

Hello

MongoDB
Sharding

Alexander
'm0use'
Kolesen

MongoDB Sharding

MongoDB

MongoDB Sharding

Alexander
'm0use'
Kolesen

- schemaless, JSON-style
- written on C++
- mmap()-based storage

Sharding

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- horizontal partitioning
- multiple nodes

How MongoDB does sharding?

MongoDB
Sharding

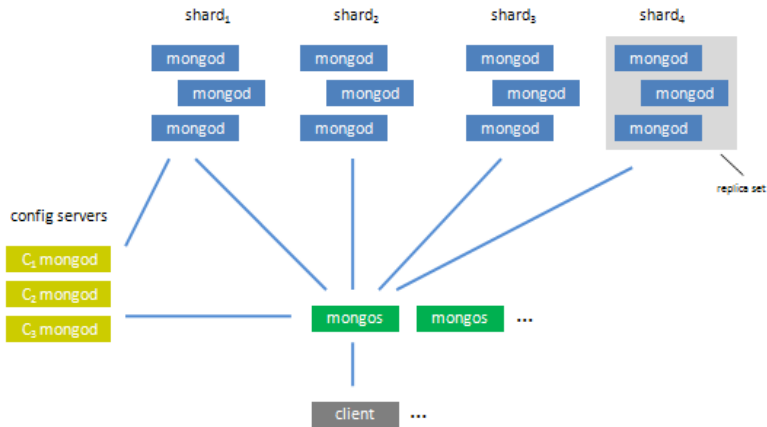
Alexander
'm0use'
Kolesen

- 'mongos' routers
- 'config' servers, exactly 3
- synchronous replication
- ReplicaSets (RS) as shards

How MongoDB does sharding?

MongoDB Sharding

Alexander 'm0use' Kolesen




ReplicaSet (RS)

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- async replication
- 3 nodes min
- 1 primary, N secondaries
- automatic failover (server, client)
- 1p-1s - need an arbiter



Balancing

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- chunks, 64M, changeable
- minimize amount of data transferred
- Nmax-Nmin 9
- manual moving, disable balancer

Choosing Shard Key

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- do it once, do it right
- monotonically increasing - :(
- so do not shard by `_id`
- random values (hash of ..) - :(

Add/Remove Shards

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- `db.runCommand(addshard: ...)`
- `db.runCommand(removeshard: ...)`

Example

MongoDB
Sharding

Alexander
'm0use'
Kolesen

```
1 | mongo> db.printShardingStatus()
2 | --- Sharding Status ---
3 |   sharding version: { "_id" : 1, "version" : 3 }
4 |   shards:
5 |     { "_id" : "mqdb-rs0", "host" : "mqdb-rs0/ip-10-128-71-123.ec2.internal:27017,ip-10-222-13-219-27017,ip-10-82-214-22.ec2.internal:27017" }
6 |     { "_id" : "mqdb-rs2", "host" : "mqdb-rs2/ip-10-121-12-92.ec2.internal:27017,ip-10-222-13-216-27017,ip-10-83-65-166.ec2.internal:27017" }
7 |     { "_id" : "mqdb-rs3", "host" : "mqdb-rs3/ip-10-177-74-253:27017,ip-10-283-76-177.ec2.internal:27017,ip-10-84-134-222.ec2.internal:27017" }
8 |   databases:
9 |     { "_id" : "admin", "partitioned" : false,
10 |       "primary" : "config" }
11 |     { "_id" : "mqdb", "partitioned" : true, "primary" : "mqdb-rs3" }
12 |   mqdb.messages chunks:
13 |   mqdb-rs3 3
14 |   mqdb-rs0 2
15 |   mqdb-rs2 2
16 | { "queue_id" : { $minKey : 1 } } --> { "queue_id" : ObjectId("4fde0e69841498f97b7e8e09" ) on : mqdb-rs3 { "t" : 1000, "i" : 1 }
17 | { "queue_id" : ObjectId("4fde0e69841498f97b7e8e09" ) } --> { "queue_id" : ObjectId("4fde0e69841498f97b7e8e0d" ) on : mqdb-rs0 { "t" : 1000, "i" : 9 }
18 | { "queue_id" : ObjectId("4fde0e69841498f97b7e8e0d" ) } --> { "queue_id" : ObjectId("4fde0e69841498f97b7e8e12" ) on : mqdb-rs2 { "t" : 1000, "i" : 10 }
19 | { "queue_id" : ObjectId("4fde0e69841498f97b7e8e12" ) } --> { "queue_id" : ObjectId("4fde39d0841498f97b7e8e34" ) on : mqdb-rs2 { "t" : 1000, "i" : 11 }
20 | { "queue_id" : ObjectId("4fde39d0841498f97b7e8e34" ) } --> { "queue_id" : ObjectId("4fde39d2841498f97b7e8e4e" ) on : mqdb-rs0 { "t" : 1000, "i" : 12 }
21 | { "queue_id" : ObjectId("4fde39d2841498f97b7e8e4e" ) } --> { "queue_id" : ObjectId("4fde39e3841498f97b7e8e82" ) on : mqdb-rs3 { "t" : 1000, "i" : 8 }
22 | { "queue_id" : ObjectId("4fde39e3841498f97b7e8e82" ) } --> { "queue_id" : { $maxKey : 1 } } on : mqdb-rs3 { "t" : 1000, "i" : 6 }
```

Limitations/Tips

MongoDB
Sharding

Alexander
'm0use'
Kolesen

- changing shard key only manually
- findAndModify - must include shardkey
- findAndModify - doesn't support '\$in'
- 'mongos' is lightweight, mostly CPU-bound
- '-smallfiles' for 'config'

The End

MongoDB
Sharding

Alexander
'm0use'
Kolesen

Q?