



# NoSQL

"Not Only SQL"

# Agenda

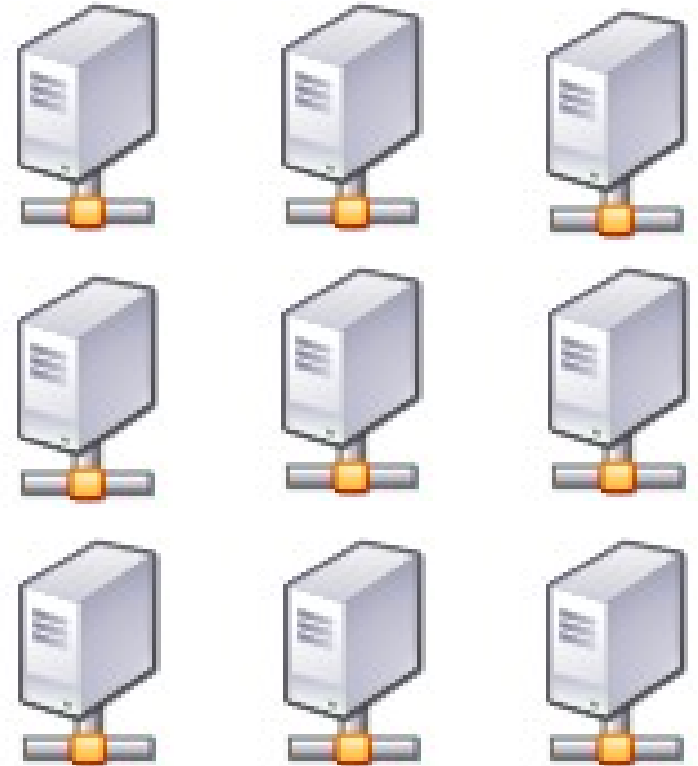
- Merkmale von NoSQL Systemen
- NoSQL System Typen
- Begriffe aus der NoSQL Welt
  - BASE, CAP, MapReduce
- Beispiele von NoSQL Systemen
  - Mit API und Live Demos
- Fazit
- Fragen



# Merkmale

- Schema frei
- Horizontal skalierbar
- Eventually consistent
- Nicht relational
- Verteilt
- Simple API

# Vertikale vs. Horizontale Skalierung





# Typen

- Document Store
- Wide Column Store
- Key Value / Tuple Store
- Graph Databases
- Object Databases
- XML Databases

# ACID vs. BASE

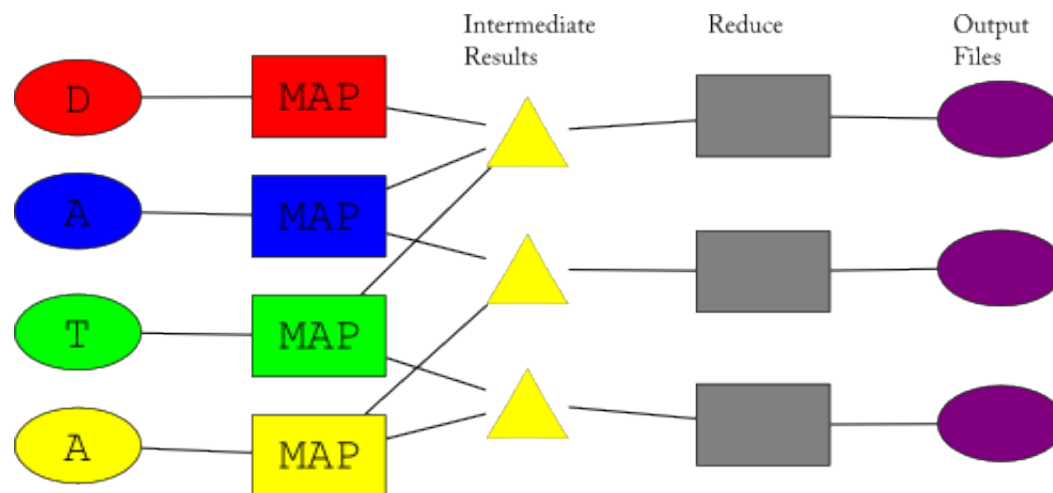
- Atomicity
    - Alles oder nichts
  - Consistency
    - Integrität
  - Isolation
    - Kapselung von Transaktionen
  - Durability
    - Persistenz aller Änderungen
- Basically Available
  - Soft-state
  - Eventual consistency
    - Zeitfenster mit Inkonsistenz

# CAP Theorem

- C – Consistency
- A – Availability
- P – Partition Tolerance
- You can only have 2 at a time
- **Make your choice!**

# Map Reduce

- Map the list of key-value pairs from one domain to another
- Group them by key
- Reduce to another list of key-value pairs





# HBase

- Unterprojekt von Apache Hadoop
- Aktuelle Version: 0.20.6
- Typ: Wide Column Store
  - „HBase uses a data model very similar to that of Bigtable. Users store data rows in labelled tables. A data row has a sortable key and an arbitrary number of columns. The table is stored sparsely, so that rows in the same table can have crazily-varying columns, if the user likes. „ - Hbase Wiki
- Geschrieben in Java
- <http://hbase.apache.org/>

# HBase API

The image shows a screenshot of an IDE (Eclipse) with the following components:

- Left Panel (Project Explorer):** Shows a project named 'HBaseExample' with a package structure including 'org.apache.hadoop.hbase', 'org.apache.hadoop.hbase.client', and 'org.apache.hadoop.hbase.mapreduce'.
- Center Panel (Code Editor):** Contains Java code for creating an HBase instance and performing operations. The code includes:

```
Wave waveProperty = (Wave) context.get("wave");
Index indexWave = (Index) context.get("index");

// Create HBase instance
HBaseConfiguration conf = new HBaseConfiguration();
conf.set("zookeeper.quorum", "zoo1-10.10.10.10:2181,zoo2-10.10.10.10:2181,zoo3-10.10.10.10:2181");
conf.set("hbase.rootdir", "hdfs://10.10.10.10:8020/hbase");
conf.set("hbase.cluster.distributed", "true");
conf.set("hbase.tmplocaldir", "/tmp/hbase-");

HBase hbase = new HBase(conf);

// Create table
TableName tableName = TableName.valueOf("mytable");
HTableUtil.createTable(hbase, tableName, new byte[][]{{"rowkey", "value"}}, 1);

// Get table
HTable table = hbase.getTable(tableName);

// Get row
byte[] rowKey = {"rowkey", "value"};
byte[] value = table.get(rowKey);
```
- Right Panel (Outline):** Shows the class hierarchy for 'org.apache.hadoop.hbase.client.HTableUtil', including methods like 'createTable', 'getTable', and 'get'. A tooltip is visible over the 'createTable' method, showing its signature: 'void createTable(HBase hbase, TableName tableName, byte[][] columnFamilies, int replicasPerRegion)'. The tooltip also includes a description: 'Creates a new table with the specified name and column families. The table is created with the specified number of replicas per region. The column families are specified as an array of byte arrays. The replicas per region is specified as an integer.' The tooltip also shows the 'Parameters' section: 'tableName: The name of the table to be created.' and 'columnFamilies: The column families to be created.'
- Bottom Panel (Problems):** Shows a list of errors and warnings. The first error is 'The project org.apache.hadoop.conf.Configuration is never used'. The second error is 'The project org.apache.hadoop.conf.Configuration is never used'. The third error is 'The project org.apache.hadoop.conf.Configuration is never used'.

# neo4j

- Typ: Graphendatenbank
- Aktuelle Version: 1.0
- Lizenz: Apache 2.0
- Geschrieben in Java
- <http://neo4j.org/>

# neo4j API

```
new neoProperty("KEY_TITLE", "Google Wave");
new neoProperty("MAX_ATTENDEES", 30);
Index index(wave, KEY_TITLE, "Google Wave");

johnson.neoRelationship(relationshipTo(undroid, NeoRelationshipTypes.ATTENDES));
johnson.neoRelationship(relationshipTo(tougrantsocialist, NeoRelationshipTypes.ATTENDES));
johnson.neoRelationship(relationshipTo(jacob, NeoRelationshipTypes.ATTENDES));
johnson.neoRelationship(relationshipTo(yanev, NeoRelationshipTypes.ATTENDES));

neoCreateRelationship(relationshipTo(jacob, NeoRelationshipTypes.ATTENDES));
neoCreateRelationship(relationshipTo(jacob, NeoRelationshipTypes.ATTENDES));
neoCreateRelationship(relationshipTo(jacob, NeoRelationshipTypes.ATTENDES));

neo.attendees();
```

**DEMO**

void java.io.PrintStream.println(String s)

| Errors        | Severity | Details                   | Resource      | Type         |
|---------------|----------|---------------------------|---------------|--------------|
| neo4j-example | Error    | Configuration is not used | neo4j-example | Java Problem |
| neo4j-example | Error    | File not found            | neo4j-example | Java Problem |
| neo4j-example | Error    | File not found            | neo4j-example | Java Problem |

# MongoDB

- Typ: Dokumentenorientiert
- Speicherformat: BSON (Binary JSON)
- Aktuelle Version: 1.4.4
- Lizenz: GNU AGPL
- Geschrieben in C
- <http://www.mongodb.org/>

# MongoDB API

The image shows an IDE window with the following content:

```
Wave.addProperty("KEY_TITLE", "Google Wave");
Wave.addProperty("MAX_ATTENDEES", 50);
Index.indexOnWave, KEY_TITLE, "Google Wave");

johndoe.createRelationshipTo(andy, MyRelationshipTypes.ATTENDEE);
johndoe.createRelationshipTo(steve, MyRelationshipTypes.ATTENDEE);
johndoe.createRelationshipTo(joe, MyRelationshipTypes.ATTENDEE);
johndoe.createRelationshipTo(amy, MyRelationshipTypes.ATTENDEE);

amy.createRelationshipTo(johndoe, MyRelationshipTypes.ATTENDEE);
amy.createRelationshipTo(steve, MyRelationshipTypes.ATTENDEE);
amy.createRelationshipTo(joe, MyRelationshipTypes.ATTENDEE);

// ... more code ...

void java.io.PrintStream.println(String x)
```

The IDE also shows a project tree on the left and a class hierarchy on the right. A large "DEMO" watermark is overlaid on the code.

# Voldemort

- Typ: Key Value Store
- Aktuelle Version: 0.81
- Lizenz: Apache License 2.0
- Geschrieben in Java
- <http://project-voldemort.com/>

# Voldemort API

The screenshot shows an IDE window with the following content:

```
Wave.setProperties(KEY_TITLE, "Google Wave");
Wave.setProperties("maxAttendees", 30);
Index.index(wave, KEY_TITLE, "Google Wave");

jobDone.addChangeListener(to(android, MyRelationshipTypes.ATTENDEE);
jobDone.addChangeListener(to(fragment, MyRelationshipTypes.ATTENDEE);
jobDone.addChangeListener(to(job, MyRelationshipTypes.ATTENDEE);
jobDone.addChangeListener(to(wave, MyRelationshipTypes.ATTENDEE);

myContext.setRelationshipType(myContext, MyRelationshipTypes.ATTENDEE);
myContext.setRelationshipType(myContext, MyRelationshipTypes.ATTENDEE);
myContext.setRelationshipType(myContext, MyRelationshipTypes.ATTENDEE);

myAttendee();

// ... more code ...

System.out.println("Found for " + wave.getTitle());
System.out.println("void java.io.PrintStream.println(String s)");
IndexHelper Prints a String and then terminate the line. This method behaves as though it invokes println
(System.out, (String) null, false);
// ... more code ...
```

**DEMO**

The IDE also shows a Package Explorer on the left with a tree structure including 'VoldemortExample'. On the right, there is an Outline view showing the class structure for 'VoldemortExample' with methods like 'getAttendees', 'setAttendees', and 'index'.

At the bottom, there is a Problems view showing a list of errors:

| Severity | Message  | Resource  | Location                            | Type         |
|----------|--|-----------|-------------------------------------|--------------|
| Warning  | Warning (3 items)  |           |                                     |              |
| Warning  | The project org.voldemort.hadoop.conf.Configuration is never used            | Wave.java | File: voldemort.../org.voldemort... | Java Problem |
| Warning  | Could not find source file for class org.voldemort.hadoop.conf.Configuration | Wave.java | File: voldemort.../org.voldemort... | Java Problem |
| Warning  | Could not find source file for class org.voldemort.hadoop.conf.Configuration | Wave.java | File: voldemort.../org.voldemort... | Java Problem |





# Weitere NoSQL Systeme

- Riak
- CouchDB
- Cassandra
- Redis
- Berkley DB
- ...

# Fazit

- Keine einheitliche API
  - Teilweise Plugins für JPA
- Blick über den Tellerrand
  - Denken jenseits von relationalem Datenmodell
- Das richtige Werkzeug für den richtigen Zweck

# Nächste Termine

- 24.08.2010 um 20:00 Uhr:
  - Java Champion Adam Bien "Stop Talking - Start Hacking ...with Java EE 6"
  
- 06.10.2010 um 18:30 Uhr:
  - Developer Evangelist Jochen Hiller "Android-Apps selbst schreiben: Techniken, Tools und Tipps"





Fragen?