

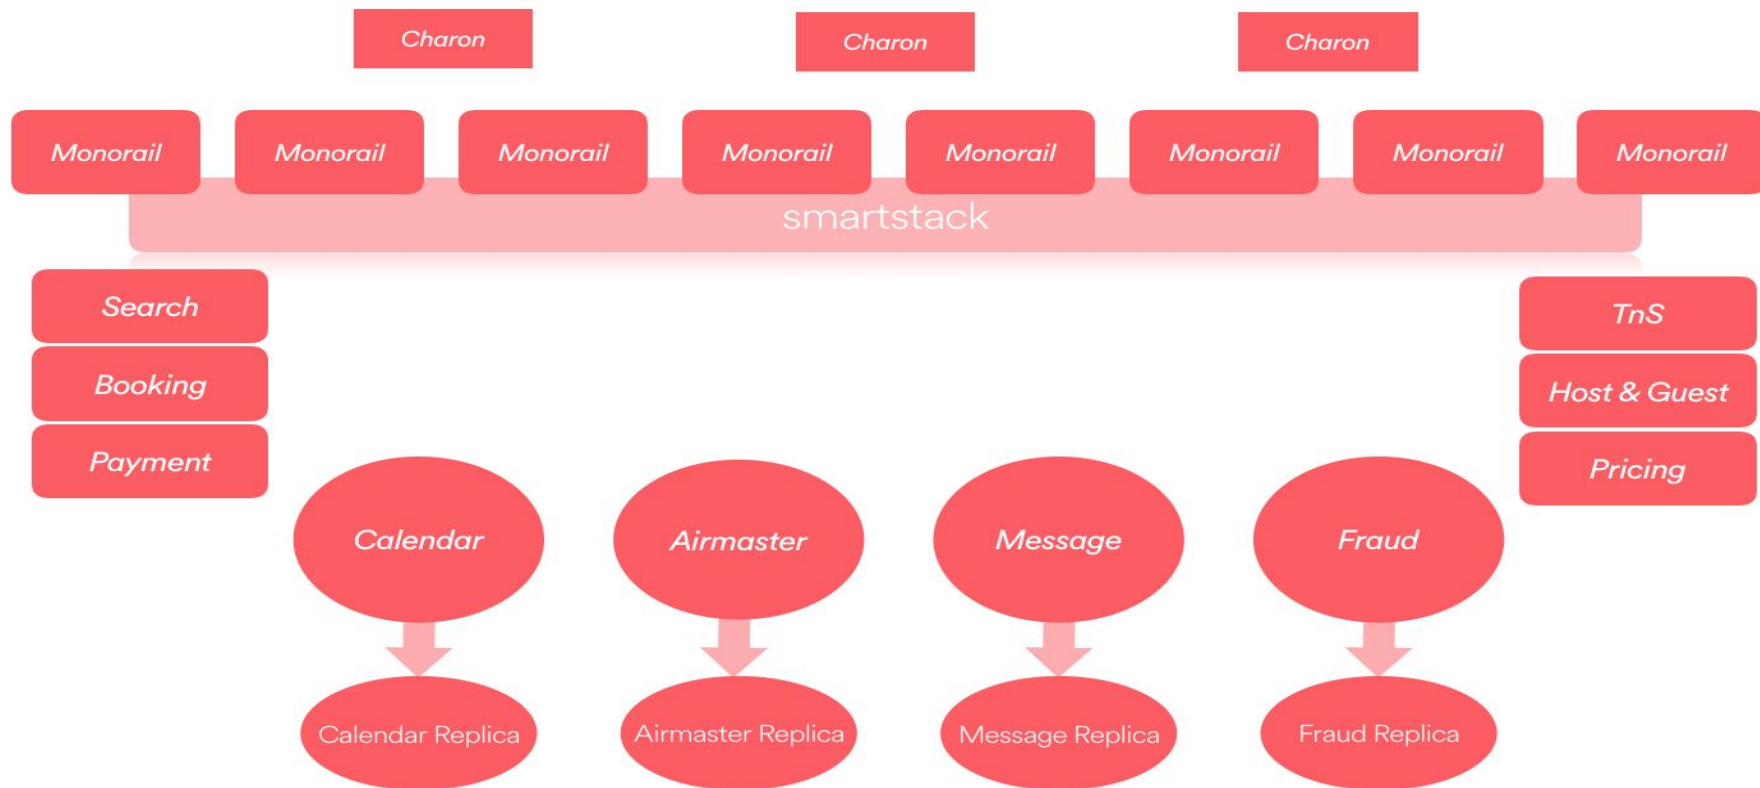
Xinyao Hu / 04/27/2017 / Percona Live 2017

Spinaltap: Airbnb's Change Data Capture System



Architecture Overview

Request Processing Path



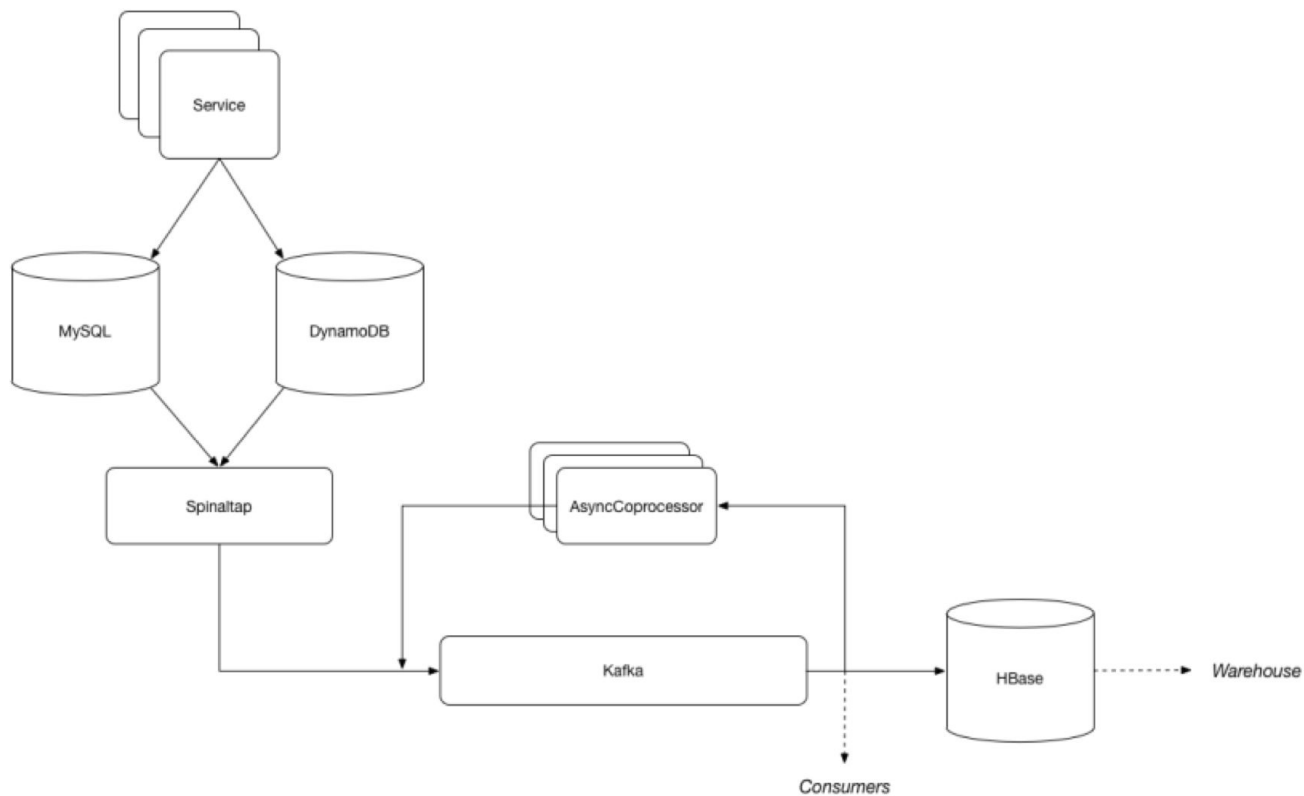
Design Principle

- Build for production
- Prepare for 5x growth
- Favor existing mature solution
- **Service should own its data.**
- **Data changes are populated through standardized events.**

Change Data Capture

- Timeline consistency
- Near realtime
- Strong consistency with source of truth data store

Architecture of Spinaltap



Requirement of Spinaltap

- Timeline consistency with at least once delivery
- Easy to add new data sources
- Support use cases with high throughput and low latency
- High Availability
- Support different types of data sources, such as MySQL / DynamoDB

MySQL as Data Source for Spinaltap

- Binlog events
 - write_rows, update_rows, delete_rows
- Use binlog + offset / GTID as logic order
- Use XidEvent to separate transaction boundary

DML (InnoDB)

Query	BEGIN
Table_map	table_id: 71 (test.person)
Write_rows	table_id: 71 flags: STMT_END_F

...

Xid	COMMIT /* xid=20 */
-----	---------------------

DDL

Query	use `test`; ALTER TABLE person DROP COLUMN first_name, ADD COLUMN last_name VARCHAR(255)
-------	--

Spinaltap for Online Traffic



Power Airbnb Search

Sync between MySQL and Lucene Index



Memcache Invalidation

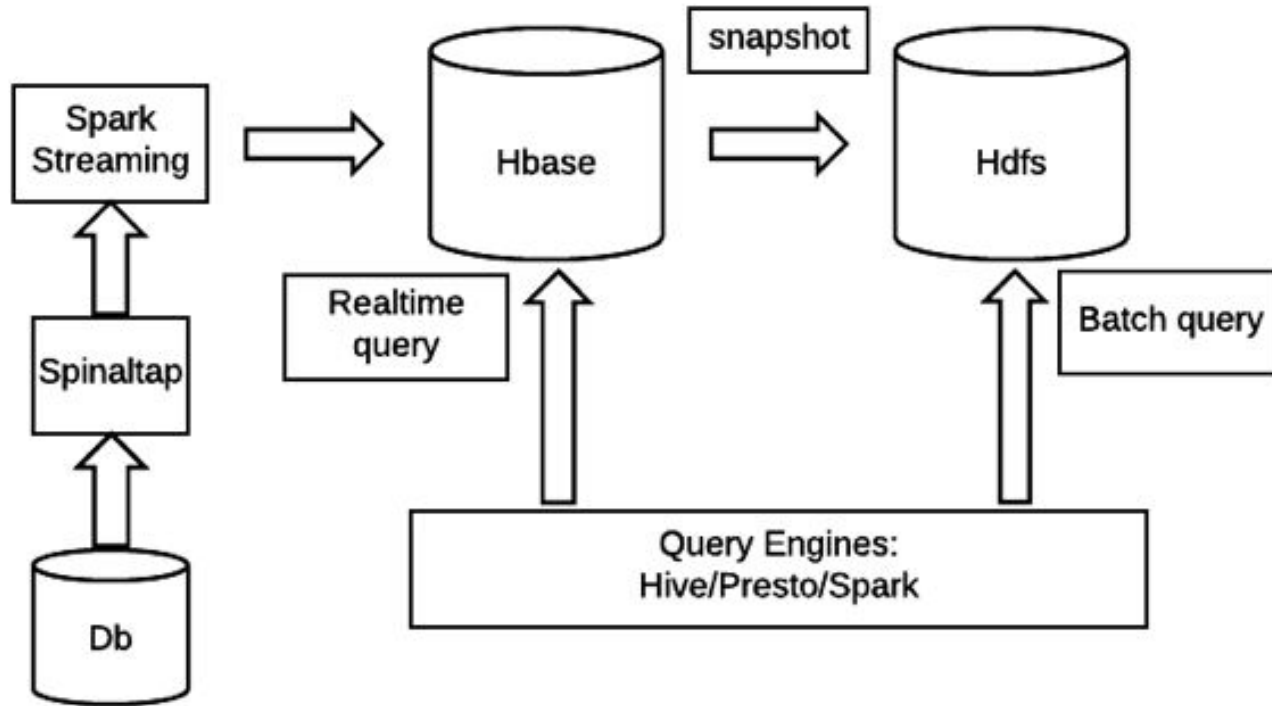
Sync between MySQL and Memcache




Review / Message Search

Sync between MySQL and ElasticSearch to power full-text search and geo based queries

Spinaltap for Offline Processing





airbnb