

A red Bugatti Veyron is shown in motion on a road. A cartoon gopher is sitting in the driver's seat, looking forward with wide eyes. The car is a two-door sports car with distinctive side air intakes and multi-spoke wheels. The background is a blurred landscape with trees and a clear sky.

Vitess (vitess.io)

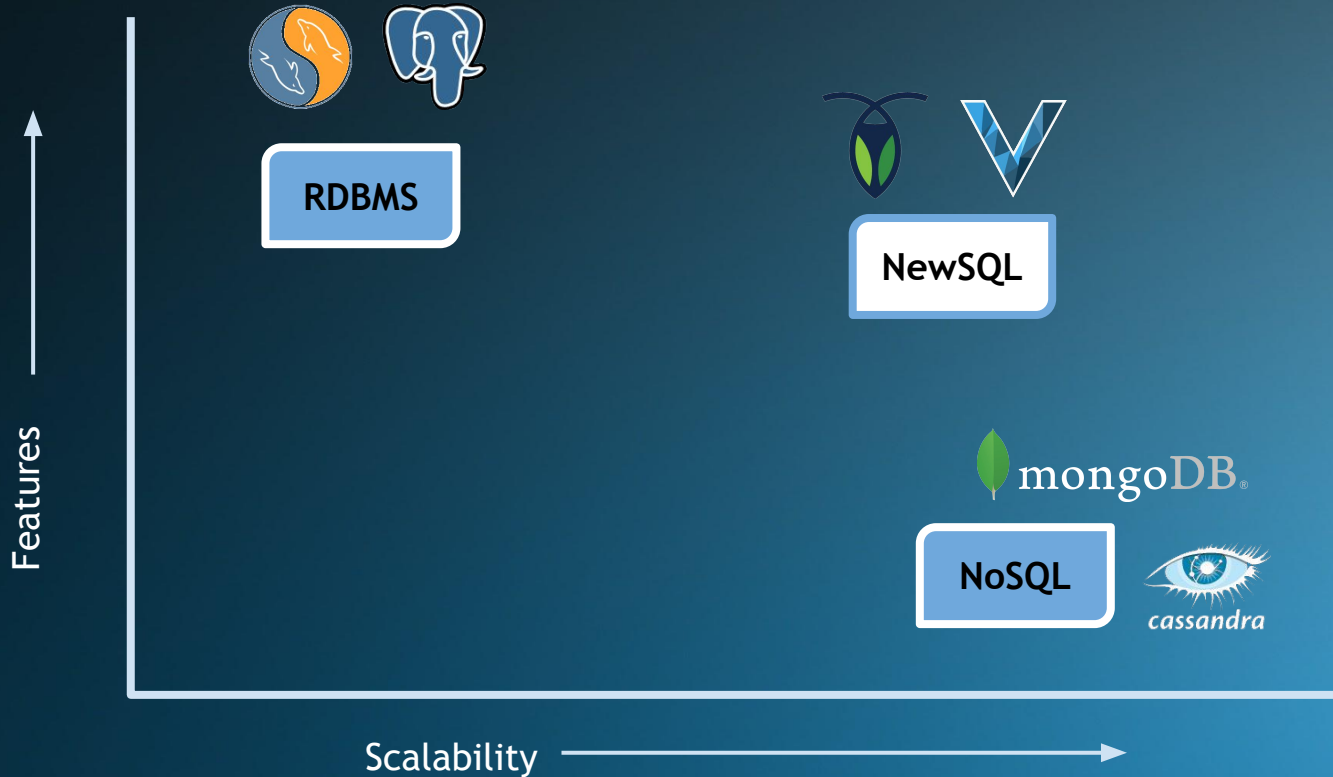
Sugu Sougoumarane
sougou@google.com
ssougou.blogspot.com

 @ssougou

Mission

To be the best performing, scalable and most available NewSQL storage solution in the Cloud based on MySQL

Database Landscape



Vitess Strengths

- Battle tested for 6 years at YouTube
- Site availability of 99.999 or greater
- Based on MySQL
- Minimize app changes
- Pluggable sharding
- Diverse environments

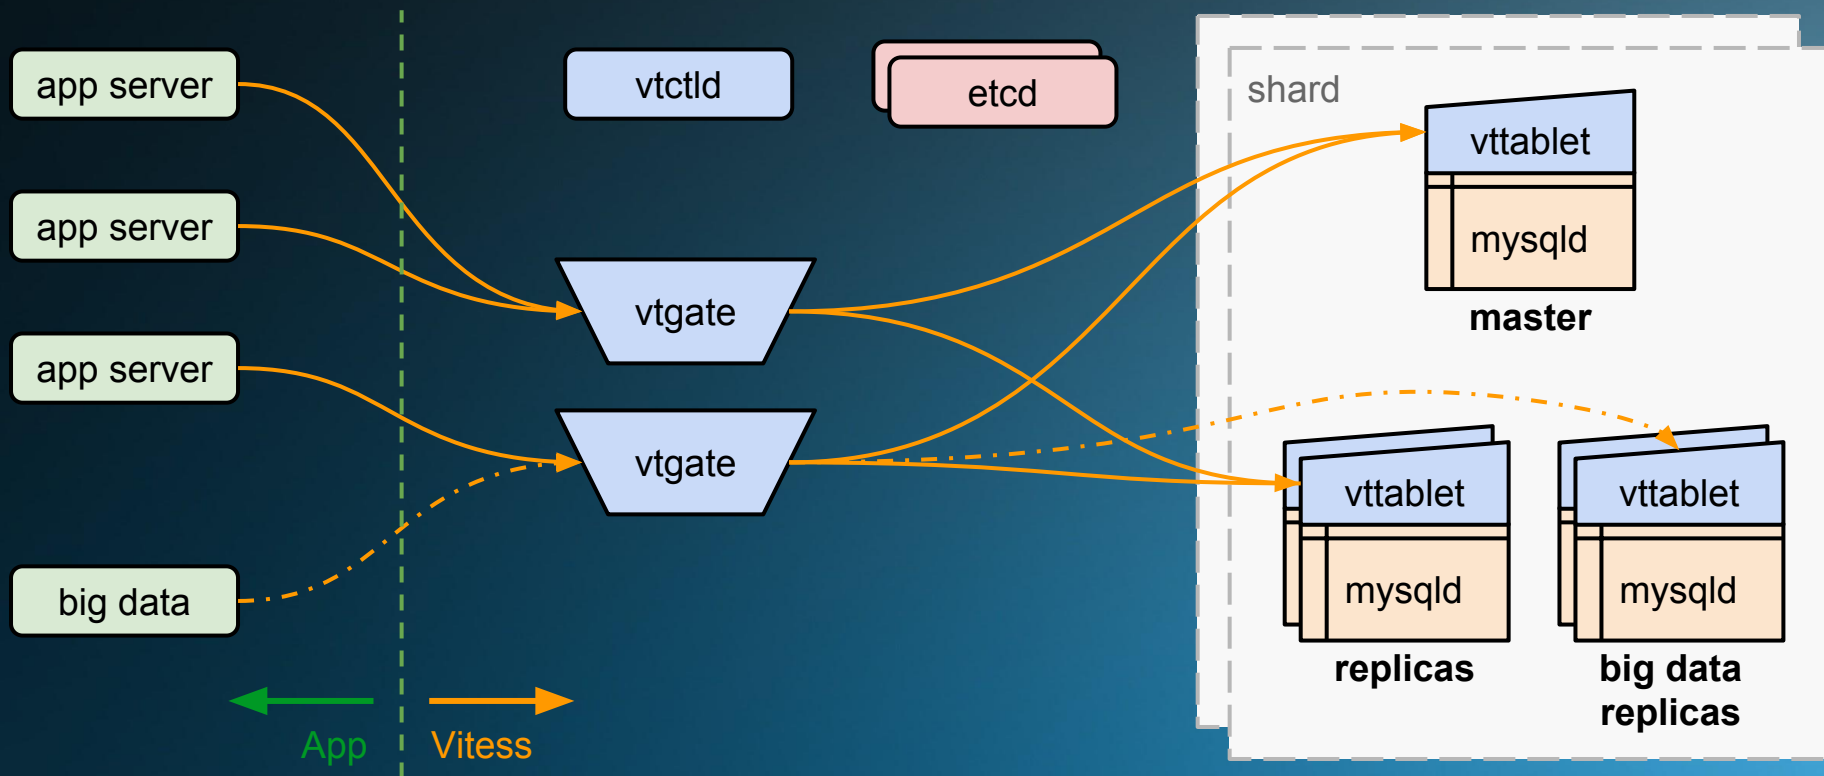


Rich Feature Set

- Cluster management
- Resharding automation
- Query services / SQL support
- Production ready



Architecture



Rethink Sharding

- Sharding Key is a NoSQL concept
 - You have only one key
- Relational parallels
 - Primary keys
 - Secondary indexes
 - Unique & non-unique indexes
 - Foreign keys



Vitess Sharding

- Range sharded by keyspace id
 - Not a physical column
- Primary Vindex maps to keyspace id
 - Multiple mapping functions
 - Pluggable
- Sequences



Secondary Indexes

- Lookup indexes backed by other tables
- Sharded/Unsharded
- Can be external



Vitess Community

Production & Testing




HubSpot

Flipkart



slack

Square



stitchlabs



BetterCloud

Contributions

Flipkart . JDBC & bulk DML features

Hubspot . MySQL compatibility

StitchLabs . Static map index

PixelFederation . PHP PDO

Slack . DBA statements



New shinies

- 2PC
- Message Queues
- RBR
- MySQL protocol
- Online schema swap

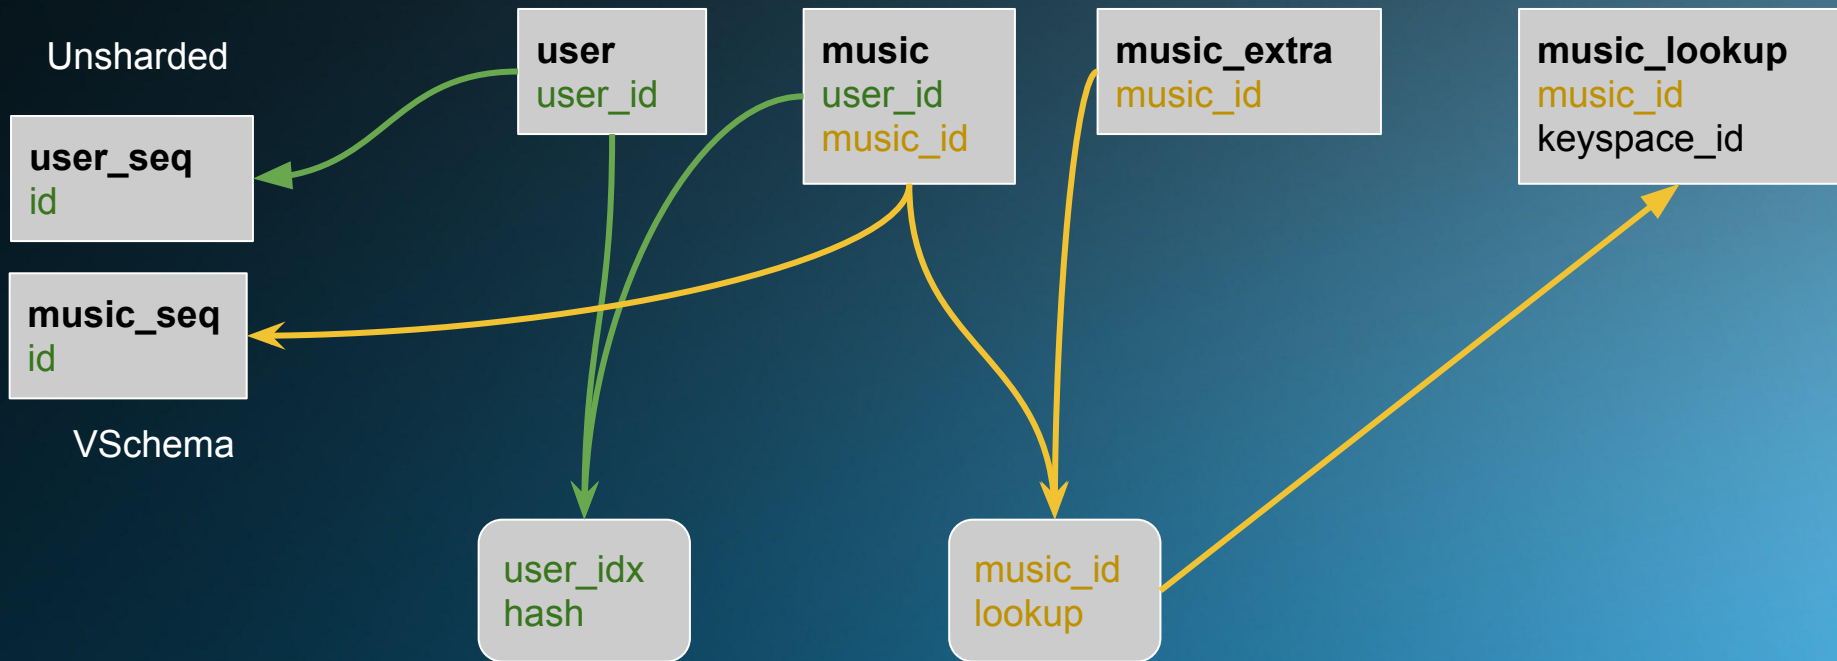


Looking ahead

- Seeking to join CNCF
- 2.1 release
- 3.0 MySQL-esque



Demo



Scaling Strategies

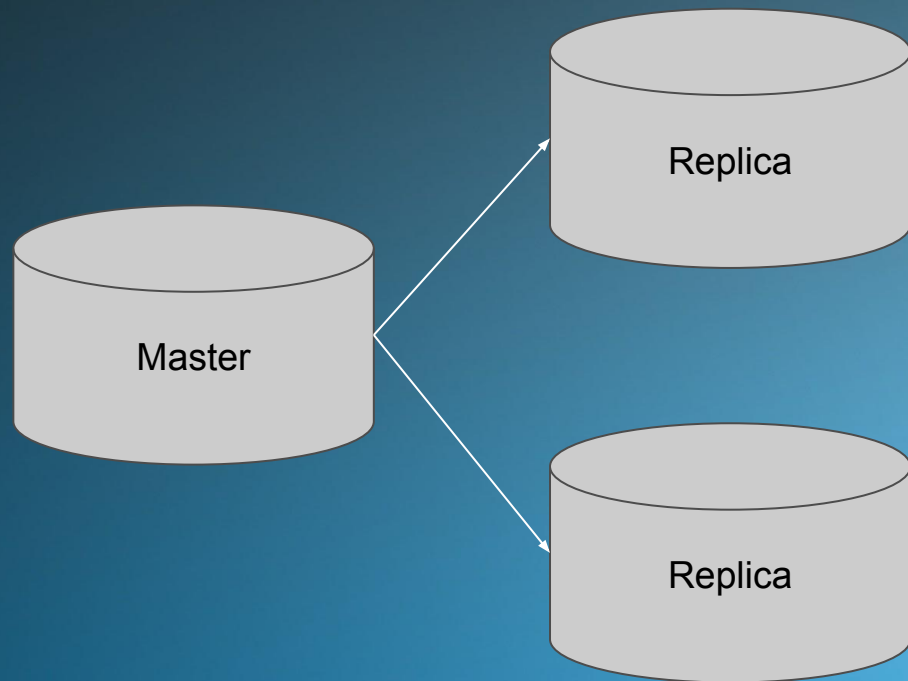
Simple Start

- **Pros**

- Simple
- Read scaling
- Durability
- Uptime: Reparents & offline backups

- **Cons**

- Write scaling
- Data size



Reshard

- Pros

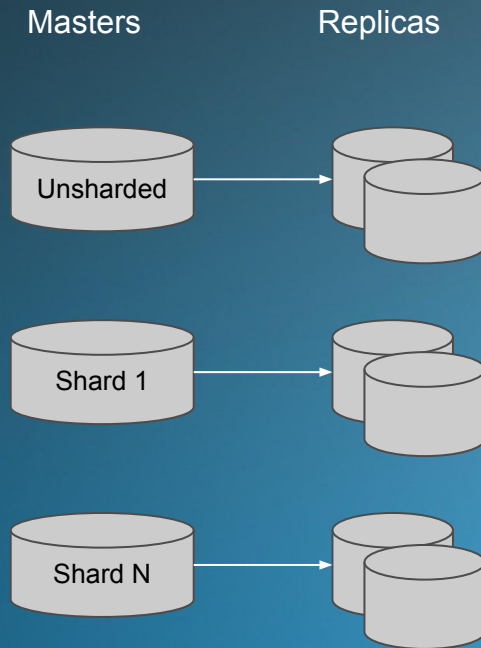
- Read and Write scaling
- Data size

- Cons

- Transactions
- Cross-shard queries
- Fat client

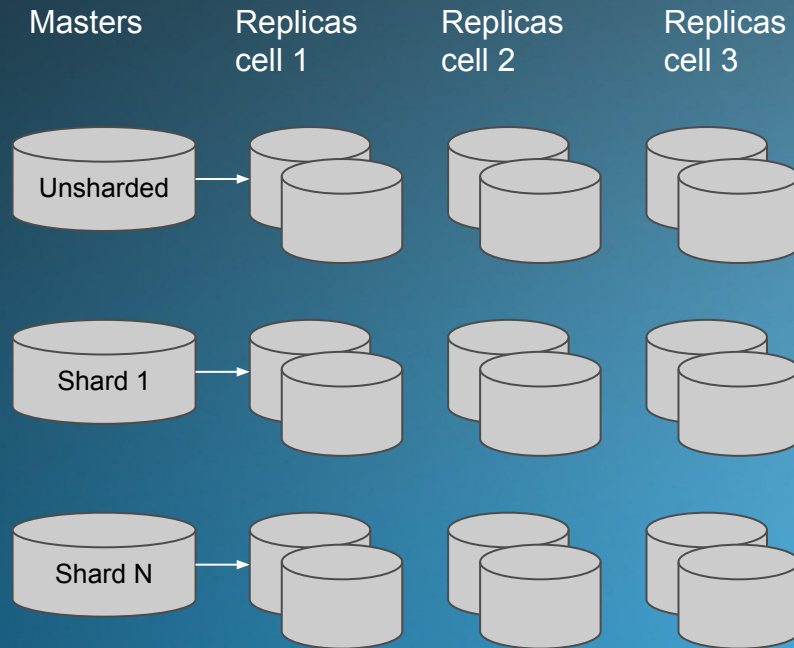
- As you grow

- Cluster management
- Worldwide latency
- DR

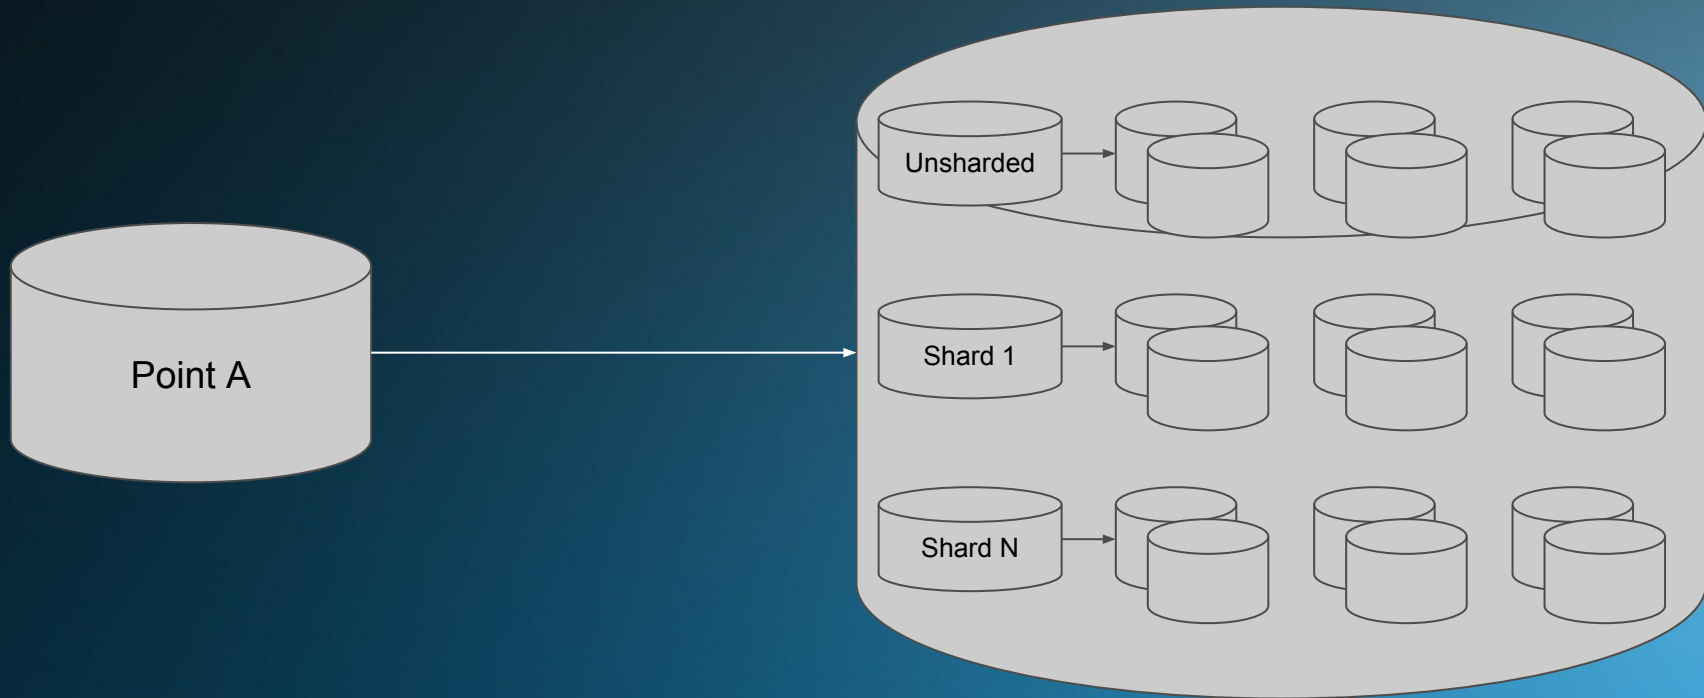


Planet Scale

- Pros
 - DR
 - Worldwide latency
- Cons
 - Cluster management
 - Resharding
 - Gets worse over time



Scaling With Vitess



Why Use Vitess?

- Sharding is a core competency
- Contribute instead of building your own
- Proven track record
- Built to scale: $O(1)$ operator overhead
- The cloud is coming

