Introduction to Exchange CDB(eCDB)

edgeyang(杨杰)

Agenda

1. Background
2. What is eCDB?
3. Key Technologies
4. Architecture

Background

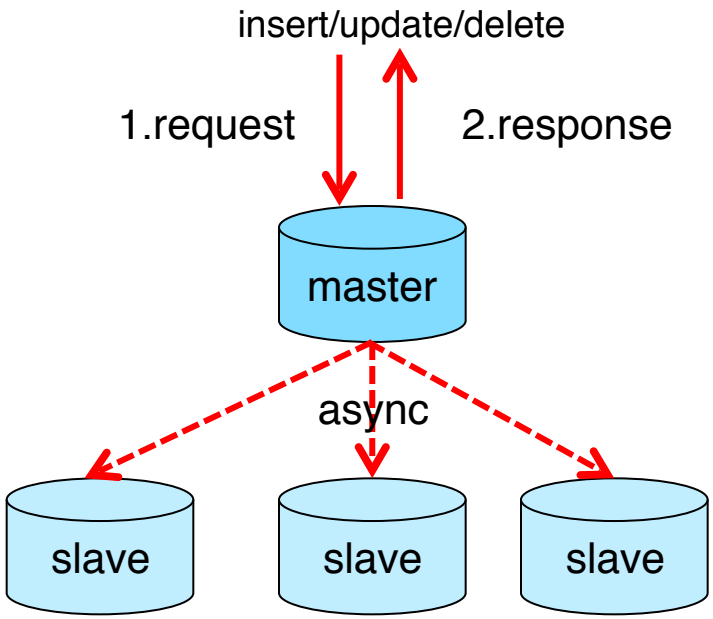

One master vs N slaves

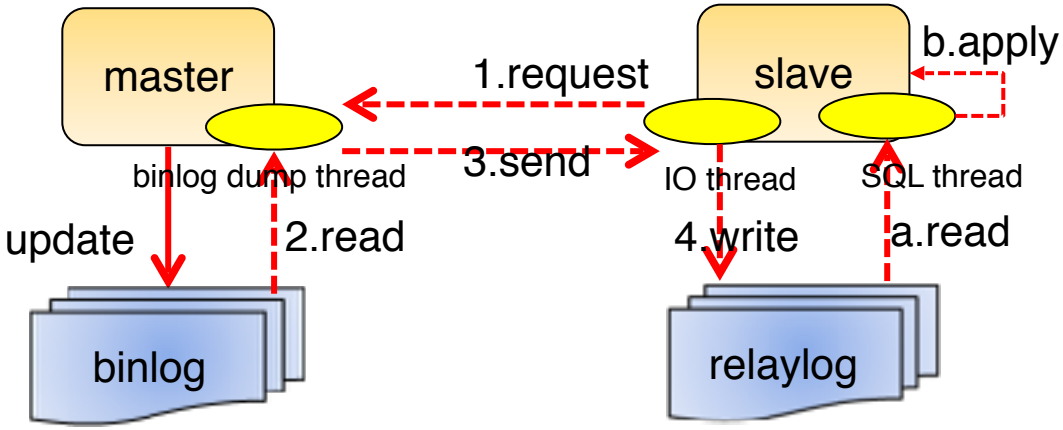

Async replication details

Data loss!

Background

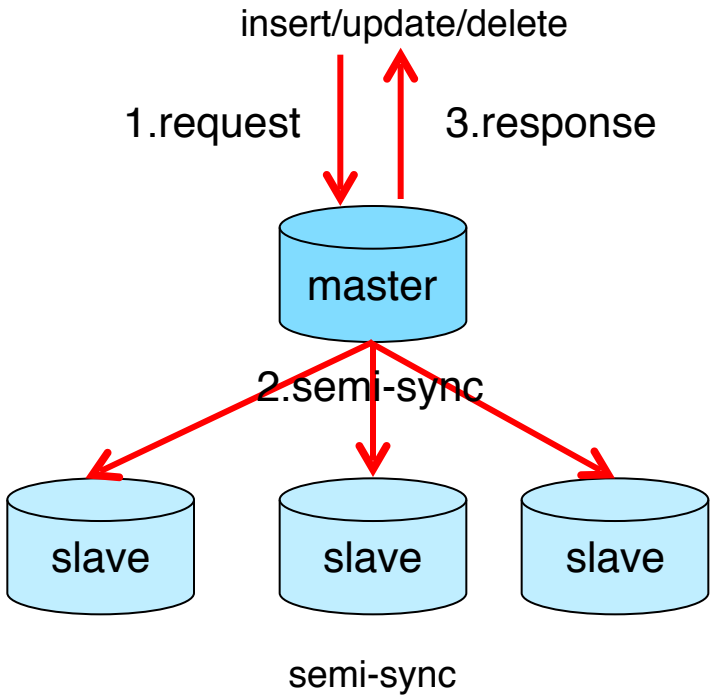

1. Weak data consistency
2. Only one slave accepts log
3. Redundant transactions
4. Low performance

Background

- p Transactional DB storage features
  - ü schema is more complex , requires transactional operation
  - ü Strong consistency
  - ü Disaster Tolerance across region

- p Pains
  - ü Data inconsistent
  - ü Longer time in service recovery (RTO)

Fault type	Operation	consistency	RTO
Recoverable	1.Reboot 2.Mysqld restart	√	10min
Unrecoverable	1.Reboot 2.Switch to slave 3. Comparison	×	>10min

What is eCDB?

eCDB = strong consistency

+ across region

+ HA

+ high performance

+ security audit

+ data sharding

+ MySQL

Customized slaves

- p Problems:
  - ü Semi-sync downgrade
  - ü Only one slave accepts log
- p Solutions：
  - ü No downgrade
  - ü Customized slaves

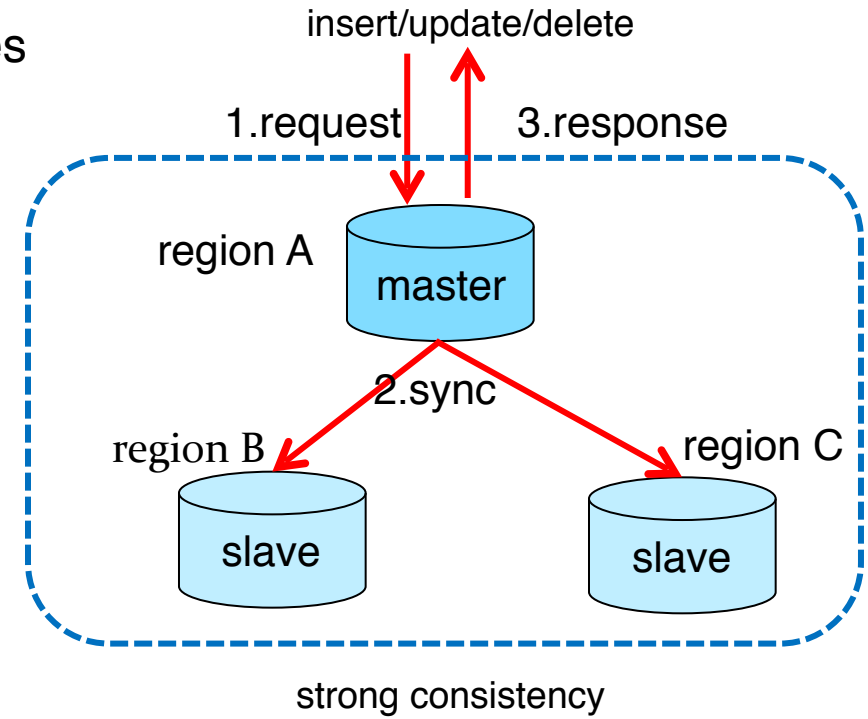

Phantom read

- Problem:
  - phantom read

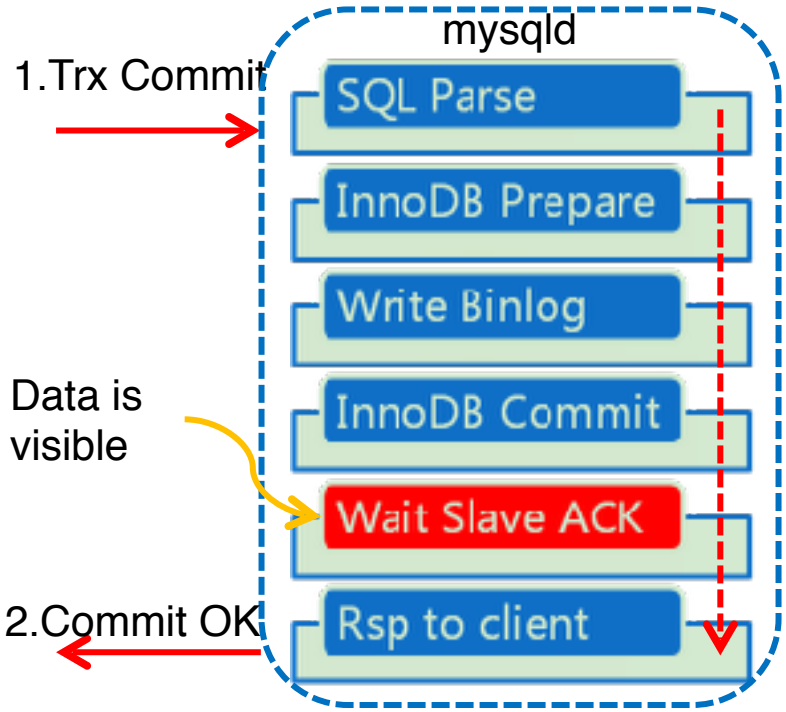

After commit

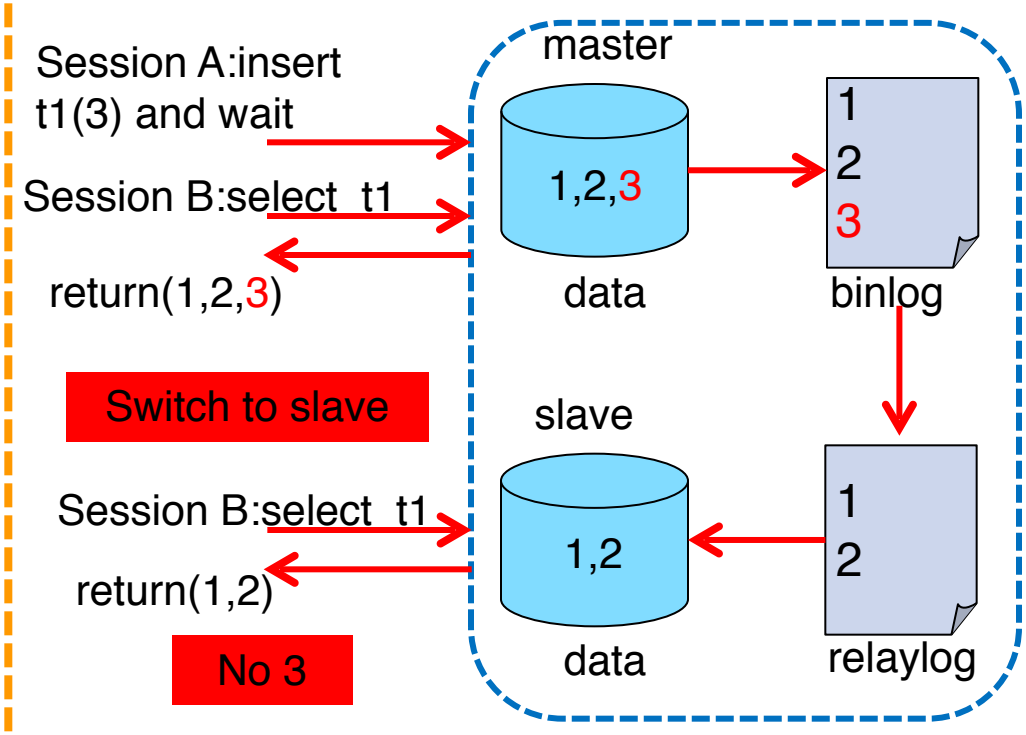

Phantom read

Phantom read

- Solution : Send to slave before commit

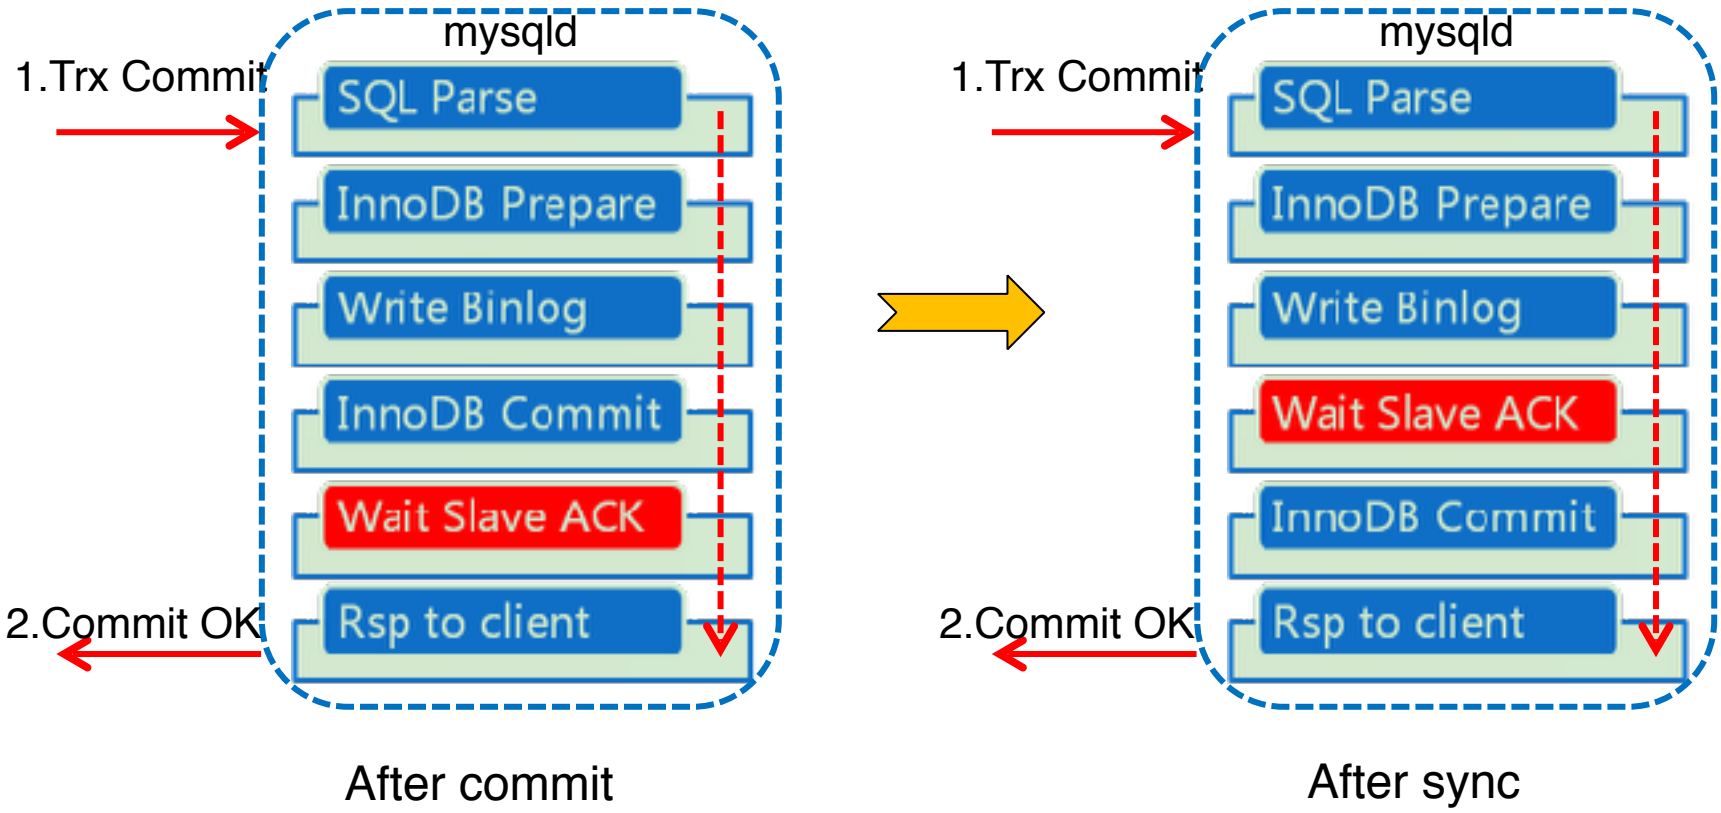

After commit

After sync

Flash back

- Problem:
  - master downgrade to slave , but has redundant transaction(s)

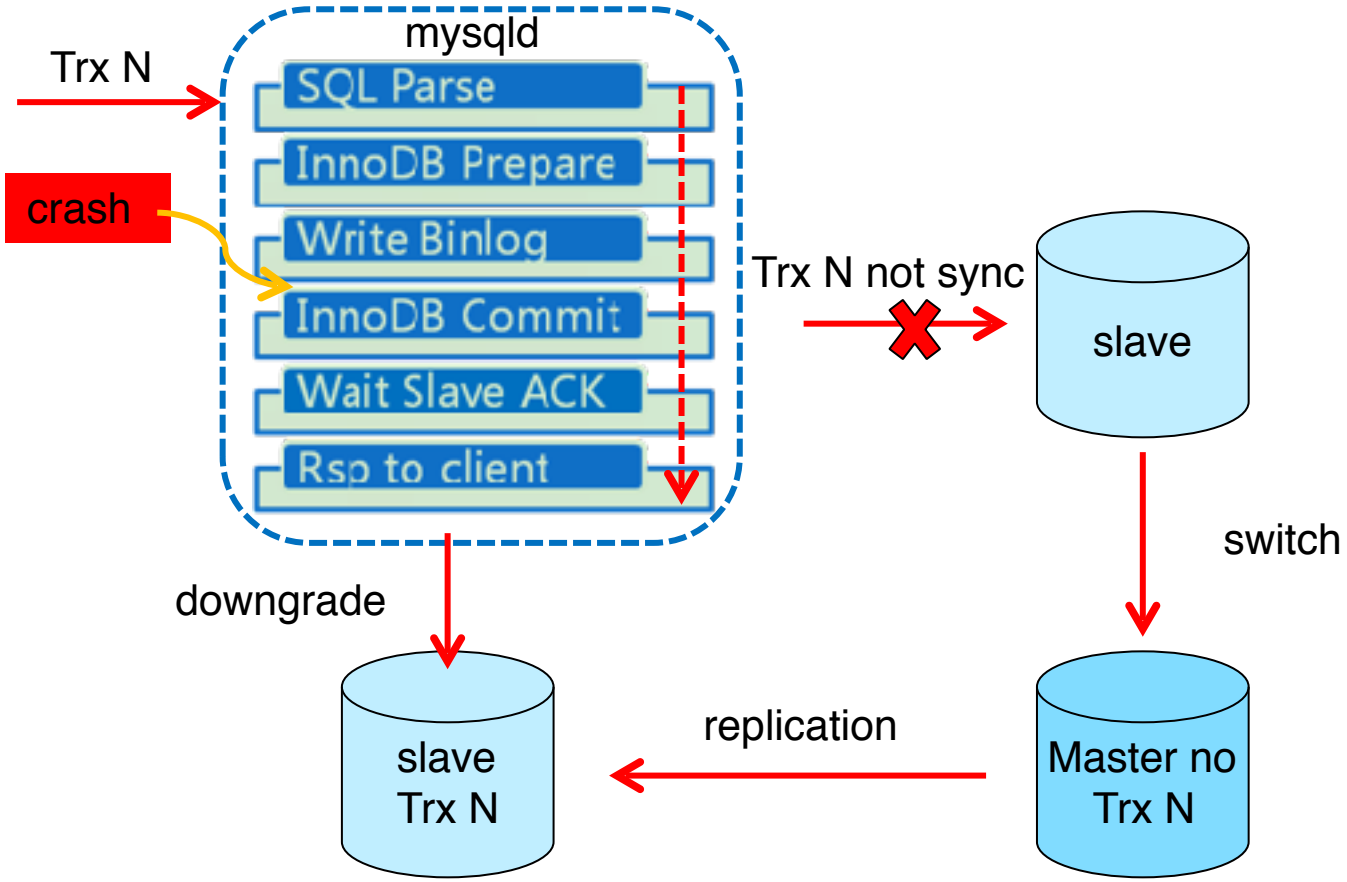

Performance

- **Key points**
  - time-consuming of Single transaction
  - Throughput(QPS)

- **Per transaction**
  - $T_{total}=T_{sql}+T_{engine}+T_{replica}$
  - $T_{replica}=T_{network}+T_{fsync}$
  - benchmark：
    - time-consuming of Full cache single transaction : 3.82ms
    - time-consuming of Semi-sync single transaction : 8.33ms including RTT 2.6ms
    - SAS disk order 512B write + fsync delay 0.13ms
    - $T_{fsync} = 8.33-3.82-2.6=1.9ms$

- **Throughput**
  - Throughput = concurrent number / time-consuming of single transaction

Optimized slave

- p Problem : IO thread delay
- ü Lock conflicts
- ü Random small IO
- ü Serialization

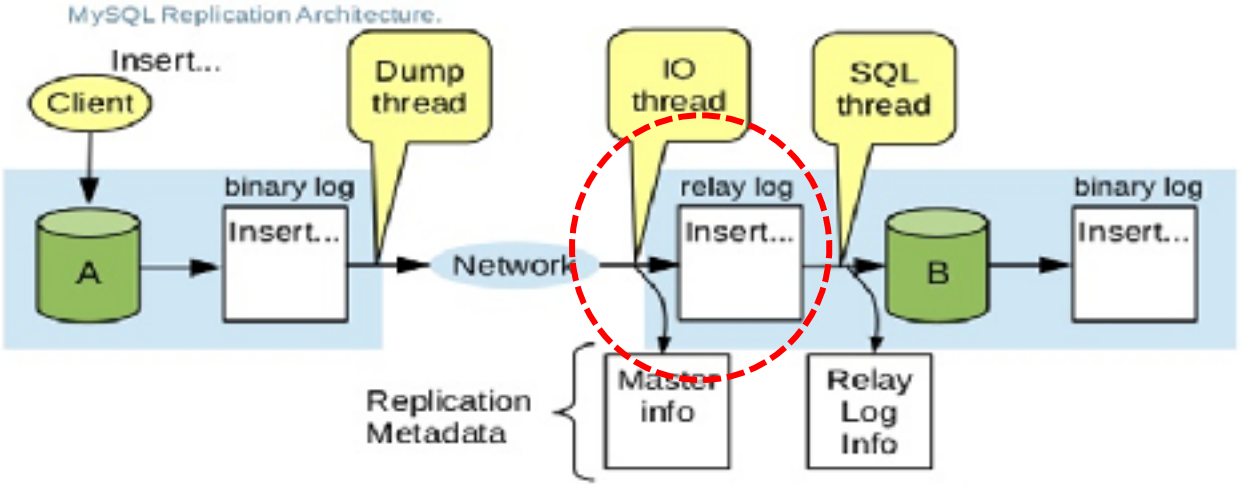

Version	Type	SGL TX(ms)	RTT	Benchmark
MySQL5.7	async	3.82	2.60	100%
MySQL5.7	semi-sync	8.33	2.60	46.30%

$T_{sql} + T_{engine} = 3.82ms$, $T_{sql} + T_{engine} + T_{replica} = 8.33ms$, $T_{replica} = T_{network} + T_{fsync}$

$T_{fsync} = T_{replica} - T_{network} = 8.33\text{-}3.82\text{-}2.60 = 1.91ms$

Optimized slave

- p Solutions : fast log transfer
- ü Less locks
- ü Merge IO
- ü multi-threaded async replay on salve

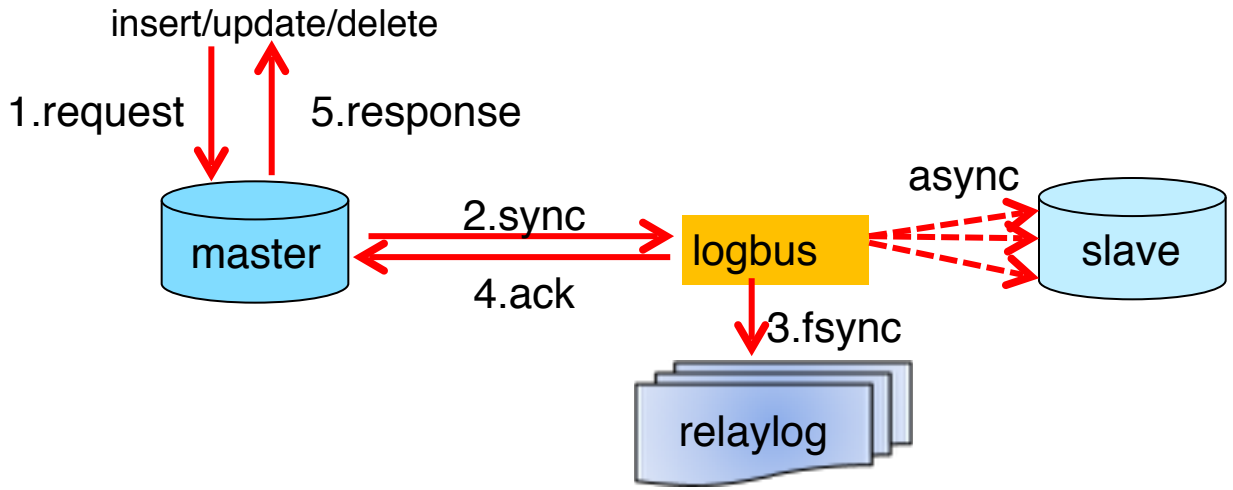

Version	Type	SGL TX(ms)	RTT	Benchmark
MySQL5.7	async	3.82	2.60	100%
MySQL5.7	semi-sync	8.32	2.60	46.30%
MySQL5.7	logbus	5.92	2.60	66.79%

optimized master

- Problem : Serialization
- ü $QPS_{(max)}$ = 1000 / RTT = 1000 / 3 = 330.

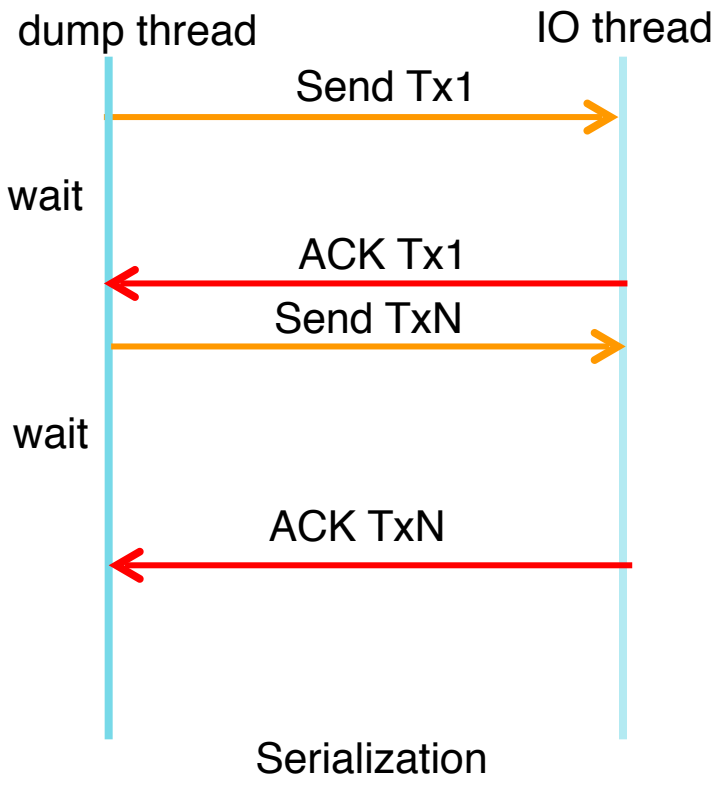

optimized master

- Solution: break into two threads
  - Dump thread
  - Ack thread

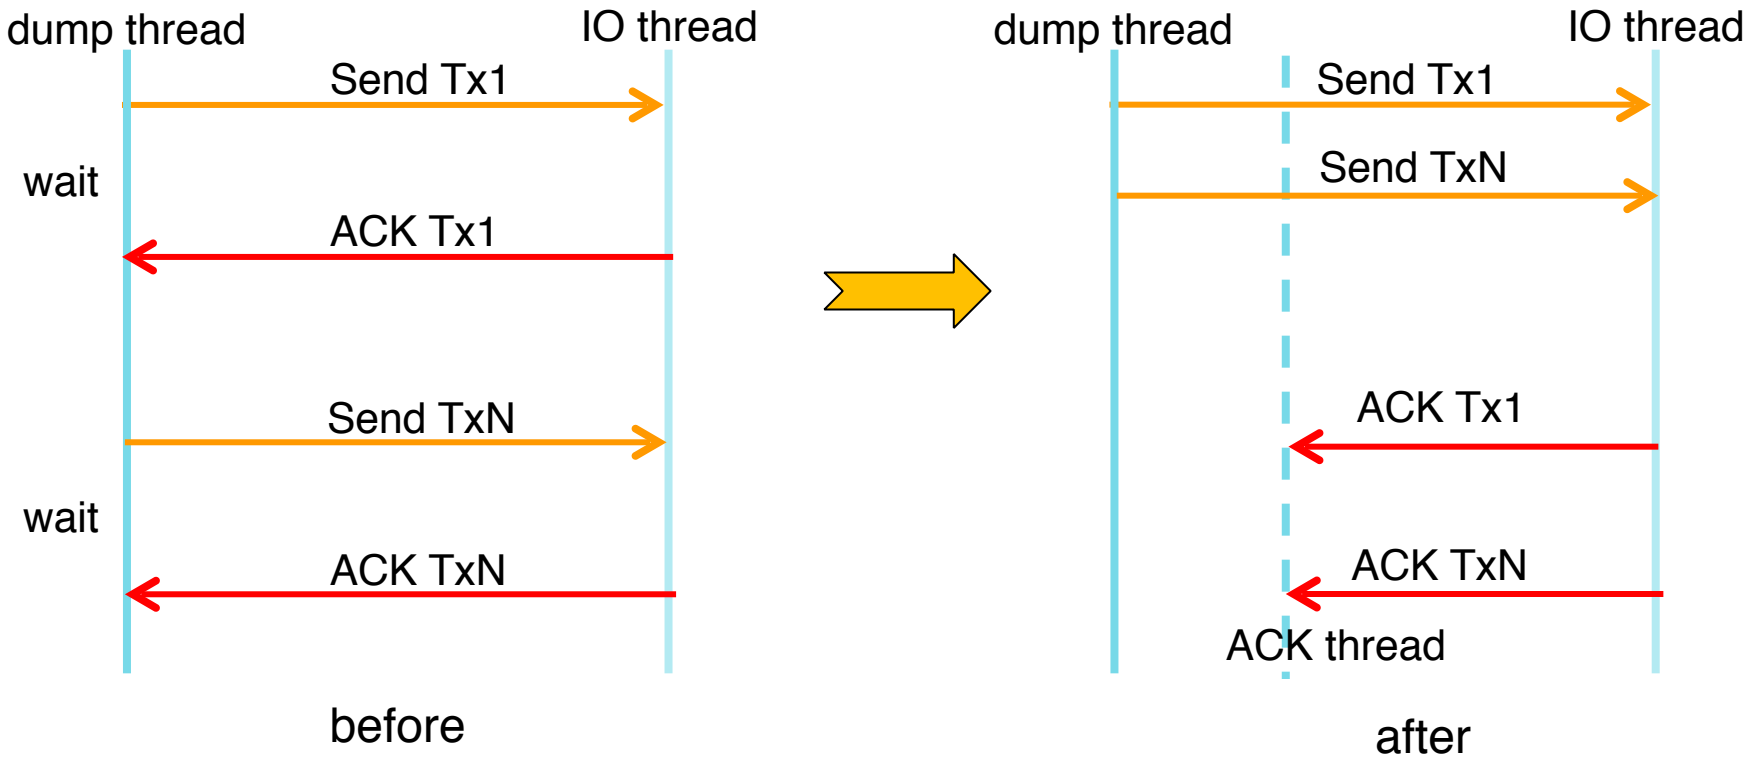

optimized master

- p Problem
- ü Cost of threads
- ü Thread blocking idle
- p Solution
- ü Thread pool(MariaDB)

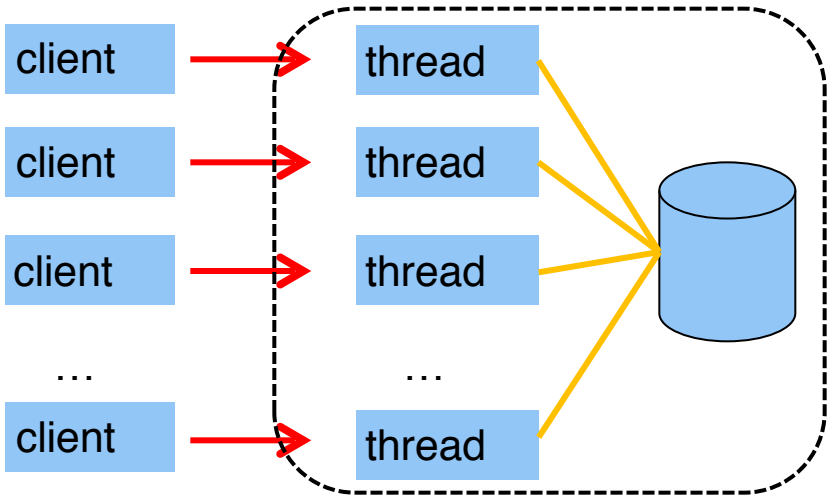

MySQL work model

```
while(1)
{
    receive sql;
    execute sql(wait response);
    rsp client;
}
```

Thread workflow

HA

- How-to less RTO?
  - Fault detection : zookeeper + lease
  - Failover : multi-thread replay
  - Ability to discover service
- RTO(aws 60-120s,ecdb <30s)

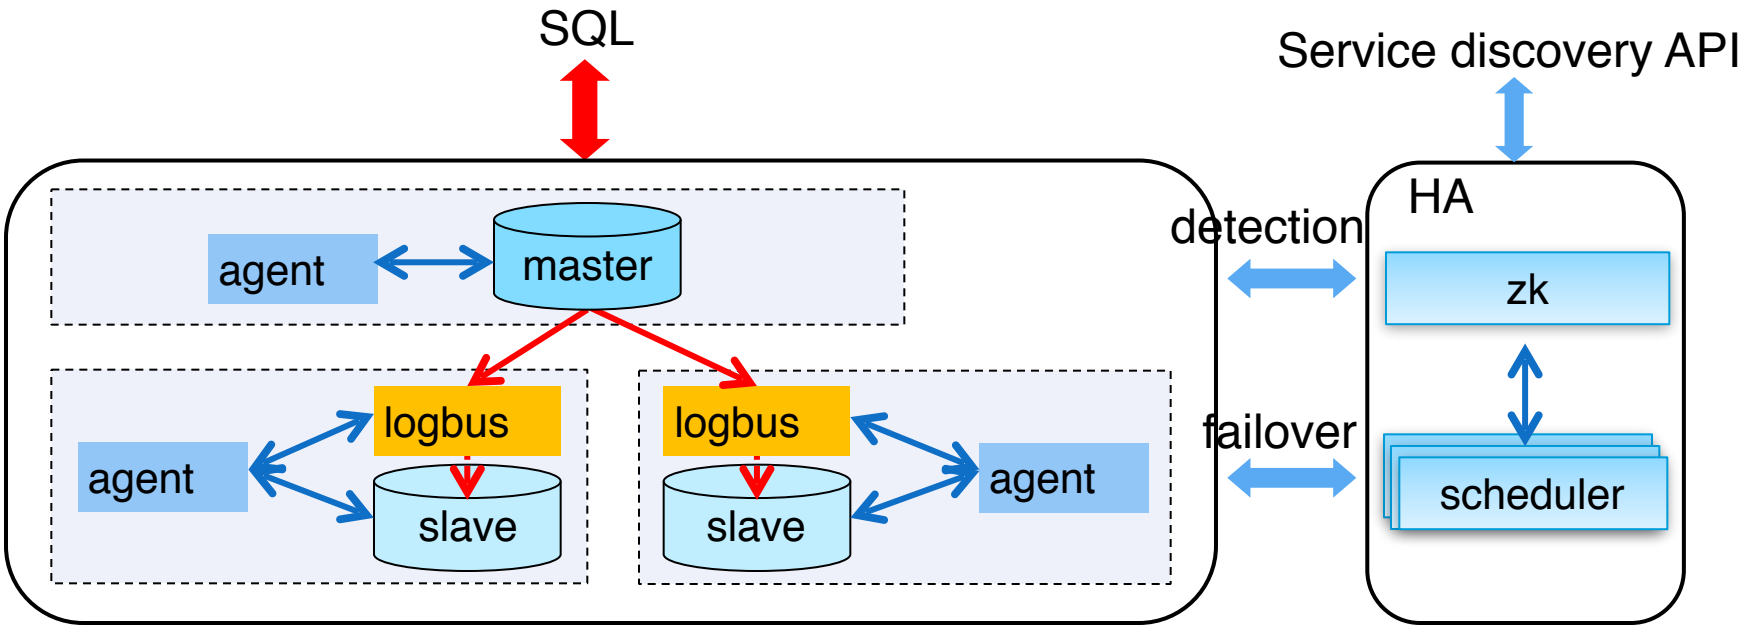

architecture

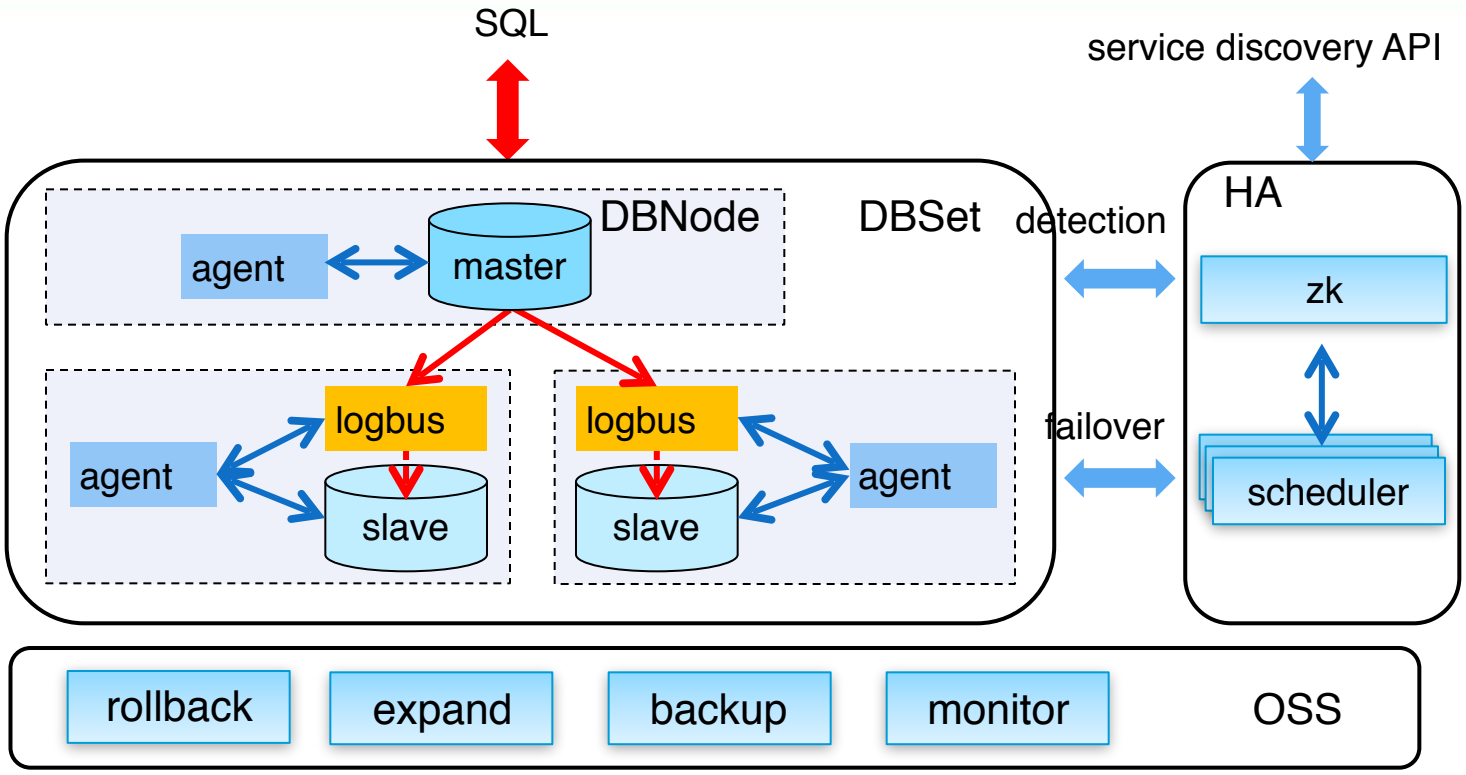

Logbus:fast log transfer

Agent : monitor mysqld/logbus

DBNode : logbus + agent + mysqld

HA： Fast fault detection & failover


Thank You

CONTACT US

- WeChat Group