CS 686: Special Topics in Big Data

Spanner

Lecture 14

Tweets

- Spanner achieved external consistency by using TrueTime API and implementing a timestamp version of BigTable
- “Make SQL Database Great Again”
  – Long-lived Leader, 7ms ago
- Spanner: "If inconsistency is a crime, let's just use TrueTime!”
- Spanner: A semi-relational, global-distributed, and multi-version database. A good choice for those who miss SQL statements.
- Google Spanner is an incredibly scalable database with strong consistency. Almost feels like it is impeccable.

Server Organization

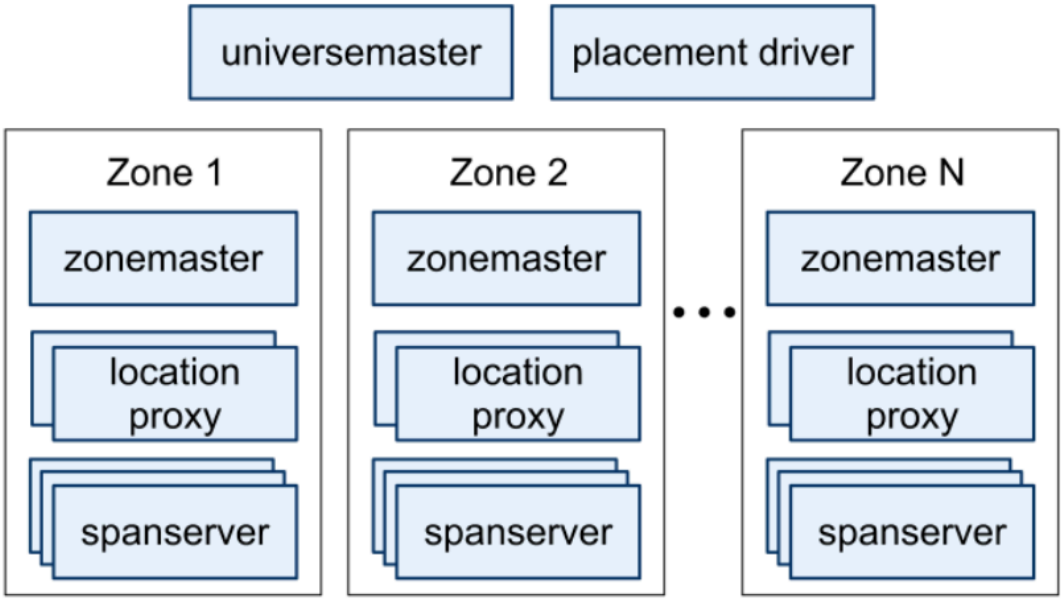

SpanServer Stack

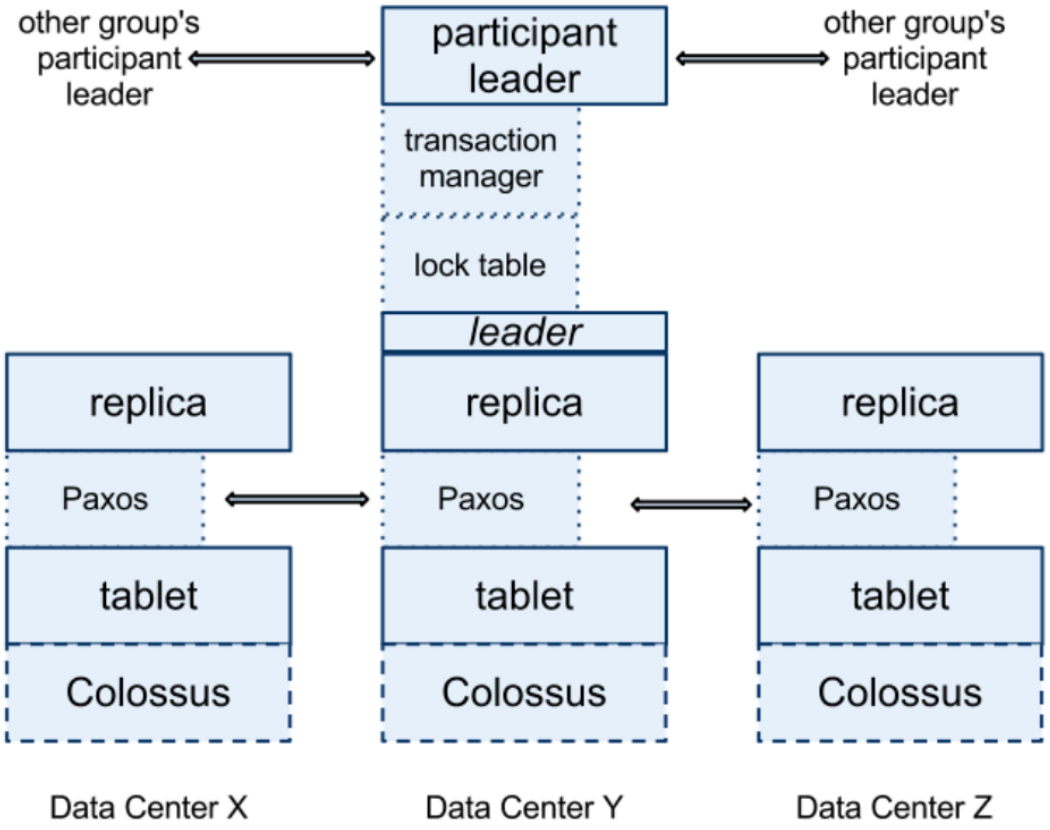

Some Thoughts

- Rows, Bags, Buckets, Tablets, Replicas, and a two-tiered consistency protocol… wow!
  - Tradeoff: describing logical model vs. actual components
- The responsibility of maintaining data locality is placed on the user
- What are the pros and cons of driving in this SQL-like direction?
  - Can it be misleading? (Performance, capabilities…)

Today's Discussion

- Split into groups of **4** (or **5** if you **really really** want to)
- Discuss your opinions on:
  - Most innovative/interesting part of the approach
  - Potential downsides
  - What parts of the paper were unclear?
  - What would you change about the paper or system?
- Sketch out your answers (**1** page) and then we'll merge for the class discussion
- **We'll present @ 12:20**