



# Everything You Know About MongoDB is Wrong

(Probably)



Mark Smith | MongoDB | @Judy2K



@Gar1t on YouTube

# Myth 0

---

You think we haven't seen this  
on YouTube





# MongoDB is Web Scale

---

We've seen it.

We've bought the T-shirts.



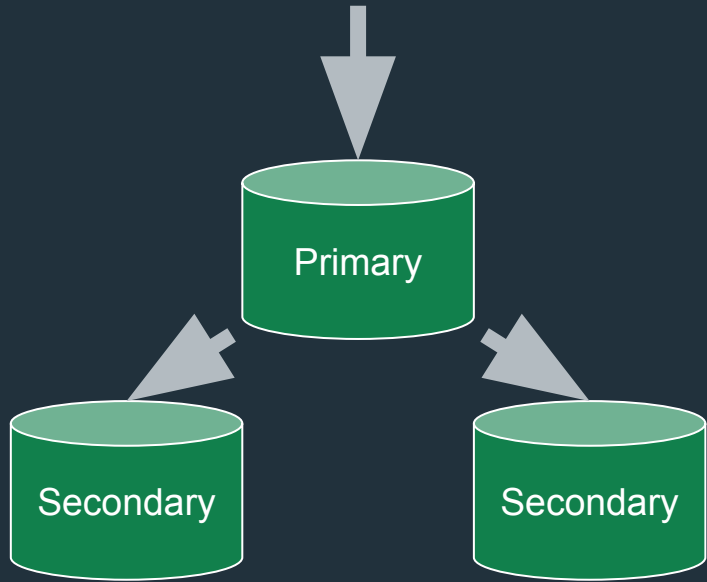
# What is MongoDB?

---




# Database Cluster

---



# Database Cluster

---




# Documents

---

```
{'_id': ObjectId('573a1390f29313caabcd4135'),  
  'title': 'Blacksmith Scene',  
  'fullplot': 'A stationary camera looks at a large anvil  
with a blacksmith behind it and one on either side.',  
  'cast': ['Charles Kayser', 'John Ott'],  
  'countries': ['USA'],  
  'directors': ['William K.L. Dickson'],  
  'genres': ['Short'],  
  'imdb': {'id': 5, 'rating': 6.2, 'votes': 1189},  
  'released': datetime.datetime(1893, 5, 9, 0, 0),  
  'runtime': 1,  
  'year': 1893}
```



# Database Cluster




# Documents

```
{'_id': ObjectId('573a1390f29313caabcd4135'),  
'title': 'Blacksmith Scene',  
'fullplot': 'A stationary camera looks at a large anvil  
with a blacksmith behind it and one on either side.',  
'cast': ['Charles Kayser', 'John Ott'],  
'countries': ['USA'],  
'directors': ['William K.L. Dickson'],  
'genres': ['Short'],  
'imdb': {'id': 5, 'rating': 6.2, 'votes': 1189},  
'released': datetime.datetime(1893, 5, 9, 0, 0),  
'runtime': 1,  
'year': 1893}
```



# Database Cluster



# Documents

```
{'_id': ObjectId('573a1390f29313caabcd4135'),  
  'title': 'Blacksmith Scene',  
  'fullplot': 'A stationary camera looks at a large anvil  
with a blacksmith behind it and one on either side.',  
  'cast': ['Charles Kayser', 'John Ott'],  
  'countries': ['USA'],  
  'directors': ['William K.L. Dickson'],  
  'genres': ['Short'],  
  'imdb': {'id': 5, 'rating': 6.2, 'votes': 1189},  
  'released': datetime.datetime(1893, 5, 9, 0, 0),  
  'runtime': 1,  
  'year': 1893}
```





Myth 1

MongoDB is **v2.4**

---

```
$ sudo apt install mongodb
$ mongod --version
db version v2.4
```



# MongoDB is **v4.4**

---

Google *mongodb* community



Myth 2

# The **JSON** database

---

MongoDB is a document database




The most popular database for x

← → ↻ 🔒 mongodb.com ☆ 🗨️ ⚙️ 🐉 ⋮

mongoDB Cloud Software Learn Solutions Docs 🔍 Contact Sign In [Try Free](#)



MongoDB is a document database, which means it stores data in JSON-like documents. We believe this is the most natural way to think about data, and is much more expressive and powerful than the traditional row/column model.

```
{
  "_id": "5cf0029caff5056591b0ce7d",
  "firstname": "Jane",
  "lastname": "Wu",
  "address": {
    "street": "1 Circle Rd",
    "city": "Los Angeles",
    "state": "CA",
    "zip": "90404"
  },
  "hobbies": ["surfing", "coding"]
}
```



## Rich JSON Documents

- The most natural and productive way to work with data.
- Supports arrays and nested objects as values.
- Allows for flexible and dynamic schemas.



# The **BSON** database

---

MongoDB *is* a document database



Myth 3

# No Transactions

---

MongoDB is a **BASE** database



# Yes Transactions

---

MongoDB is an **ACID** database



Myth 4

# No Relationships

---

You can't have relationships between documents.





# **Yes** Relationships

---

You totally can.



# An Aggregation Pipeline

---

```
db.orders.aggregate([  
  $lookup: {  
    from: "inventory",  
    localField: "item",  
    foreignField: "sku",  
    as: "inventory_docs"  
  }  
])
```



# An Aggregation Pipeline

---

```
{  
  "_id" : 1,  
  "item" : "almonds",  
  "price" : 12,  
  "quantity" : 2,  
  "inventory_docs" : [  
    { "_id" : 1, "sku" : "almonds", "instock" : 120 }  
  ]  
}
```



# A Very Simple Aggregation Pipeline

---

```
{  
  "_id" : 1,  
  "item" : "almonds",  
  "price" : 12,  
  "quantity" : 2,  
  "inventory_docs" : [  
    { "_id" : 1, "sku" : "almonds", "instock" : 120 }  
  ]  
}
```



Myth 5

# MongoDB is About **Sharding**

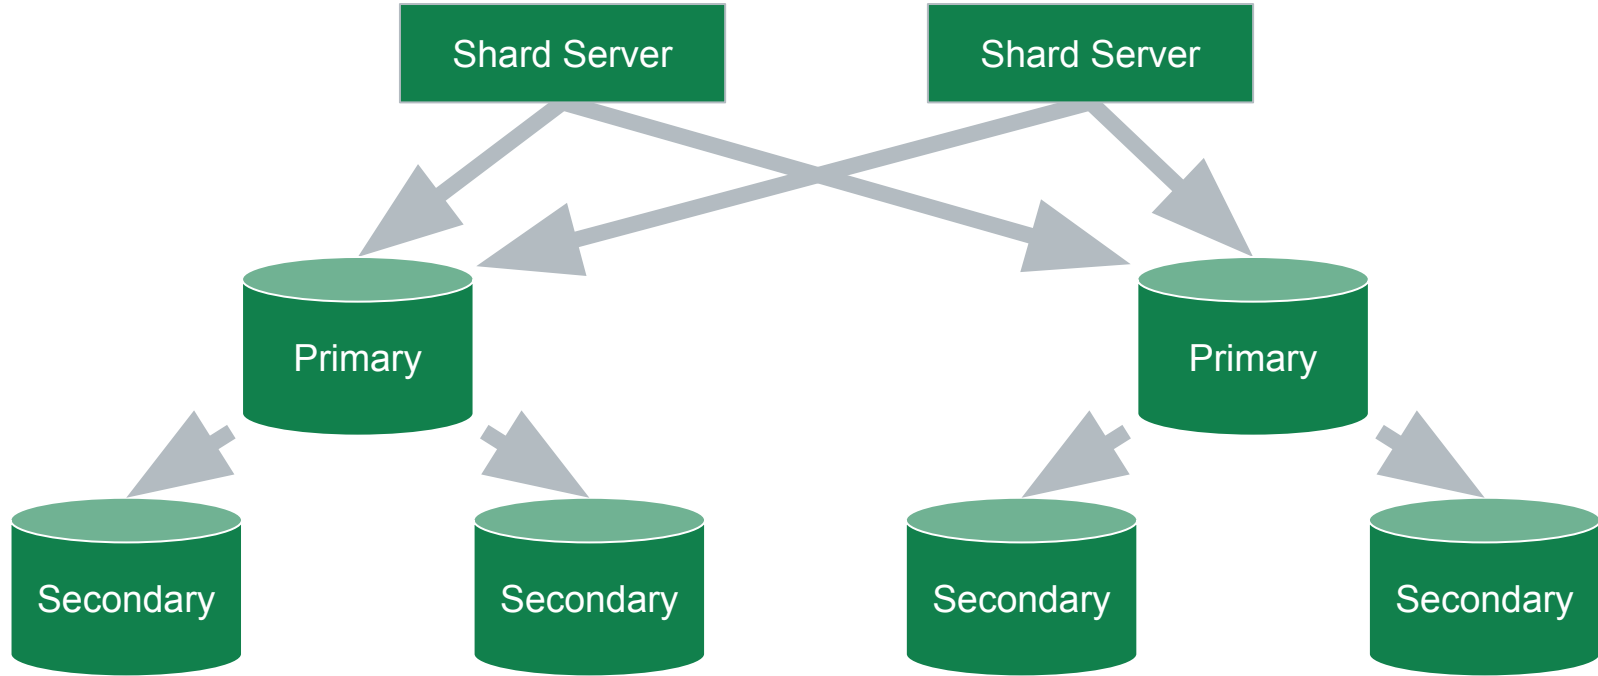
---

That's how you scale



# Quick Sharding Primer - 2 shards

---



# Just Buy a **Bigger Machine**

---

MongoDB *can* do sharding,  
but you probably shouldn't



Actually, use **Atlas**

---

We will host your database for you





# Microsharding

---

... isn't a thing any more



# Map-Reduce

---

... with JavaScript.



## Fault-tolerance



# Aggregation Pipelines

---



# A Very Simple Aggregation Pipeline

---

```
db.orders.aggregate([
  { $match: { status: "A" } },
  { $group: { _id: "$cust_id", total: { $sum: "$amount" }}}
])
```



# A Very Simple Aggregation Pipeline

---

```
db.orders.aggregate([
  { $match: { status: "A" } },
  { $group: { _id: "$cust_id", total: { $sum: "$amount" }}}
])
```



# A Very Simple Aggregation Pipeline

---

```
db.orders.aggregate([
  { $match: { status: "A" } },
  { $group: { _id: "$cust_id", total: { $sum: "$amount" }}}
])
```



Myth 6

# MongoDB is Insecure

---

Haven't there been breaches?





# MongoDB is **Secure**

---

Industry-standard security.  
Bad defaults have been fixed.



Myth 7

# MongoDB **Loses** Data

---

... doesn't it?



# MongoDB Is **Durable**


---

... but we used to have some silly defaults



# Database Cluster

---



Myth 8

# MongoDB is Really **Easy**

---

Well ... yes ... and no.



# COVID-19 Open Data

---

<https://bit.ly/mongodb-covid>





# Some stuff You Know About MongoDB is *Right*

(Probably)



Mark Smith | MongoDB | @Judy2K



# Thank you

That's all folks

Mark Smith | MongoDB | @Judy2k