

Three pieces
mongod, mongos and client driver library

Mongod is the process

Mongos is a router, it routes writes to correct mongod instance

Shares writing

Autosharding increases
Writes (chunks – distributed reads / writes),
helps with scaling

Replication sets are for high
availability and read scaling
not write scaling

Each shard / partition has
It's own replication set

Stores sharding configuration

Infrequently stores small amounts
of data

Different workloads require
different storage strategies

Exposed by a storage engine
API

Provides more flexible to your
Deployments

WT

MMAP V1

Encrypted

RAID 10 i/o

MongoDB Architecture

Community supported
Hadoop integration

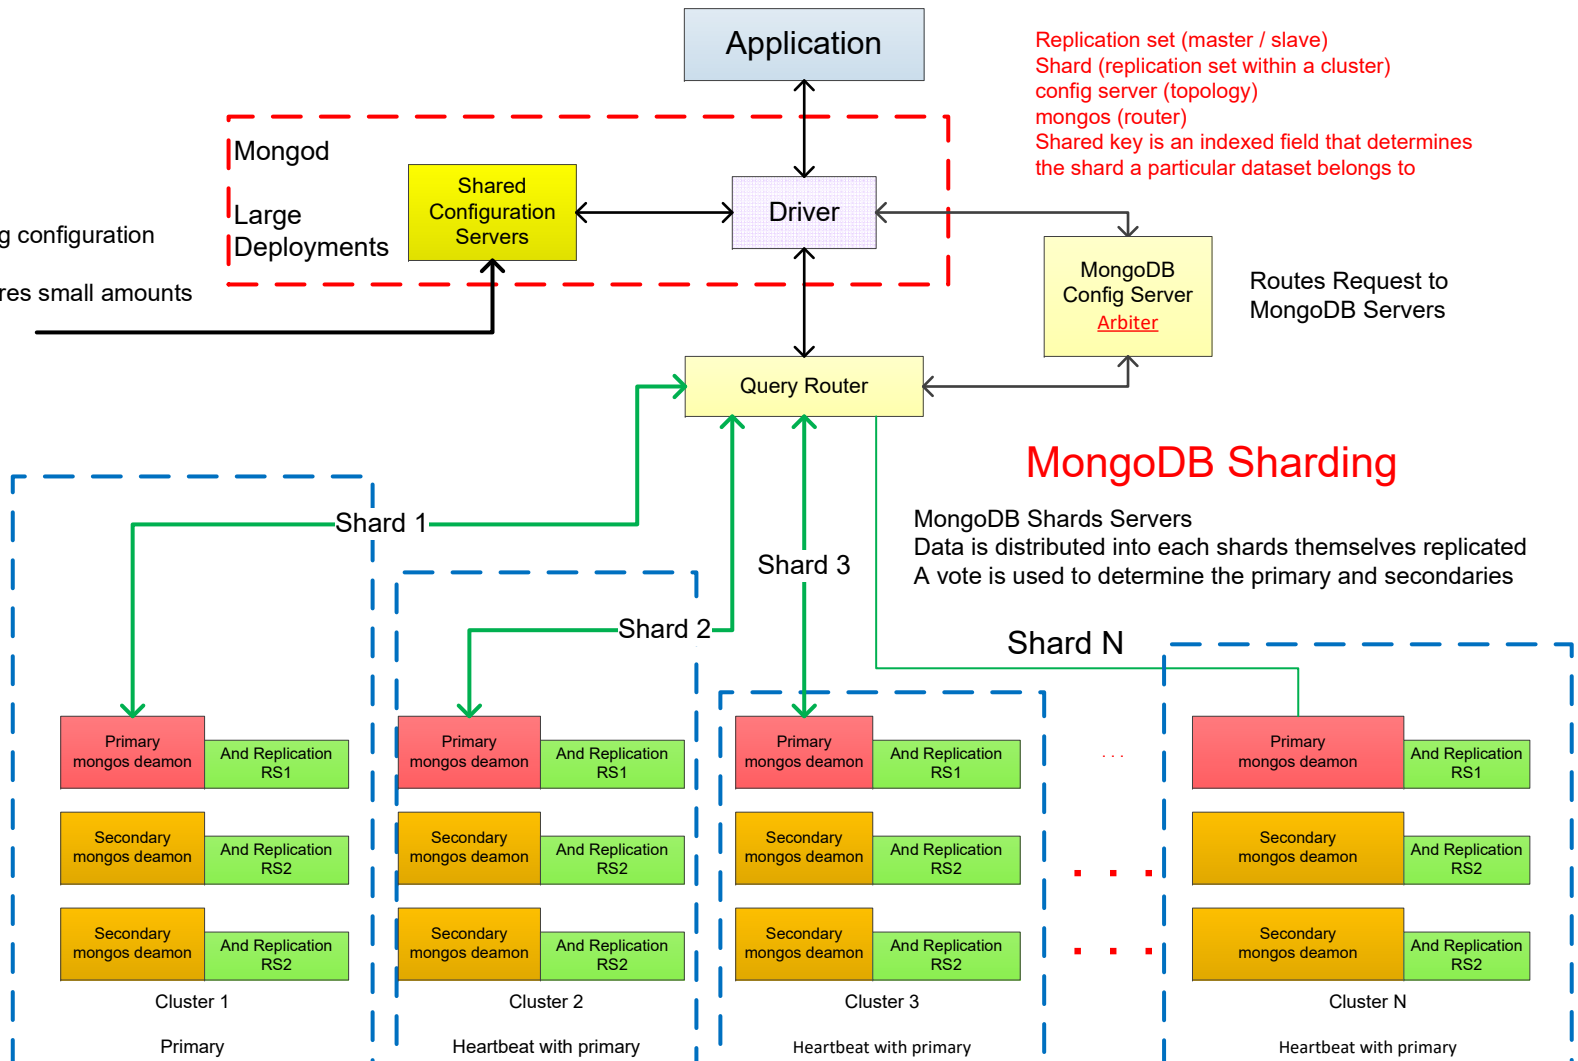
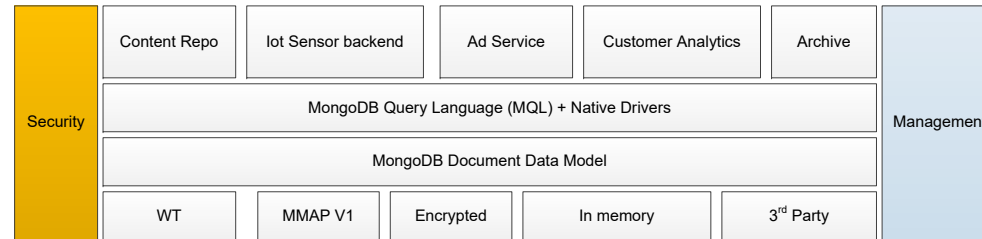
Aggregation framework & native MapReduce

Rick secondary indexes, including geospatial and TTL indexes

Built in replication for high availability

Auto-sharding for horizontal scalability

JSON data model with dynamic schemas



MongoDB Sharding

MongoDB Shards Servers
Data is distributed into each shards themselves replicated
A vote is used to determine the primary and secondaries

Routes Request to
MongoDB Servers