


# HOW TO WRITE A CV



## THE 10 WHYS OF NOSQL

NOSQL ROADSHOW  
MARCH 2012 COPENHAGEN



[nosqlberlin.de](http://nosqlberlin.de)  
[nosqlfrankfurt.de](http://nosqlfrankfurt.de)  
[nosql powerdays](http://nosqlpowerdays.com)



[HTTP://NOSQL-DATABASE.ORG](http://nosql-database.org)



## Your Ultimate Guide to the Non - Relational Universe!

[the **biggest** nosql link [Archive](#) in the web]  
...never miss a [conceptual](#) article again...  
News Feed covering all changes [here](#) !

**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. [based on 5 sources, 10 constructive feedback emails (thanks!) and 1 insulting comment. Agree / Disagree? [Tell](#) me so! ]

### LIST OF NOSQL DATABASES [beta]

#### Core NoSQL Systems:

##### Wide Column Store / Column Families

**Hadoop / HBase:** API: **Java / any writer**, Protocol: **any write call**, Query Method: **MapReduce Java / any exec**, Replication: **HDFS Replication**, Written in: **Java**, Concurrency: ?, Misc: **Links:** 3 Books [\[1, 2, 3\]](#)

**Cassandra:** API: **many Thrift » languages**, Protocol: ?, Query Method: **MapReduce**, Replicaton: , Written in: **Java**, Concurrency: **eventually consistent**, Misc: like "Big-Table on Amazon Dynamo alike", initiated by Facebook, Slides [»](#), Clients [»](#), Installation [»](#)

**Hypertable:** API: **Thrift** (Java, PHP, Perl, Python, Ruby, etc.), Protocol: **Thrift**, Query Method: **HQL, native Thrift API**, Replication: **HDFS Replication**, Concurrency: **MVCC**, Consistency Model: **Fully consistent** Misc: High performance C++ implementation of Google's Bigtable. Commercial support [»](#)

**Cloudera:** Professional Software & Services based on Hadoop.

**Amazon SimpleDB:** Misc: not open source / part of AWS, Book [»](#)

[**SciDB:** **Array** Data Model for Scientists, paper [»](#), poster [»](#), HiScaBlog [»](#)]  
[OpenNeptune, Qbase, KDI]:

### EVENTS

17th JAN **Membase** Seattle [»](#)  
5th JAN **Membase** San Diego  
3rd DEC **Mongo** Mountain View  
20th NOV **NoSQL Dundee** [»](#)  
All past NoSQL Conferences [»](#)  
register your event here! [»](#)

**NoSQL Summer in 27 cities** [»](#)

Conceptual quality articles,  
links, research papers, etc.

## NoSQL



## ARCHIVE

Announcing the availability of  
the **worlds first NoSQL Book!**

### NoSQL FORUMS [3]

- \* Global NOSQL Forum [»](#)
- \* Forum Berlin [»](#)
- \* Forum France [»](#)
- \* Forum Japan [»](#)

### GREAT NoSQL NEWS FEEDS

- \* MyNoSQL by Alex P [»](#)  
[He is definitely bigger...]
- \* On Twitter: nosqlupdate [»](#)
- \* HighScalability Blog [»](#)

### BREAKING NoSQL NEWS

**2011**



**The NoSQL  
Year!**

**Developer Skills**

- 1. HTML5**
- 2. MongoDB**
- 3. iOS**
- 4. Android**
- 5. Mobile app**
- 6. Puppet**
- 7. Hadoop**
- 8. jQuery**
- 9. PaaS**
- 10. Social Media**







**Redis Cluster soon**

**Hadoop + Cassandra 1.0**

**MongoDB 2.0 + MMS**


**Riak 1.1**

**>1000 hosting Services**

**NoSQL is PaaS Standard!**









Oracle's Thomas Kurian says no to SQL. Sorta. (Photo: Oracle/Flickr)

**white paper:**  
**„debunking the (NoSQL) hype“**

**“The NoSQL databases are beginning to feel like an **ice cream** store that entices you with a **new flavor of the month,**” the white paper read. “[But] you shouldn’t get too attached to any of the flavors because it **may not be around for too long.**”**

# ORACLE®


## Oracle NoSQL Database

### Oracle Big Data Appliance

New | Integrated "Big Data" Platform

- Engineered System to Acquire, Organize, Analyze "Big Data"
- Super-Fast – Massively Parallel Processing & Loading into Oracle
- Complete – Software engineered with Hardware


ORACLE




**ConsHash**  
**config ACID**  
**no single PF**  
**DataC Replication**  
**Top Admin**  
**“BerkleyDB reloaded”**



Hadoop + Manager







**user defined functions in C++ & Java**  
**⇒ 10x faster than SQL or Stored Procs**

**UDF connector for Hadoop**



**C++ APIs for Map Reduce**





**Greenplum, Pervasive  
and 100 others too...**



*WHY ?*



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

*NoSQL is specialization!*

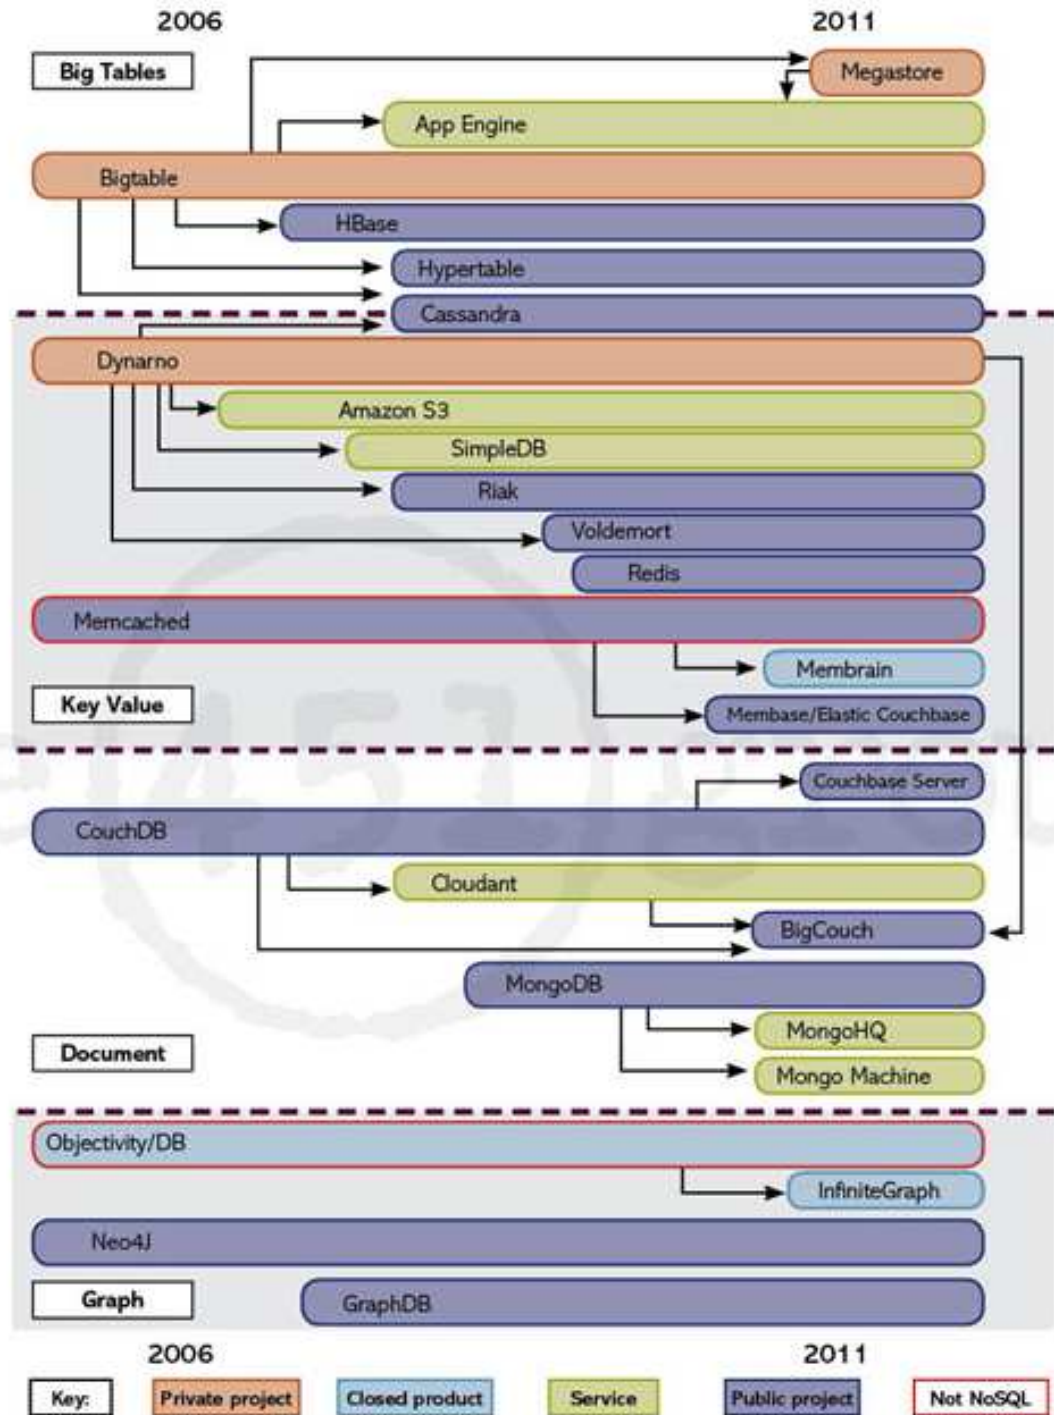
**JAP  N**



*interaction vs. transaction is 1000 : 1*

*WHY?*

*Needs!*



*80s / 90s*



- *Megabytes, Gigabytes*
- *Scale Out worked fine*
- *Golden Age for Oracle, DB2, MS, MySQL...*



*WHY #1*

*Big Data!*

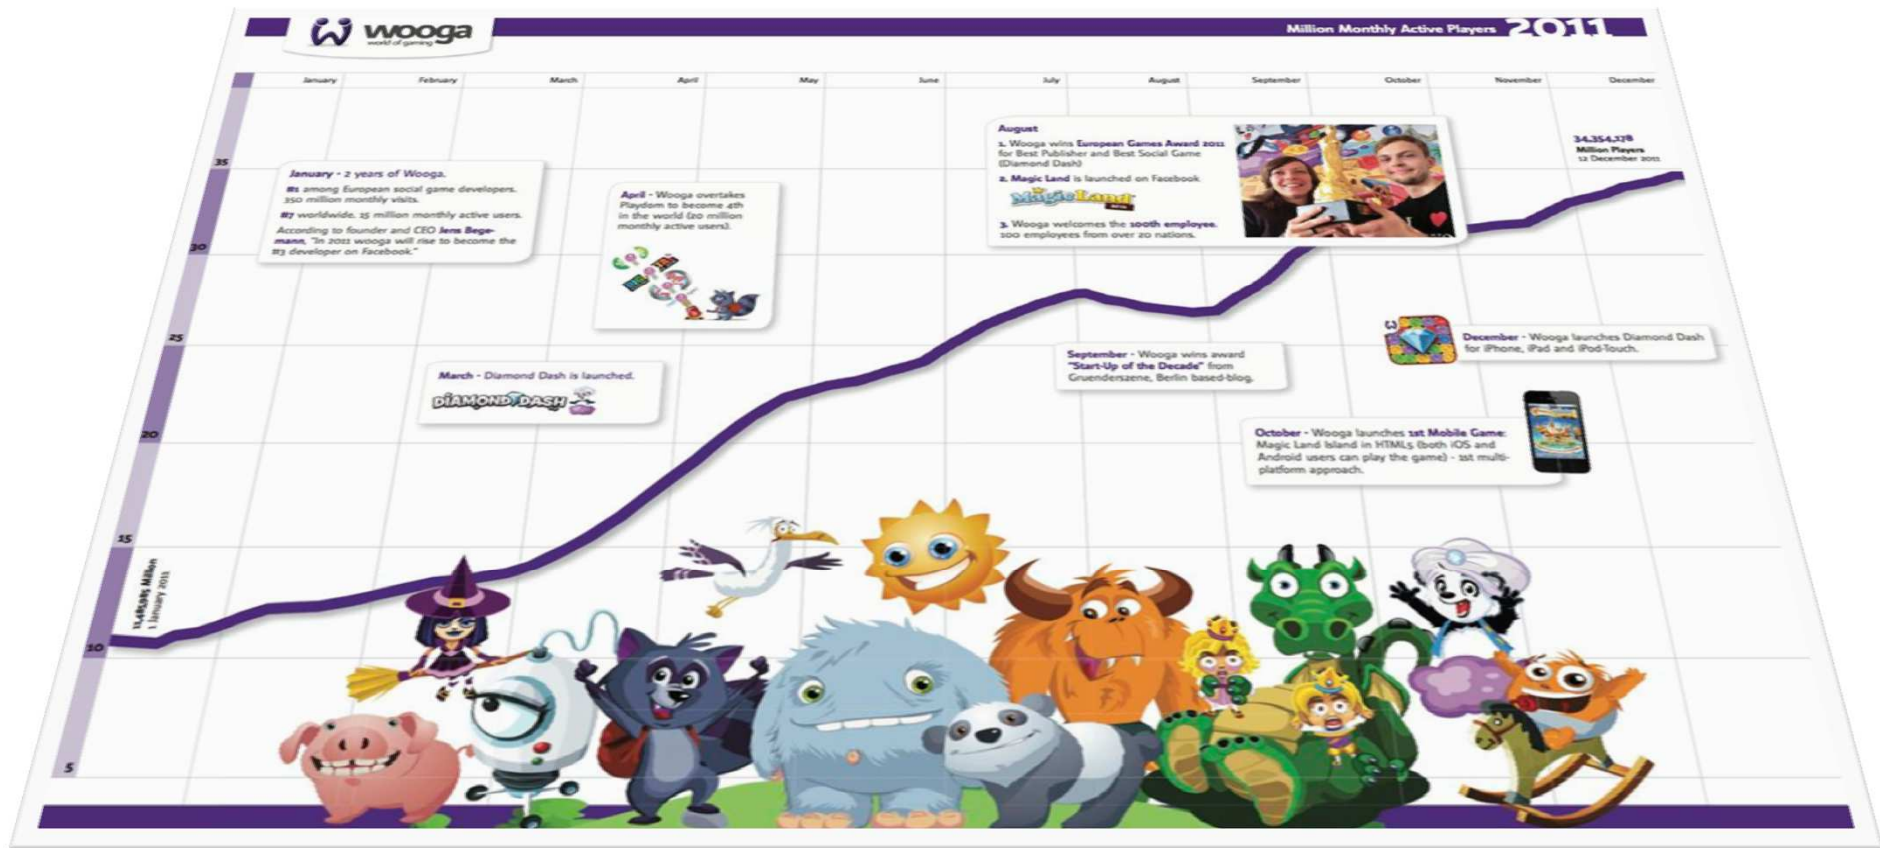
*Terrabyte-, Petabyte-Age  
Simple Scale-Out needed*



*I am not Amazon?!*

# WHY #2

# Massive Write Performance! Write Availability!



# WHY #3

# Sub Second Responses!

## cyberport

Apple & Zubehör | Notebook & Tablet | PC & Zubehör | Smartphone & Foto | TV & Audio | Haushalt

Bitte Suchbegriff eingeben ... **SUCHEN** [Zum Geschäftskunden-Shop](#)

> Notebook und Tablet > Notebooks

### 7 Artikel gefunden

**Marken:**  Acer  Alienware  Asus  Dell  Fujitsu  HP  Lenovo

**Preis:** € 1502 - € 2910

**Displaygröße:** 17 Zoll - 18 Zoll

**Leistung (0=min. - 5=max.):** 5 - 5

Kauftipp (1)  Neuheit (1)  
 Preishit (1)  Testsieger (0)

Verfügbarkeit für Versand ([Ändern](#)):  
 Auf Lager, innerhalb 24 h versandfertig (5)

**weitere Filter ausblenden**

**Arbeitsspeicher:** 16 GB - 16 GB **Festplattenkapazität:** 256 GB - 2 TB **Max. Akkulaufzeit:** 1,5 Stunden - 5,5 Stunden

**Max. Prozessortakt:** 2,9 GHz - 3,6 GHz **Gewicht:** 3,5 kg - 3,9 kg **Garantie:** 2 Jahre - 2 Jahre

Prozessortyp:  Grafiktyp:  Festplattentyp:   USB 3.0 (8)

Betriebssystem:  Displayoberfläche:  Farbe:   UMTS (0)

[Suchergebnis mitteln](#) [alle Filter aufheben](#)

Artikel pro Seite: [12](#) | [20](#) | [40](#) Sortieren nach:  Darstellung:

**acer** **Acer Aspire 8951G-2671687Wikk Ethos Big Mama II i7-2670QM SSD GT555M WiDi** PREIS TIP KAUFTIPP

**Versand-Preis** **Abholung vor Ort**

**€ 1.549,00\*** ~~€ 1.649,00\*\*~~

**Versandkosten:** **Gratisversand möglich**

**Auf Lager, innerhalb 24 h versandfertig**

**Finanzierung:** ab € 33,60 / Monat


[Angebot anfordern](#)

- ★★★★★
- Intel® Core™ i7-2670QM Prozessor (bis 3,1 GHz), Quad-Core
- 46,7 cm (18") Full HD 16:9 LED Display (glänzend), Webcam
- 16 GB RAM, 120 GB SSD + 750 GB, Blu-ray Disc Brenner, Fingerprint

# WHY #4

# Fast KV Access


amazon.com<sup>®</sup>



>100 Mio PVs  
**YOU\*PORN**

# WHY #5

## Flexible Schema (Migration) Flexible Datatypes

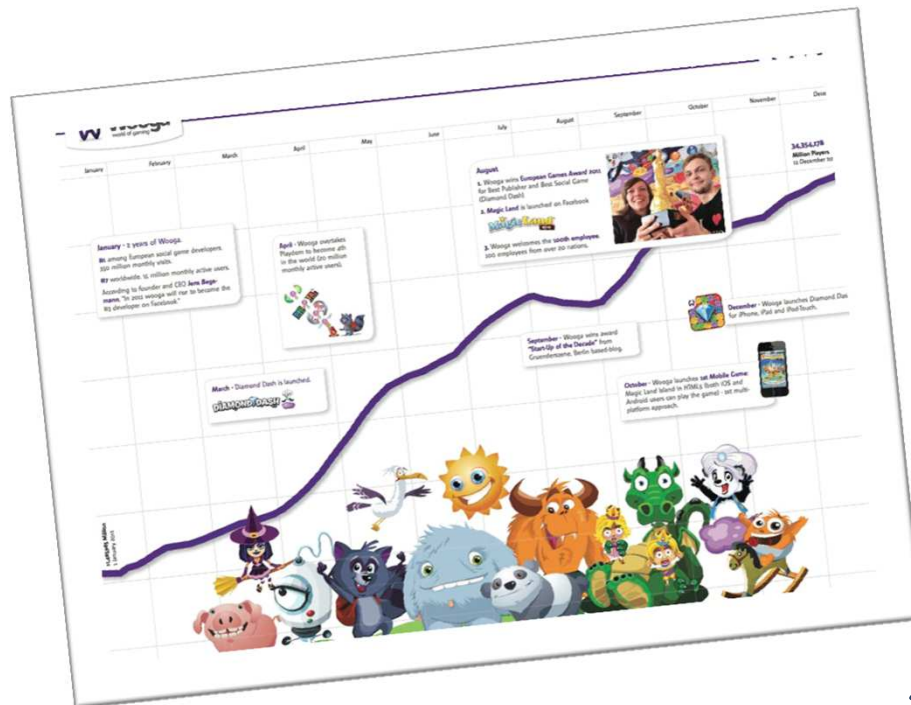


ALTER TABLE => Shut Down for 3 hours in "the night?!"



# WHY #6 *easier -maintainability*

*-administration and operations*



**1 Game =  
2 REDIS instances  
+ 2 Replicas**



*WHY #7*

*no single point of failure*

*#8*

*up 24/7*

*Google 40.000\$ /Min*

*Amazon 35.000\$ /Min*

*Yahoo 10.000\$ /Min*

*in 2008 !!*

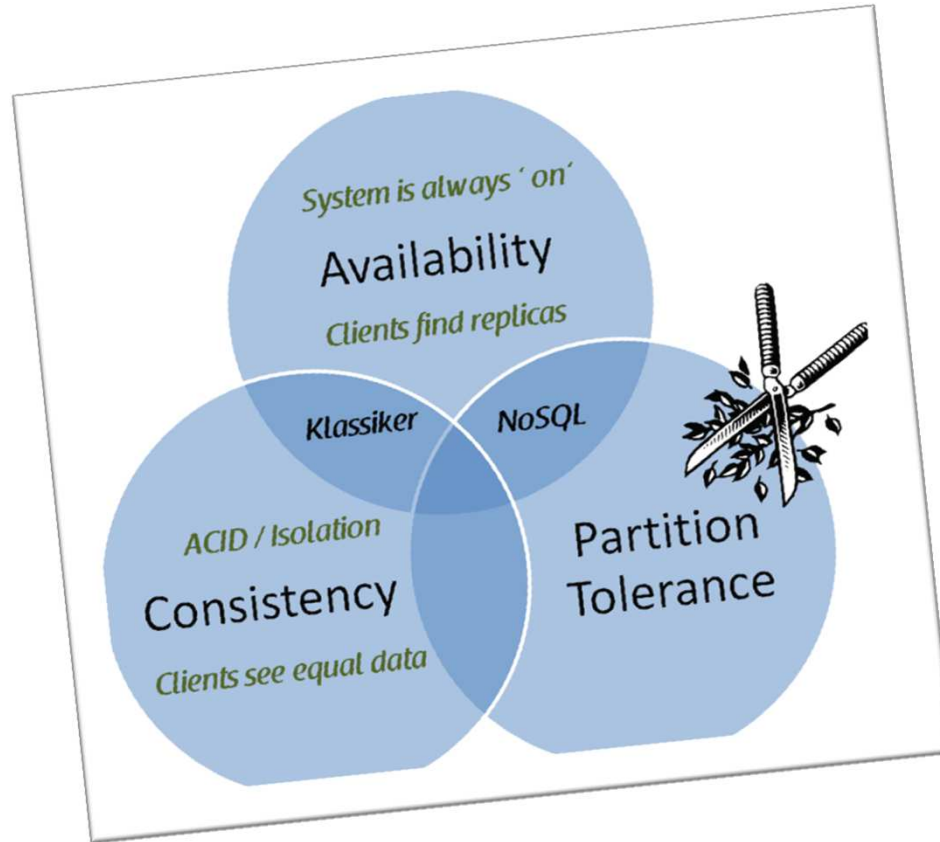
**WHY #9**

*Programmer ease of use*

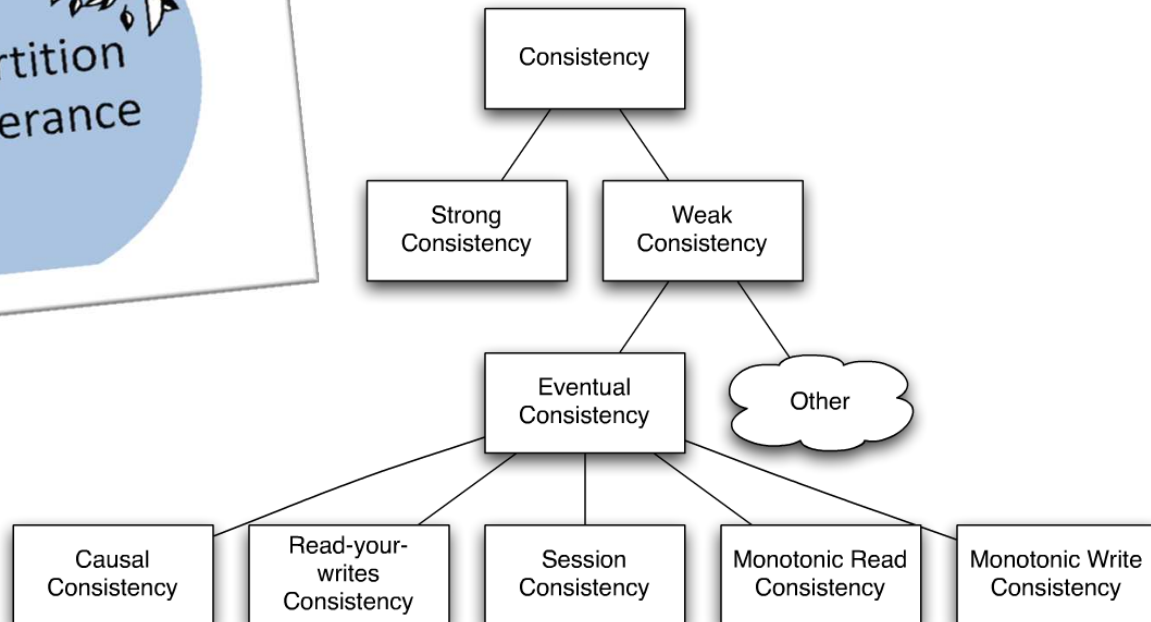


# WHY #10

# Changing CAP needs



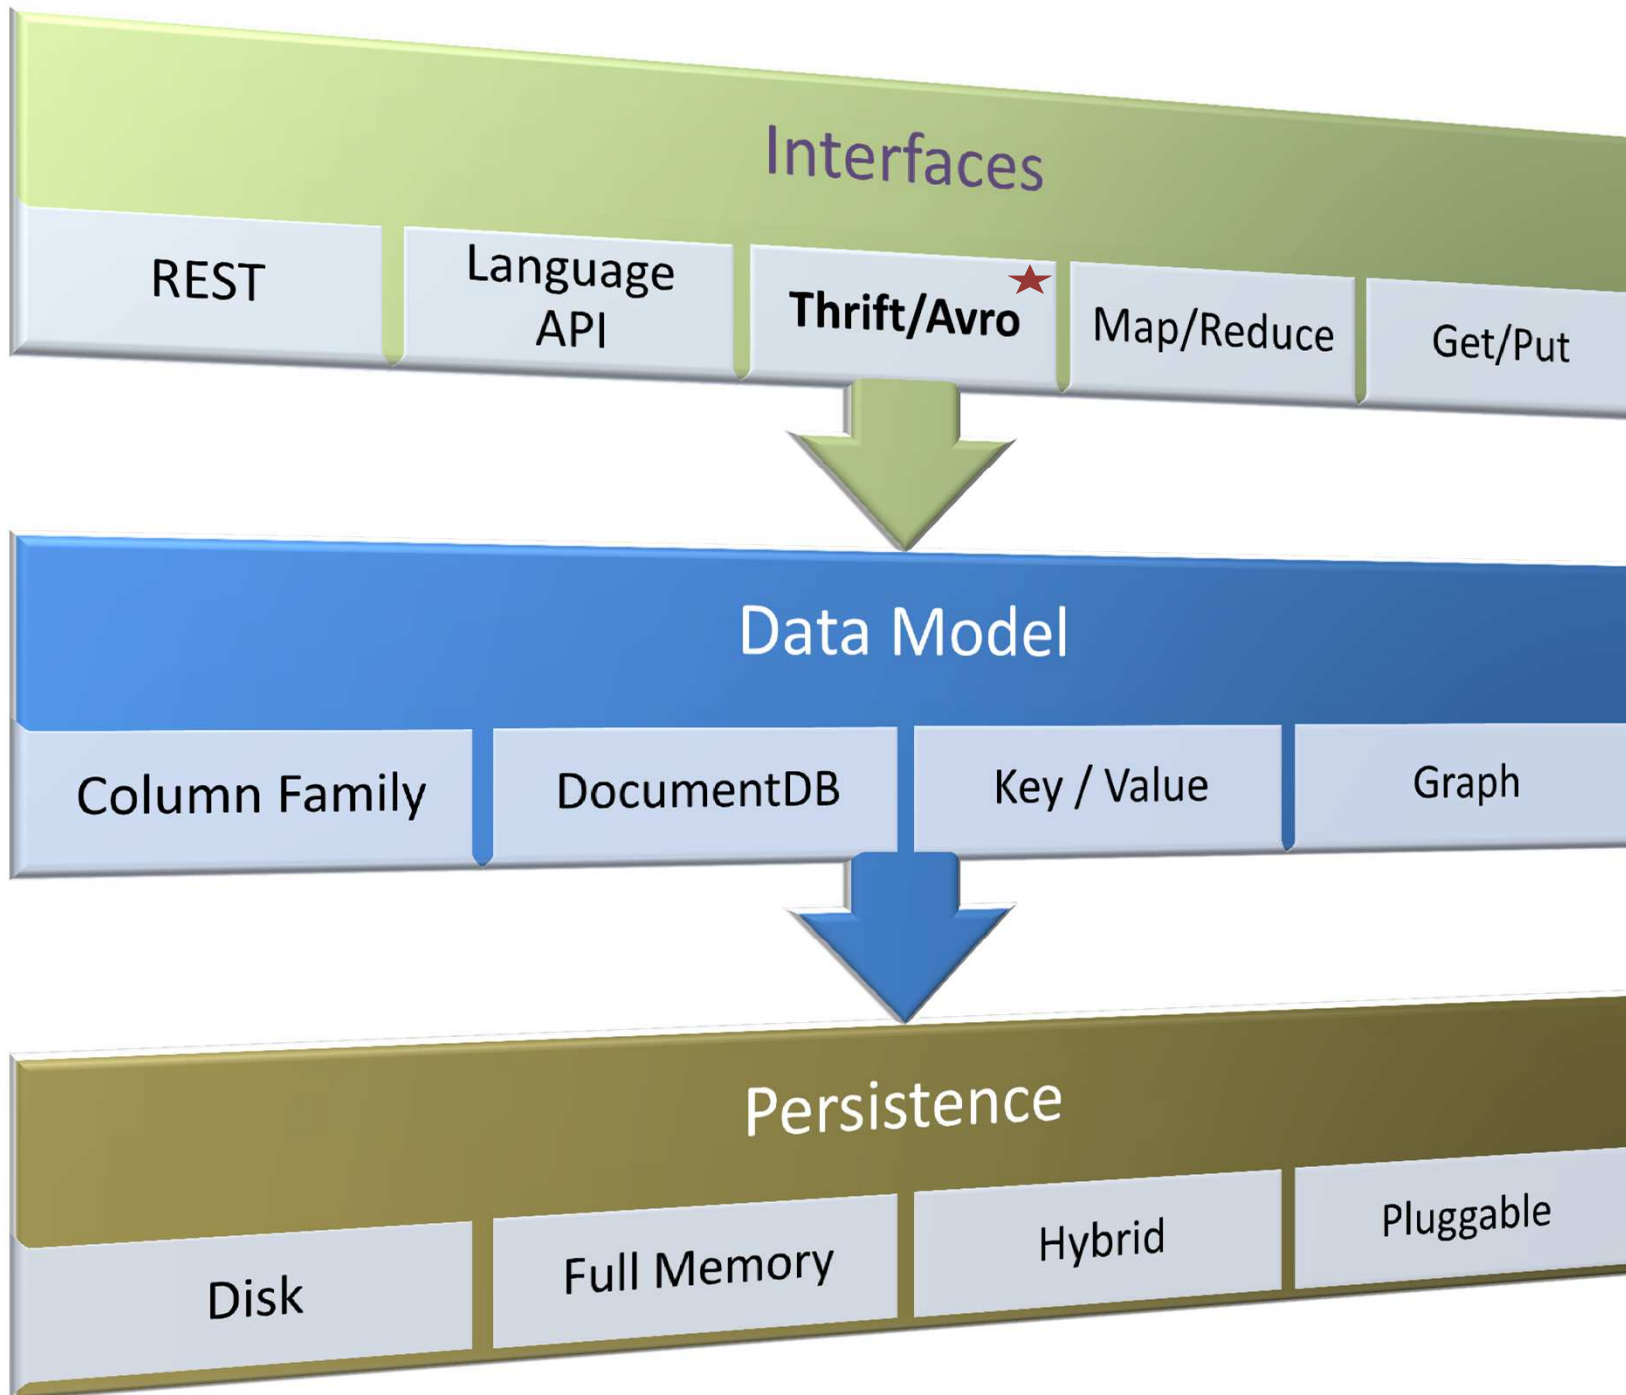
## CONSISTENCY MODELS



next talk!

<b>Relational DB</b>	Server	Database	Table	Primary Key			
<b>Key Value DB</b>	Cluster	Keyspace		Key	Value		
<b>Column Family DB</b>	Cluster	Table / Keyspace	Column Family	Key	Column Name	Column Value	Super Column optional
<b>Document DB</b>	Cluster	Docspace		Doc Name	Doc Content		
<b>GraphDB</b>	Server	Nospace	Nodes & Links				

# DATA MODEL





*Column Family*

---



*DocumentDBs*


---



*Key/ValueDBs*

---

*Voldemort, Chordless, Scalaris, Dynamo / Dynamite*



*HYPERGRAPHDB*

*GraphDBs*

---

*others*

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



**> 220 DBs**




*tough consulting...*



- + Scaling= new node
- + Community
- + API
- Replication
- SetUp, Optimize, Management



- + Scaling= new node
- + Replication
- + Configuration (r, w)
- Documentation
- Query
- (storage-conf.xml)



	HBase	Cassand.	DynDB	Mongo	CouchBS	Riak	Redis	ES
schema free	Chunks	TAB	TAB	JSON	JSON	JSON	K/V	JSON
,realtime'								
performance	mass data			100k			100k	
scaling							V3	
ring / shardRepl	chunks	ring	invisible	ShRe	ring	ring	repl	ShRe
self tuning		In progress	SaaS	hard				
prod / tools			why?	MMS			Why?	
aggregations						?	DSS	
queries	Hive,Pig	CQL?					DSS	
full text s		?					DSS	
filemanagem								
community								
APIs				++	-		++	REST
support			why?					
docs				++	+/-			

misc

Geo,  
StoredProcs

Geo

persist  
config

limited  
sharding

persist  
config

Performance  
Transactions  
Queries  
Architecture  
Data

non-funct-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 Komplexity (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:**

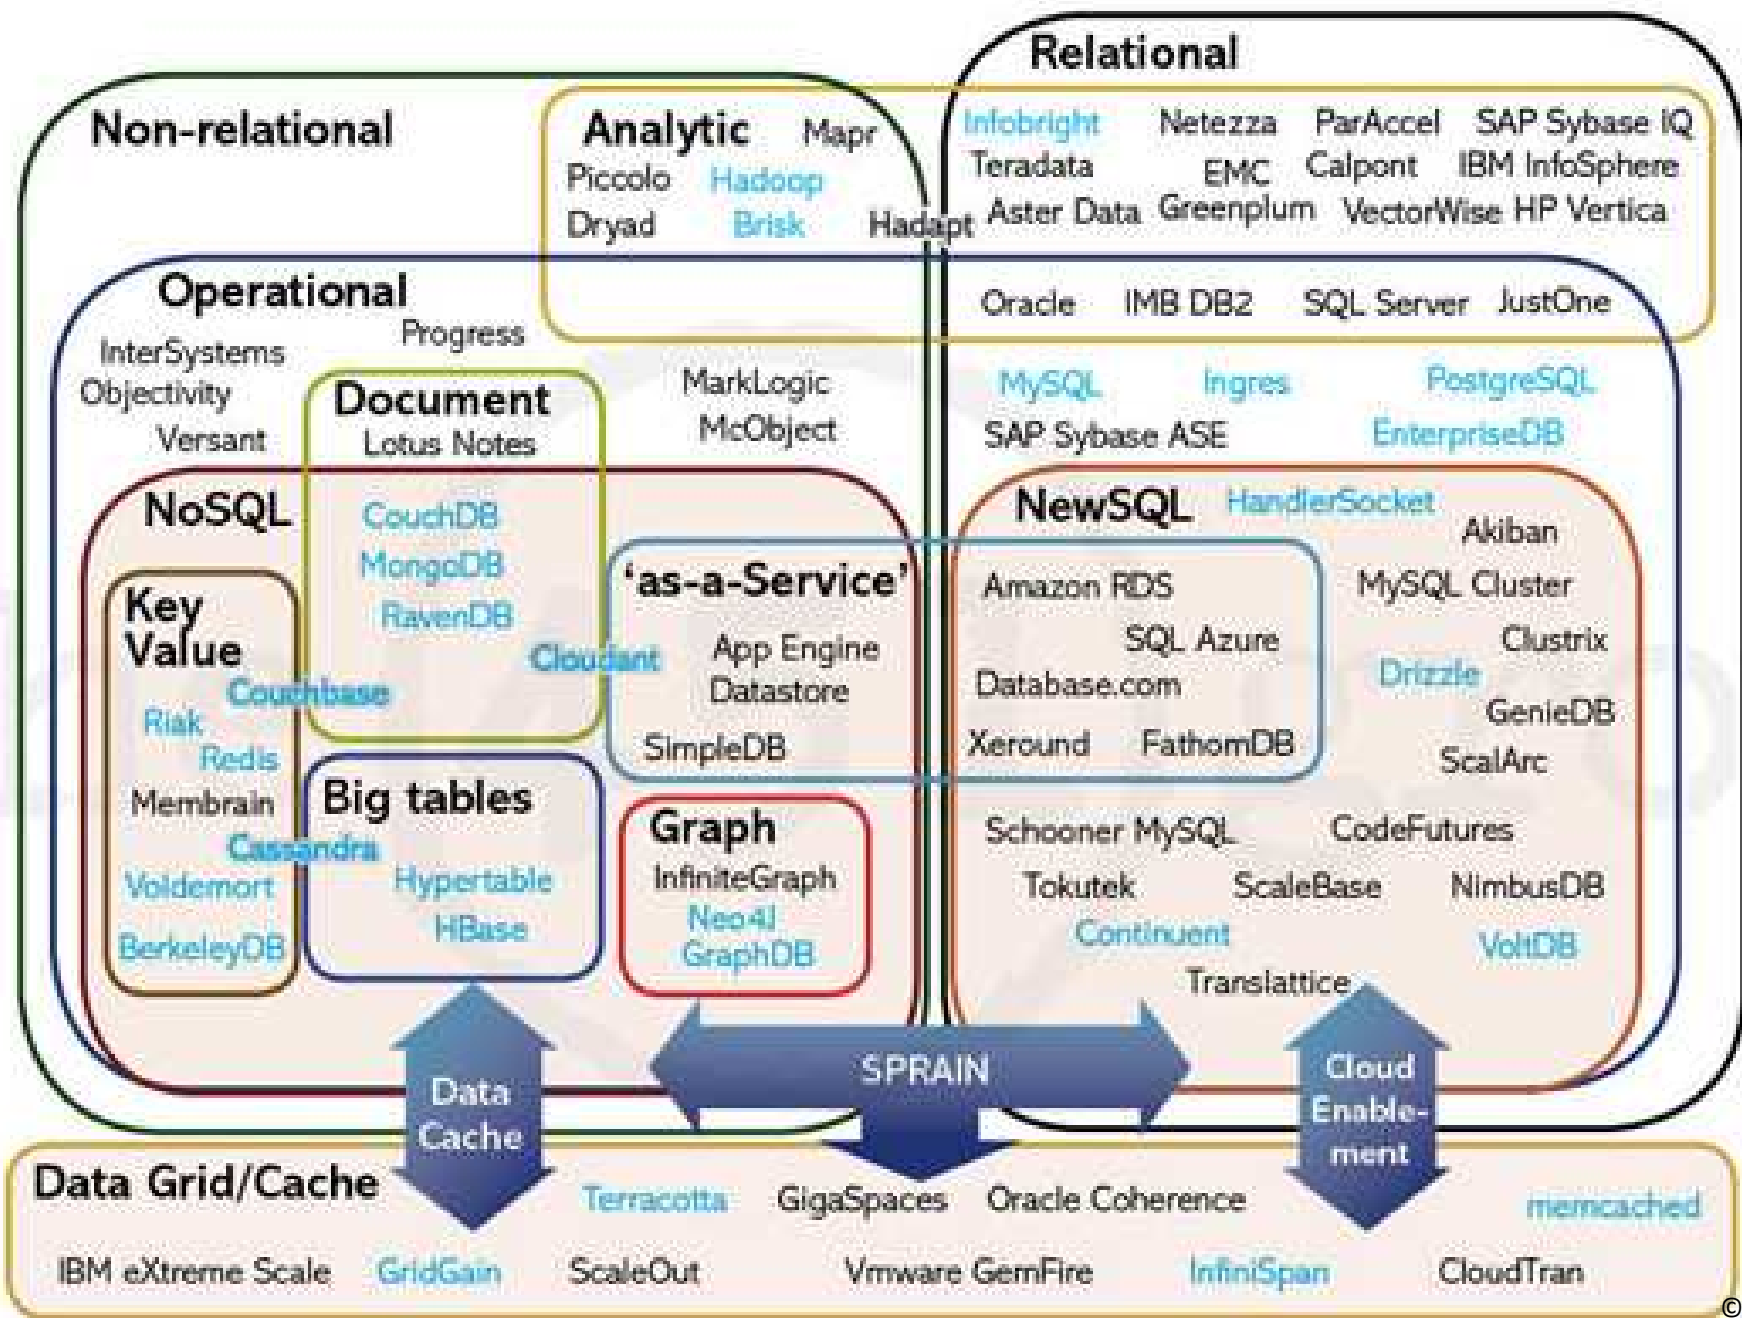
*Replication, Refactoring Frequency, DB-Support, Qualification / simplicity, Company restrictions, DB diversity (allowed?), Security, Safety / Backup & Restore, Crash Resistance, Licence...*





MoreSQL 

**strikes back**



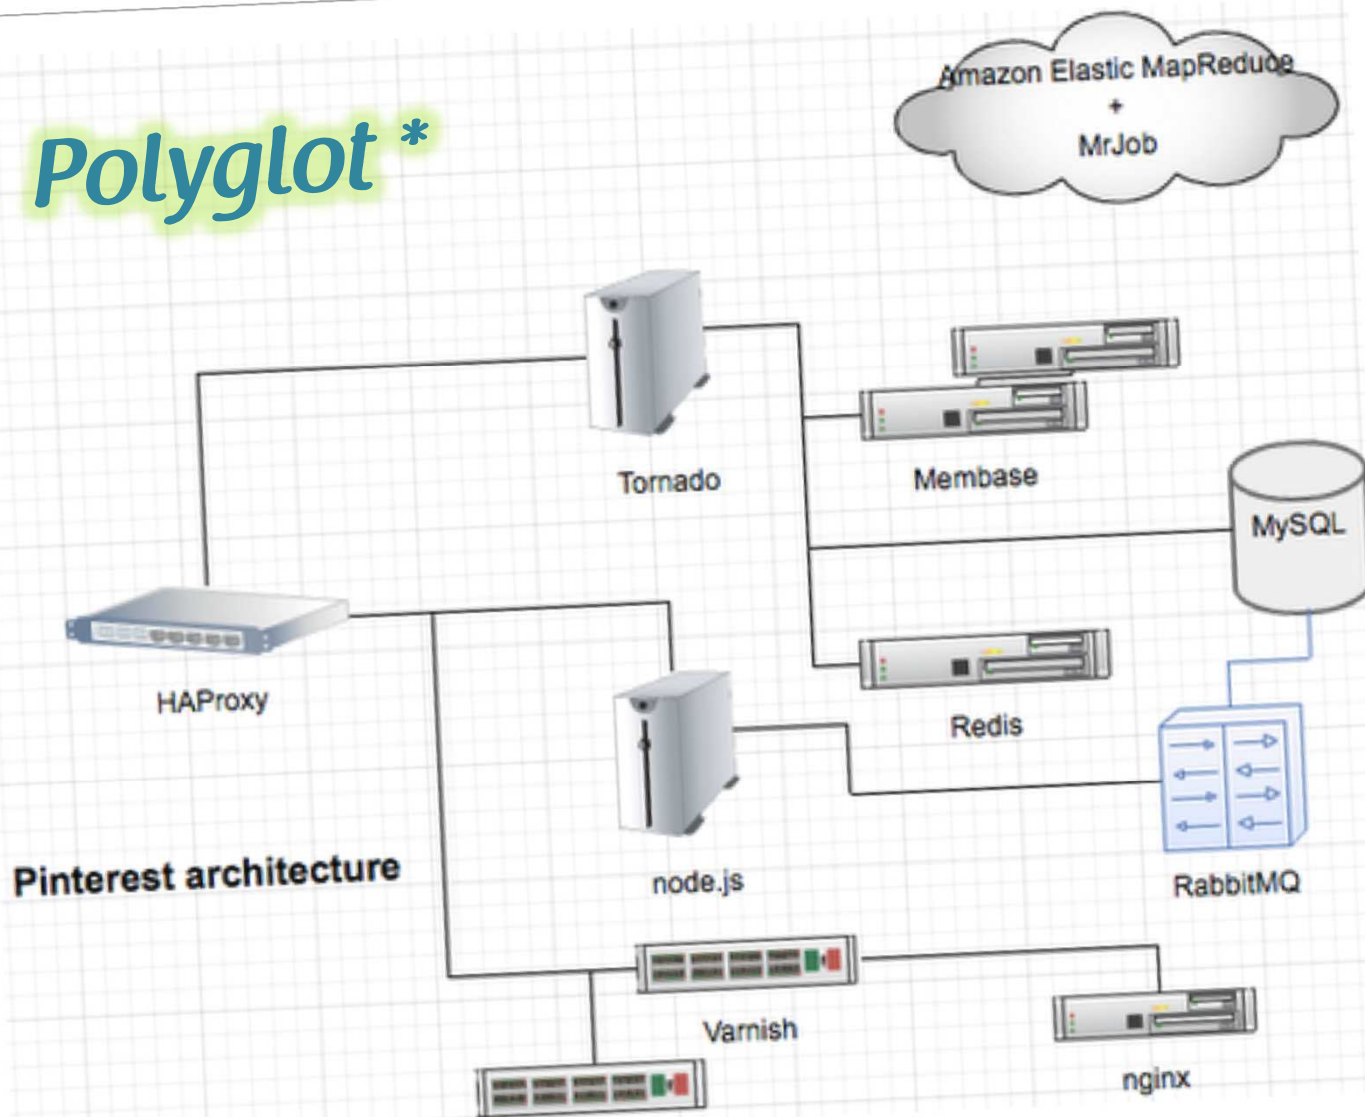


**Check **hybrid** solutions!**

**easier & better than memcache + RDBMS**



# Polyglot\*



# Conclusion #1

**There is no  
“one perfect solution”**

**Check hybrid solutions  
and NewSQL DBs too!**

# Conclusion #2



# Conclusion #3

why



***THANKS!***

***Q&A***

***please contact  
edlich.de***