



UNIVERSITY OF TEXAS AT ARLINGTON

Introduction to mongoDB

*Adopted from slides and/or materials from
mongoDB.org, W. Bo, R. Gabor, C. Scheich and K. Avery*

Outline



University of Texas
at Arlington

- ▶ Background
- ▶ Data model
- ▶ Comparison to RDBMS
- ▶ ***CRUD*** operations
- ▶ Replication
- ▶ Sharding

Background



University of Texas
at Arlington

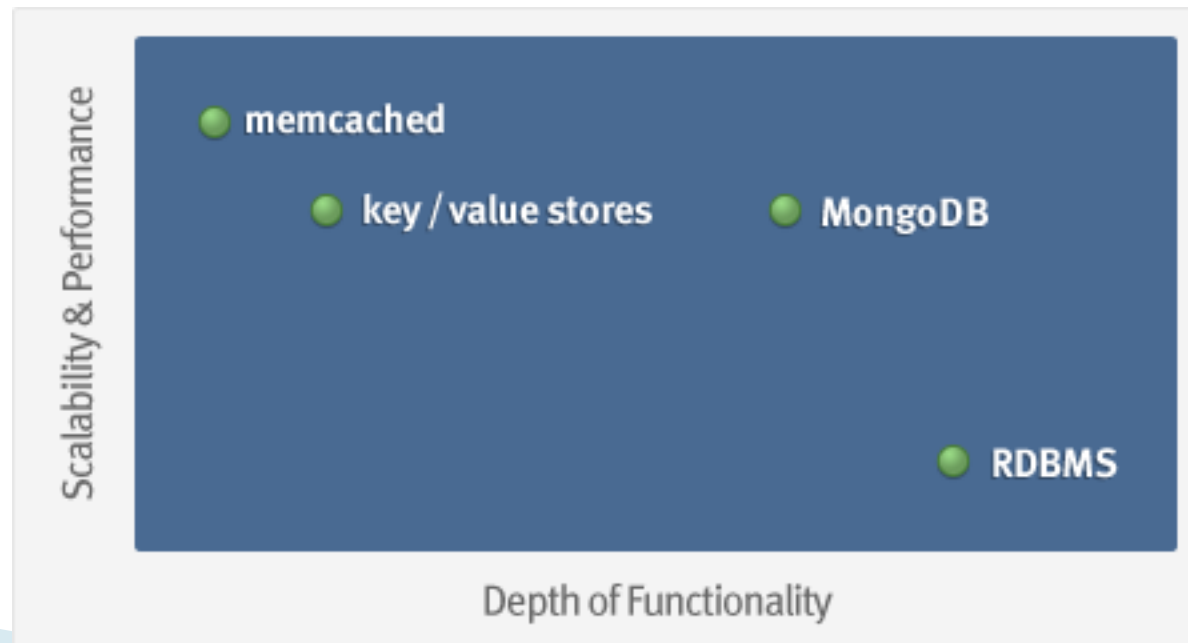
- ▶ stands for humongous
- ▶ created by 10gen
- ▶ open source (Implemented in C++)
- ▶ document-oriented database designed with
 - 1) scalability
 - 2) developer agility in mind
- ▶ Instead of storing data in tables and rows as we would do with a relational database, in MongoDB we store **BSON** (binary representation of JSON) documents with dynamic schemas (schema-less).

Background



University of Texas
at Arlington

- ▶ **Goal:**
 - bridge the gap between key–value stores (which are fast and scalable) and relational databases (which have rich functionality).



Background



University of Texas
at Arlington

- Key features:
 - scale horizontally over commodity hardware
 - support RDBMS features:
 - Ad hoc queries
 - Fully featured indexes
 - Single-field
 - Compound
 - Multi-key
 - Geospatial
 - Text
 - Aggregation operations (through map-reduce)

Data model



University of Texas
at Arlington

- DB ← Collections ← Documents
- Collection = set of “related” documents sharing common indexes
- Primary key (*object id*) is automatically created and indexed for a document
 - Help narrow search span; otherwise mongoDB needs to scan all documents in every shards

Data model



University of Texas
at Arlington

- Schema design:
 - Embedding
 - nesting of objects and arrays inside a BSON document (pre-joined) → fast
 - Link (referencing)
 - references between documents → requires follow-up query
- Principle:
 - embedding when you can, link when you must

Comparison to RDBMS



University of Texas
at Arlington

RDBMS	MongoDB
Database	Database
Table, View	Collection
Row	Document (BSON → JSON)
Column	Field
Index	Index
Join	Embedding/Referencing
Foreign Key	Referencing
Primary Key	ObjectID

CRUD operations



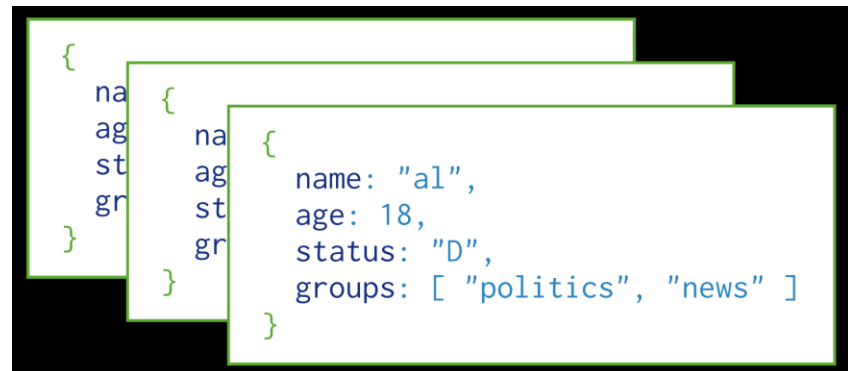
University of Texas
at Arlington

▶ CRUD

- stands for **C**reate, **R**ead, **U**ppdate, and **D**elate
 - used for reading and manipulating data
- ## ▶ Example: JSON documents and collection

```
{  
  name: "sue",  
  age: 26,  
  status: "A",  
  groups: [ "news", "sports" ]  
}
```

← field: value
← field: value
← field: value
← field: value



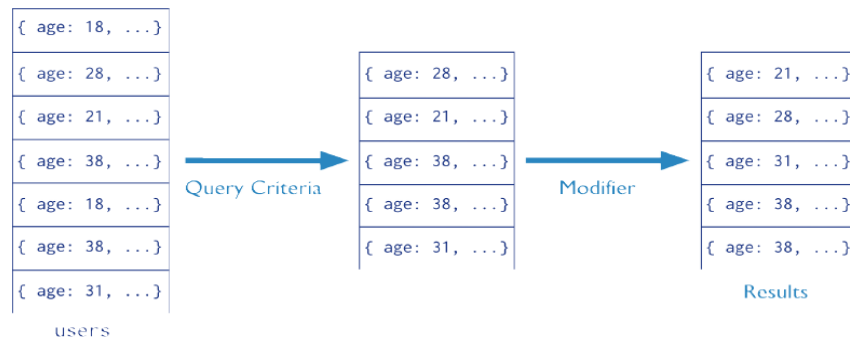
CRUD operations



University of Texas
at Arlington

- ▶ target a specific collection at a time
 - NOT support JOIN operations

Collection Query Criteria Modifier
`db.users.find({ age: { $gt: 18 } }).sort({age: 1 })`



- ▶ Note:
 - Read op.: users have options to choose preferred member in a replica set; otherwise primary
 - Write op.: users specify consistency level (*write concern*)

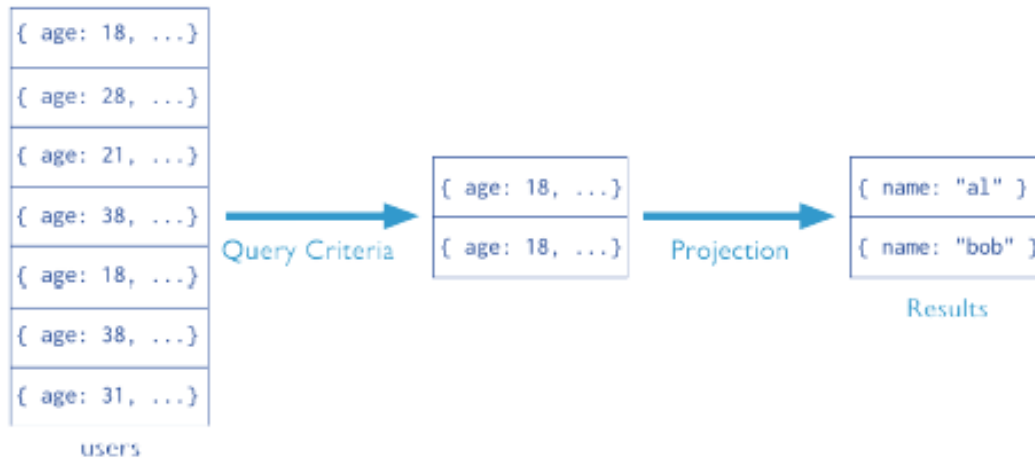
CRUD operations



University of Texas
at Arlington

► More example:

Collection Query Criteria Projection
`db.users.find({ age: 18 }, { name: 1, _id: 0 })`



CRUD operations



University of Texas
at Arlington

Collection
↓
db.users.insert(
Document
↓
{
 name: "sue",
 age: 26,
 status: "A",
 groups: ["news", "sports"]
}

Document

```
{  
  name: "sue",  
  age: 26,  
  status: "A",  
  groups: [ "news", "sports" ]  
}
```

insert →

Collection

{ name: "al", age: 18, ... }
{ name: "lee", age: 28, ... }
{ name: "jan", age: 21, ... }
{ name: "kai", age: 38, ... }
{ name: "sam", age: 18, ... }
{ name: "mel", age: 38, ... }
{ name: "ryan", age: 31, ... }
{ name: "sue", age: 26, ... }

users

CRUD operations



University of Texas
at Arlington

```
SELECT _id, name, address
FROM users
WHERE age > 18
LIMIT 5
```

← projection
← table
← select criteria
← cursor modifier



```
db.users.find(
  { age: { $gt: 18 } },
  { name: 1, address: 1 }
).limit(5)
```

← collection
← query criteria
← projection
← cursor modifier

- ▶ Results always include object id (*_id*) unless explicitly specify (*_id:0*)

CRUD operations



University of Texas
at Arlington

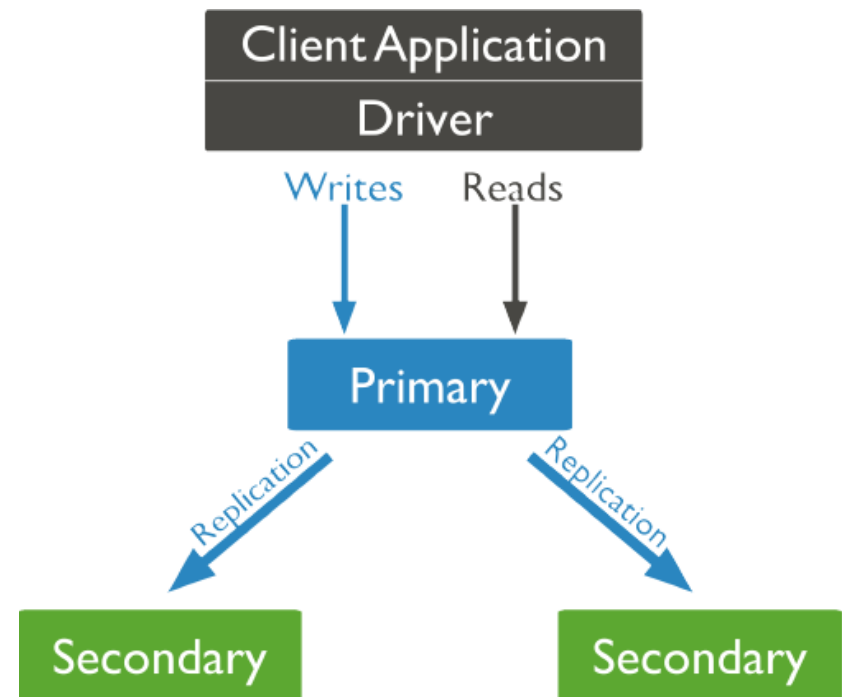
- ▶ Create
 - `db.collection.insert(<document>)`
 - `db.collection.update(<query>, <update>, { upsert: true })`
- ▶ Read
 - `db.collection.find(<query>, <projection>)`
- ▶ Update
 - `db.collection.update(<query>, <update>, <options>)`
- ▶ Delete
 - `db.collection.remove(<query>, <justOne>)`

Replication



University of Texas
at Arlington

- ▶ Replica set
 - group of servers (or *mongod* processes) that maintain the same data set.
 - Only one server is active for writes (the primary, or master) at a given time

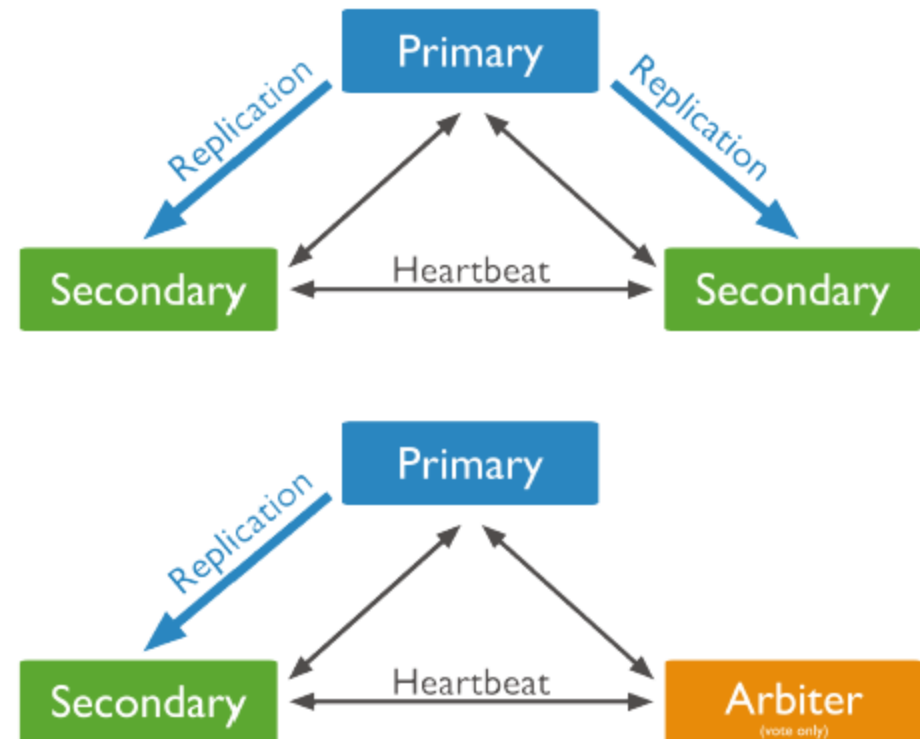


Replication



University of Texas
at Arlington

- ▶ Replica set
 - Replica set members send heartbeats (pings) to each other every two seconds. If a heartbeat does not return within 10 seconds, the other members mark the delinquent member as inaccessible.
 - Arbiter doesn't have copy of data (vote only)




Sharding



University of Texas
at Arlington

- ▶ divides the data set and distributes the data over multiple servers, or shards (order-preserving manner)
- ▶ Each shard is an independent database, and collectively, the shards make up a single logical database.

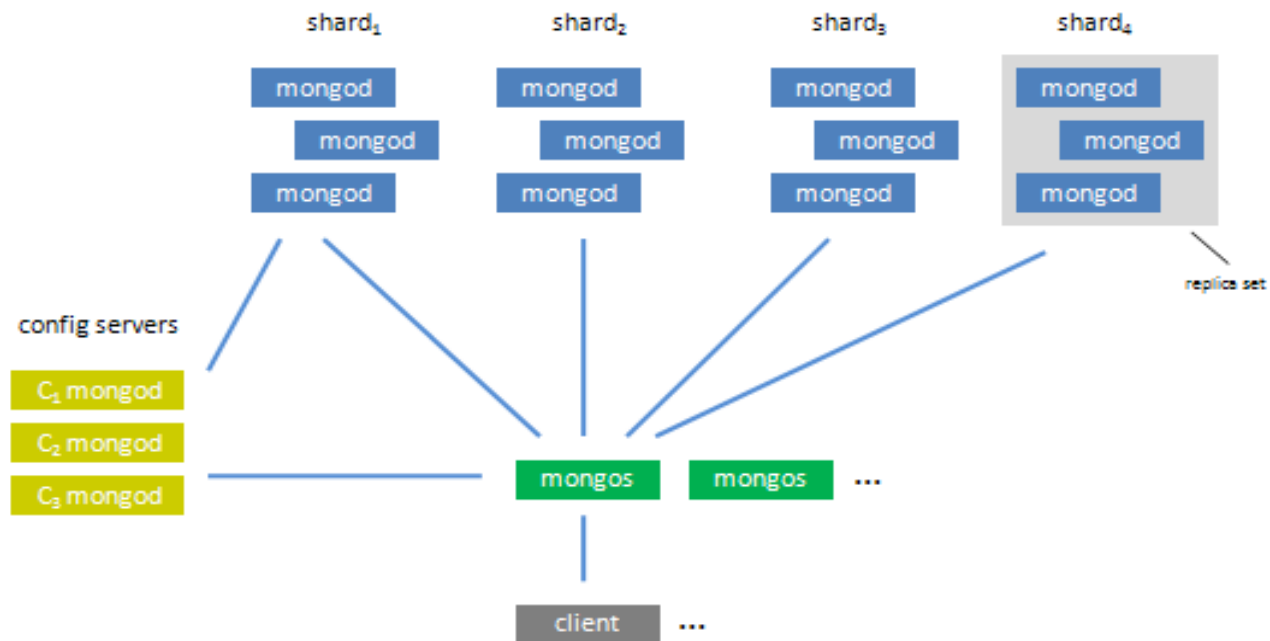


Sharding



University of Texas
at Arlington

- ▶ The set of servers/mongod process within the shard comprise a replica set




Sharding



University of Texas
at Arlington

▶ Example:




Machine 1	Machine 2	Machine 3
Alabama → Arizona	Colorado → Florida	Arkansas → California
Indiana → Kansas	Idaho → Illinois	Georgia → Hawaii
Maryland → Michigan	Kentucky → Maine	Minnesota → Missouri
Montana → Montana	Nebraska → New Jersey	Ohio → Pennsylvania
New Mexico → North Dakota	Rhode Island → South Dakota	Tennessee → Utah
	Vermont → West Virginia	Wisconsin → Wyoming

Sharding



University of Texas
at Arlington

- ▶ When the distribution of a sharded collection in a cluster is uneven, the *balancer* process will perform chunks migration

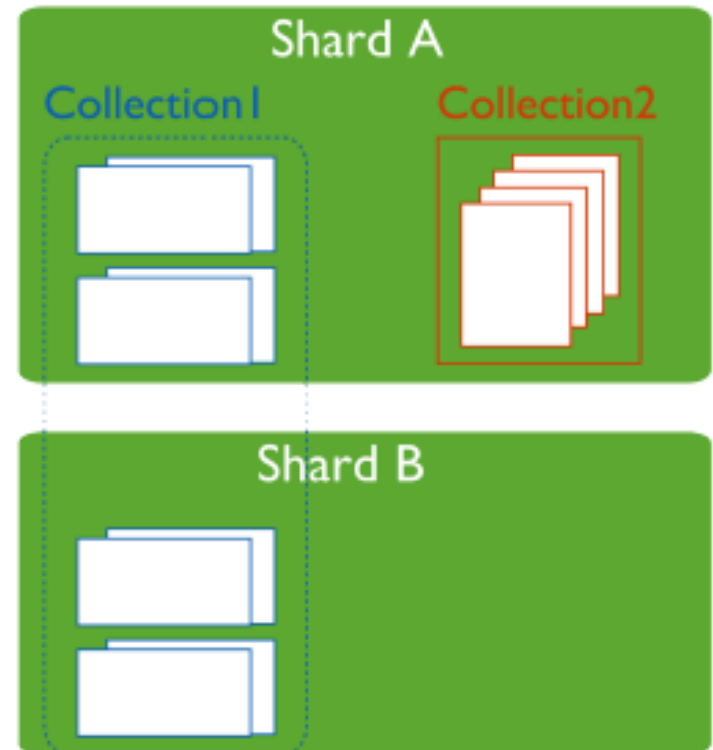


Sharding



University of Texas
at Arlington

- ▶ Every database has a “primary” shard that holds all the un-sharded collections in that database
- ▶ Example:
 - shard A = primary shard



Replication & Sharding



University of Texas
at Arlington

- ▶ Sharding is the tool for scaling a system
- ▶ Replication is the tool for data safety, high availability, and disaster recovery.
- ▶ The two work in tandem yet are orthogonal concepts in the design.

Questions?



University of Texas
at Arlington

