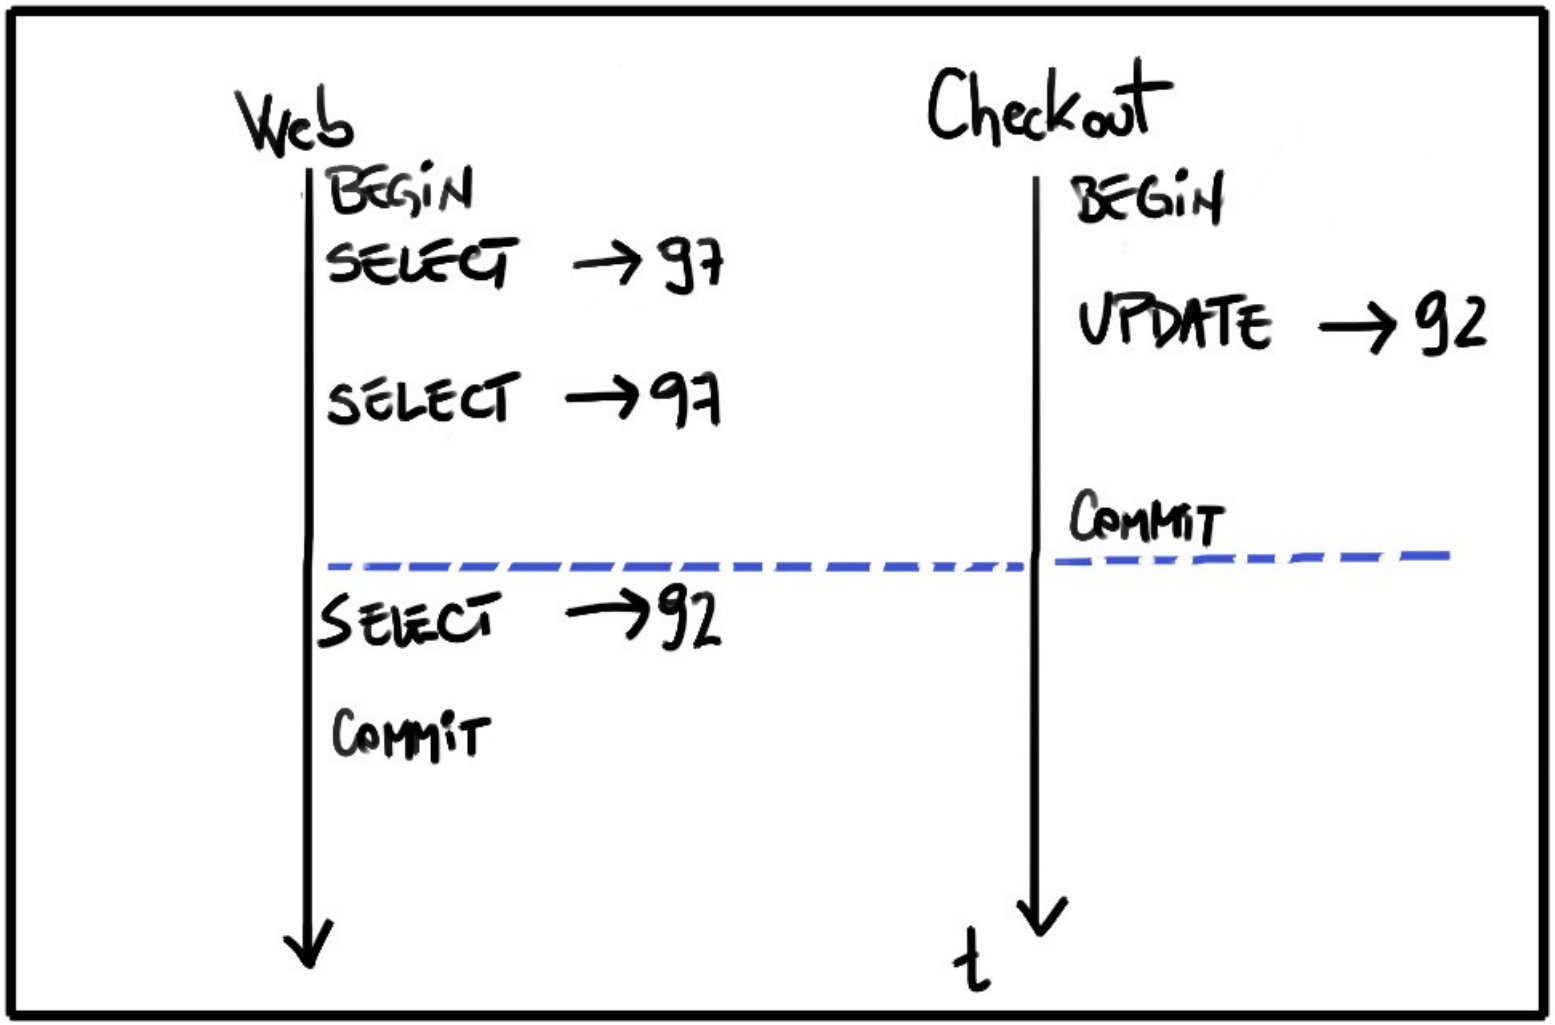

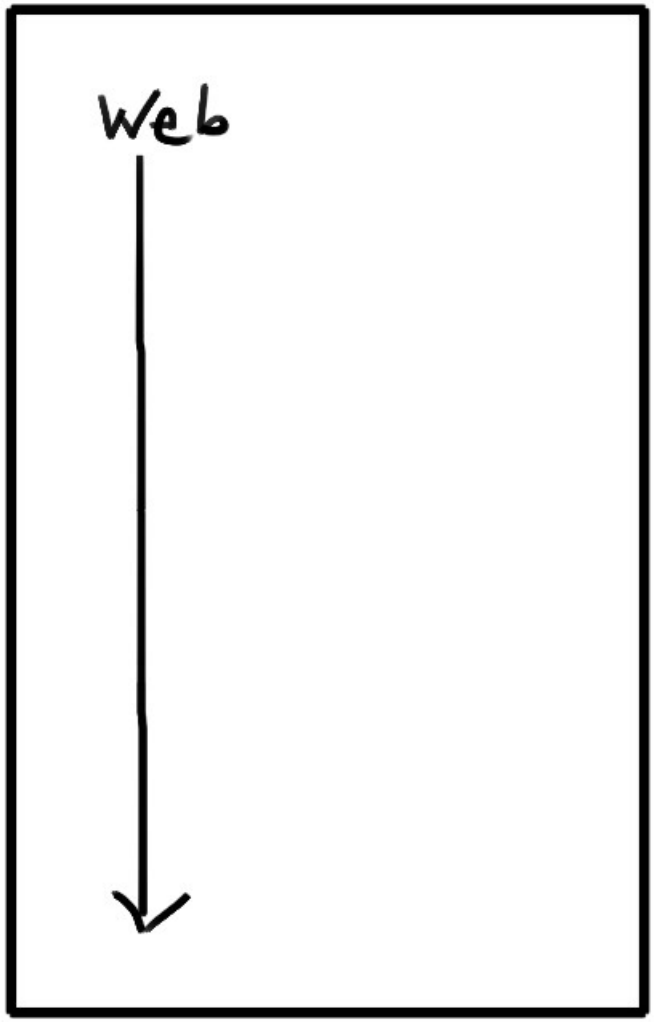

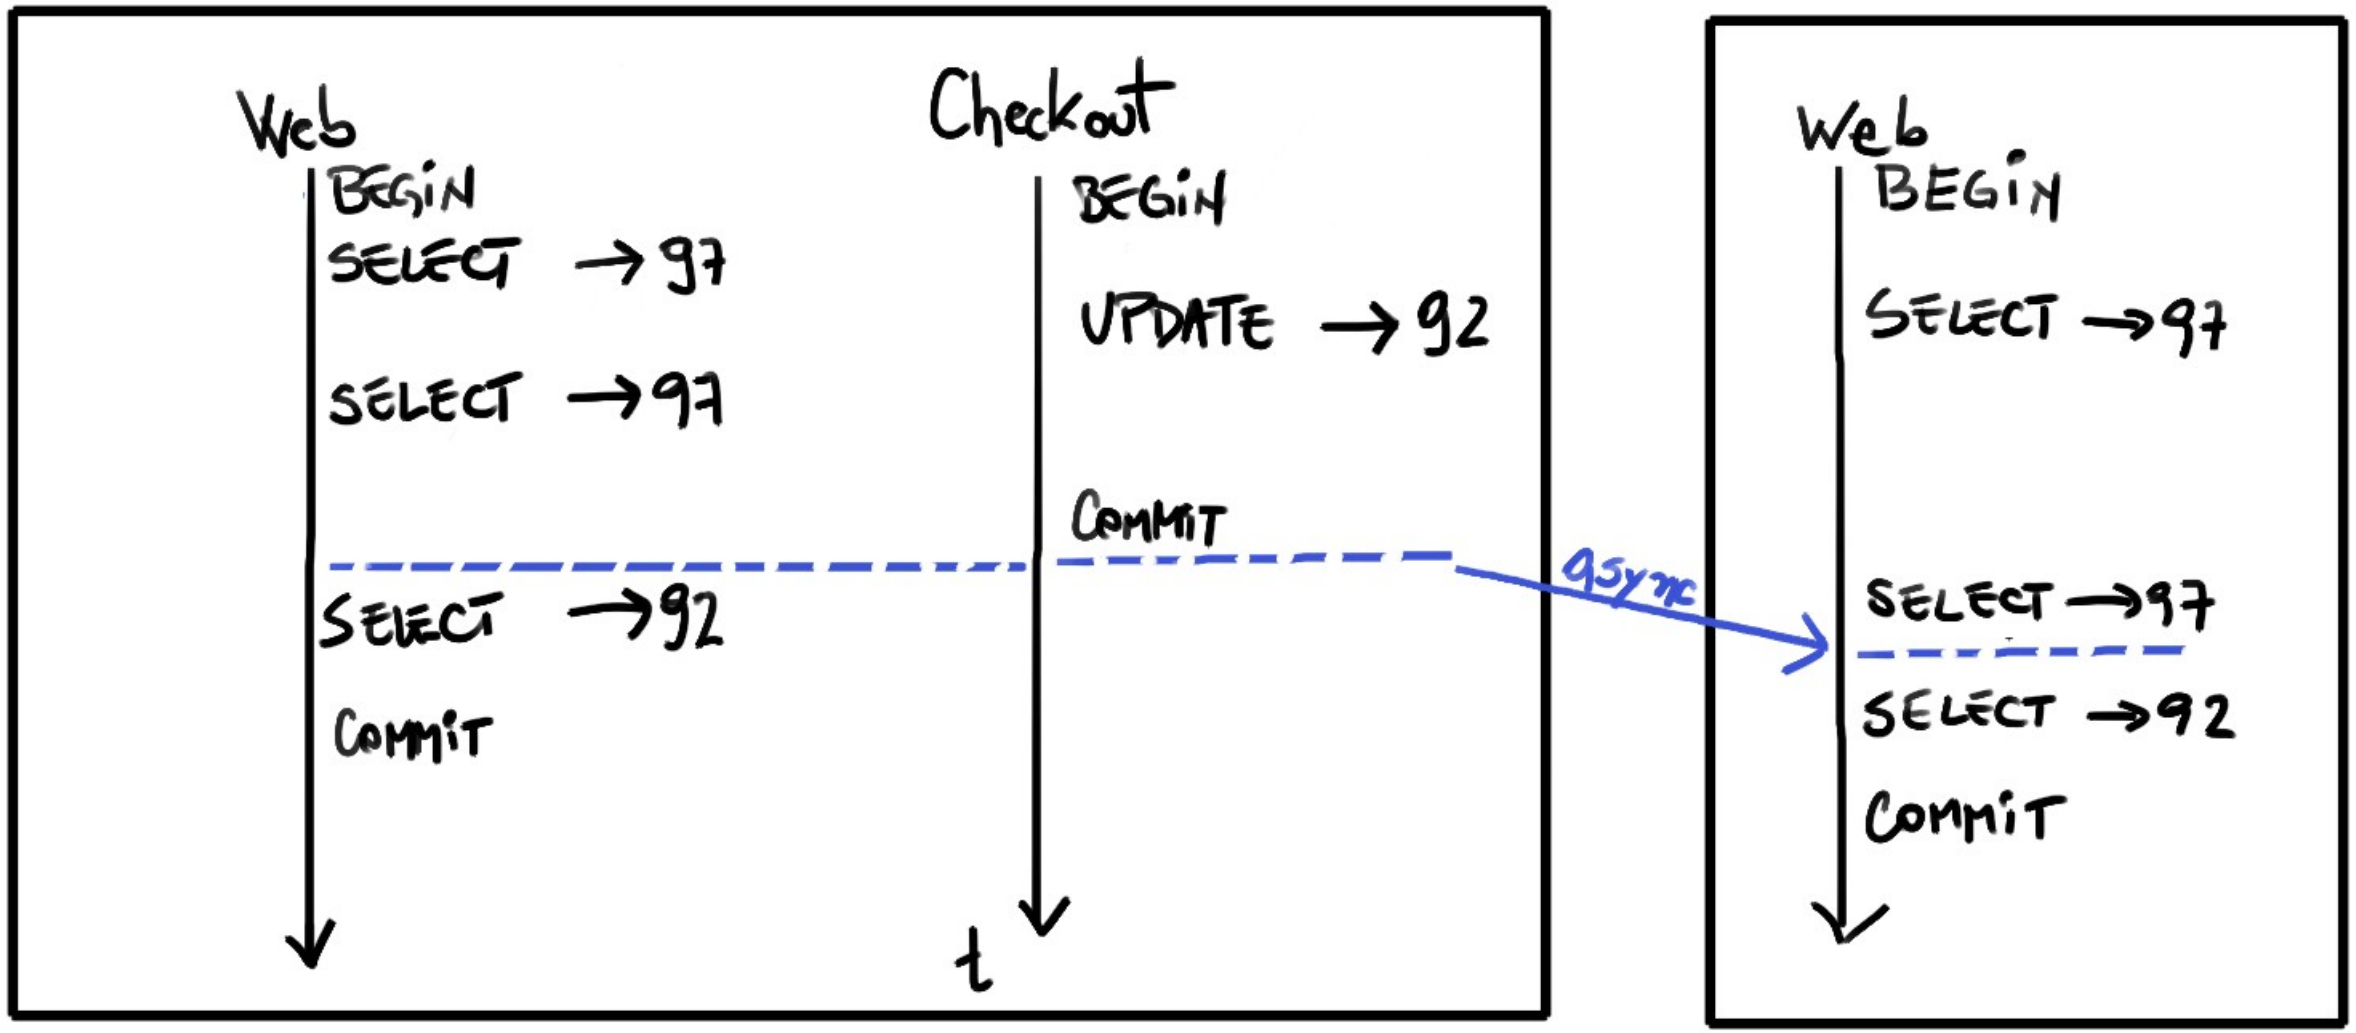

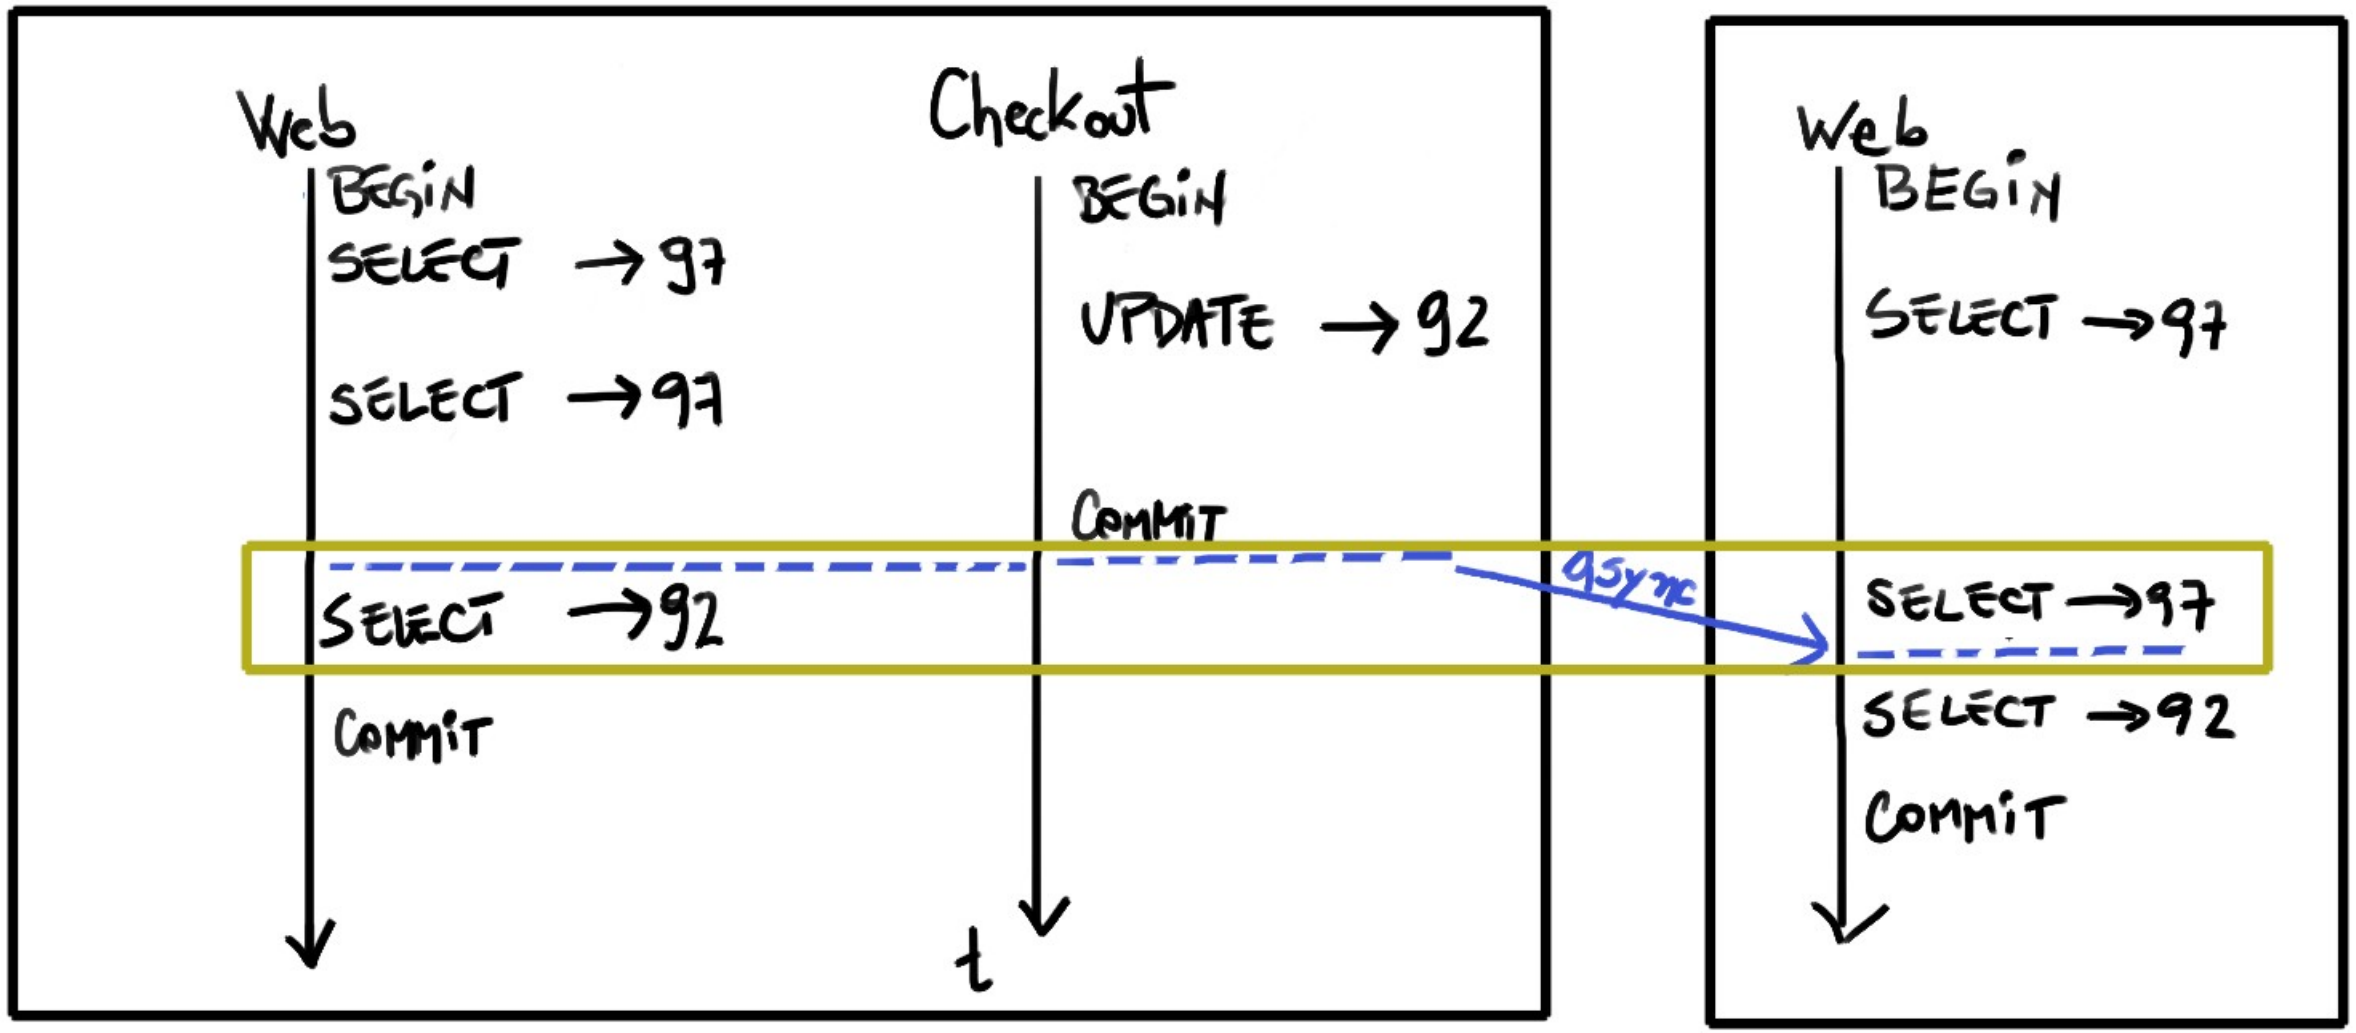

Consistency

Every read receives the most recent write or an error

Eventual Consistency

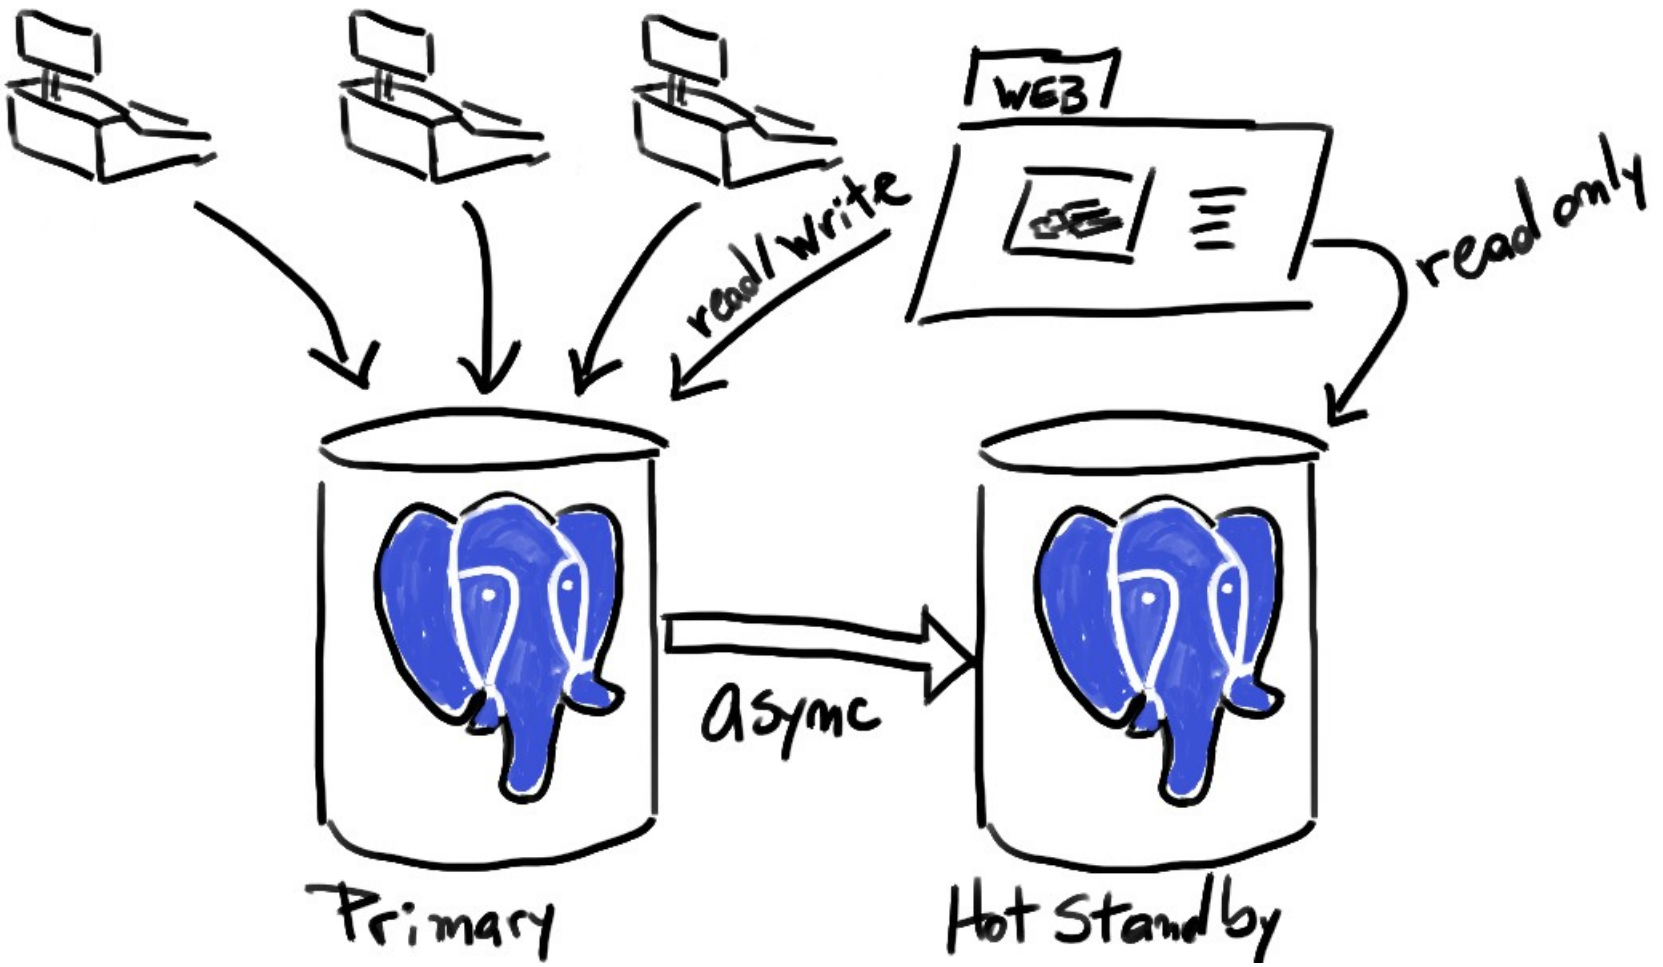

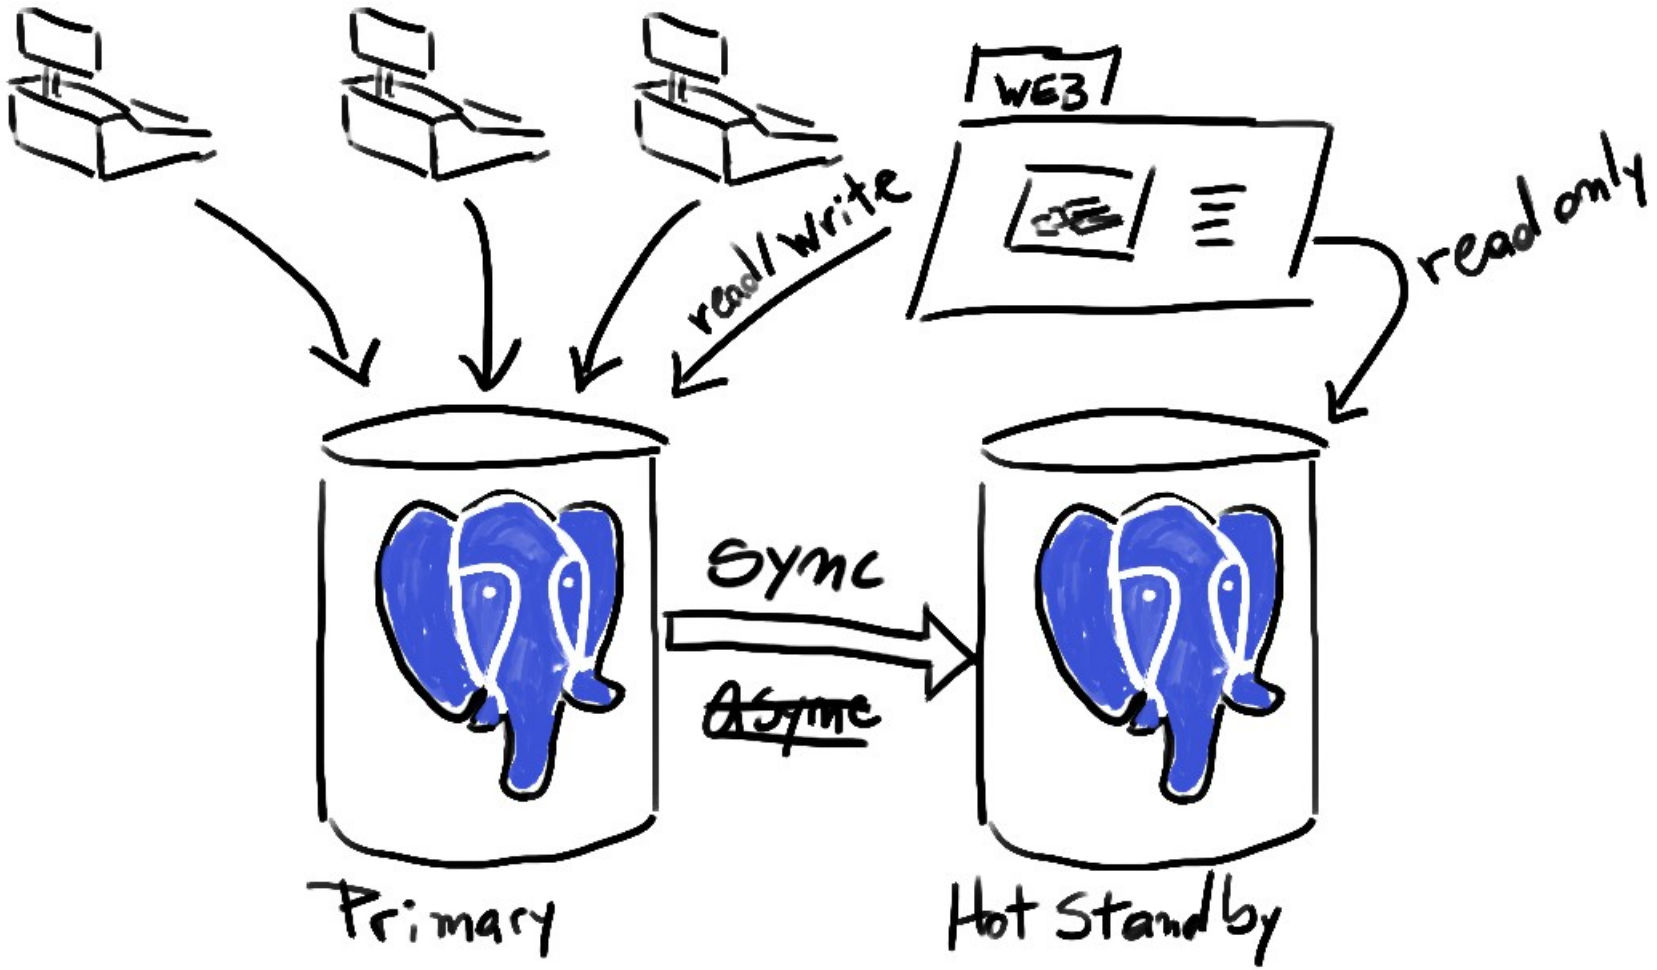

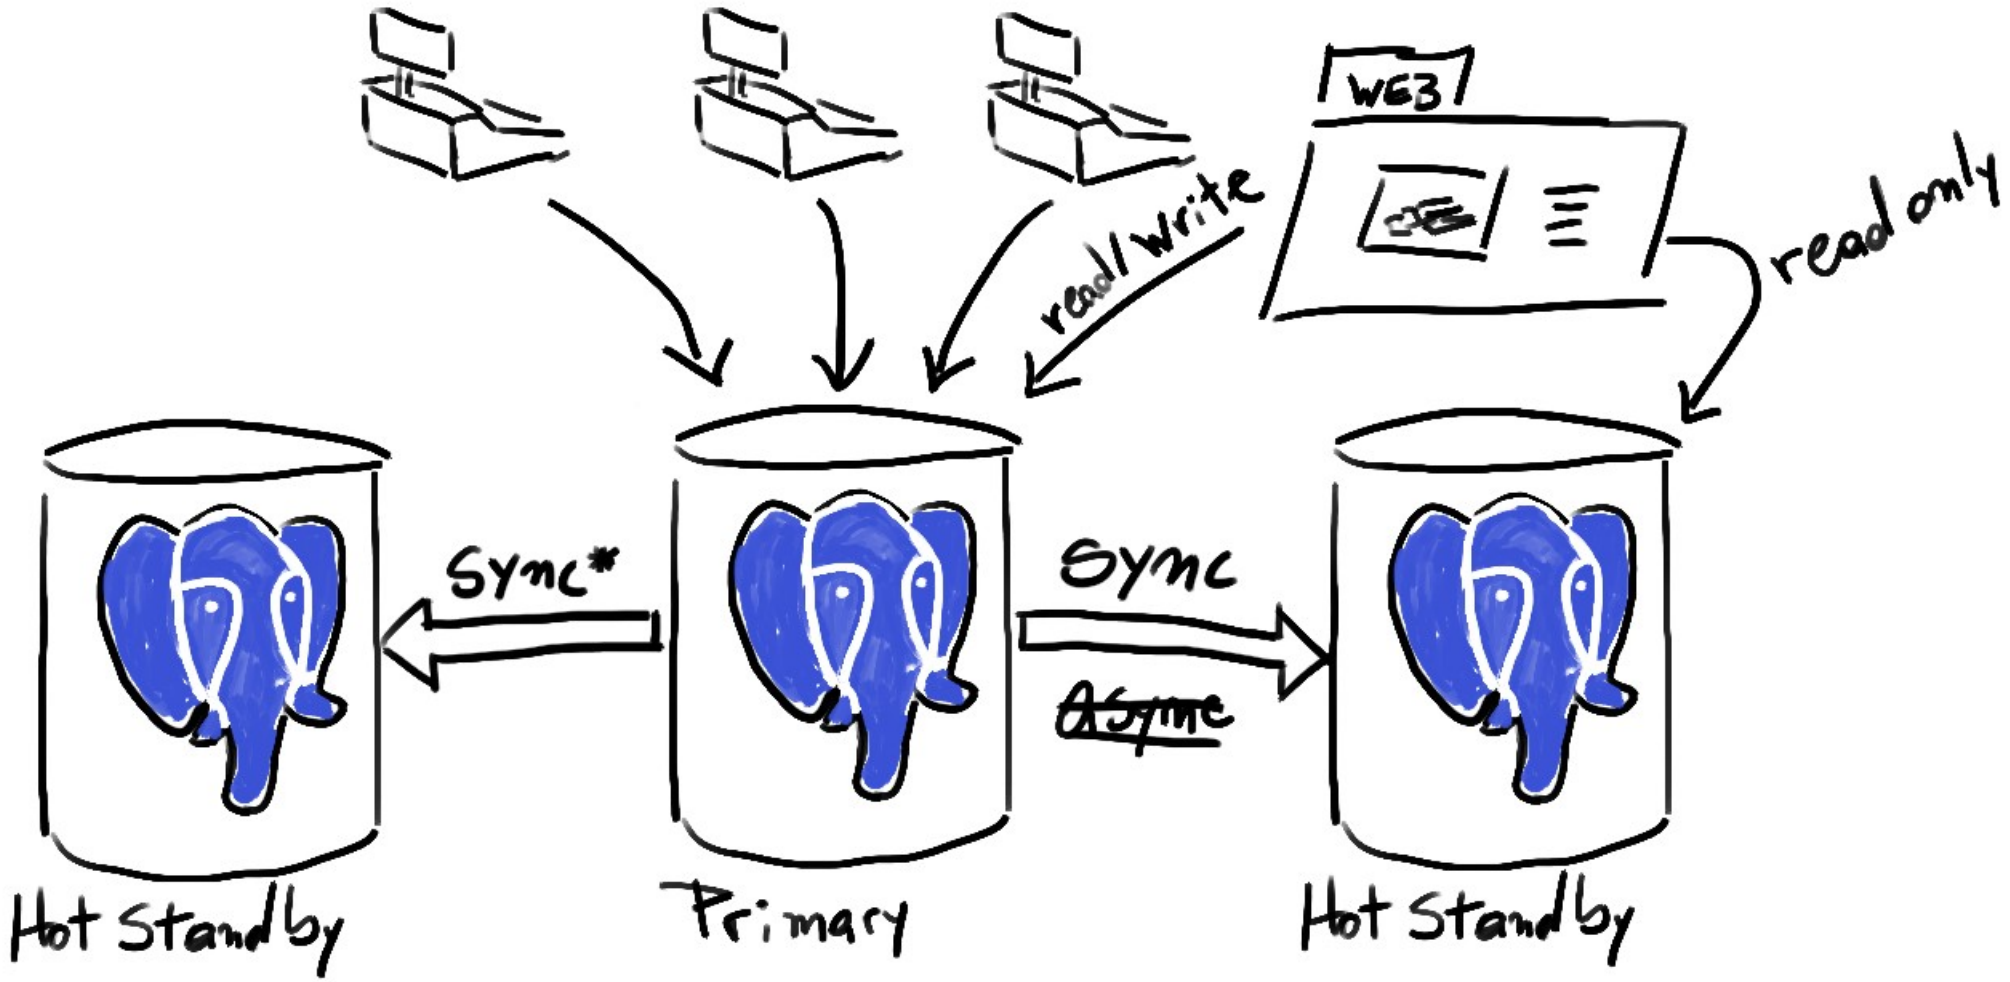

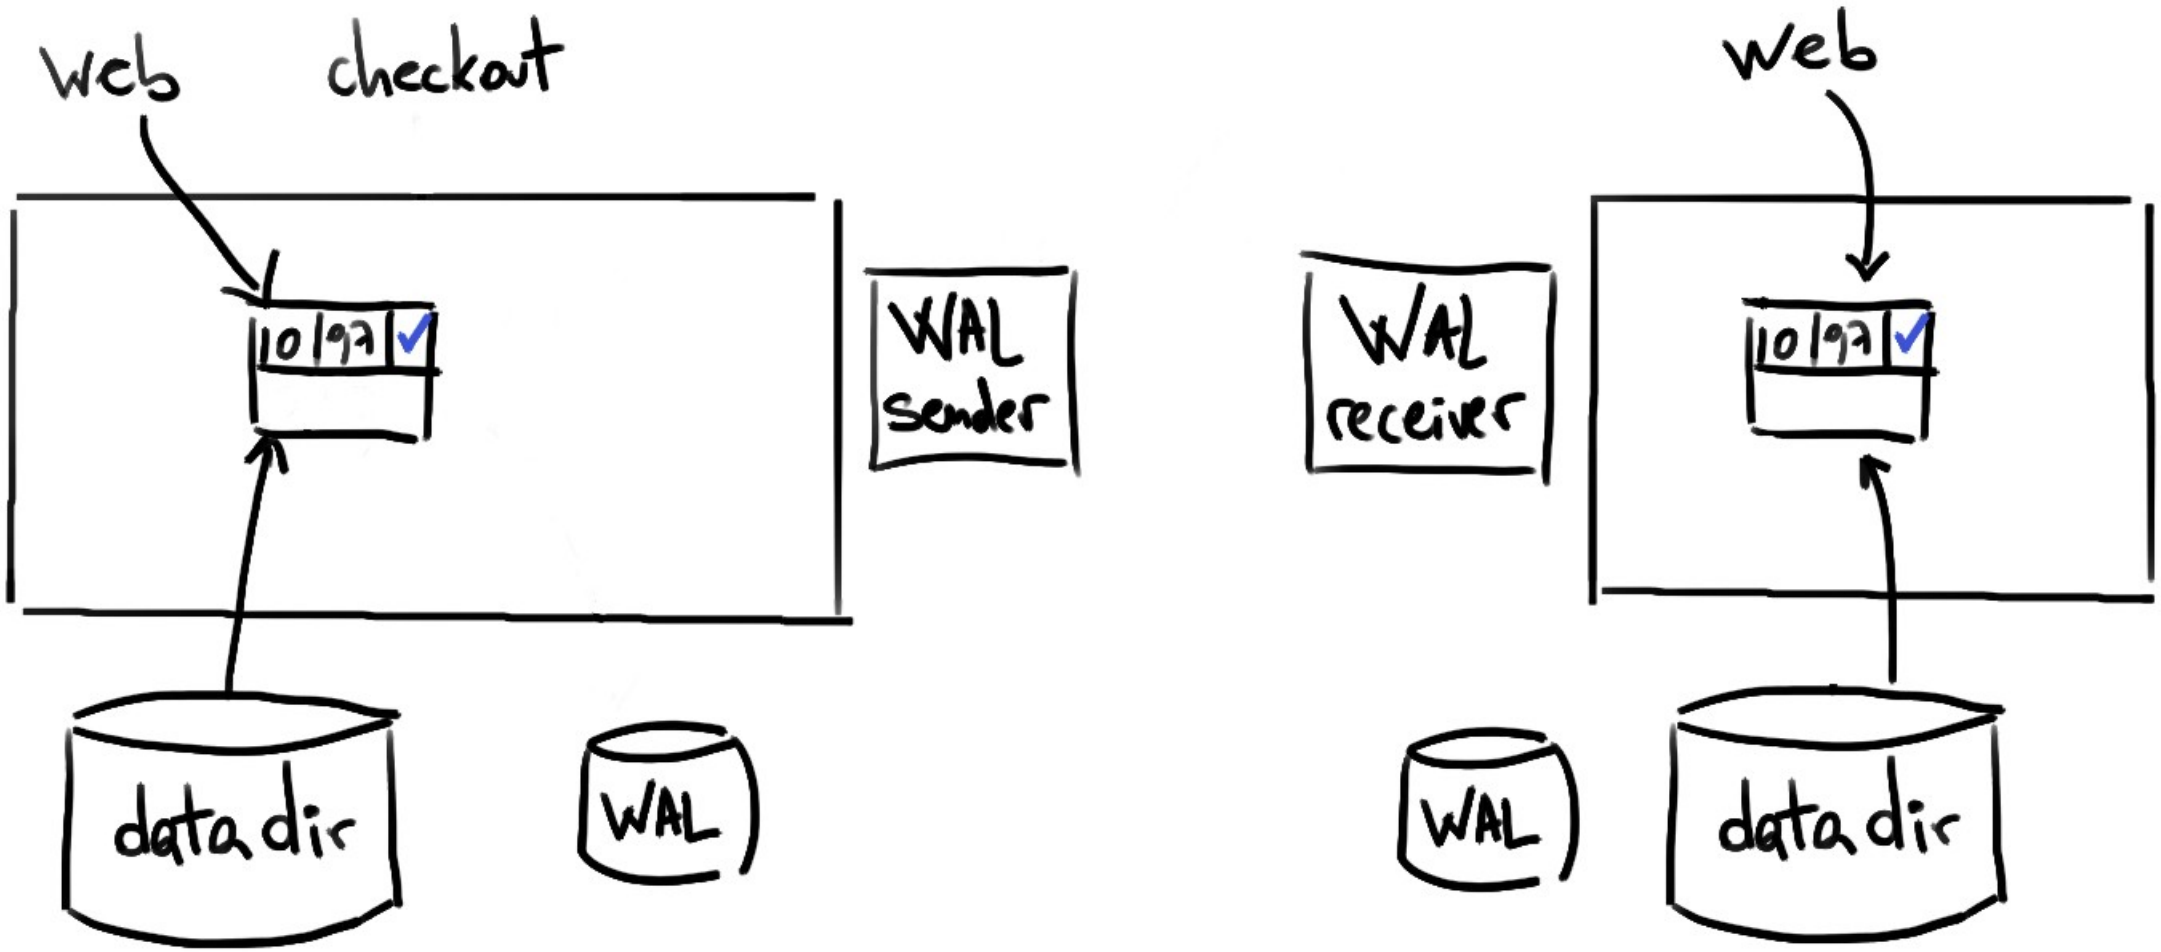

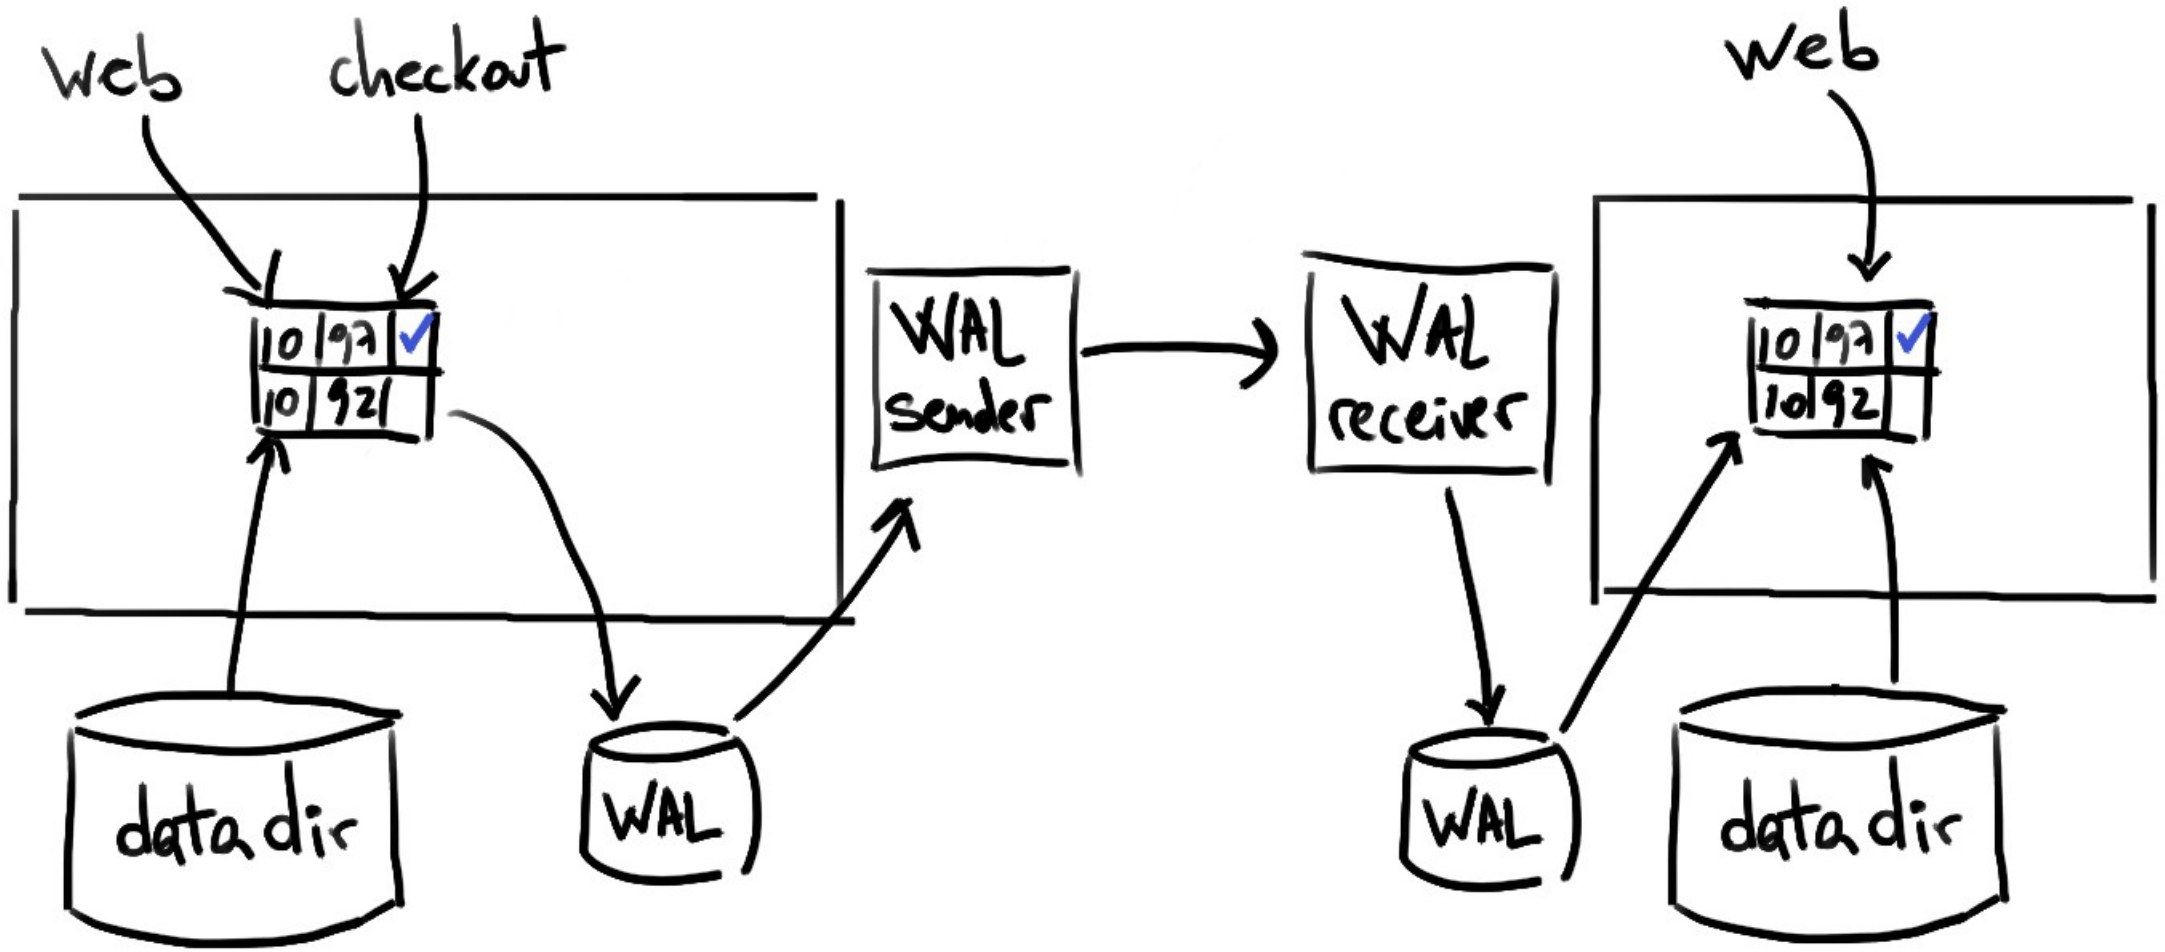

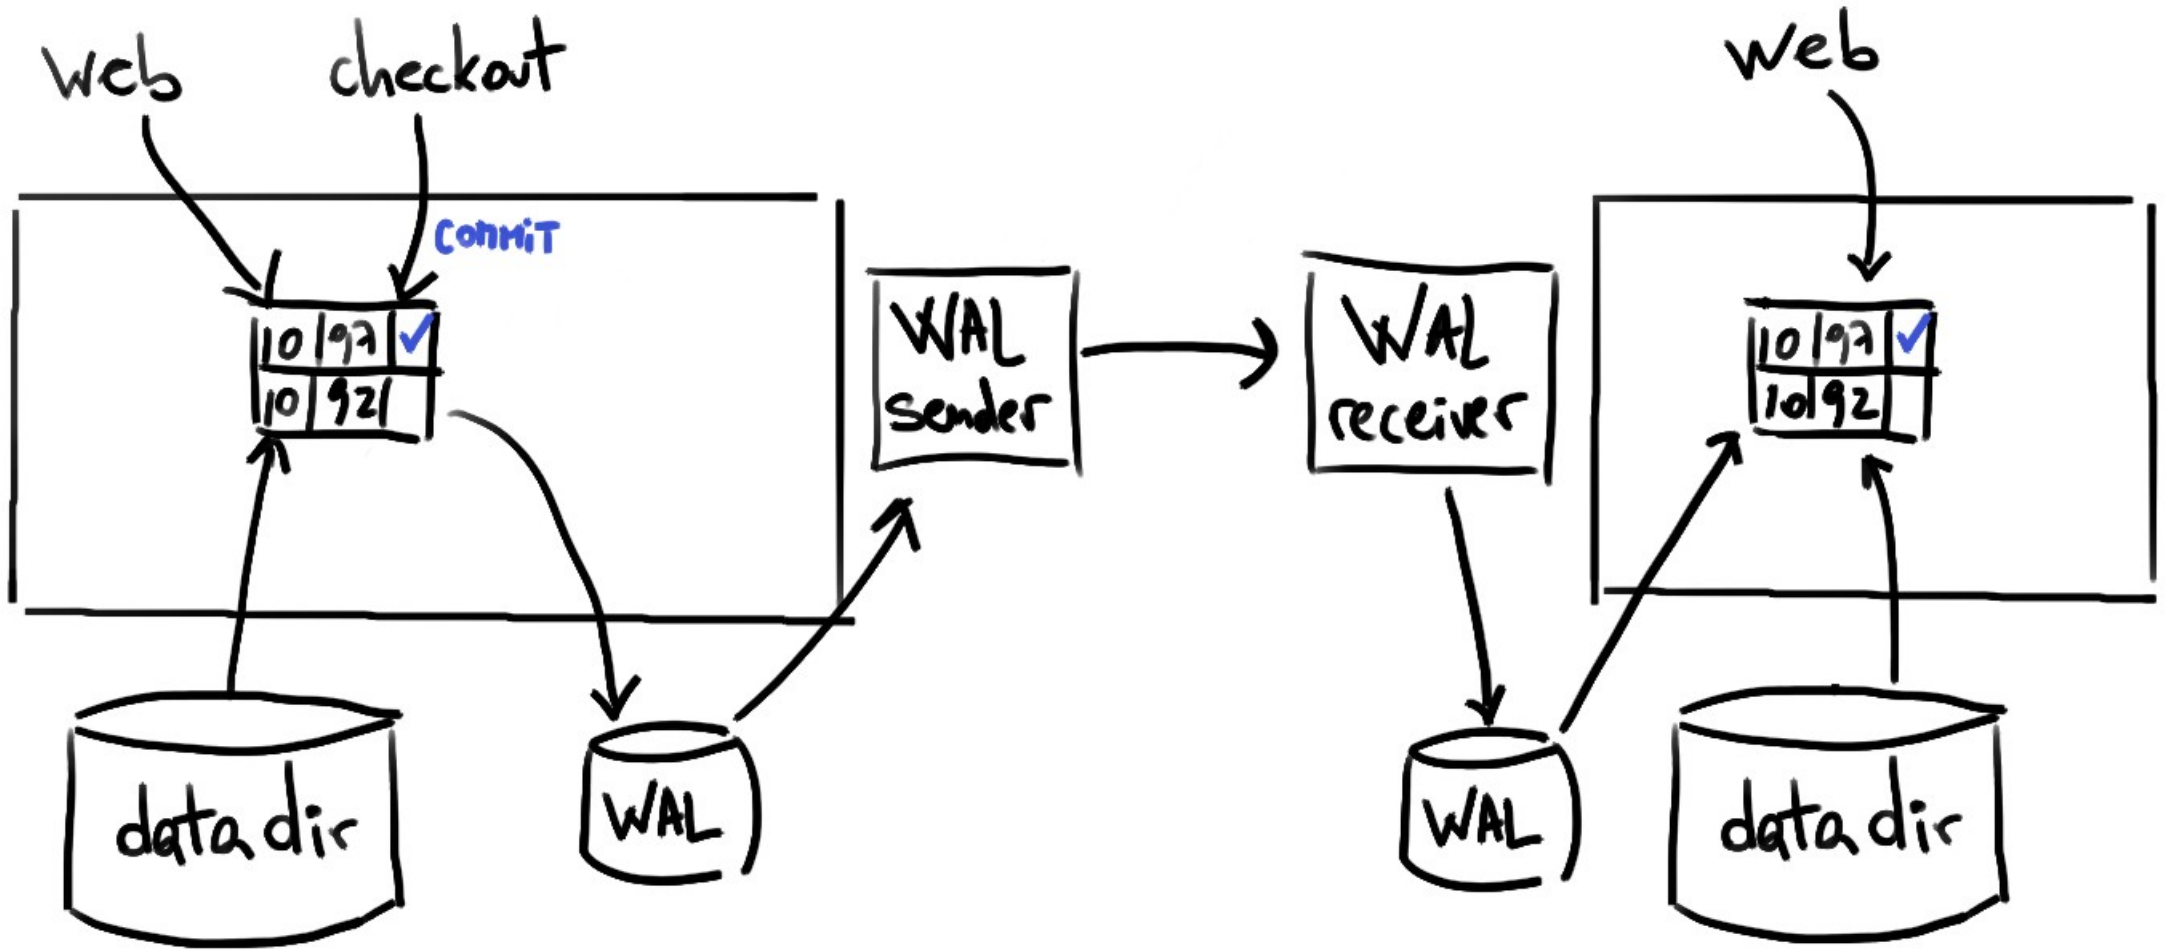

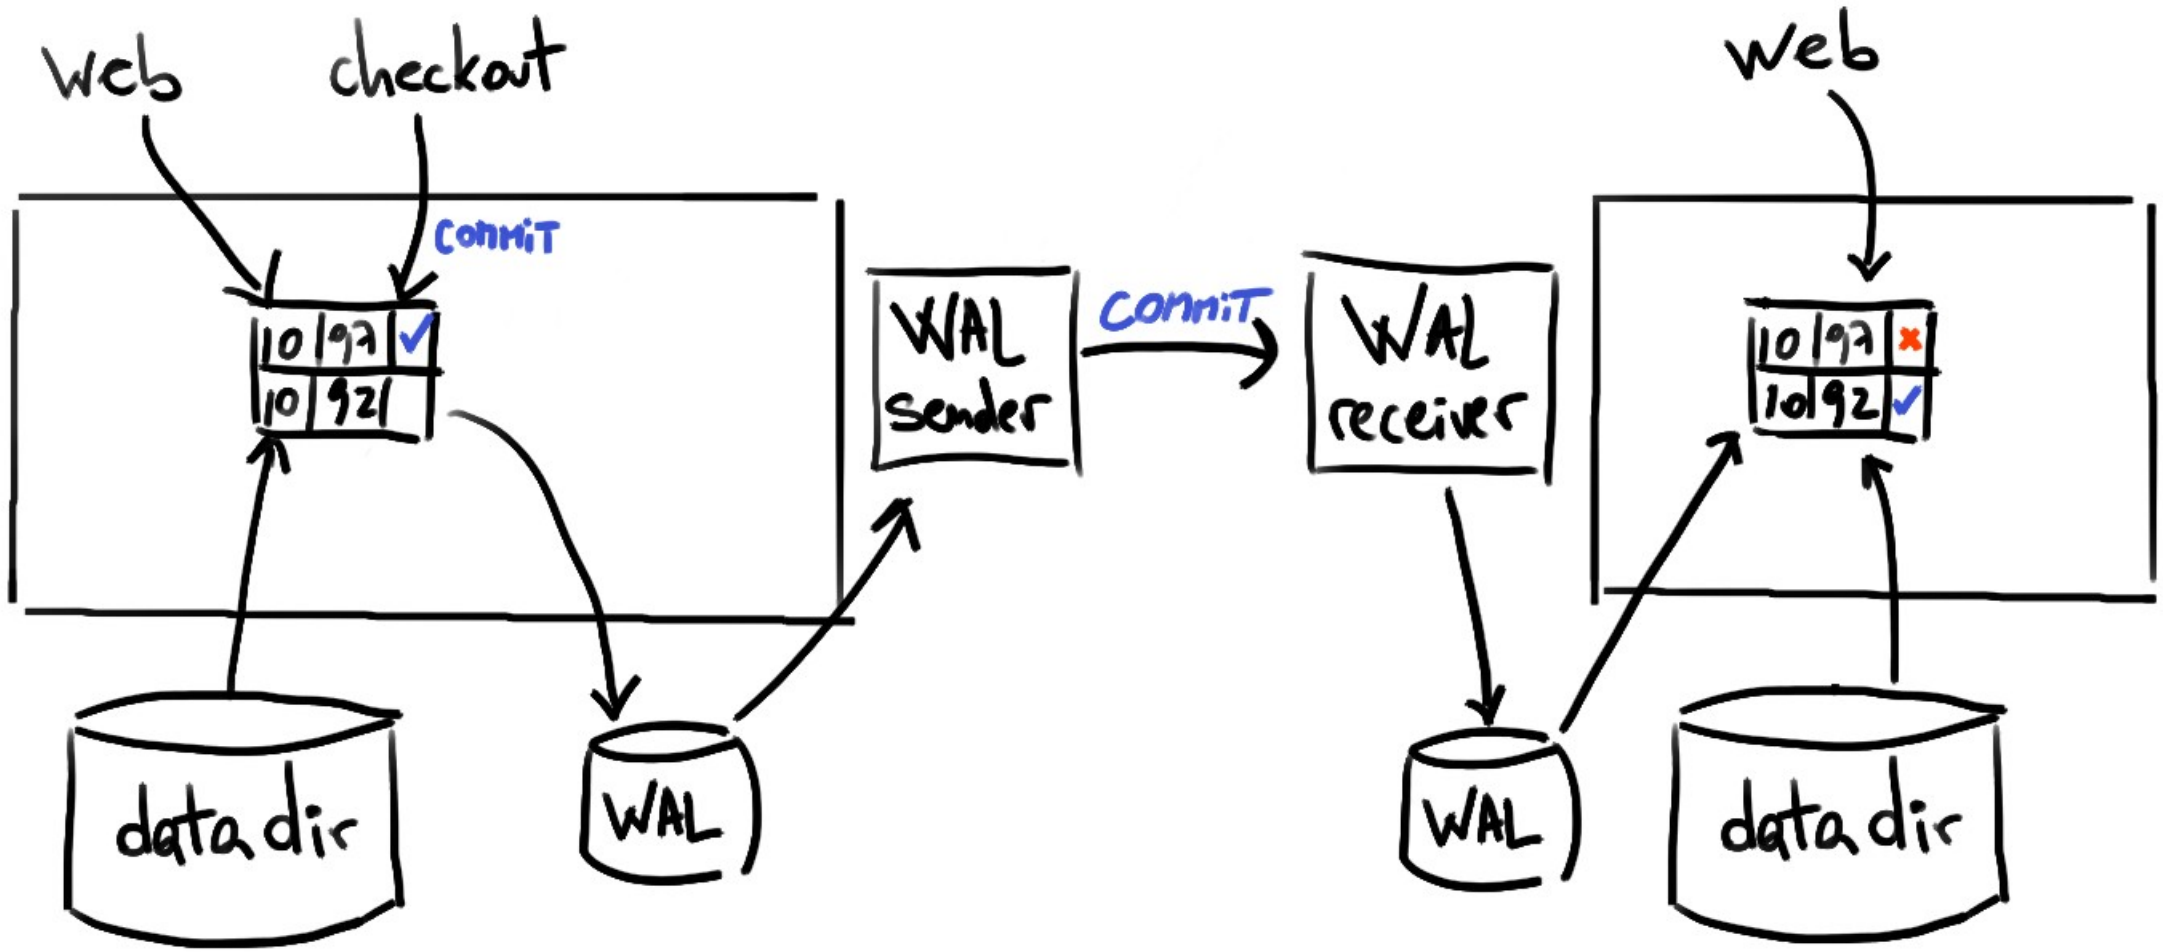

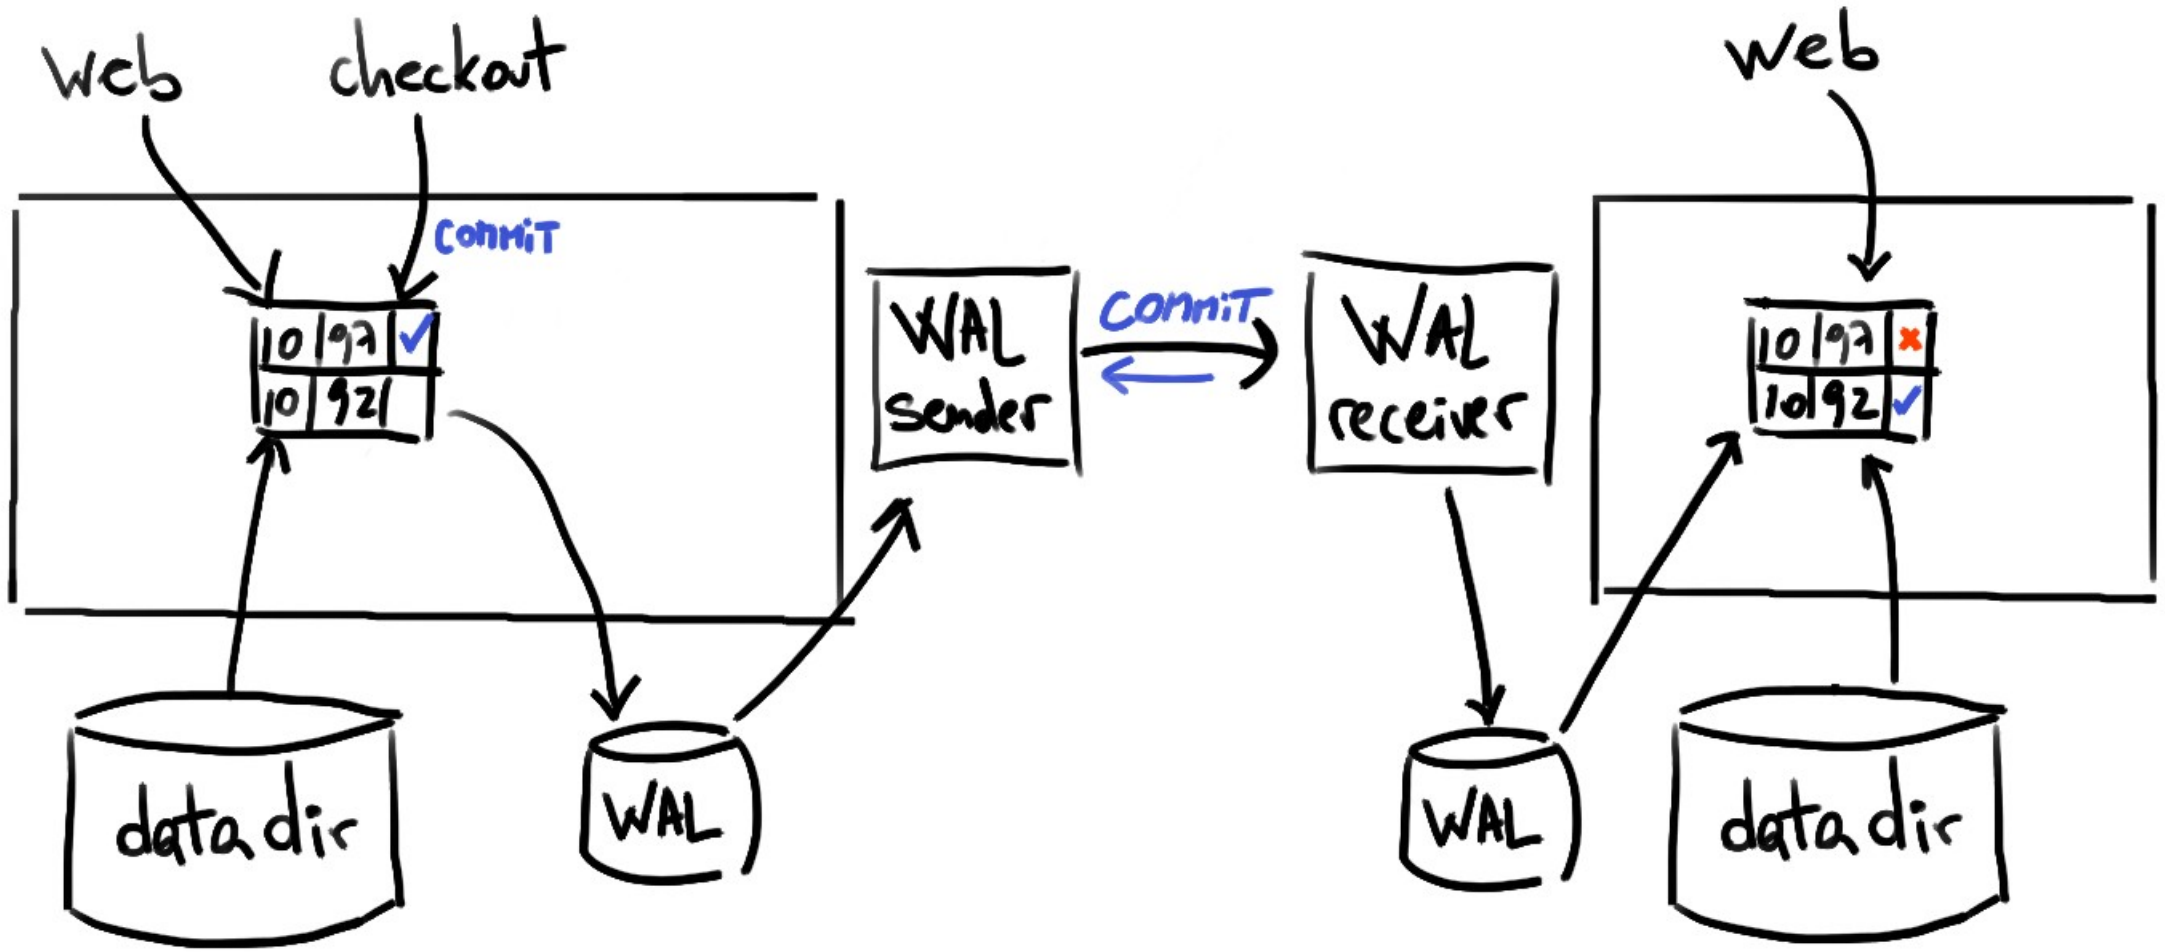

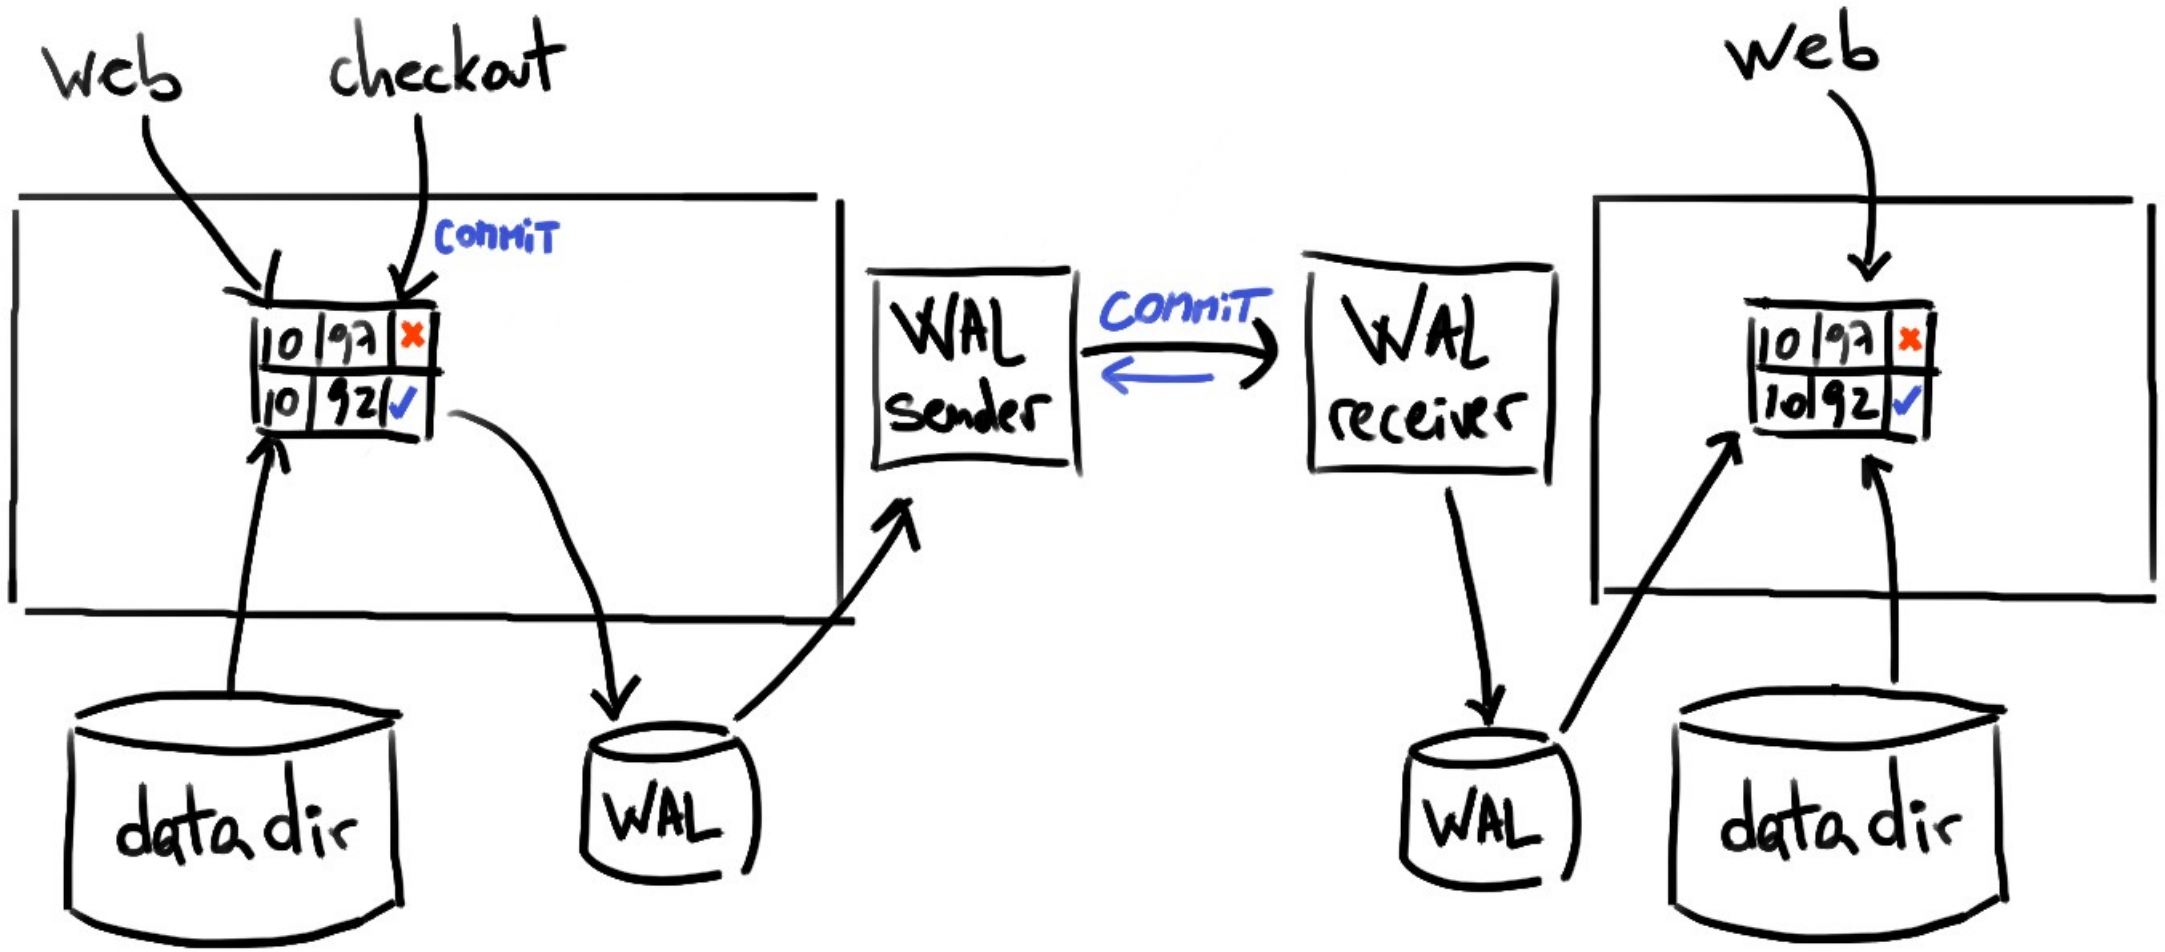

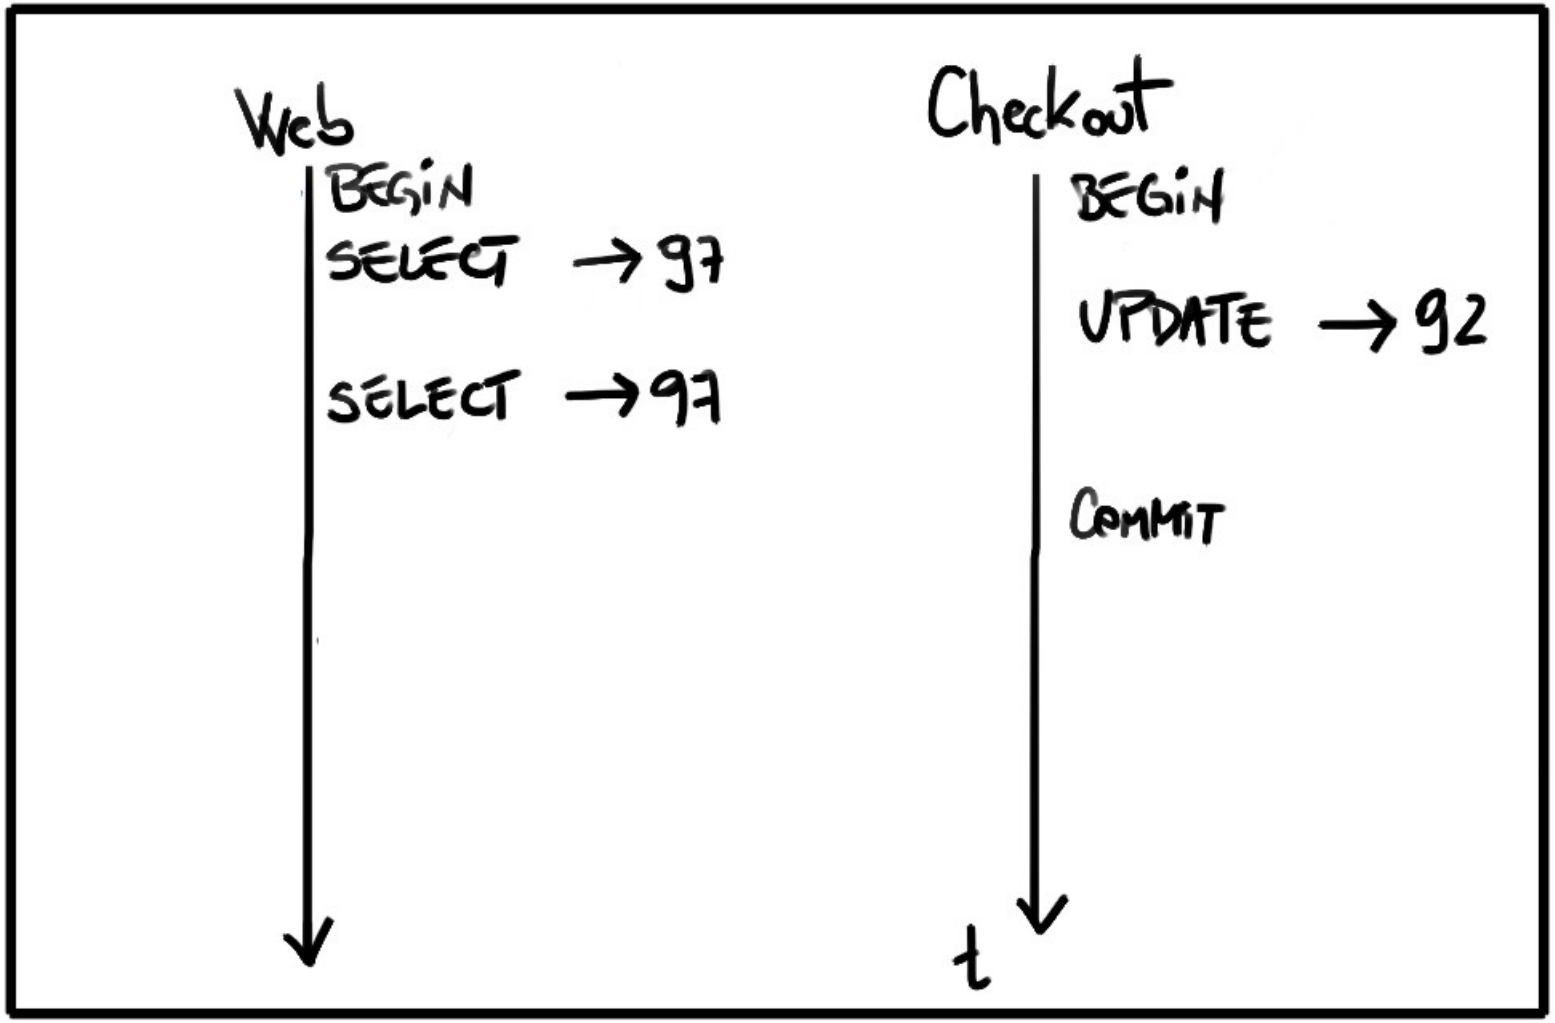

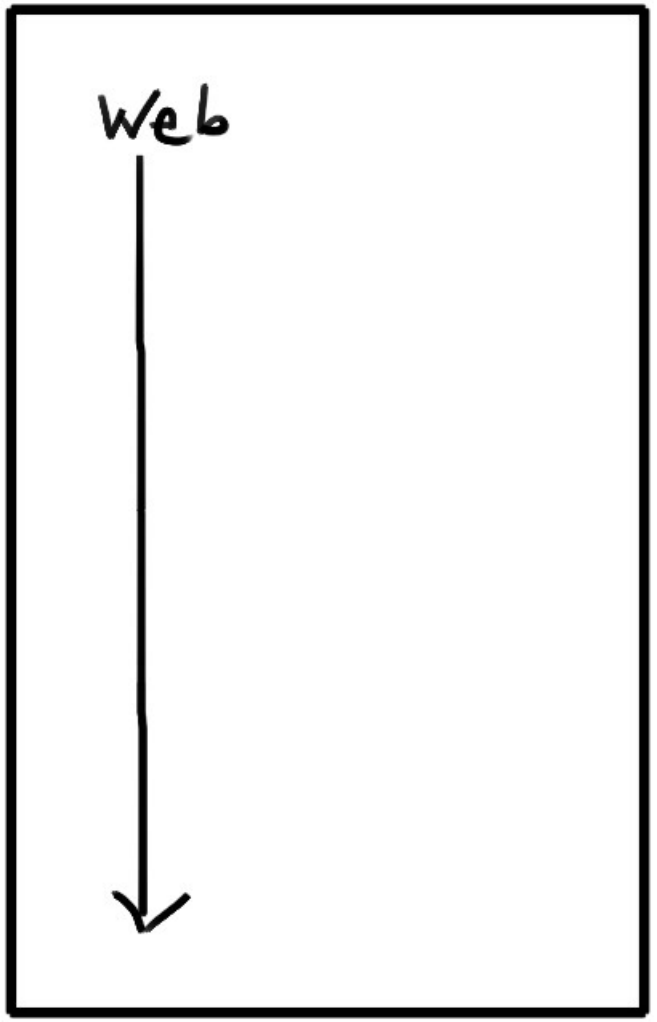

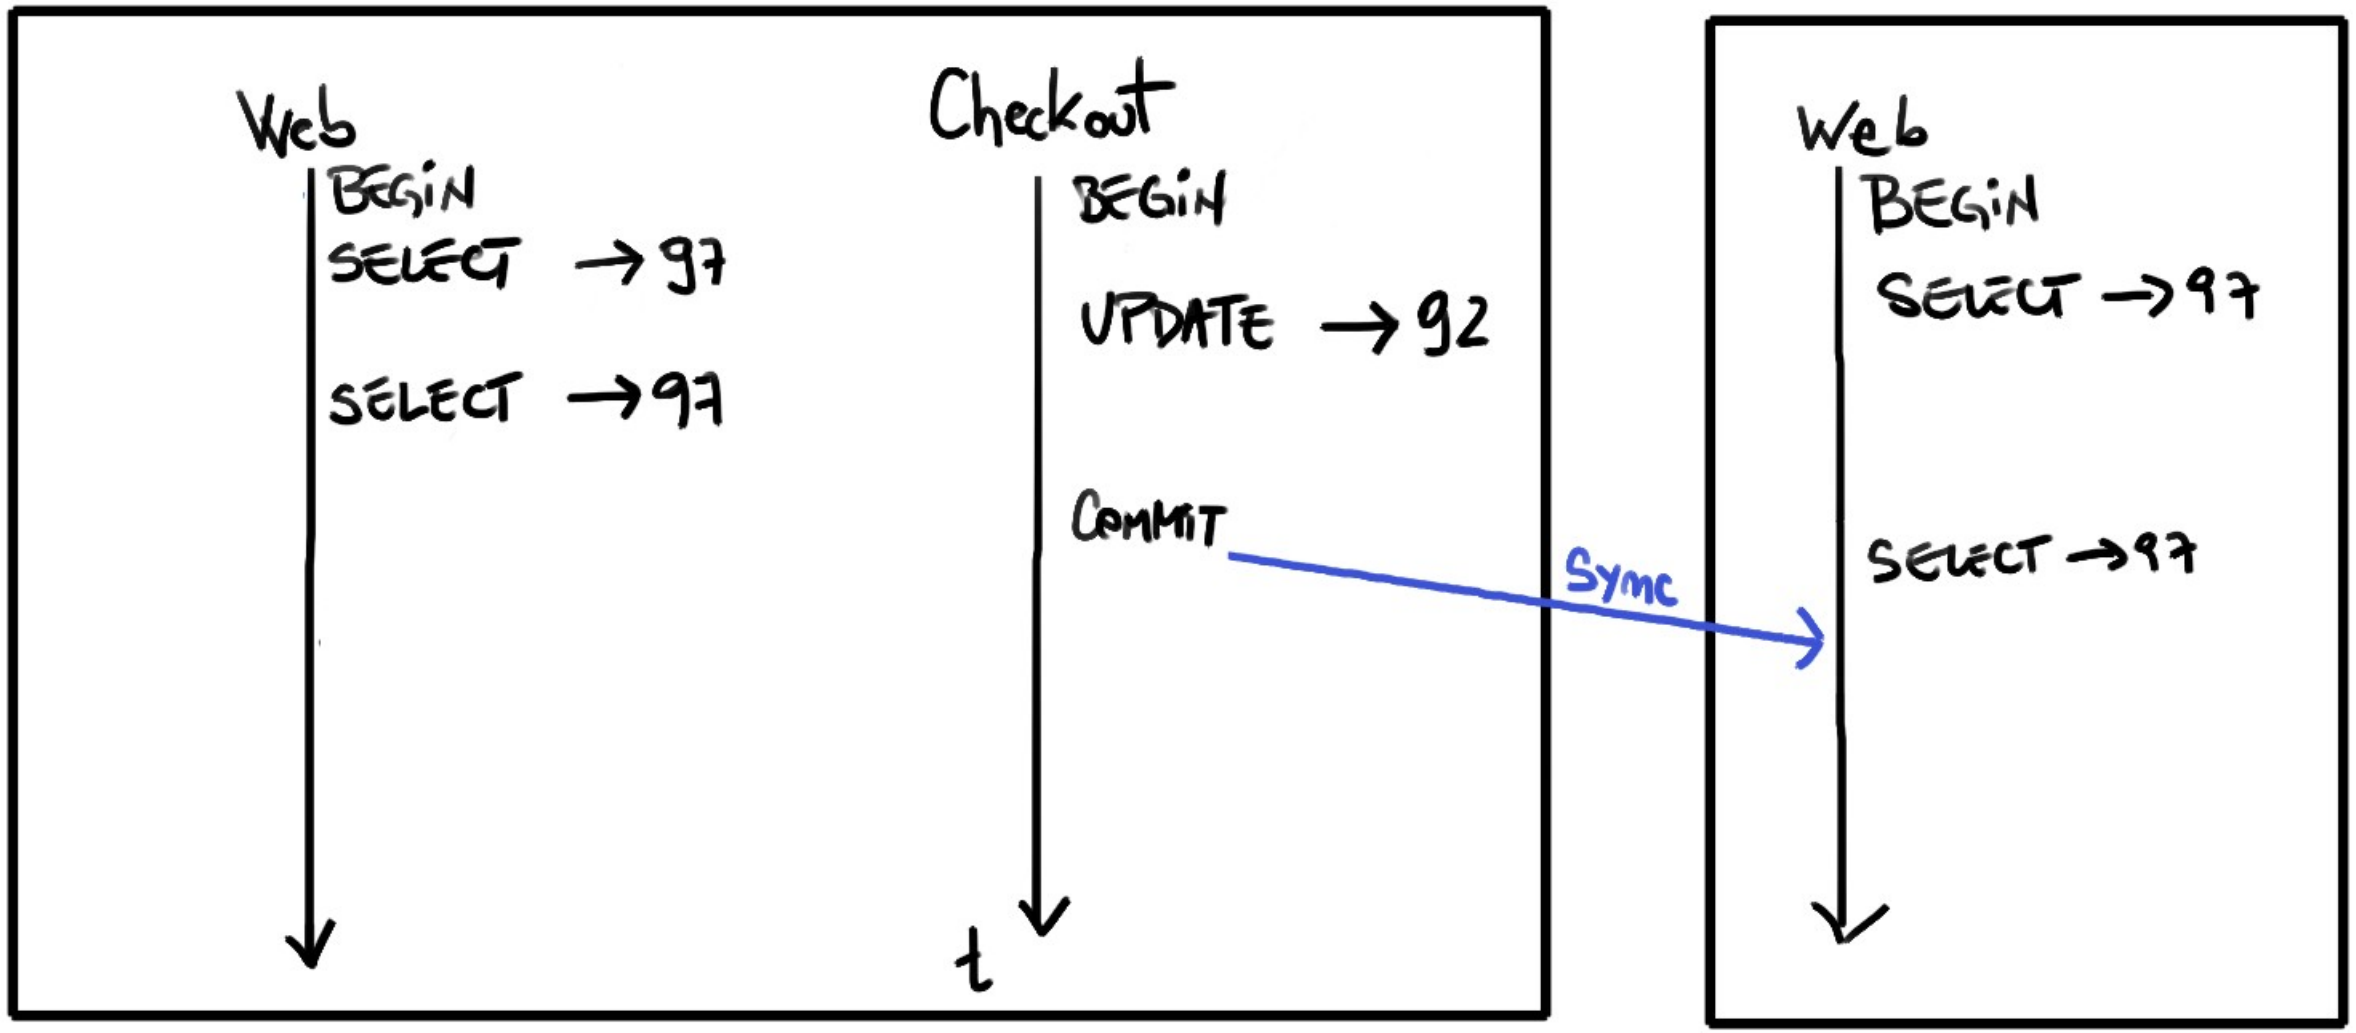

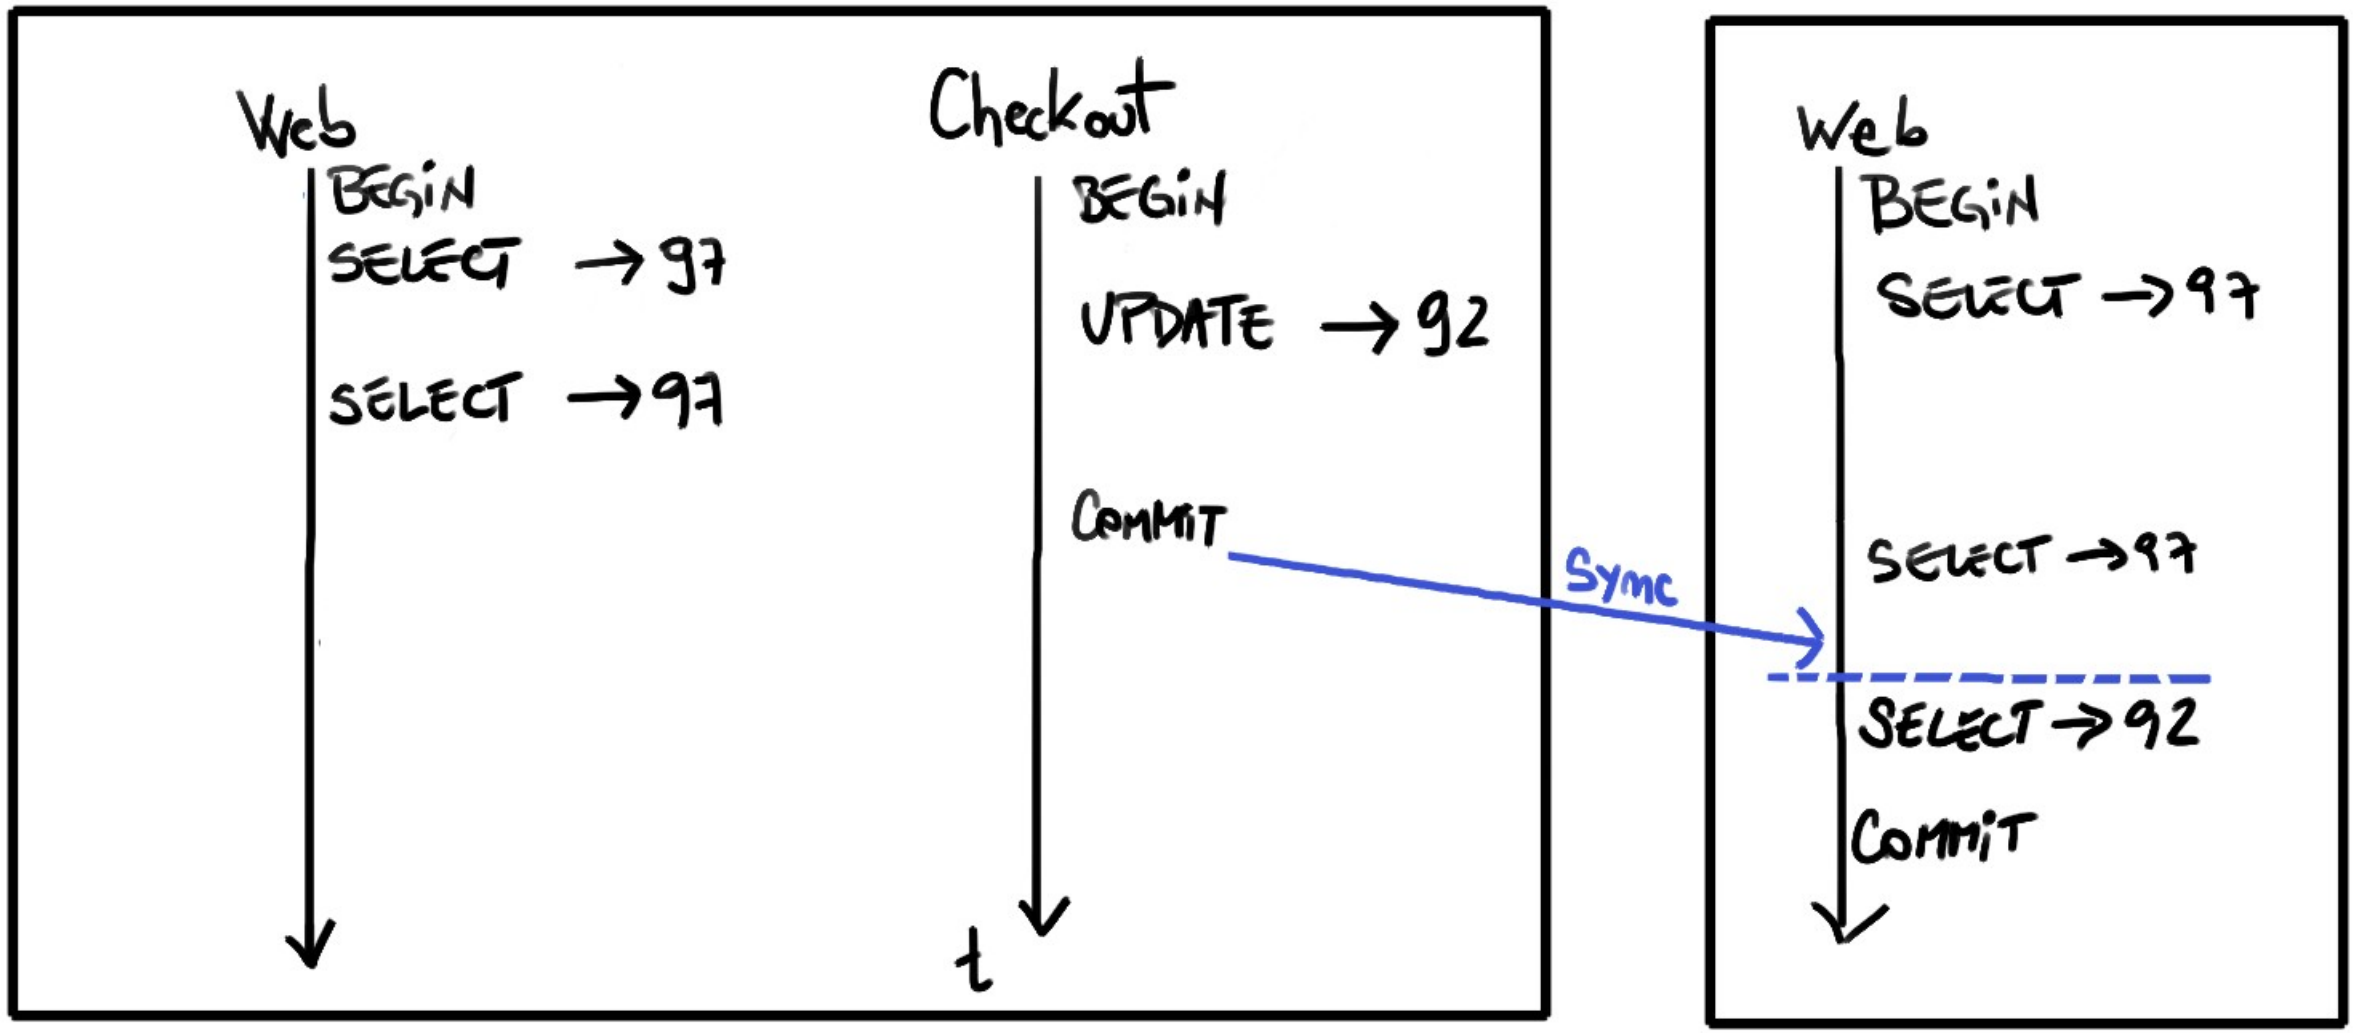

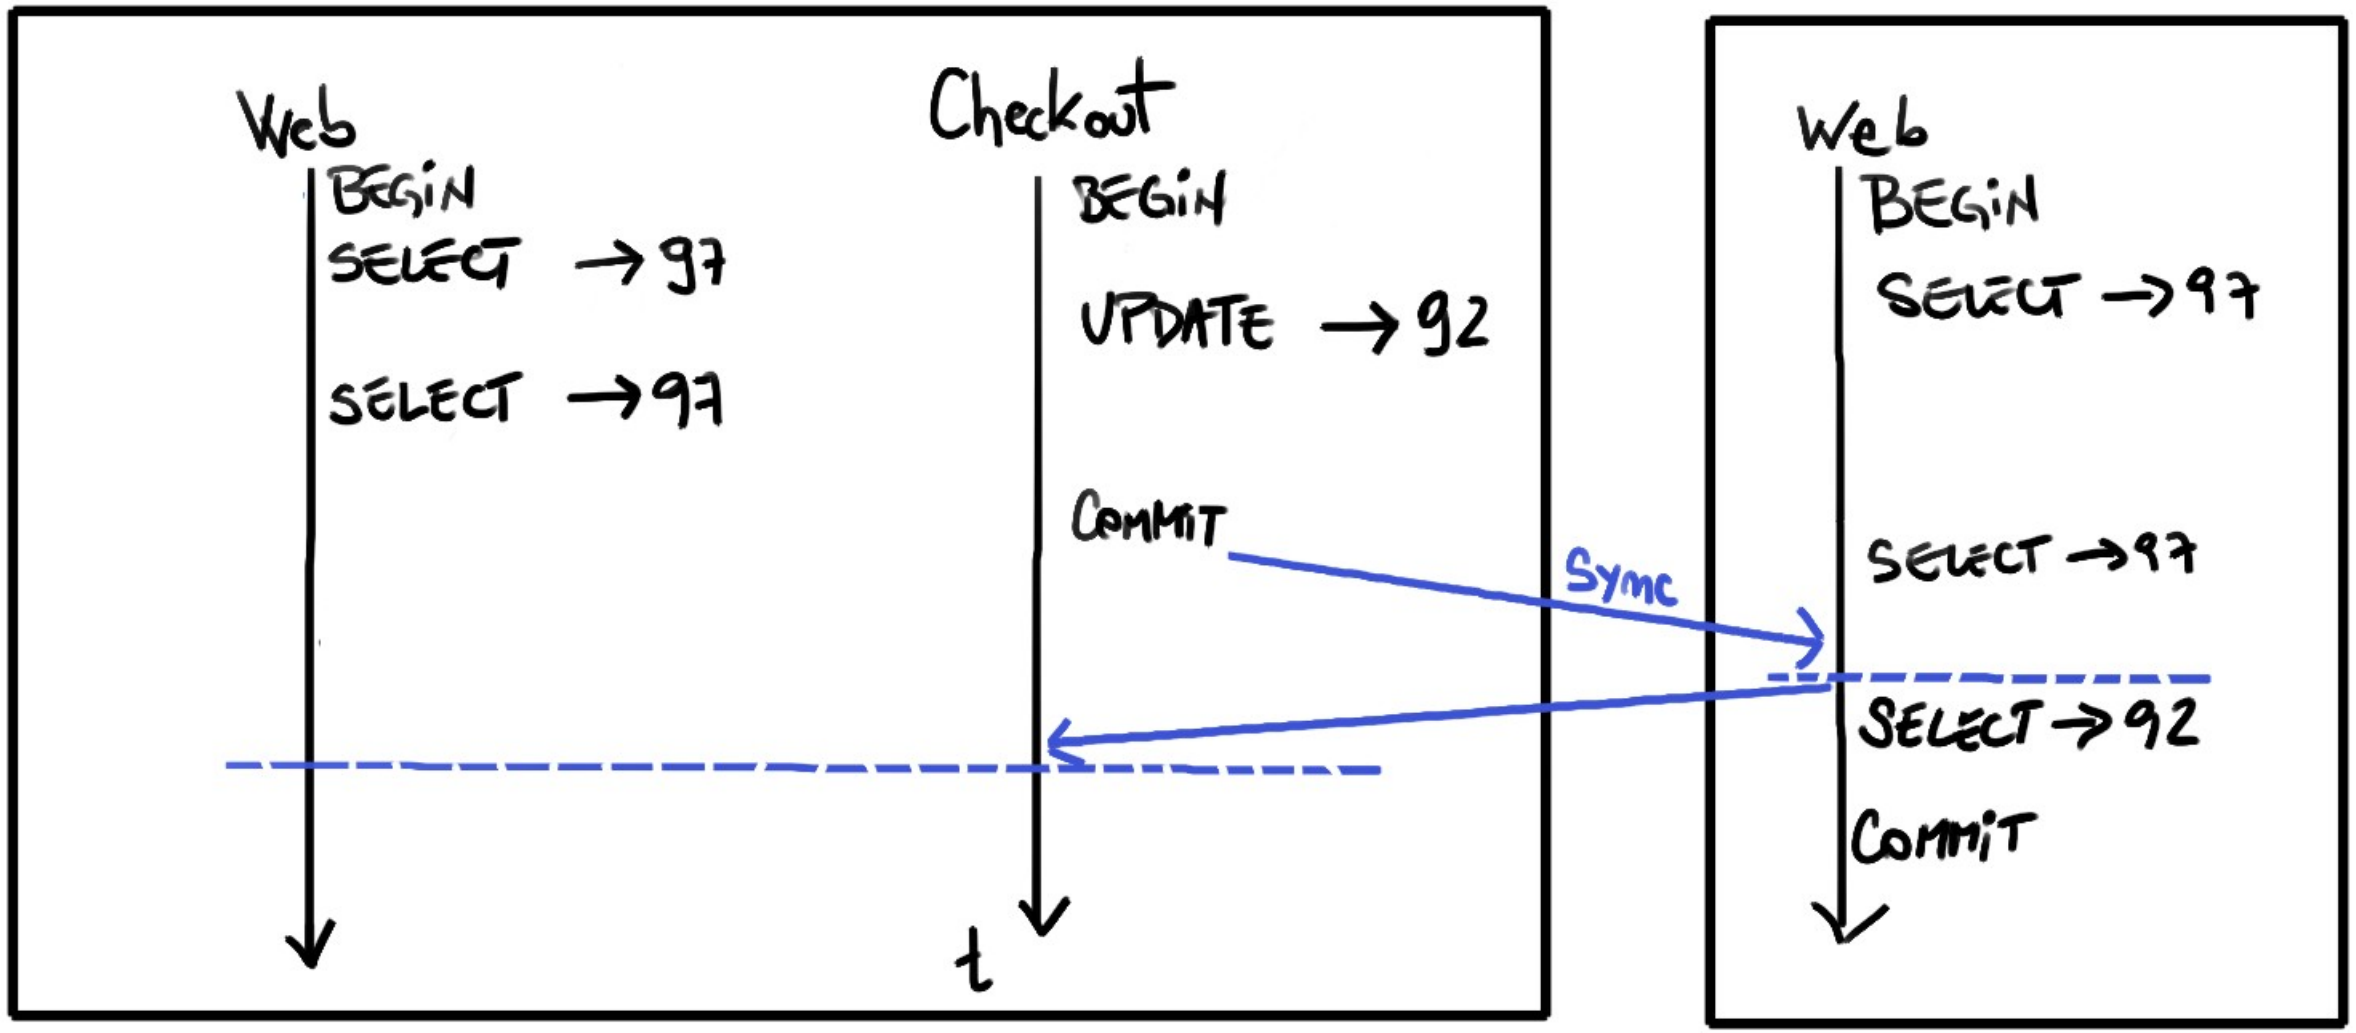

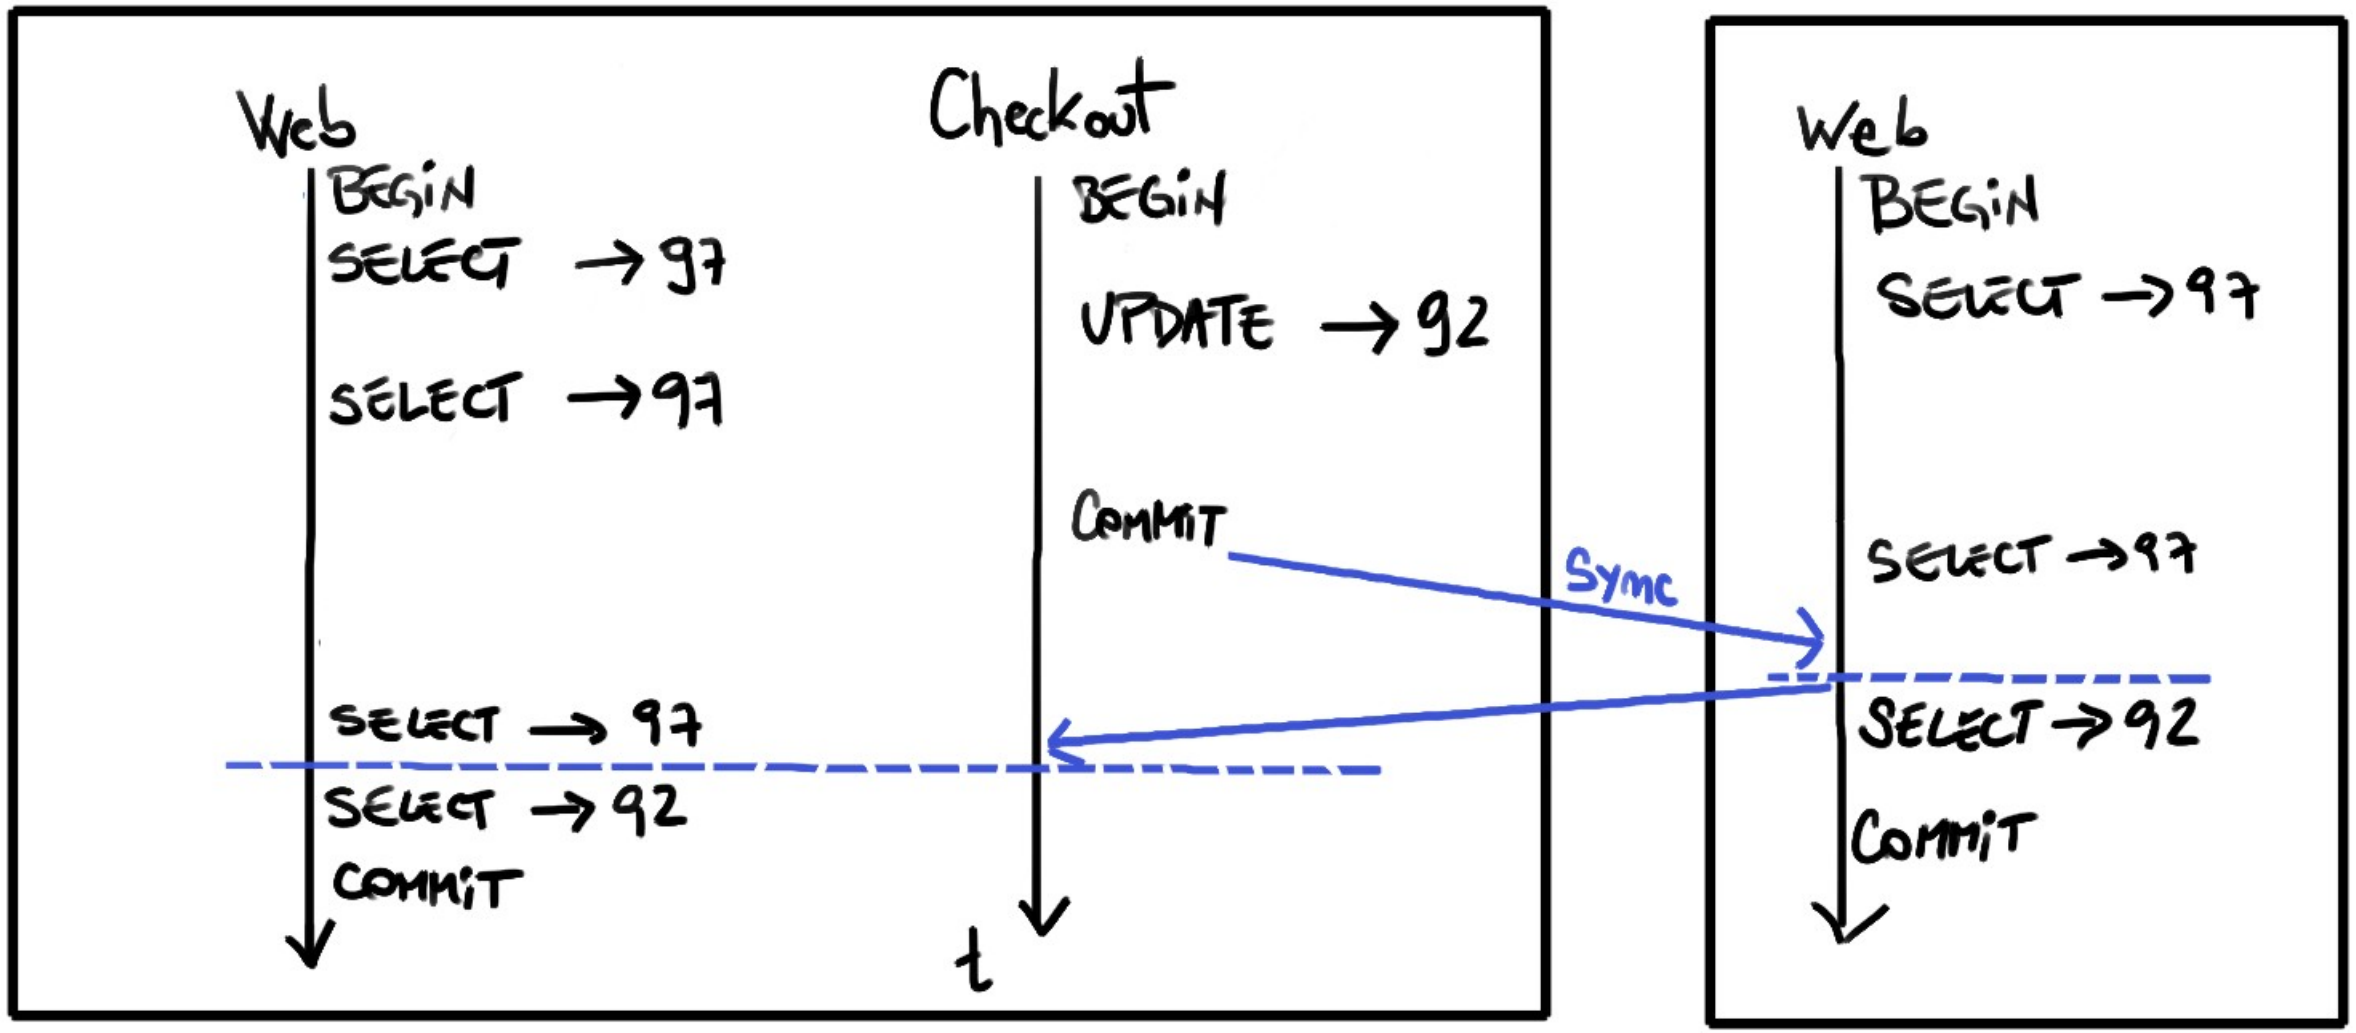

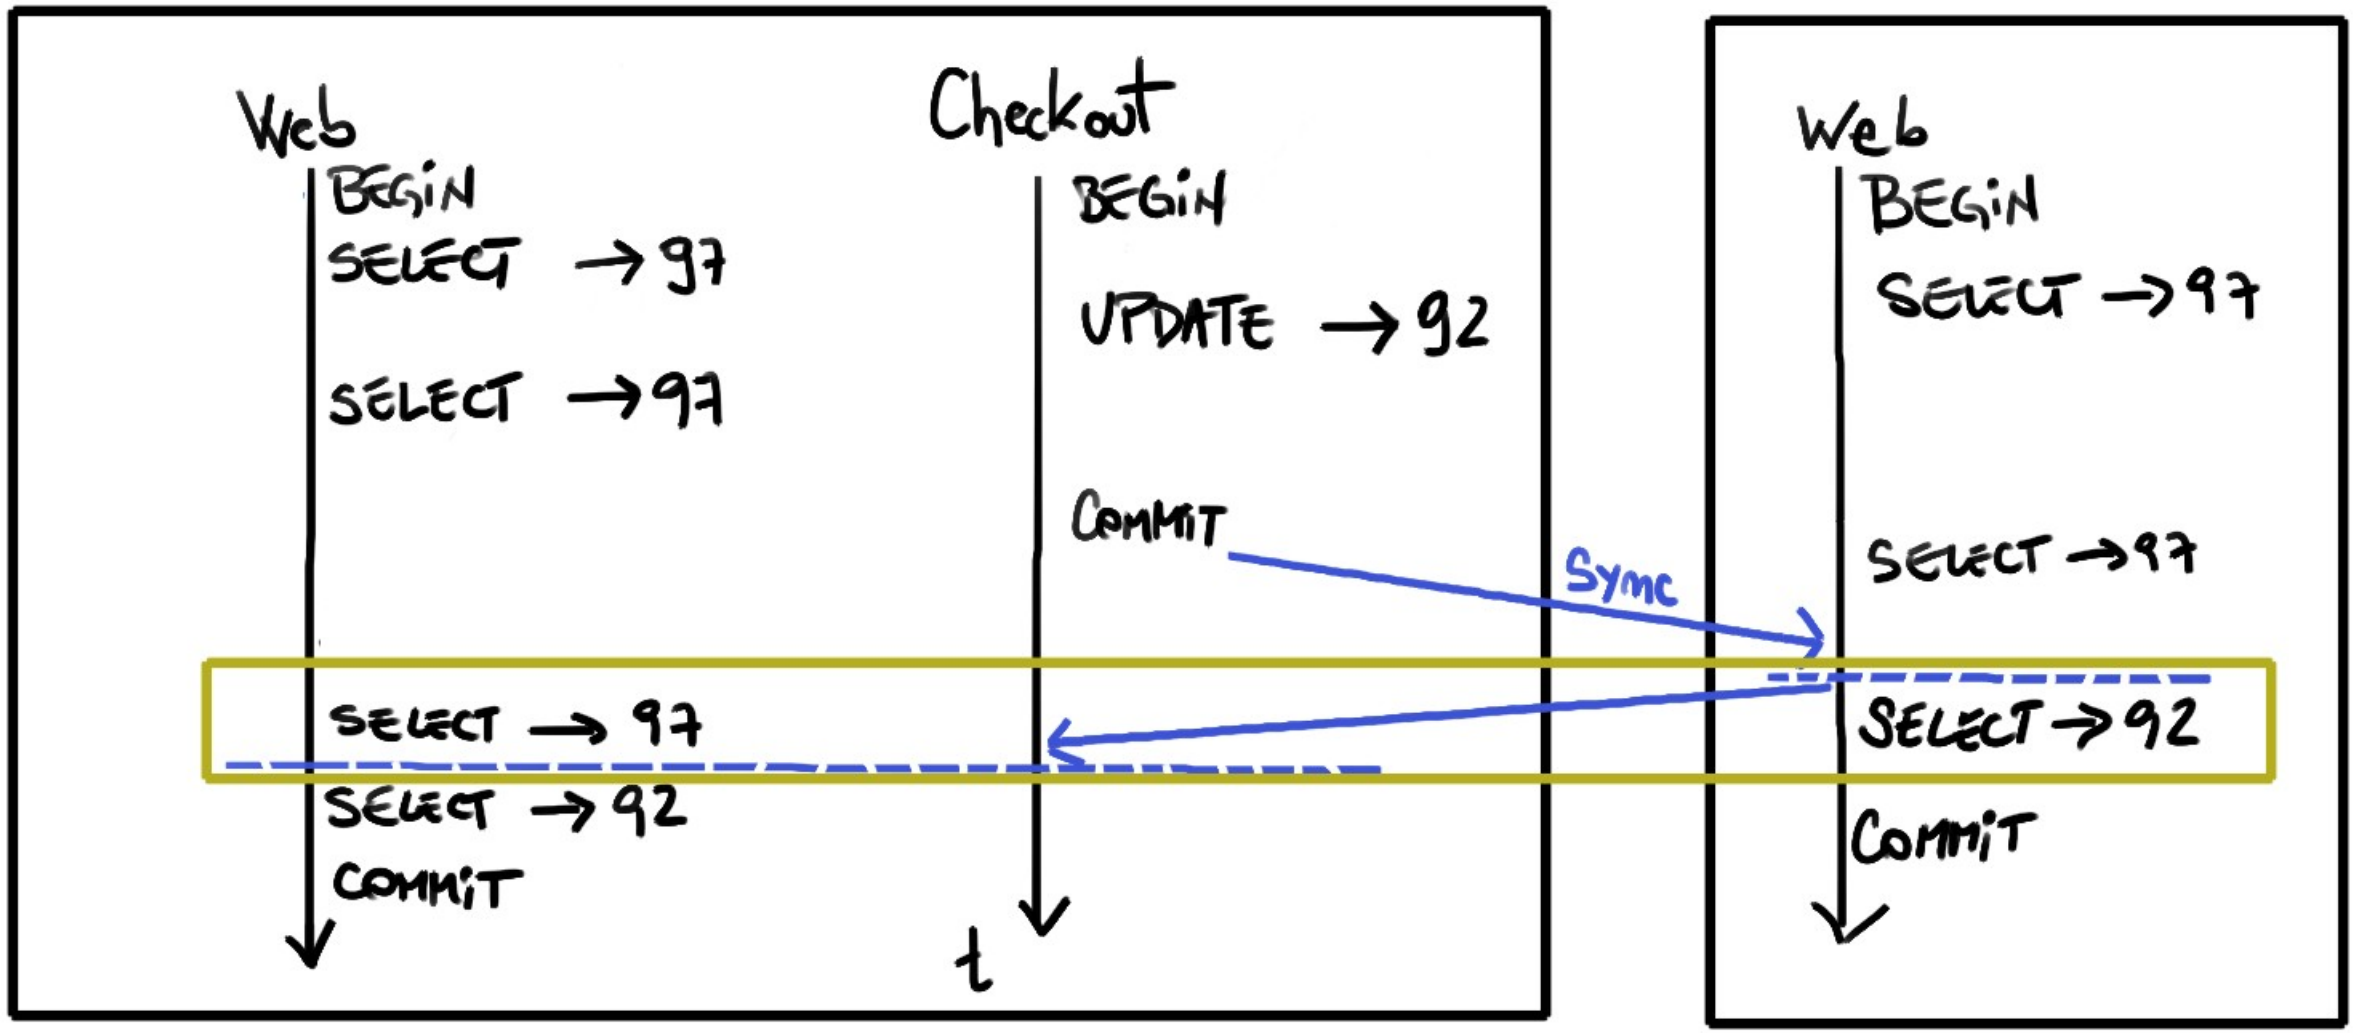

Consistency

Every read receives the most recent write or an error

Eventual Consistency

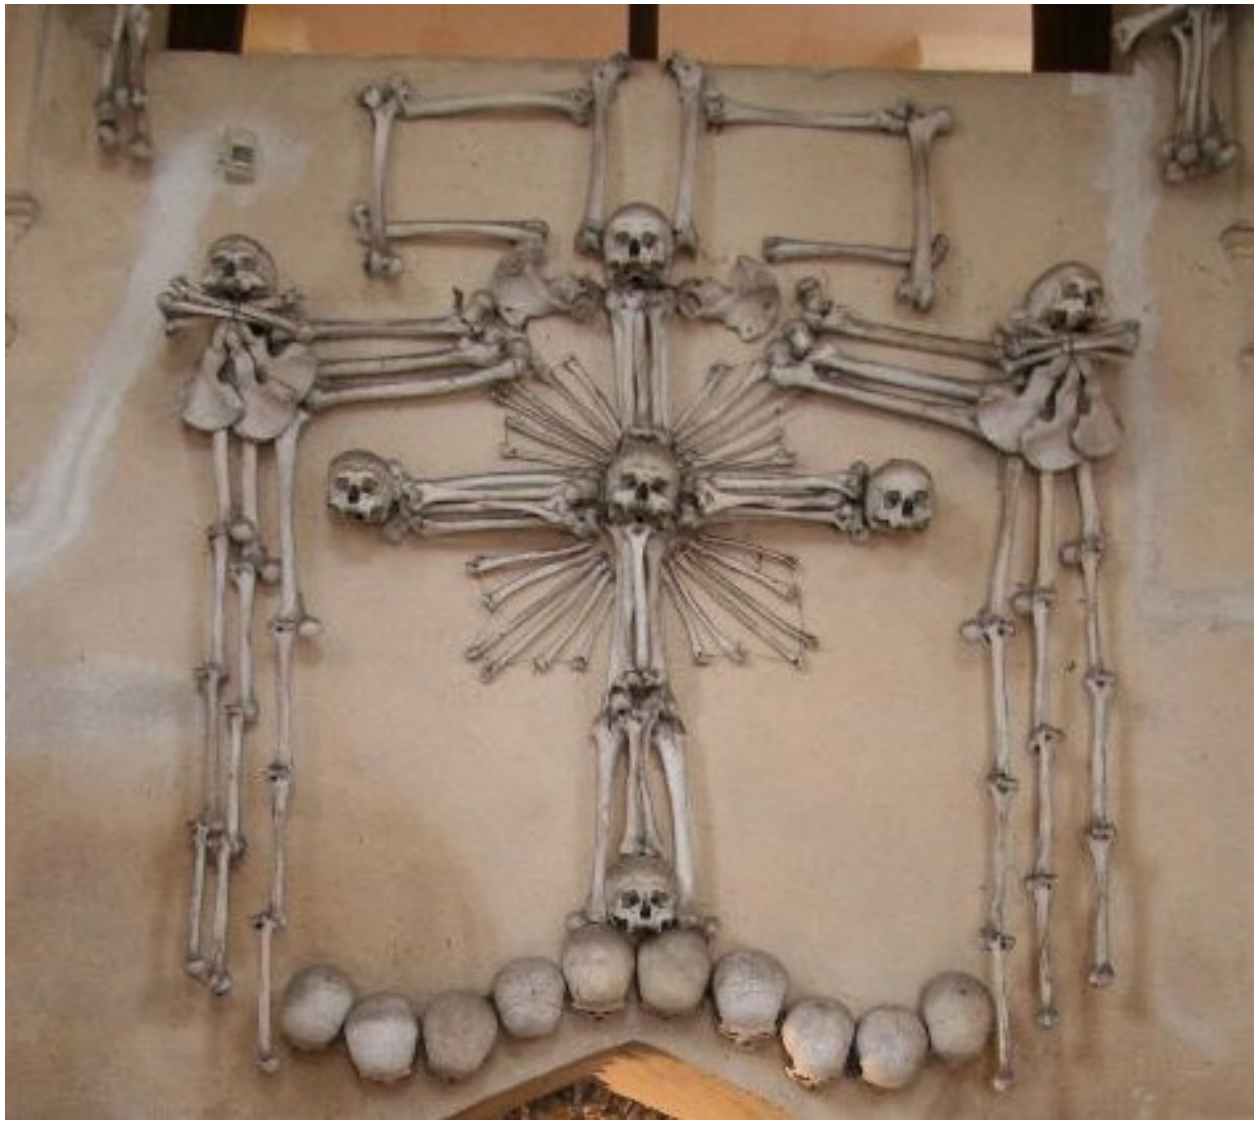

Synchronous Replication reduces data-loss

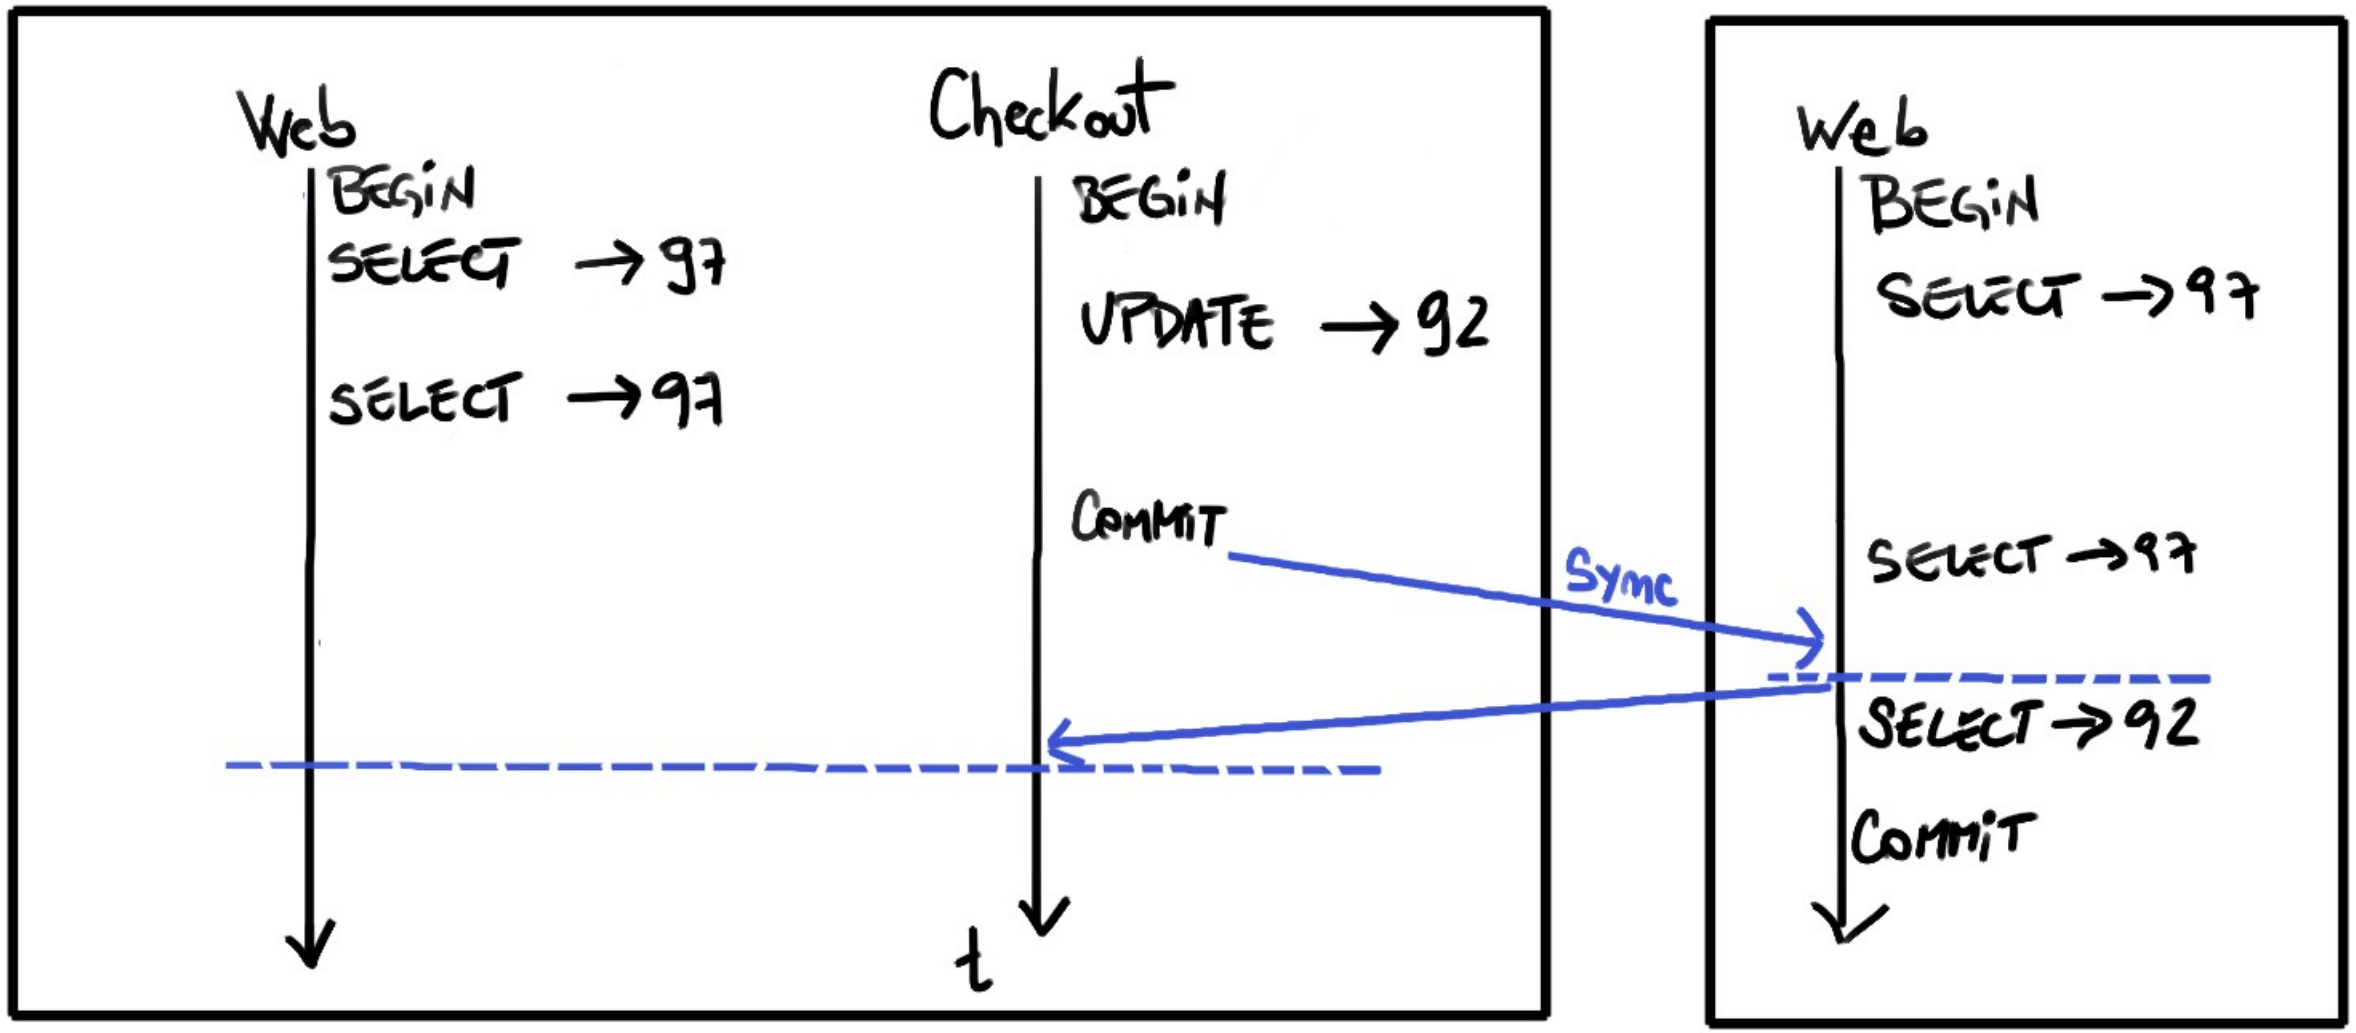

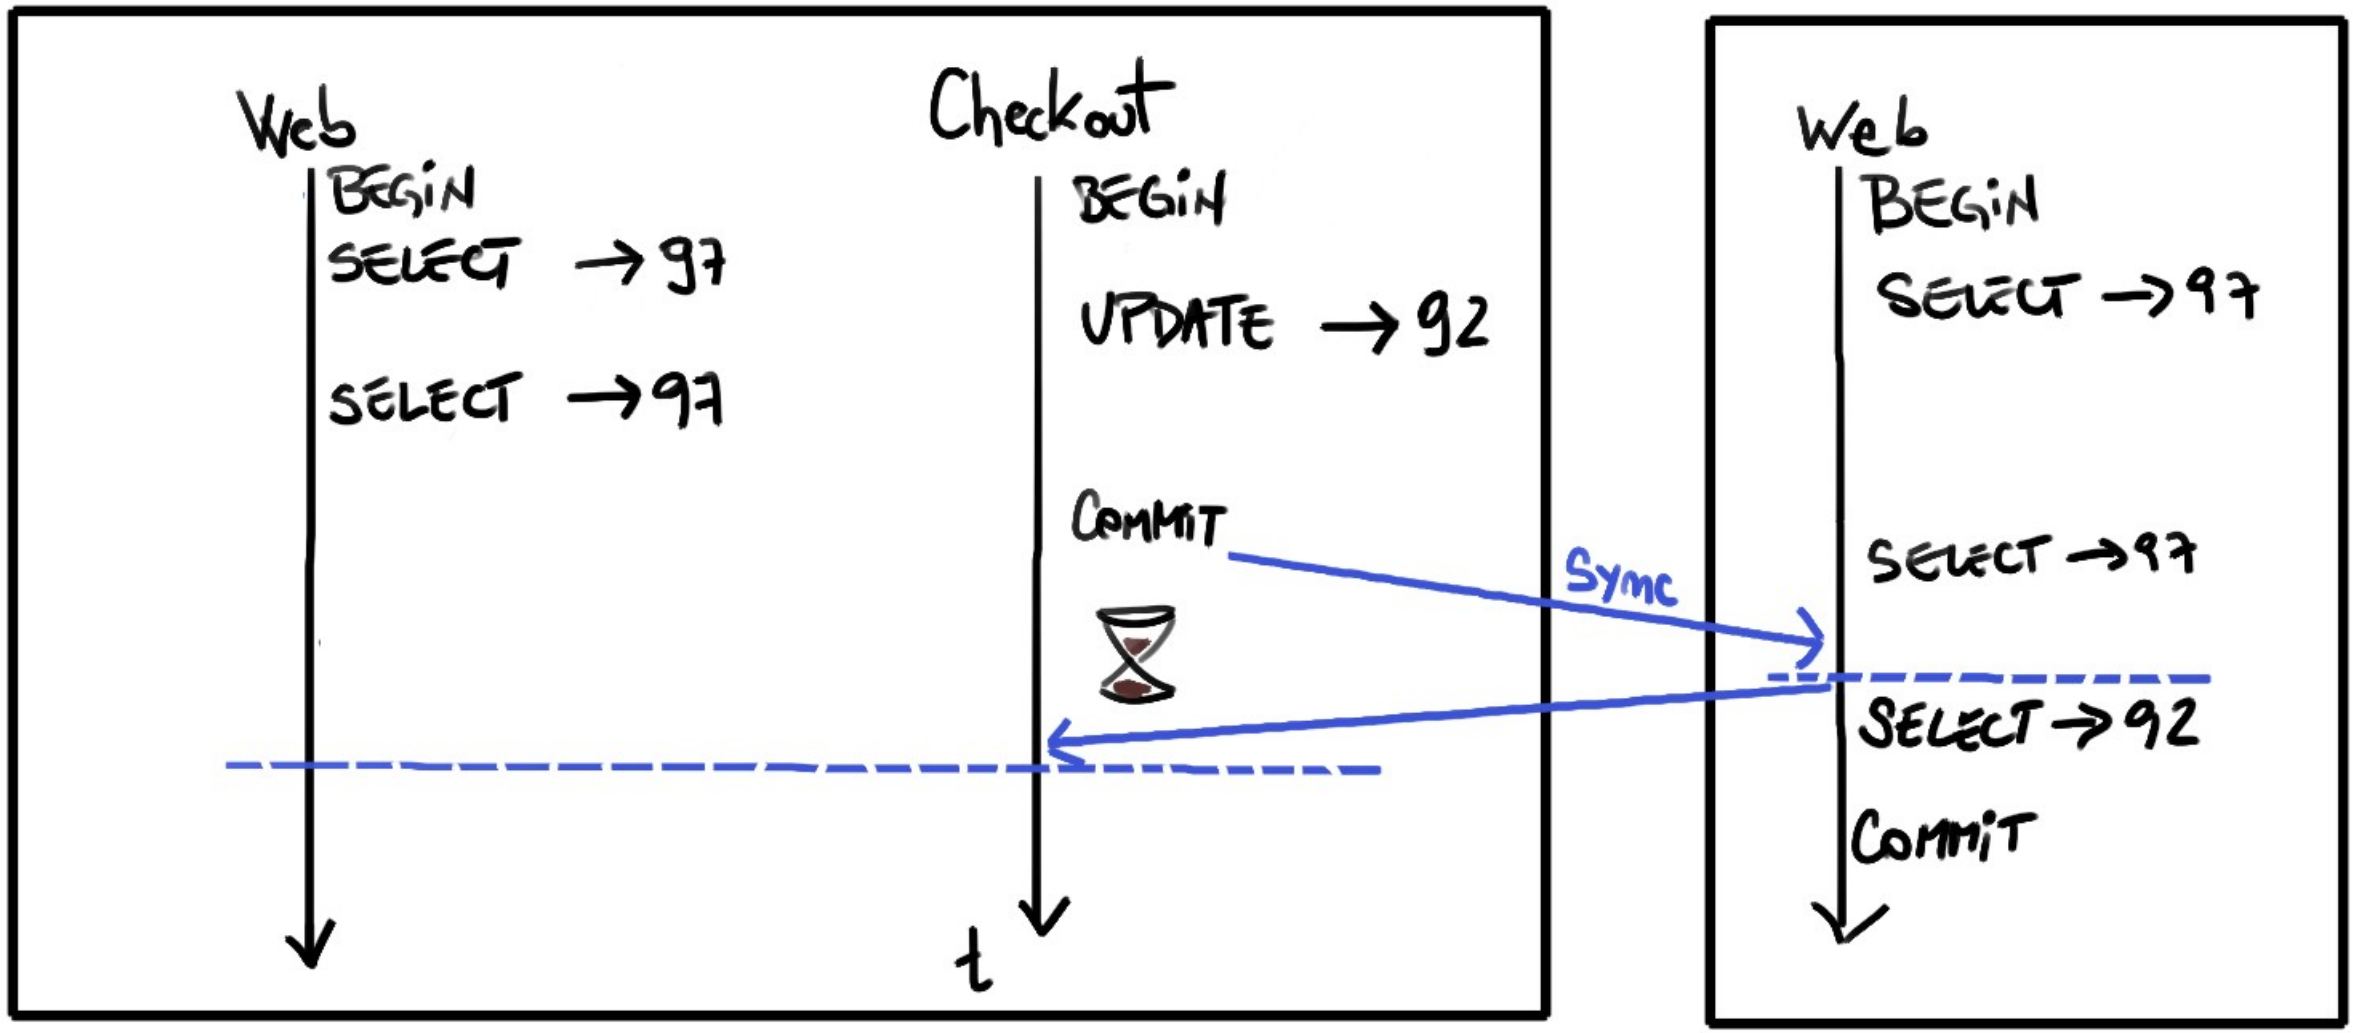

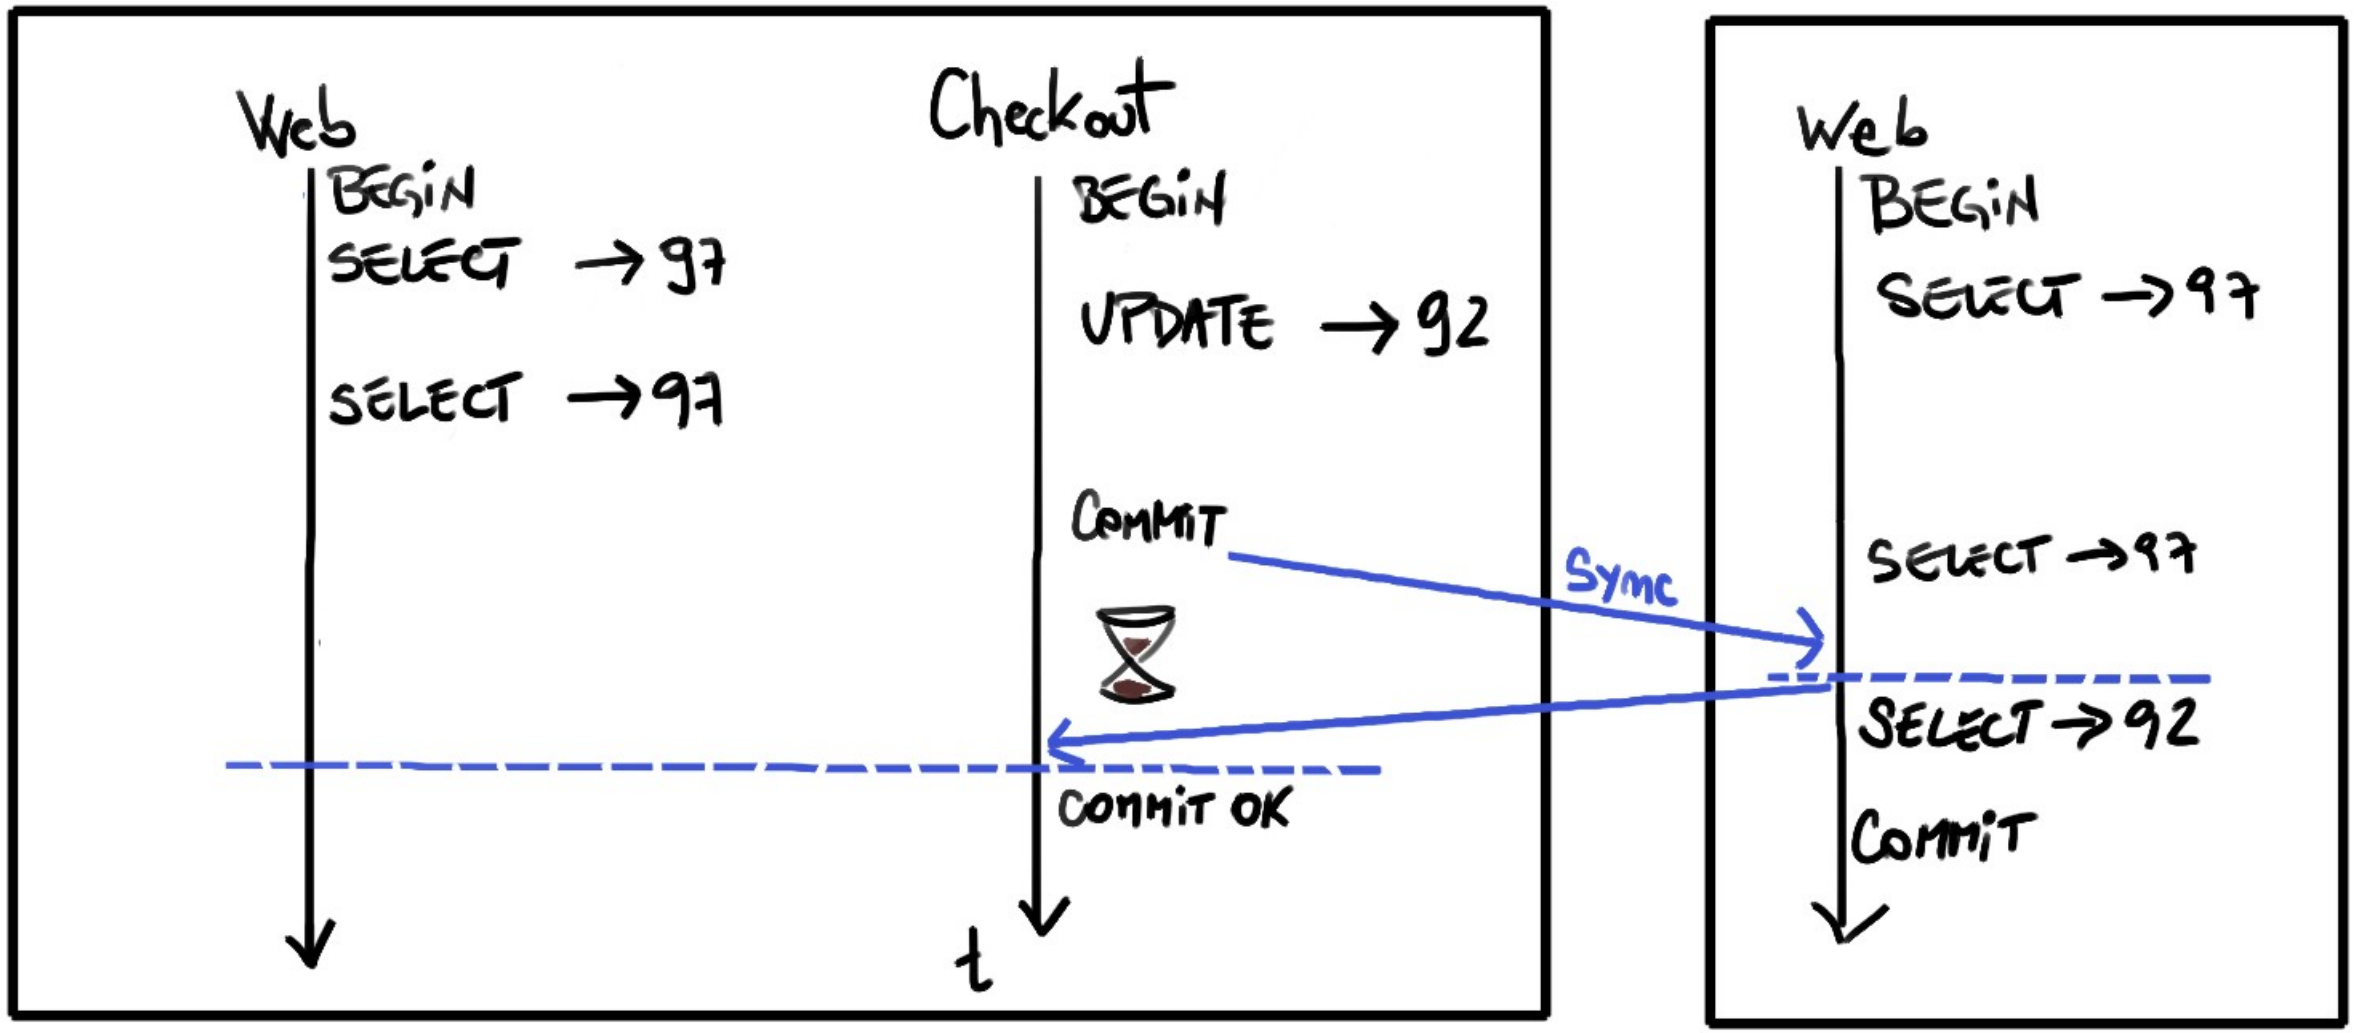

Monitor replication state and lag

Monitor replication state and lag

`pg_catalog.pg_stat_replication`

- `application_name`
- `state`

Monitor replication state and lag

`pg_catalog.pg_stat_replication`

- `application_name`
- `state`
- `sent_lsn`
- `write_lsn`
- `flush_lsn`
- `replay_lsn`

Monitor replication state and lag

`pg_catalog.pg_stat_replication`

```
application_name
state

sent_lsn

write_lsn
flush_lsn
replay_lsn

write_lag
flush_lag
replay_lag
```

Monitor replication state and lag

`pg_catalog.pg_stat_replication`

```
application_name
state

sent_lsn

write_lsn
flush_lsn
replay_lsn

write_lag
flush_lag
replay_lag

sync_priority
sync_state
reply_time
```

Logical Replication

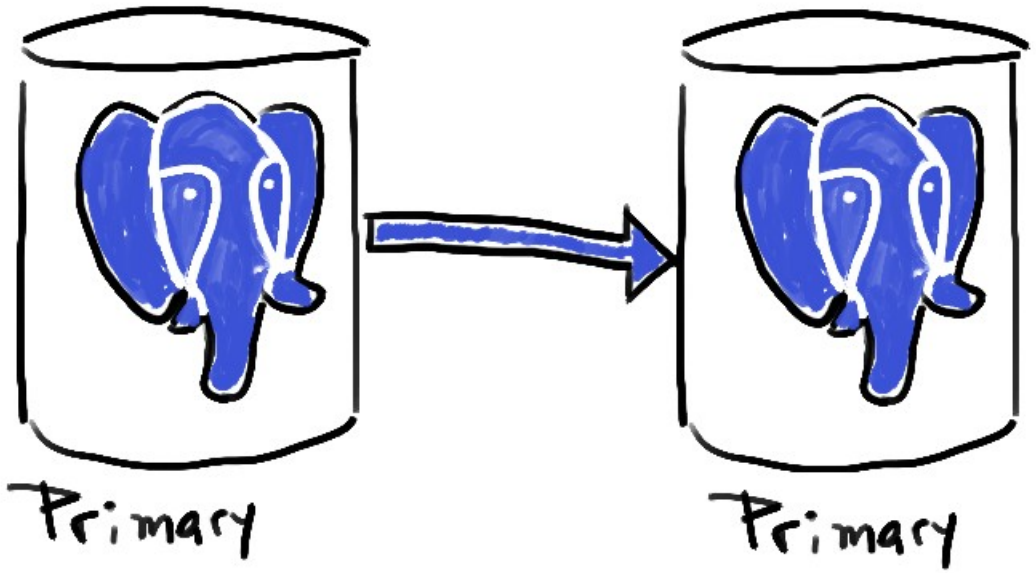

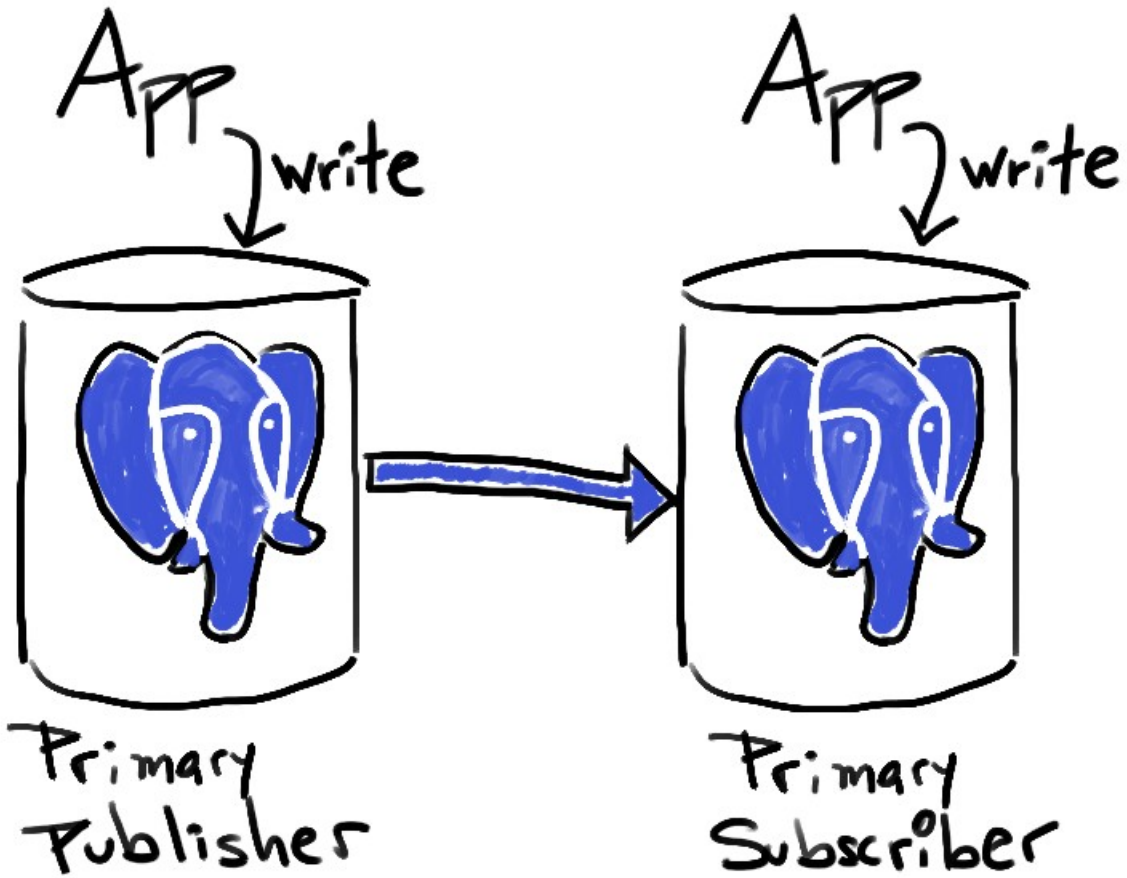

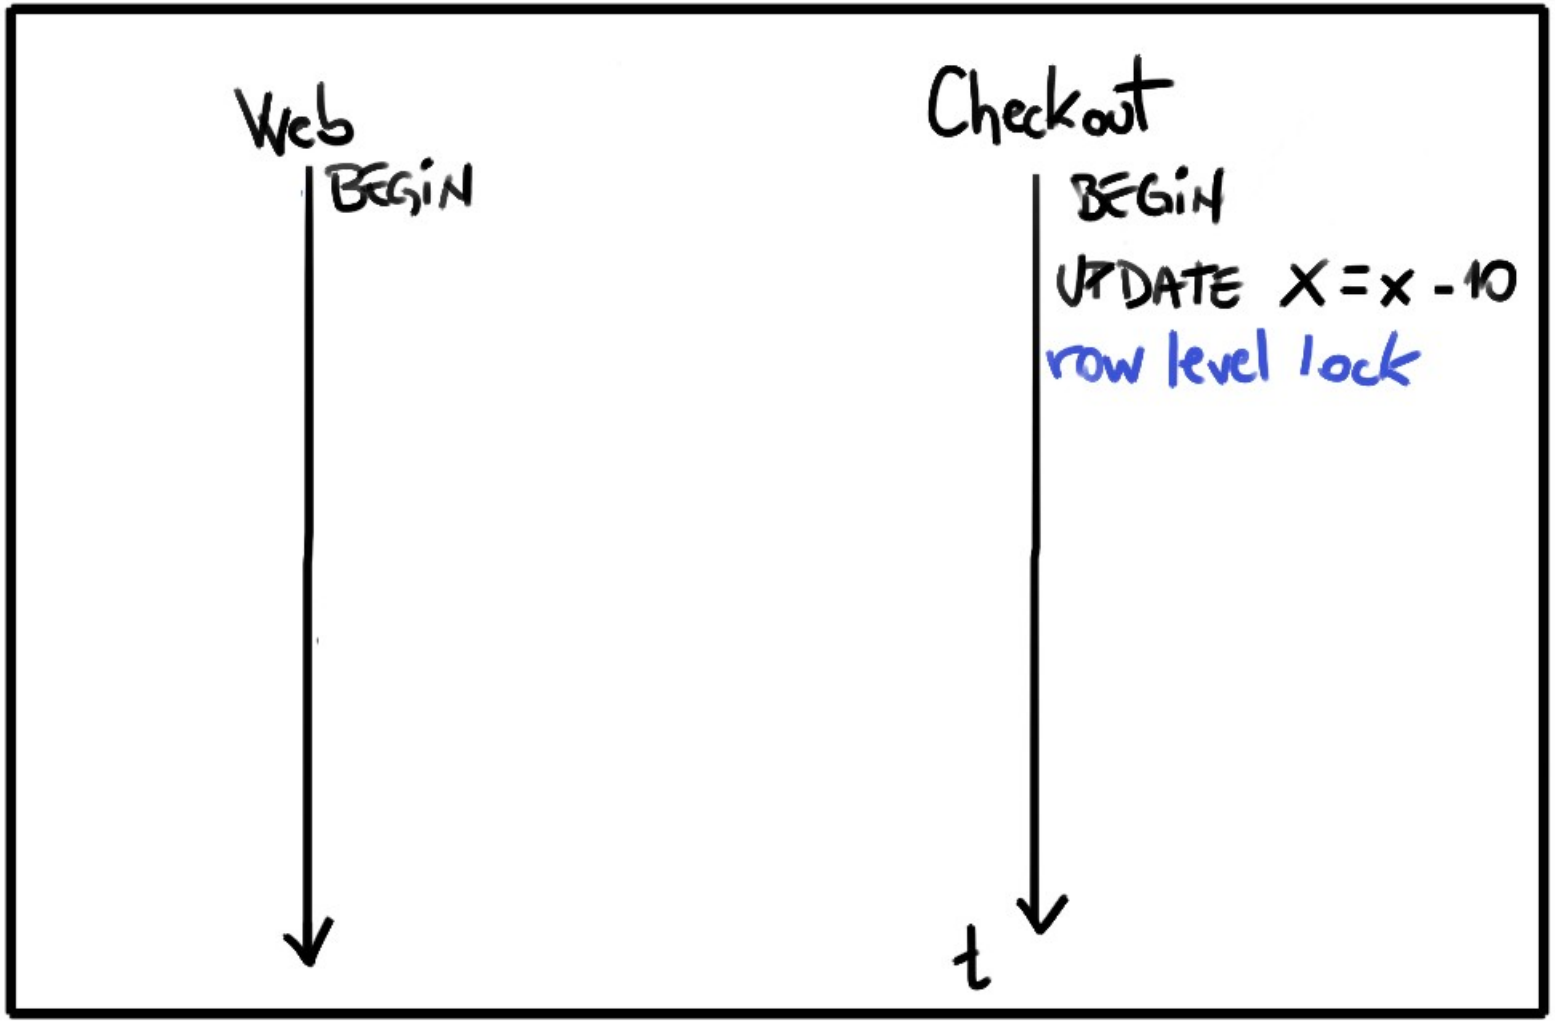

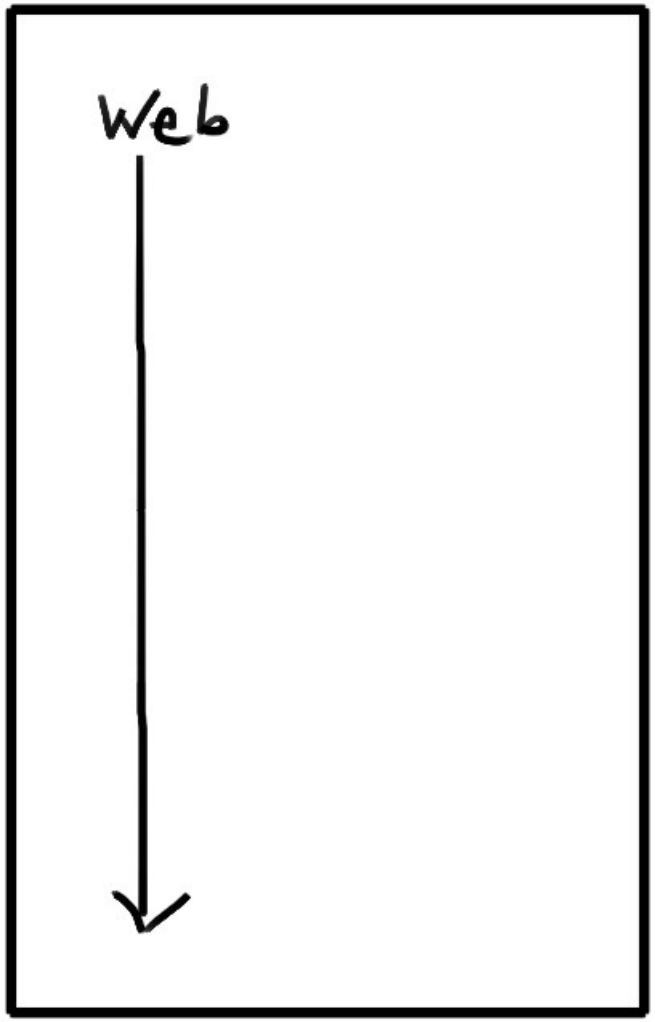

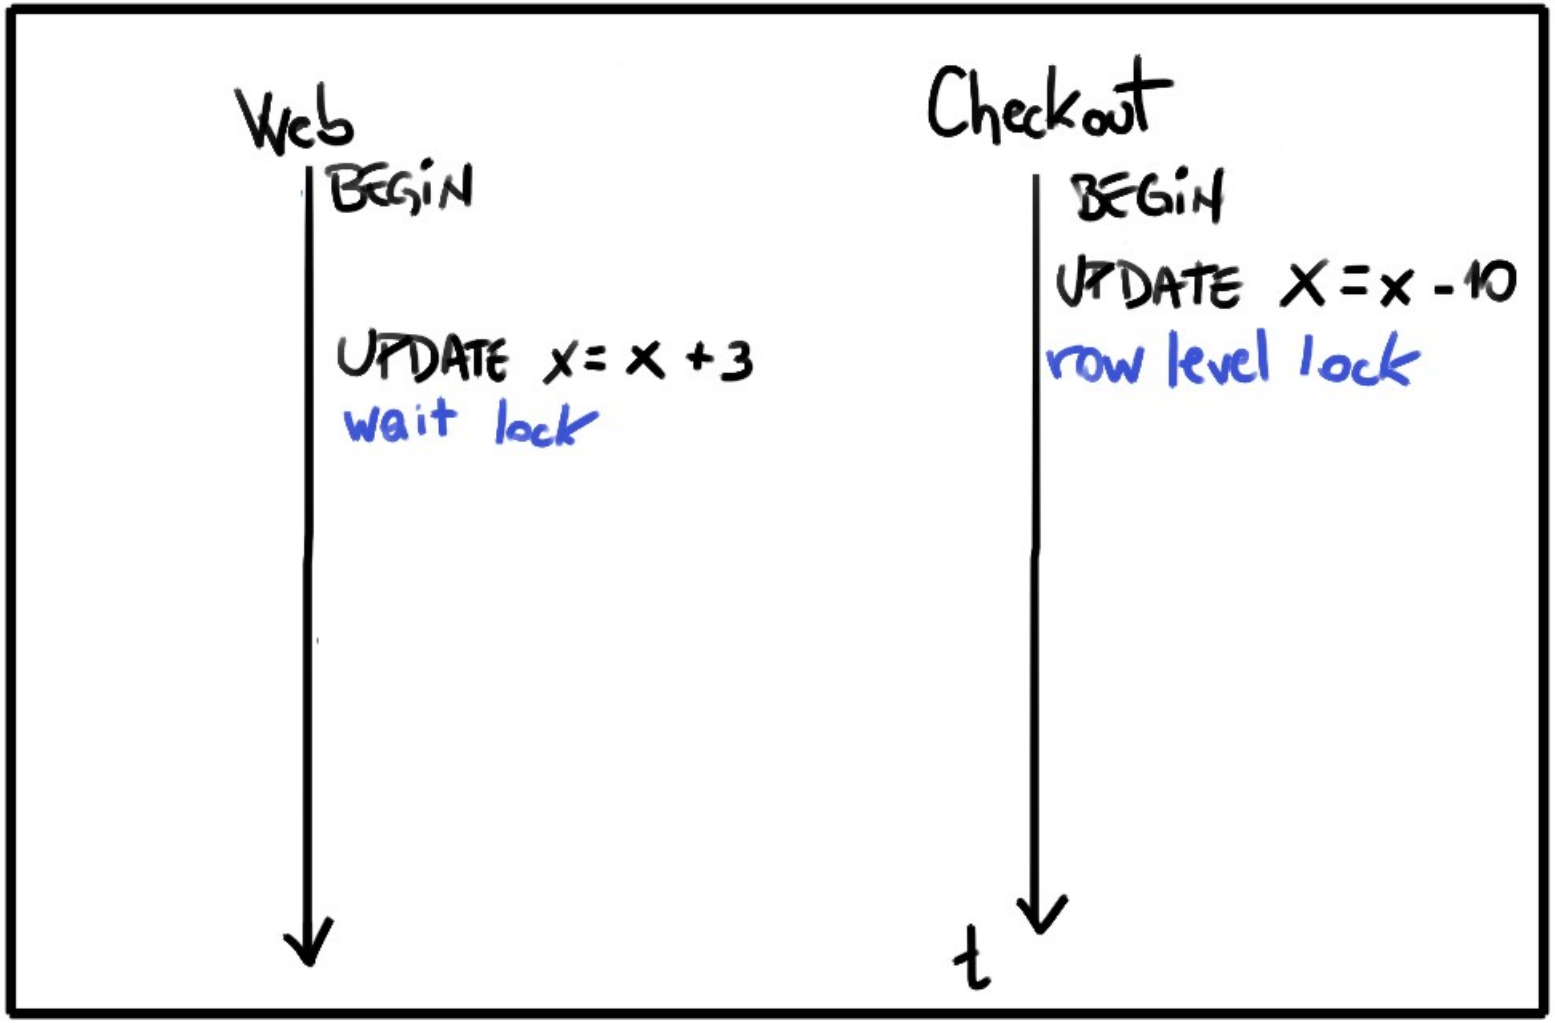

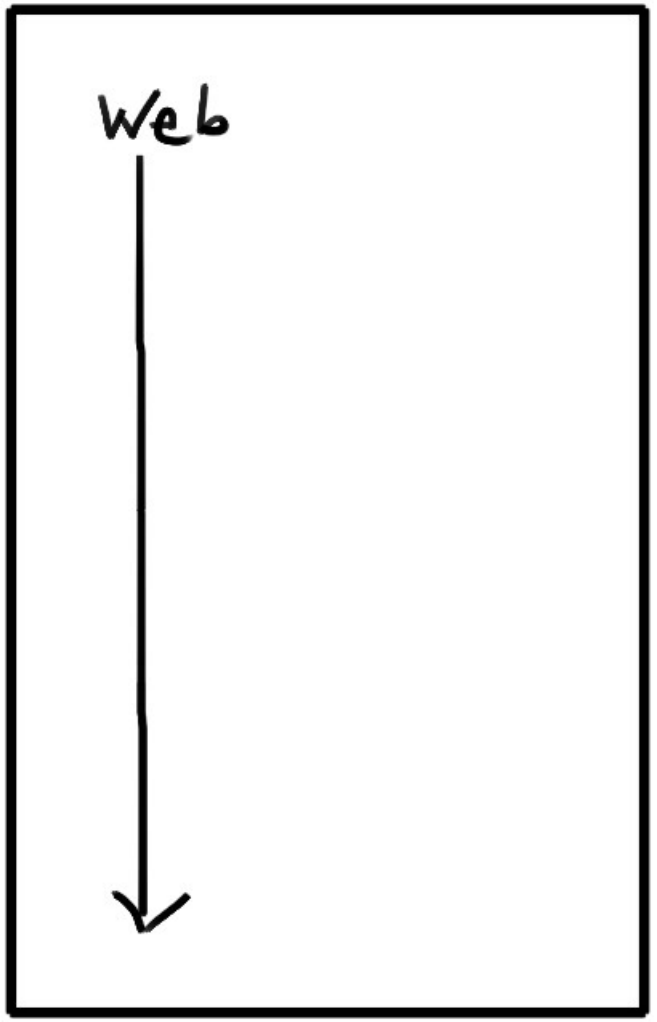

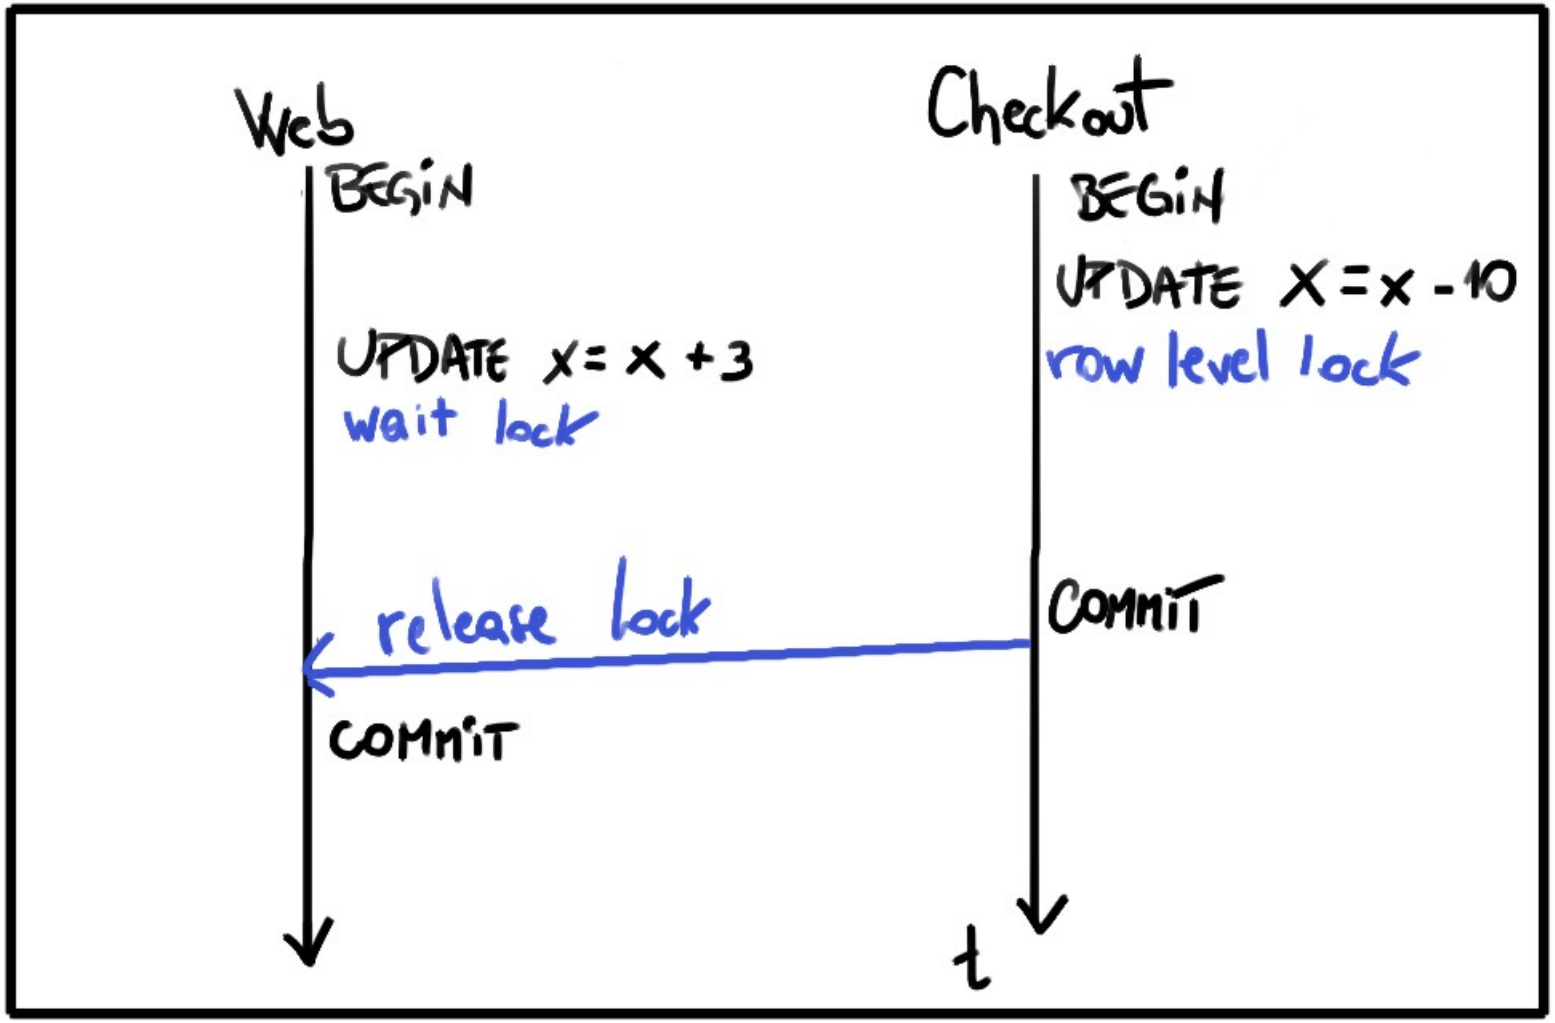

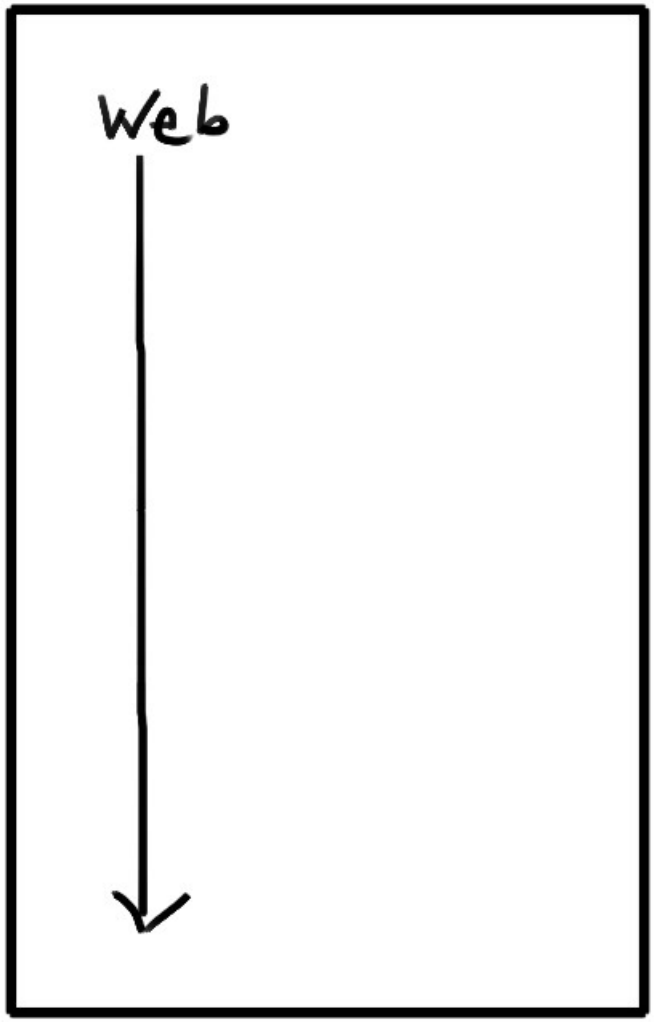

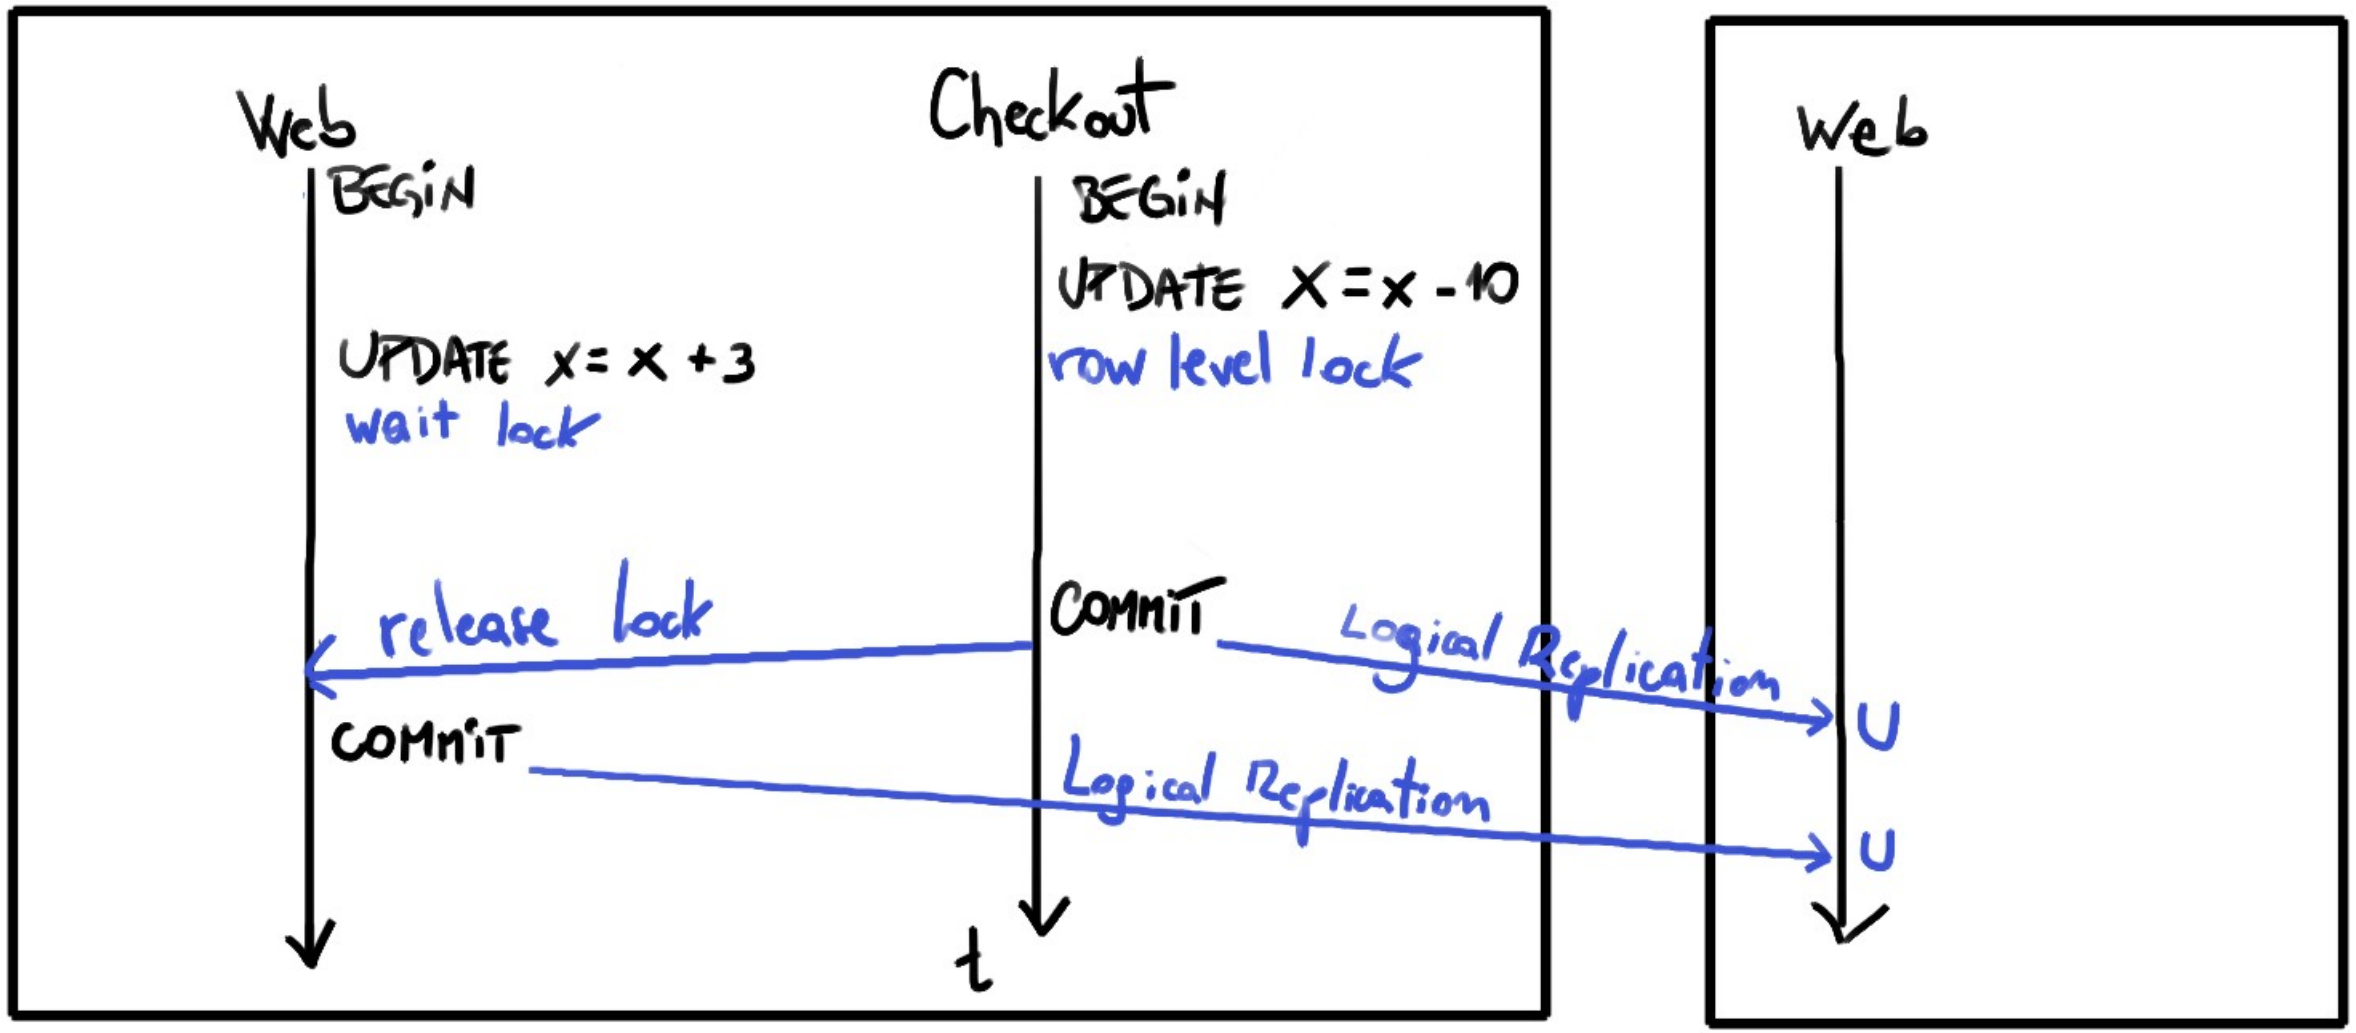

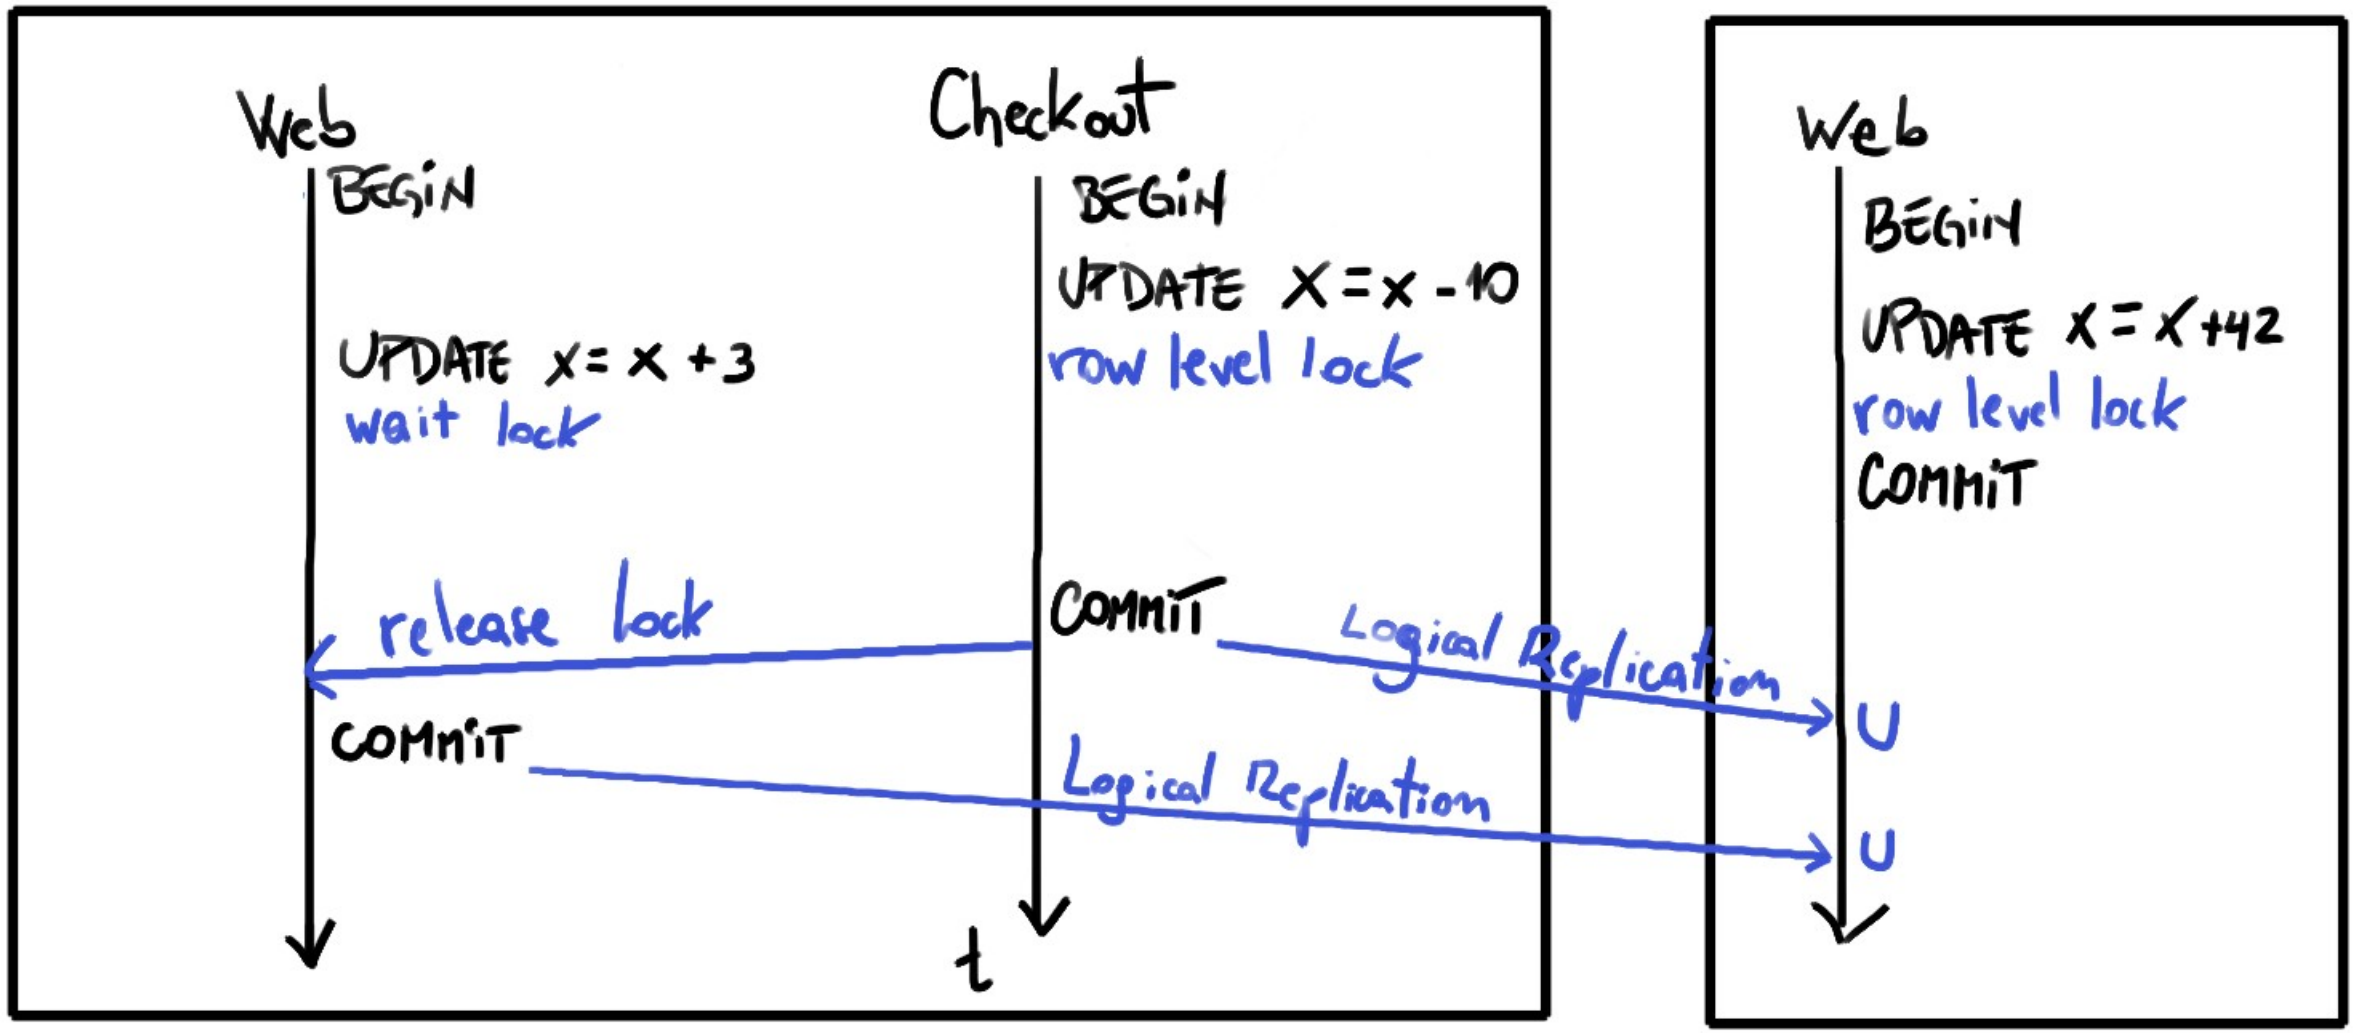

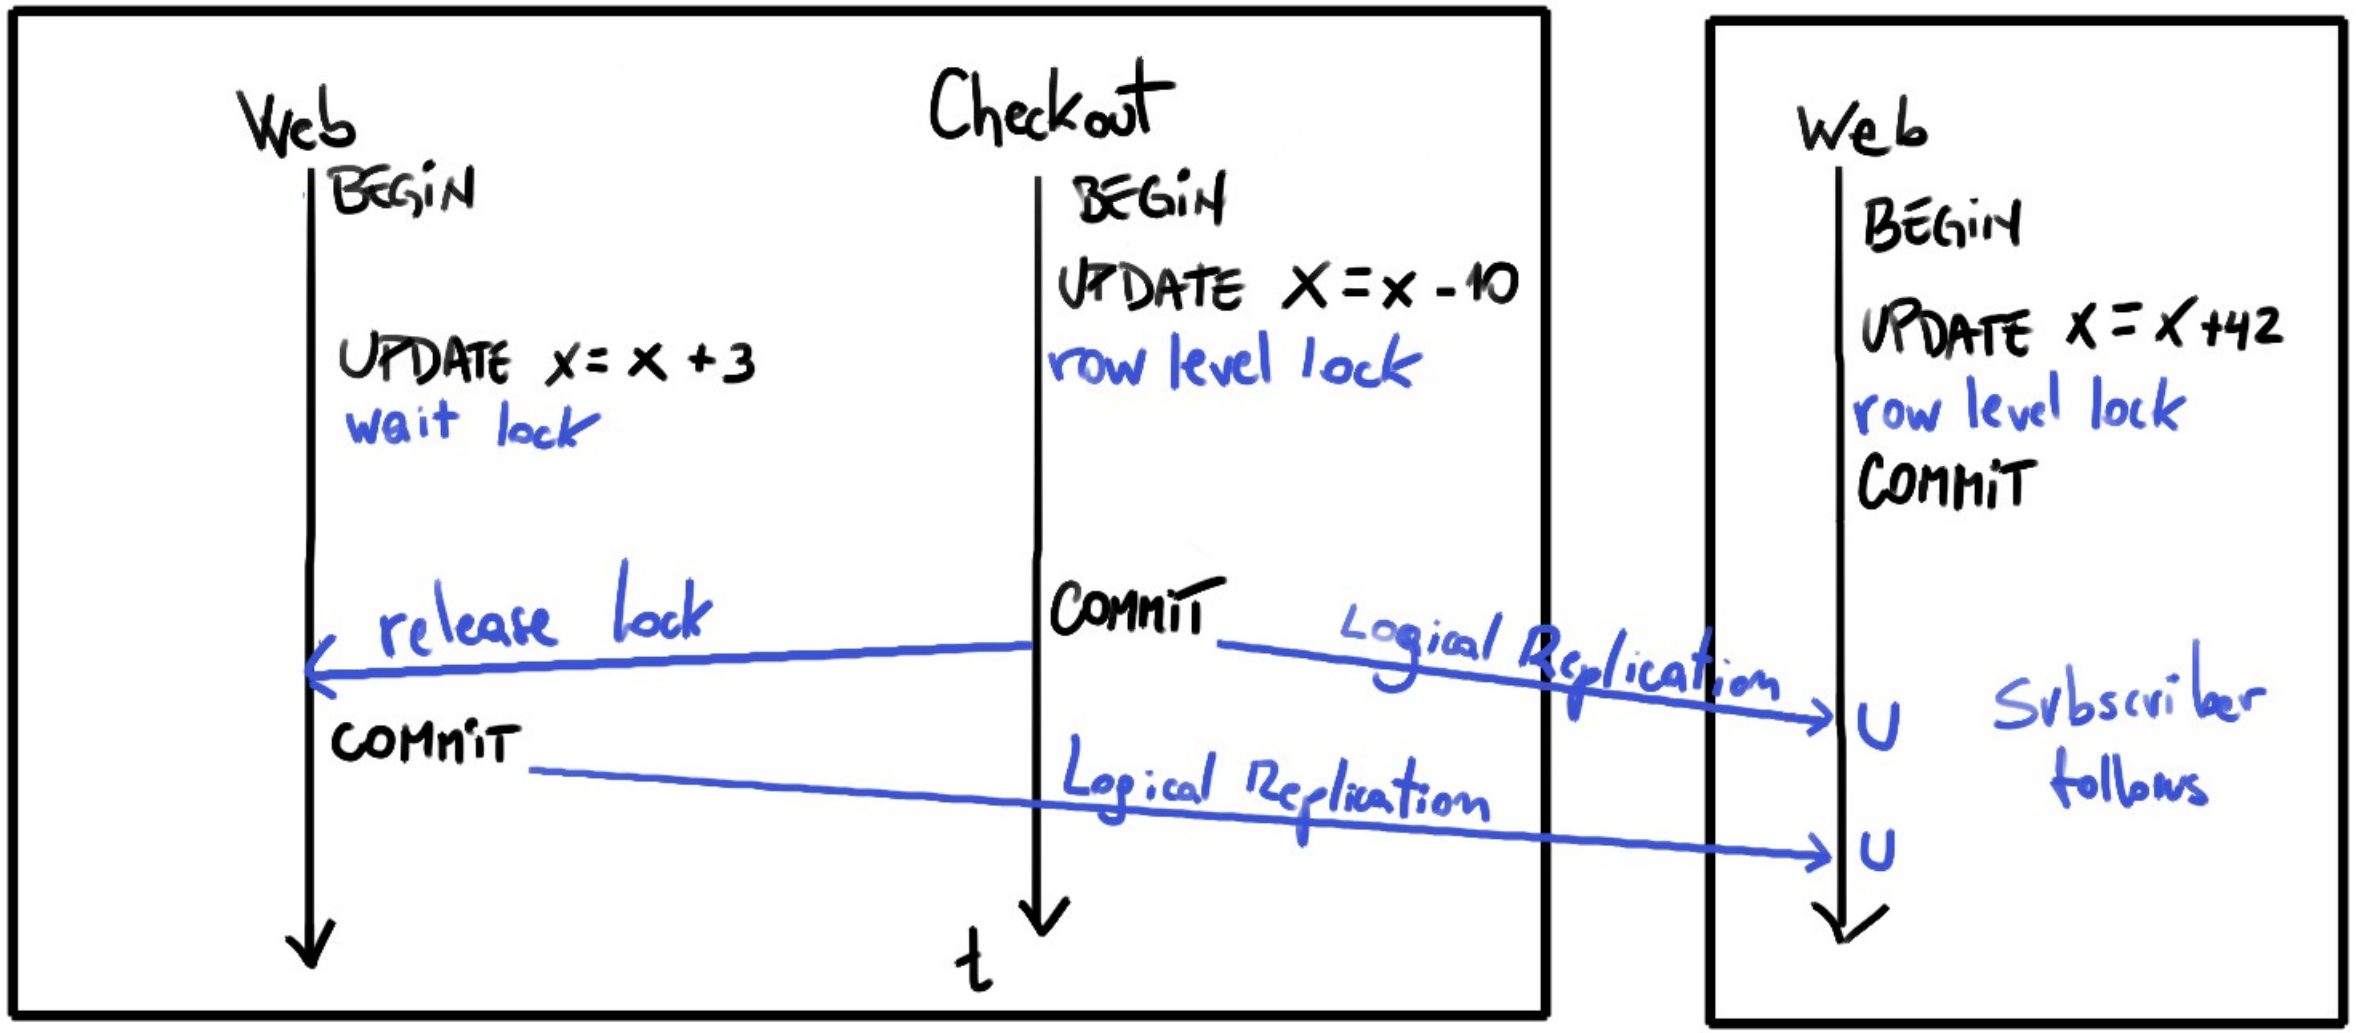

Logical Replication

Two Primary nodes

No remote locking

Logical Replication

Two Primary nodes

No remote locking

Subscriber follows

Logical Replication

Two Primary nodes

No remote locking

Subscriber follows

but not blindly

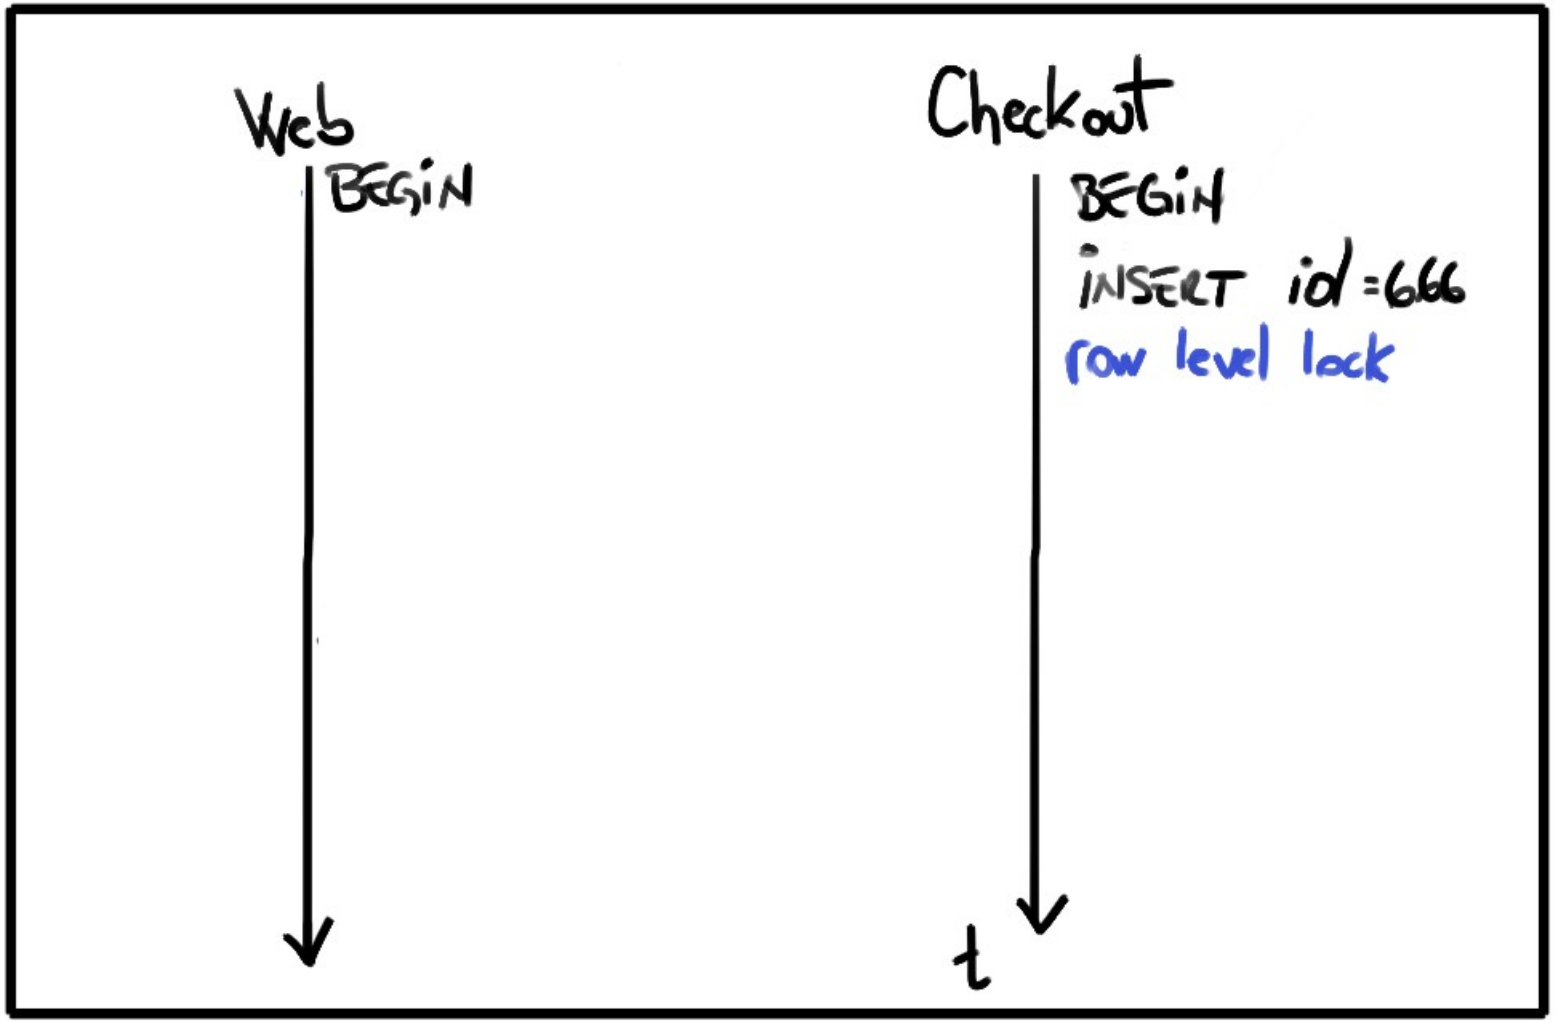

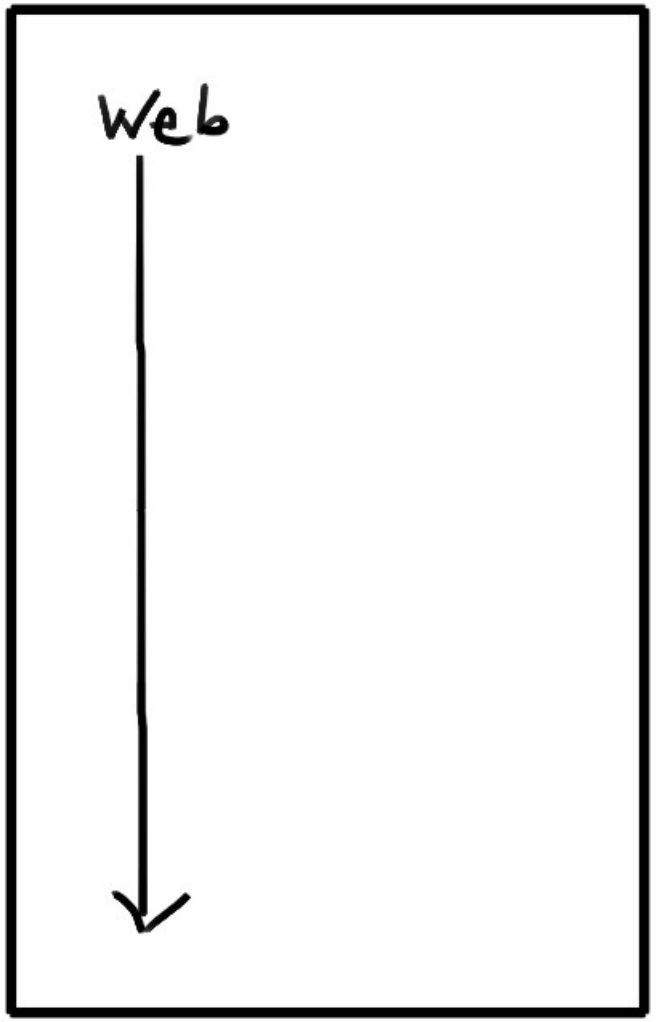

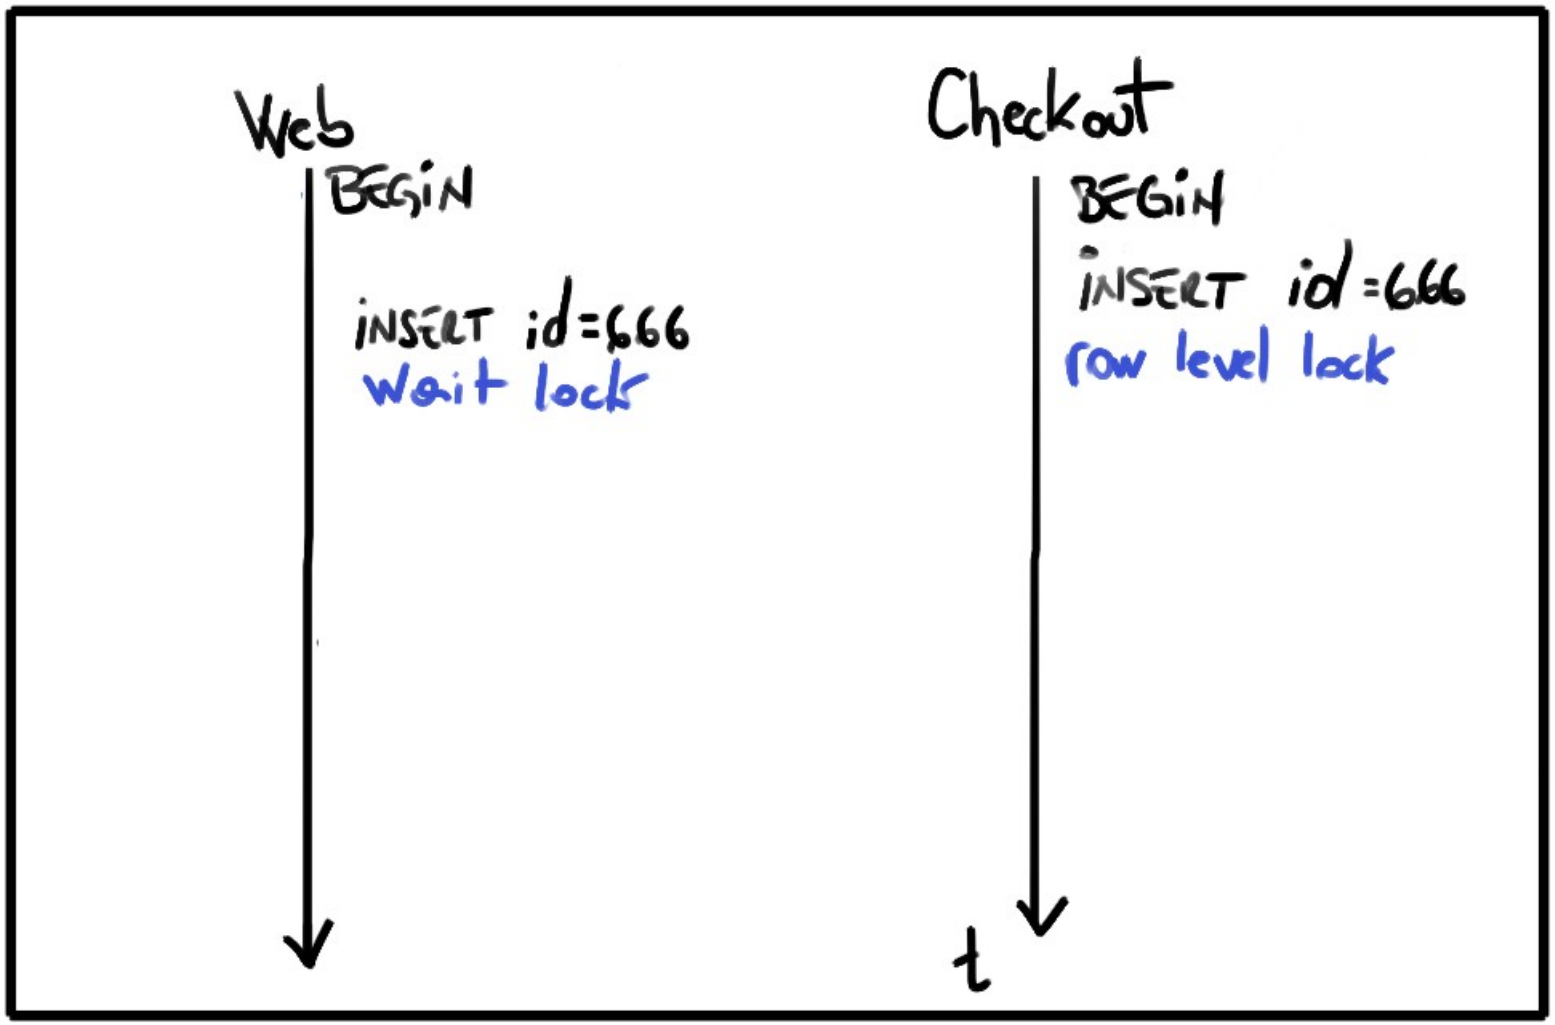

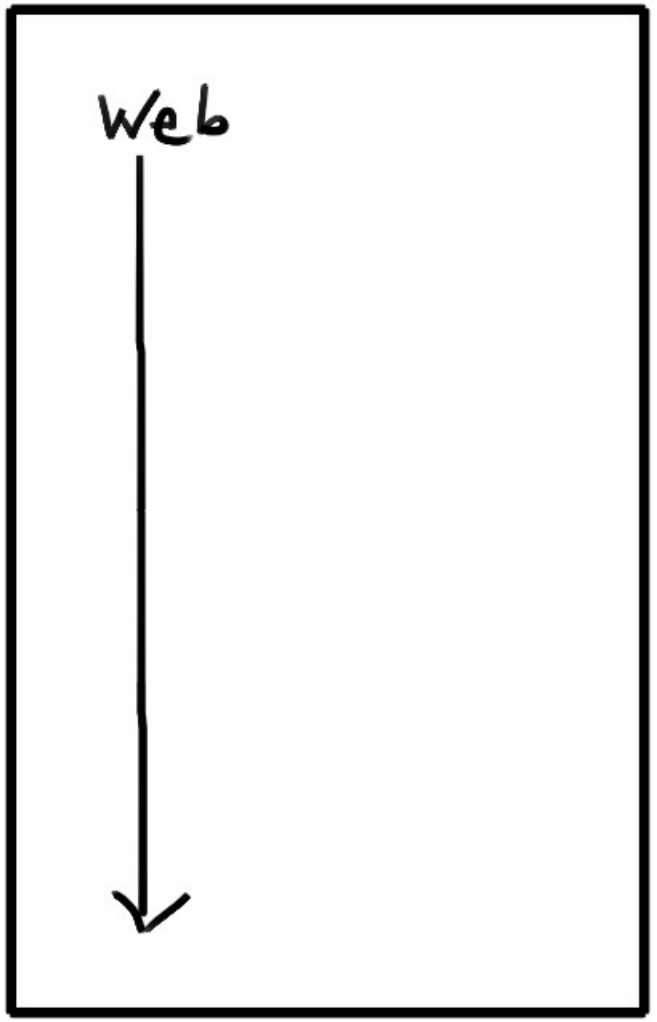

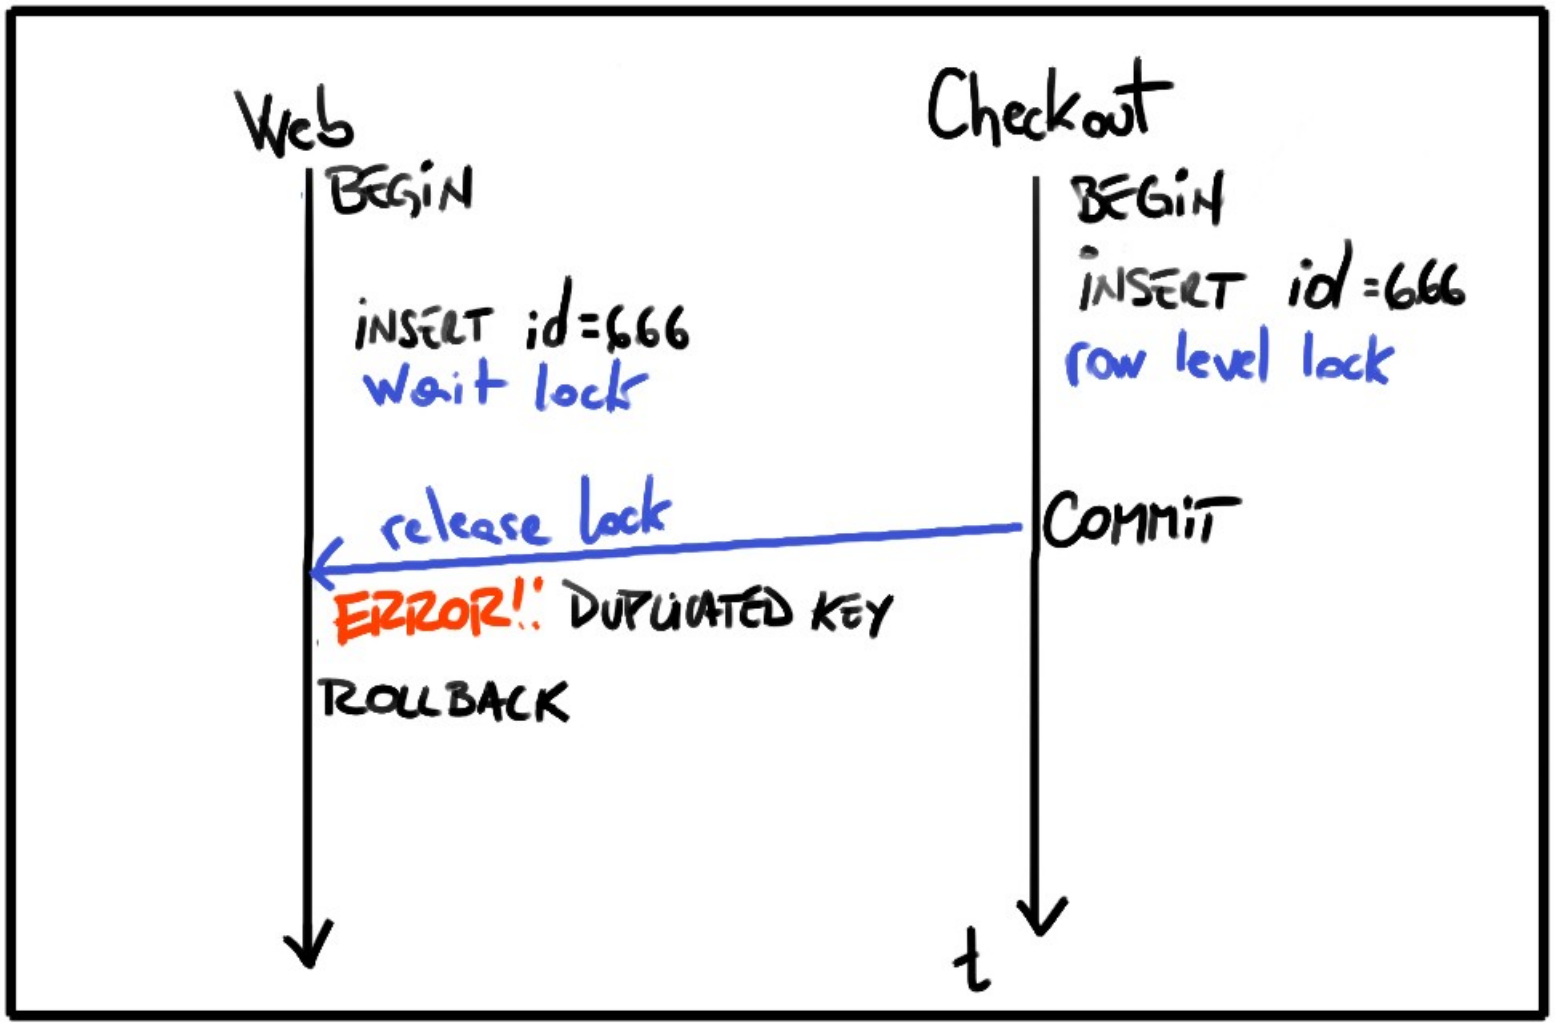

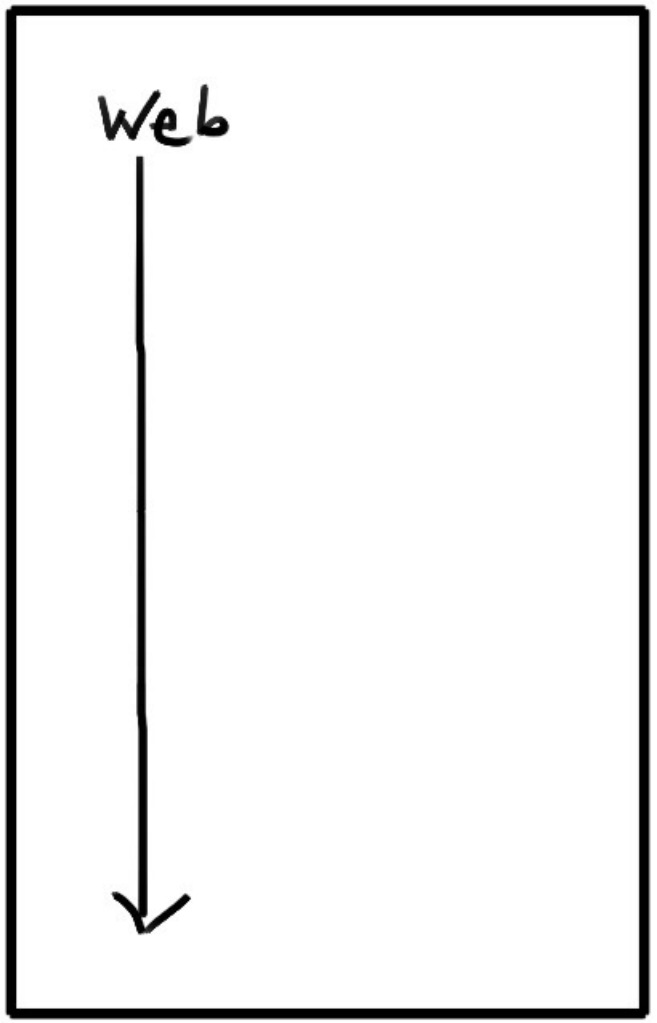

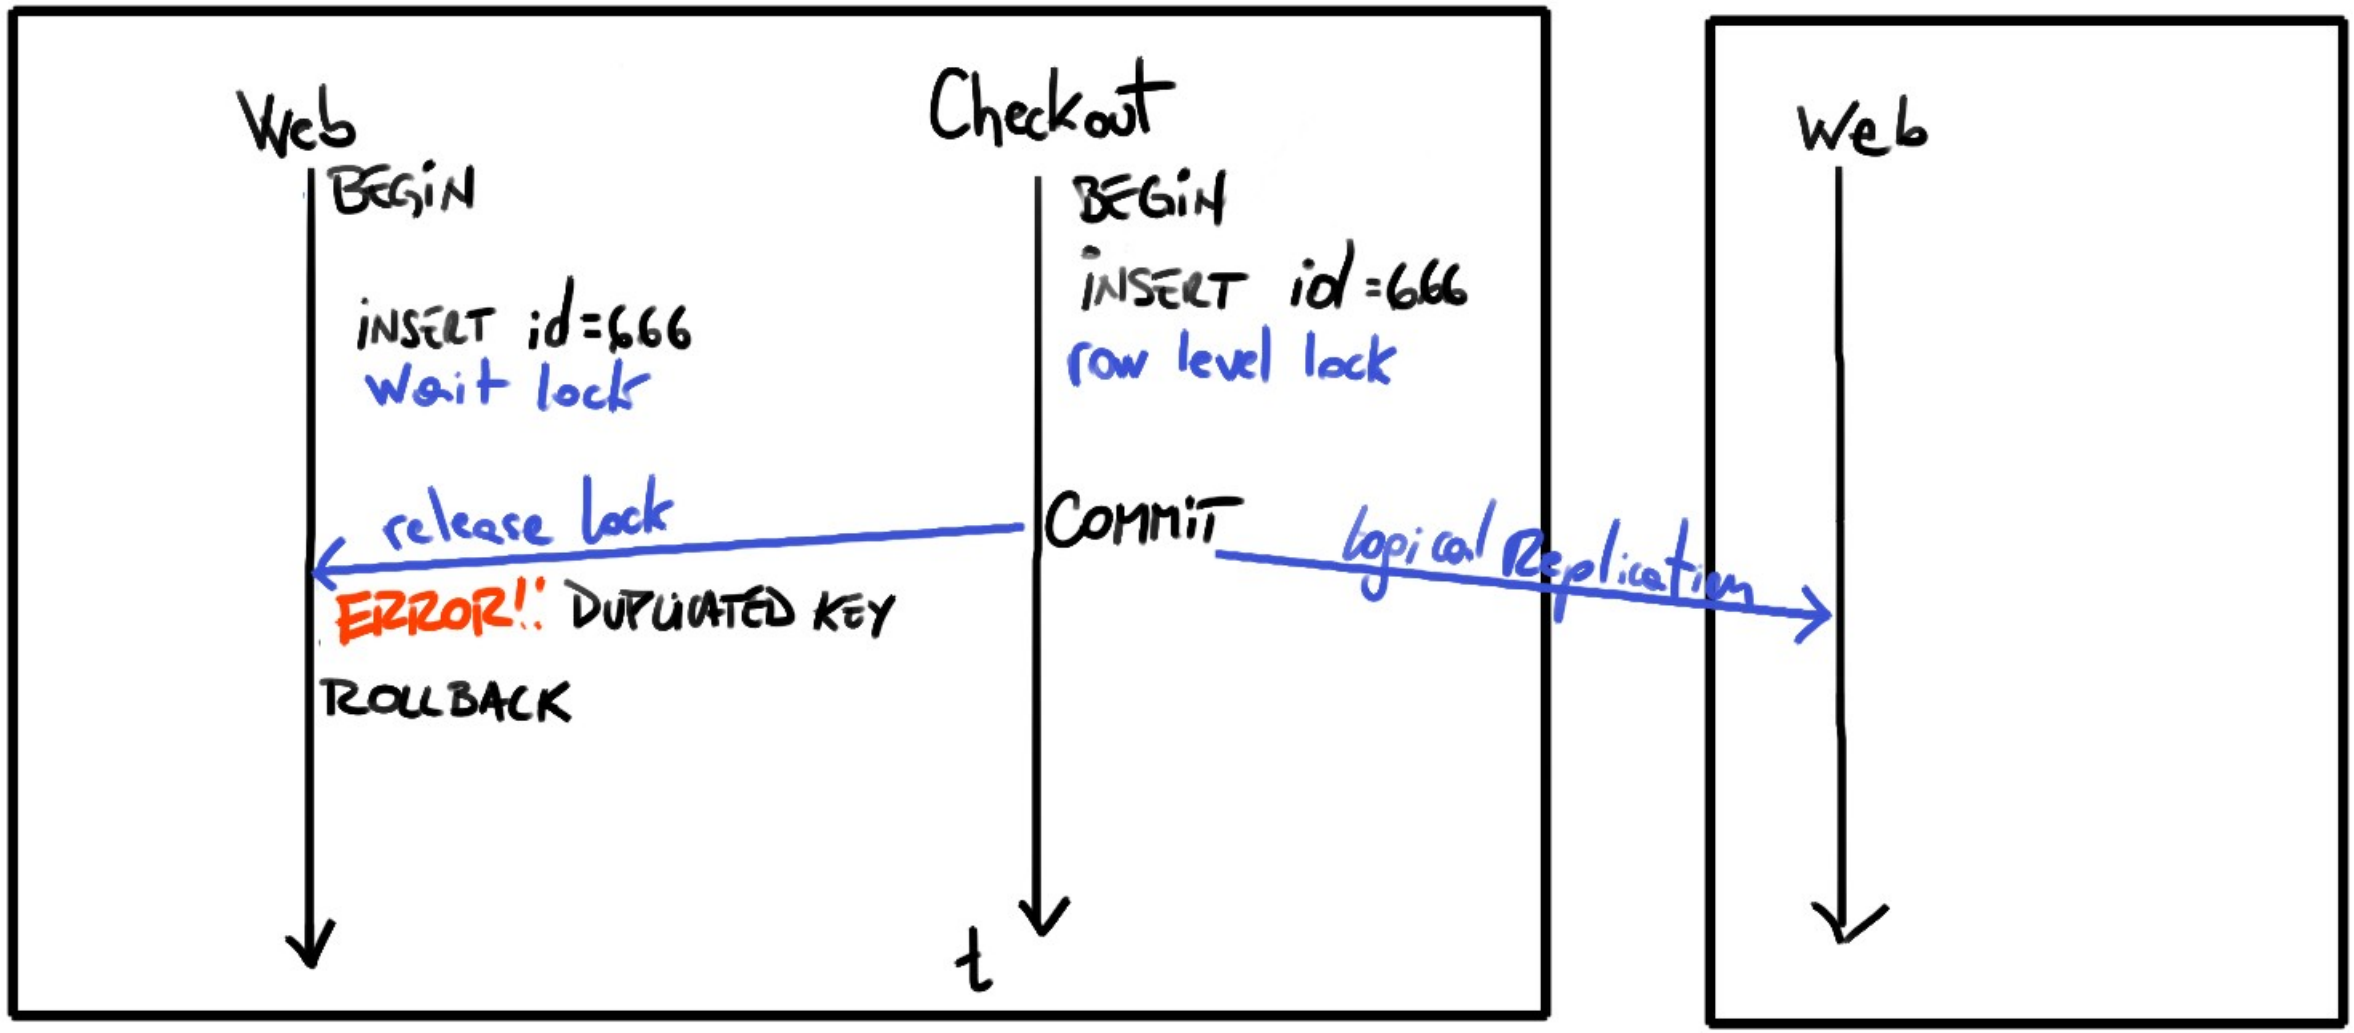

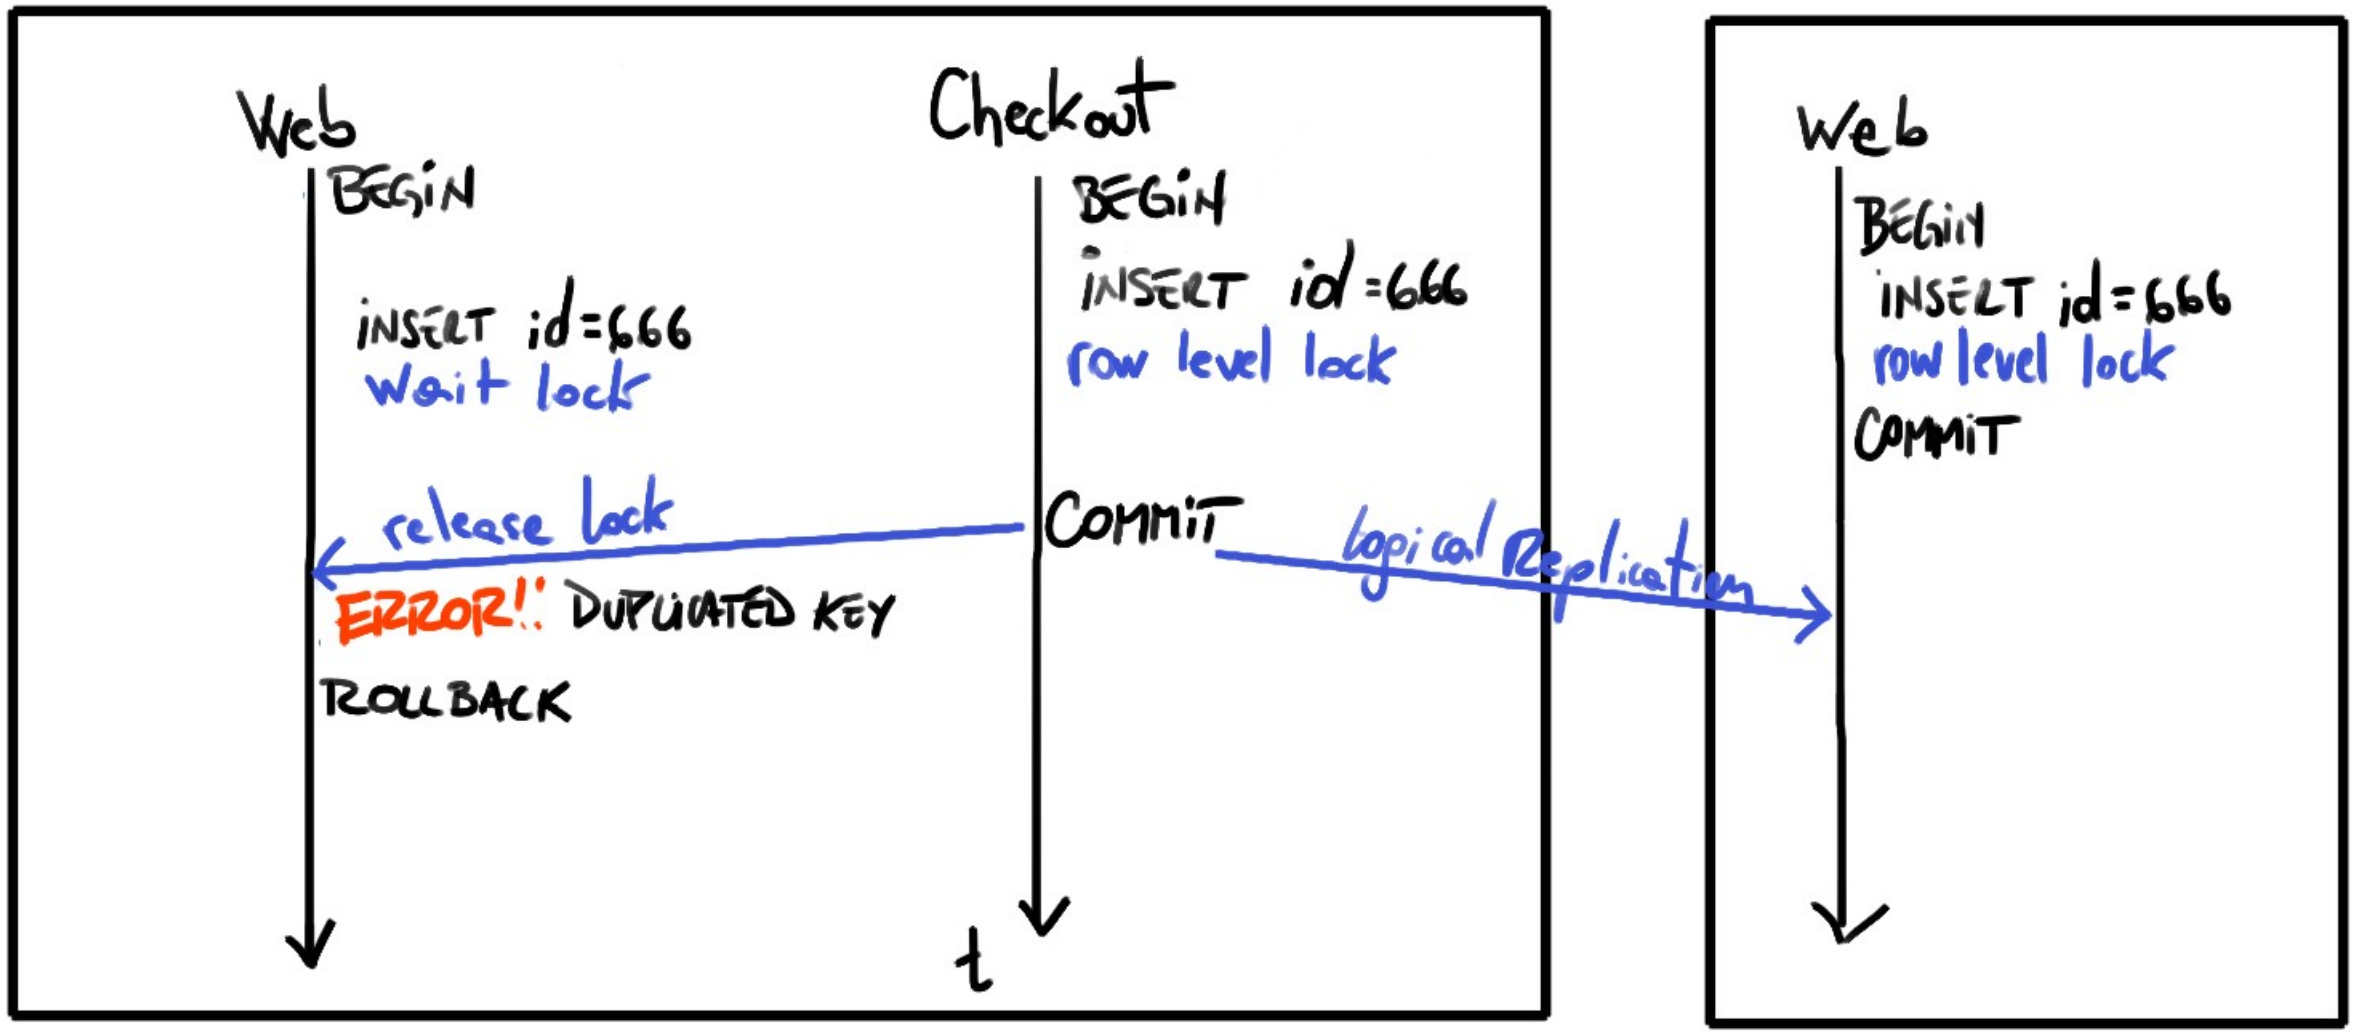

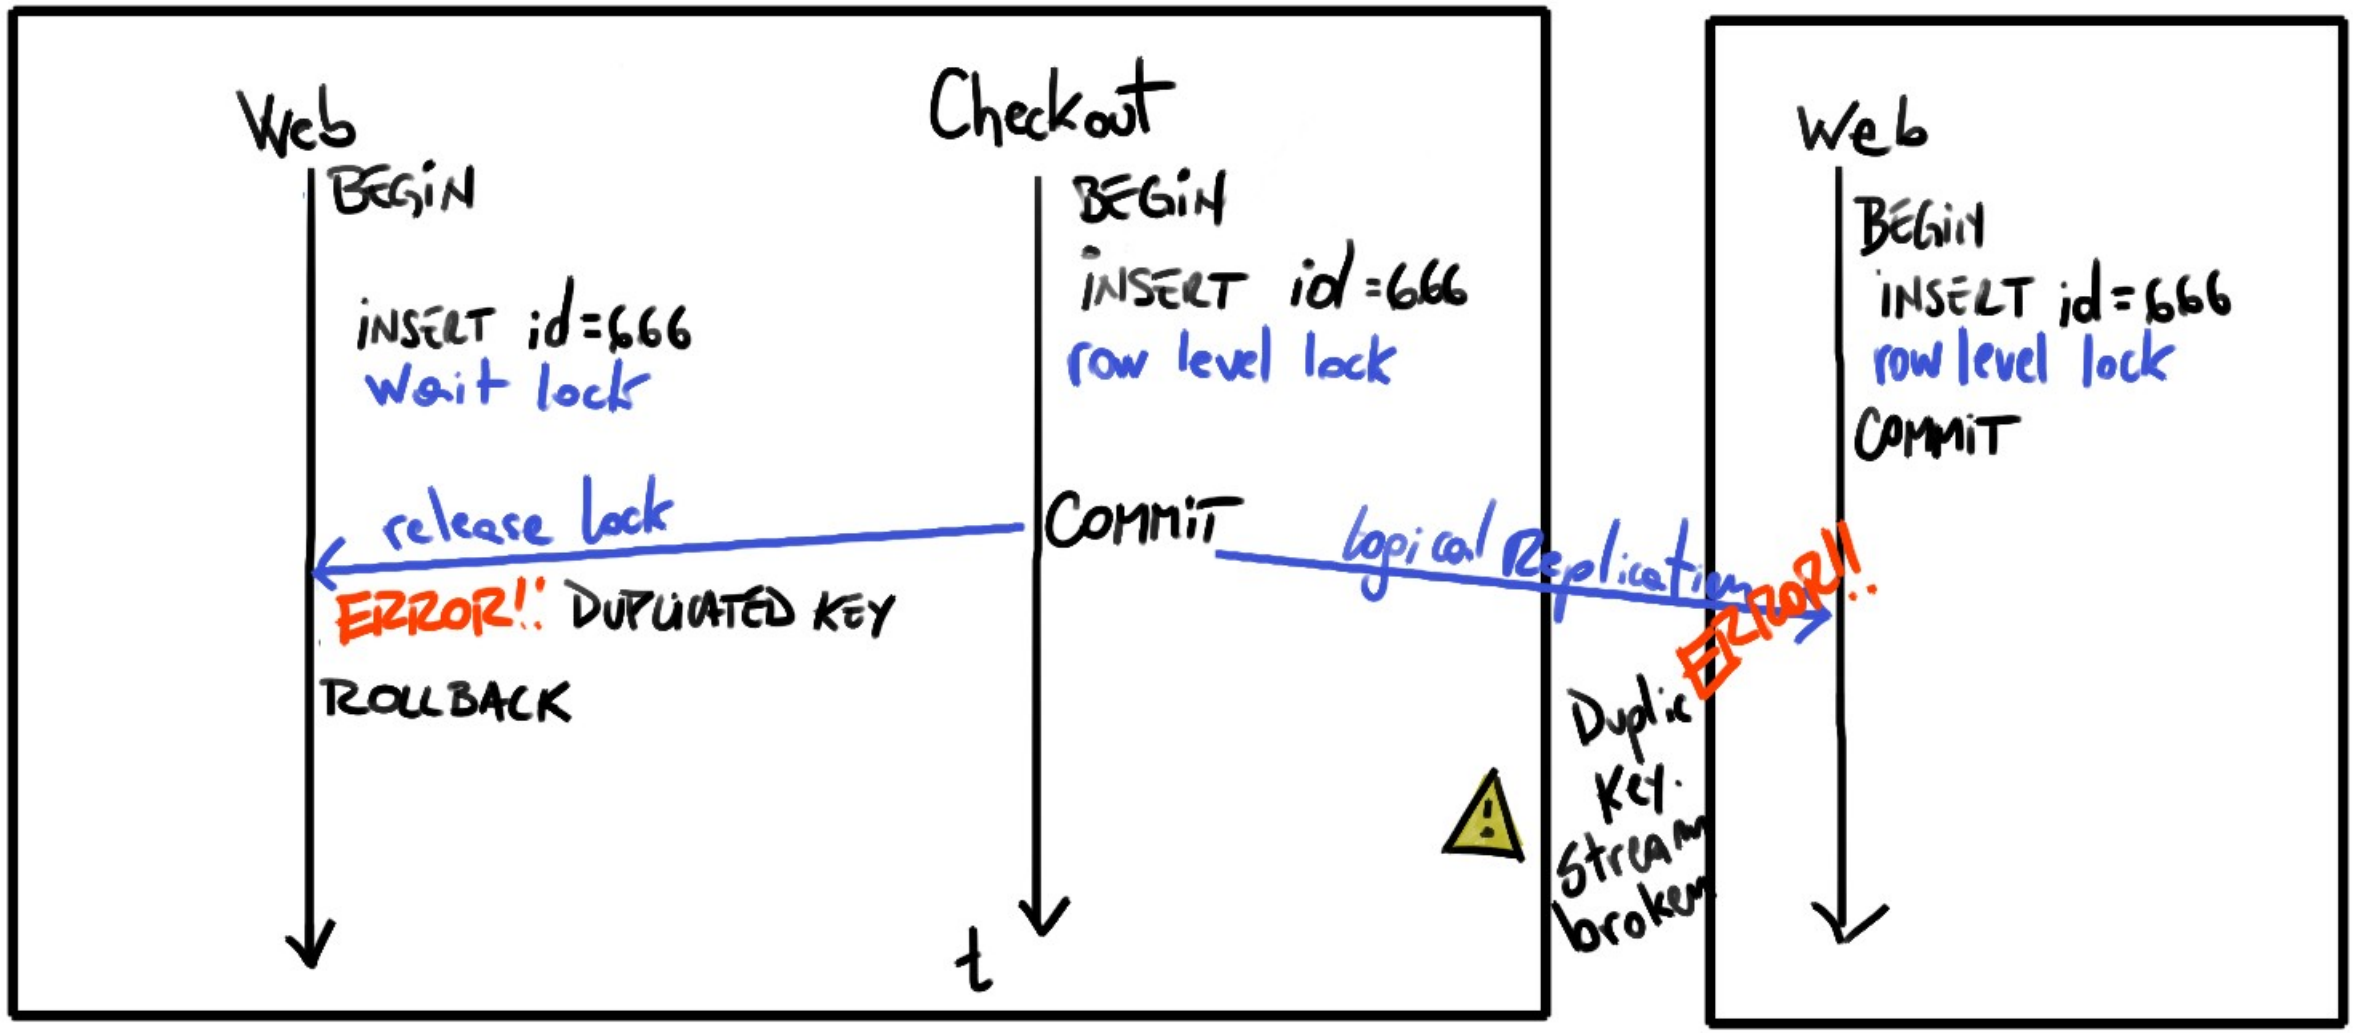

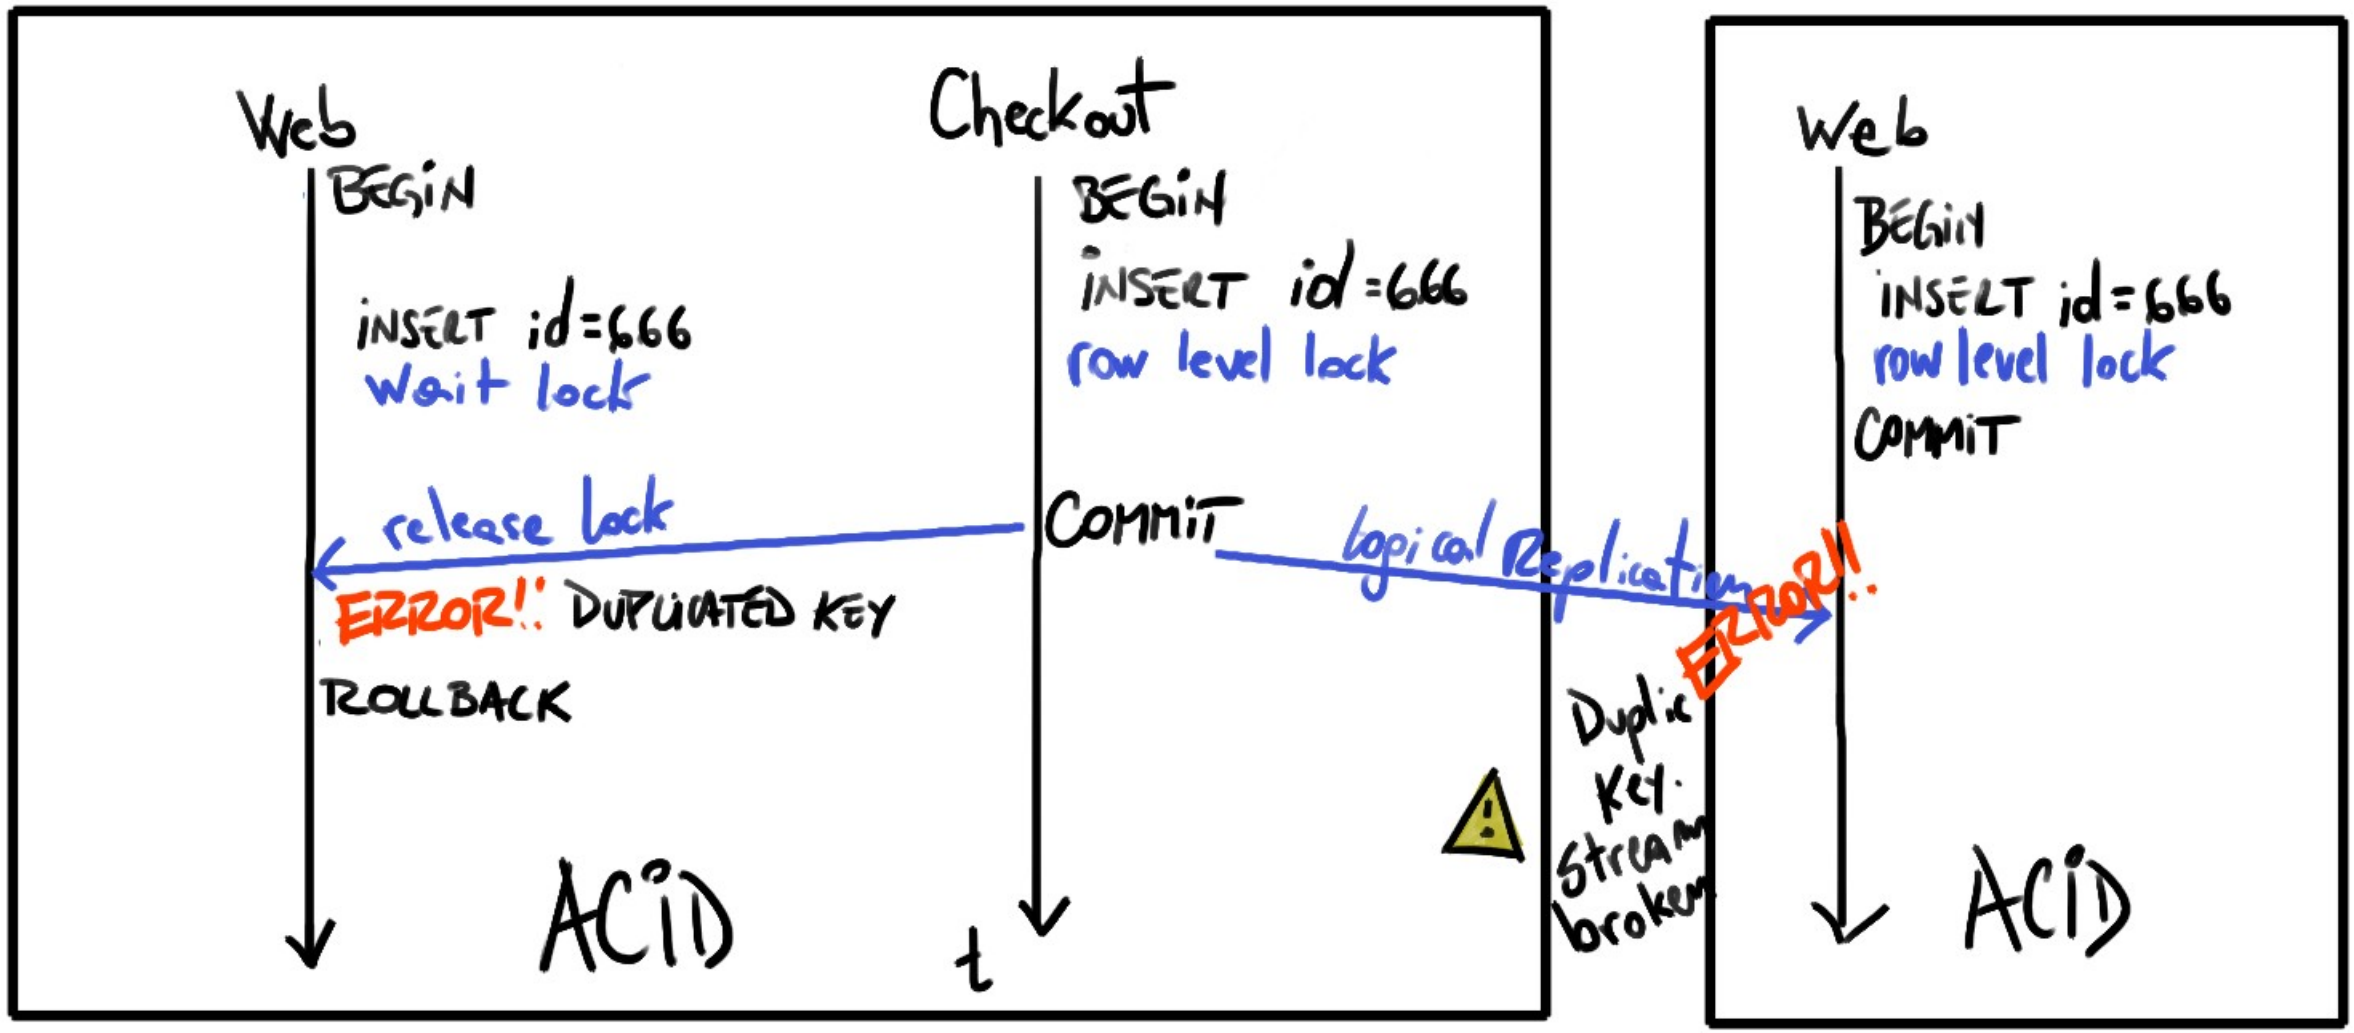

Closing Words

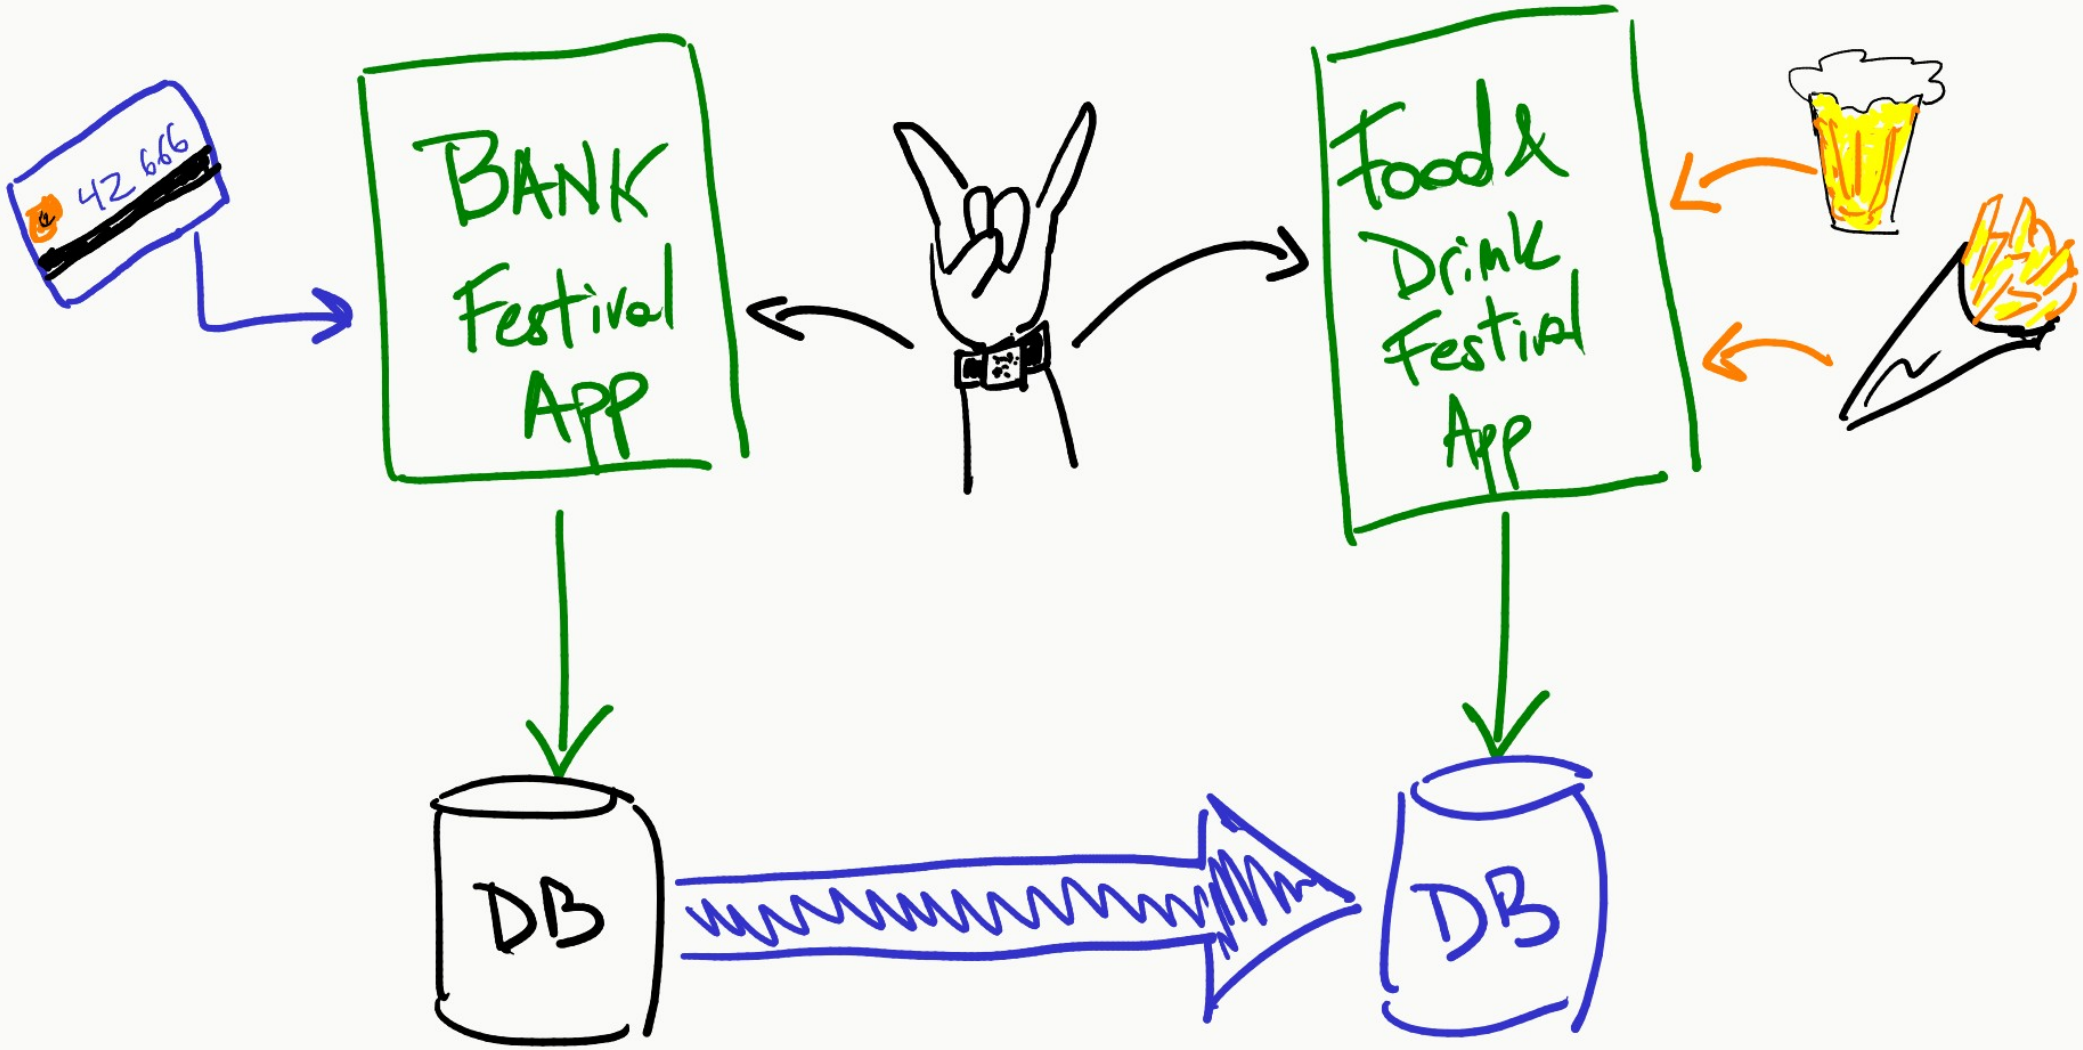
42 666
BANK
Festival
App
Food &
Drink
Festival
App
DB
DB

Consistency

Every read receives the most recent write or an error

Consistenct in PostgreSQL Replication

Consistenct in PostgreSQL Replication

- ACID properties on each node

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Eventually consistency globally

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Very good eventually consistency globally

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Very good eventually consistency globally
  - Monitor replication lag

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Very good eventually consistency globally
  - Monitor replication lag
- Synchronous replication reduces data loss

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Very good eventually consistency globally
  - Monitor replication lag
- Synchronous replication reduces data loss
- Logical Replication will not compromise ACID

Consistenct in PostgreSQL Replication

- ACID properties on each node
- Very good eventually consistency globally
  - Monitor replication lag
- Synchronous replication reduces data loss
- Logical Replication will not compromise ACID
  - No remote locking

Test your assumptions

Thank you