

BigData-Systeme mit Apache Cassandra

Taking databases to a next level

im Juni 2014





CURSED
ORACLE®

1. Wenn relationale Datenbanken an ihre Grenzen kommen

„Meine MySQL Datenbank funktioniert doch!“

2. Was Cassandra anders macht

„Wo sind meine Joins?“

3. Der Datastax Cassandra Treiber

„Das ist ja wie JDBC!“



„Oracle SQL verhindert den weltweiten Datenaustausch“

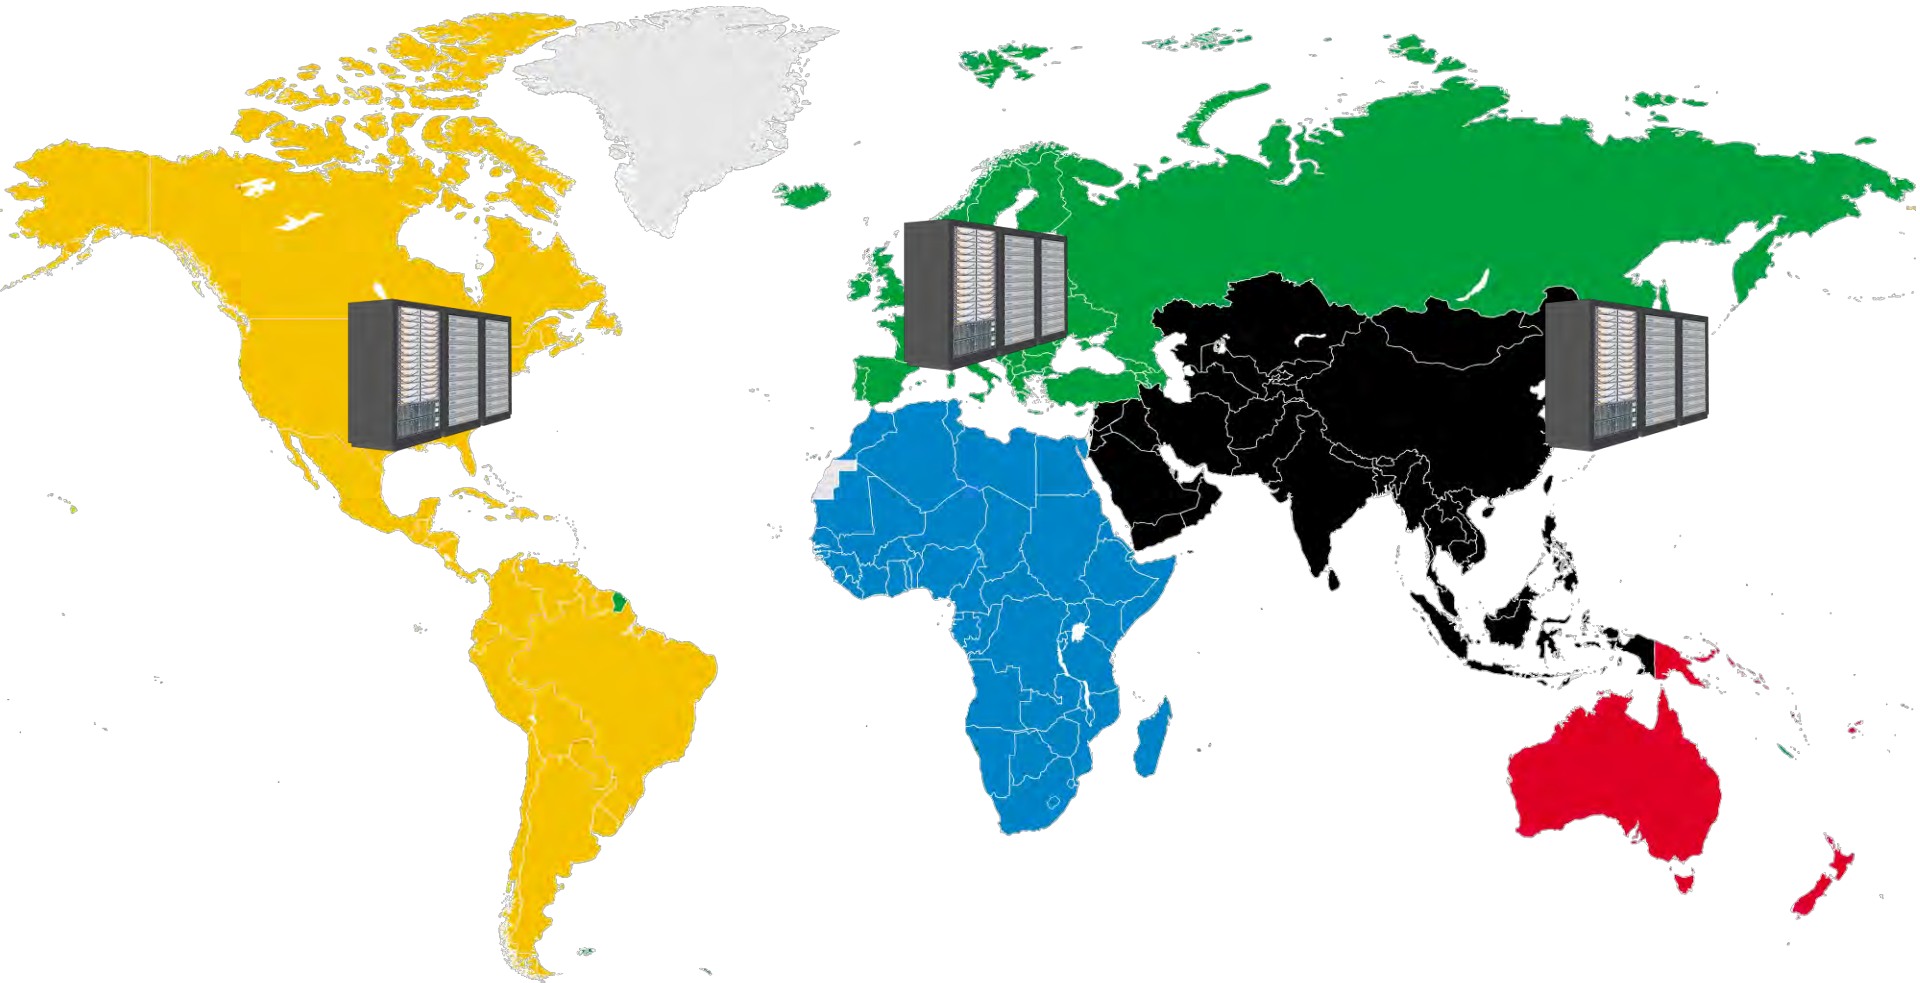
„Wir benötigen regelmäßige Downtimes für Schema Änderungen“



„Relationale Datenbanken sind nicht ausgelegt auf unsere Datenmengen“

Kosteneffizienz?

*„Oracle SQL verhindert
den weltweiten Datenaustausch“*



„Wir benötigen regelmäßige Downtimes für Schema Änderungen“

user_id	name	first_name
4711	Skywalker	Luke
4712	Solo	Han
4713	Chewbacca	null

user_id	name	first_name	homeworld
4711	Skywalker	Luke	Polis Massa
4712	Solo	Han	null
4713	Chewbacca	null	Kashyyyk

„Wir benötigen regelmäßige Downtimes für Schema Änderungen“

KEY

COLUMNS

4711	„name“ : „Skywalker“	„first_name“ : „Luke“	„homeworld“ : „Polis Massa“
4712	„name“ : „Solo“	„first_name“ : „Han“	
4713	„name“ : „Chewbacca“	„homeworld“ : „Kashyyyk“	

„Kosteneffizienz?“




4711

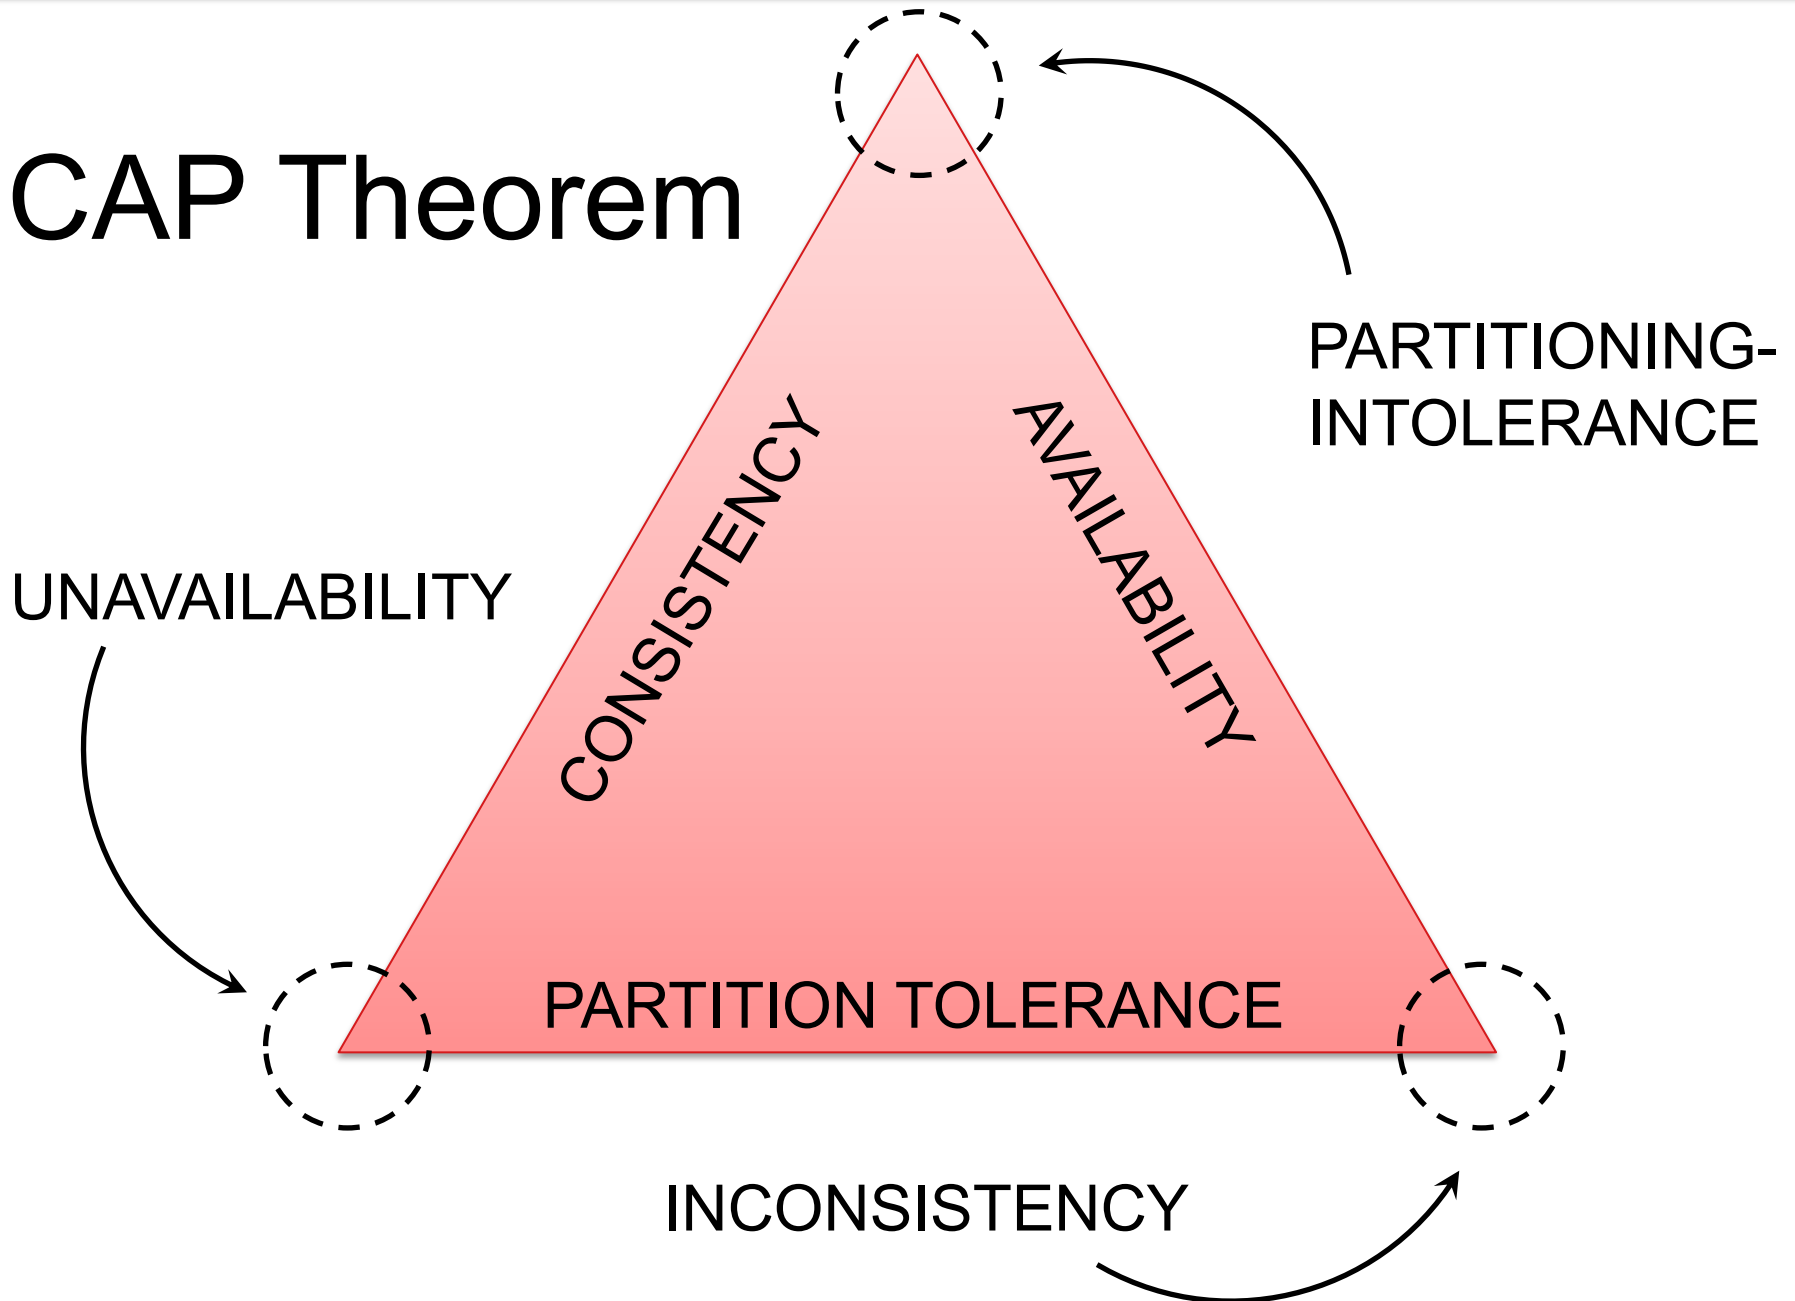
„name“ : „Skywalker“

„first_name“ : „Luke“

„homeworld“ : „Polis Massa“



CAP Theorem



ACID



BASE



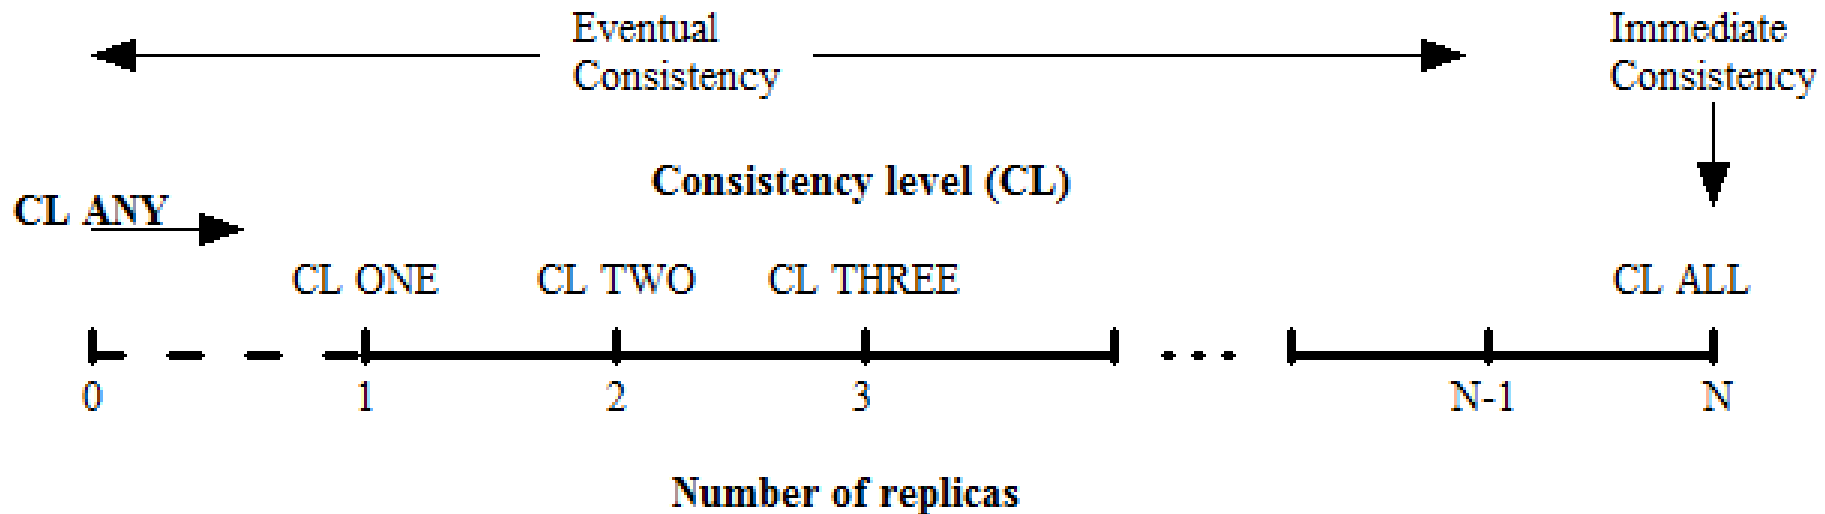
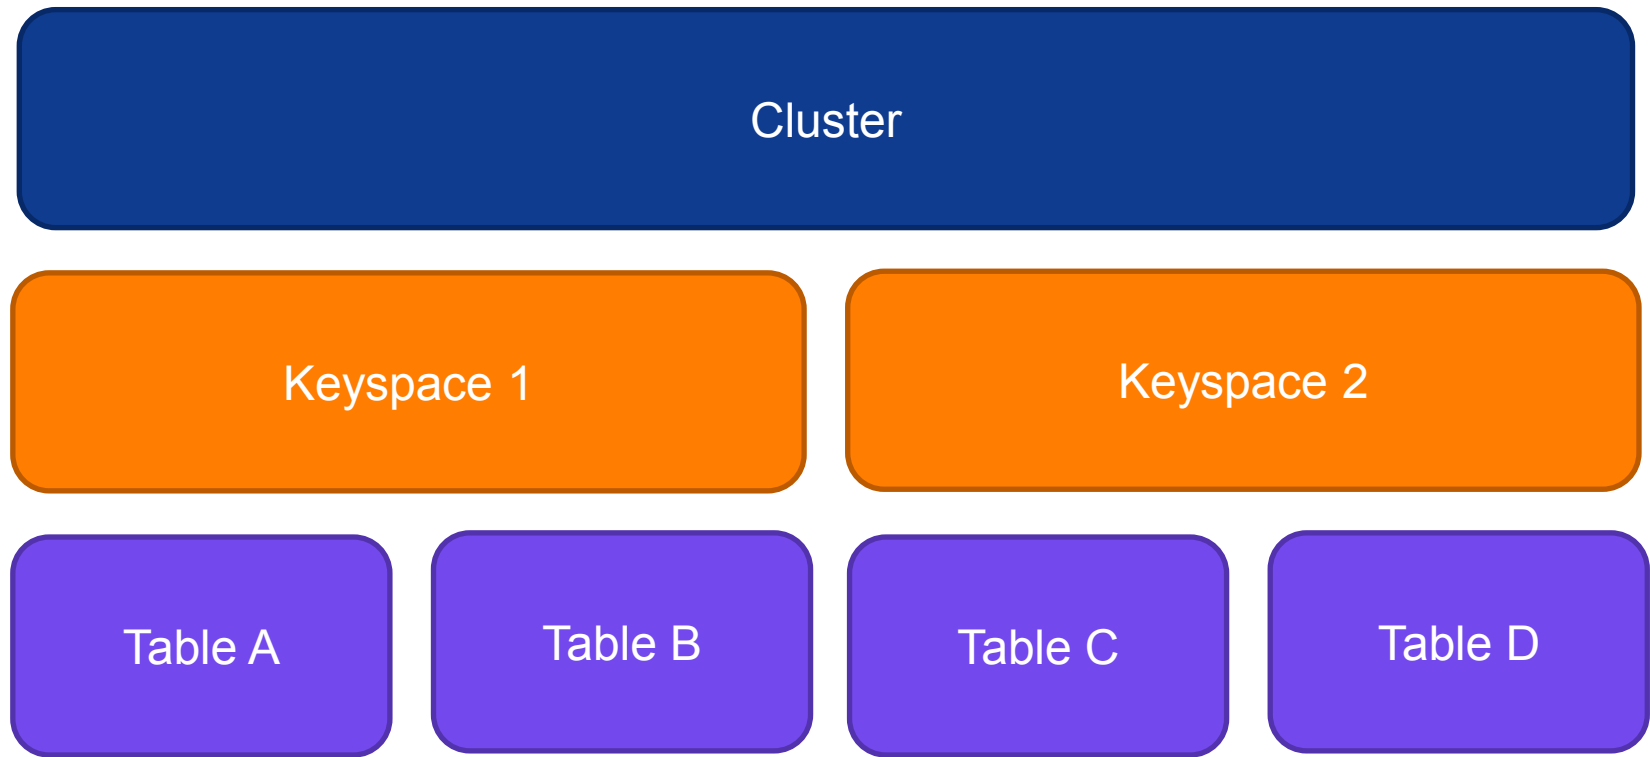


Illustration of Tunable Consistency

„CQL is exactly like SQL“
(except where it is not)



Keyspaces

```
CREATE KEYSPACE starwars
  WITH replication = {
    'class': 'SimpleStrategy',
    'replication_factor': 3
  }
```

Tabellen

```
CREATE TABLE user (  
    user_id uuid,           CQL  
    name text,  
    first_name text,  
    homeworld text,  
    PRIMARY KEY (user_id));
```

```
CREATE TABLE user (  
    user_id int,           SQL  
    name varchar(255),  
    first_name varchar(255),  
    homeworld varchar(255),  
    PRIMARY KEY (user_id));
```

INSERT INTO

SQL

```
INSERT INTO user (  
    user_id, name, first_name)  
VALUES (  
    4712, 'Solo', 'Han');
```

CQL

```
INSERT INTO user (  
    user_id, name, first_name)  
VALUES (  
    4712, 'Solo', 'Han');
```

SQL

SELECT

```
SELECT name, homeplanet  
FROM user  
WHERE user_id = 4712;
```

CQL

```
SELECT name, homeplanet  
FROM user  
WHERE user_id = 4712;
```



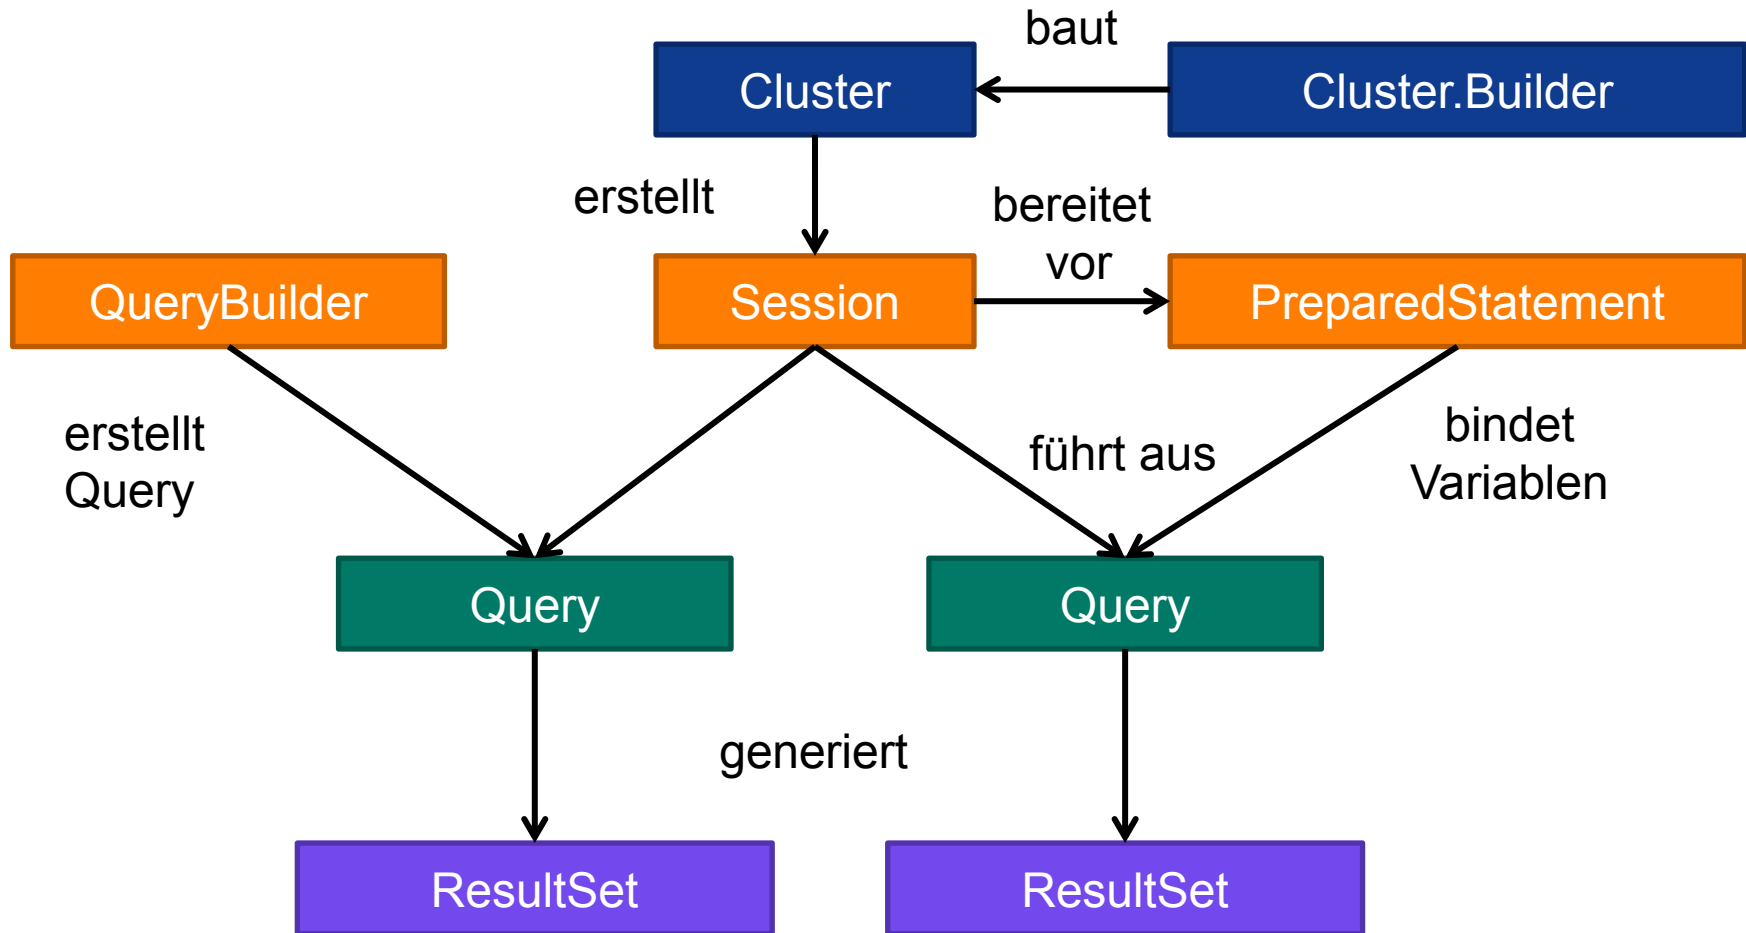
DB API

CQL API

OO API

CQL Native Protocol





Simple Cassandra Client

user_id	name	first_name	homeworld
4711	Skywalker	Luke	Polis Massa
4712	Solo	Han	null
4713	Chewbacca	null	Kashyyyk

maven

```
<dependency>  
  <groupId>com.datastax.cassandra</groupId>  
  <artifactId>  
    cassandra-driver-core</artifactId>  
  <version>2.0.1</version>  
</dependency>
```



1. Aufbau der Verbindung

```
import com.datastax.driver.core.Cluster;
import com.datastax.driver.core.Host;
import com.datastax.driver.core.Metadata;

public class SimpleClient {
    // ...
}
```

```
private Cluster cluster;
```

```
public void connect(String node) {  
    // ...  
}
```

```
public void close() {  
    // ...  
}
```

```
public static void main(String[] args) {  
    // ...  
}
```

```
public void connect(String node) {  
    cluster = Cluster.builder()  
        .addContactPoint(node)  
        .build();  
}
```

```
public void close() {  
    cluster.close();  
}
```

```
public static void main(String[] args) {  
    SimpleClient client = new SimpleClient();  
    client.connect("127.0.0.1");  
    client.close();  
}
```

2. Absetzen von Abfragen

```
private Session session;
```

```
public void connect(String node) {  
    cluster = Cluster.builder()  
        .addContactPoint(node)  
        .build();  
    // NEU  
    session = cluster.connect();  
}
```

```
public void createSchema() {}
```

```
public void loadData() {}
```

```
public void queryData() {}
```

```
public void createSchema() {
    session.execute(
        "CREATE KEYSPACE sw WITH replication "
        + "= {'class':'SimpleStrategy', "
        + "'replication_factor':3};"
    );

    session.execute(
        "CREATE TABLE sw.user (" +
        "user_id uuid PRIMARY KEY," +
        "name text," +
        "first_name text," +
        "homeworld text);"
    );
}
```

```
public void loadData() {  
    session.execute(  
        "INSERT INTO sw.user (user_id, name, " +  
        "first_name, homeworld) " +  
        "VALUES (" +  
        "756716f7-2e54-4715-9f00-91dcbea6cf50," +  
        "'Skywalker'," +  
        "'Luke'," +  
        "'Polis Massa')";  
    );  
}
```

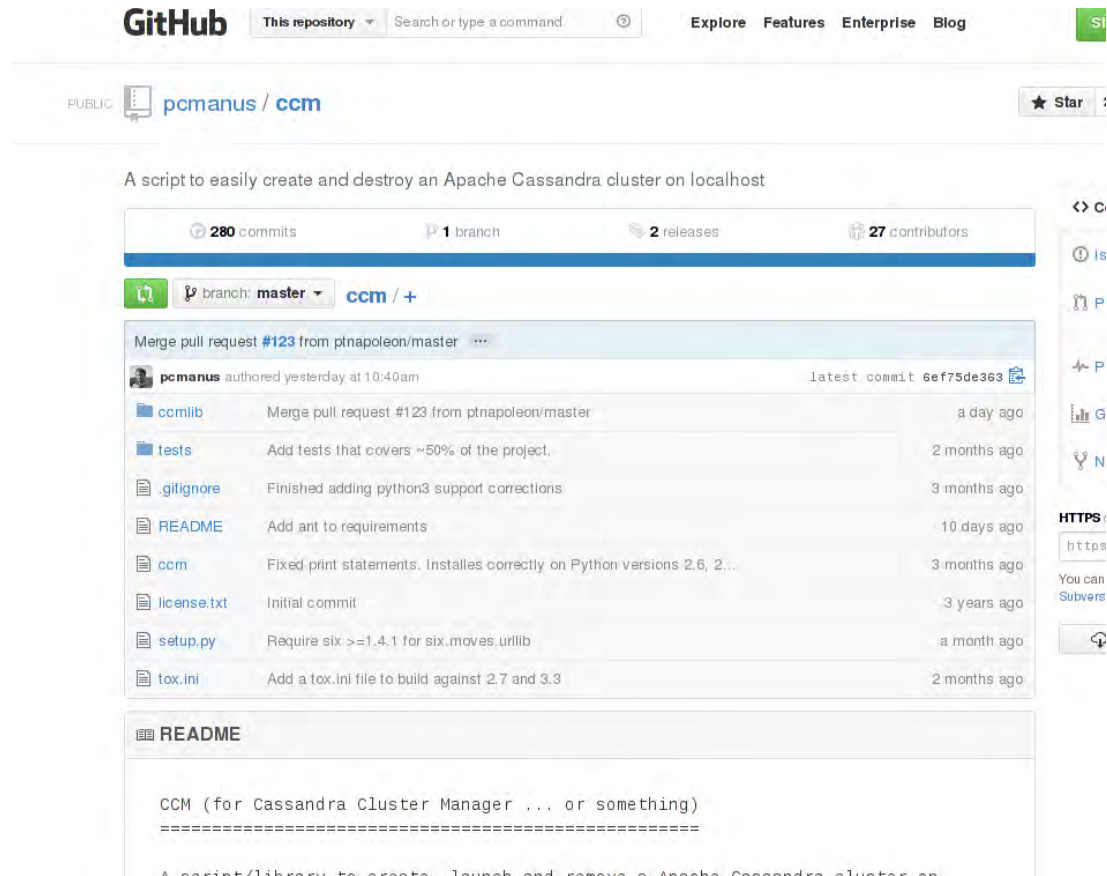


```
public void queryData() {  
    ResultSet results = session.execute(  
        "SELECT * FROM sw.user " +  
        "WHERE id = " +  
        "756716f7-2e54-4715-9f00-91dcbea6cf50;");  
  
    for (Row row : results) {  
        System.out.println(  
            row.getString("first_name") + " | " +  
            row.getString("name") + " | " +  
            row.getString("homeworld")  
        );  
    }  
}
```

3. Simple Client

```
public static void main(String[] args) {  
    SimpleClient client = new SimpleClient();  
    client.connect("127.0.0.1");  
  
    client.createSchema();  
    client.loadData();  
    client.queryData();  
  
    client.close();  
}
```

CCM – Cassandra Cluster Manager



GitHub This repository Search or type a command Explore Features Enterprise Blog

PUBLIC pcmanus / ccm ★ Star

A script to easily create and destroy an Apache Cassandra cluster on localhost

280 commits 1 branch 2 releases 27 contributors

branch: master ccm / +

Merge pull request #123 from ptnapoleon/master

pcmanus authored yesterday at 10:40am latest commit 6ef75de363

File	Description	Time ago
ccmlib	Merge pull request #123 from ptnapoleon/master	a day ago
tests	Add tests that covers ~50% of the project.	2 months ago
.gitignore	Finished adding python3 support corrections	3 months ago
README	Add ant to requirements	10 days ago
ccm	Fixed print statements. Installs correctly on Python versions 2.6, 2...	3 months ago
license.txt	Initial commit	3 years ago
setup.py	Require six >=1.4.1 for six.moves.urllib	a month ago
tox.ini	Add a tox.ini file to build against 2.7 and 3.3	2 months ago

README

```
CCM (for Cassandra Cluster Manager ... or something)
=====
A script/library to create, launch and remove a Apache Cassandra cluster on
```

<https://github.com/pcmanus/ccm>

QueryBuilder

```
public List<Row> getRows() {  
    Statement statement = QueryBuilder  
        .select()  
        .all()  
        .from("sw", "user")  
        .where(eq("user_id", "xyz"));  
    return getSession()  
        .execute(statement)  
        .all();  
}
```



Philipp Stussak, Consultant

Telefon: +49 (7044) 95103 – 100

Mail: philipp.stussak@widas.de

www.widas.de

