Build better.
Dream bigger.

#RoachFest22

RoachFest22
A CockroachDB Celebration

NETFLIX

The journey of CockroachDB in Netflix

Shahar Zimmerman, Senior Software Engineer
Shengwei Wang, Senior Software Engineer

RoachFest22
A CockroachDB Celebration

History of databases @ Netflix

2008

On-premise datacenter issue causes 3-day outage

2010

On-premise to AWS migration begins

2014

“Data everywhere” using Cassandra with global replication

2019

Cassandra is still the only Netflix-managed choice

2019: some use-cases

Cloud Drive: filesystem-like service for Netflix media assets.

Content Delivery: global control plane service for network devices.

Spinnaker: cloud continuous-delivery platform.

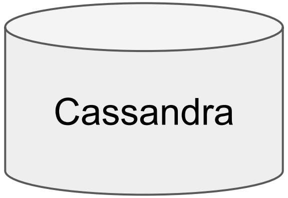

Lacks rich transactions
Difficult secondary indices

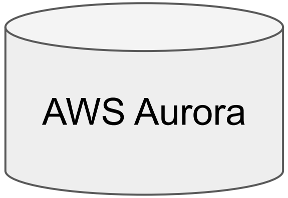

Single-active topology

2019: Why CockroachDB?

Database	CockroachDB	AWS DynamoDB	AWS Aurora
Multi-active topology?	✔	✔	Single-region writer
Global consistent secondary indices?	✔	Eventually consistent	✔
Global txns?	✔	Single-region only	✔
Full control of the system?	✔ (on-premise)	Limited	Limited
SQL?	✔	No	✔

Quick snapshot

- Since 2020
- 100+ prod clusters
- 150+ test clusters
- Most clusters are using single region topology with 3 AZ (zone failure)
- some clusters are exploring multi-region topologies
- The biggest cluster 60 nodes in one region, 26.5 TB

How do we build CockroachDB?

- We maintain a Netflix fork, manual version upgrade
- Reason for the customized build:
    - Netflix Authentication
    - Netflix logging/alarming system
    - CDC Netflix integration
- Building is a hands-on process
    - Code change
    - Testing
    - Building environment
- Version Strategy
    - couple months delay with major release, usually after X.X.5 released
    - minor release only for relevant patching

How do we deploy/manage clusters

- AWS EC2 instances
- Jenkins Build + Spinnaker delivery
- Control Panel - Antigravity, a shared control plane for multiple data solutions
- Heavily rely on the automation script
    - On instance script

Problems & future work

- Data Range Unavailable
    - Transient Instability
    - No online repairing mechanism
- Discrepancy between transaction latency and statement latency
- Using EBS
- Fully Secure Mode
- Critical Business User Case Support

Backup and restore

- Separate backup application maintaining the backup
- Restore
- Prod to Test Restore/Refresh mechanism
  - Customer want to use prod data to do test
  - In-Place Restore
  - Red/Black Refresh

Thank you