

# NoSQL

Marián Rusnák

Seminar of DISA Laboratory

9.12.2013

# Outline

- Motivation
- What is NoSQL
- Key-Value stores
- Wide-Column stores
- Document stores
- Summary

# THE WORLD OF DATA

NUMBER OF EMAILS SENT EVERY SECOND

2.9  
MILLION

DATA CONSUMED BY HOUSEHOLDS EACH DAY

375  
MEGABYTES

VIDEO UPLOADED TO YOUTUBE EVERY MINUTE

20  
HOURS

DATA PER DAY PROCESSED BY GOOGLE

24  
PETABYTES

TWEETS PER DAY

50  
MILLION

TOTAL MINUTES SPENT ON FACEBOOK EACH MONTH


700  
BILLION

DATA SENT AND RECEIVED BY MOBILE INTERNET USERS

1.3  
EXABYTES

PRODUCTS ORDERED ON AMAZON PER SECOND

72.9  
ITEMS



SOURCES: Cisco, comScore, Statista, Nielsen, YouTube

IN THE 21ST CENTURY, we live a large part of our lives online. Almost everything we do is reduced to bits and sent through cables around the world at light speed. But just how much data are we generating? This is a look at just some of the massive amounts of information that human beings create every single day.

# Motivation

- Massive data volumes
- Extreme query workload
- Schema evolution



# Traditional RDBMS

- Transactions – ACID
- Integrity
- Complex queries



ORACLE®



# But...

- Modern web apps may have different needs
  - Low latency
  - Scalability & elasticity
  - High availability
  - Flexible schemas / ser
  - Distributed




# What is NoSQL

- Term NoSQL
  - Not Only SQL
  - Non-relational database
- Databases different from traditional RDBMS

**N**ot  
**O**nly **SQL** OR ~~**SQL**~~

# What is NoSQL

- Meet modern web apps requirement
- Scalable
  - huge volumes of data
- Distributed
  - multiple nodes
  - multiple datacenters
- Flexible
  - no strict schema





# What is NoSQL

- Umbrella term for different types of datastores
  - Key-Value stores
  - Wide column stores / column families
  - Document stores
  - Graph databases
  - Other...

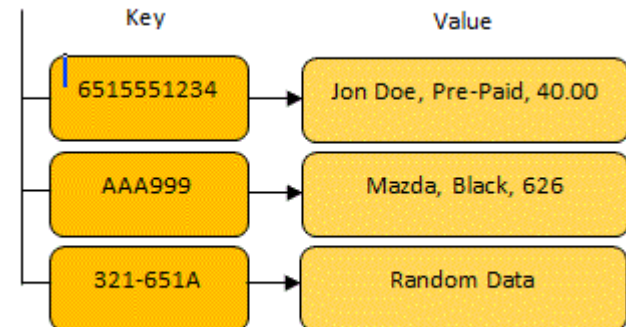
# What is NoSQL

- <http://nosql-database.org/>
  - 150 NoSQL databases
- Almost every big company has own solution



# Key-Value Stores

- Simplest
- Largest group
- Something like distributed hash tables
- Examples
  - Amazon Dynamo
  - Amazon DynamoDB
  - Voldemort
  - Redis
  - Oracle NoSQL Database
  - SimpleDB



# Dynamo



- Key-Value distributed storage system
- Pioneer in the area
- Many features incorporated by others
- Peer-to-Peer
- Simple data model – unique keys, no schema

# Dynamo



- Consistent hashing
  - “Ring” of nodes
  - Virtual nodes – node has several positions in the ring
- Replication
  - Successors in the ring

# Dynamo



- Quorum
  - Coordinator
  - $R+W > N$
- Gossip
  - Failure detection
- Object versioning

# Voldemort




- LinkedIn, now open-source
- Based on Dynamo
- Written in Java



# Wide-Column Stores

- Key is associated with multiple attributes
- Inspired by Google BigTable
- Examples
  - Google BigTable
  - HBase
  - Cassandra

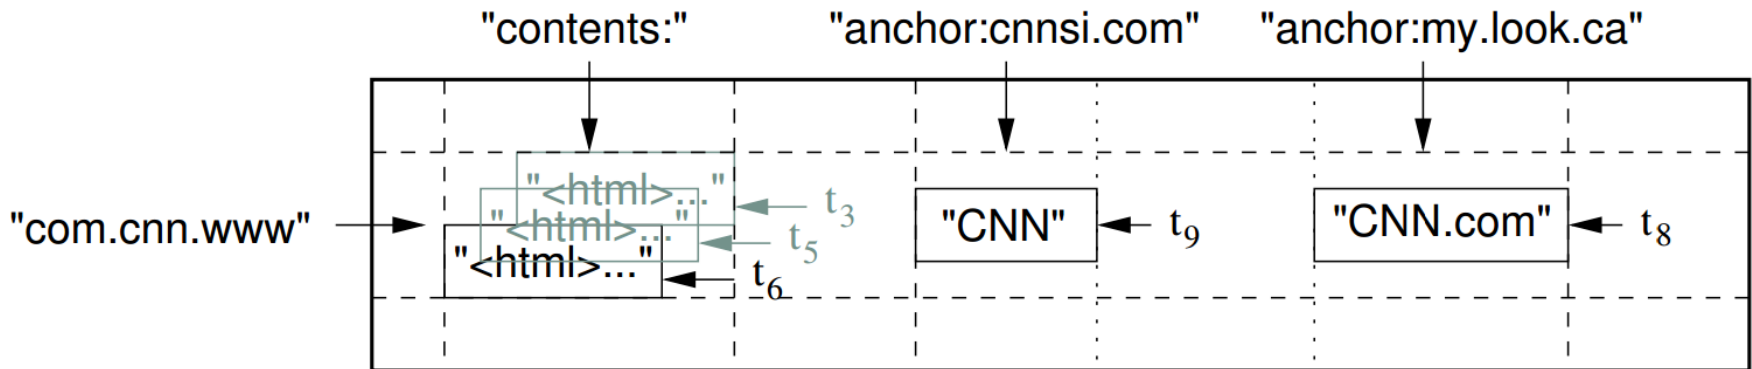




# BigTable



- Proprietary
  - Only short info paper released
- Sparse, distributed multi-dimensional sorted map
- $\langle \text{row\_key}, \text{column\_key}, \text{timestamp} \rangle \rightarrow \langle \text{value} \rangle$



# BigTable



- Rows
  - Data sorted by row key
- Tablets
  - Sequence of rows
  - Distributed
- Columns, Column families
  - Unlimited numbers of columns

# BigTable



- Versioning
  - Timestamps
  - Garbage collection
- Stores data on Google File System

# HBase



- Inspired by BigTable
- Open-source – Apache
- Uses Hadoop Distributed File System



# Cassandra *Cassandra*

- Facebook
  - Fast inbox search
- Now open-source Apache project
- Data model of BigTable
- Infrastructure of Dynamo



facebook

# Document Stores

- Value is more than a string
  - JSON
  - BSON
- Inspired by IBM Lotus Notes
- Very flexible schema
- Examples
  - MongoDB
  - CouchDB



# CouchDB



- Open-source Apache project
- Schema free
- JSON format
- B-Tree storage
- MVCC, no locking
- No joins, no primary or foreign keys
- Written in Erlang
- REST API

# MongoDB



- Open-source
- BSON format – similar to JSON
- Queries can be objects
- Multiple types of indexing
- Master/slave replication
- Written in C++
- Drivers in many languages
- Most popular
  - 23 000 questions on stackoverflow




# Trends




# Trends

- Gaining in popularity...



# Trends

- ...but still got a long way to go



# Summary

- NoSQL is suitable for modern web apps
  - Scalable
  - Distributed
  - Flexible
- 4 main types
  - Key-Value stores
    - Key to value mapping
  - Wide column stores / column families
    - More attributes associated with key
  - Document stores
    - Key to document mapping, not only to string
  - Graph databases
    - Connection between objects

# References

- SlideShare
  - [http://www.slideshare.net/marin\\_dimitrov/nosql-databases-3584443](http://www.slideshare.net/marin_dimitrov/nosql-databases-3584443)
  - <http://www.slideshare.net/dstainer/introduction-to-nosql-databases>
  - <http://www.slideshare.net/quipo/nosql-databases-why-what-and-when>
- Google BigTable paper
  - <http://static.googleusercontent.com/media/research.google.com/sk//archive/bigtable-osdi06.pdf>
- Stanford course
  - [http://infolab.stanford.edu/~widom/cs145/NoSQL\\_activity.pdf](http://infolab.stanford.edu/~widom/cs145/NoSQL_activity.pdf)
- <http://www.mongodb.com/learn/nosql>