



Couchbase

Maroš Kopec, Bibiána Ťureková, Nikoleta Češková, Iuliia Gurianova

Agenda

- Evolution
- Main features
- Core principles
- Architecture
- Comparison with MongoDB
- Demo



Couchbase



membase

- Memcached technology
- Sharding
- Replication
- Persistence



- Document oriented model based on JSON

Main Features

- **Flexible data model**
 - Json documents with no fixed schema
- **Easy scalability**
 - Grow cluster without downtime or application changes
 - Automatic sharding and rebalancing
- **High performance**
 - consistently high throughput
 - sub-millisecond response times
- **Always on**
 - No downtime for upgrades, maintenance...

Developer integration

- Provide libraries for:
 - Net
 - PHP
 - Ruby
 - Python
 - C
 - Node.js
 - Java
 - GO



Key customers



Real use cases

- Real-time user activity
- User preferences
- Caching
- Promotions tracking
- Users targeting

Core principles

- Documents stored in **buckets** (“db”)
- **vbuckets**
 - Owner of keys
 - Responsible for distribution and replication
- Document = key + value

Couchbase Data

Server stores metadata
with each key/value pair
(document)
Unique and Kept in RAM



```
meta
{
  "id": "u::tesla",
  "rev": "1-0002bce0000000000",
  "flags": 0,
  "expiration": 0,
  "type": "json"
}
```

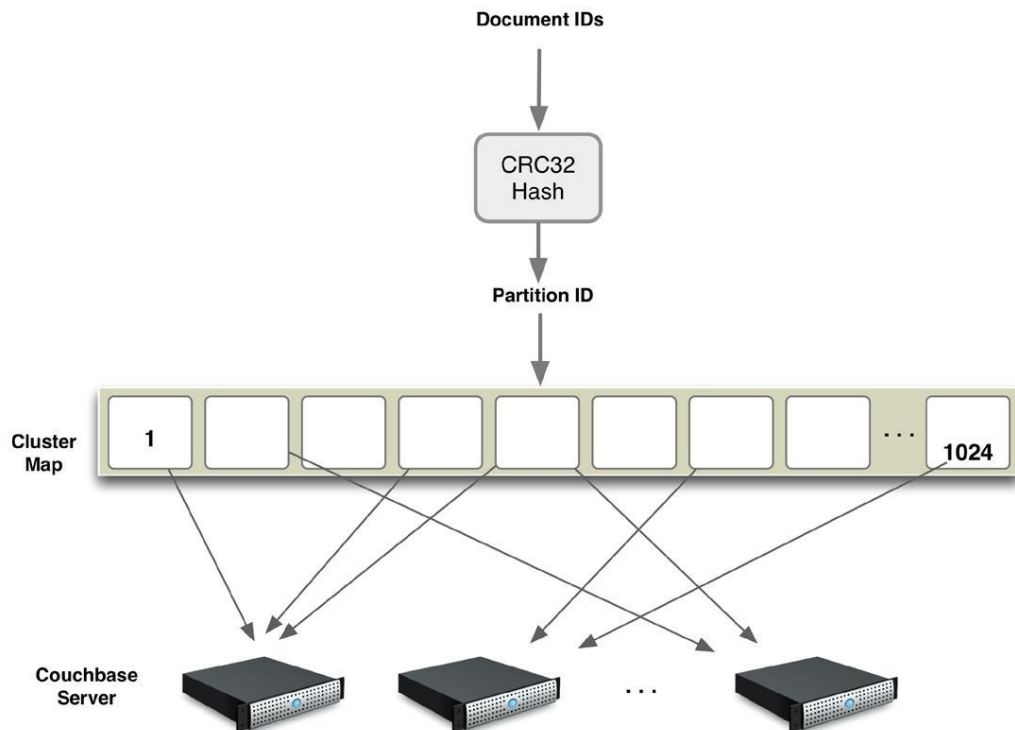
Document Value
Most Recent In RAM And
Persisted To Disk



```
document
{
  "sellerid": 123456,
  "type": "car",
  "style": "sedan",
  "year": 2013,
  "trim": "performance",
  "model": "s"
}
```

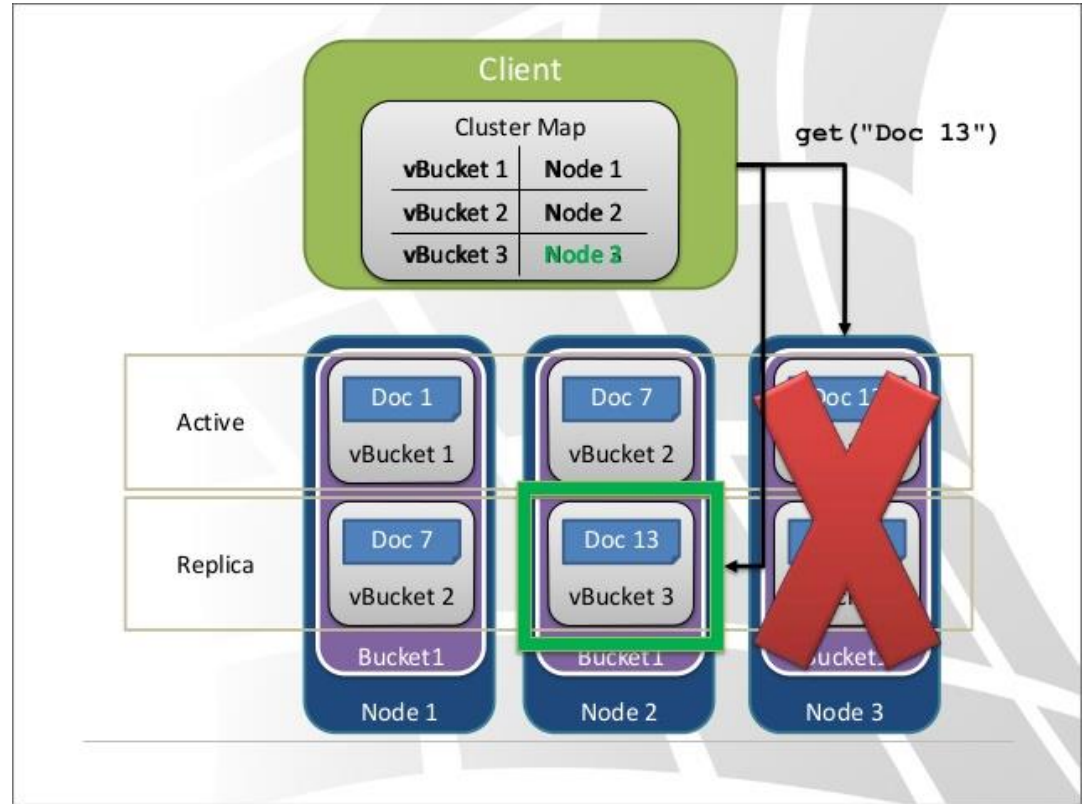

Architecture

- Highly distributed architecture
- Horizontal scaling by hash sharding (by document ID)
- Servers are equal
- Server
 - Data manager
 - Cluster manager



Replication

- documents in a bucket are isolated
- configurable number of replicas (up to three)
- if a server crashes, the system will locate the replicas of the documents, and promote them to active status



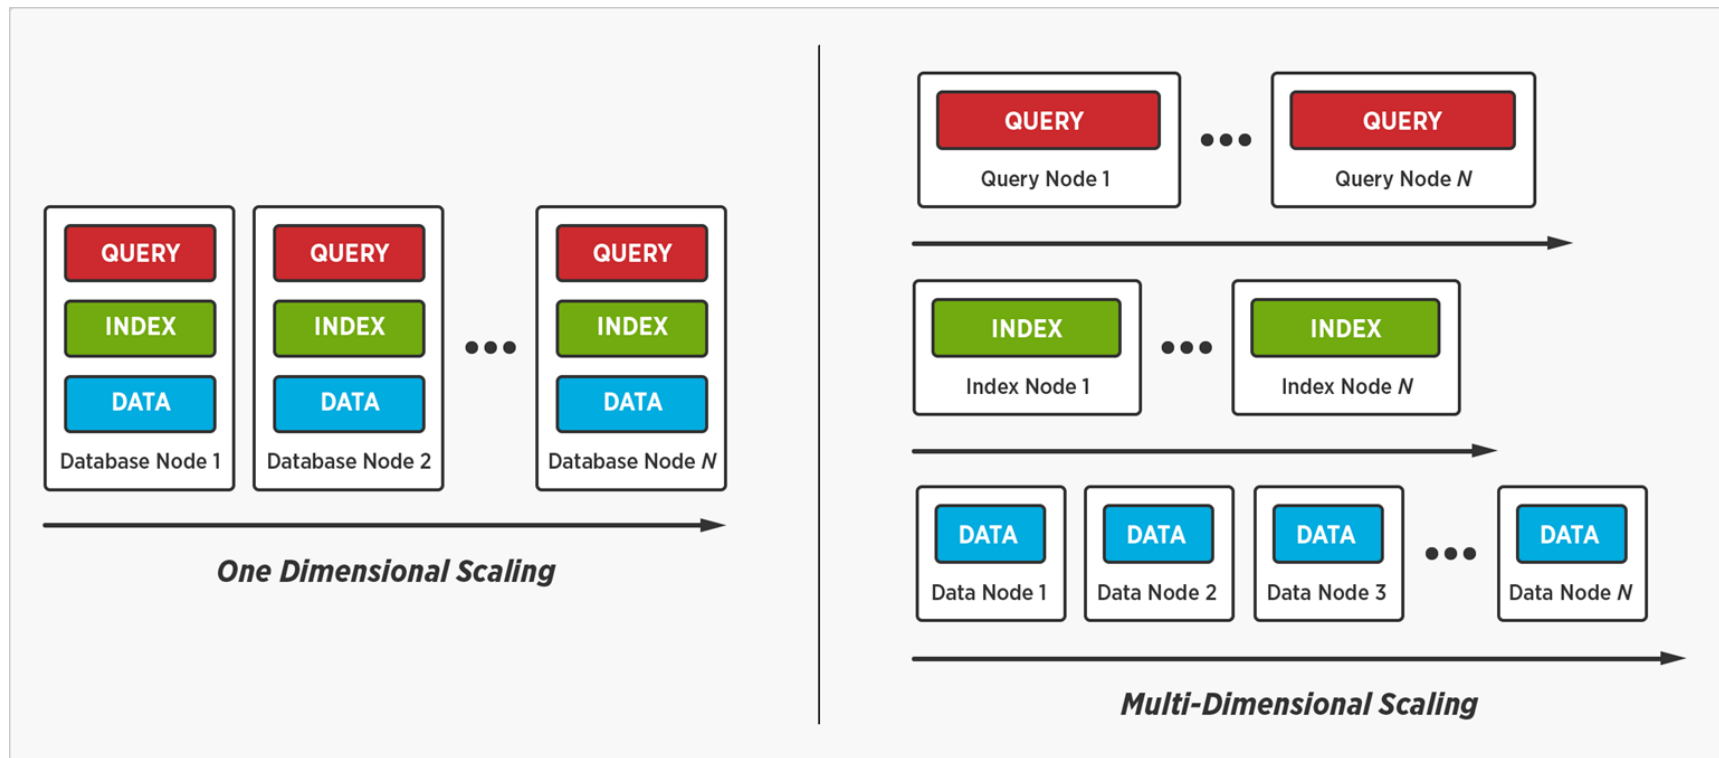
MongoDB



Couchbase

- master-slave; Single-node and sharded environments
 - third-party caches
 - unique proprietary language
 - CLI tools for maintenance
- masterless, clustered, and replicated distributed database; Multi-Dimensional Scaling (MDS)
 - caching layer for both data and indexes
 - SQL-based query language
 - Web UI, CLI, REST API

Multi-Dimensional Scaling



Queries


MongoDB™ Query

```
db.stocks.aggregate([
  { "$match": {
    "$and": [
      {"symbol": {
        "$in": [
          "AAPL",
          "GOOG"]}},
      {"value": {
        "$gt": 0 }}]}},
  { "$group": {
    "_id": {
      "symbol": "$symbol" },
    "sum(value * volume)": {
      "$sum": {
        "$multiply": [
          "$value",
          "$volume"]}}}},
  { "$project": {
    "_id": 0,
    "sum(value * volume)": "$sum(value * volume)",
    "symbol": "$_id.symbol"}}
  { "$sort": {
    "sum(value * volume)": -1,
    "symbol": 1 }}]})
```


Couchbase N1QL

```
SELECT SUM(value * volume) AS val, symbol
FROM db.stocks
WHERE symbol IN ( "AAPL", "GOOG" ) AND value > 0
GROUP BY symbol
ORDER BY val DESC, symbol ASC
```

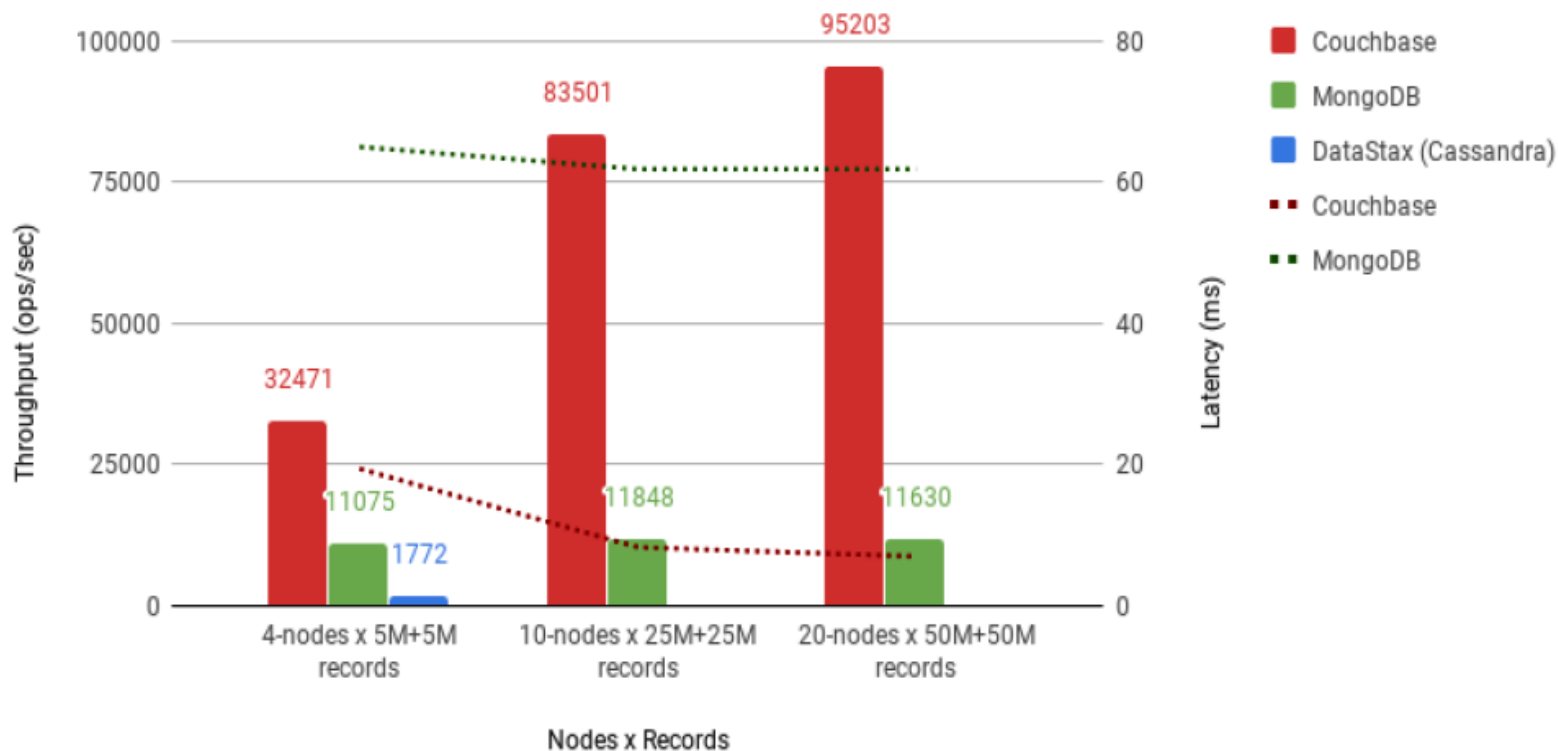
Update workload



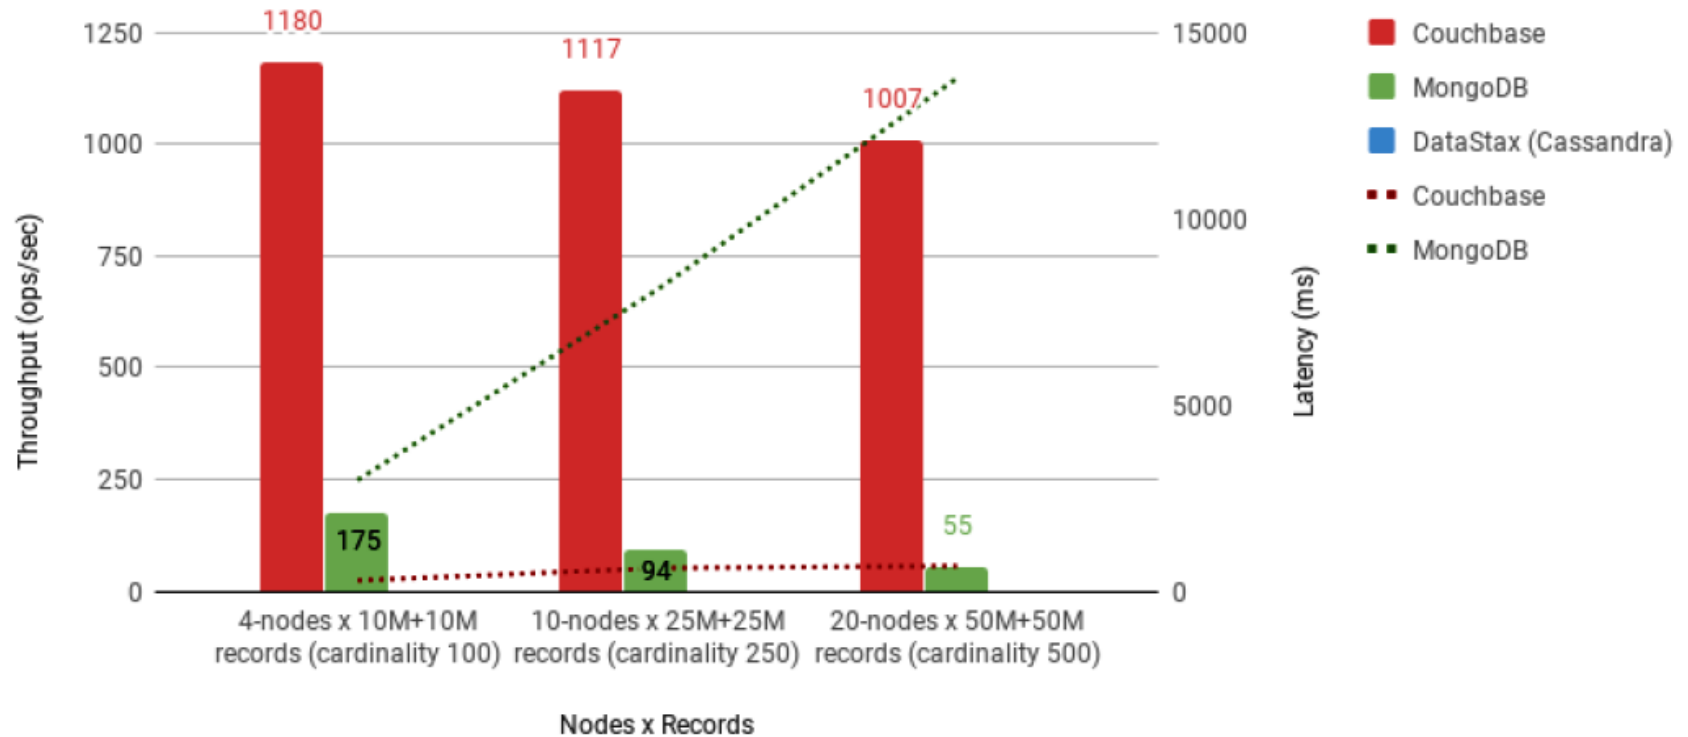
Short-range scan workload



Pagination Workload: Filter with OFFSET and LIMIT



Join Workload: JOIN with GROUP BY and ORDER BY



DEMO

Questions