

Making choices:

What kind of relationship are you seeking with your database?

February 13, 2014

J.R. Arredondo

Director, Data Services Product Marketing

@jrarredondo



This is your chance to leave the room!

You paid a lot of money to be here! Nobody will be offended if there is a better session for you!

Ballroom AB

- MLbase: Distributed Machine Learning Made Easy

Ballroom CD

- Real-time Analytics with Open Source Technologies

Ballroom E

- The Great Debate: Technology Creates More Jobs than it Destroys

GA Ballroom J

- Making Data Move: Stream Processing in Go

GA Ballroom K

- StatusWolf: Creating Dashboards That Don't Suck Using Art and Engineering

Mission City M

- Driving the Future of Smart Cities - How to Beat the Traffic

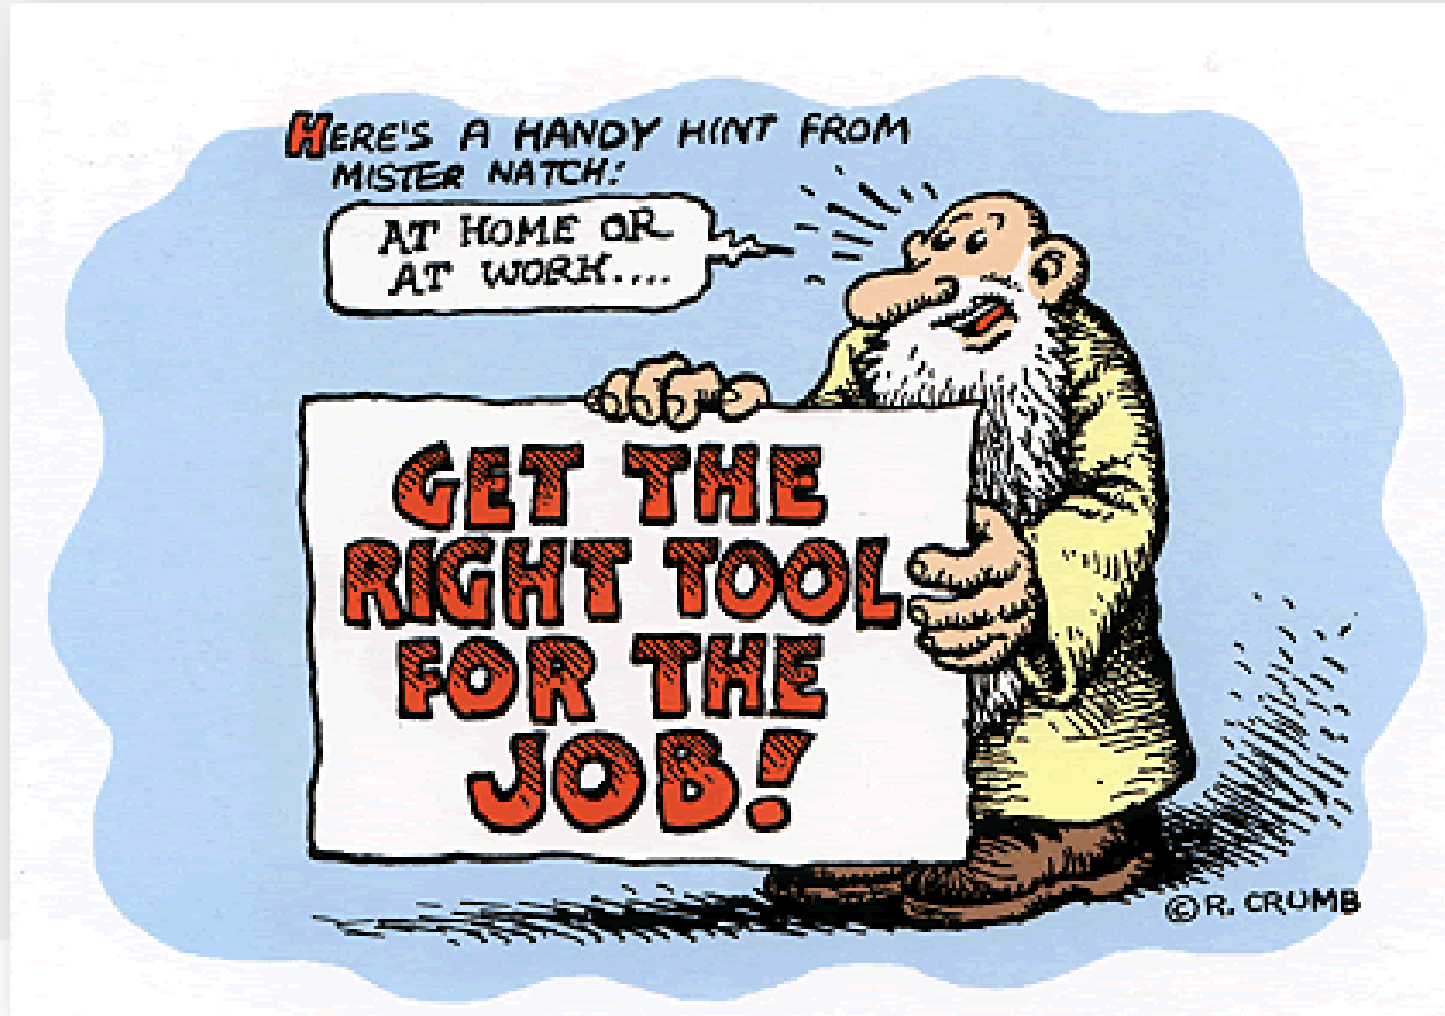
Ballroom F

- Real-Time Analytics with NewSQL: Why Hadoop is not enough

Ballroom H

- Think Big Data is the Answer? Think Again.

Common advice these days from smart people



This is what they think you are about to do

(not that you are about to)



Let's take a step back



Databases are not simple, single purpose tools



What kind of relationship are you seeking with your database?



The image shows a screenshot of the Facebook 'People I may know' interface. The main heading is 'Relationship' with an 'Edit' button. Below this, there are two sections: 'Relationship Status' and 'Family'. The 'Relationship Status' dropdown menu is open, showing a list of options: '---', 'Single', 'In a relationship', 'Engaged', 'Married', 'In a civil union', 'In a domestic partnership', 'In an open relationship', 'It's complicated' (highlighted in blue), 'Separated', 'Divorced', and 'Widowed'. To the right of the dropdown, there is a 'Choose Relationship' button. Below the 'Family' section, there is a heart icon and the text 'Add Your Family'.

How did we get here?

App development is changing

	Traditional apps (CRM, HR, Finance apps)	Modern apps (mobile, social, media, games)
Infrastructure	Custom-built <u>for</u> the app	Programmable <u>by</u> the app
Data	Mostly resides on premise	Mostly resides on cloud



Applications are becoming systems of engagement

	Traditional apps (CRM, HR, Finance apps)	Modern apps (mobile, social, media, games)
Characteristics of the system	Systems of Record Highly structured Slow to change Transactional Stable Core to the business Not very social	Systems of Engagement Loosely structured Quick to adapt Conversational Dynamic and in flux Edge of the business Fundamentally social
Data	Mostly resides on premise	Mostly resides on cloud



We are building different kinds of applications

MEDIA



GAMING



M2M



MOBILE



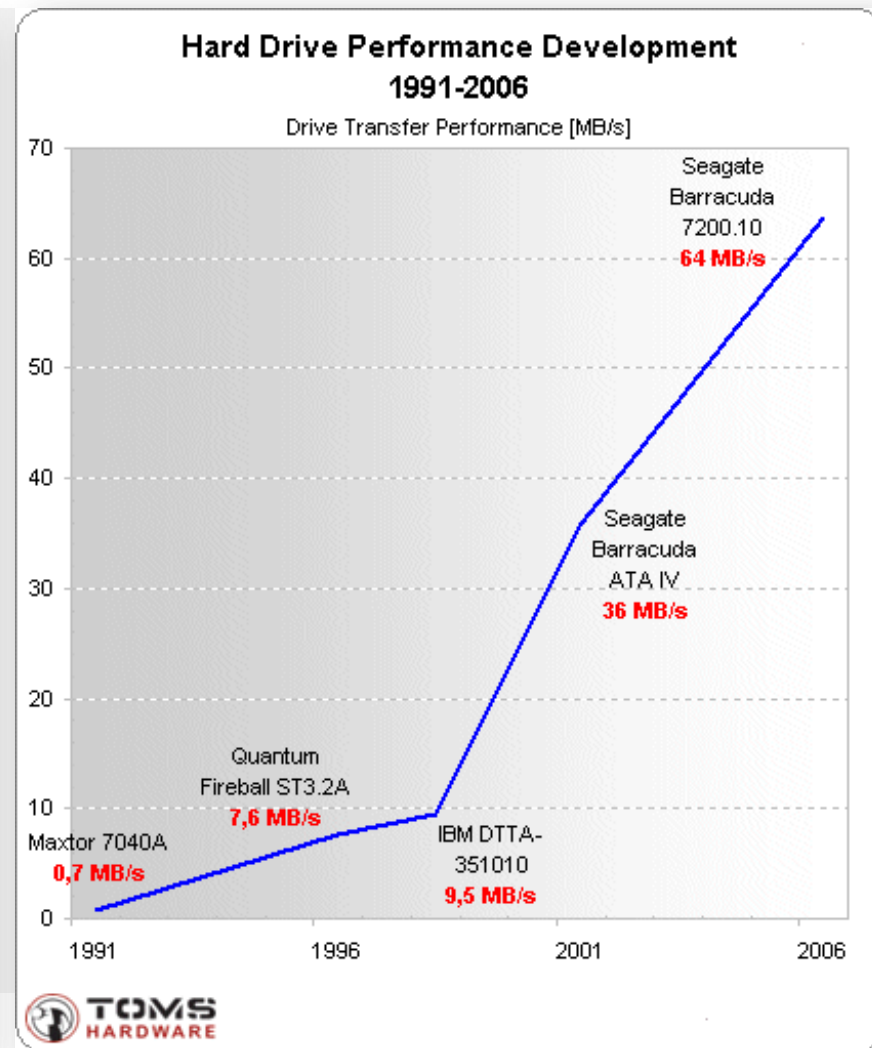
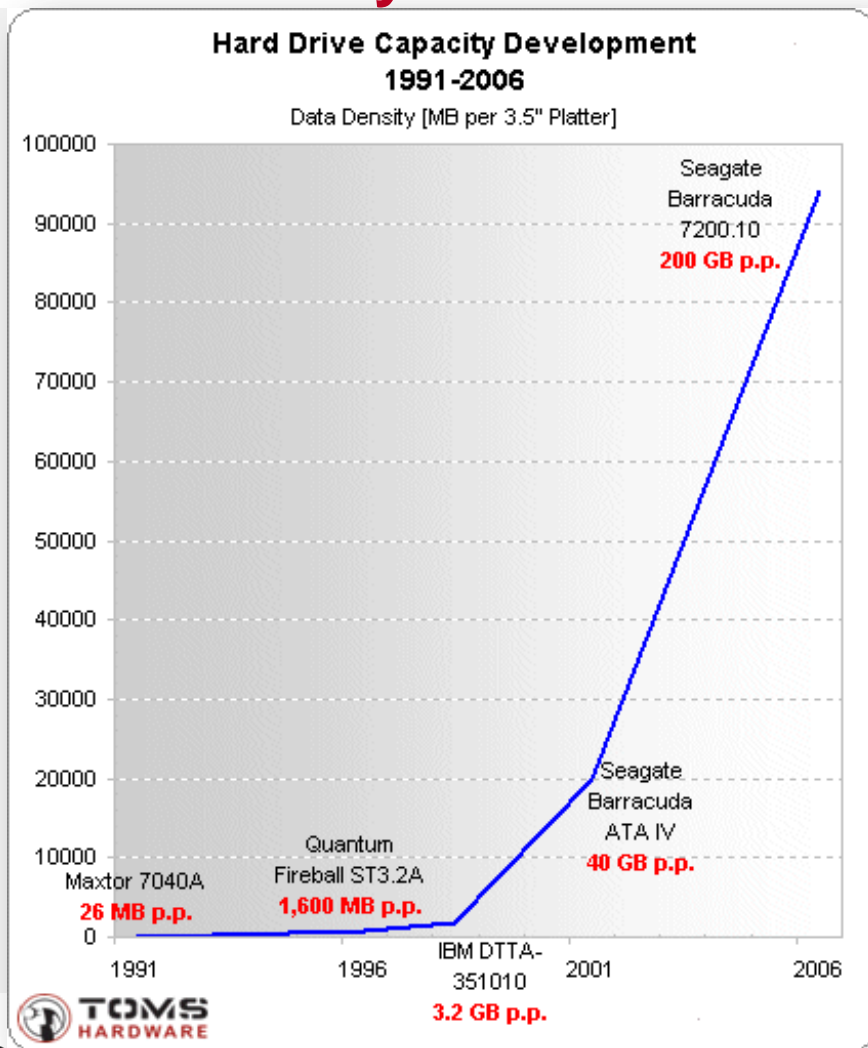
SOCIAL



SOME UNIQUE SCENARIOS

- Frequently written, rarely read
- Binary files
- Short term data
- Multi-location access
- Zero downtime needs
- Dynamic or object oriented models
- Trying to avoid RAID / storage limits
- High speed data retrieval needs
- Large files

In the 15 year period before 2006, storage density increased 10,000x, but performance only increased about 100x



As a result, a revolution ensued in the world of Data Services

There are about 150+ choices **just** in the “NoSQL” subset



How should we pick?

What are you, fundamentally, trying to do?

Transact

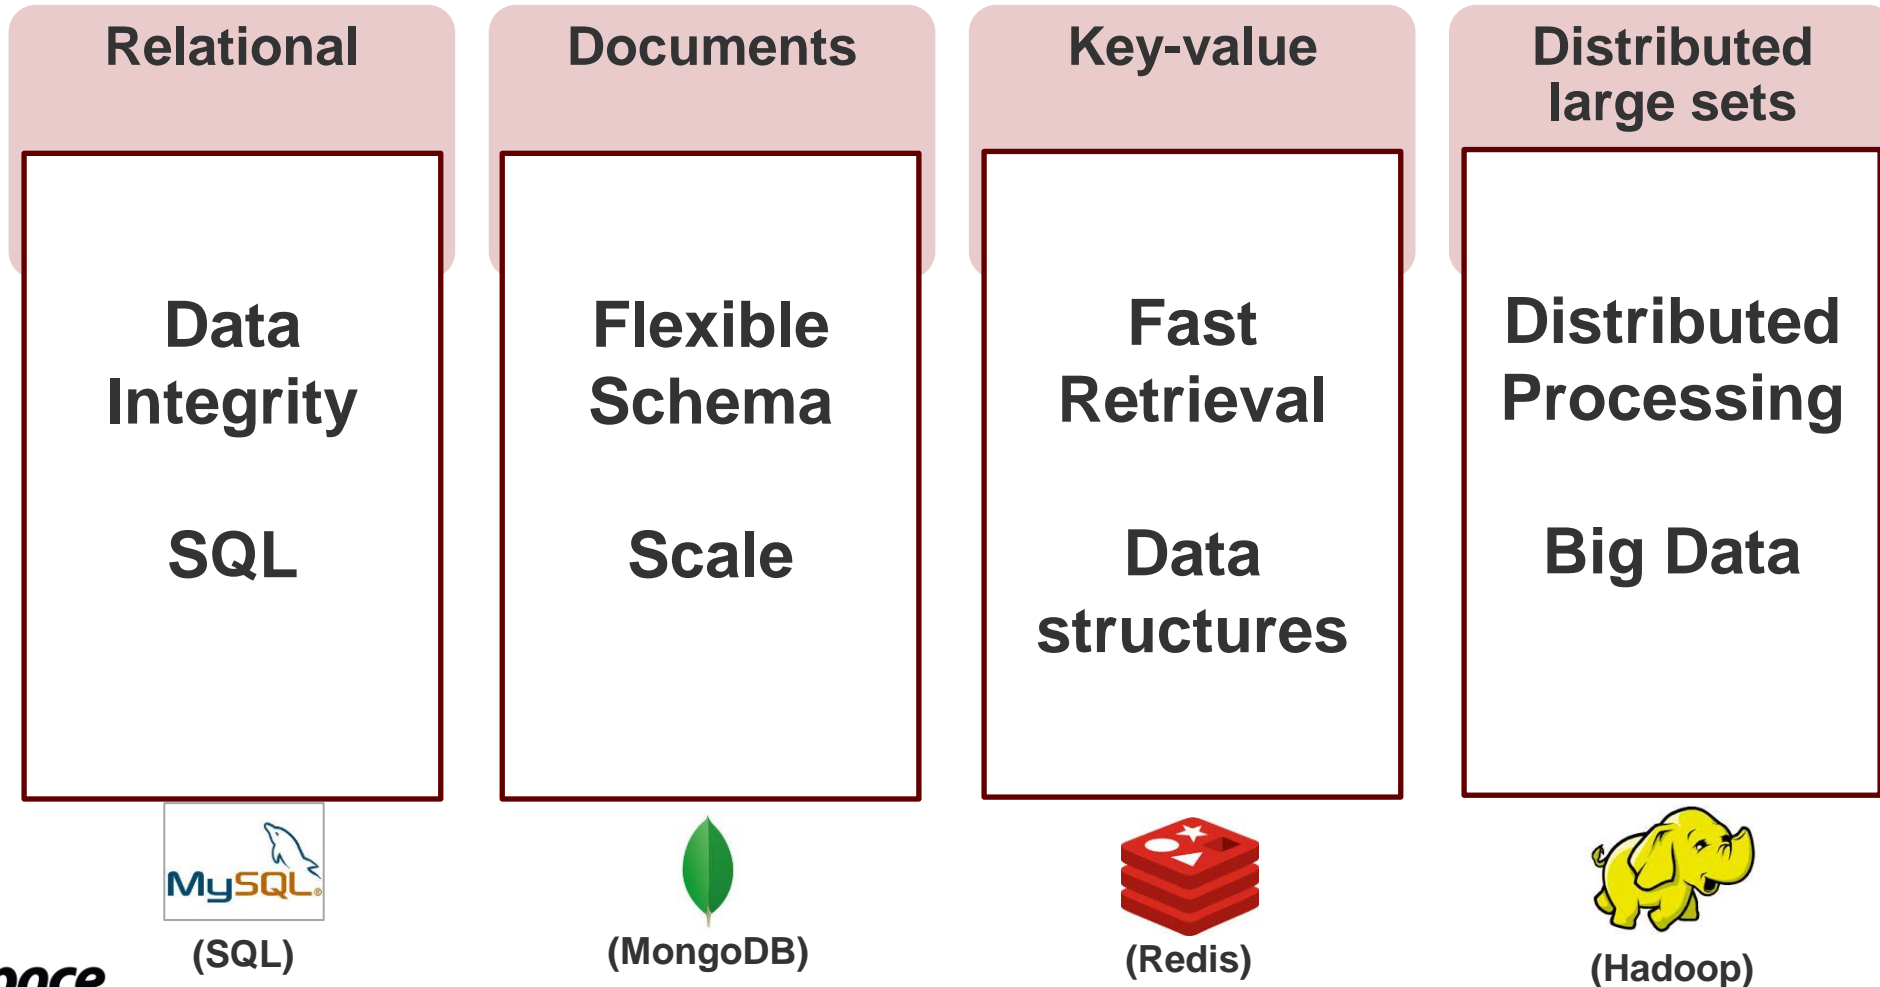
Improve
an
implementation

Analyze

Or both (operational intelligence!)

Understand the personality of your database

Let's use these examples (things may be different in your implementation)



Transacting

Relational databases

They literally saved the world from running major operations using paper

Strengths

- **Data integrity** through data types and semantic rules
 - AGE \geq 0
 - Person must have a NAME
- **Querying**
- **Aggregation**
- **SQL**

```
SELECT SUM(VALUE)  
FROM CAR  
GROUP BY MODEL;
```

“Weaknesses”

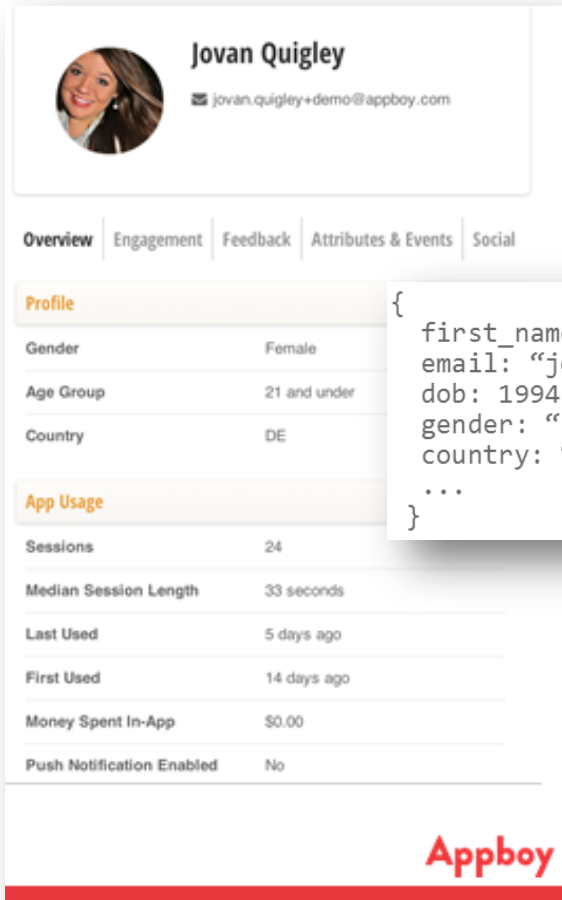
- **Complex development** as developer needs to map relational model with object oriented code
- **Complexity grows** exponentially as relational model grows
- **Difficult to scale**
- **Expensive** (hardware, software)

If your operation depends on the integrity of your business rules, the relational model rules.

Scaling is a little difficult.

Flexibility of data model (and its problems) with document databases

Appboy: App marketing automation platform for mobile apps



The screenshot shows a user profile for Jovan Quigley. The profile includes a name, email address, and a navigation menu with tabs for Overview, Engagement, Feedback, Attributes & Events, and Social. Below the navigation menu, there are two sections: 'Profile' and 'App Usage'. The 'Profile' section contains a table with fields like Gender, Age Group, and Country. The 'App Usage' section contains a table with fields like Sessions, Median Session Length, Last Used, First Used, Money Spent In-App, and Push Notification Enabled. The Appboy logo is visible in the bottom right corner of the screenshot.

Profile	
Gender	Female
Age Group	21 and under
Country	DE

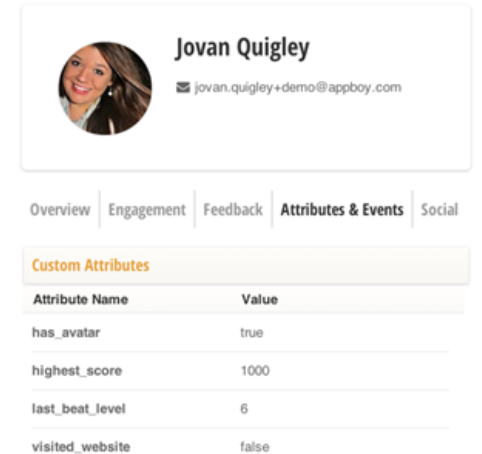
App Usage	
Sessions	24
Median Session Length	33 seconds
Last Used	5 days ago
First Used	14 days ago
Money Spent In-App	\$0.00
Push Notification Enabled	No

```
{  
  first_name: "Jovan",  
  email: "jovan+demo@appboy.com",  
  dob: 1994-10-24,  
  gender: "F",  
  country: "DE",  
  ...  
}
```

Extensible User Profiles

Custom attributes can go alongside other fields!

```
{  
  first_name: "Jovan",  
  email: "jovan+demo@appboy.com",  
  dob: 1994-10-24,  
  gender: "F",  
  custom: {  
    has_avatar: true,  
    highest_score: 1000,  
    visited_website: false,  
    ...  
  },  
  ...  
}
```



The screenshot shows a user profile for Jovan Quigley, similar to the one on the left. Below the navigation menu, there is a section titled 'Custom Attributes' which contains a table with columns for Attribute Name and Value. The table lists several custom attributes: has_avatar (true), highest_score (1000), last_beat_level (6), and visited_website (false). The Appboy logo is visible in the bottom right corner of the screenshot.

Custom Attributes	
Attribute Name	Value
has_avatar	true
highest_score	1000
last_beat_level	6
visited_website	false

```
db.users.find(...).update({$set: {
```

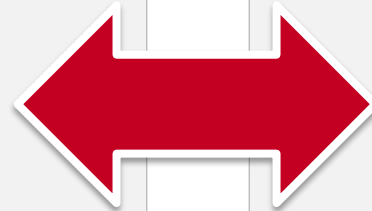
```
{ visited_website: true }  
{ visited_website: "yes" }
```

Tradeoffs

Data integrity
Business rules
Consistency
Transaction isolation
Atomicity

and

Rigidity



Flexibility of schema
Dynamic data models
Horizontal scale
Easier to get started

and

Inconsistency of data

It's good to understand the fundamental "theory"

What does your problem really need?

ACID

- **Atomicity:** A transactions either happens completely, or not at all
 - No partial transactions
- **Consistency:** Transactions end in a "valid" state
 - No violation of rules
- **Isolation:** Transaction appears as if it is the only thing happening to the database
 - Relaxed most times
 - Deals with phantom, dirty reads or non repeatable reads
- **Durability:** Committed transactions are permanent
 - Even after failure

BASE

- **Basically available:**
 - Supporting partial failures without complete system failure
 - Design as if users would end up in different partitions
- **Soft state:**
 - Things can be in flux for a little bit of time
- **Eventual consistency:**
 - Things right themselves

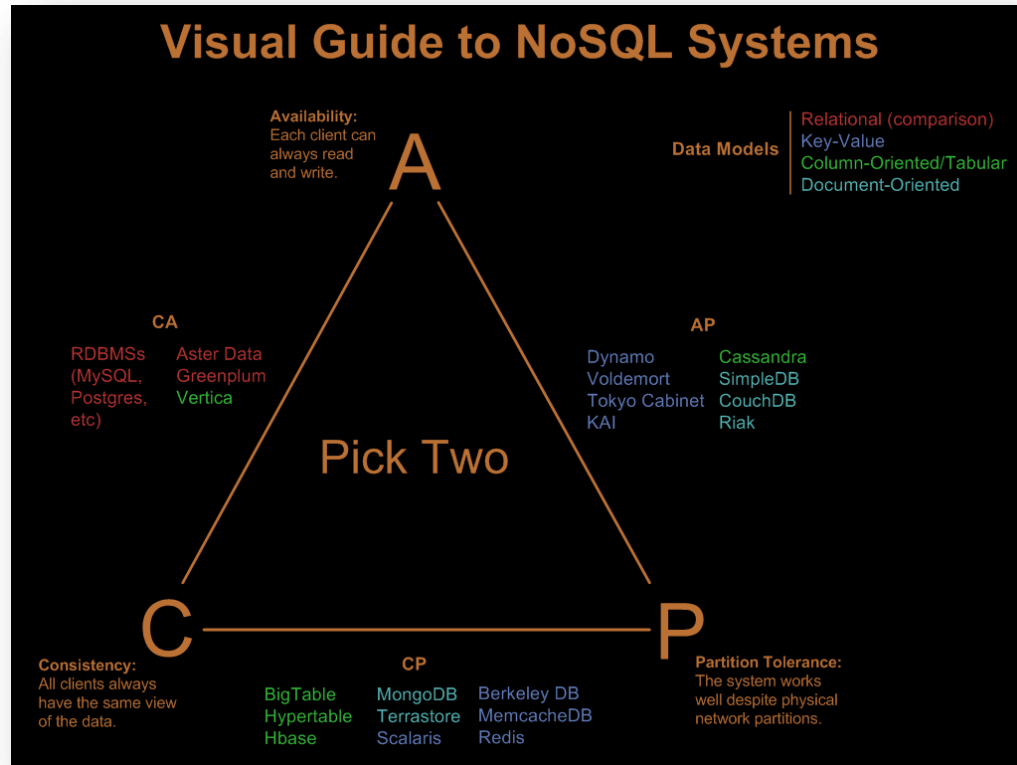
New ways of thinking:

Do customers really need to know the level of inventory of a product to place an order? Maybe all they want is to know that it is not zero

Know your CAP, really

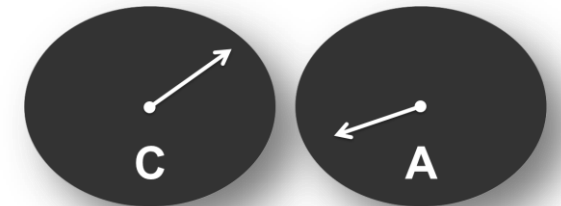
Consistency, Availability and Partition Tolerance

You can only have 2 out of 3 in CAP!



Wait! It's not that simple

- Partitions are not generally common
- Choosing Consistency or Availability is not final
- “It depends”
 - Maybe on user
 - Maybe on system
 - Maybe on type of data
- Just think:



- How am I going to detect a problem in the network? (P)
- How am I going to limit operations once I detect that?
- How am I going to compensate to recover?

Sometimes...
... the answer is both

What is Untappd?

A social discovery and sharing network for beer drinkers



- Heavily used during weekends and at night
- Complex SQL queries
- “What are my friends drinking?”
- “Where can I find this beer?”



MySQL and MongoDB together

It's not one or the other

- What works best for the workflow?
 - MySQL worked best for reference data for us
 - Not everything moved to MongoDB

What stayed in MySQL?

Check-ins
Users
Relationships Data
Primary Datastore

What moved to MongoDB?

Activity Feed (Friend's Graph)
Recommendation Data
Location-based Check-ins



What relationship are you seeking with your transactional database?

Be clear, upfront and intentional

Don't let a database just happen to you

Analytics

This conference was full of good sessions on Hadoop analytics

Just two thoughts

It is possible to be “too” data driven

We are
dealing with
people here

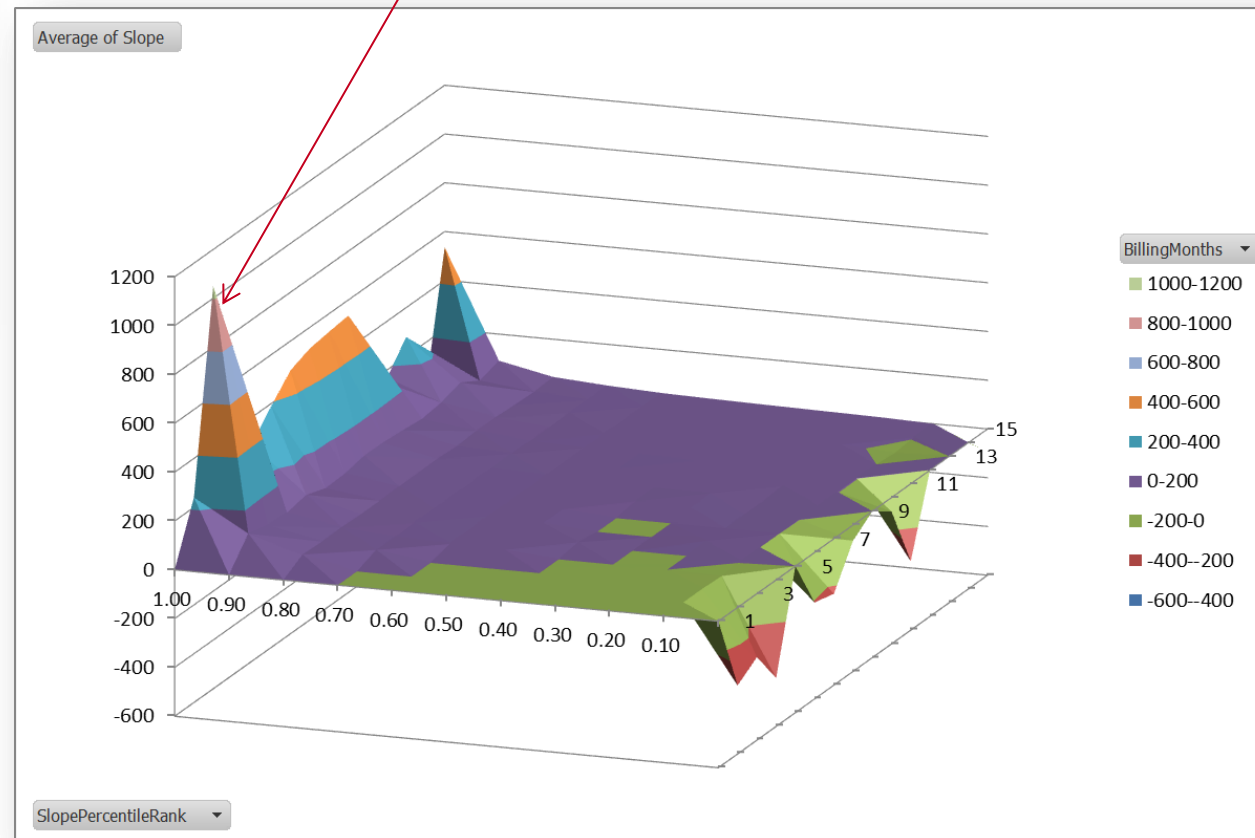
Data vs.
Intuition

The role of visualizations and end user tools

Excel still rules in many places (*gasp*)

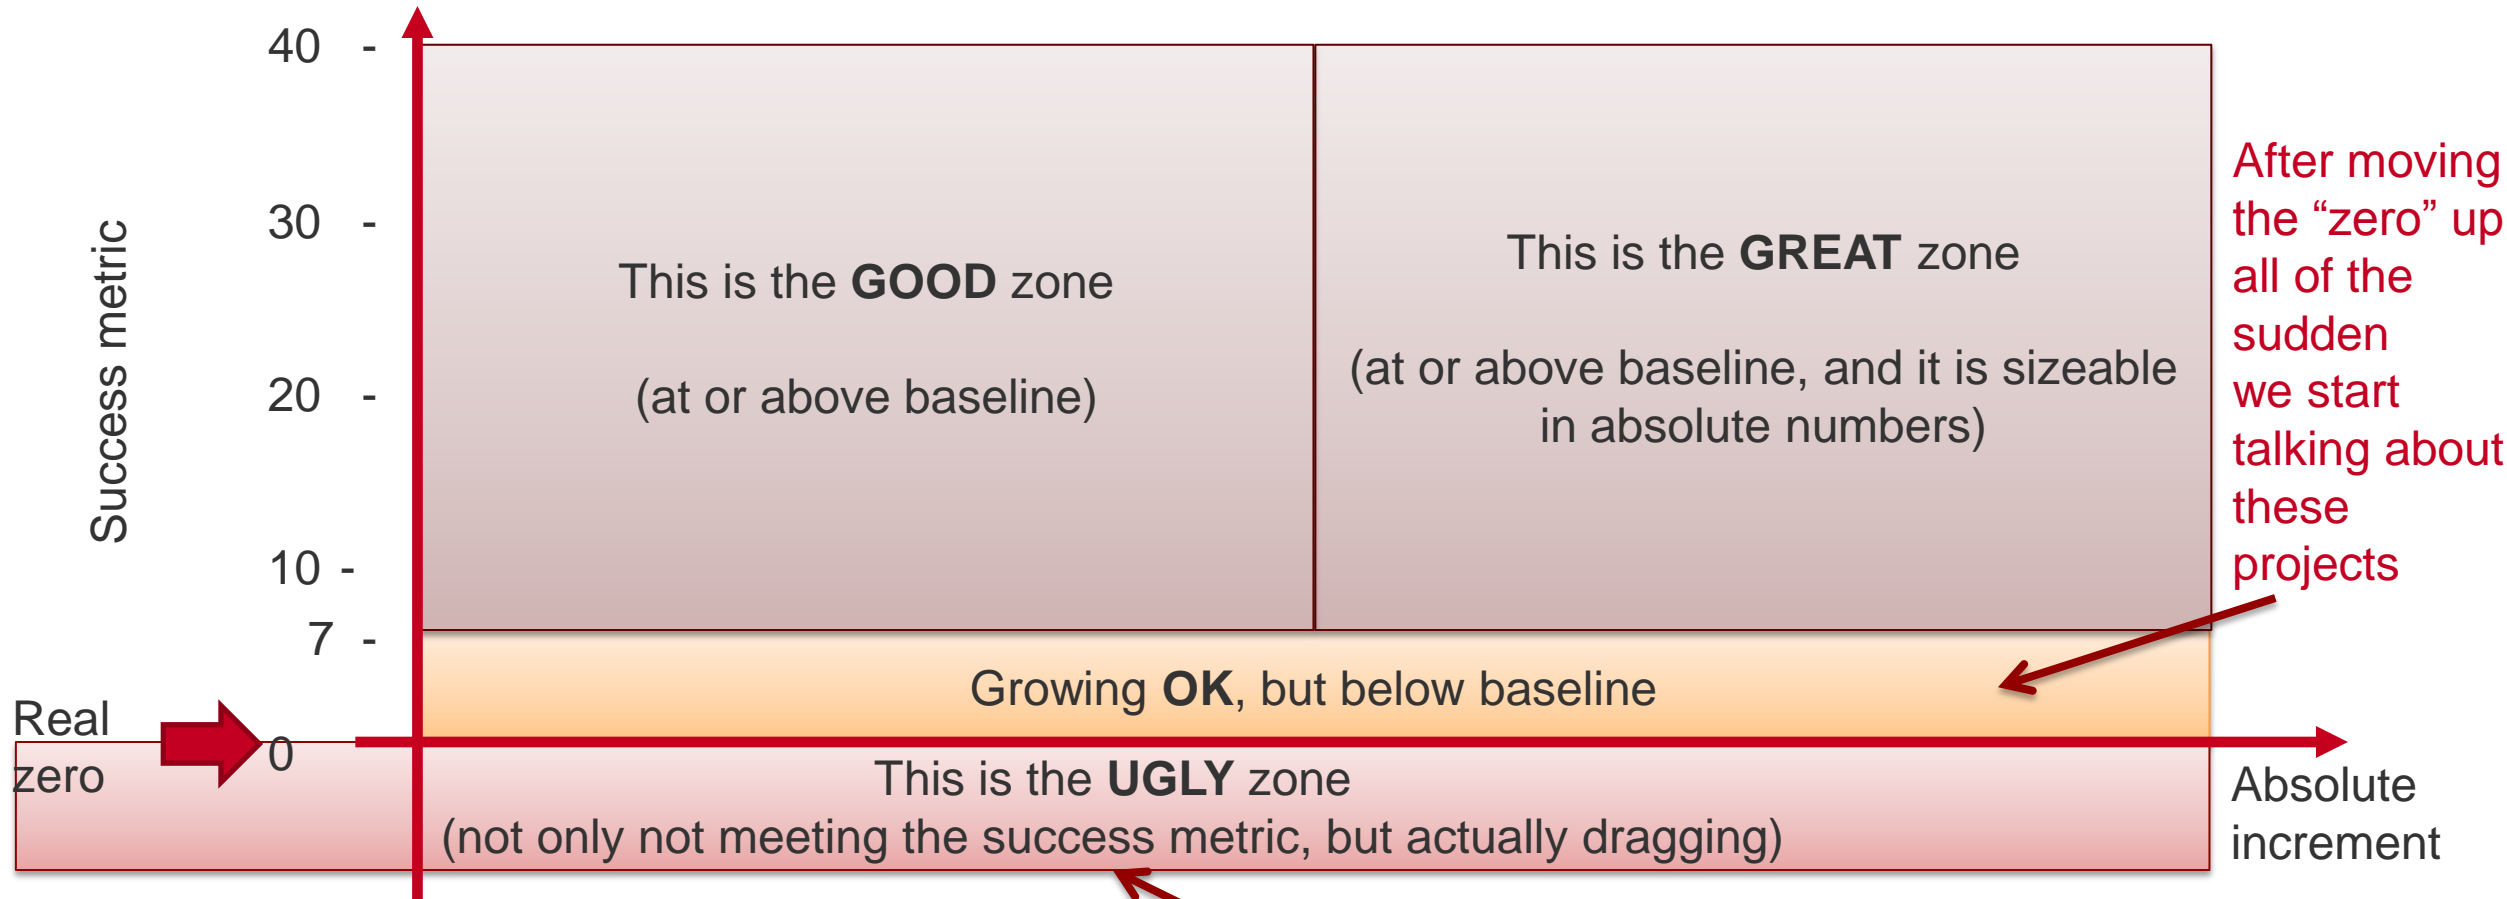
201203	Marker2	Trend	R
33440.59			
18007.59			
10901.37			
10721.84			
10077.15			
7597.64			
6355.67			
6025.06			
5400.66			
5291.82			
4785.9			
4521.8			

What is happening at this time?



Which things do you pay attention to?

“Redefining zero” changed the set of “things that we should discuss”



Data vs. Intuition: when do you need to use one vs. the other?

The iPhone would not exist if Apple would have had focus groups or made it a “data driven decision”

What relationship do you seek with your analytics backend, say, Hadoop + others?

**Hadoop, please just be there
for me, in the background**

**So that I can do what we
really need to do in this place
(drive insights to the
business)**

“Improving” an app

Just one simple example before we go

The “ilities” and their cousins: Improving an app

These are some of the challenges indirectly related to data that we must deal with

- Stability
- Fit for core scenarios
- Configurability to different scenarios
- Integration with development languages
- Integration with other databases
- SQL compatibility
- End user vs. Developer skillset
- Conceptual changes
- Platform availability
- Data type and semantic needs
- Security

- Performance
- Scalability
- Consistency
- Resiliency
- Data model
- Flexibility
- Cost
- Training
- Tools availability
- Development experience

Really understand the personality of your database

Let's talk about Redis

“Redis is a cache”

- SET
- GET

Redis is a server for data structures

- Keys
- Strings
- Hashes
- Lists
- Sets / Sorted Sets
- Publish / Subscribe



Huge difference!

What relationship do you seek with your database?

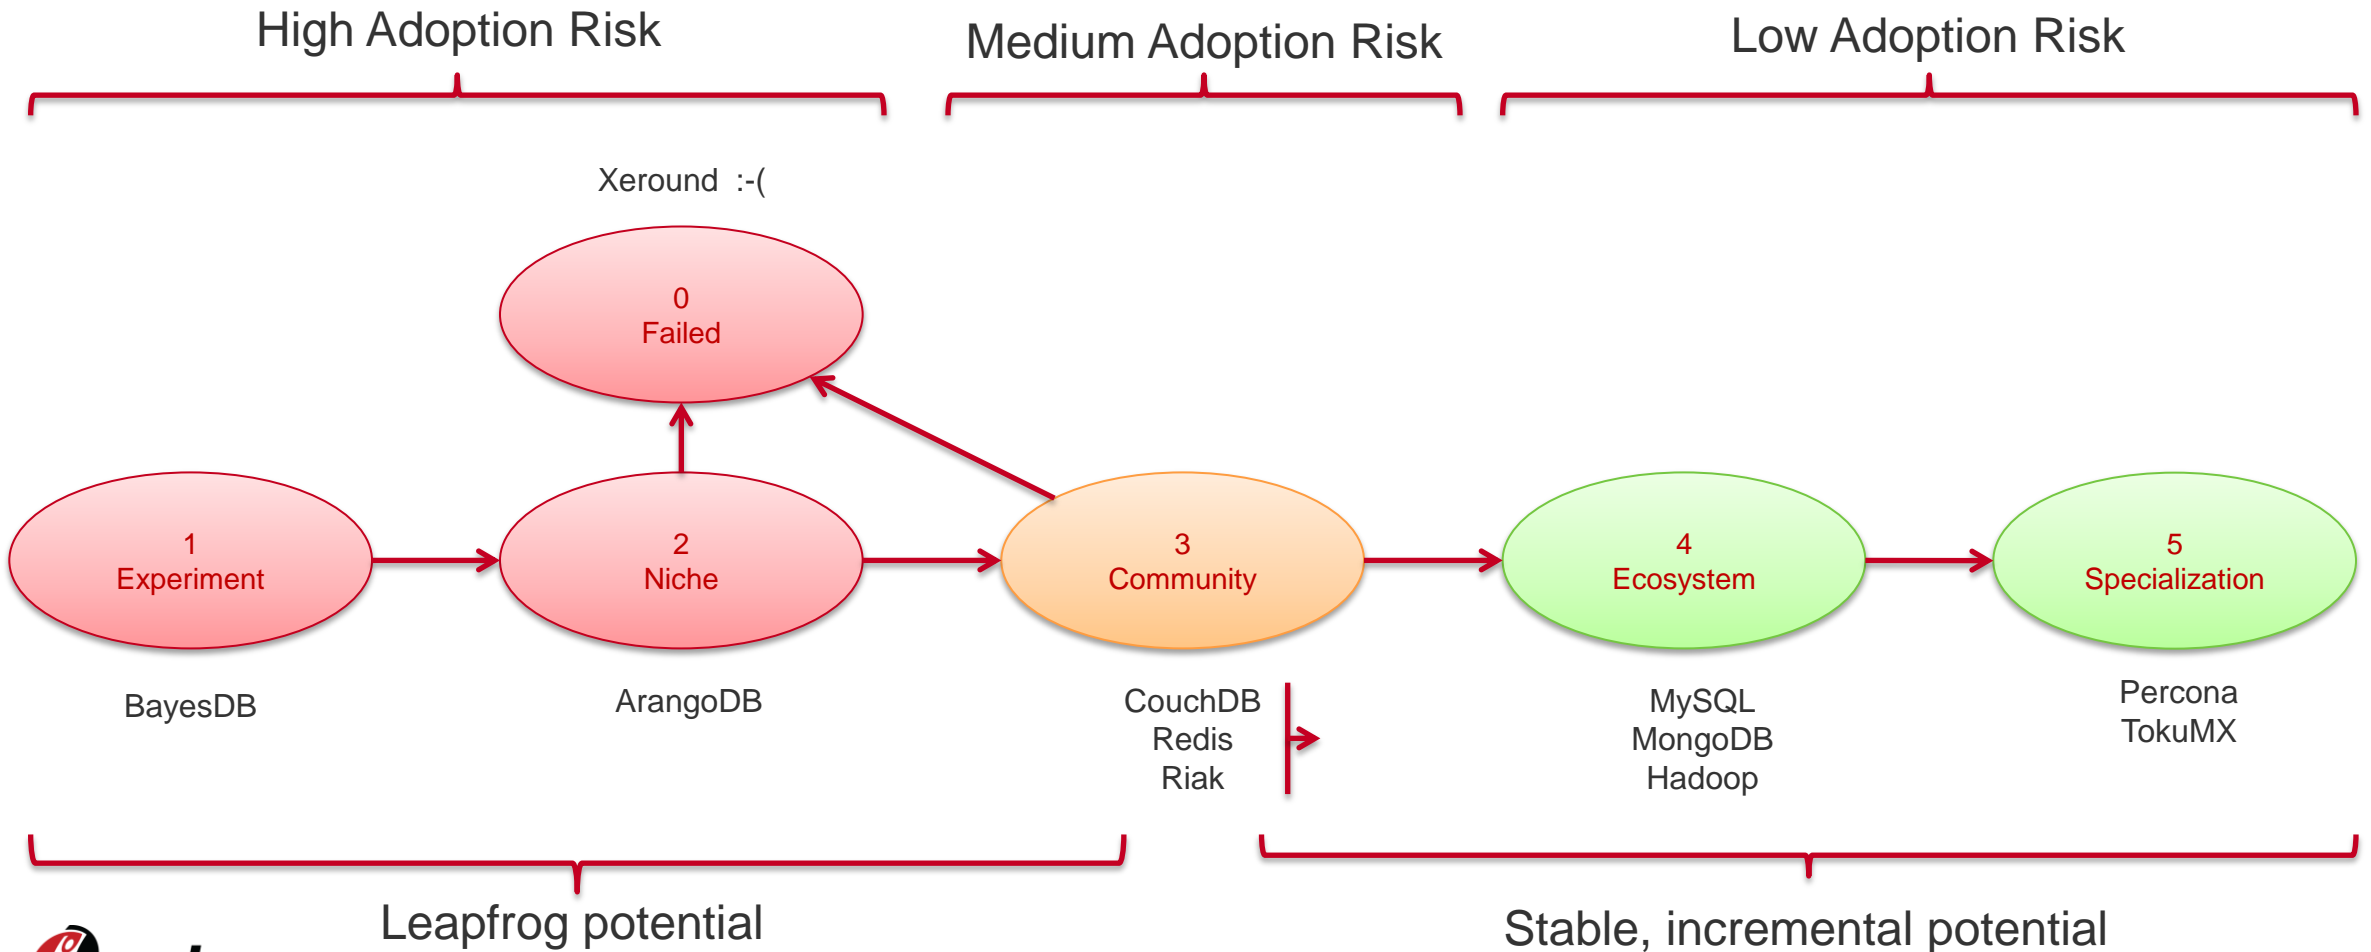
**One built on real
understanding of what you
bring to the table**

No first impressions please

Data Services also have a lifecycle

Be aware of the “state machine” of Data Services evolution

Know what you are getting yourself into

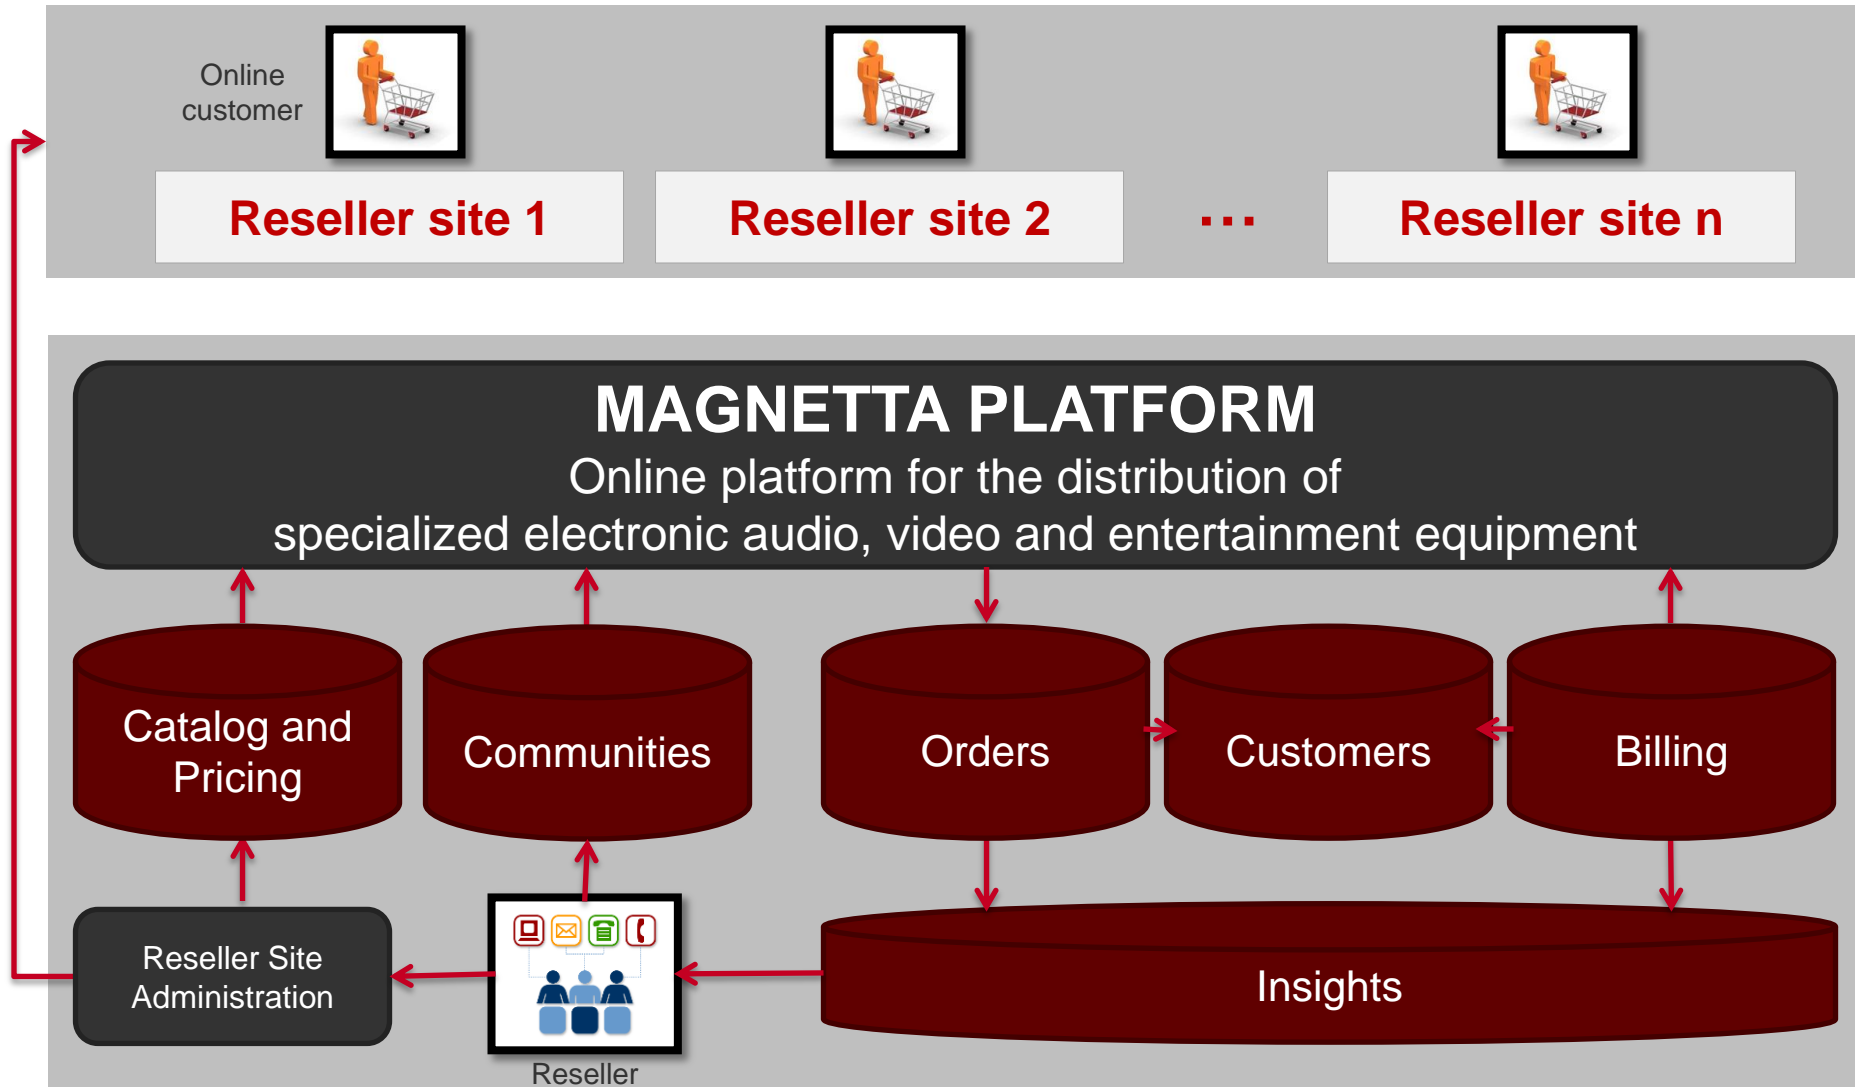


Simple things work



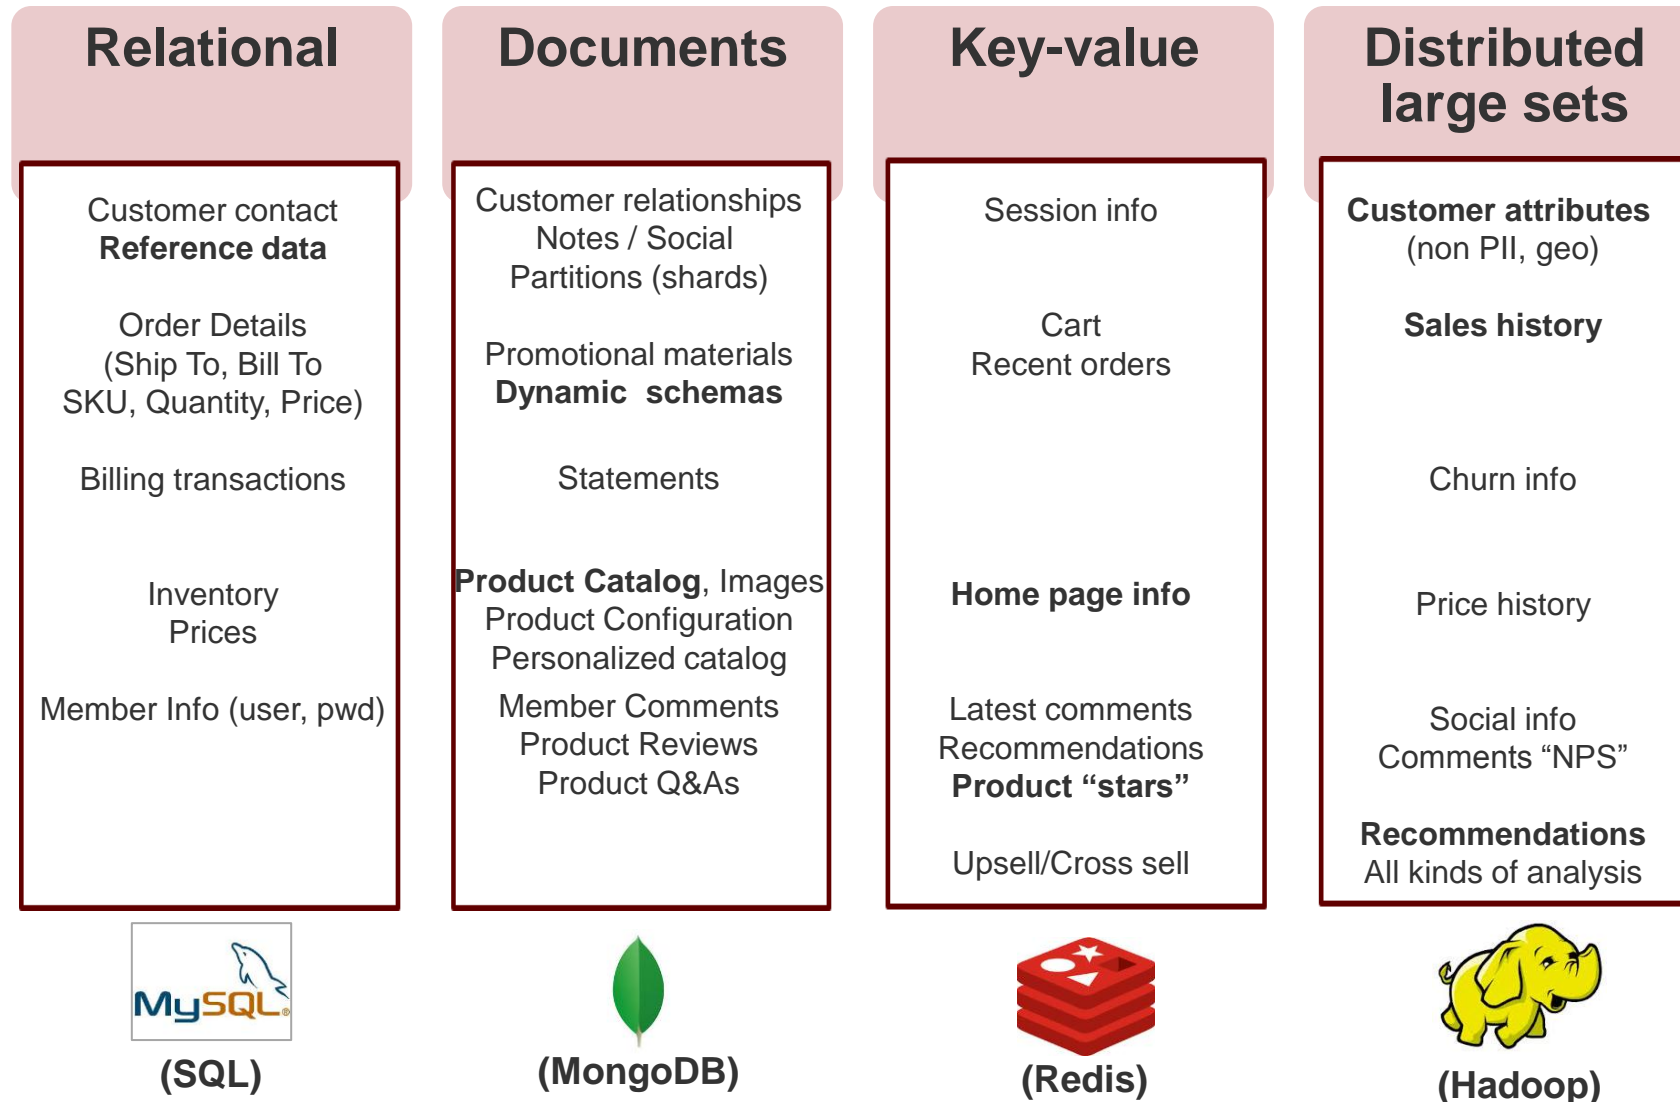
Being pragmatic: seeing it all in practice

We will use a fictional eCommerce company called Magnetta



What kind of info could Magnetta store in each type of database?

(remember that it always “depends” and use it as the foundation for your data access layer)



Rackspace's vision is Data as a Service

From databases to data as a service

From Database-as-a-Service to Data-as-a-Service

Focus on building your app, not managing databases

Highest value activity for your application

Build your application
(i.e. game, startup, mobile app, site)

**PRESUMABLY YOU WANT
TO BE FOCUSED HERE**

This is the only job that YOU MUST DO without anybody's help because this is your intellectual property

Manage software infrastructure
(i.e. databases)

**PRESUMABLY YOU DON'T
WANT TO HAVE TO
MANAGE DATABASES
OR SERVERS**

Manage hardware infrastructure

It only takes away from time building your application



The next vision for databases: Data-as-a-Service

Applications just access the data as a service, while the database is transparent

Highest value activity for your application

Build your application
and
manage your data

**PRESUMABLY YOU WANT
TO BE FOCUSED HERE**

This is the only job that YOU MUST
DO without anybody's help because
this is your intellectual property

hostname, port number

Data

as a Service

**The app just
interacts with
THE DATA**

The application does not see the
infrastructure

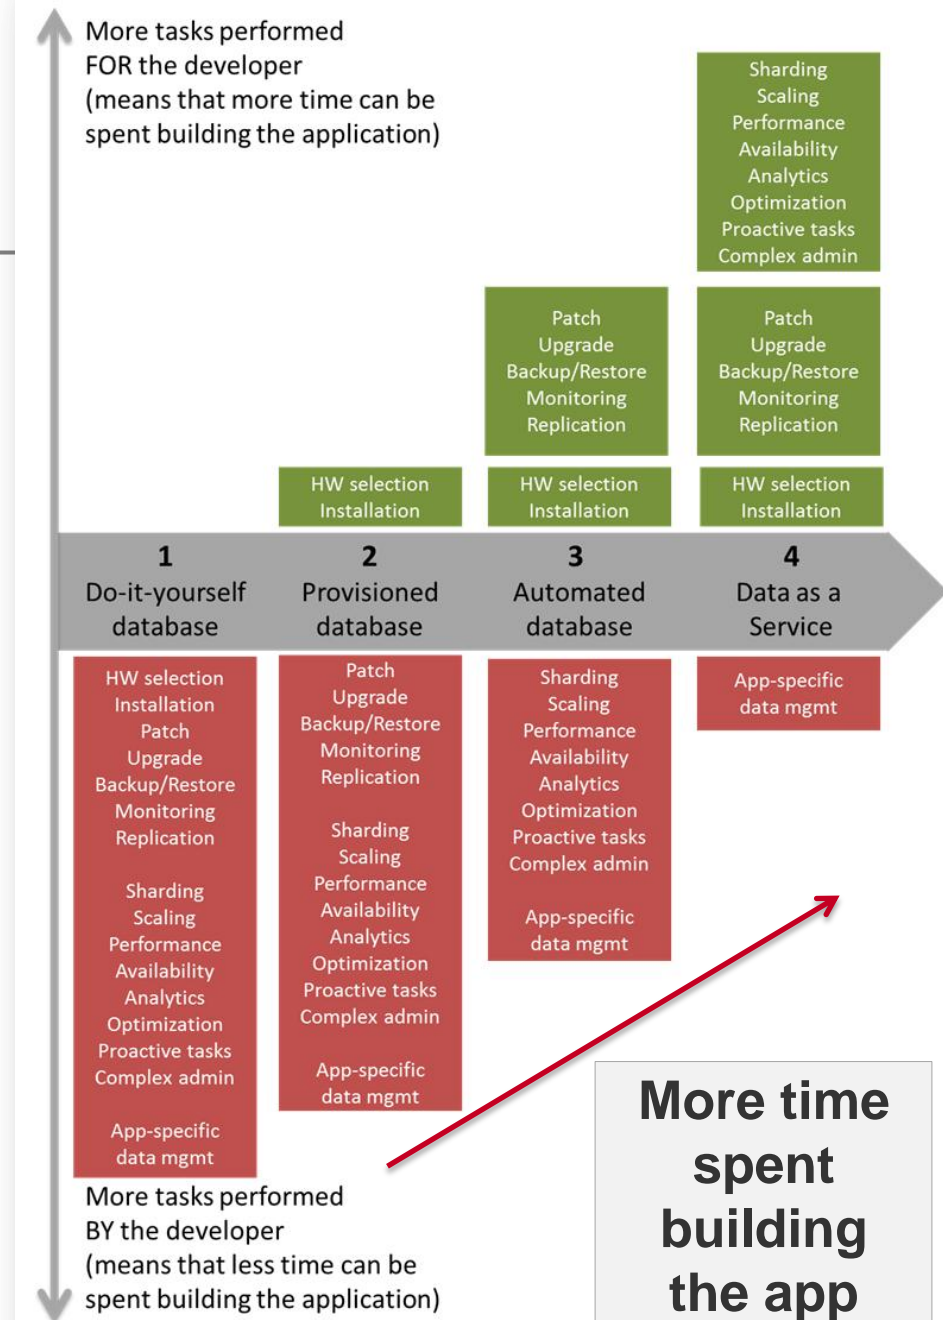
Towards transparent databases



Data-as-a-Service: more time building, less time managing databases

Four levels of DaaS transparency

- For some businesses, database or infrastructure management **IS core of the business**
- For most software-based businesses, database or infrastructure management represents time and resources **not spent building the application**
- **You must answer for yourself:** are you in the business of managing infrastructure, or in the business of [your market here]?



Data has mass and gravity: you need choices

(Or: “Divorces are expensive”)

Your Private
Cloud on prem

Private Cloud at
(your favorite
provider)

Public Cloud

Managed Cloud

Data Services at Rackspace



2 offerings in partnership
with Hortonworks for Hadoop-based
applications



2 acquisitions for
MongoDB and Redis apps



REDIS TO GO

Strong portfolio of
traditional offerings



In Summary, databases are like relationships

Don't let a database just happen to you (be intentional)

Don't lose sight of the business goal

- Databases are best when invisible

Really understand the true nature of the database

Database management = time not spend building your business app

THANK YOU

@jrarredondo



RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78218
US SALES: 1-800-961-2888 | US SUPPORT: 1-800-961-4454 | WWW.RACKSPACE.COM