




```
{name: "mongo", type: "DB"}
```

```
db.NLJUG.insert(  
{ "presenters" : "Maikel Alderhout & Bas van Oudenaarde",  
  "type" : "JFALL",  
  "datum" : ISODate(31/10/2012) } )
```



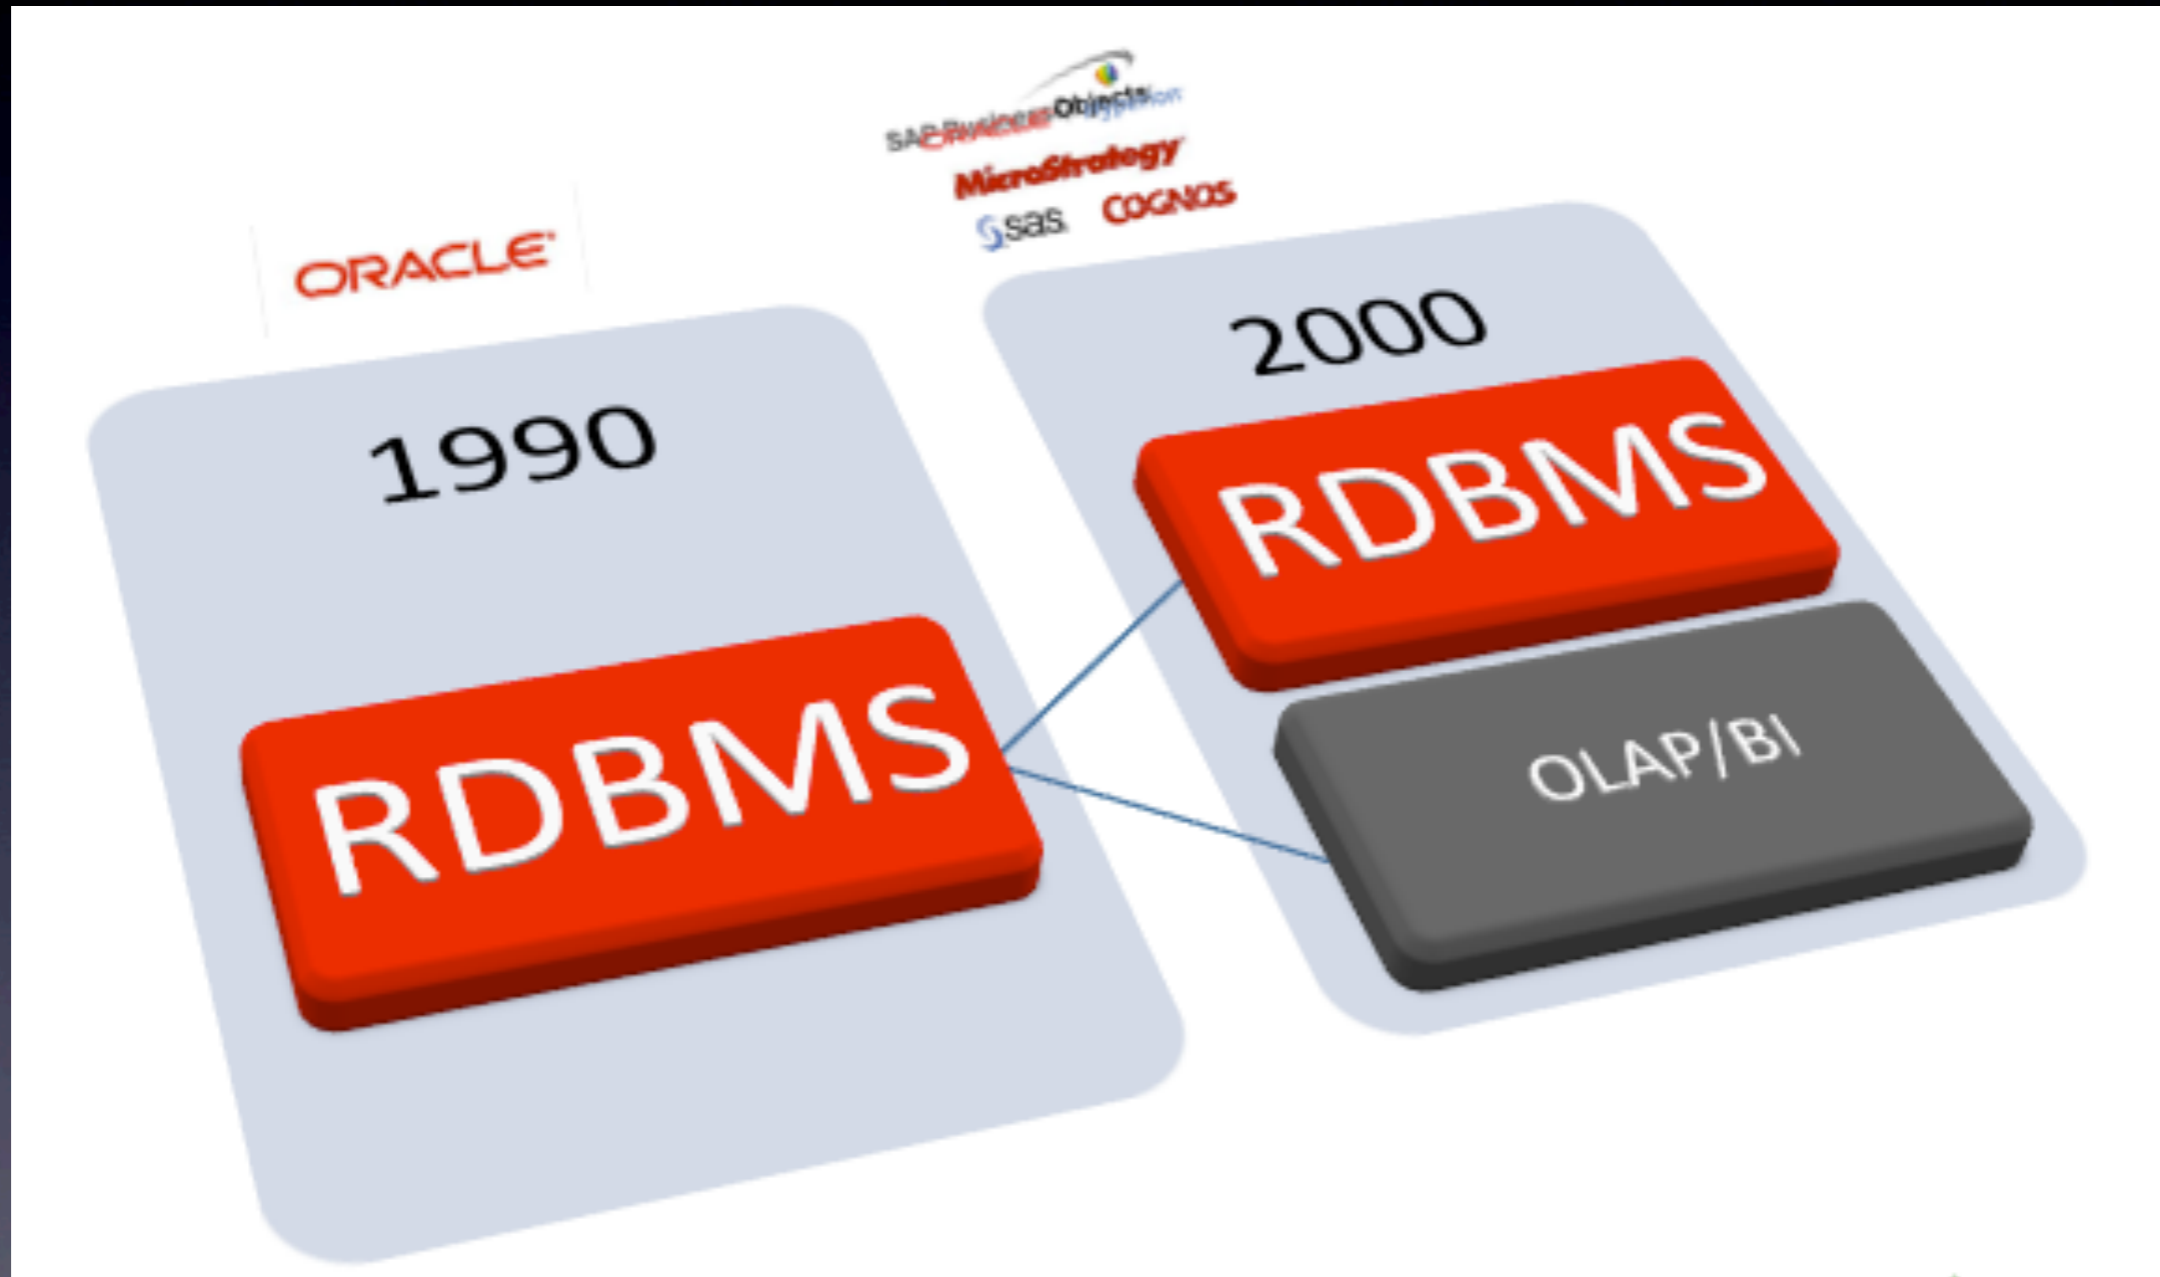
# NoSQL Hype

- **Not only** ... of **No SQL**  
‘**not only**’ relationele data opslag

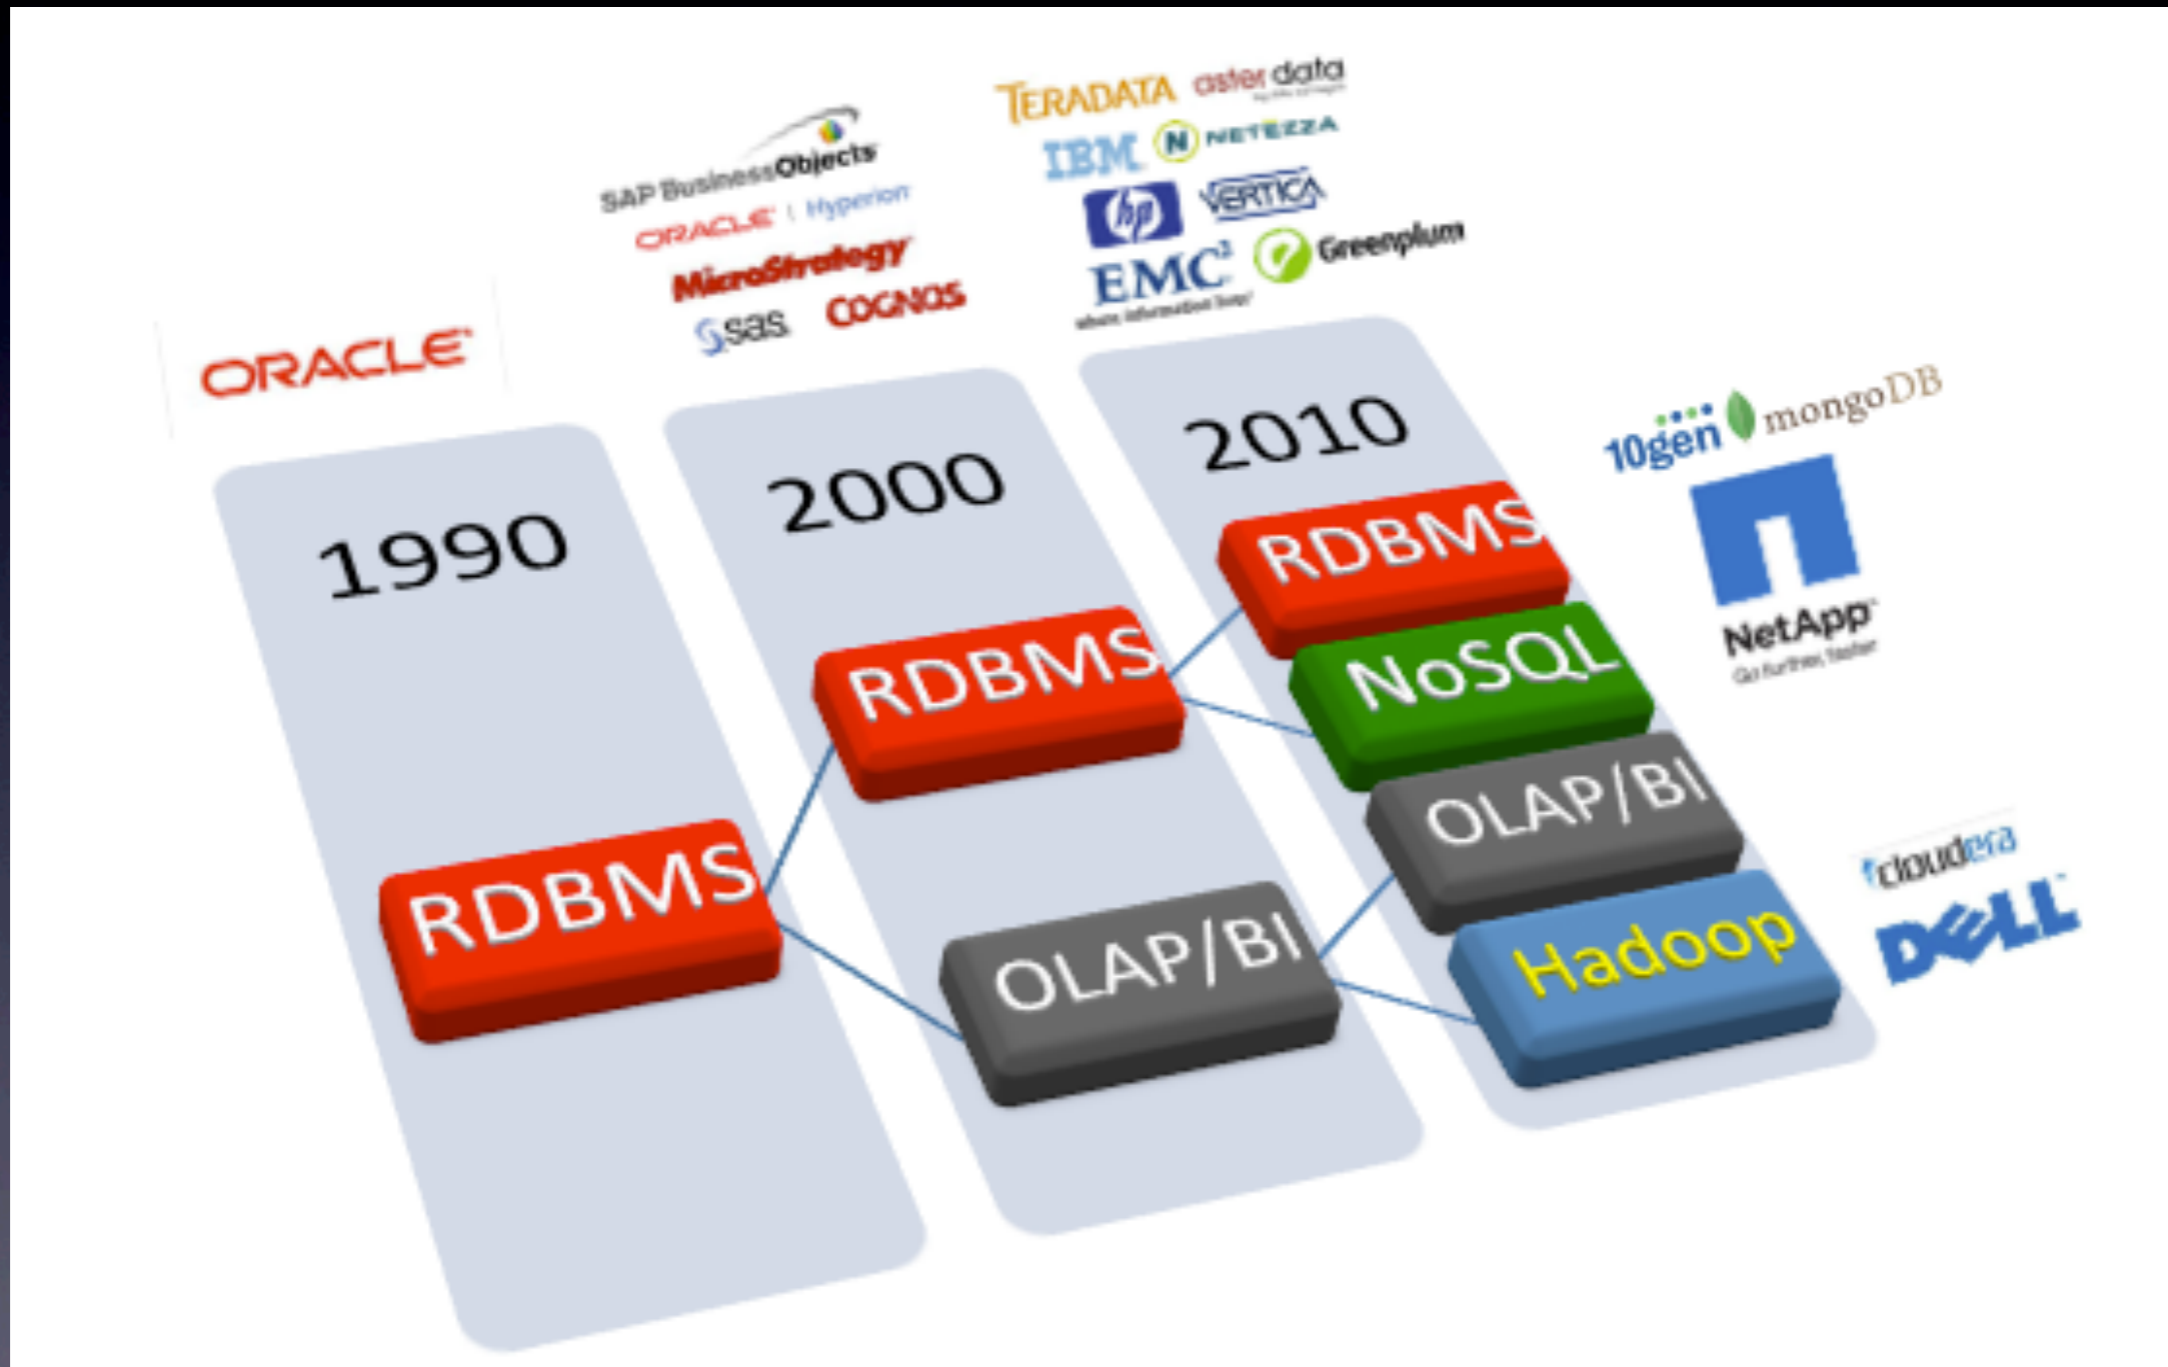
# Database evolution #1



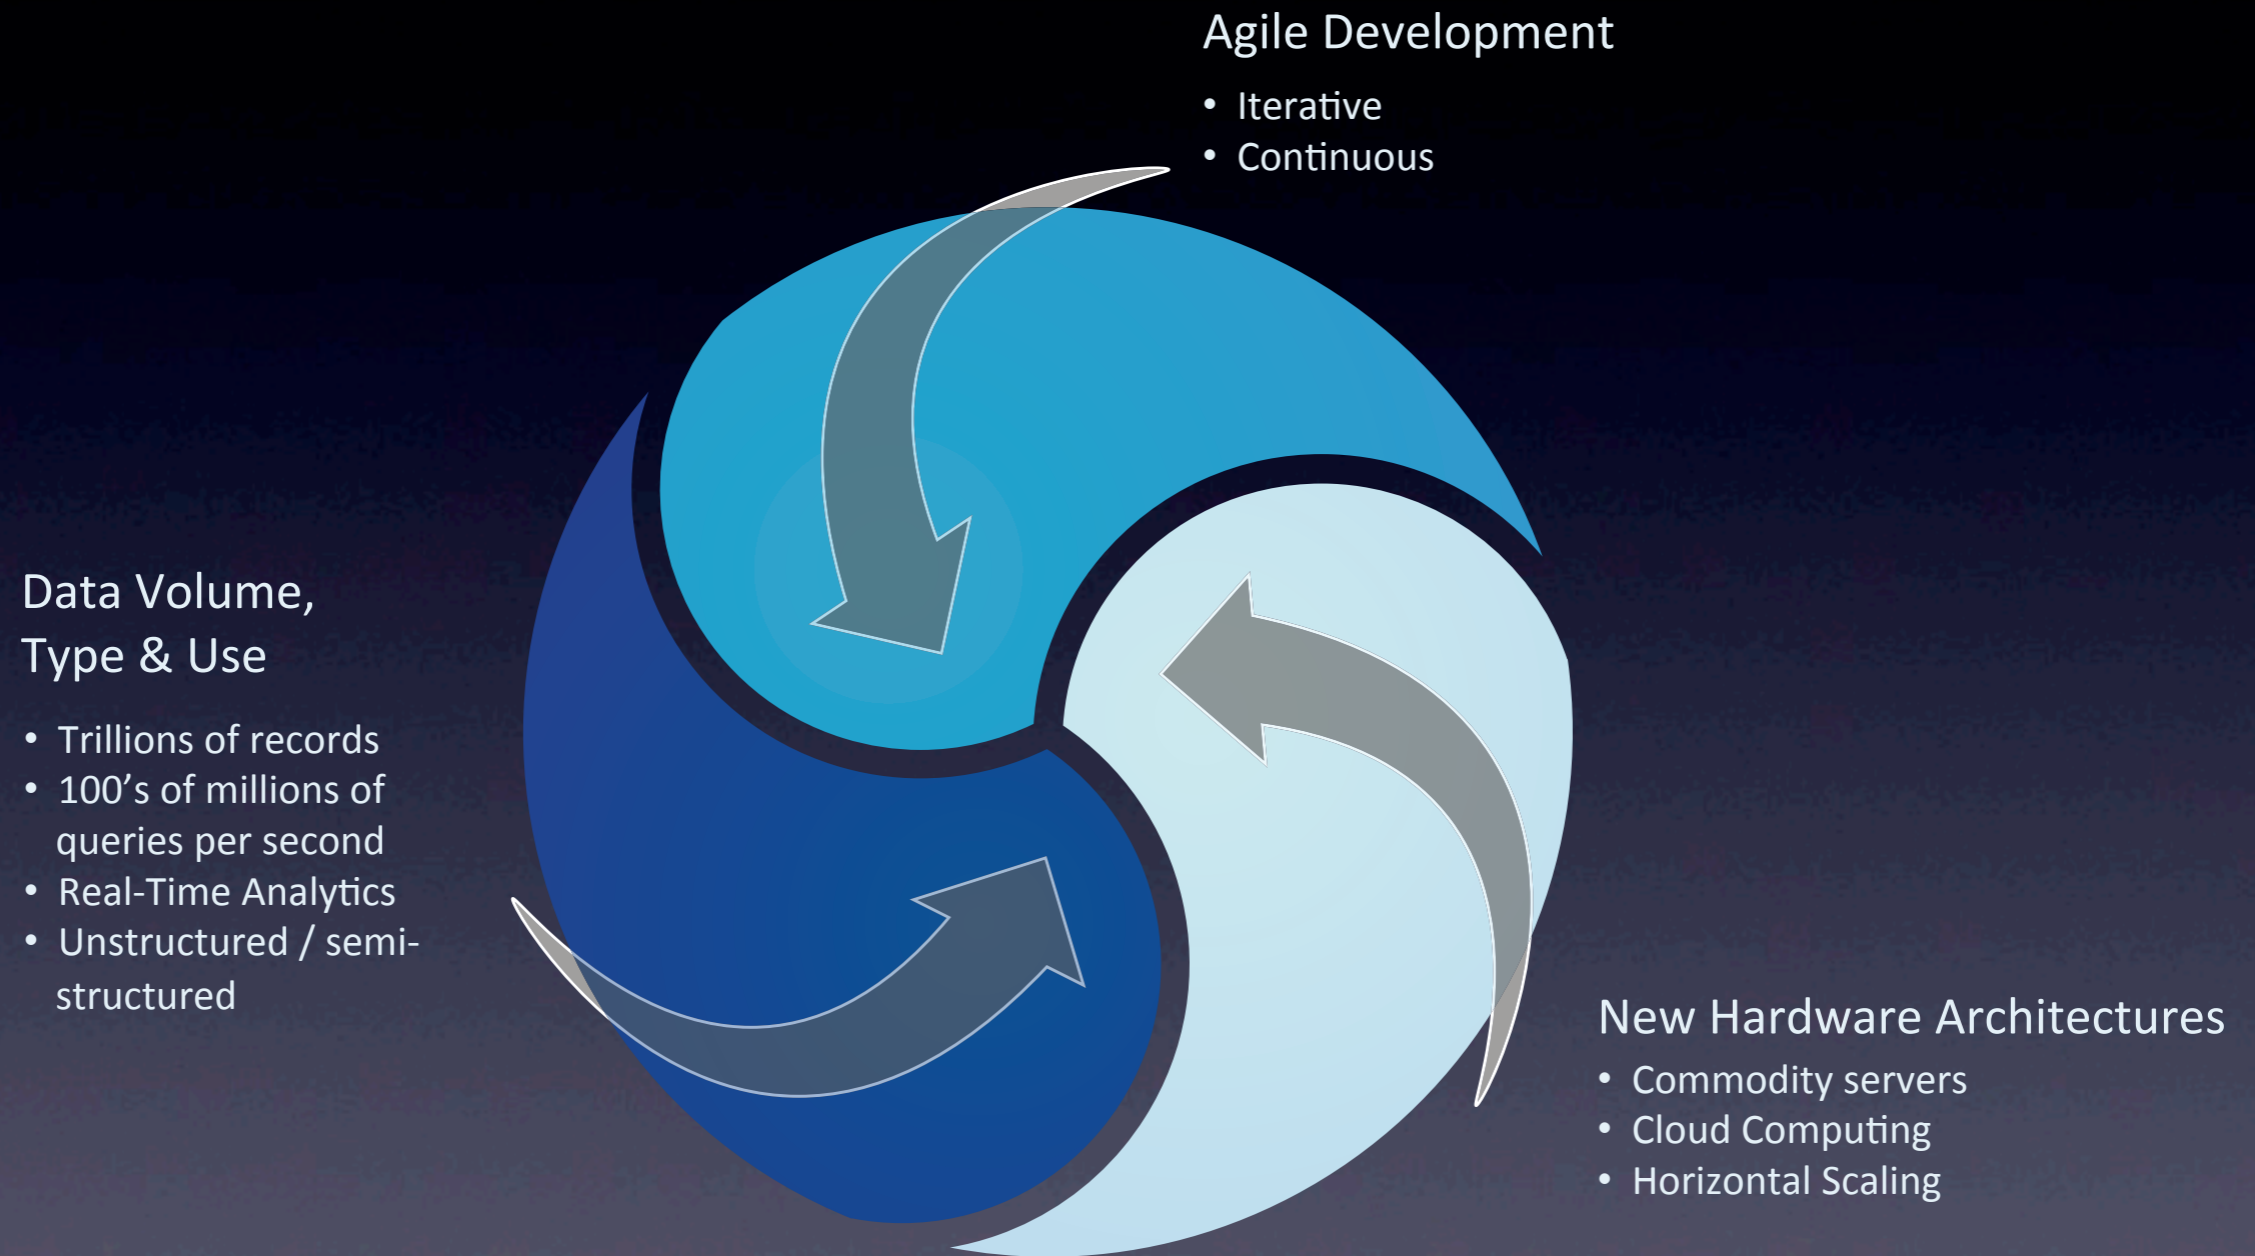
# Database evolution #2



# Database evolution #3



# De huidige trends hebben een behoorlijke impact in het traditionele database landschap



# Hoeho Java problemen?

- Schaalbaarheid 3-tier architecturen
- ORM, mapping vanuit domain objecten.  
Lekker alles aan elkaar knopen
- Caching, niet te veel naar de database
- Alles in memory, dan maar?

# NoSQL oplossingen

## Key-Value

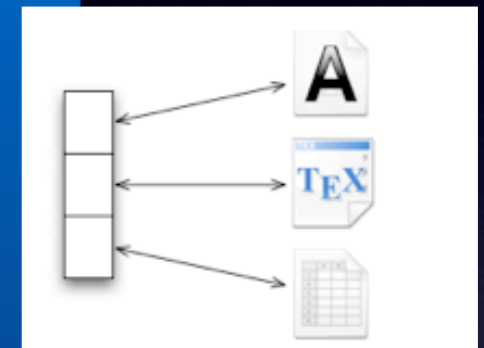


## Key-value store

bv. Voldemort, Dynamite, Tokyo


## Document

bv CouchDB, **MongoDB**, Riak



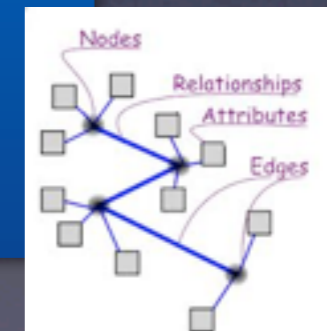
## ColumnFamily / BigTable clones

bv Google BT, HBase, Hypertable, Cassandra



## Graph databases

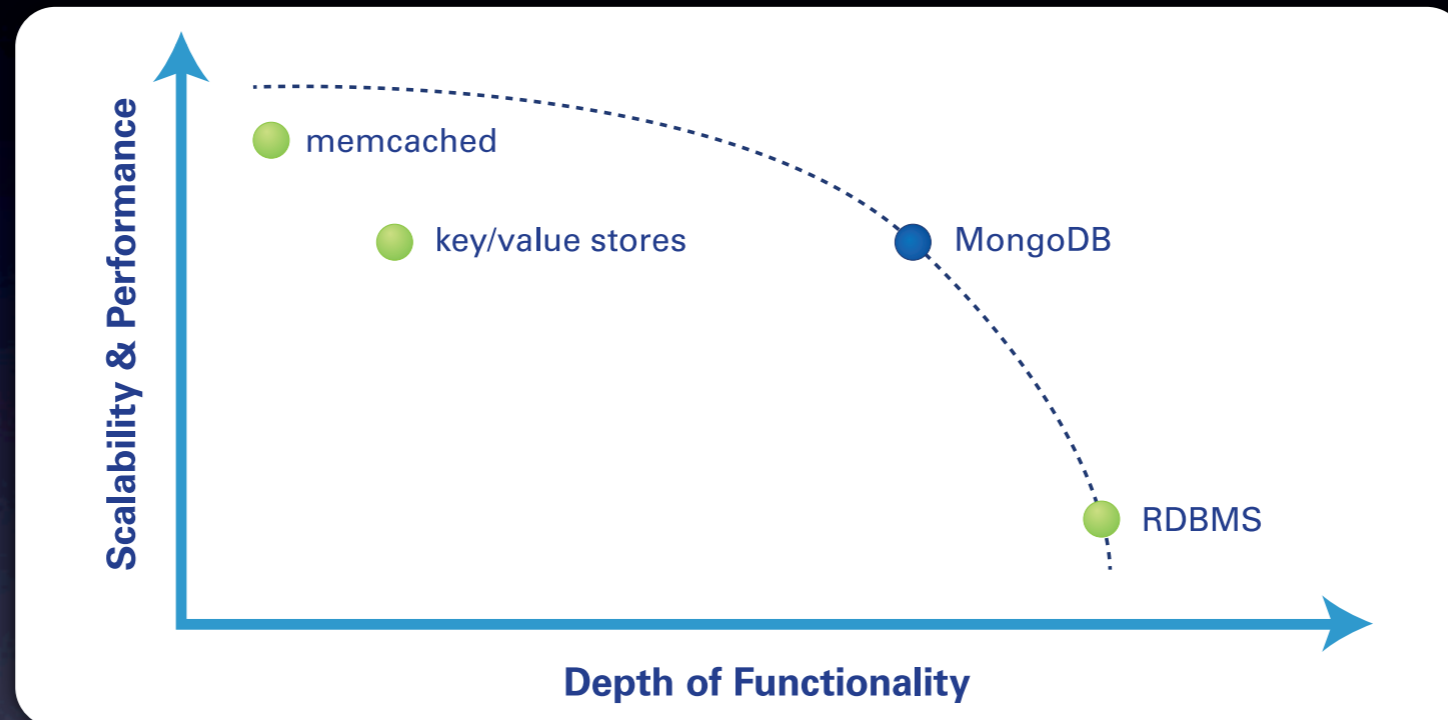
Neo4J, Sones, AllegroGraph







# MongoDB is a scalable, high-performance NoSQL database.



- Open source, written in C++
- Document-oriented Storage
  - Based on JSON Documents
  - Schema-less
- Full featured indexes, query language
- Replication & High Availability
- Auto-sharding

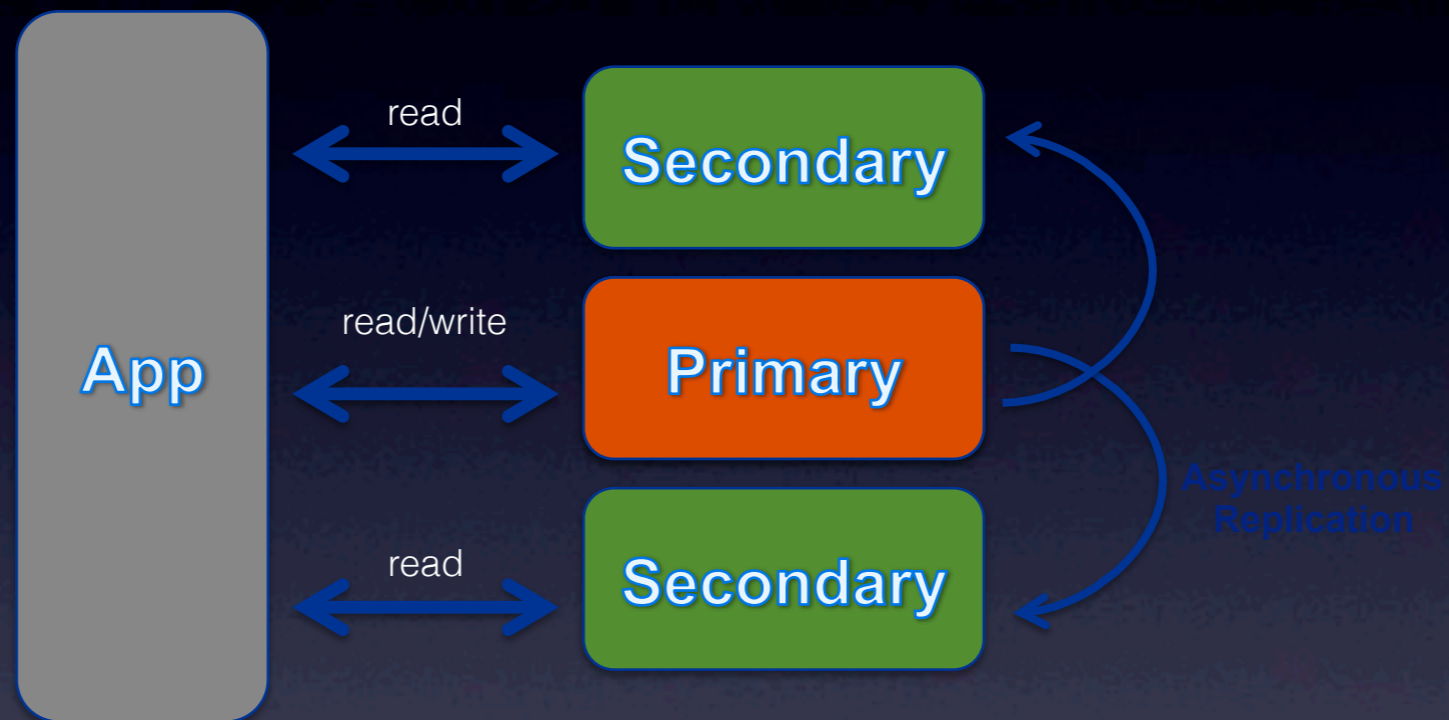
# JSON Document

## Person

```
{  
  "name" : "Bas",  
  "group" : [ "Open Source", "mongoDB", "Big Data" ]  
}
```

Replication  
Sharding  
Durability

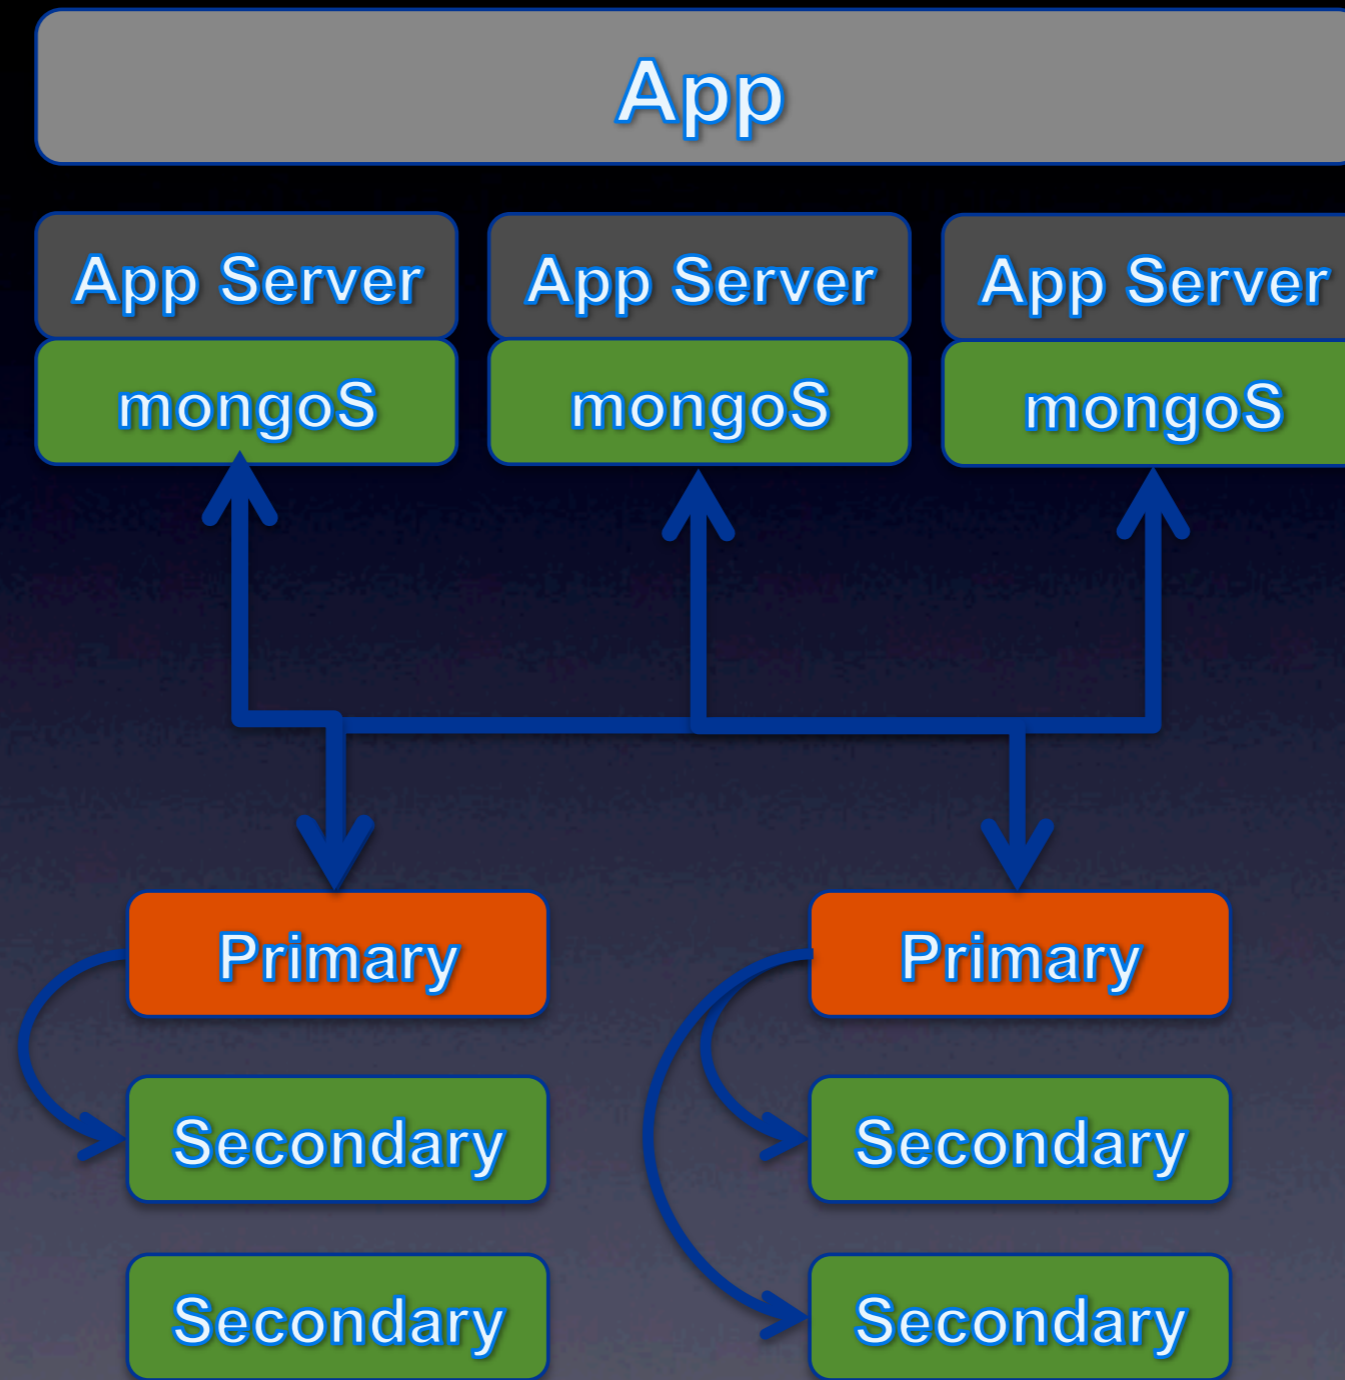
# Replicatie



**Automatic election  
of new Primary**

# Sharding



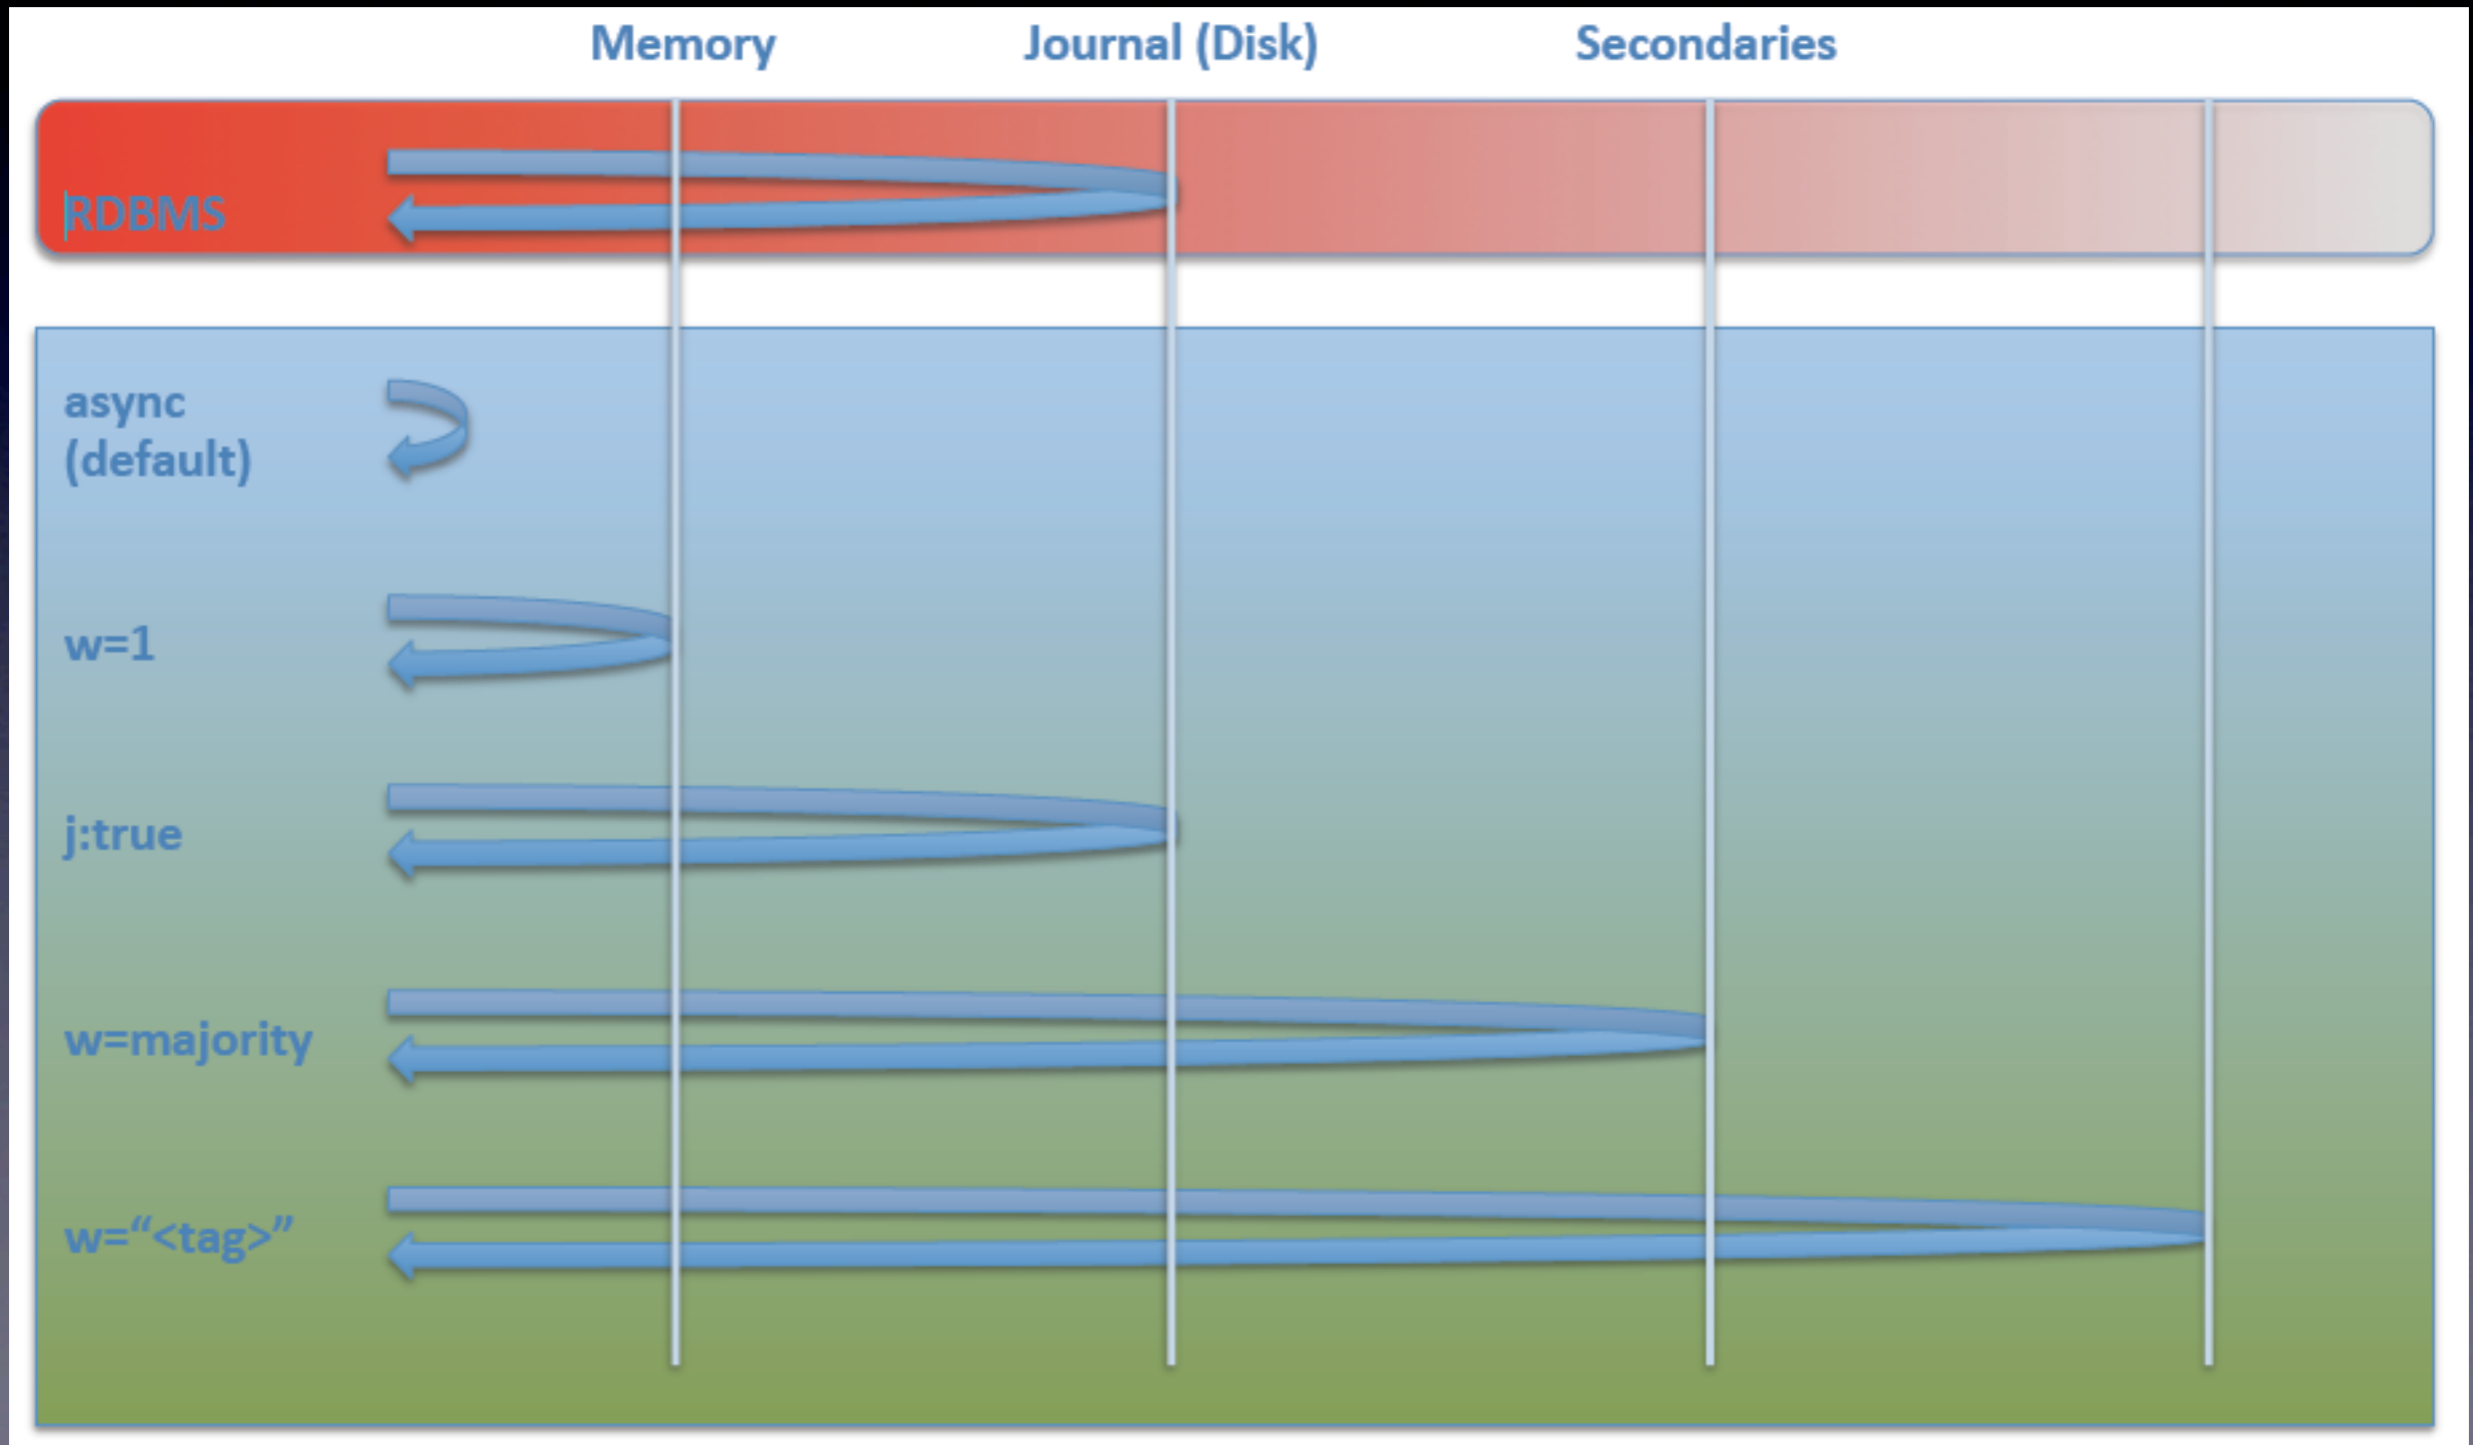


# Sharding features

- Automatic partitioning
- Automatic Load-Balancing across shards
- Range-based
- Convert to sharded system with no downtime
- Fully consistent
- Application code unaware of data location
- Zero code changes



# Durability





## Lab00 Install MongoDB

### OSX

In a terminal session, begin by downloading the latest release. In most cases you will want to download the 64-bit version of MongoDB. Replace x.y.z with the current stable version (i.e. 2.2.0).

```
1 | curl http://downloads.mongodb.org/osx/mongodb-osx-x86_64-x.y.z.tgz > mongo.tgz
```

Once you've downloaded the release, issue the following command to extract the files from the archive:

```
1 | tar -zxvf mongo.tgz
```

Before you start mongod for the first time, you will need to create the data directory. By default, mongod writes data to the /data/db/ directory. To create this directory, and set the appropriate permissions use the following commands:

```
1 | sudo mkdir -p /data/db
2 | sudo chown `id -u` /data/db
```

You can start a mongod directly in the terminal without creating a control script. This command assumes that the mongod binary is accessible via your system's search path, and that you have created a default configuration file located at /etc/mongod.conf.

```
1 | mongod --config /etc/mongod.conf
```

Now you can start MongoDB from another shell

```
1 | ./bin/mongo
```

[More info](#)

[Install MongoDB on onther Linux environments](#)

[Install MongoDB on Windows](#)

Labs

Lab00 Install MongoDB Lab01 Basic

CyberChimps 

© MongoDB.Info



## Lab01 Basic CRUD

Type some commands like

```
1 > help
2 db.help() help on db methods
3 db.mycoll.help() help on collection methods
4 sh.help() sharding helpers
5 rs.help() replica set helpers
6 help admin administrative help
7 help connect connecting to a db help
8 help keys key shortcuts
9 help misc misc things to know
10 help mr mapreduce
11 show dbs show database names
12 show collections show collections in current database
13 show users show users in current database
14 show profile show most recent system.profile entries with time >= 1ms
15 show logs show the accessible logger names
16 show log [name] prints out the last segment of log in memory, 'global' is default
17 use <db_name> set current database
18 db.foo.find() list objects in collection foo
19 db.foo.find( { a : 1 } ) list objects in foo where a == 1
20 it result of the last line evaluated; use to further iterate
```

```
21 | DBQuery.shellBatchSize = x set default number of items to display on shell
22 | exit quit the mongo shell
```

```
1 | > version()
2 | version: 2.2.0
```

```
1 | > show databases
```

```
1 | > show collections
```

### Switch database to test

```
1 | > use test
```

### Show database statistics

```
1 | > db.stats()
```

### Show Javascript function of db.stats()

```
1 | > db.stats
```

### Make a new database

```
1 | > use portraitGallery
2 | switched to db portraitGallery
```

### Create first documents in collection person

```
1 | db.person.insert (
2 | {
3 |   "name" : "Bas",
4 |   "group" : [ "Open Source", "mongoDB", "Big Data" ],
5 | } )
```

and another one

```
1 db.person.save(  
2 {  
3   "name" : "Maikel",  
4   "group" : [ "Oracle", "ExaData", "Big Data"],  
5 } )
```

```
1 Functionally, save and insert are very similar, especially if no _id  
2 value is passed. However, if an _id key is passed, save() will update  
3 the document, while insert() will throw a duplicate key error.
```

```
1 > show collections  
2 person  
3 system.indexes
```

```
1 > db.person.count()  
2 2
```

## Querying

```
1 > db.person.find()  
2 { "_id" : ObjectId("508ae77091896857c075af71"), "name" : "Bas", "group" : [ "Open  
3 Source", "mongoDB", "Big Data" ] }  
3 { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel", "group" : [  
   "Oracle", "ExaData", "Big Data" ] }
```

```
> db.person.find({"group" : "Big Data" })
```

```
1 { "_id" : ObjectId("508ae77091896857c075af71"), "name" : "Bas", "group" : [ "Open  
2 Source", "mongoDB", "Big Data" ] }  
2 { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel", "group" : [  
   "Oracle", "ExaData", "Big Data" ] }
```

```
1 > db.person.find({"group" : "Oracle" })
```

```
2 | { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel", "group" : [
  | "Oracle", "ExaData", "Big Data" ] }
```

## Update

If the person Maikel later decides that he no longer wants his groups stored in his document, he can remove the value just as easily using the \$unset operator:

```
1 | > db.person.update({name: "Maikel"}, {$unset: {group: 1}})
2 | > db.person.find()
3 | { "_id" : ObjectId("508ae77091896857c075af71"), "name" : "Bas", "group" : [ "Open
  | Source", "mongoDB", "Big Data" ] }
4 | { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel" }
```

1 | Replace all group members:

```
1 | > db.person.update({name: "Bas" },{ $set: {"group" : ["EOSS", "SQL", "OpenStack"]}})
2 | > db.person.find()
3 | { "_id" : ObjectId("508ae77091896857c075af71"), "group" : [ "EOSS", "SQL",
  | "OpenStack" ] }
4 | { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel" }
```

1 | Add one group member

```
1 | > db.person.update({name: "Bas" },{ $push: { group: "Oracle" } })
2 | > db.person.find()
3 | { "_id" : ObjectId("508aecf791896857c075af73"), "group" : [ "EOSS", "SQL",
  | "OpenStack", "Oracle" ], "name" : "Bas" }
4 | { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel" }
```

1 |

## Upsert

An upsert is a special type of update. If no document is found that matches the update criteria, a new document will be created by combining the criteria and update documents. If a matching document is found, it will be updated normally. Add true parameter.

```
1 | > db.person.update({name: "Jan" }, { $set: {"group" : ["OpenStack"]}}, true)
2 | > db.person.find()
3 | { "_id" : ObjectId("508aecf791896857c075af73"), "group" : [ "EOSS", "SQL",
  | "OpenStack" ], "name" : "Bas" }
4 | { "_id" : ObjectId("508ae77d91896857c075af72"), "name" : "Maikel" }
5 | { "_id" : ObjectId("508af10eb97c8d1dc77ab660"), "group" : [ "OpenStack" ], "name" :
  | "Jan" }
```

## Multiple Updates

Add true parameter

```
1 | > db.person.update( { }, { $set: {"group" : ["EOSS", "mongoDB", "OpenStack"]}}, true,
  | true )
2 | > db.person.find()
3 | { "_id" : ObjectId("508aecf791896857c075af73"), "group" : [ "EOSS", "mongoDB",
  | "OpenStack" ], "name" : "Bas" }
4 | { "_id" : ObjectId("508ae77d91896857c075af72"), "group" : [ "EOSS", "mongoDB",
  | "OpenStack" ], "name" : "Maikel" }
5 | { "_id" : ObjectId("508af10eb97c8d1dc77ab660"), "group" : [ "EOSS", "mongoDB",
  | "OpenStack" ], "name" : "Jan" }
```

## Delete

If given no parameters, a remove operation will clear a collection of all its documents. To get rid of Maikel type:

```
> db.person.remove({"name": "Maikel"})
```

```
1 | > db.person.find()
2 | { "_id" : ObjectId("508aecf791896857c075af73"), "group" : [ "EOSS", "mongoDB",
```



```
3 | "OpenStack" ], "name" : "Bas" }  
  | { "_id" : ObjectId("508af10eb97c8d1dc77ab660"), "group" : [ "EOSS", "mongoDB",  
  | "OpenStack" ], "name" : "Jan" }
```

If your intent is to delete the collection along with all of its indexes, use the drop() method:

```
1 | > db.person.drop()
```

Labs

Lab00 Install MongoDB Lab01 Basic

CyberChimps 

© MongoDB.Info



## Lab02 Replication

Start by creating a data directory for each replica set member, one for the primary and one for the secondary:

```
1 | mkdir /data/node1
2 | mkdir /data/node2
```

Next, start each member as a separate mongod. Since you'll be running each process on the same machine, it's probably easiest to start each mongod in a separate terminal window:

```
1 | mongod --replSet person --dbpath /data/node1 --port 40001
2 | mongod --replSet person --dbpath /data/node2 --port 40002
```

Logon on the primary node to proceed, you need to configure the replica set, because if you examine the mongod log output, the first thing you'll notice are error messages saying that the configuration can't be found.

```
1 | mongo localhost:40001
2 | MongoDB shell version: 2.2.0
3 | connecting to: localhost:40001/test
```

```
1 | > rs.initiate()
2 | {
3 |   "info2" : "no configuration explicitly specified -- making one",
```

```
4  "me" : "Computername.local:40001",
5  "info" : "Config now saved locally. Should come online in about a minute.",
6  "ok" : 1
7  }
```

Now connect again to the primary node, and add the secondary node:

```
1  person:PRIMARY> rs.add("Computername:40002")
2  { "ok" : 1 }
```

Check if the configuration is ok, with rs.status():

```
1  person:PRIMARY> rs.status()
2  {
3    "set" : "person",
4    "date" : ISODate("2012-10-28T19:50:52Z"),
5    "myState" : 1,
6    "members" : [
7      {
8        "_id" : 0,
9        "name" : "Computername.local:40001",
10       "health" : 1,
11       "state" : 1,
12       "stateStr" : "PRIMARY",
13       "uptime" : 1266,
14       "optime" : Timestamp(1351453811000, 1),
15       "optimeDate" : ISODate("2012-10-28T19:50:11Z"),
16       "self" : true
17     },
18     {
19       "_id" : 1,
20       "name" : "Computername.local:40002",
21       "health" : 1,
22       "state" : 2,
23       "stateStr" : "SECONDARY",
24       "uptime" : 41,
25       "optime" : Timestamp(1351453811000, 1),
```

```

26 | "optimeDate" : ISODate("2012-10-28T19:50:11Z"),
27 | "lastHeartbeat" : ISODate("2012-10-28T19:50:51Z"),
28 | "pingMs" : 0
29 | }
30 | ],
31 | "ok" : 1
32 | }

```

And now its time to check if it works. We put a person in our primary database:

```

1 | person:PRIMARY> use portraitGallery
2 | switched to db portraitGallery
3 | person:PRIMARY> db.person.save(
4 | {
5 |   "name" : "Maikel",
6 |   "group" : [ "Oracle", "ExaData", "Big Data"],
7 | } )

```

Logon on the secondary and check if the data is there, and don't forget to enable reading with `rs.slaveOk()` or `db.getMongo().setSlaveOk()`

```

1 | mongo localhost:40002
2 | MongoDB shell version: 2.2.0
3 | connecting to: localhost:40002/test
4 | person:SECONDARY> rs.slaveOk()
5 | person:SECONDARY> use portraitGallery
6 | switched to db portraitGallery
7 | person:SECONDARY> db.person.find()
8 | { "_id" : ObjectId("508d971dda0730903bcbb612"), "name" : "Maikel", "group" : [
9 |   "Oracle", "ExaData", "Big Data" ] }
person:SECONDARY>

```

Now we can test it with a filler script. Type in the primary something like:

```

1 | person:PRIMARY> for(i=0; i<1000000; i++) { db.person.save({person: i}); }

```

And in the secondary check if the collection is filled:

```
1 person:SECONDARY> db.person.find()
2 { "_id" : ObjectId("508f95e9e38917f43ae20db3"), "person" : 0 }
3 { "_id" : ObjectId("508f95e9e38917f43ae20db4"), "person" : 1 }
4 { "_id" : ObjectId("508f95e9e38917f43ae20db5"), "person" : 2 }
5 { "_id" : ObjectId("508f95e9e38917f43ae20db6"), "person" : 3 }
6 { "_id" : ObjectId("508f95e9e38917f43ae20db7"), "person" : 4 }
7 { "_id" : ObjectId("508f95e9e38917f43ae20db8"), "person" : 5 }
8 { "_id" : ObjectId("508f95e9e38917f43ae20db9"), "person" : 6 }
9 { "_id" : ObjectId("508f95e9e38917f43ae20dba"), "person" : 7 }
10 { "_id" : ObjectId("508f95e9e38917f43ae20dbb"), "person" : 8 }
11 { "_id" : ObjectId("508f95e9e38917f43ae20dbc"), "person" : 9 }
12 { "_id" : ObjectId("508f95e9e38917f43ae20dbd"), "person" : 10 }
13 { "_id" : ObjectId("508f95e9e38917f43ae20dbe"), "person" : 11 }
14 { "_id" : ObjectId("508f95e9e38917f43ae20dbf"), "person" : 12 }
15 { "_id" : ObjectId("508f95e9e38917f43ae20dc0"), "person" : 13 }
16 { "_id" : ObjectId("508f95e9e38917f43ae20dc1"), "person" : 14 }
17 { "_id" : ObjectId("508f95e9e38917f43ae20dc2"), "person" : 15 }
18 { "_id" : ObjectId("508f95e9e38917f43ae20dc3"), "person" : 16 }
19 { "_id" : ObjectId("508f95e9e38917f43ae20dc4"), "person" : 17 }
20 { "_id" : ObjectId("508f95e9e38917f43ae20dc5"), "person" : 18 }
21 { "_id" : ObjectId("508f95e9e38917f43ae20dc6"), "person" : 19 }
22 Type "it" for more
23 person:SECONDARY> db.person.count()
24 194079
25 person:SECONDARY> db.person.count()
26 215657
27 person:SECONDARY> db.person.count()
28 228488
29 person:SECONDARY> db.person.count()
30 239528
31 person:SECONDARY>
```

Works!!!

Labs



© MongoDB.Info

Lab00 Install MongoDB Lab01 Basic