Eventual Consistency – Data Model Implications

Open Tech School

Kore Nordmann / @koredn / <kore@qafoo.com>

June, 11th 2015

Qafoo

passion for software quality

Open TechSchool

The Usecase

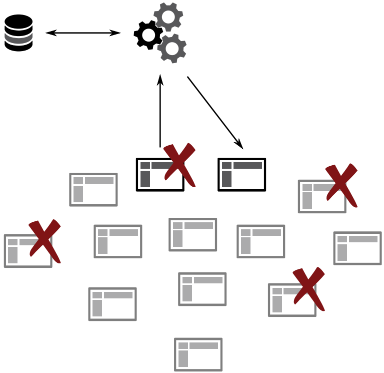

Solutions

- Transactions
  - Write is only ACK'd if all nodes ACK'd
    - Not possible if nodes do not ACK properly (Solr, MongoDB, ElasticSearch, . . .)
    - Two / three phase commits take time. . .
  - Rollback and deny writes entirely if one node does not ACK
    - Omitted rollback requires full-sync
    - Requires re-transmitting all data
    - Checking which IDs are transmitted requires iterating all IDs
- Eventual Consistency

CAP Theorem

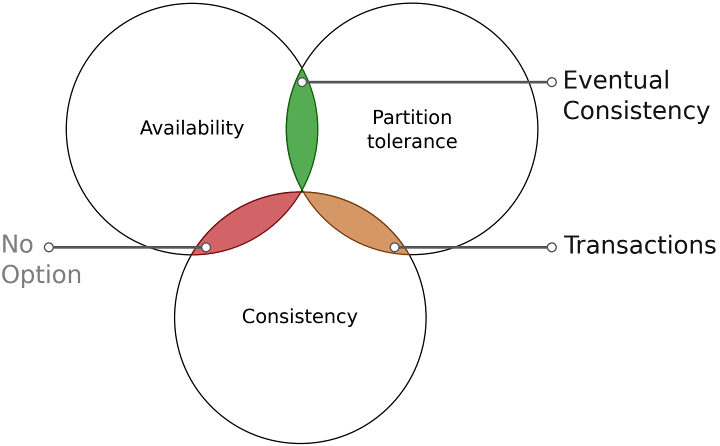
Availability
Partition tolerance
Consistency
Eventual Consistency
No Option
Transactions

Eventual Consistency

Sounds good – but how?

Eventual Consistency

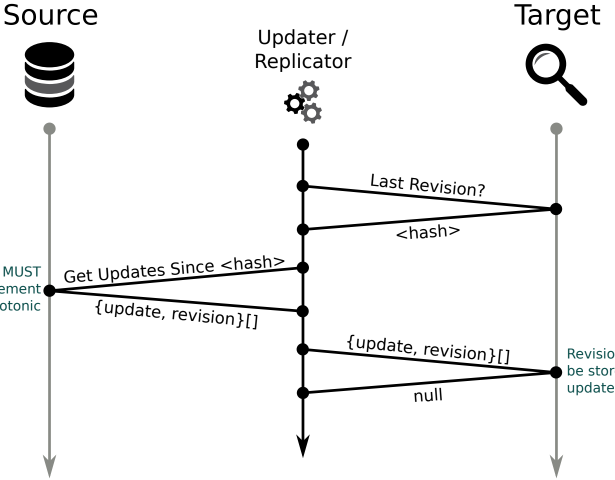

Revisions MUST
increment
strictly monotonic

Revisions MUST NOT
be stored if an
update fails.

Data Modelling on Source

- Store denormalized “updates”
  - Revision (globally strictly monotonic)
  - Store denormalized data
  - Keep deletes
- Maintaining referential integrity is hard – but not impossible

CouchDB – Multi Master

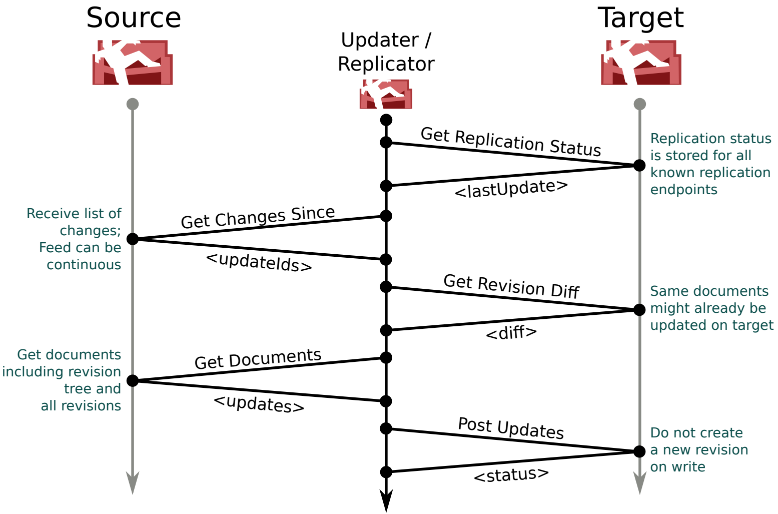

Replication status is stored for all known replication endpoints

Receive list of changes; Feed can be continuous

Same documents might already be updated on target

Get documents including revision tree and all revisions

Do not create a new revision on write

CouchDB – Merging

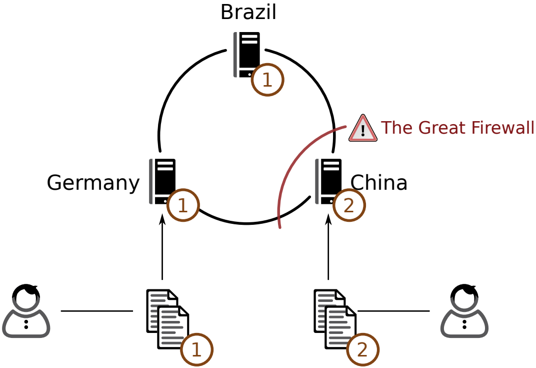

Relation Modelling

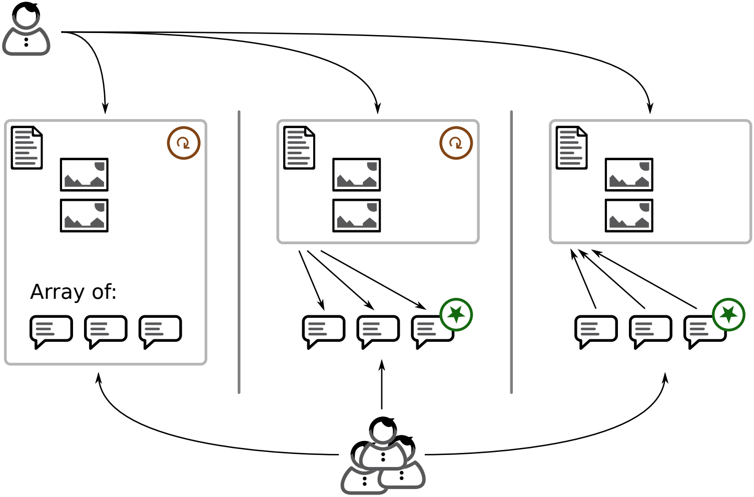

Simpel JOIN Query

```
function (doc) {
    if (doc.type === "blog_post") {
        emit([doc._id], null);
    }

    if (doc.type === "blog_comment") {
        emit([doc.blog_post_id, doc.date],
             null);
    }
}
```

```
?starkey=["post"]&endkey=["post", {}]

["post"]                          => null
["post", "2015-06-11␣13:23"] => null
["post", "2015-06-11␣13:37"] => null
["post", "2015-06-11␣13:42"] => null
```

passion for software quality

Implications

- Split data into small documents
  - … depending on how often data changes
- Reference parent, not children
  - A blog post referencing its comments, vs. …
  - comments referencing their blog post.

Summary

- Embrace Eventual Consistency
- Play with `http://hood.ie/`

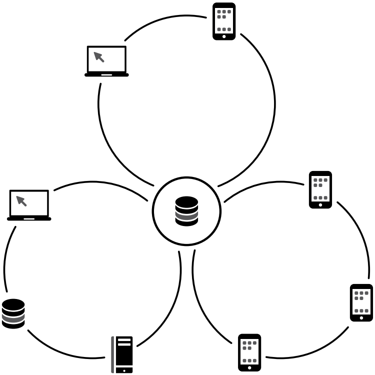

Qafoo
passion for software quality

Qafoo
passion for software quality

THANK YOU

Rent a quality expert
qafoo.com