NoSQL: An Architect's Perspective

Eberhard Wolff
Architecture and Technology Manager, adesso AG
@ewolff

adesso

Which Database Shall I use?

Or: The NoSQL Game

DBA vs. Architect

NoSQL: An Architect's Perspective

Eberhard Wolff
Architecture and Technology Manager, adesso AG
@ewolff

adesso

Which Database Shall I use?

Or: The NoSQL Game

DBA vs. Architect

NoSQL: An Architect's Perspective

Eberhard Wolff

Architecture and Technology Manager, adesso AG

@ewolff

adesso

ich Database

Or: Th

Me

- Architecture & Technology Manager at adesso
- adesso is a leading IT consultancy in the German speaking region
- Speaker
- Author (e.g. first German Spring book)
- Blog: http://ewolff.com
- Twitter: @ewolff

Back in the Days....

?

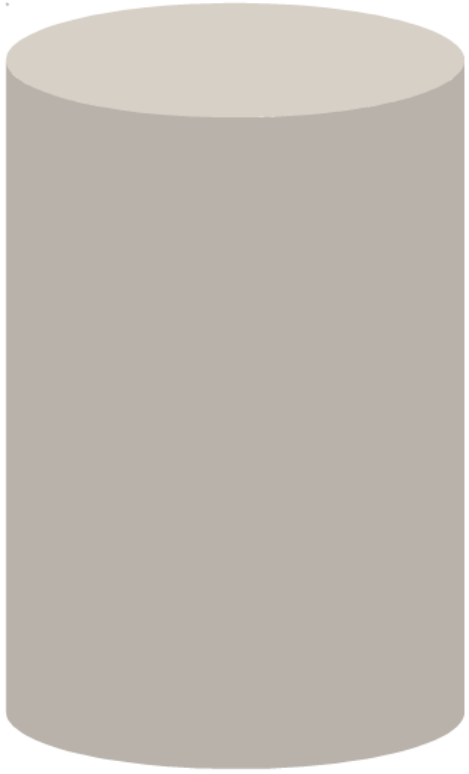

RDBMS!

NoSQL Is All About the Persistence Question

?

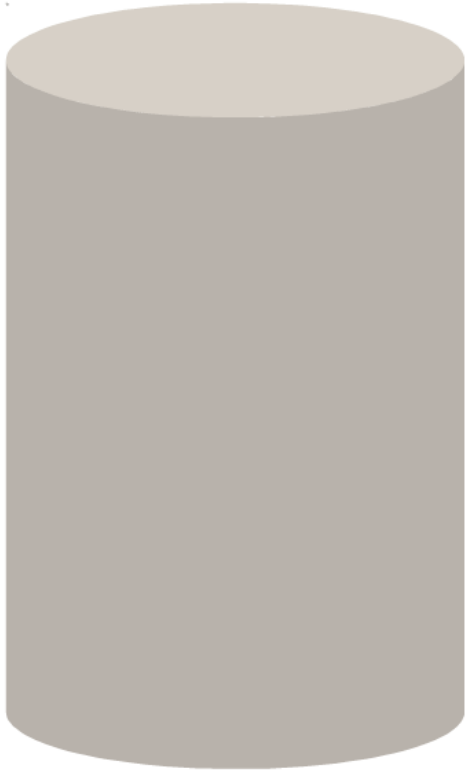

Eberhard Wolff

Architecture and Technology M

@ewolff

- Architecture & Technology Manager at adesso
- adesso is a leading IT consultancy in the German speaking region
- Speaker
- Author (e.g. first German Spring book)
- Blog: http://ewolff.com
- Twitter: @ewolff

RDBMS!

Which Database Shall I use?

DBA vs.

adesso

Or: The NoSQL Game

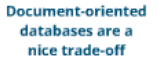

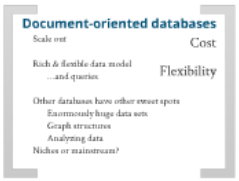

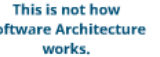

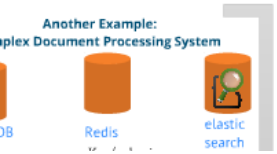

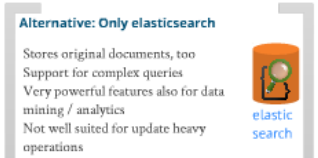

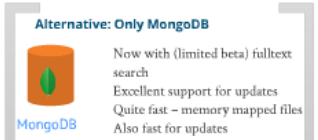

NoSQL Benefits

Costs

- Scale out instead of Scale Up
- Cheap Hardware
- Usually Open Source

Ops

Dev

Flexibility

- Schema in code not in database
- Easier to upgrade schema
- Easier to handle heterogeneous data

No Object/relational impedance mismatch

- NoSQL database are more OO like

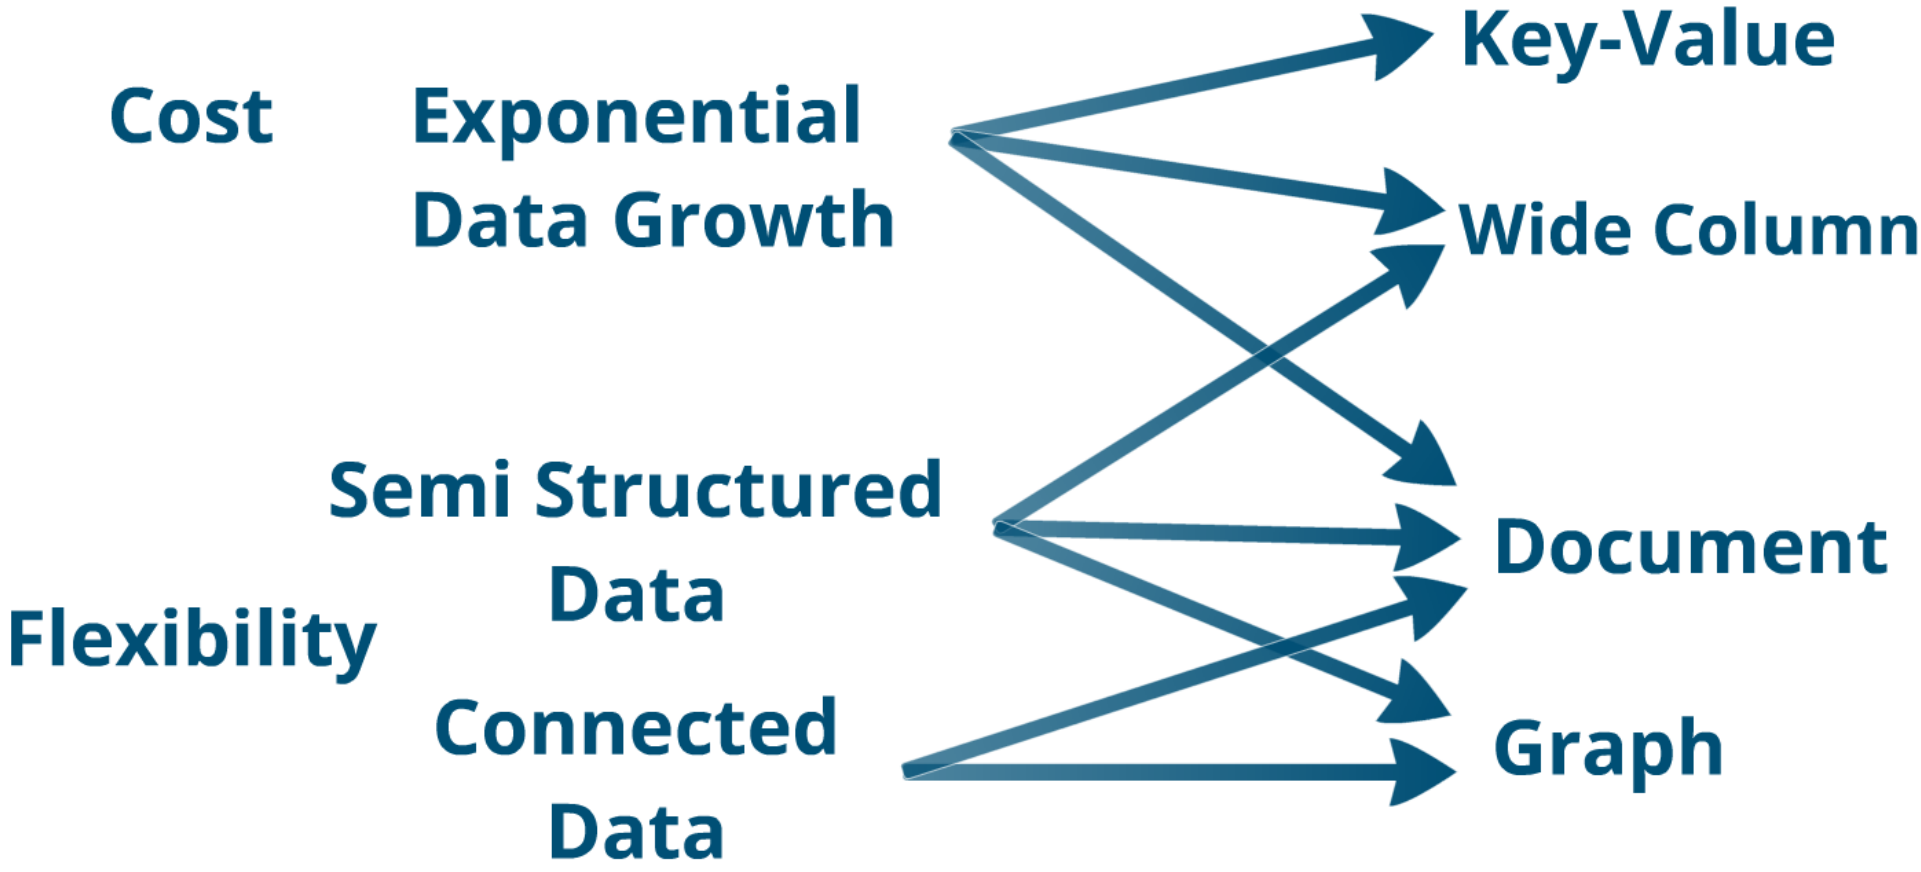
Cost
Exponential Data Growth
Semi Structured Data
Connected Data
Flexibility
Key-Value
Wide Column
Document
Graph

Document-oriented databases are a nice trade-off

Document-oriented databases

Scale out

Cost

Rich & flexible data model

…and queries

Flexibility

Other databases have other sweet spots

- Enormously huge data sets
- Graph structures
- Analyzing data

Niches or mainstream?

The NoSQL Game

Needs tran tions
& report Da fit well in
tables.

0

Financial Data

1000

Com plex ocu e -lik d
stru ure and or plex u ies

Product Catalog

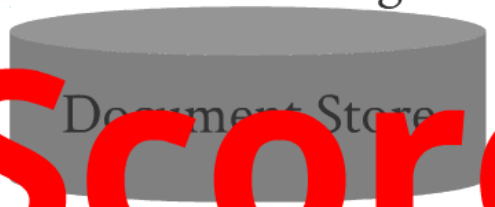

High Score!

2250

High Performance
Scalability
No complex q

B on f ends, their
rcha and reviews

Shopping Cart

750

Recommendation

500

Just Like the Patterns Game!

Points for each Pattern used

Extra points if one class implements multiple Pattern

org.springframework.aop.framework

Class AbstractSingletonProxyFactoryBean

This is not how Software Architecture works.

Why not?

More is worse!

More hardware

More Developer Skillz

Not necessarily bad

More Ops Trouble

- Installation
- Backup
- Disaster Recovery
- Monitoring
- Optimizations

But: Polyglott Persistence Has a Point

- Replace RDBMS often not an option
  - Enterprises will stick to RDBMS
  - Pure technology migration basically never happens
  - …only vendors think differently

Example: Archive Database

- Store current data in RDBMS
- Store archive in NoSQL (MongoDB)
- Archive contains mainframe data

Benefit: Use flexibility to allow for many data formats

Benefit: No need to convert mainframe data

Benefit: Store lots of data cheaply

Another Example:
Complex Document Processing System

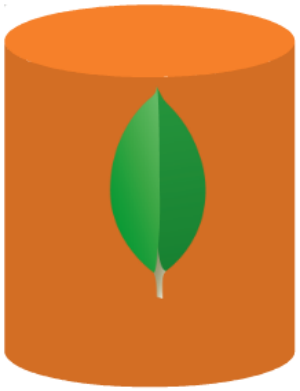

MongoDB

Document-oriented

Documents

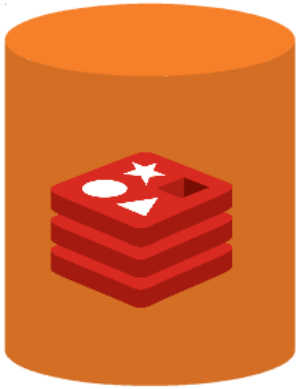

Redis

Key/value in memory

Meta Data for quick access

elastic search

Search engine

Search index

Alternative: Only elasticsearch

Stores original documents, too

Support for complex queries

Very powerful features also for data mining / analytics

Not well suited for update heavy operations

Backup / disaster recovery?

elastic search

Alternative: Only MongoDB

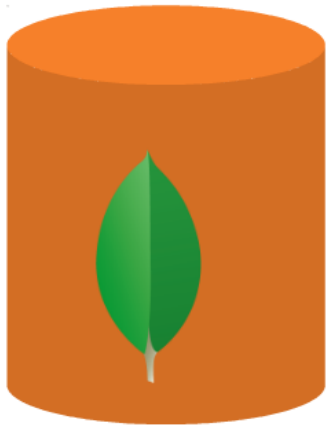

MongoDB

Now with (limited beta) fulltext search

Excellent support for updates

Quite fast – memory mapped files

Also fast for updates

Disaster recovery possible

Map/Reduce support

What about Redis?

MongoDB uses memory mapped files

Why Redis?

Like a Swiss Knife

- Small & handy
- Cache
- Messaging
- Central coordination in a distributed environment

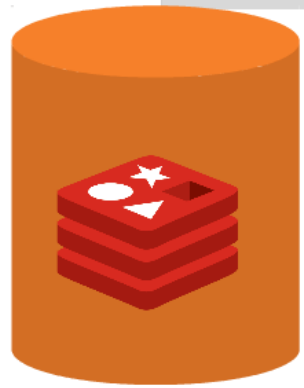

Your Choice: a trade off!

Typical architecture decision

Shall I use?

N

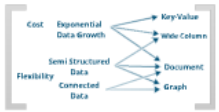

Document-oriented databases are a nice trade-off

DBA vs. Architect

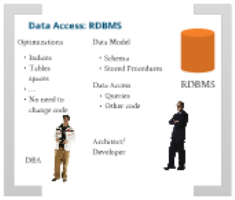

RDBMS separates data from data access

Indices, joins and normalization allow flexible data access patterns

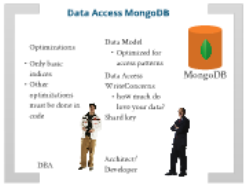

Data Access: RDBMS

Optimizations

- Indices
- Tables spaces
- …
- No need to change code

DBA

Data Model

- Schema
- Stored Procedures

Data Access

- Queries
- Other code

Architect/ Developer

RDBMS

RDBMS separates data from data access

Indices, joins and normalization allow

data access

Indices, joins and normalization allow flexible data access patterns

Data Access MongoDB

Optimizations

- Only basic indices
- Other optimizations must be done in code

DBA

Data Model

- Optimized for access patterns

Data Access WriteConcerns

- how much do love your data?

Shard key

Architect/ Developer

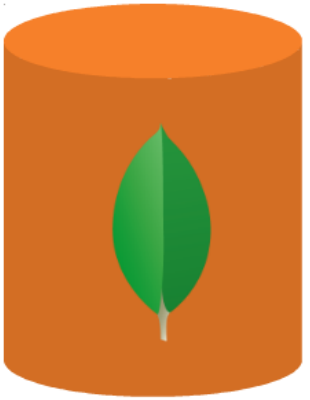

MongoDB

Cluster: RDBMS

Works somehow

i.e. transparent to developers

A special setup of hardware and RDBMS software

DBA

Cluster: MongoDB

CAP theorem

If the network is down choose Consistency xor Availabilty

Write Concerns:

- Unacknowledge
- Acknowledged
- Journaled
- Some nodes in the replica set

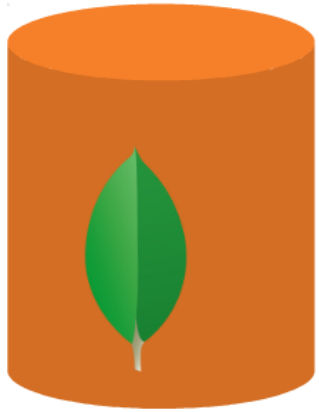

Influences replication

MongoDB has master / slave replication

Queries might go to master only or also slaves

Influences consistency

MongoDB

Architect/
Developer

More Power and more Responsibility

Another driver for cross functional teams

Architect

DB Admin

Architects

Architecture has always been a multi-dimensional problem

Need to choose persistence technology

Need to think about operations

Needs to do DBA work

NoSQL Is All About the Persistence Question

?