L22: NoSQL

CS3200 Database design (sp18 s2)

https://course.ccs.neu.edu/cs3200sp18s2/

4/5/2018

Several slides courtesy of Benny Kimelfeld

Outline

- Introduction
- Transaction Consistency
- 4 main data models
  - Document Stores (e.g., MongoDB)
  - Key-Value Stores (e.g., Redis)
  - Graph Databases (e.g., Neo4j)
  - Column-Family Stores
- Concluding Remarks

SQL Means More than SQL

- SQL stands for the query language
- But commonly refers to the traditional RDBMS:
  - Relational storage of data
    - Each tuple is stored consecutively
  - Joins as first-class citizens
    - In fact, normal forms prefer joins to maintenance
  - Strong guarantees on transaction management
    - No consistency worries when many transactions operate simultaneously on common data
- Focus on *scaling up*
  - That is, make a single machine do more, faster

Vertical vs. Horizontal Scaling

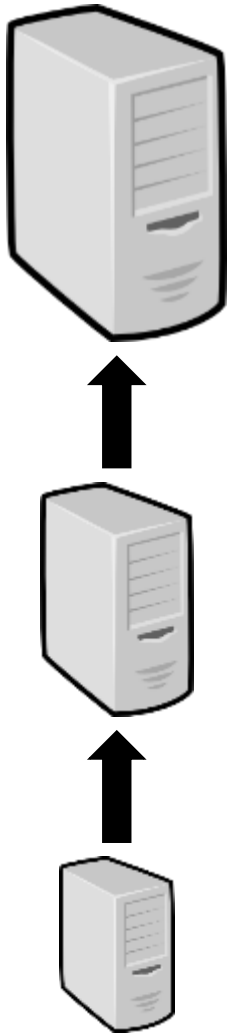

"scale up"

- **Vertical scaling ("scale up")**: you scale by adding more power (CPU, RAM)
- **Horizontal scaling ("scale out")**: you scale by adding more machines

"scale out"

Trends Drive Common Requirements

Social media + mobile computing

- Explosion in data, always available, constantly read and updated
- High load of simple requests of a common nature
- Some consistency can be compromised (e.g., 👍)

Cloud computing + open source

- Affordable resources for management / analysis of data
- People of various skills / budgets need software solutions for distributed analysis of massive data

Database solutions need to *scale out*
(utilize distribution, “scale horizontally”)

Compromises Required

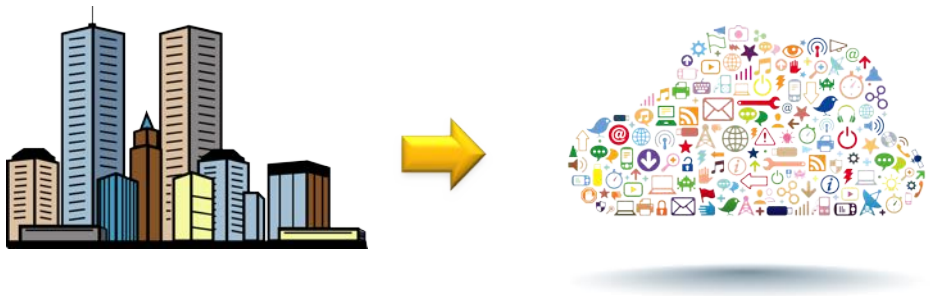

What is needed for effective distributed, data- and user-intensive applications?

1. Use data models and storage that allow to avoid joins of big objects
2. Relax the guarantees on consistency

NoSQL

- Not Only SQL
  - Not the other thing!
  - Term introduced by Carlo Strozzi in 1998 to describe an alternative database model
  - Became the name of a movement following Eric Evans's reuse for a distributed-database event
- Seminal papers:
  - Google's BigTable
    - Chang, Dean, Ghemawat, Hsieh, Wallach, Burrows, Chandra, Fikes, Gruber: Bigtable: A Distributed Storage System for Structured Data. OSDI 2006: 205-218
  - Amazon's DynamoDB
    - DeCandia, Hastorun, Jampani, Kakulapati, Lakshman, Pilchin, Sivasubramanian, Vosshall, Vogels: Dynamo: amazon's highly available key-value store. SOSP 2007: 205-220

NoSQL from nosql-database.org

“

- Next Generation Databases mostly addressing some of the points: being *non-relational*, *distributed*, *open-source* and *horizontally scalable*.
- The original intention has been modern web-scale databases. The movement began early 2009 and is growing rapidly. Often more characteristics apply such as: schema-free, easy replication support, simple API, eventually consistent / BASE (not ACID), a huge amount of data 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.

”

What is NoSQL?

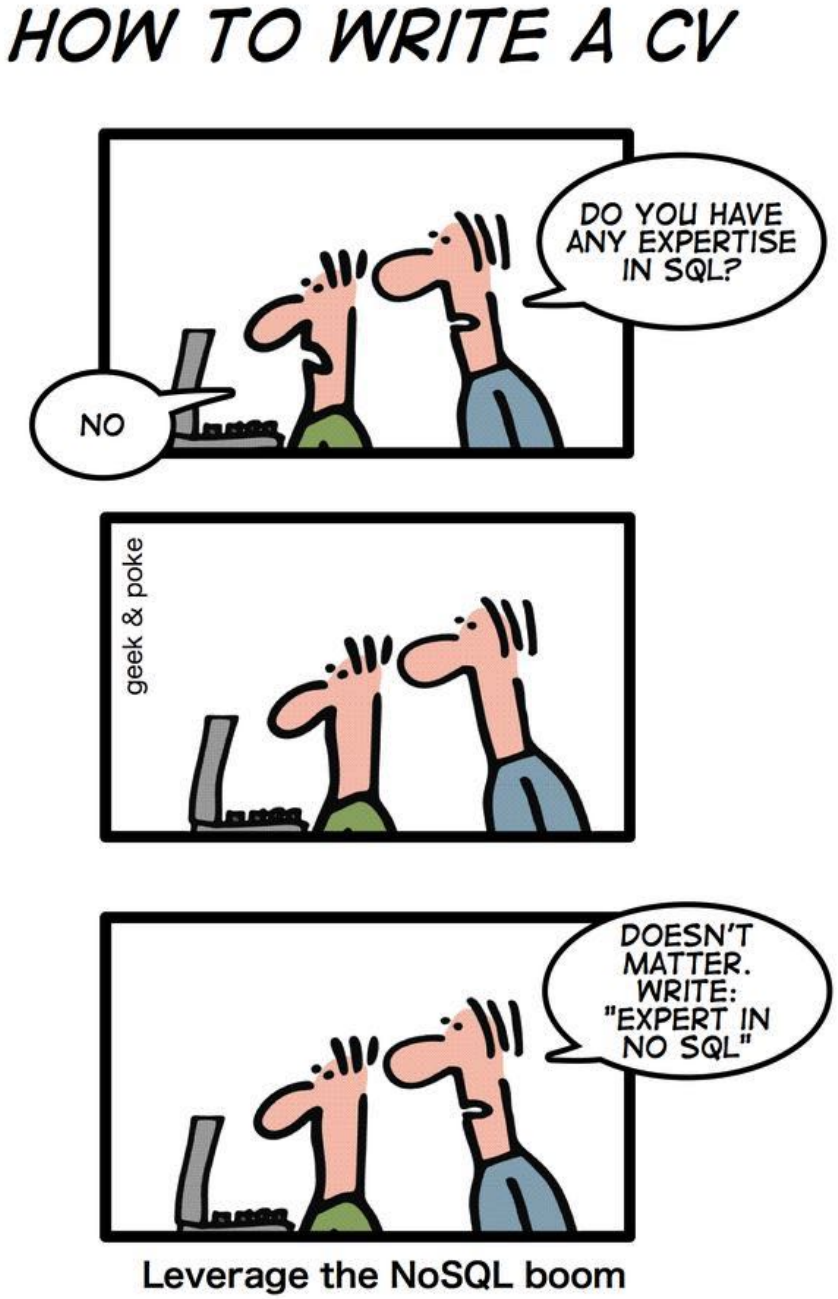

Source: Geek and Poke: http://geekandpoke.typepad.com/geekandpoke/2011/01/nosql.html

Common NoSQL Features

- Non-relational data models
- Flexible structure
  - No need to fix a schema, attributes can be added and replaced on the fly
- Massive read/write performance; availability via horizontal scaling
  - Replication and sharding (data partitioning)
  - Potentially thousands of machines worldwide
- Open source (very often)
- APIs to impose locality

When the database grows: Partitioning Tables

Key	Product Name	Short Description	Review	Picture
01	Americano @ Starbucks	Black, no sugar	I'd buy again	
02	BB @ Seattle's Best	Black, no sugar	The best	
03	TB @ Zoka Coffee	Black, no sugar	It's okay	
04	BC @ Coffee	Black, no sugar	Never again	

Source: http://cloudgirl.tech/data-partitioning-vertical-horizontal-hybrid-partitioning/

Vertical vs. Horisontal Partitioning

Key	Product Name	Short Description	Review	Picture
01	Americano @ Starbucks	Black, no sugar	I'd buy again	
02	BB @ Seattle's Best	Black, no sugar	The best	
03	TB @ Zoka Coffee	Black, no sugar	It's okay	
04	BC @ Coffee	Black, no sugar	Never again	

Key	Product Name	Review
01	Americano @ Starbucks	I'd buy again
02	BB @ Seattle's Best	The best
03	TB @ Zoka Coffee	It's okay
04	BC @ Coffee	Never again

Key	Short Description	Picture
01	Black, no sugar	
02	Black, no sugar	
03	Black, no sugar	
04	Black, no sugar	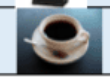

Source: http://cloudgirl.tech/data-partitio

Horizontal Partitioning ("sharding")

Key	Product Name	Short Description	Review	Picture
01	Americano @ Starbucks	Black, no sugar	I'd buy again	
02	BB @ Seattle's Best	Black, no sugar	The best	
03	TB @ Zoka Coffee	Black, no sugar	It's okay	
04	BC @ Coffee	Black, no sugar	Never again	

Key	Product Name	Short Description	Review	Picture
01	Americano @ Starbucks	Black, no sugar	I'd buy again	
02	BB @ Seattle's Best	Black, no sugar	The best	

Key	Product Name	Short Description	Review	Picture
03	TB @ Zoka Coffee	Black, no sugar	It's okay	
04	BC @ Coffee	Black, no sugar	Never again	

Source: http://cloudgirl.tech/data-partitioning-vertical-horizontal-hybrid-partitioning/

Vertical vs. Horizontal partitioning

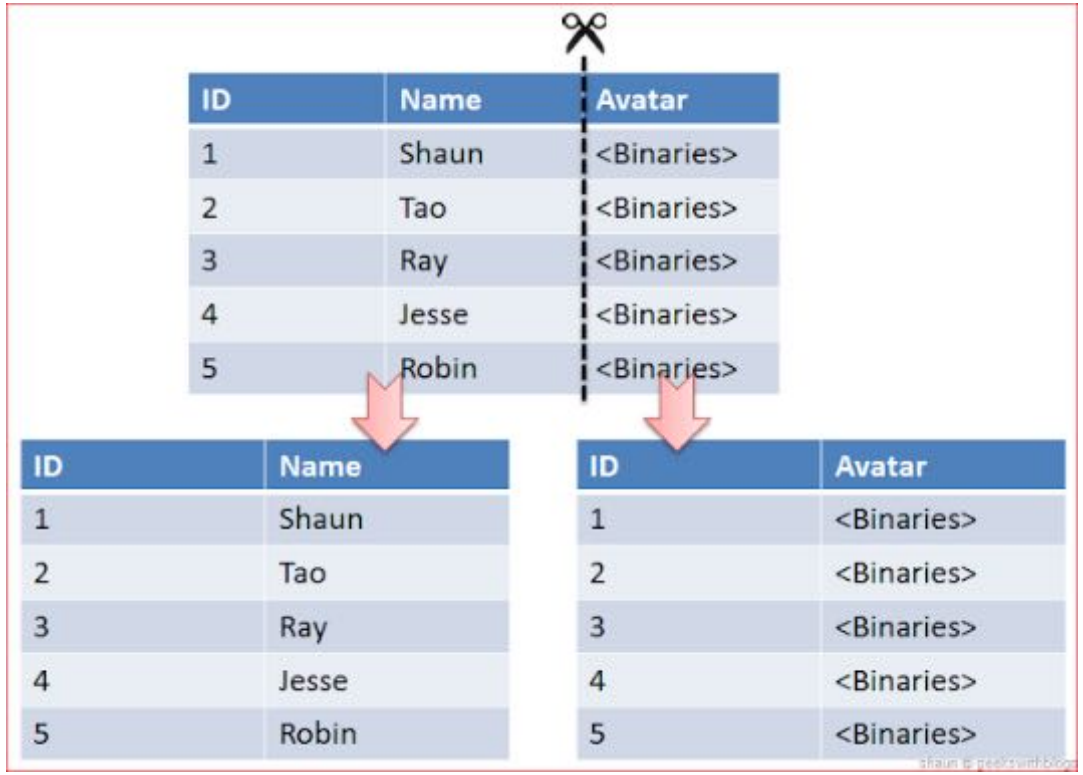

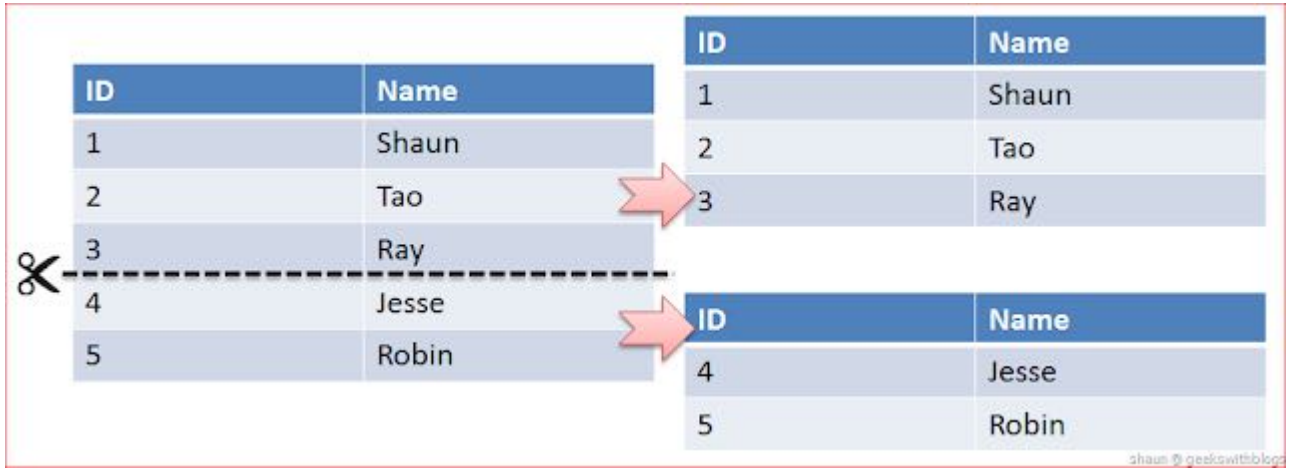

Source: http://www.piyushgupta.co.uk/2016/04/database-scaling-jargons.html, http://slideplayer.com/slide/12131436/70/images/17/SQL+Azure+Azure+Custom+Sharding.jpg

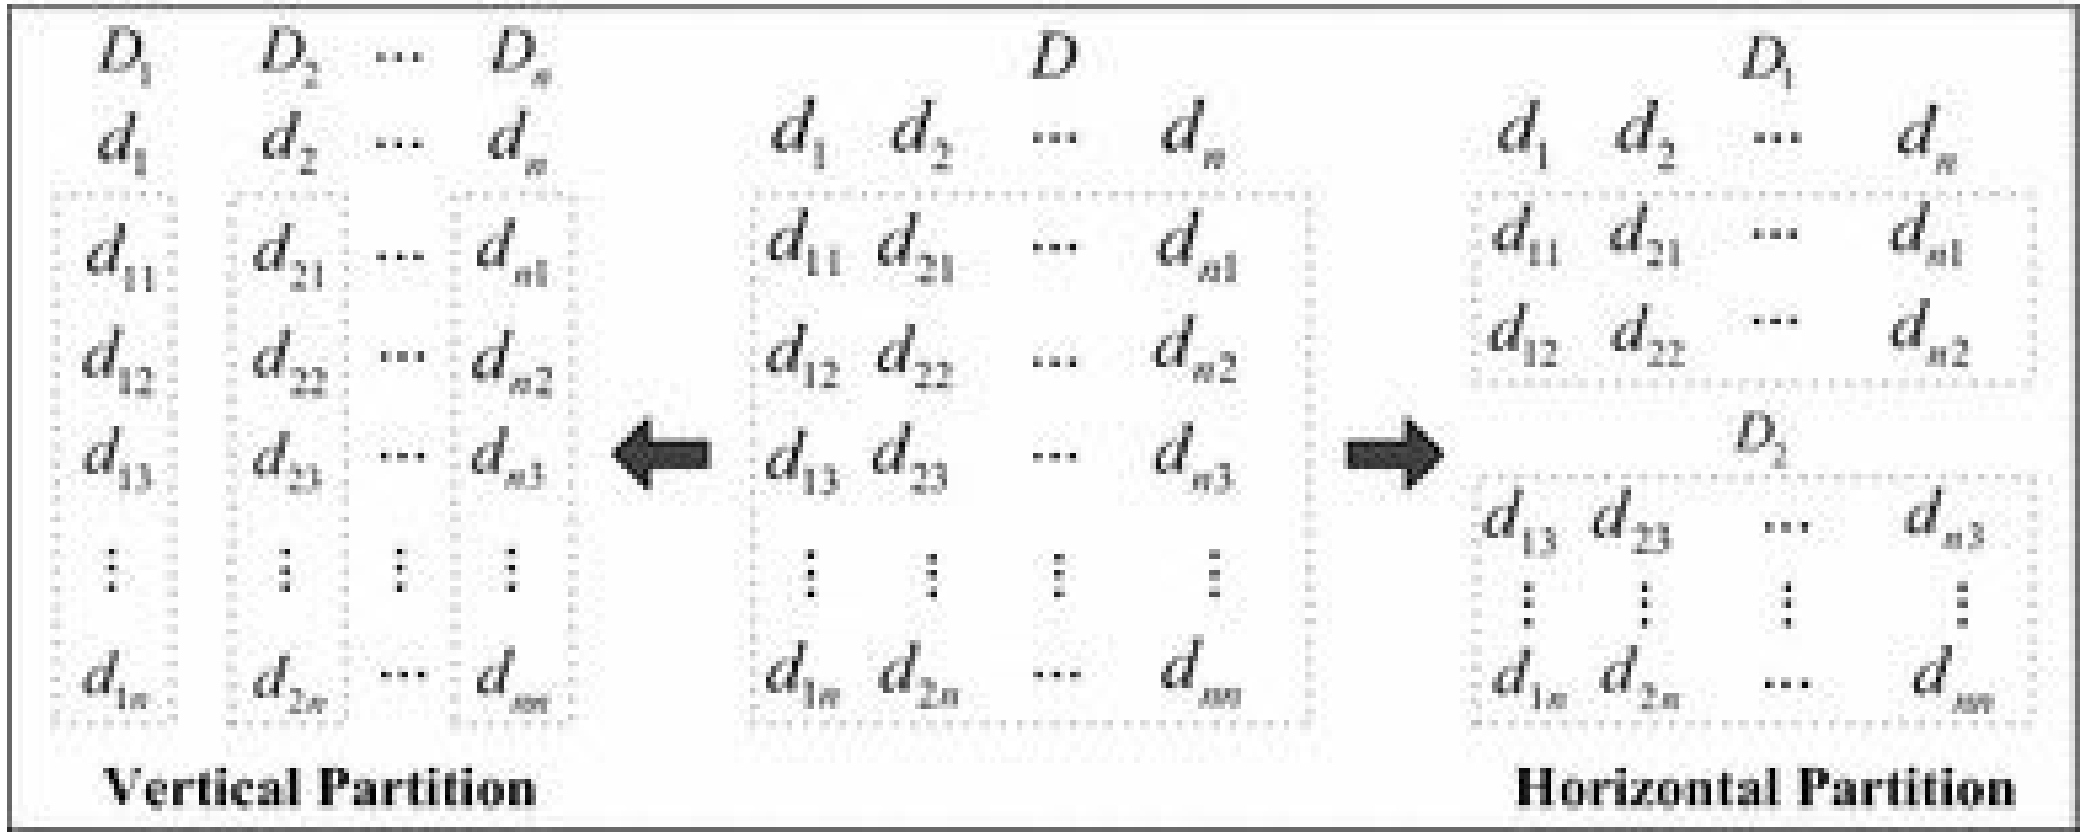

Source: An efficient scheme for probabilistic skyline queries over distributed uncertain data, Xiaoyong Li, Yijie Wang, Jie Yu