


# DATABASES = FUN?

```
select fun, profit from real_world where rational = false
```

# Thanks!

## Teching Software-Engineering

- **20 years with *Relational Databases***
- **PostgreSQL, SQLite, Oracle ...**



2006



2004



Really great stuff!

Sorry. No books available!

2008



http://nosql-database.org 2010



nosqlberlin.de  
nosqlfrankfurt.de  
nosql powerdays 2010

Oracle, IBM, etc. wollen auch...

**No SQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source and horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE (not ACID), a huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

non-relational

**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source and horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE (not ACID), a huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

Und nicht vergessen...

Einfache Installation / Handling + Fun!

# Scale-out!



**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

# Open-Source

nur indirekte finanzielle Interessen



**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

# Web-Scale

Historie



**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

# Schema free

Strange Loop Pattern:

Design for **Crash!**



```
{
  "firstName": "John",
  "lastName": "Smith",
  "age": 25,
  "address": {
    "streetAddress": "21 2nd Street",
    "city": "New York",
    "state": "NY",
    "postalCode": "10021"
  },
  "phoneNumber": [
    {
      "type": "home",
      "number": "212 555-1234"
    },
    {
      "type": "fax",
      "number": "646 555-4567"
    }
  ]
}
```

**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

# Replication fun?

```
POST /_replicate HTTP/1.1
{"source": "database", "target": "http://example.org/database"}
```

**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

**NoSQL DEFINITION:** Next Generation Databases mostly addressing some of the points: being **non-relational, distributed, open-source** and **horizontal scalable**. The original intention has been **modern web-scale databases**. The movement began early 2009 and is growing rapidly. Often more characteristics apply as: **schema-free, easy replication support, simple API, eventually consistent / BASE** (not ACID), a **huge data amount**, and more. So the misleading term "*nosql*" (the community now translates it mostly with "**not only sql**") should be seen as an alias to something like the definition above.

*kommt noch...*




a **NOSQL**  
summer



[nosqltapes.com](http://nosqltapes.com) !

JAP N

NoSQL is specialization!



NoSQL foundations:

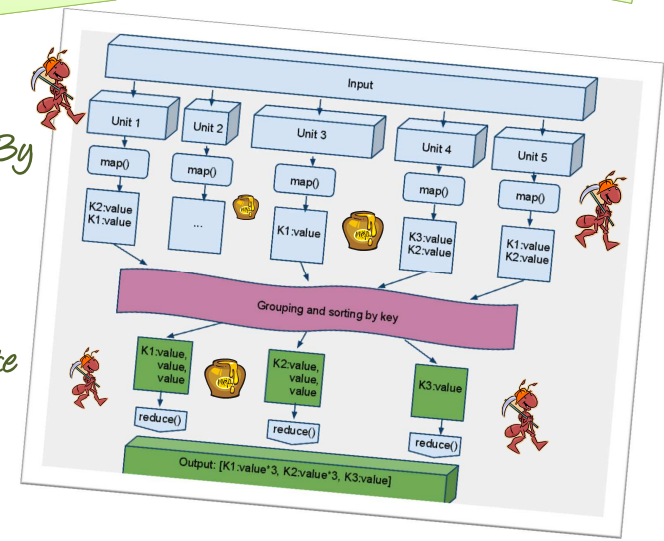
# Theorie

funktionale (graph) Dekomposition? Oder...

Schutzpatent 😊

Group By

Aggregate



Programmierung top!  
herrlich parallelisierbar


Programmierung nervt!

Nur 'large data indexing'



„A giant step back!  
Incompatible, missing  
features, not new, ...”

Starke Konkurrenz: Stratosphere (TUB), ePic, SwissBox, etc.



## Paralellization Contracts

⇒  
compile, analyze, optimize

auf einer atmenden Cloud!

# Eventually Consistent

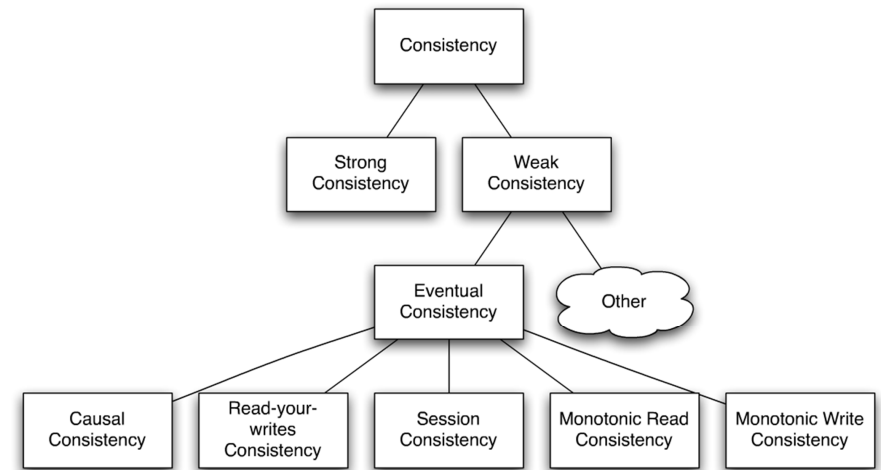
## ACID

- Amazon Dynamo
- MySQL Replication

## BASE



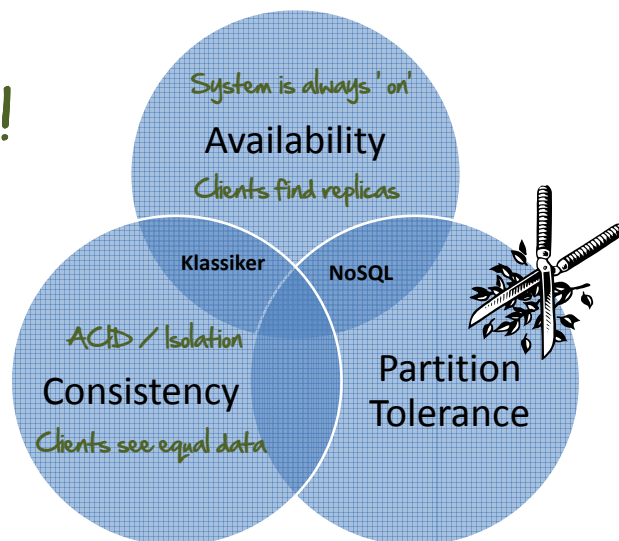
# Consistency Models



© Wilfried Springer NoSQL Rollercoaster

# CAP Theoreme

Pick 2!



„Don't throw **C** away so easy! It's complex.“



What you really have is:

1. Application errors \*
2. Repeatable DBMS errors \*
3. Unrepeatable DBMS errors
4. Operating System errors
5. Hardware failure in cluster
6. Network partition in local cluster
7. A disaster \*
8. WAN failure



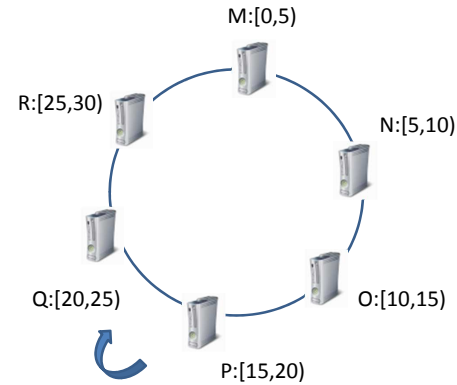
- 6 = Network Partition is rare
- 3,4,5,6 is mostly a Single Node
- Algorithms can help!

„give up P rather than sacrificing C. Use VoltDB or NimbusDB“



## Consistent Hashing

HASH	KNOTEN	REPLIKAT
2	M	N,O
8	N	O,P
10	O	P,Q
17	P	Q,R
22	Q	R,M
26	R	M,N





$$W \checkmark = 2 * W$$

$$R \checkmark = 1 * R$$



- ausfallsicher ✓
- leicht erweiterbar ✓
- gut verteilt / vnodes ✓

## MVCC

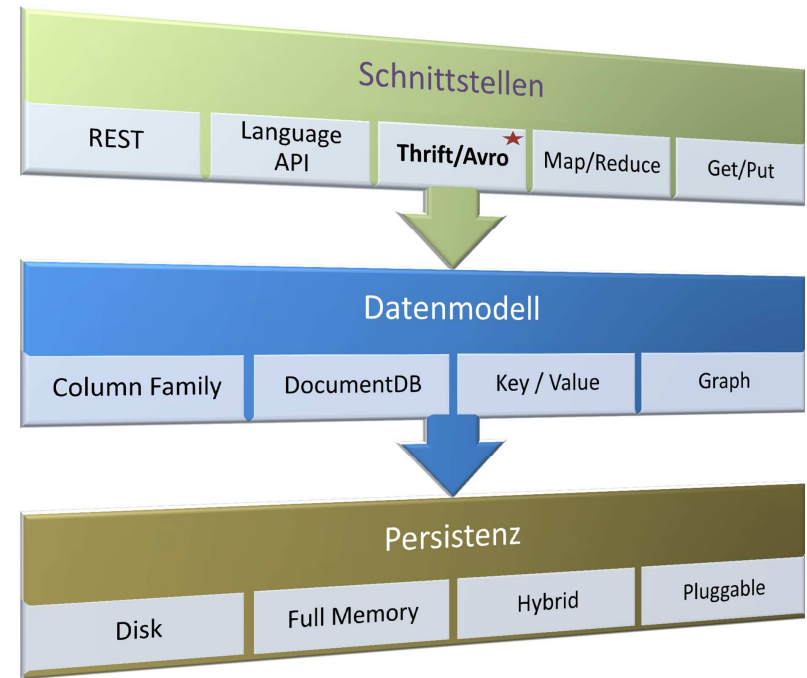
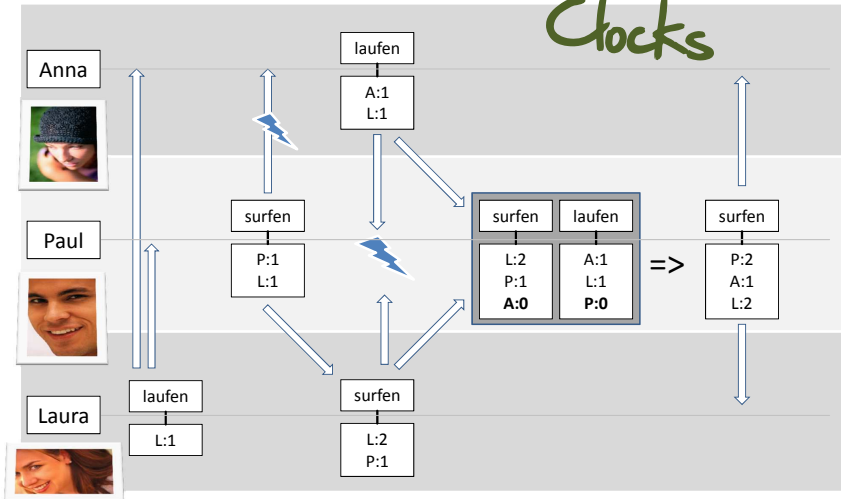
Multi Version Concurrency Control



pessimistisches Locking?



# Vector Clocks



you write a .proto file like this:

```
message Person {
  required int32 id = 1;
  required string name = 2;
  optional string email = 3;
}
```

Then you compile it with `protoc`, the protocol buffer compiler, to produce code in  
then, if you are using C++, you use that code like this:

```
Person person;
person.set_id(123);
person.set_name("Bob");
```

# Google Protocol Buffers



```
struct UserProfile {
  1: i32 uid,
  2: string name,
  3: string blurb
}
service UserStorage {
  void store(1: UserProfile user,
  UserProfile retrieve(1: i32 uid)
}
```



```
# Make an object
up = UserProfile(uid=1,
  name="Mark Slee",
  blurb="I'll find something to put here.")

# Talk to a server via TCP sockets, using a binary protocol
transport = TSocket.TSocket("localhost", 9090)
transport.open()
protocol = TBinaryProtocol.TBinaryProtocol(transport)

# Use the service we already defined
service = UserStorage.Client(protocol)
service.store(up)

# Retrieve something as well
up2 = service.retrieve(2)
```



JSON!

• Binary data transfer

• automatic RPC generation

• no code generation

• Client + Server tauschen Schema bei Änderung

# Daten-Modelle

Relationale DB	Server	Database	Table	Primary Key			
Key Value DB	Cluster	Keyspace		Key	Value		
Column Family DB	Cluster	Table / Keyspace	Column Family	Key	Column Name	Column Value	Super Column optional
Document DB	Cluster	Docspace		Doc Name	Doc Content		
GraphDB	Server	Nodespace	Nodes & Links				

Column Family

DocumentDBs

Key/ValueDBs

GraphDBs

andere

Voldemort, Chordless, Scalaris, Dynamo / Dynamite

db4o, Versant, Objectivity, Gemstone, Progress, Mark Logic, EMC Momentum, Tamino, GigaSpaces, Hazelcast, Terracotta ...

## Wide Column Stores / Column Families

```
KEY_1 => COLUMN_FAMILY_ONE => A => Value
=> B => Value
=> C => Value
=> COLUMN_FAMILY_TWO => D => Value
=> E => Value
KEY_2 => COLUMN_FAMILY_ONE => A => Value
=> B => Value
=> COLUMN_FAMILY_TWO => D => Value
=> E => Value
KEY_3 => COLUMN_FAMILY_ONE => A => Value
=> B => Value
=> C => Value
=> F => Value
=> G => Value
=> COLUMN_FAMILY_TWO => D => Value
=> E => Value
```

Keyspace #1:1:1:1:1	CF:Users
"Jim"	"reactivity": "2"
"Joe"	"reactivity": "9"
CF:Heroes	"nick_name": "Little Joe"
"Tally"	"color": "white"
CF:Weapons	"owner": "Jim"
"D31"	"ammo": "15"
CF:Boats	"owner": "Joe"
"H3H-3000"	"player_1": "Jim"
CF:Messages	"player_2": "Joe"
"H3H-3000-L"	"from": "Joe"
"H3H-3000-L"	"to": "Jim"
"H3H-3000-L"	"winner": "Joe"
"H3H-3000-L"	"text": "Loooooooo!"

Kunden-ID	Vorname	Nachname	Straße	Stadt	Land	PLZ	Telefon	E-Mail
123	Robert	Schmidt	Hauptstr. 123	Düsseldorf	Deutschland	40210	222-333-4444	
456	Jens	Hansen	Bahnhofstraße 456	Köln	Deutschland	50441	333-444-5555	
789	Doris	Thomas	Schillerstraße 789	München	Deutschland	80331	444-555-6666	dthomas@xyz.com




- + Skalierung = new node
- + Community
- + API
- Replikation
- Aufsetzen, Optimierung, Wartung



- + Skalierung = new node
- + Replikation
- + Konfiguration (r, w)
- Dokumentation
- Abfragen
- (storage-conf.xml)



- + stressfreie SaaS Lösung
- + transparent scaling
- UTF-8 String
- Daten liegen bei Amazon
- kein tuning / config




# Document DBs



# mongoDB

2. Runde  
+= 6,5 Mio \$

- + memory mapped, indexes, queries, marketing
- durability, single instance design



NoSQL Divergenz: CouchDB = NoNoSQL

NoSQL Konvergenz:

Liquid Data at the bottom (JSON)

besser suchbar

Easy to run,  
develop, manage

# K/V-Stores



Data Structure Server ->

- + sehr schnell > 100.000 / sek
- + konfigurierbarer Disc sync
- + API für eigene Anbindung
- + einfache Replikation
- + hash, list, set, sorted set, messages
- + Installation

UNIX: 38 sek

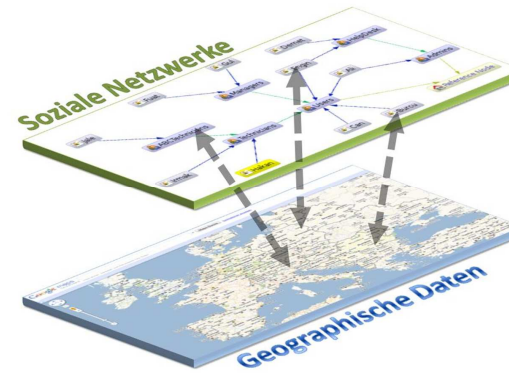
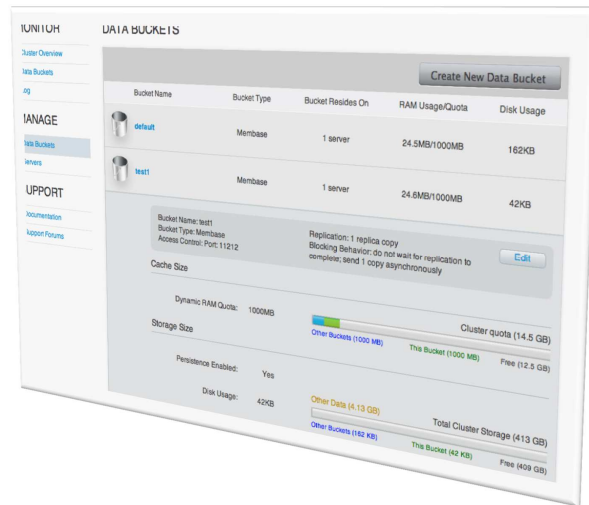
Windows: 18 sek

- noch nicht skalierbar (2.\*)






memcached API



GraphDBs

Property Graph



Neo4j, Sones, HyperGraphDB, InfiniteGraph, InfoGrid, Dex, VertexDB, Filament, OrientDB

x

DataModel, Query Methods, Languages, License, Protocol



**Infinite Graph**

- URL <http://www.infinitegraph.com>
- Goals **Very large graph databases**
- Transactions Distributed transactions
- Repl./Scaling Objectivity/DB
- Persistency Objectivity/DB

```

// Note: Website(), Tag() extend BaseVertex; Edge() extends BaseEdge
Website xkcd = new Website("xkcd", "http://xkcd.com/");
Tag good = new Tag("good");
_xGDB.addVertex(xkcd);
_xkcd.addEdge(new Edge(), good, EdgeKind.BIDIRECTIONAL);
  
```



Key / ValueDB + DocumentDB + ObjectDB + GraphDB +schneller...

> 220 DBs



Start-Ups 😊

etablierte Unternehmen ☹️

1. Data
2. Transactions
3. Performance
4. Queries
5. Architecture
6. other Non-Functional Requirements

### Analyse your Data

Domain-Data, Log-Data, Event-Data, Message-Data, critical Data, Business-Data, Meta-Data, temp Data, Session-Data, Geo Data, etc.

### Data- / Storage-Model:

relational, column-o, doc-alike, graphs, objects, etc.

### What Types / Type-System?

Data-Navigation, Data Amount, Data Complexity (Deep XML?)

**ACID vs. BASE vs. Mixture?**  
**CAP decisions**

### Performance Dimension Analysis

Latency, Request behaviour, Throughput

**Scale-Up vs Scale-Out**

### Query Requirements

Typical queries, Tools, Ad-Hoc Queries, SQL / LINQ needed, Map/Reduce? ...

### Distribution Architecture

local, parallel, distributed / grid, service, cloud, mobile, p2p, ...

### Data Access Patterns

read / write distribution, random / sequential, Access Design Patterns

### Non Functional Requirements:

Replication, Refactoring Frequency, DB-Support, Qualification / simplicity, Company restrictions, DB diversity (allowed?), Security, Safety / Backup & Restore, Crash Resistance, Licence...



# NoSQL

## Lessons learned

## Think outside the MyOracleSql Box ✓

- VoltDB
- Vertica
- GenieDB
- MonetDB
- Hadoop++
- ...

## RAM + SSD rocks

RethinkDB



Gustavo Alonso

Lot's of >1PT RAM DBs in CA!

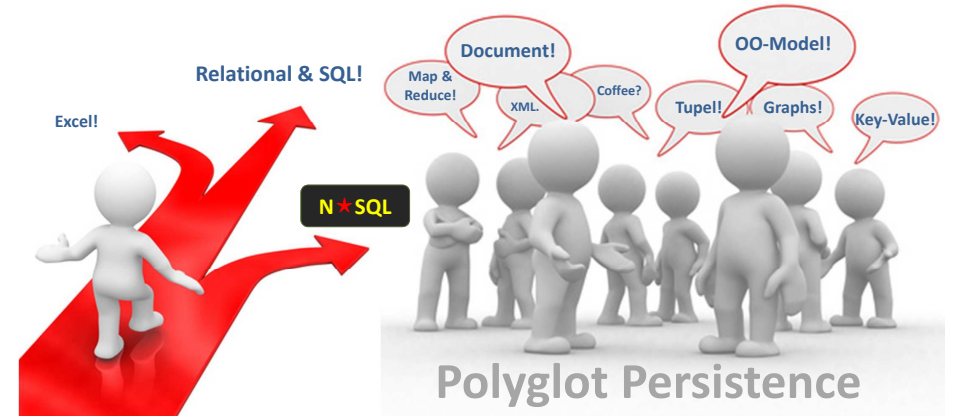
Queries in FPGA!



Tonnen cleverer *hybrid* Lösungen da!



The world is *diverse!* Act accordingly!



Thanks for listening!

<http://edlich.de>

Diskussion!

