Scaling to 920M users

Migrating One of the Largest Streaming User Management Platforms to TiDB

Empowering Scalability, Performance with Simplicity

Mydbops MyWebinar : 48

Kabilesh PR
Founding Partner, Mydbops

Mydbops
Scaling Databases

Mydbops
Scaling Databases

Your Trusted Open Source Database Management Partner

With 9+ Years of Expertise

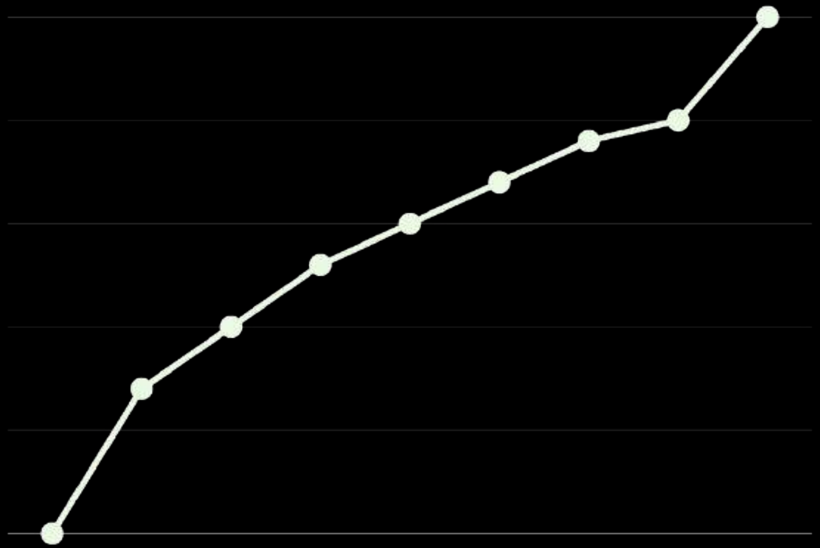

Mydbops by the Numbers

9+ years
Of Expertise

800 +
Happy Clients

10 B +
DB Transactions Handled per Day

6000 +
Servers Monitored

3000 +
Tickets Handled per Day

Database Technologies

Mydbops Services

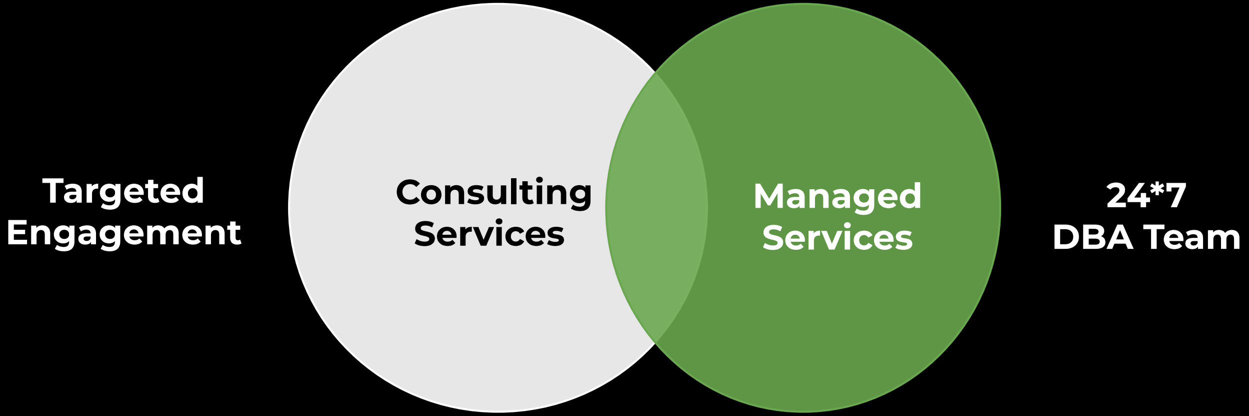

Focus on MySQL, MariaDB, MongoDB, PostgreSQL & TiDB

About the Customer

Mydbops
Scaling Databases

About the Client:

- **Business Domain:** SaaS / Cloud software for subscription billing, customer lifecycle management, monetization.
- **Target Industries:** Digital media, telecommunications, video service providers, OTT, streaming, digital subscription platforms.
- **User Base :** Serving in 180 Countries with 920 Million end users

DB Architecture

Running with MySQL 8.0.x

- **Writer**: 48vCPU & 384GBRAM
- **Reader 1 Instance:** 48vCPU & 384GBRAM
- **Reader 2 Instance:** 96vCPU & 768GB RAM
- Using provisioned IOPS 16k
- Datasize 7TB
- Multi AZ deployment.

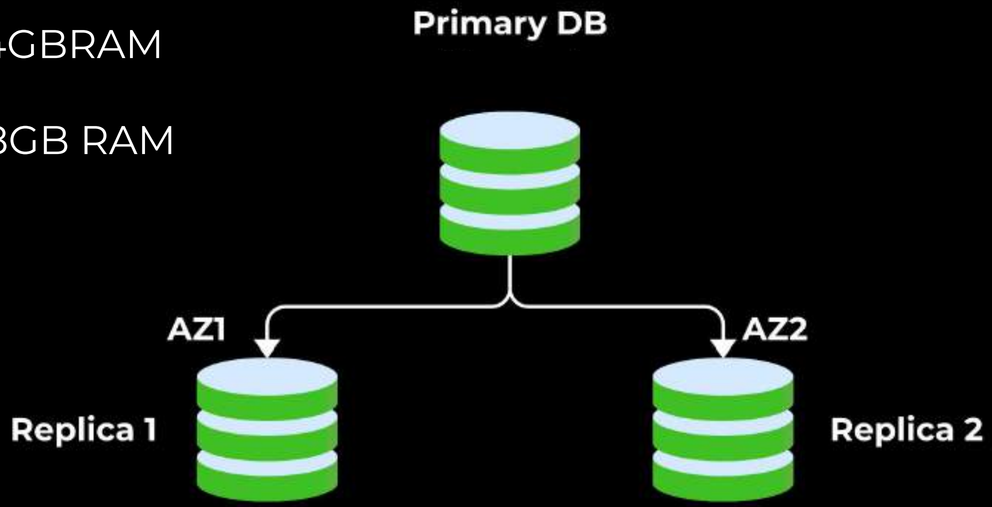

Mydbops
Scaling Databases

Scaling Requirement

Requirement

For upcoming, greatest rivalry game event on their platform, They were anticipating traffic at a scale 3 Million QPS importantly the DB p99 latency to be within 20ms.

Application TPS : 25k

Mydbops
Scaling Databases

Initial Scaling

Phase 1

- All Instances was scaled to R6g.48xlarge
- 2 replica were added more to have read scalability

Result :

With the Maxed-out Instances, it was able to handle 10k TPS ie. 300k QPS with expected latency range. Remaining tests were failed.

Mydbops
Scaling Databases

Phase 2

- Engineers worked on optimising MySQL compatible db
- Multiple config changes and patch was applied as well
- **Redis Cache layer:** Introduced to reduce read pressure
- **Keyspaces(Cassandra):** Two heavy write, log based tables pushed.

Result :

With supporting components, It reached 12k TPS, ie., 365k QPS

Mydbops
Scaling Databases

Database Performance Bottlenecks

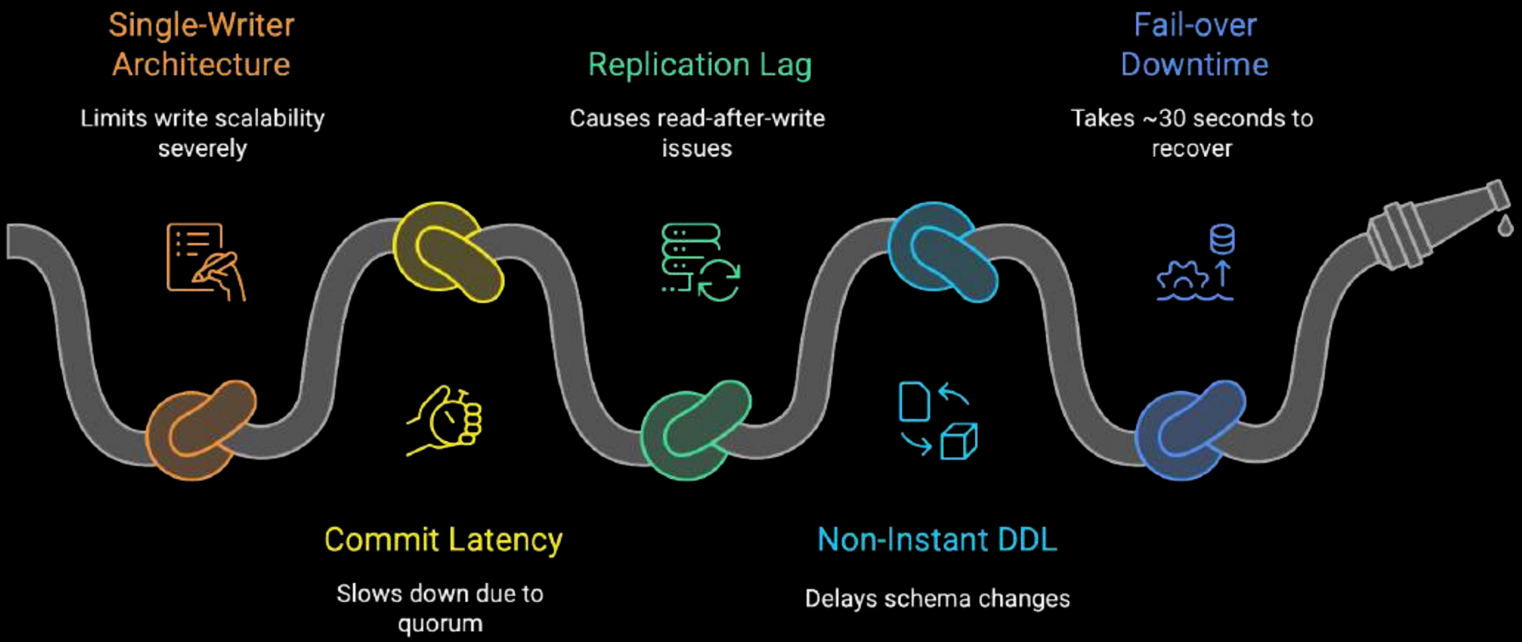

Mydbops
Scaling Databases

The TiDB Introduction

Mydbops
Scaling Databases

Why TiDB?

- Evaluated as MySQL compatible, Distributed SQL database build for Horizontal scalability.
- Vertical & Horizontal scaling with Automatic Sharding
- Multi-master distributed Architecture.
- HTAP capabilities.
- Zero code changes
- Open Source

Why TiDB?

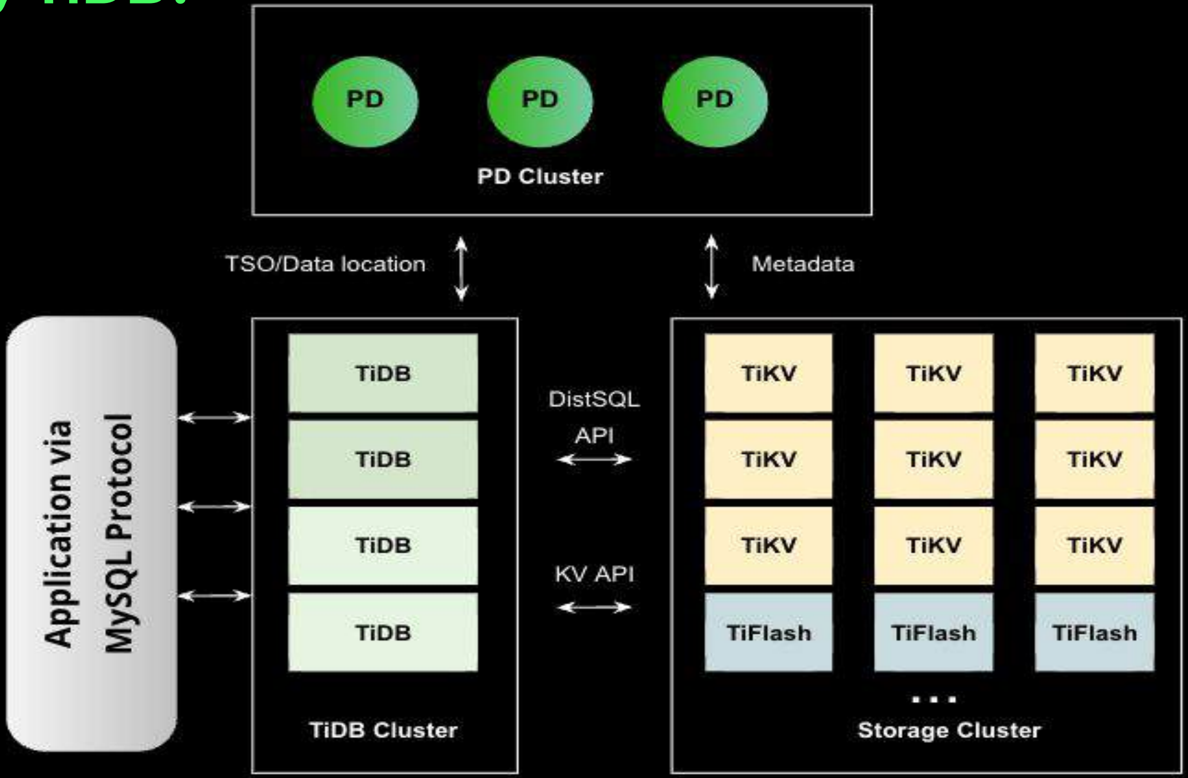

PD
PD
PD
PD Cluster
TSO/Data location
Metadata
Application via MySQL Protocol
TiDB
TiDB
TiDB
TiDB
TiDB Cluster
DistSQL API
KV API
TiKV
TiKV
TiKV
TiKV
TiKV
TiKV
TiKV
TiKV
TiKV
TiFlash
TiFlash
TiFlash
...
Storage Cluster

The Migration Journey

Migration

We adopted Zero Downtime Migration Plan, ie., TiDB would act as replica for MySQL.

Steps:

- **Schema Migration:** Recreated skeletal structure with TiDB DDL Compatibility.
- Data loading done with DM cluster followed by CDC using binlogs

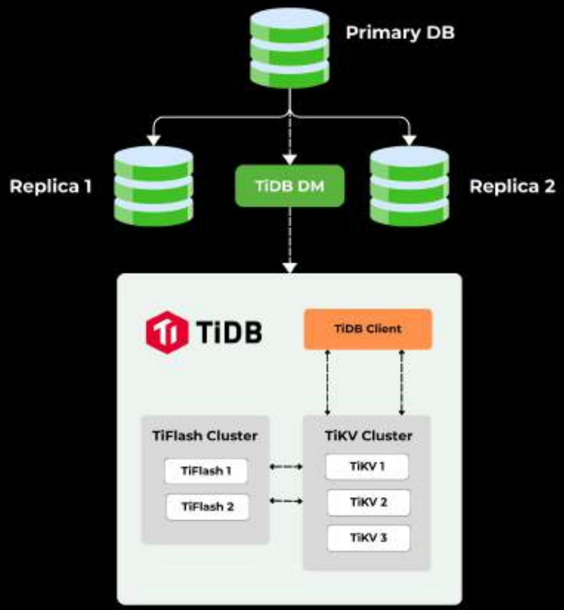

Mydbops
Scaling Databases

Migration

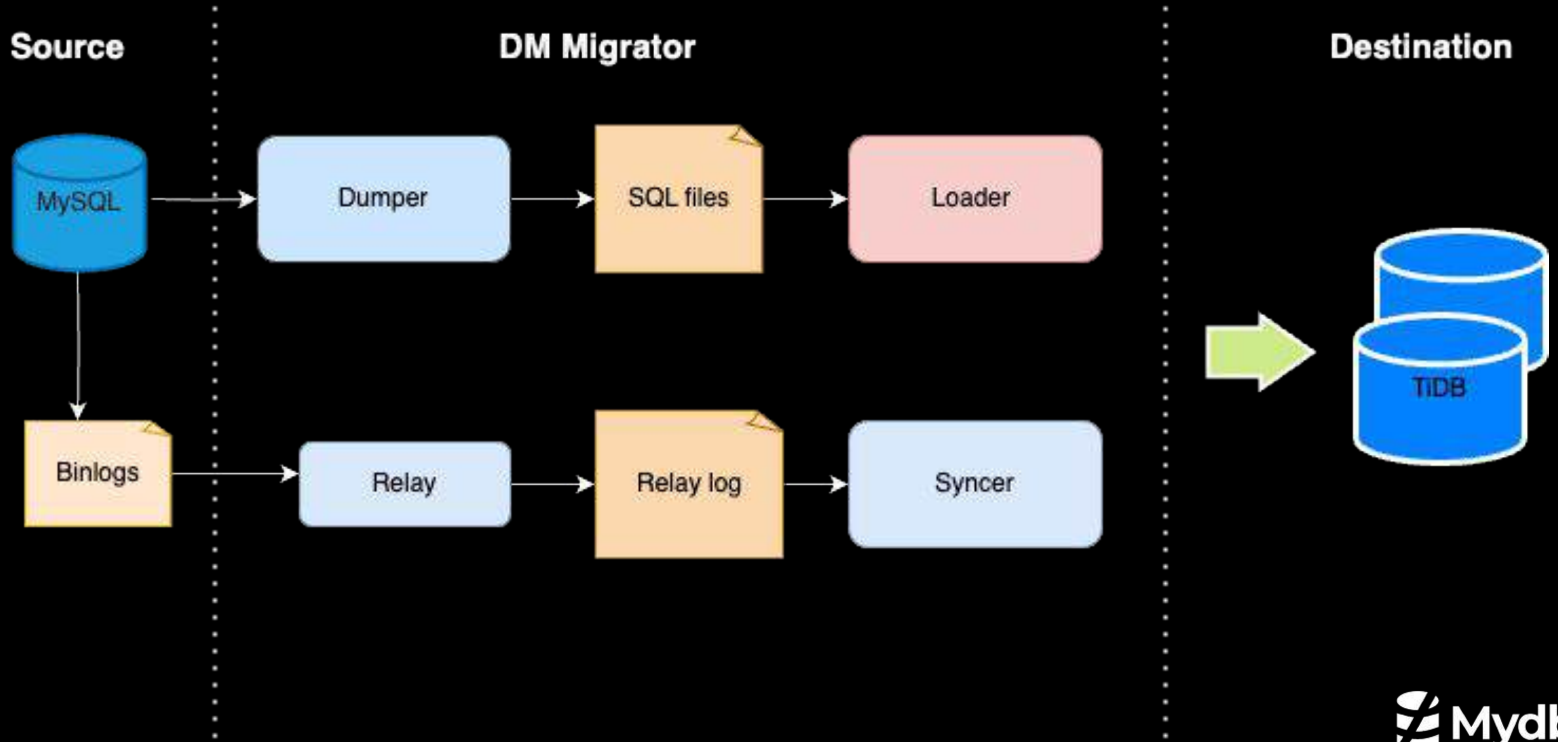
Source
DM Migrator
Destination
MySQL
Dumper
SQL files
Loader
Binlogs
Relay
Relay log
Syncer
TiDB

Migration

- Total 7 TB of data from MySQL Compatible database was Migrated in less than 28 Hours, with replication – Thanks for Physical loading of DM (Lightning)
- Storage was reduced by 79% with default TiDB with no perf Impact, 1.5 TB
- Primary Key changes done along with TiDB restore

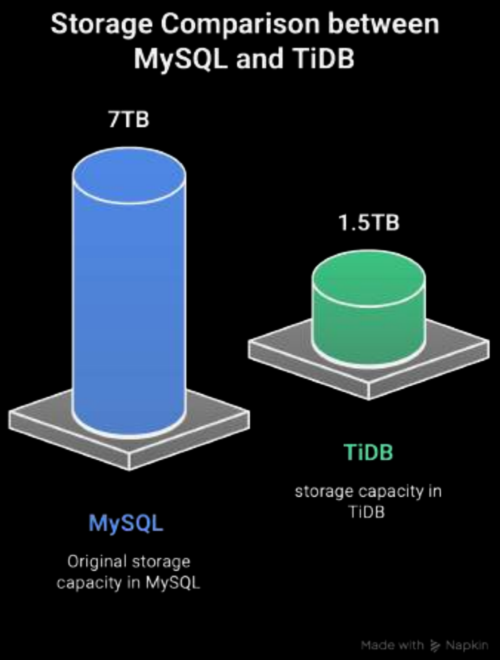

Mydbops
Scaling Databases

BenchMarking TiDB

Load Testing

- Series of Load test runs 12K → 15K → 20K → 25K TPS
- Dynamic scaleout of TiDB and TiKV to manage High concurrency and low latency.

TPS(k)	PD / TiDB / TiKV	CPU per node PD / TiDB / TiKV	CPU - PD / TiDB / TiKV (avg%)	P99	Read / Write / QPS(k)
25+redis+cassendra	3 / 90 / 56	16 / 32 / 48	65 / 60 / 65	14	1.44M / 1.82M / 3.26M
20+cassendra	3 / 60 / 45	16 / 32 / 48	55 / 55 / 60	15.6	140 / 837 / 977
25	3 / 60 / 45	16 / 32 / 48	48 / 52 / 45	8	384 / 404 / 788
22	3 / 60 / 30	16 / 32 / 48	45 / 50 / 45	7.6	327 / 347 / 674
20	3 / 60 / 30	16 / 32 / 48	45 / 35 / 40	10	300/ 315 / 615
18	3 / 60 / 21	16 / 32 / 48	33 / 35 / 38	13	164/ 390 / 554
15	3 / 60 / 21	16 / 32 / 48	35 / 32 / 35	7.4	223 / 232 / 455
12	3 / 60 / 15	16 / 32 / 48	40 / 30/ 33	7.3	179 / 186 / 365

Cluster Summary

Version

v8.5.2

Instances

121

Hosts that instances deployed

121

Σ Memory capacity (of all hosts)

35.6 TiB

Σ CPU physical cores (of all hosts)

4752

Σ CPU logical cores (of all hosts)

4752

It's All About Performance At Scale !!

Mydbops
Scaling Databases

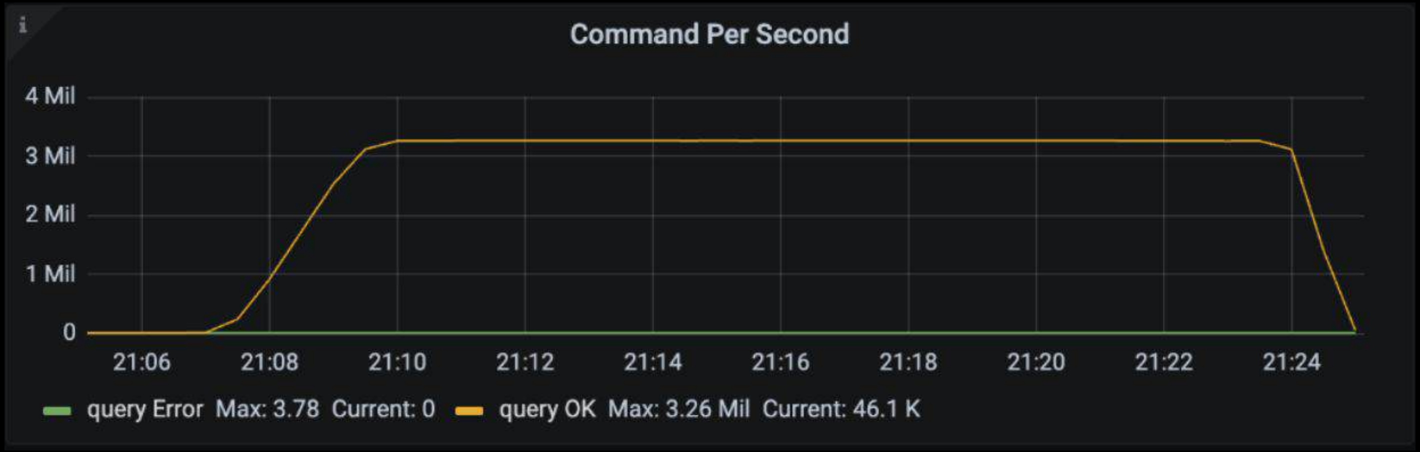

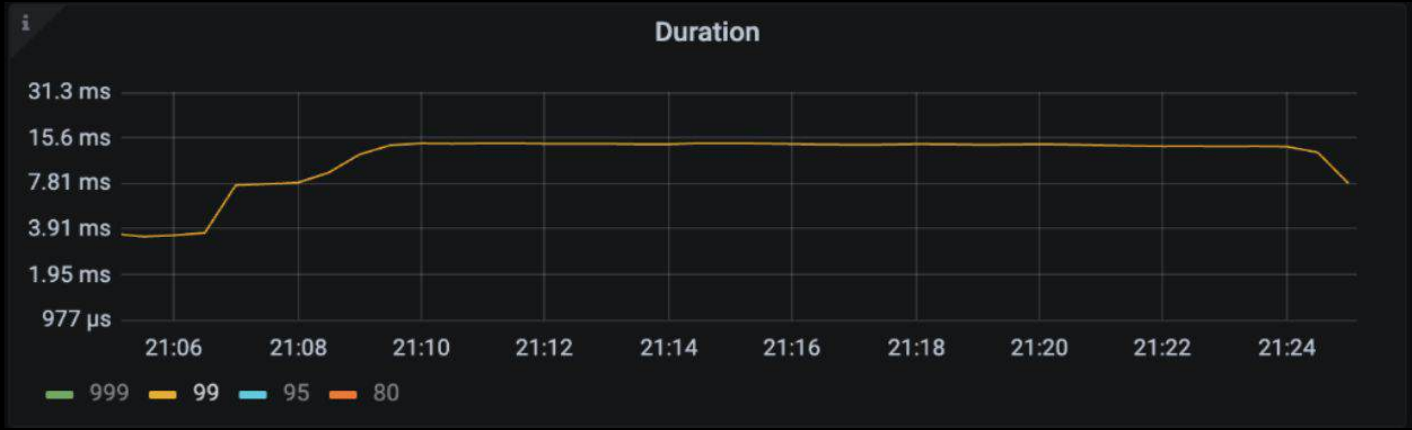

Mydbops
Scaling Databases

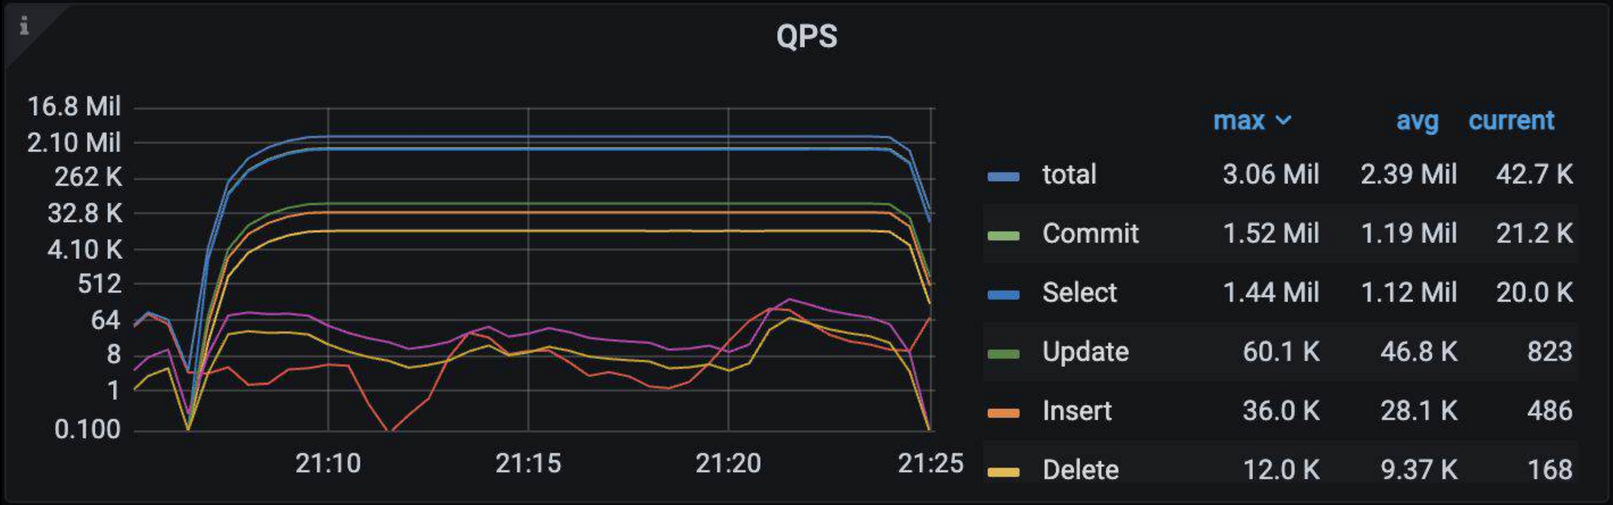

Mydbops
Scaling Databases

Writes

Mydbops
Scaling Databases

Issues & workaround

Floodgates : When Redis Fails

- When Redis fails (or evicts hot keys), every cache miss turns into a direct DB query, floods database.

- Cache Table as Saviour !!
  Master/config tables(46) were cached in TiDB (SQL layer) ie., on the wire caching, there by we not only addressed floods,but completely replaced/removed redis use case

Mydbops
Scaling Databases

Write Hotspot Issue

Write hotspot observed for , heavy write table causing high CPU with some TiKV nodes.

Hash partitions:

- Identified Hot / Heavy write tables, Enabled Hash-based partition on the PK to distribute evenly.
- With Hash partitions we were able to remove/replace keyspace.

Mydbops
Scaling Databases

Connection Imbalance With NLB

All TiDB nodes(sql) were placed under a Network Load Balancer (NLB) for connection-balancing. Due to stickiness and skewness connections were unevenly distributed at high concurrency.

Tiproxy replaced NLB:

- Even connection distribution
- Connection failover
- Auto node discovery
- Handled 60K connections, a significant increase from the previous limit of 16k on the primary DB.

Mydbops
Scaling Databases

SQL Binding - Zero application changes

Under certain scenarios optimiser tends to choose wrong query plan resulting in higher latency.

SQL Bindings:

Pin a stable execution plan to a SQL pattern without changing application code.

```
CREATE GLOBAL BINDING
FOR   SELECT c FROM t WHERE a = ? AND b = ?
USING SELECT /*+ USE_INDEX(t idx_a) */ c FROM t WHERE a = ? AND b = ?;
```

Mydbops
Scaling Databases

Slow Realtime Reporting

Real Time reports on login counts, transactions were delayed.

TiFlash for Real Time Analytics:

- Batch reports were done in less than 7 mins before 1 hour
- Removed dependencies on external analytics
- No more replication lag

Outcomes

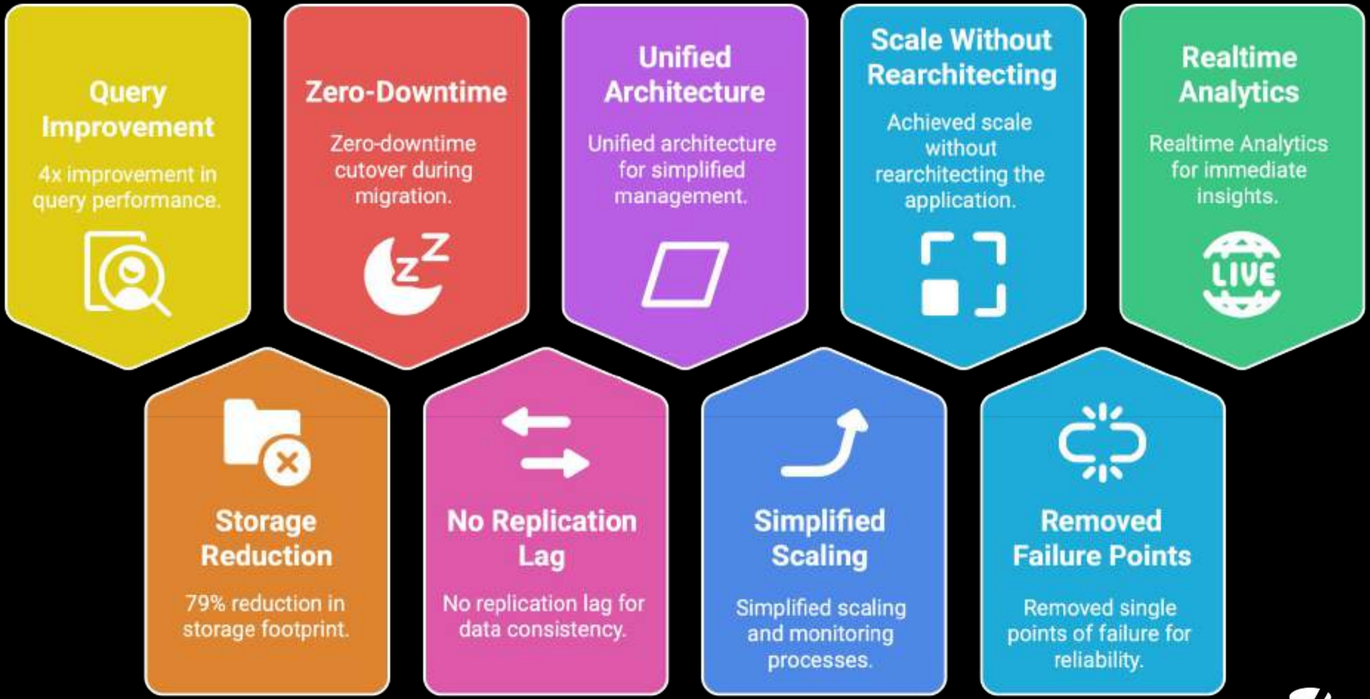

Mydbops
Scaling Databases

RollBack

Mydbops
Scaling Databases

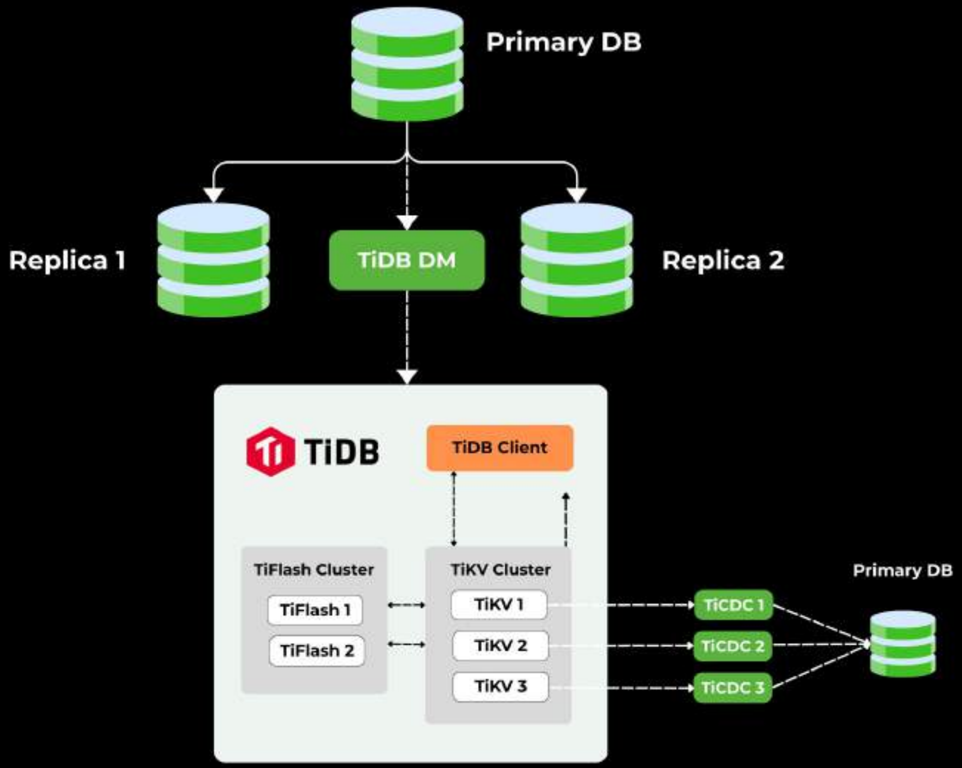
Primary DB
Replica 1
TiDB DM
Replica 2
TiDB
TiDB Client
TiFlash Cluster
TiFlash 1
TiFlash 2
TiKV Cluster
TiKV 1
TiKV 2
TiKV 3
TiCDC 1
TiCDC 2
TiCDC 3
Primary DB

Any Questions?

Mydbops
Scaling Databases

Connect with us

info@mydbops.com

+91-9686032223

www.mydbops.com

Mydbops
Scaling Databases

Thank You

Mydbops
Scaling Databases