FoundationDB

A Distributed Unbundled Transactional Key Value Store

Yuhan Wang

Agenda

- Background
- Design
- Contributions
- Limitations

Background: why a transactional key value store?

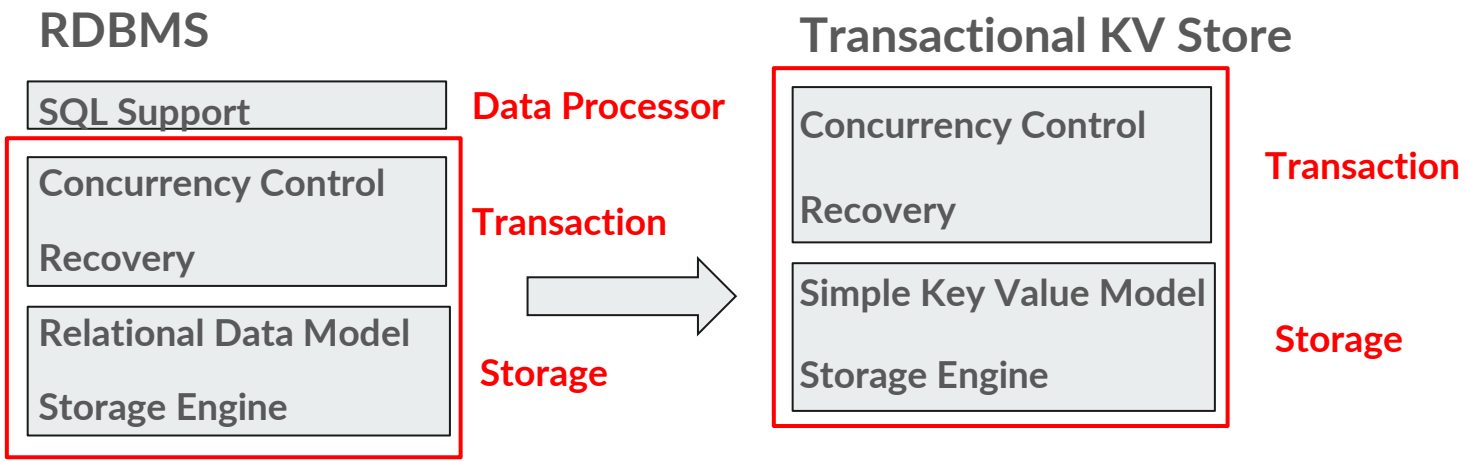

Design: Architecture

- Coordinators: Paxos
- Transaction System: Stateless
- Log System: Stateful
- Storage System: Stateful

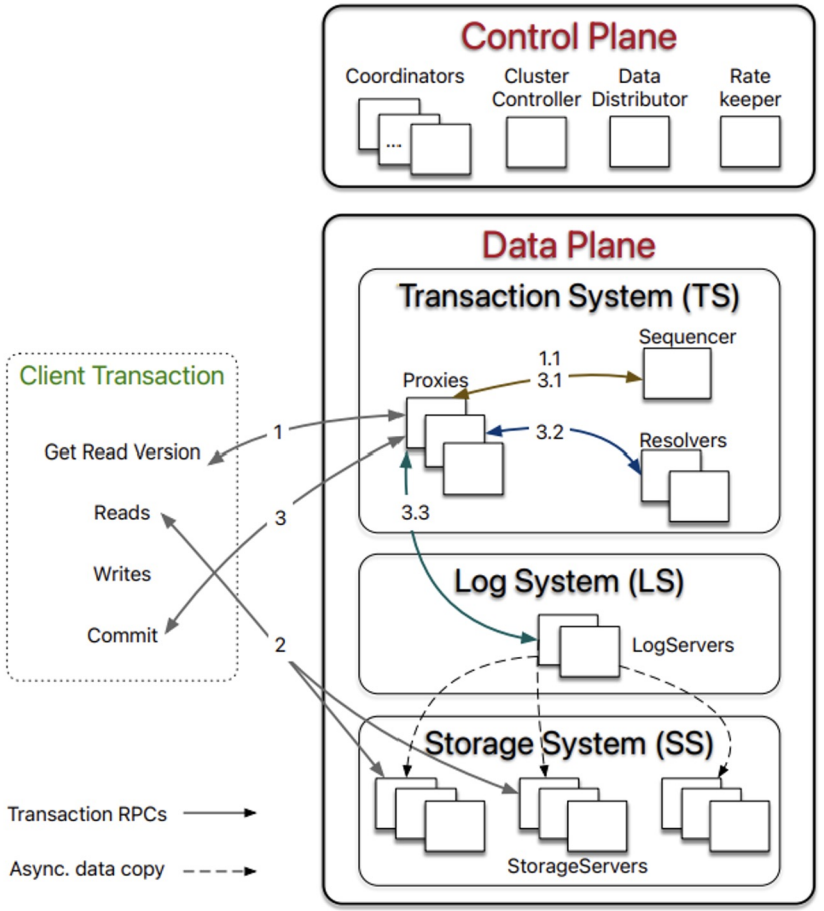

Design: Concurrency Control

- Strict Serializability
- Backward Validation OCC
- MVCC
- Commit Version：
  - prev version
  - current version

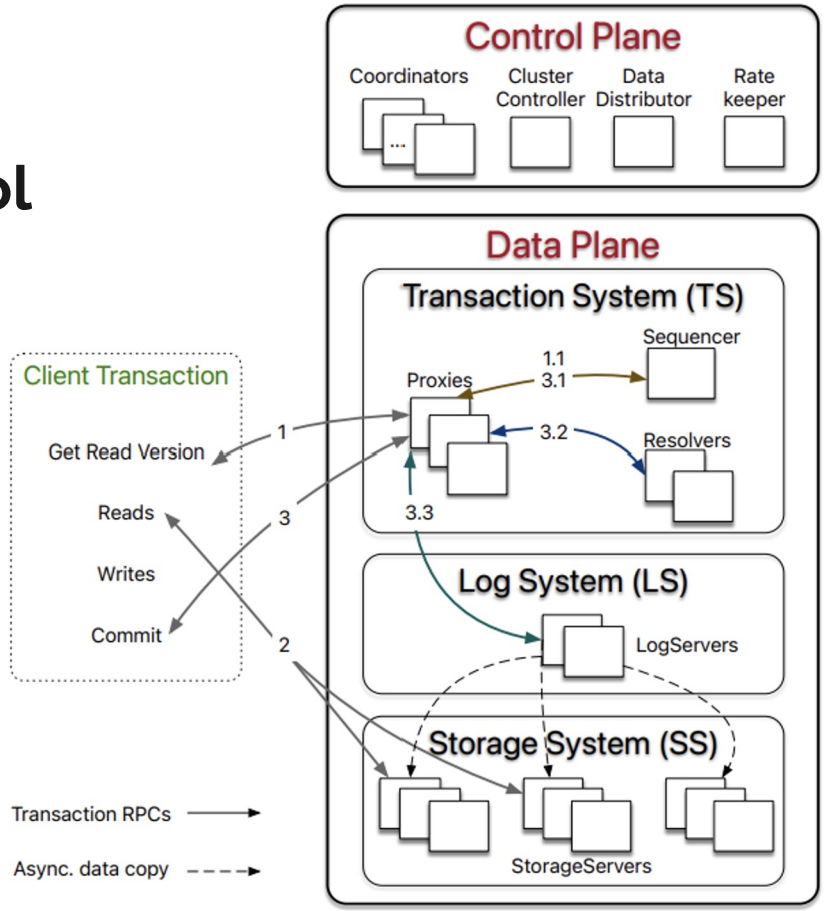

Design: Concurrency Control

Algorithm 1: Check conflicts for transaction T_x.

Require: *lastCommit*: a map of key range → last commit version

```
1 for each range ∈ R_r do
2     ranges = lastCommit.intersect(range)
3     for each r ∈ ranges do
4         if lastCommit[r] > T_x.readVersion then
5             return abort;
  // commit path
6 for each range ∈ R_w do
7     lastCommit[range] = T_x.commitVersion;
8 return commit;
```

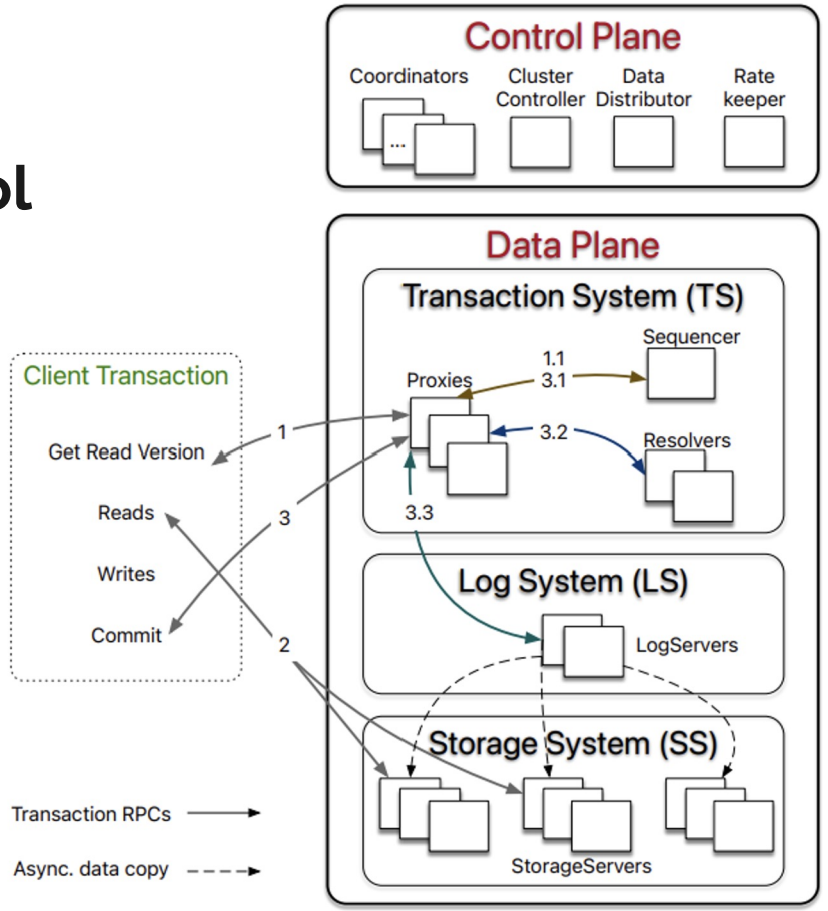

Design: Log System

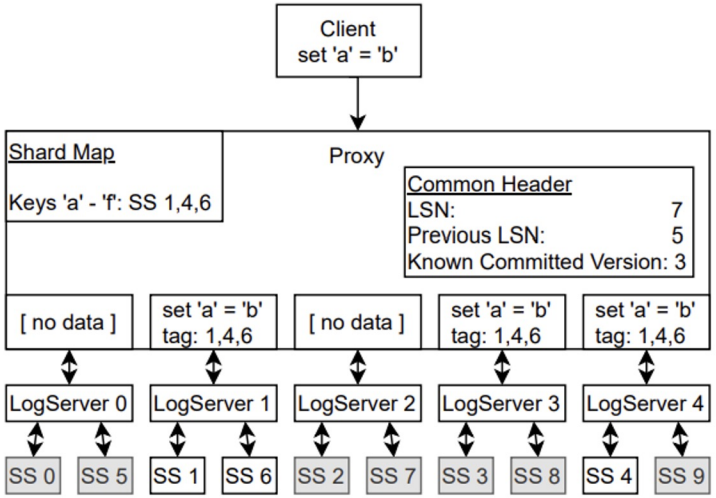

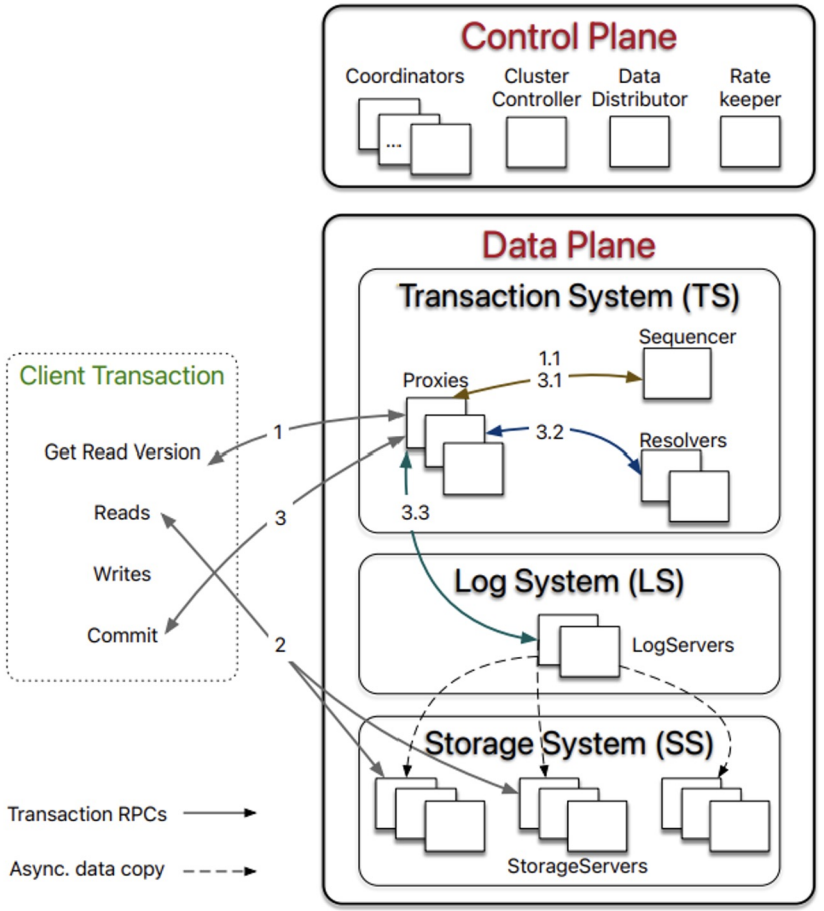

Design: Recovery

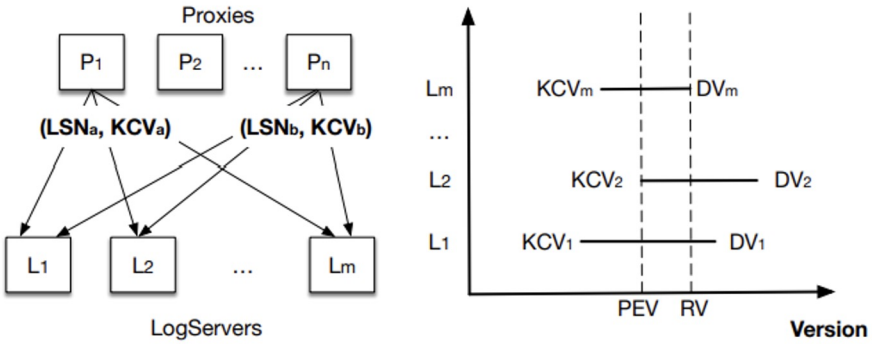

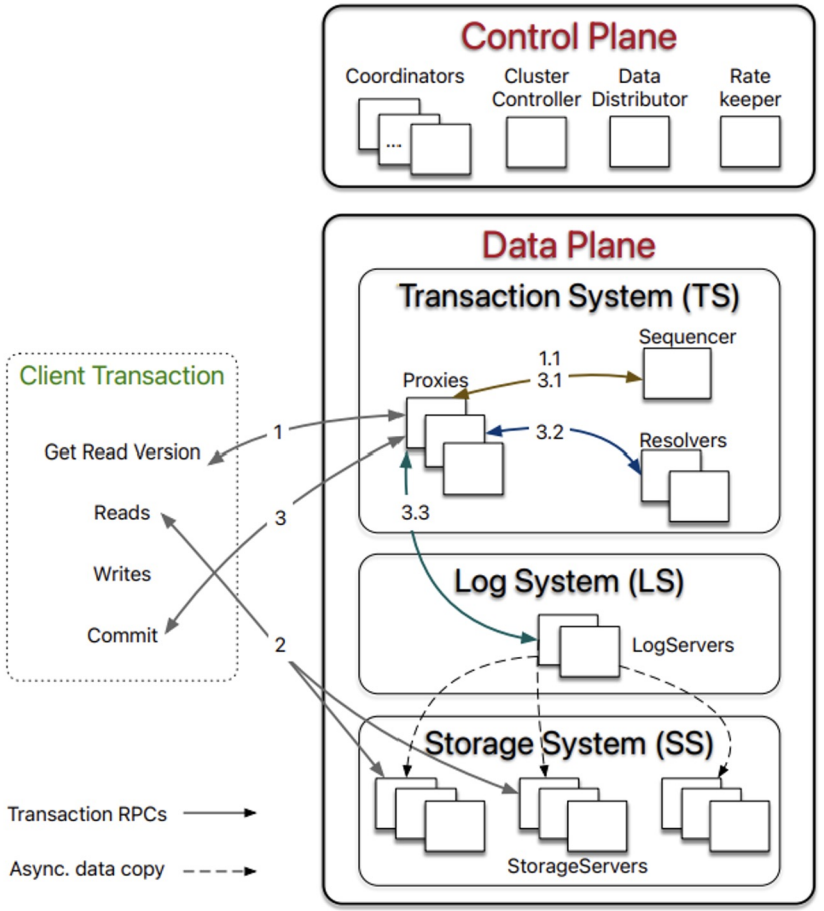

Design: Tricky Problem

Assign versions in Sequencer?

- Strict Serializability
- SS Aggressive Fetching

Solution:

- Increasing Commit Versions
- Lasted Commited Version (as Read Version)

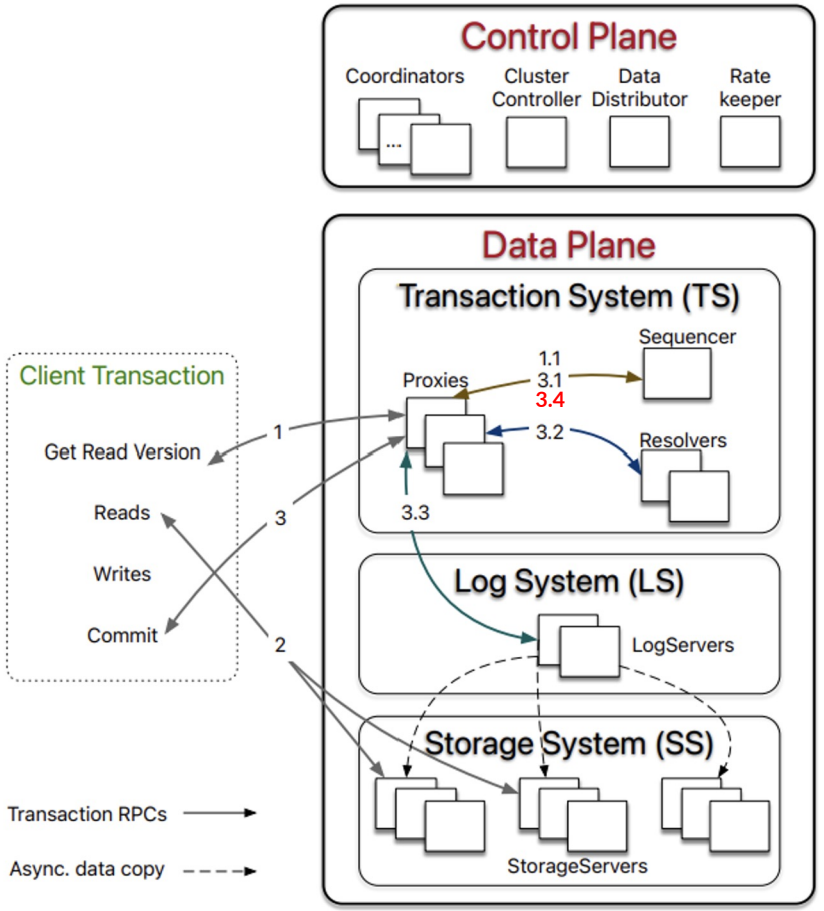

Contributions

- OCC vs Two Phase Locking
- Fast Recovery vs Always On
- Optimize For Happy Case

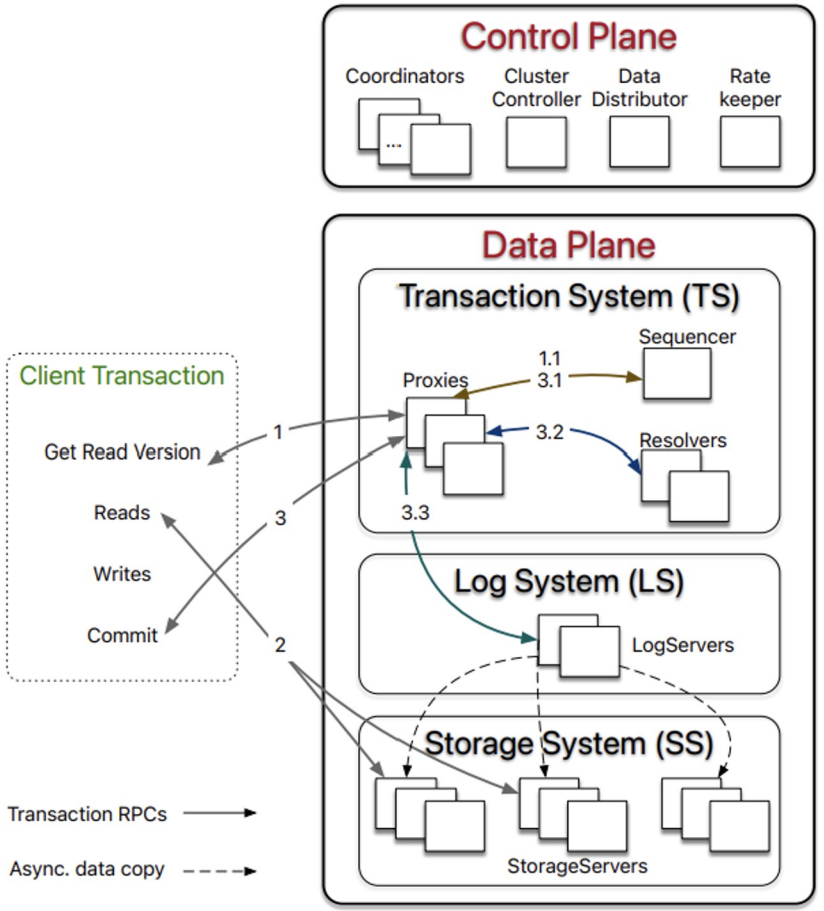

Limitations

- Transaction Size (OCC)
- 5s MVCC Window
- No RW Transactions in Recovery

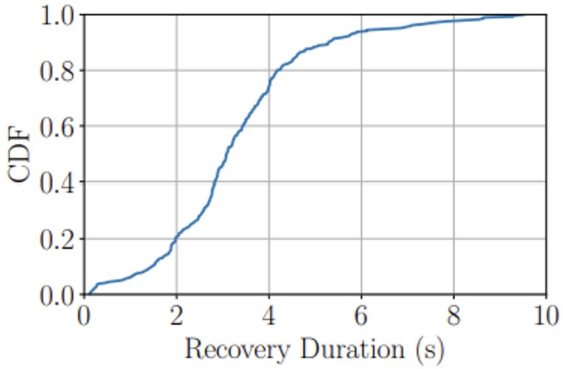

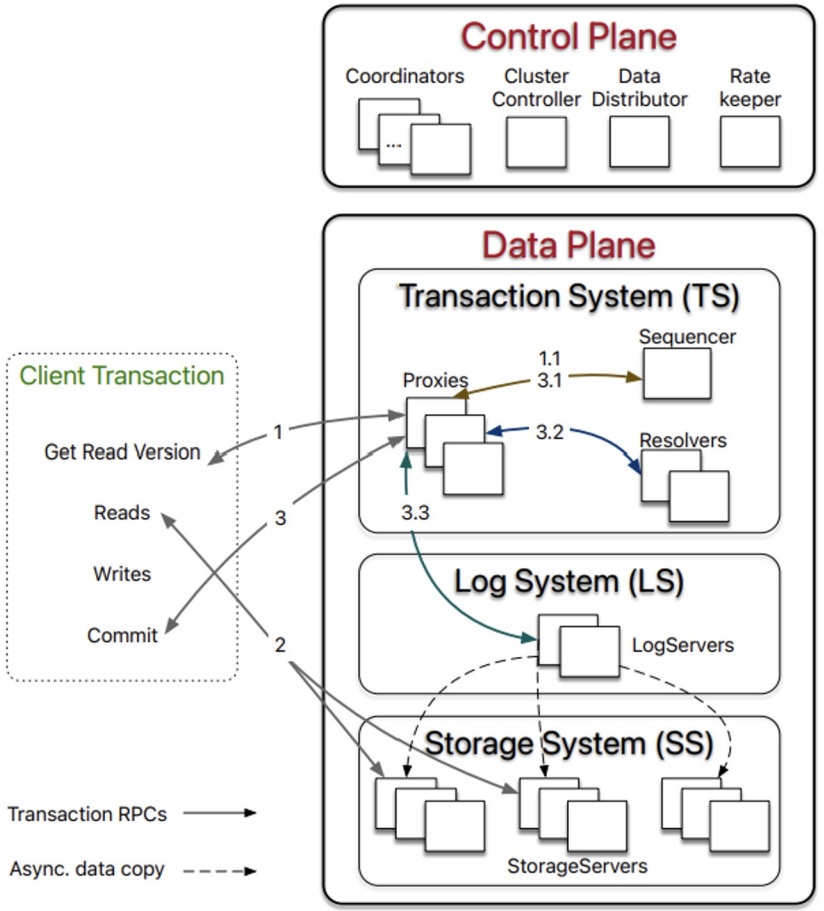

Thanks!

And Questions?