

MongoDB

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

Database paradigms

- Relational (RDBMS)
- NoSQL
 - Key-value stores
 - Document databases
 - Wide column stores (BigTable and clones)
 - Graph databaseS
- Other

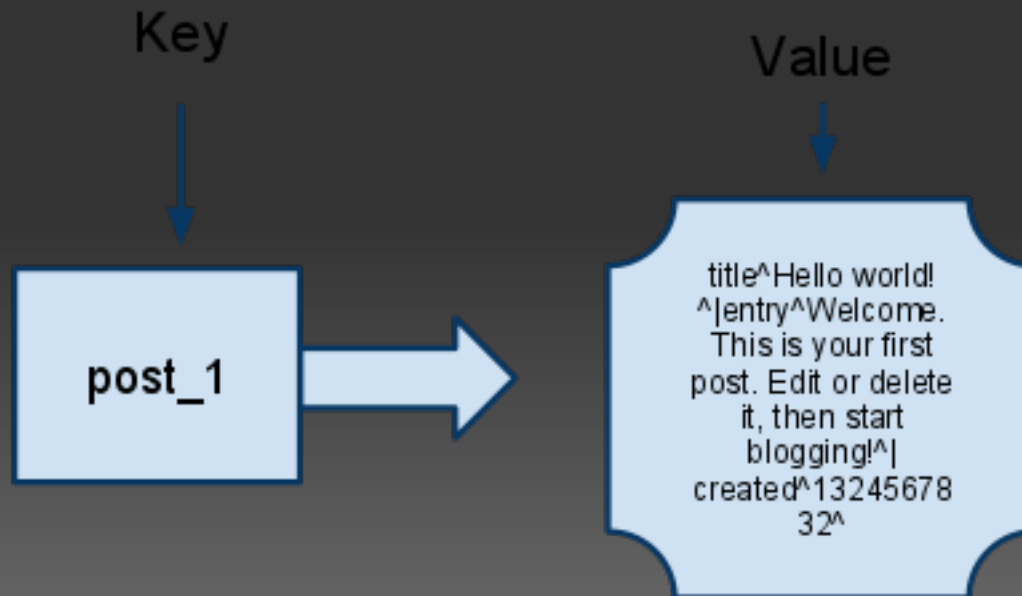
Relational Databases

- ACID (Atomicity Consistency Isolation and Durability)
- SQL
- MySQL, PostgreSQL, Oracle, ...

id	title	entry	created
1	Hello world!	Welcome. This is your first post. Edit or delete it, then start blogging!	13224677834

Key-value stores

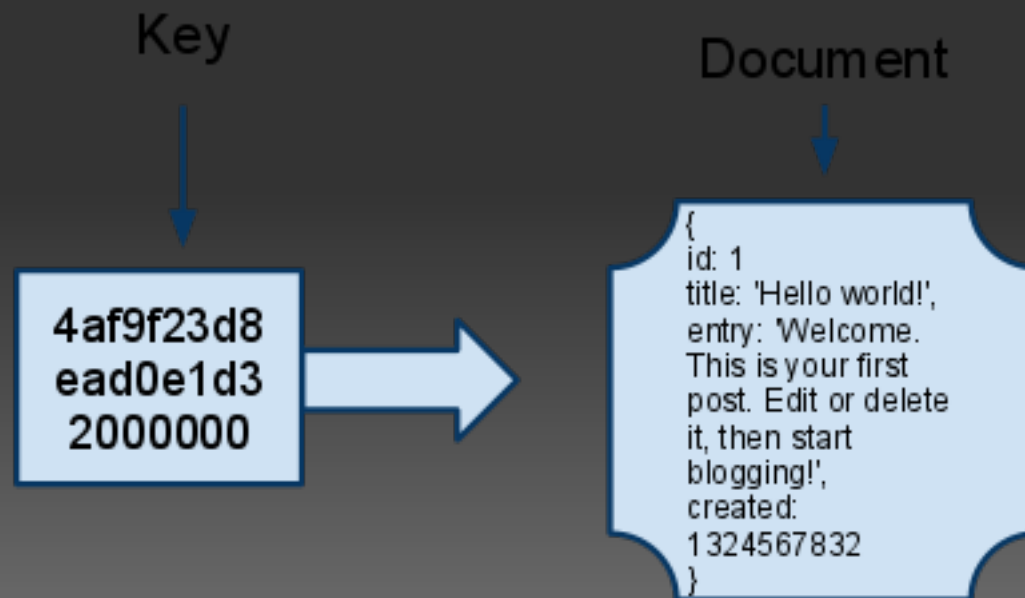
- "One Key, one value, no duplicates, and crazy fast."
- It is a Hash!
- The value is binary object aka."blob" - the DB does not understand it and does not want to understand it.
- Amazon Dynamo, MemcahceDB, ...



Document databases

Key-value store, but the value is (usually) structured and "understood" by the DB.

Querying data is possible (by other means than just a key).
Amazon SimpleDB, CouchDB, MongoDB, Riak, ...



Wide column stores



Often referred as "BigTable clones"

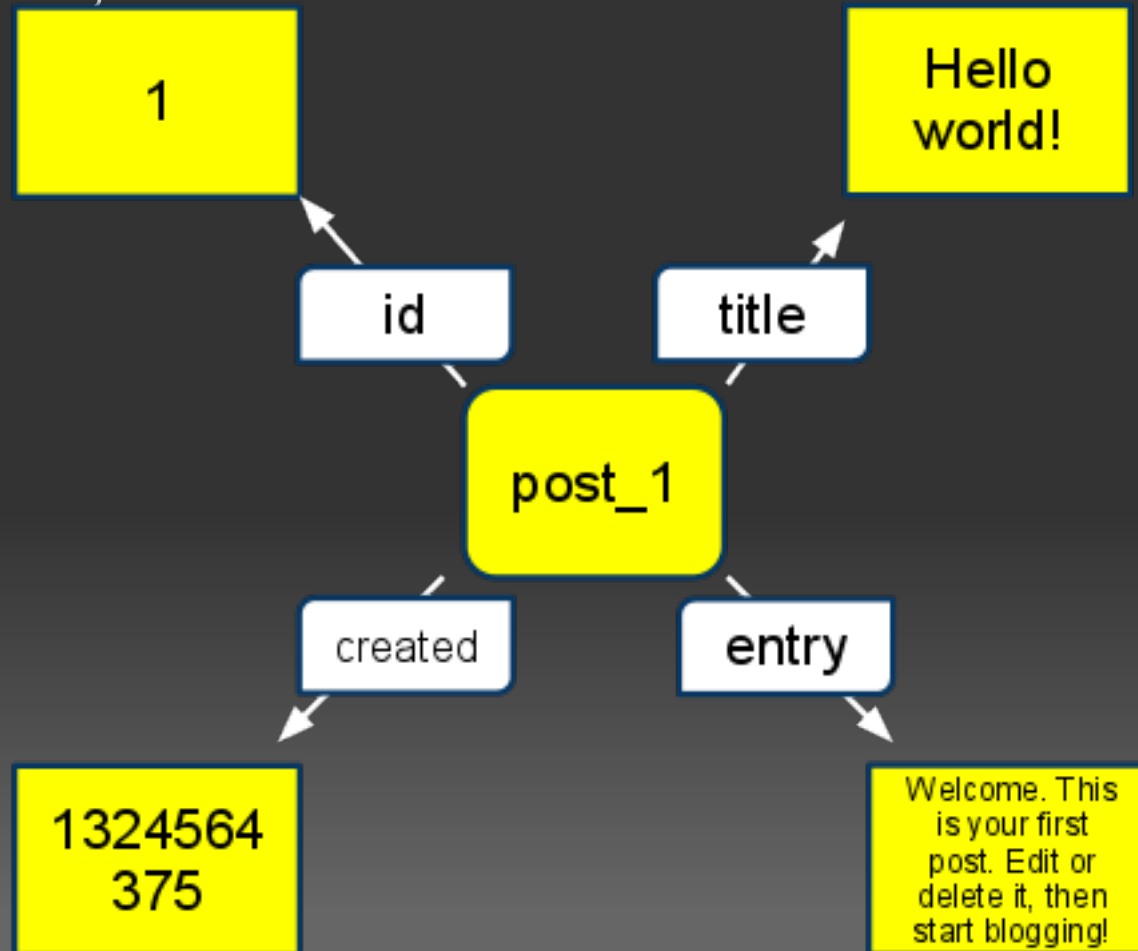
"a sparse, distributed multi-dimensional sorted map"

Google BigTable, Cassandra (Facebook), HBase, ...

row-id	column family	title	time	value
post_1	1	title	14	Hello world!
post_1	1	entry	12	Welcome. This is your first post. Edit or delete it, then start blogging!
post_1	1	created	16	1324067832

Graph databases

- "Relation database is a collection loosely connected tables" whereas "Graph database is a multi-relational graph".
- Neo4j, InfoGrid, ...



Relational databases have almost limitless indexing, and a very strong language for dynamic, cross-table, queries (SQL)

That's why they handle all kinds of relationships well and dynamically.

NoSQL databases

might have limited support for dynamic queries and indexing

don't support JOIN like operations of SQL

but you can store some relationships into document itself

Why NoSQL?

- Schema-free
- Massive data stores
- Scalability
- Some services simpler to implement than using RDBMS
- Great fit for many "web 2.0" services

Why NOT NoSQL?

DRBMS databases and tool are mature

NoSQL implementations often "alpha"

Data consistency, transactions

"Don't scale until you need it"

RDBMS vs, NoSQL

Strong consistency vs, Eventual consistency

Big dataset vs. HUGE datasets

Scaling is possible vs. Scaling is easy

SQL vs. Map-Reduce

Good availability vs. Very high availability

Questions?