SIGNAL 18
REPLICATION MANAGER

ShardProxy Replication-Manager

DataOps - Juin 2018
Kentoku SHIBA - Stephane VAROQUI

Replication Manager
MRM

Replication-Manager - Sharding

SIGNAL 18 REPLICATION MANAGER

About Sharding Frameworks

- They mostly don’t do efficient joins against shard group
- They mostly don’t preserve ACID properties
- They mostly relies on limited SQL parser
- They mostly force you to change application code

About Clustering Databases

- They mostly don’t scale writes
- They mostly duplicating datas before commit
- They mostly don’t allow long running transactions

About Scale Up

- xTBytes DB is hard to backup and transport
- Million tables issues on metadata
- Losing it all
- Different workload, requires different dataops (memory based, fragmentation, storage engine , hardware resources, maintenance policy)

Replication-Manager - Election on async replicas

SIGNAL 18 REPLICATION MANAGER

What is replication-manager ?

- State machine
- Event scheduler
- Monitoring framework (repl,status,var, schema)
- DB Job cron sender & receiver (backups,logs)
- HA failover
- **Multi topology**

  Master-slave, Multi-master, ring

  Multi-source, Gtid, Pseudo-gtid
- **Multi route**

  Haproxy, ProxySQL, Vitess, Maxscale, Consul

  Scripts, **Shardproxy**
- Multi clients (Rest API, HTTP , Cmd line)

What is Signal 18 ?

- Packages supported
- Non regression tests
- Continuous integration build

What is OpenSvc ?

- Services Orchestrator

  Dockers, KVM, Zone, LXC

  Service failover and placement
- Infrastructure Monitor
- Configuration Manager

Replication-Manager - Routing

ShardProxy

- MariaDB 10.3 and spider
- Table discovery on multiple shards clusters

Maxscale - ProxySQL

- Database protocol aware
- Database topology aware
- SQL parser for complex filtering
- Funneling, Multiplexing, Pooling
- Need scaleup depending on the route and filter complexity
- Pluggable for routes, parsers, protocols, monitors and filters
- BSL licence vs BSD

HaProxy

- Protocol agnostic
- Authentication agnostic using server that Support Proxy Protocol (MariaDB 10.3,Percona, Amazon)
- Long time open source: well known, tested, documented
- Proven minimal resources usage at layer 7

Consul DNS

- More and more used for mico services

Scripts or others ….

Infrastructure - Provision Shard Groups

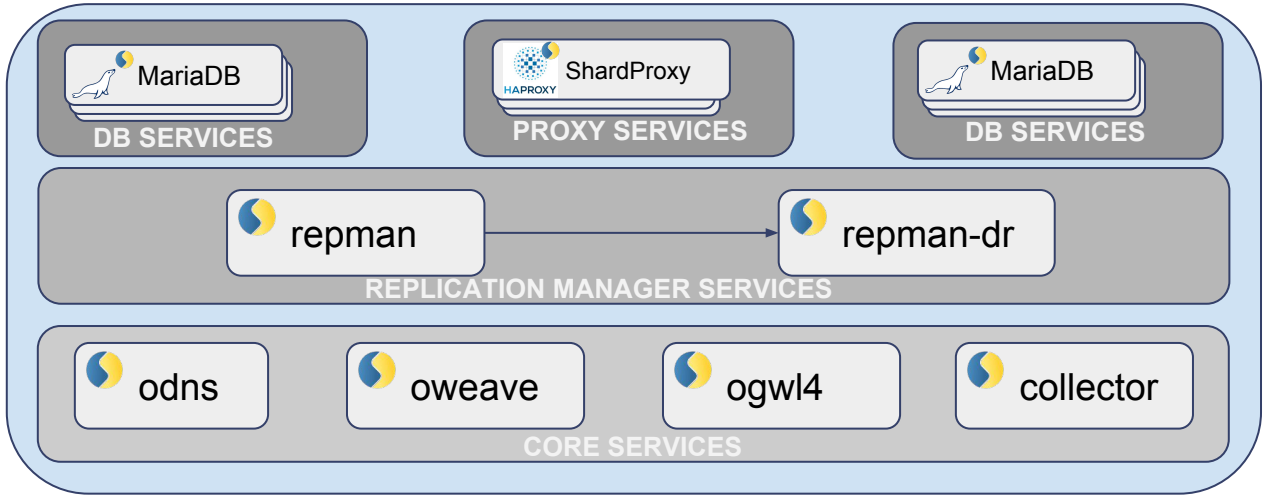

Replication-Manager - Sharding

Multi Cluster Schema Detection -VSchema

- For each cluster a monitoring schema generic framework is enabled

```
monitoring-schema-change =true
monitoring-schema-change-script=""
```

- When schema or table is discovered VSchema framework is lookink up for duplicates in all clusters configured with same ShardProxy

- Unique tables are push down to ShardProxy with a spider table federated to the master node table of the cluster

- Duplicate tables are are push down to ShardProxy creating a spider table where each partition point to a duplicate cluster master node, a hash of primary key is default partitioning function

Replication-Manager - Election on async replicas

SIGNAL 18 REPLICATION MANAGER

Failover Monitoring Workflow

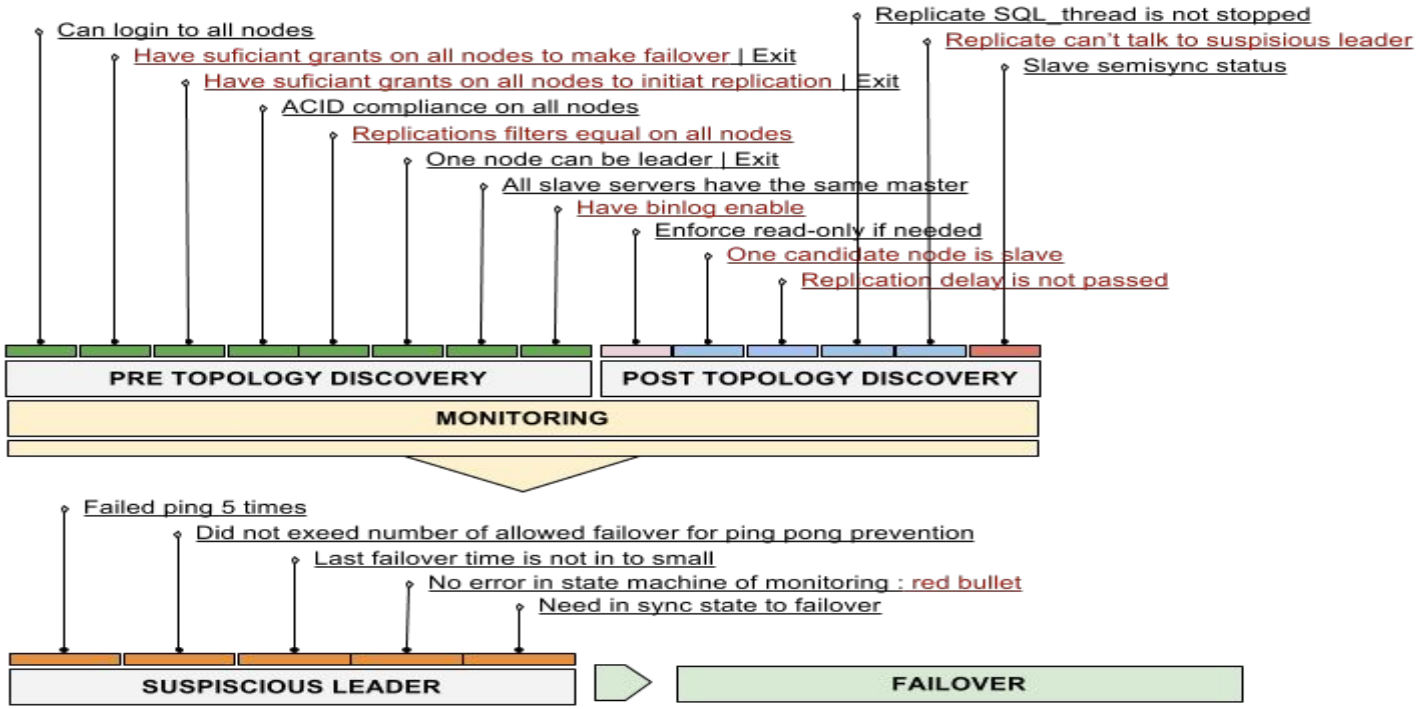

Replication-Manager - Election on async replicas

SIGNAL 18 REPLICATION MANAGER

Switchover Workflow

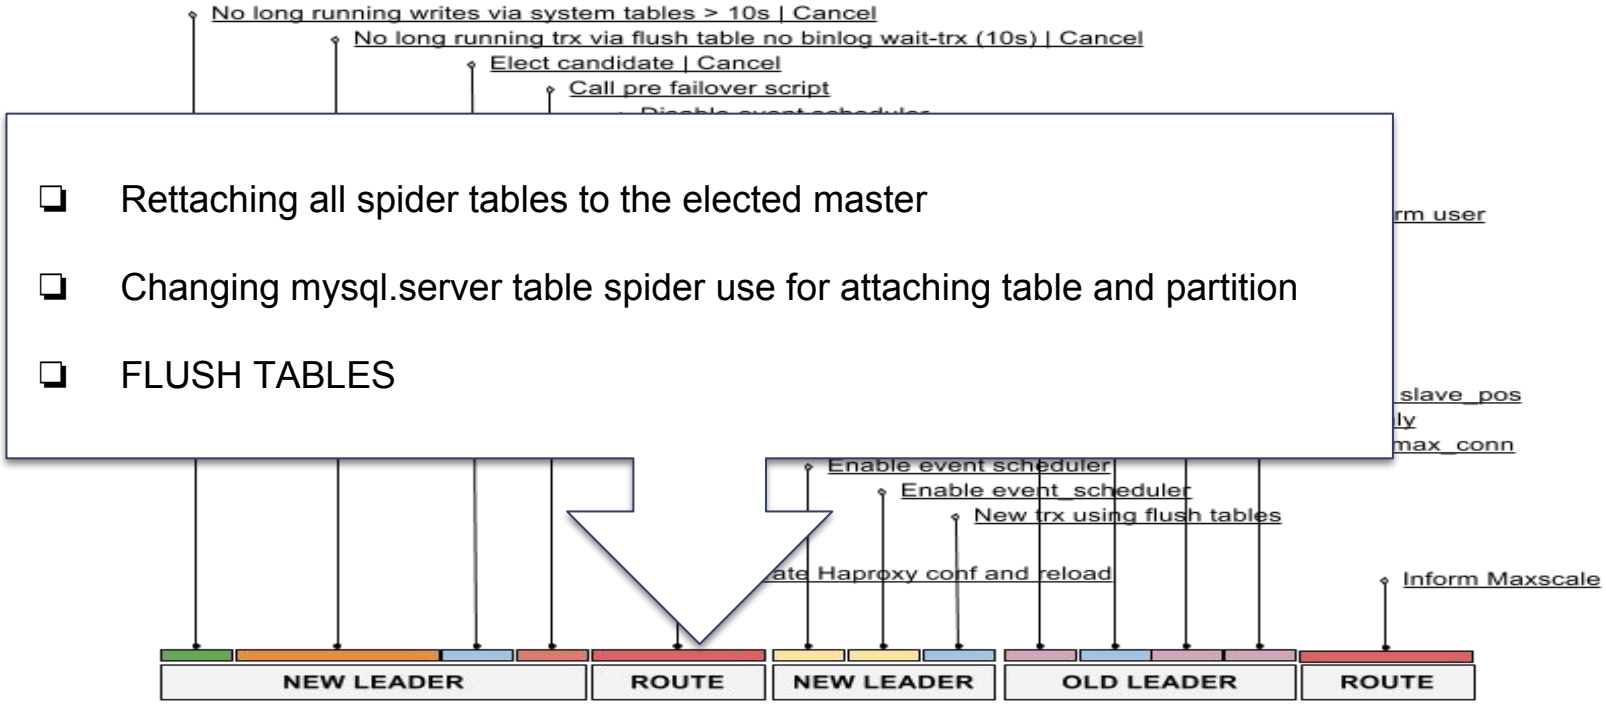

- Rettaching all spider tables to the elected master
- Changing mysql.server table spider use for attaching table and partition
- FLUSH TABLES

Replication-Manager - Sharding

SIGNAL 18 REPLICATION MANAGER

Multi Cluster Schema Detection -VSchema

- ./replication-manager-cli api --url="https://127.0.0.1:10005/api/clusters/cluster_mdbshardproxy_shard1/schema"


```
"world.City": {
        "tableSchema": "world",
        "tableName": "City",
        "engine": "MyISAM",
        "tableRows": 4079,
        "dataLength": 273293,
        "indexLength": 43008,
        "tableCrc": 17295449481889899010,
        "tableClusters": "cluster_mdbshardproxy_shard2"
    },
```

Replication-Manager - Sharding

SIGNAL 18 REPLICATION MANAGER

Multi Cluster Schema resharding

- ./replication-manager-cli api --url="https://127.0.0.1:10005/api/clusters/cluster_mdbshardproxy_shard2/schema/world/City/actions/reshard-table"

Replication-Manager - Settings

Minimum settings

```
[Cluster_Mdbshardproxy_Shard1]
title = "Shard1"
db-servers-hosts = "127.0.0.1:3331,127.0.0.1:3332"
db-servers-prefered-master = "127.0.0.1:3331"
db-servers-credential = "root:mariadb"
db-servers-connect-timeout = 1
replication-credential = "root:mariadb"
[Cluster_Mdbshardproxy_Shard2]
title = "Shard2"
db-servers-hosts = "127.0.0.1:3333,127.0.0.1:3334"
db-servers-prefered-master = "127.0.0.1:3333"
db-servers-credential = "root:mariadb"
db-servers-connect-timeout = 1
```

```
[Default]
shardproxy = true
shardproxy-servers = "127.0.0.1:3336"
shardproxy-user = "root:mariadb"
```

Replication-Manager - Ressources

SIGNAL 18 REPLICATION MANAGER

Downloads

- https://www.signal18.io

Source

- https://github.com/signal18/replication-manager

Roadmap & Documentation

- https://docs.signal18.io/

Talk to us

- https://gitter.im/mariadb-corporation/replication-manager

Q&A