



ROUTE

How Route Enables Always-On Data For over 1 Billion Orders and Counting with CockroachDB

Bryan Call, Senior Principal Software Engineer



RoachFest24

A tale of two Storage Failures

1

Typical SQL DB

Ruined weekend, lots of stress, lots of downtime and data loss

2

Cockroach DB

No data loss, no downtime, and a chill evening

1

Typical SQL DB

When you have a catastrophic disk failure beyond what your hardware redundancy can handle you are in for a long stressful weekend.

- Storage failure means restore from backup
- Restore from backup means downtime
- Database servers are more often pets and not cattle

2

CockroachDB

When your cloud provider's storage fails, replacing a node is no big deal and no downtime required.

- Storage failure means range lease transfers, no big deal
- Replacing storage/node involves no downtime
- Cockroach servers are like roaches, not pets

Who am I ?

Stereotypical geek

Yes I have a photo from high school with my personal SQL Server box

It all started in a .com basement startup


Postgres95, Sybase, MSSQL, MySQL, DynamoDB and more

Love code and love infrastructure

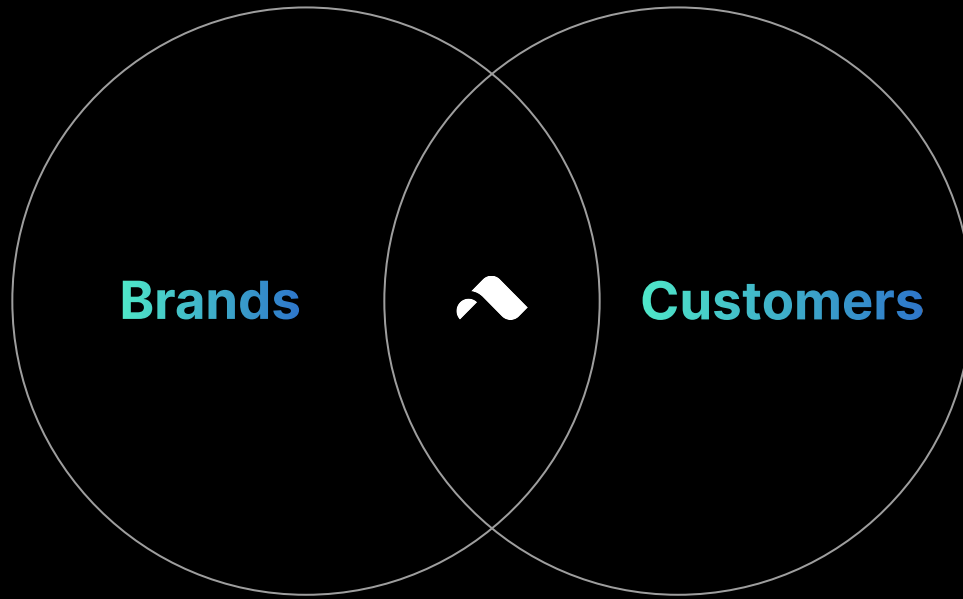
I've always preferred being truly full stack

Route has been home for 4 years

Worked with Cockroach both as a leader and as a grunt



“**Route** is the leading post-purchase customer experience solution, protecting brands and their customers with shipping insurance, package tracking, fast issue resolution, carbon neutral shipping, and remarketing.”



Founded only 5 years ago, Route serves **13K merchants** and deeply connects brands with their customers.

Delivering on Route's Mission

How has CockroachDB contributed to accomplishing Route's mission

1

Everything you order in one place

Fulfilling our mission of everything you order in one place means all your data has to be, well in one place, that can handle it all.

Our data platform has to scale!

2

Increase Revenue

We deliver solutions to increase revenue for merchants. That is unless we are down or even worse take them down.

Our data platform has to be "always-on"!

3

Increased Customer Satisfaction

Route delivers positive customer experiences for merchants when the worse happens. Spinning load icons are not positive experiences.

Our data platform has to perform!

1

Everything you order in one place

- From the beginning we knew we'd need to deal with large amounts of data. Our use cases meant it would be relational in nature.
- Having everything in one place physically and logically in a single database makes things easier.
- The dual merchant and consumer focus meant many traditional relational DB sharding techniques wouldn't work for us.
- Our CTO at the time Ryan Debenham wanted a database that would handle scaling and sharding for us without requiring we build it ourselves.

5K to Marathon

Route's Cockroach journey mirrors my own journey as a runner.

Marathon Ready

Ryan knew we would need a marathon grade database and pushed us to implement one early on.

Our Cockroach journey has involved sore muscles, and exhausting long runs.



First 5K



Second Half Marathon

Runners High

14 Half marathons later and I truly understand the term runner's high.

4 years later and I understand why choosing Cockroach was worth the effort.

How we've Scaled

Route has experienced tremendous growth. CockroachDB has scaled with us every step of the way.

3x

Node
Count

12 x

CPU Count

52 TB

Storage
Available

No Painful Data Migrations
No Required Downtime

2

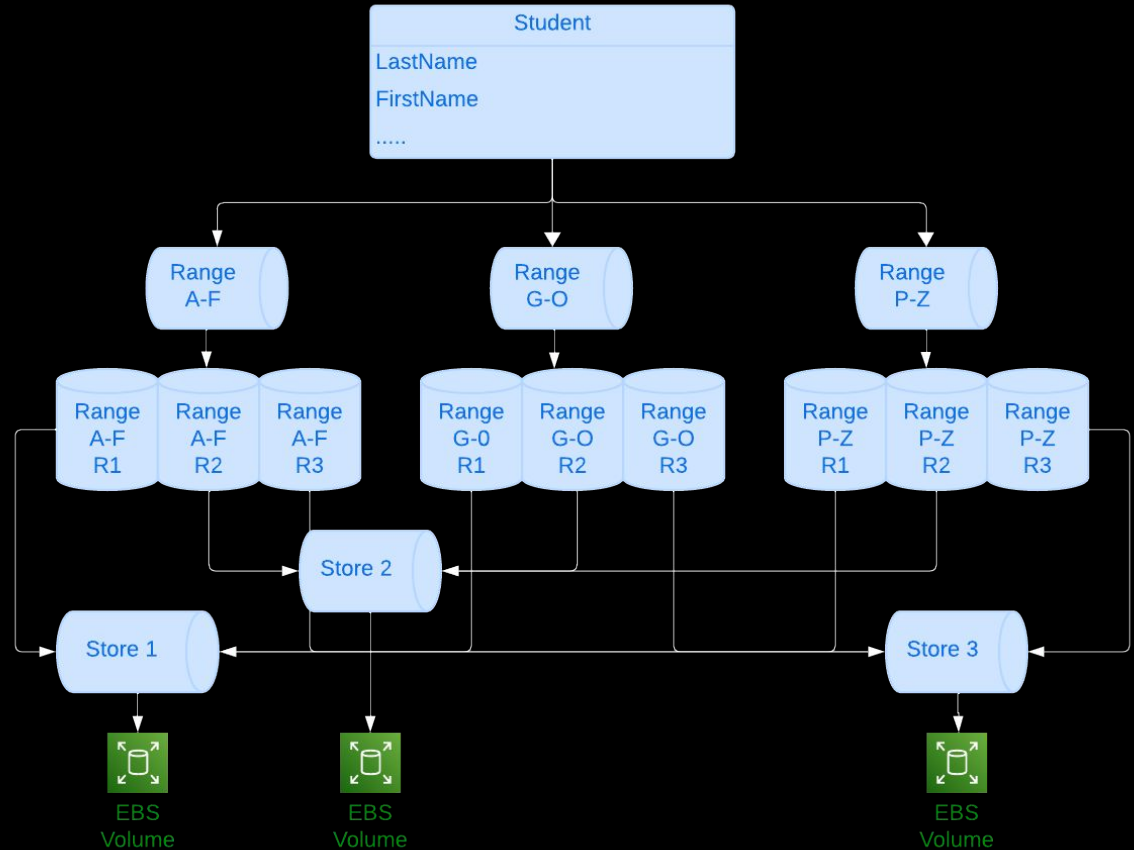
Increase Revenue

Delivering on the promise to increase ecommerce revenue means we have to be up and not interfere with the checkout process.

- Never had a major Black Friday / Cyber Monday Outage
- Survived multiple “hardware” failure issues in AWS
- Performed numerous scaling operations, hardware migrations and even a complete infrastructure re-architecture without downtime
- Patch and upgrade regularly without traditional maintenance windows

Cockroach Architecture

Cockroach divides tables into ranges and replicates them to provide a highly redundant architecture. These range replicas are grouped into stores which correspond to actual disks.




Route Cockroach Architecture

Self hosted in AWS using ECS




Beam Aboard

Our "Transporter" leverages EC2 User Init script which runs custom tooling to find a tagged group of EBS volumes and mount them on start up.

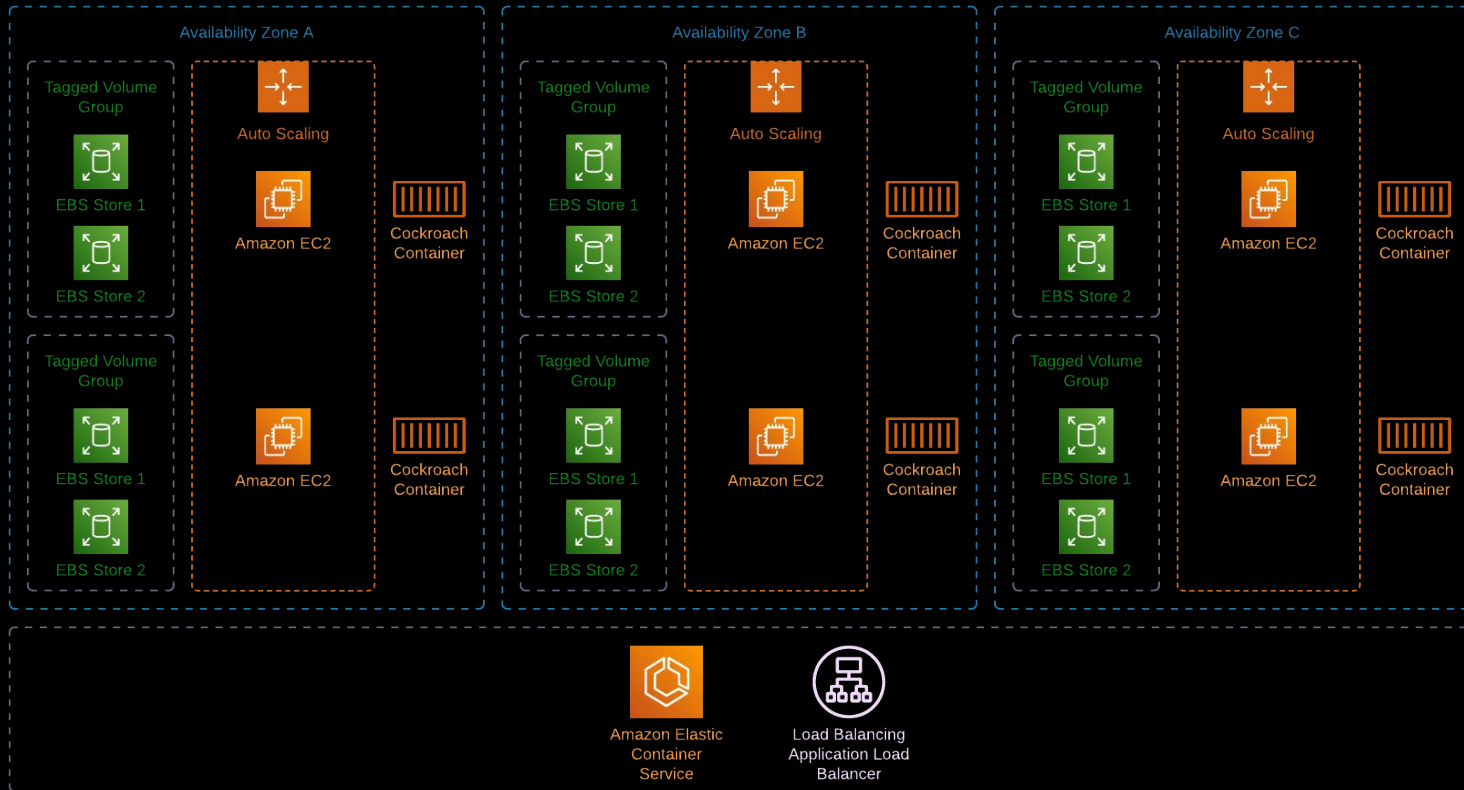


Crew



Route Cockroach Architecture on AWS

Can “step on any roach” and survive



Key Learnings

How we've learned to fully leverage Cockroach's "Always-On" capabilities.

- ✓ **Servers are roaches not pets**
Make sure your servers don't require special treatment and are easily replaced
- ✓ **Automate as much failure recovery as possible**
Minimize what humans have to do when things aren't going well
- ✓ **Train for the pain**
Regularly simulate failures so you have confidence in how things will go
- ✓ **Don't ignore a pulled muscle**
If something doesn't feel right, check it out

3

Increased Custom Satisfaction

Route helps our merchants increase customer satisfaction. Cockroach helps us give our internal data clients a positive experience.

- Separation of transactional analytics workloads made easy with changefeeds
- SQL language provides flexibility in meeting client query needs
- Understanding query plans is essential to keeping query performance high

Comparing Databases

- ✓ DynamoDB
- ✓ Traditional SQL
- ✓ CockroachDB

- You are the optimizer
 - Specify query by hash key
 - Query by hash + sort key
 - Scan
- Query performance is consistent

Comparing Databases

- ✓ DynamoDB
- ✓ Traditional SQL
- ✓ CockroachDB

- Flexible query language
- Optimizer figures out how to answer your question in an effective manner
- Variable query performance
- Limited to single server

Comparing Databases

- ✓ DynamoDB
- ✓ Traditional SQL
- ✓ CockroachDB

- Flexible query language
- Queries can leverage power of multiple servers
- A sophisticated optimizer is required

How do you optimize the OPTIMIZER

The key to query performance

✓ Just ask Glenn or Seung

Having access to an EA is an invaluable resource. Unlike many vendors Cockroach support delivers!

✓ Master the tools

Metrics console, SQL Insights and SQL Diagnostics are all invaluable resources

✓ Ode to covering indexes

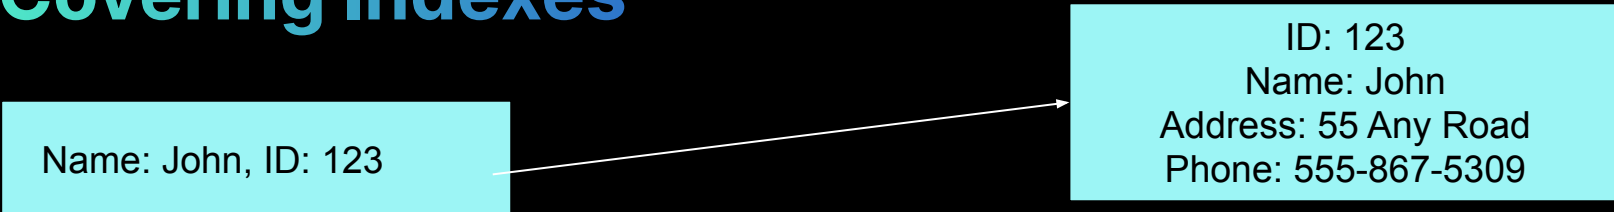
Not unique to Cockroach but extra valuable to understand given Cockroach's distributed nature

✓ Master the stats

Table stats are key to the optimizer making good decisions

Ode to Covering Indexes

Name: John, ID: 123



ID: 123
Name: John
Address: 55 Any Road
Phone: 555-867-5309

- With a normal index you look up by your indexed column and then have to look up the record by its primary key

Name: John, ID: 122,
Phone: 555-867-5309

- With a covering index the relevant data is stored with the index so no record look up by primary key is needed
- Used in other SQL databases but the cost of that primary key look up becomes potentially more expensive with a distributed database
- This makes understanding and using them more worthwhile with CockroachDB

Covering Indexes

Super useful performance tool but like anything there are tradeoffs

- ✓ **Syntax is STORING (CoIA, CoIB ...)**

IE CREATE INDEX IDX_People_Name on People (Name) STORING (Phone)

- ✓ **You are duplicating storage and IO for stored columns on updates**

Not great for highly volatile columns or if you will be writing more than reading

- ✓ **SQL Insights like to recommend them but be wise**

The recommendations sometimes include a ton of columns and may include highly volatile columns

Mile 10 Moments

Our Cockroach journey has not been pain free. We have had moments where it was hard. We've had issues to work through.



Pushing Through

By pushing through, taking the time to learn the tech, staying engaged with our account team, keeping things well maintained and investing in our decision we have totally reaped the rewards and seen the benefits of using CockroachDB.



Reaping the Benefits

A final story



Cockroach has given me personally the liberty to not worry about our most critical database going down due to some random hardware failure.



While I was enjoying Bermuda's pink sand beaches, an EBS volume was failing. No one had to try and reach me and no one was anxious to tell me about it when I got back. It was a "non event" for the team.

Just another example of how Cockroach provides "always-on" data and the freedom to rest easily.

The End

~ That's all folks ~