

Objectives

- Project Prep
- Data: MongoDB

March 22, 2017

Sprenkle - CSCI397

1

Tool Teams

Monday (2)
Wednesday (1)
Friday (2)

GraphQL & Falcor

- Patrick
- Perry
- Phil
- Sarah Anne

Fri

JMeter

- Brandon
- Joe
- Steve

Wed

React

- Keith
- Mina
- Rob

Mon

Firebase

- Janie
- Nika

Wed

Flask

- Emily
- Hammad
- Sima

Wed

March 22, 2017

Sprenkle - CSCI397

2

Review

- What data storage/search mechanisms have we talked about so far?
 - How do we access them?
 - What can we do with them?
 - What are their relative strengths and limitations?

March 22, 2017

Sprenkle - CSCI397

3

MONGODB

March 22, 2017

Sprenkle - CSCI397

4

MongoDB Overview

- NoSQL
 - Documents rather than records
- Represents JSON documents in BSON (binary-encoded format)
 - Benefits: provide additional data types, ordered fields, efficient for encoding and decoding within different languages

March 22, 2017

Sprenkle - CSCI397

5

MongoDB

- `db.getCollectionNames()`

March 22, 2017

Sprenkle - CSCI397

6

CRUD Examples

```
> db.user.insert({
  first: "John",
  last : "Doe",
  age: 39
})
```

```
> db.user.find ({
  "first" : "John",
  "last" : "Doe",
  "age" : 39
})
```

```
> db.user.update(
  {"_id" :
  ObjectId("51...")},
  {
    $set: {
      age: 40,
      salary: 7000}
  }
)
```

```
> db.user.remove({
  "first": /^J/
})
```

```
DELETE * FROM user WHERE
first LIKE "J%";
```

March 22, 2017

Sprenkle - CSCI397

7

Inserting

- `db.collectionname.insert(document)`
 - Document → JSON
 - `db.customers.insert({"name": "Jane Doe"})`

March 22, 2017

Sprenkle - CSCI397

8

Search: find()

- `db.bank.findOne()`
- `db.bank.find()`
- `db.bank.find().pretty()`
- `db.bank.find({lastname: "Johnson"})`

March 22, 2017

Sprenkle - CSCI397

9

What are these queries finding?

- `db.bank.find({ age : { $gt : 35 }})`
- `db.bank.find({ age : { $gte : 40 }})`
- `db.bank.find({ state : { $in : ["VA", "PA"] }})`
- `db.bank.find({ $and : [{ balance : {$gt : 40000} }, {age : {$lt : 25} }] })`
- `db.bank.find({ $and : [{ balance : {$gt : 40000} }, {age : {$lt : 25} }] }).count()`
Now, find just the women

March 22, 2017

Sprenkle - CSCI397

10

Functions

- `count(query, options)`
 - <https://docs.mongodb.com/manual/reference/method/db.collection.count/#db.collection.count>
- `distinct(field, query, options)`
 - <https://docs.mongodb.com/manual/reference/method/db.collection.distinct/#db.collection.distinct>

March 22, 2017

Sprenkle - CSCI397

11

Comparing Schema Design

- SQL: Optimizing how data is stored
- MongoDB: Optimize how data is used
- SQL: What answers do I have?
- MongoDB: What questions do I have?

March 22, 2017

Sprenkle - CSCI397

12

Buzzword Bingo: CAP Theorem

- Many nodes
- Nodes contain *replicas of partitions* of data
- Consider what happens
 - Under high load
 - During updates

March 22, 2017 Sprengle - CSCI397 13

Buzzword Bingo: CAP Theorem

- Consistency
 - all replicas contain the same version of data
- Availability
 - system remains operational (reads AND writes) on failing nodes
- Partition tolerance
 - multiple entry points
 - system remains operational on network split

CAP Theorem: satisfying all three at the same time is impossible

March 22, 2017 Sprengle - CSCI397 14

Theory of noSQL

- Consistency
 - all replicas contain the same version of data
- Availability
 - system remains operational (reads and writes) on failing nodes
- Partition tolerance
 - multiple entry points
 - system remains operational on network split

CAP Theorem: satisfying all three at the same time is impossible

March 22, 2017 Sprengle - CSCI397 15

Theory of noSQL

- Consistency
 - all replicas contain the same version of data
- Availability
 - system remains operational (reads and writes) on failing nodes
- Partition tolerance
 - multiple entry points
 - system remains operational on system split

CAP Theorem: satisfying all three at the same time is impossible

New work at Google says “We can do all three!”

March 22, 2017 Sprengle - CSCI397 16

Benefits

- Speed!
- Rich dynamic queries
- Lazy creation
- Schema-less
- Returns JSON
- Easy Replication and Failover
- Auto-Sharding
- MapReduce

March 22, 2017

Sprenkle - CSCI397

17

Limitations

- No Transactions
 - Not good for purchases, banking, inventory
- No Joins
- RAM intensive
- No referential integrity
- Eventual Consistency

March 22, 2017

Sprenkle - CSCI397

18

Looking Ahead

- Fri: Tool Team work
- Mon: Tool Team work
- Wed
 - Richard Marmorstein's talk
 - Due: Assign 7
 - Tool: What do we need installed?

March 22, 2017

Sprenkle - CSCI397

19