

From RDBMS to NoSQL



Efficient implementations of table **joins** and of **transactional** processing **require centralized** system.

NoSQL Databases:

- Database **schema** tailored for a **specific application**
 - **keep together** data pieces that are often accessed together
- Write operations might be slower but **read is fast**
- **Weaker consistency** guarantees

=> **efficiency** and horizontal **scalability**

Example (4): Aggregates



```
// collection "Customer"
{
  "customerID": 1,
  "name": "Jan Novák",
  "address": {
    "city": "Praha",
    "street": "Krásná 5",
    "ZIP": "111 00"
  }
}
```

```
// collection "Invoice"
{
  "invoiceID": 2015003,
  "orderNumber": 11,
  "bankAccount": "64640439/0100",
  "paymentDate": "2015-04-16",
  "address": {
    "city": "Brno",
    "street": "Slunečná 7",
    "ZIP": "602 00"
  }
}
```

```
// collection "Order"
{
  "orderNumber": 11,
  "date": "2015-04-01",
  "customerID": 1,
  "orderItems": [
    {
      "productID": 111,
      "name": "Vysavač ETA E1490",
      "quantity": 1,
      "price": 1300
    },
    {
      "productID": 112,
      "name": "Sáček k ETA E1490",
      "quantity": 10,
      "price": 300
    }
  ],
  "invoice": { "bankAccount": ..., ...}
}
```

NoSQL Databases: Aggregate-oriented



Many **NoSQL** stores are **aggregate-oriented**:

- There is **no general strategy** to set **aggregate** boundaries
- **Aggregates** give the database information about which bits of data will be **manipulated together**
 - What should be stored on the same node
- **Minimize** the number of nodes **accessed** during a search
- Impact on **concurrency** control:
 - NoSQL databases typically support **atomic** manipulation of a single **aggregate** at a time

Agenda



- Fundamentals of RDBMs and NoSQL Databases
- Data Model of Aggregates
- Models of **Data Distribution**
 - scalability
 - **sharding** vs. **replication**: master-slave, peer-to-peer
 - combination
- Consistency
 - write-write vs. read-write conflict
 - strategies and techniques
 - relaxing consistency

Vertical Scalability (Scaling up)



- Involve **larger** and more **powerful** machines
 - large disk storage using **disk arrays**
 - massively **parallel** architectures
 - large **main memories**
- Traditional choice
 - in favour of **strong consistency**
 - very **simple** to realize (**no** handling of data **distribution**)
- Works in many cases **but...**

Vertical Scalability: Drawbacks



- Higher **costs**
 - Large machines **cost more** than equivalent commodity HW
- Data growth **limit**
 - Large machine works well until the data grows to fill it
 - Even the **largest** of machines **has a limit**
- **Proactive** provisioning
 - In the beginning, **no idea** of the **final scale** of the application
 - An **upfront budget** is needed when scaling vertically
- **Vendor** lock-in
 - Large machines are produced by **a few vendors**
 - Customer is **dependent on** a single vendor (proprietary HW)

Distribution Models: Overview



- for horizontal scalability
- **Two** generic ways of data **distribution**:
 - **Replication** – the same data is copied over multiple nodes
 - Master-slave vs. peer-to-peer
 - **Sharding** – different data chunks are put on different nodes (data partitioning)
 - Master-master
- We can use either or **combine them**
 - **Distribution models** = specific ways to do **sharding**, **replication** or combination of both

Distribution Model: Single Server



- Running the database on a **single machine** is always the **preferred** scenario
 - it **saves** us a lot of **problems**
- It can **make sense** to use a NoSQL database on a single server
 - Other **advantages remain**: Flexible data model, simplicity
 - **Graph databases**: If the graph is “almost” complete, it is difficult to distribute it

Transaction Processing in NoSQL




- Basically, no problem if the DB is **centralized**
 - ACID can be implemented
 - Various **levels of isolation** (details later in the course)
 - read uncommitted
 - read committed
 - repeatable reads
 - serializable
- **Distributed transactions** (details later in the course)
 - X/Open Distributed Trans. Processing Model (X/Open XA)
 - Two-phase Commit Protocol (**2PC**)
 - Strong Strict Two-phase Locking (SS2PL)

PC: Partition Tolerance & Consistency




Example: two **users**, two **masters**, two **write** attempts



- **Strong** consistency:
 - Before the write is committed, **both** nodes have to **agree** on the order of the writes

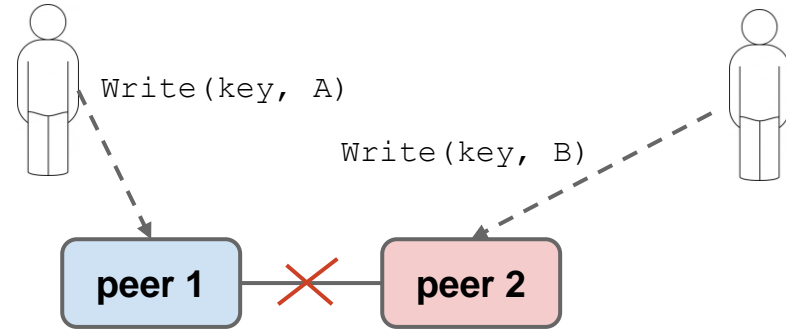
- If the nodes are **partitioned**, we are **losing Availability**
 - (but reads are still available)



PA: Partition Tolerance & Availability



- Choosing **Availability**:
 - Peer-to-peer replication
 - Eventual consistency



- In case of Partitioning
 - All requests are answered (full Availability)
 - We risk **losing consistency** guarantees completely
- But we can do something in the middle: **Quorum**
 - for replication consistency

Relaxing Durability II



- **Replication durability** (of a write operation)
 - The writing node can either
 1. **acknowledge** (answer) the write operation **immediately**
 - not wait until spread to other replicas
 - if the writing node crashes before spreading, durability fails
 - write-**behind** (write-back)
 2. or it can first spread the update to other replicas
 - operation is answered only after acknowledgement from the others
 - write-**through**
 - both variants are possible for P2P repl., master-slave replication, quora...

References



- I. Holubová, J. Kosek, K. Minařík, D. Novák. Big Data a NoSQL databáze. Praha: Grada Publishing, 2015. 288 p.
- Sadalage, P. J., & Fowler, M. (2012). NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence. Addison-Wesley Professional, 192 p.
- doc. RNDr. Irena Holubova, Ph.D. MMF UK course NDBI040: Big Data Management and NoSQL Databases
- Eric Brewer: Towards Robust Distributed Systems.
www.cs.berkeley.edu/~brewer/cs262b-2004/PODC-keynote.pdf