YugabyteDB's Multi-Region Deployment Options

Franck Pachot

Friday, Feb/11/2022

YFTT
YUGABYTEDB
FRIDAY
TECH TALKS

yugabyteDB

Needs to go beyond one region

Cloud availability zones

- Redundant power, networking, and connectivity
- Geographically located to avoid risks of simultaneous fire, earthquake...
- High-bandwidth, low-latency (1-2 milliseconds)

Writes in sync to the quorum, RPO=0, RTO<=3s

But…

- A region can experience total disaster: need for DR
- Public cloud region failure happens: need HA
- And on-premises deployment have rarely more than 2 AZ
- And some users may query from another country
- Some data must in a specific country for legal reasons

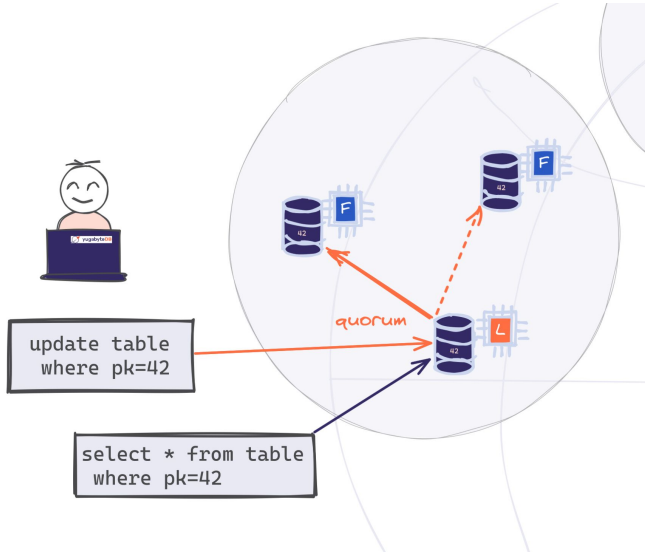

Yugabyte © 2022 – All Rights Reserved

Demo

eu-west-1a	Ireland
us-west-1a	California
eu-central-1a	Germany

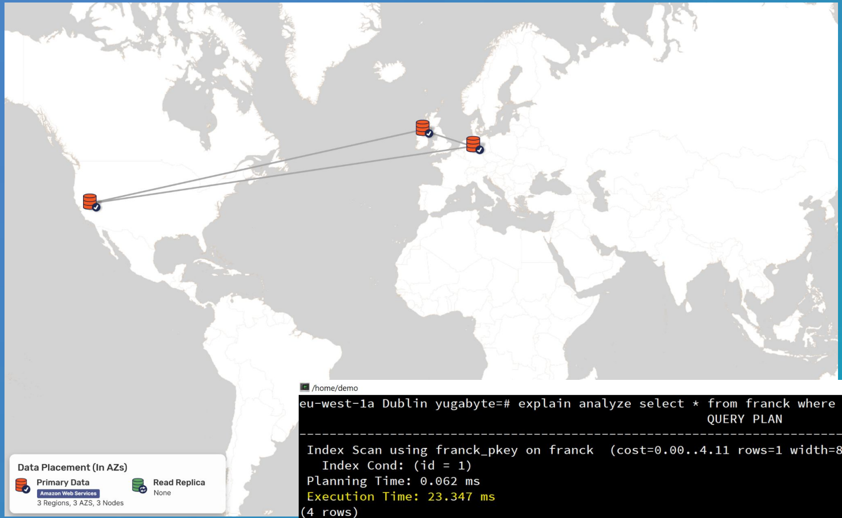

/home/demo

```
eu-west-1a Dublin yugabyte=# explain analyze select * from franck where id=1;
                                                    QUERY PLAN
-----------------------------------------------------------------------------------------------------------------------
 Index Scan using franck_pkey on franck  (cost=0.00..4.11 rows=1 width=8) (actual time=23.316..23.318 rows=1 loops=1)
   Index Cond: (id = 1)
 Planning Time: 0.062 ms
 Execution Time: 23.347 ms
(4 rows)

Time: 70.020 ms
eu-west-1a Dublin yugabyte=#
```

```
us-west-1a San Jose yugabyte=#  explain analyze select * from franck where id=1;
                                                     QUERY PLAN
-------------------------------------------------------------------------------------------------------------------------
 Index Scan using franck_pkey on franck  (cost=0.00..4.11 rows=1 width=8) (actual time=149.161..149.164 rows=1 loops=1)
   Index Cond: (id = 1)
 Planning Time: 0.134 ms
 Execution Time: 149.205 ms
(4 rows)

Time: 328.126 ms
us-west-1a San Jose yugabyte=#
```

```
eu-central-1a Frankfurt am Main yugabyte=#  explain analyze select * from franck where id=1;
                                                  QUERY PLAN
---------------------------------------------------------------------------------------------------------------------
 Index Scan using franck_pkey on franck  (cost=0.00..4.11 rows=1 width=8) (actual time=0.460..0.462 rows=1 loops=1)
   Index Cond: (id = 1)
 Planning Time: 0.056 ms
 Execution Time: 0.499 ms
(4 rows)

Time: 21.766 ms
eu-central-1a Frankfurt am Main yugabyte=#
```

yugabyteDB

Yugabyte © 2022 – All Rights Reserved

Multi-regions need latency considerations

Cloud regions

- distance = latency
  slower and less predictable

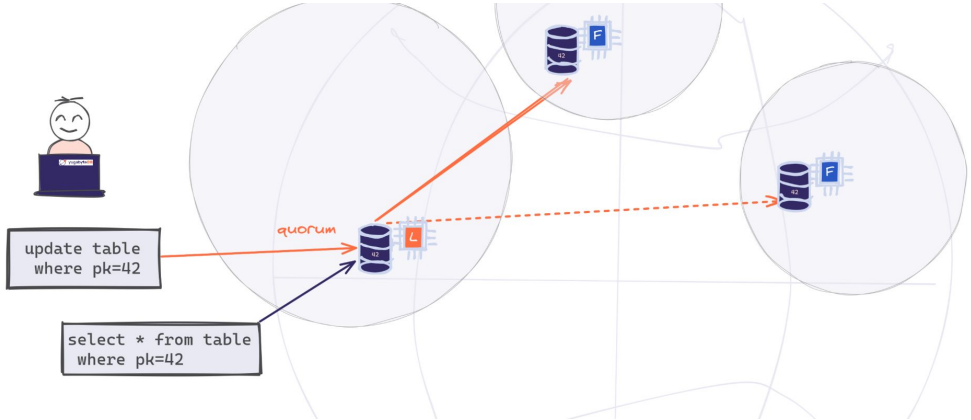

Reads from the leader: fast if pinned to the user region

Writes: the leader is local, but waits on quorum, fast if nearby region

One region should be close enough for write latency
Example:
13ms us-east1 (N. Virginia) - us-east-2 (Ohio)
60ms us-east2 (N. Virginia) - us-west-2 (Oregon)

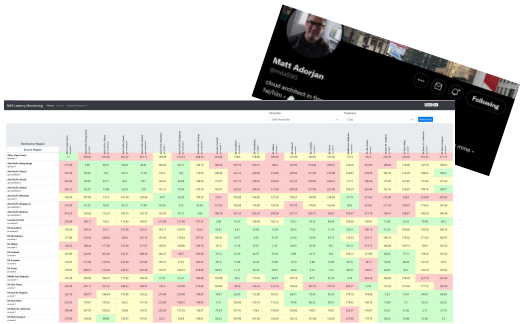

Yugabyte © 2022 – All Rights Reserved

Multi-regions topology options

More details in next YFTT episodes

- Placement per table (duplicated reference tables), per partition (geo-partitioning)
- Allow reads from followers (snapshot consistency)
- Add asynchronous read replicas (not in quorum)
- xCluster bi-directional asynchronous replication (2 DC)
- Cluster-aware client driver

https://docs.yugabyte.com/latest/explore/multi-region-deployments/

Yugabyte © 2022 – All Rights Reserved

Thank You

Join us on Slack: yugabyte.com/slack (#yftt channel)

Star us on Github: github.com/yugabyte/yugabyte-db

YFTT YUGABYTEDB FRIDAY TECH TALKS

yugabyteDB