



Continuous Database Design

Entwicklertag Karlsruhe
Mai 2017

Orientation in Objects GmbH

Weinheimer Str. 68
68309 Mannheim

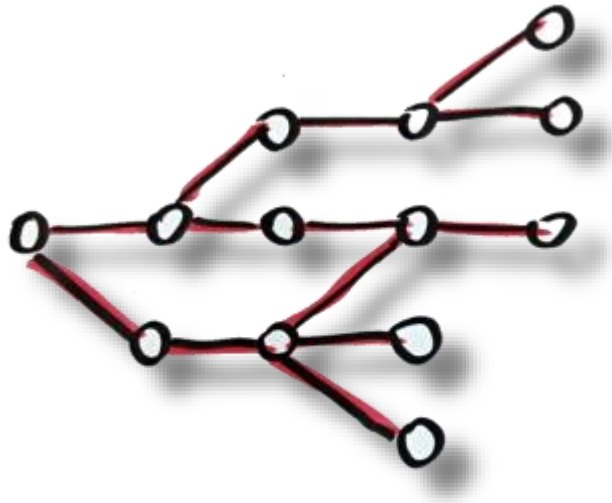
www.oio.de
info@oio.de

Thorsten Maier

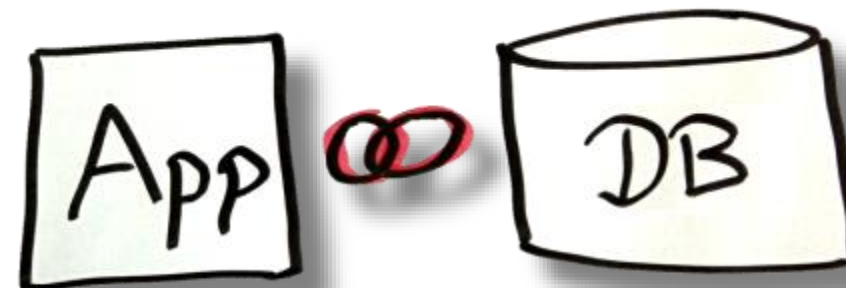
Trainer, Berater, Entwickler



 [@ThorstenMaier](#)

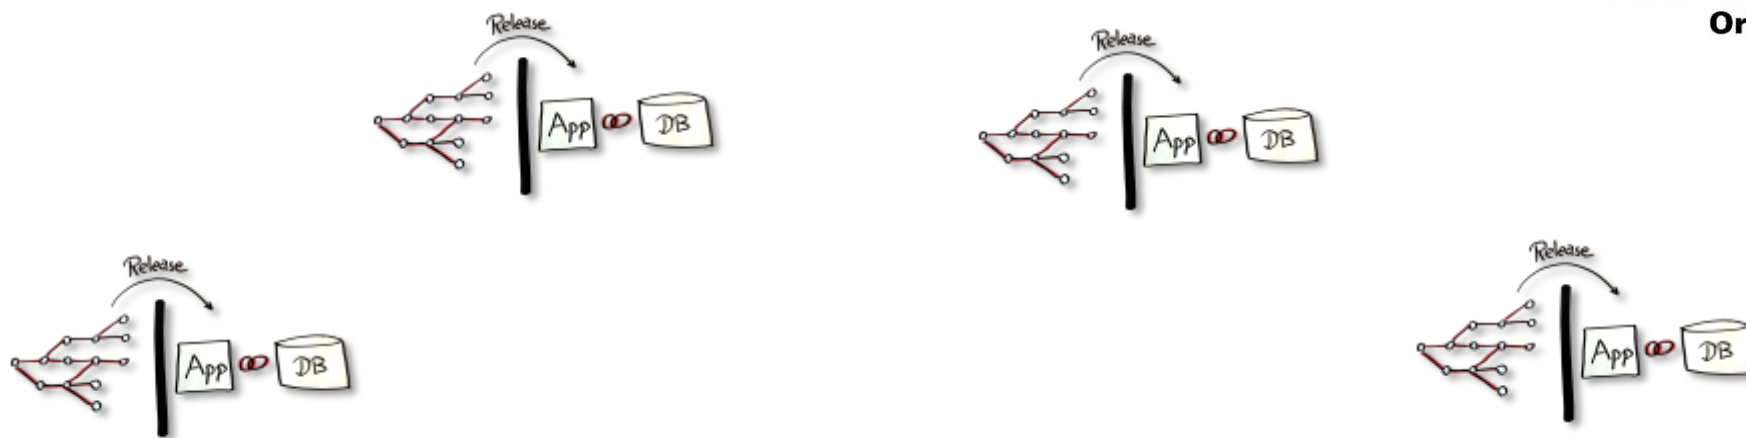


Die Welt der Entwickler

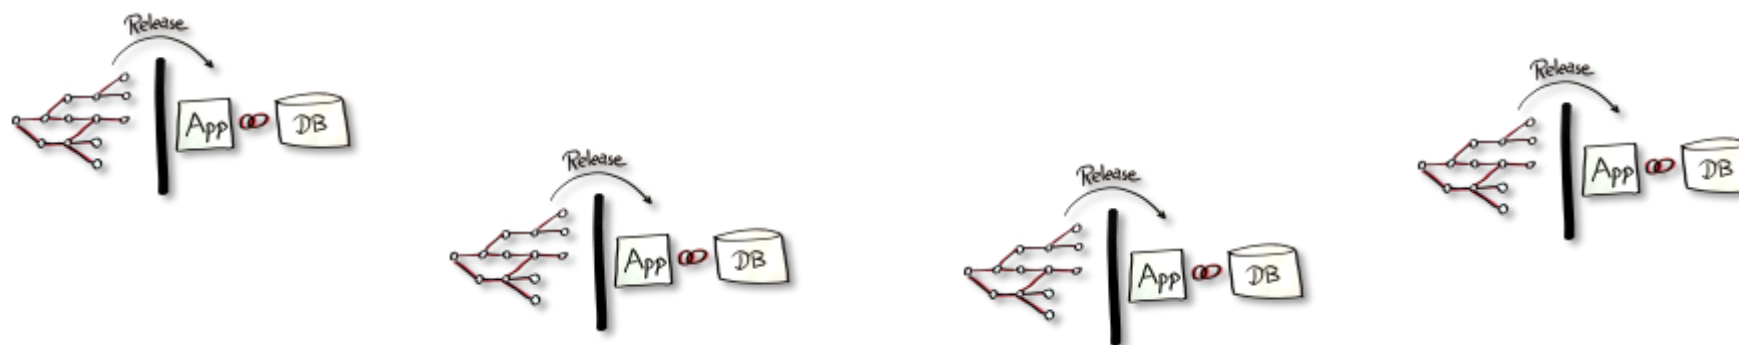


Die Welt des Betriebs





Continuous Delivery



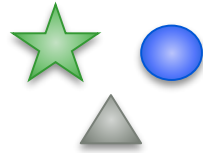
Wann Continuous Delivery?

Schnelle Reaktion auf Anforderungsänderungen

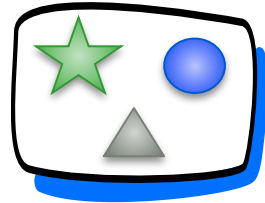
Was brauchen wir dazu?

Automatisierung

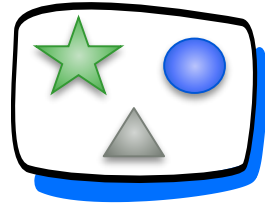
Modularisierung



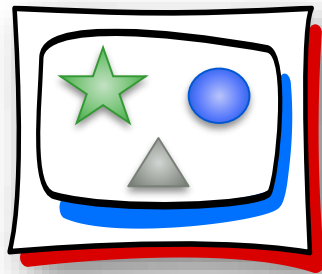
3 Bausteine



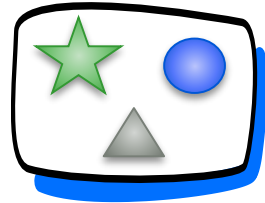
Anwendung mit 3 Bausteinen



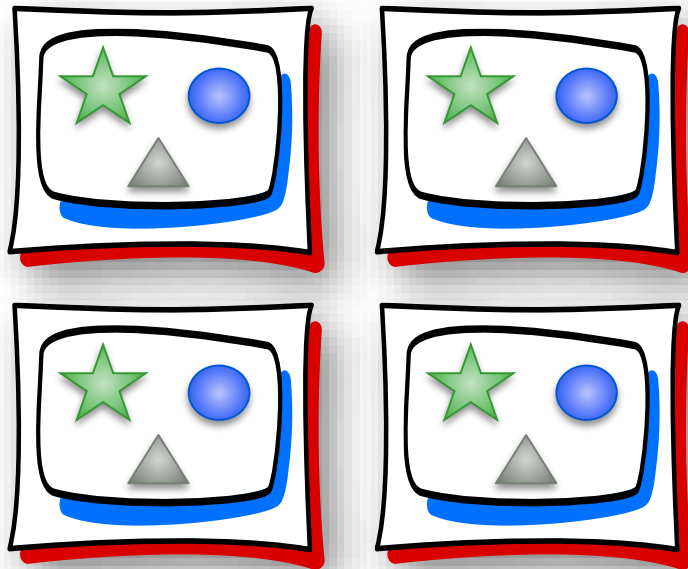
Anwendung mit 3 Bausteinen




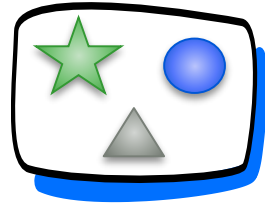
Deployment auf einem Server



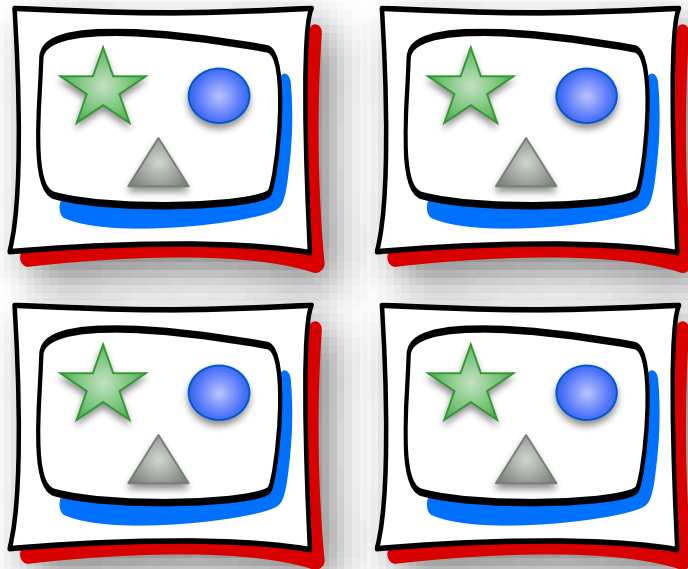
Anwendung mit 3 Bausteinen




Baustein  wird 4 mal benötigt



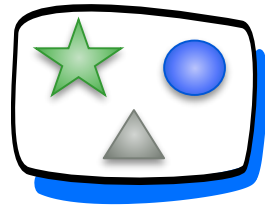
Anwendung mit 3 Bausteinen



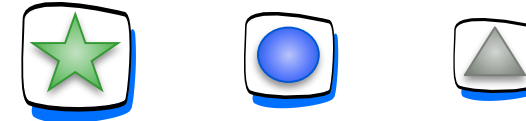
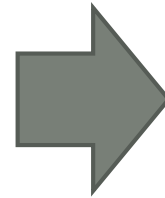
Baustein  wird 4 mal benötigt

Nachteile

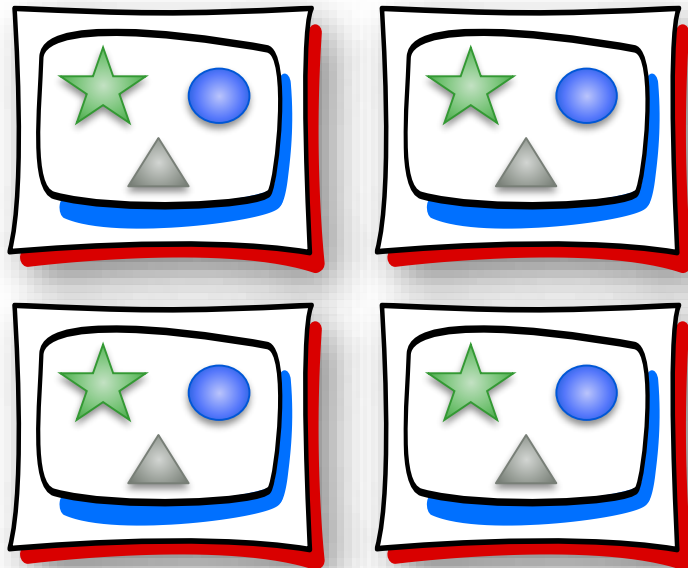
 wird nur 2 mal benötigt
Änderung an  erfordert Reploy von 




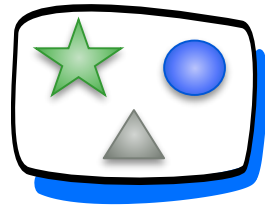
Anwendung mit 3 Bausteinen



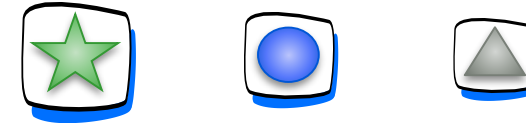
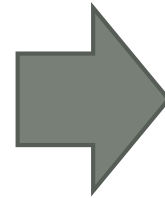
Unabhängige Artefakte



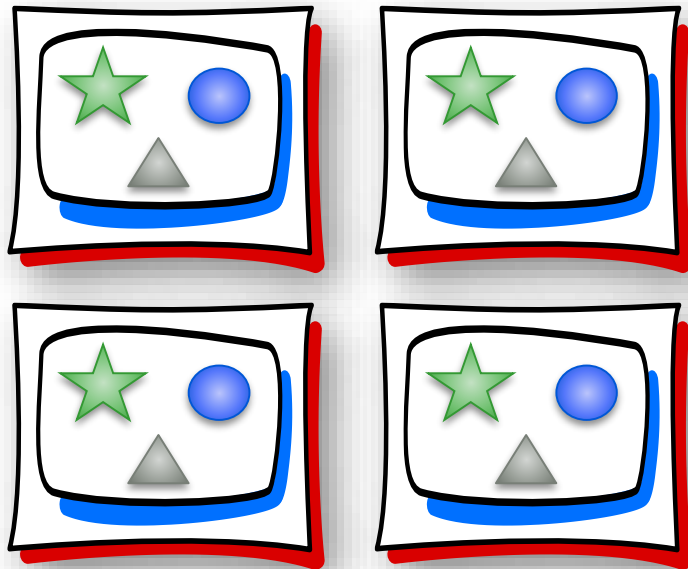
Baustein  wird 4 mal benötigt




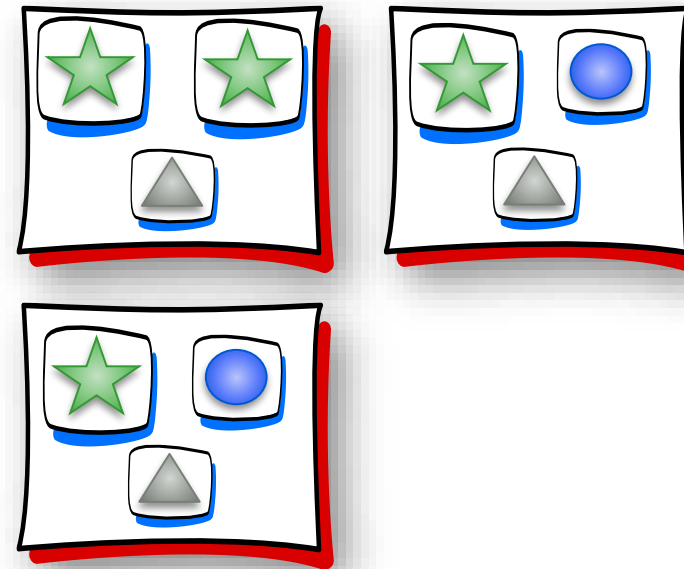
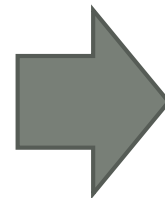
Anwendung mit 3 Bausteinen



Unabhängige Artefakte



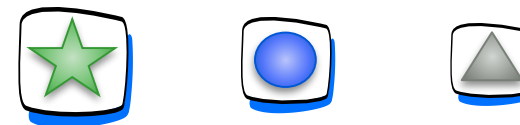
Baustein  wird 4 mal benötigt



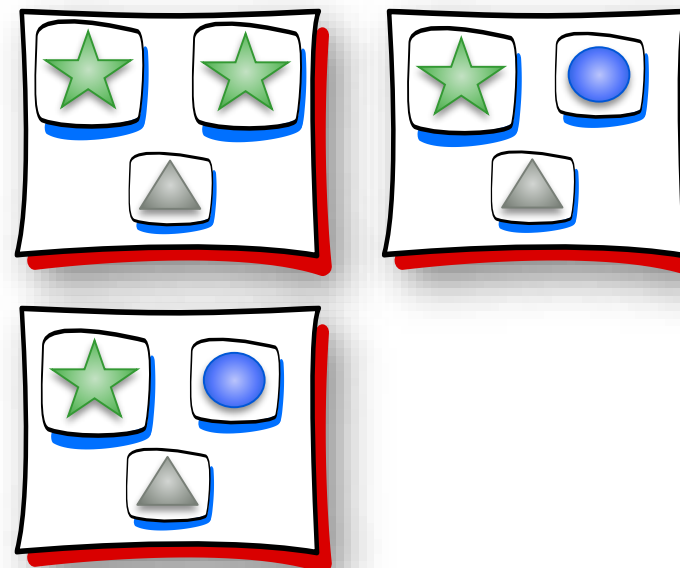
Flexible Skalierung auf 3 Server

Vorteile

Isolierte Entwicklung
Schnellere Reaktionszeiten
Austauschen im laufenden Betrieb
Continuous Delivery wird möglich

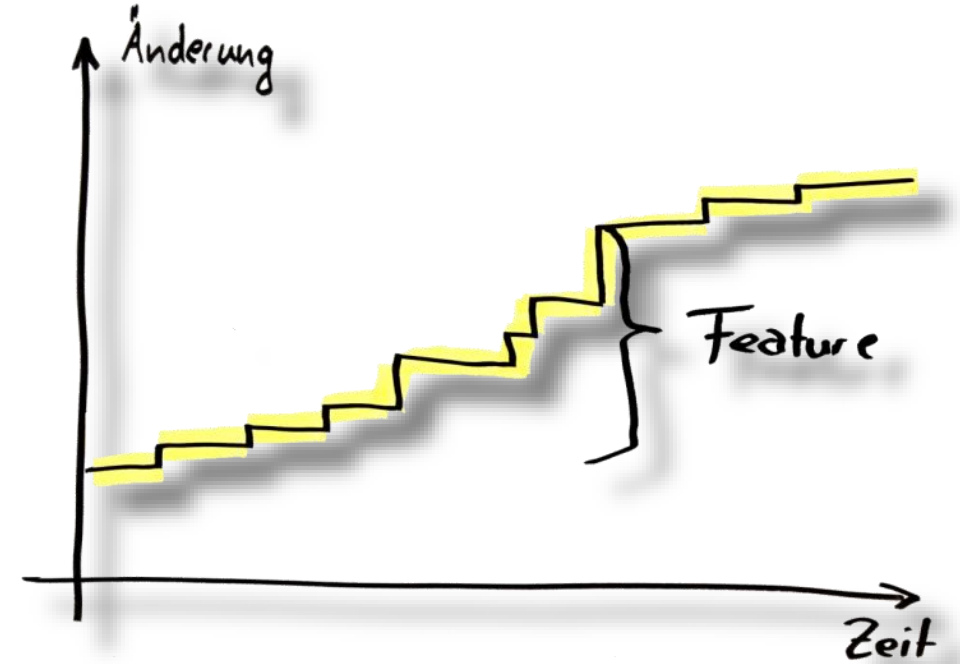
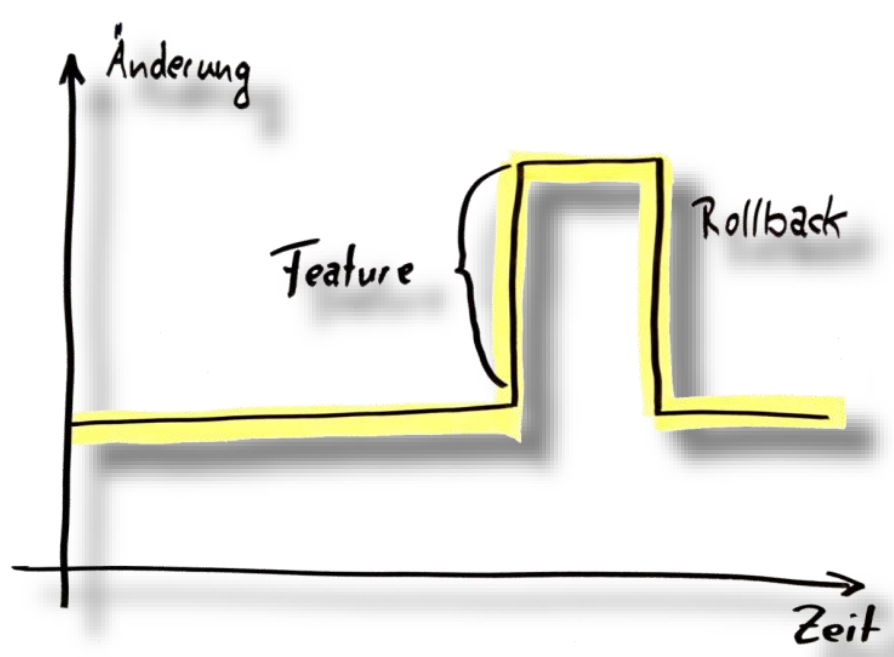


Unabhängige Artefakte



Flexible Skalierung auf 3 Server

Was brauchen wir noch?
Umdenken was ein Feature ist
30 Feature pro Tag?

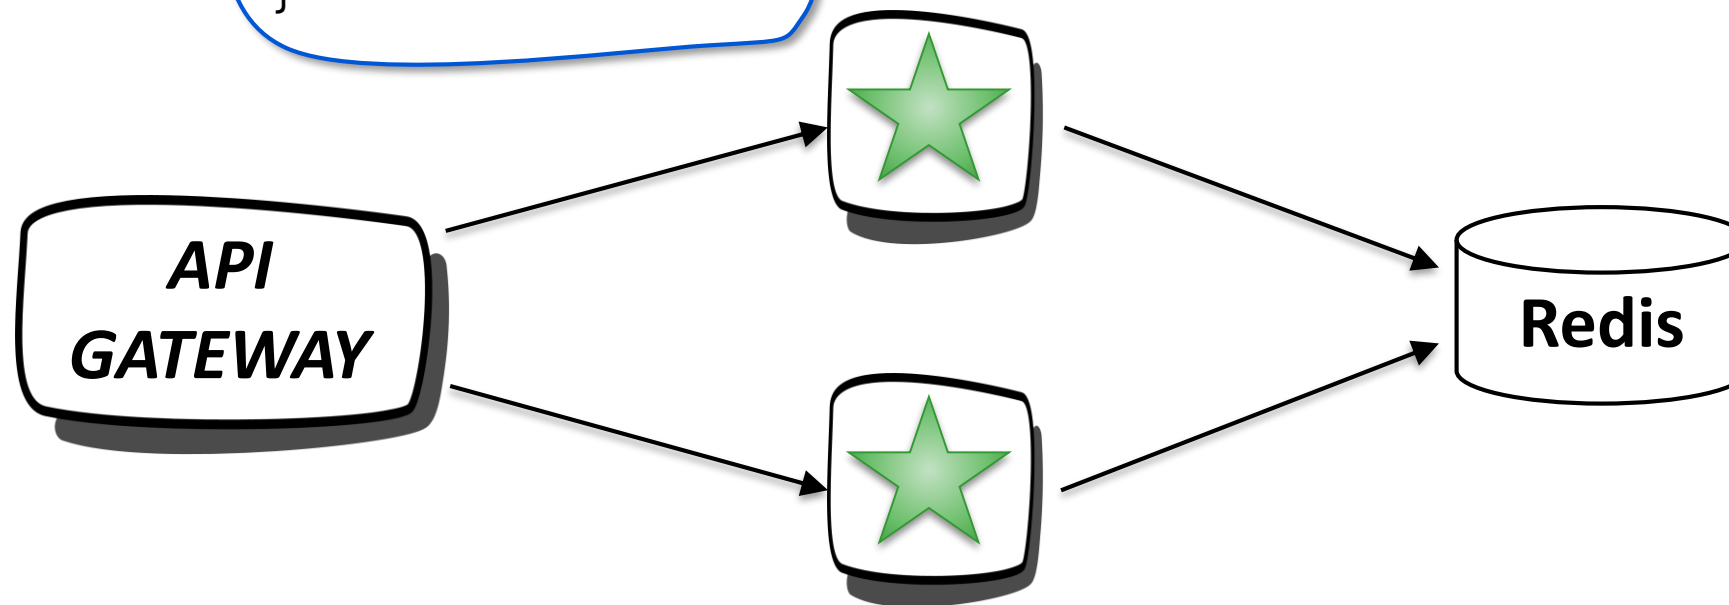


Deploy \neq Feature Launch

Wir brauchen noch mehr

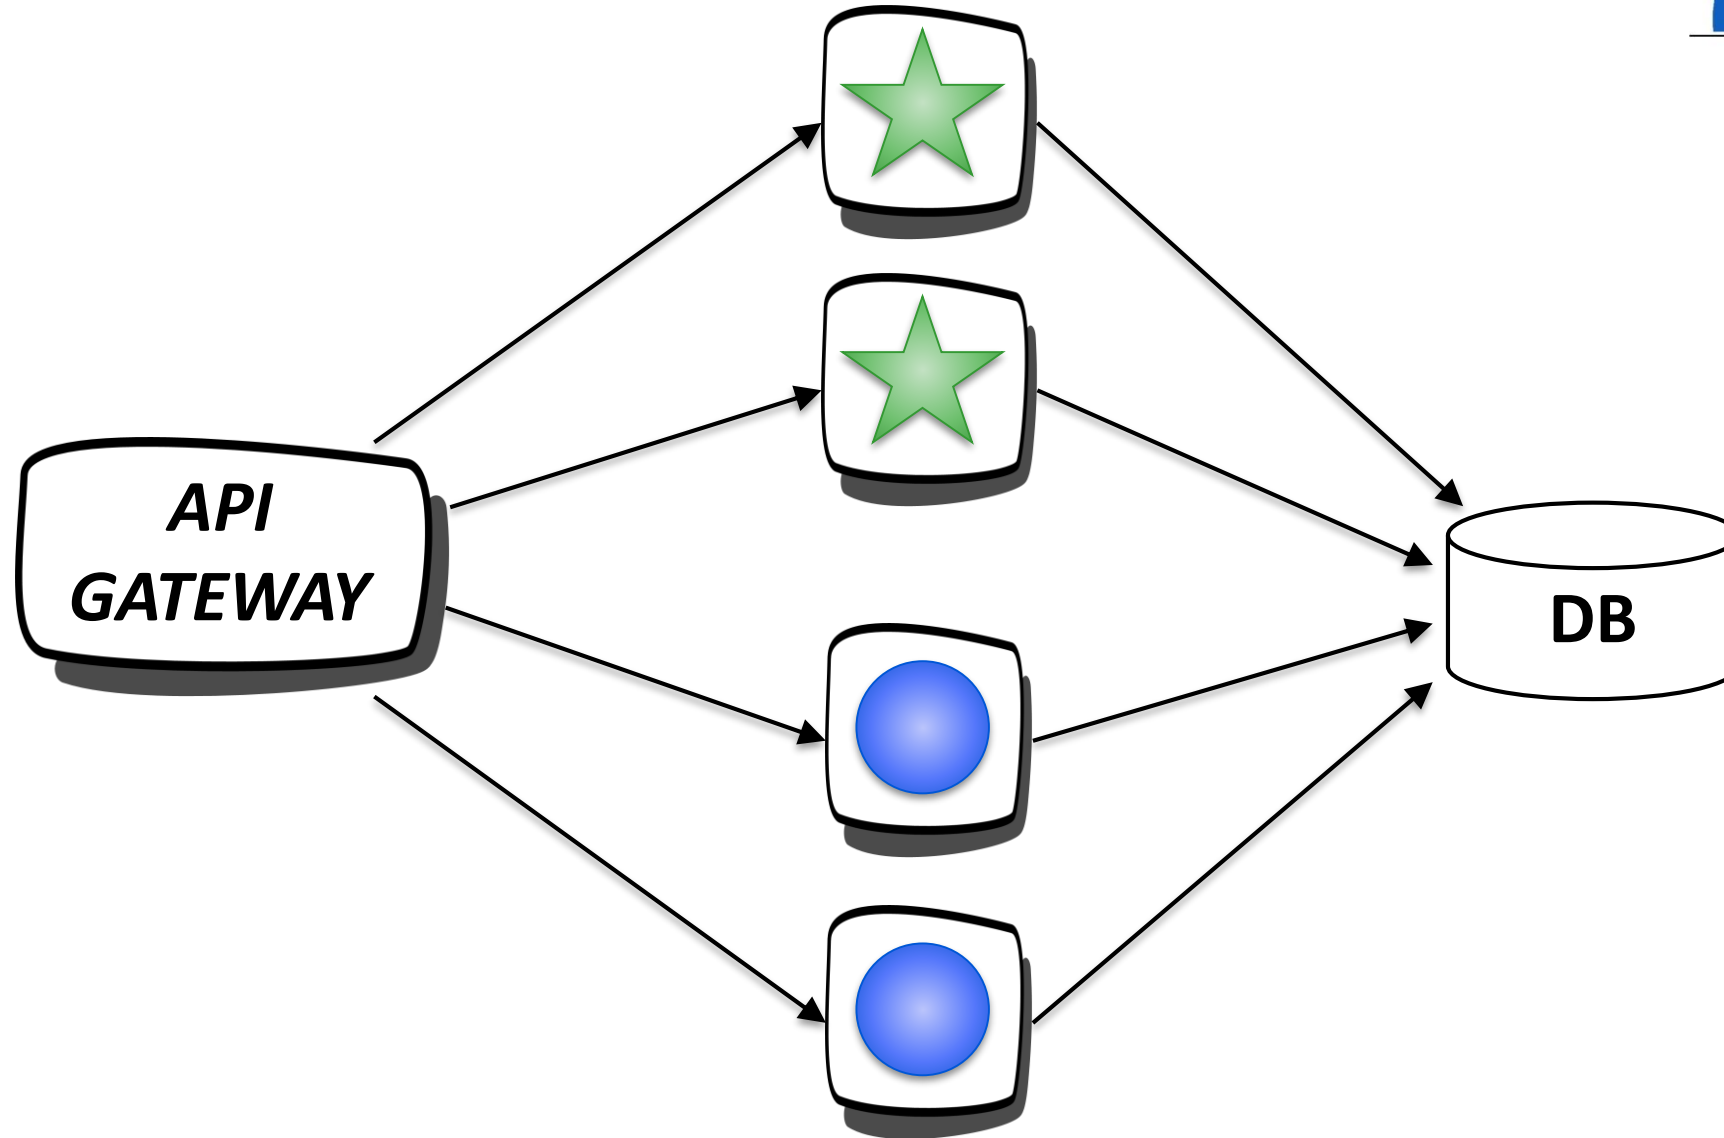
Zustandslose Anwendung

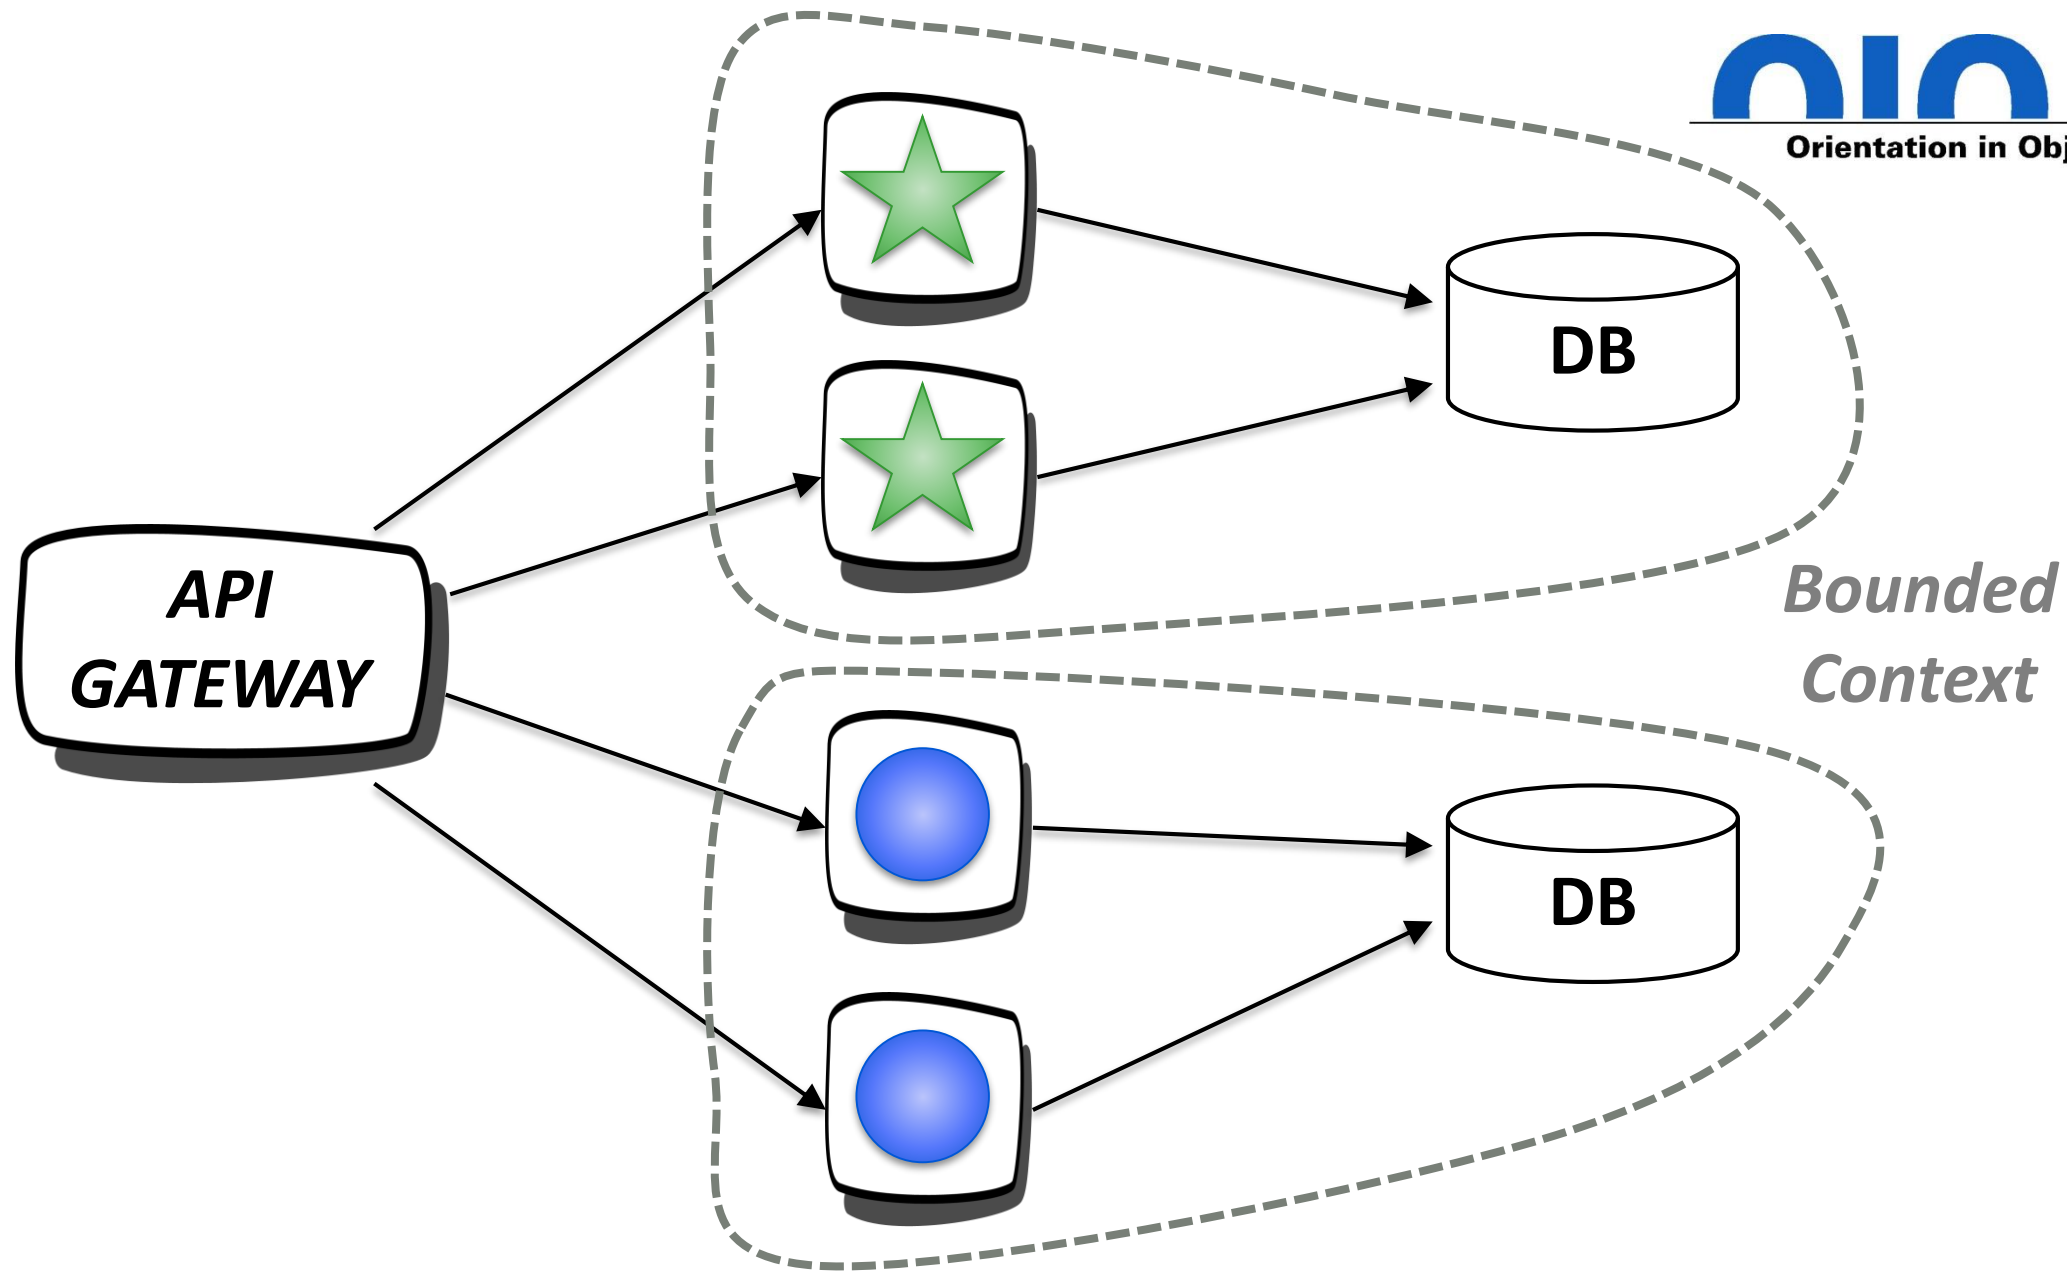
```
@EnableRedisHttpSession  
public class SpringSessionConfig {  
    @Bean  
    public JedisConnectionFactory connectionFactory() {  
        return new JedisConnectionFactory();  
    }  
}
```



Wie viele Datenbanken haben wir eigentlich?

Eine? Mehrere?





Schemaänderungen

Wie ändern wir das Schema?


```
public class Person {
    private String firstname;
}
```



	id [PK] bigint	firstname character vari
1	1	Thorsten
2	2	Max
3	3	John
4	4	Dieter

```
public class Person {
    private String firstname;
    private String lastname;
}
```



	id [PK] bigint	firstname character varying(255)	lastname character varying(255)
1	1	Thorsten	Maier
2	2	Max	Mustermann
3	3	John	Doe
4	4	Dieter	Develop



Hibernate: `hbm2ddl.auto=update` in production?

▲
52
▼

Is it okay to run Hibernate applications configured with `hbm2ddl.auto=update` to update the database schema in a production environment?



Hibernate: hbm2ddl.auto=update in production?

▲
52
▼

Is it okay to run Hibernate applications configured with `hbm2ddl.auto=update` to update the database schema in a production environment?

13 Answers

active

oldest

votes

▲
70
▼
✓

NOOOOOOOOOOOOOOOOOOOOOOO! :)

It's unsafe.

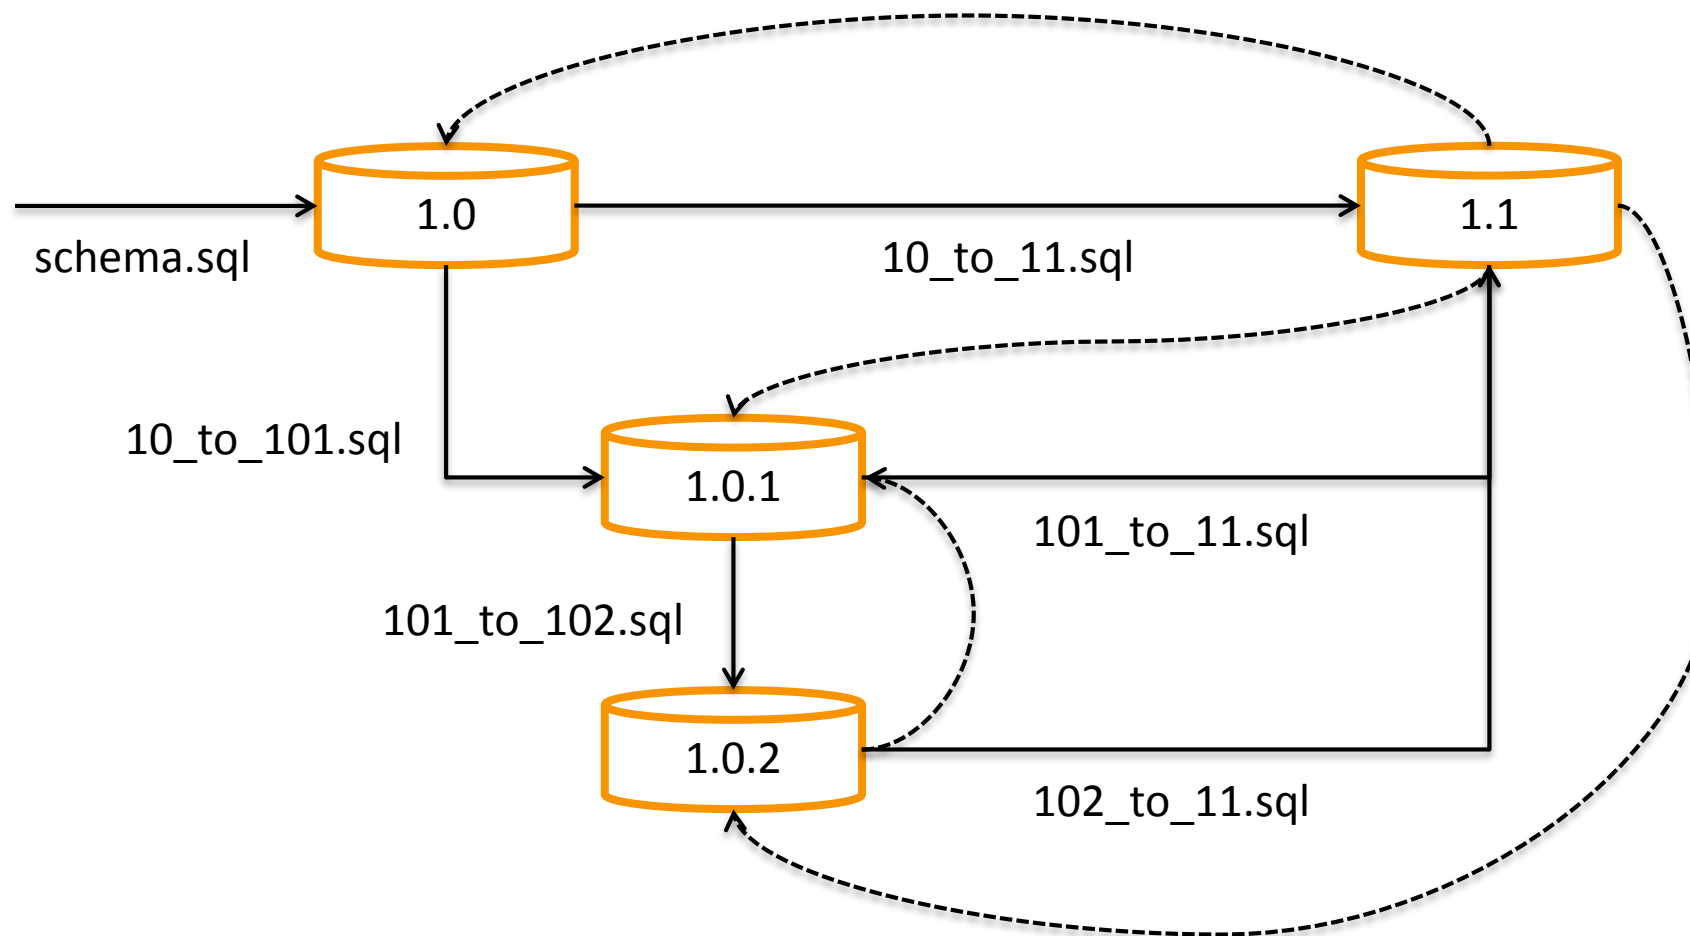
Despite folks in Hib do their best, you simply cannot rely on automatic updates in production. Write your own patches, review them with DBA, test them, then apply them manually.

Theoretically, if hbm2ddl update worked in development, it should work in production too. But in reality, it's not always the case. :-)

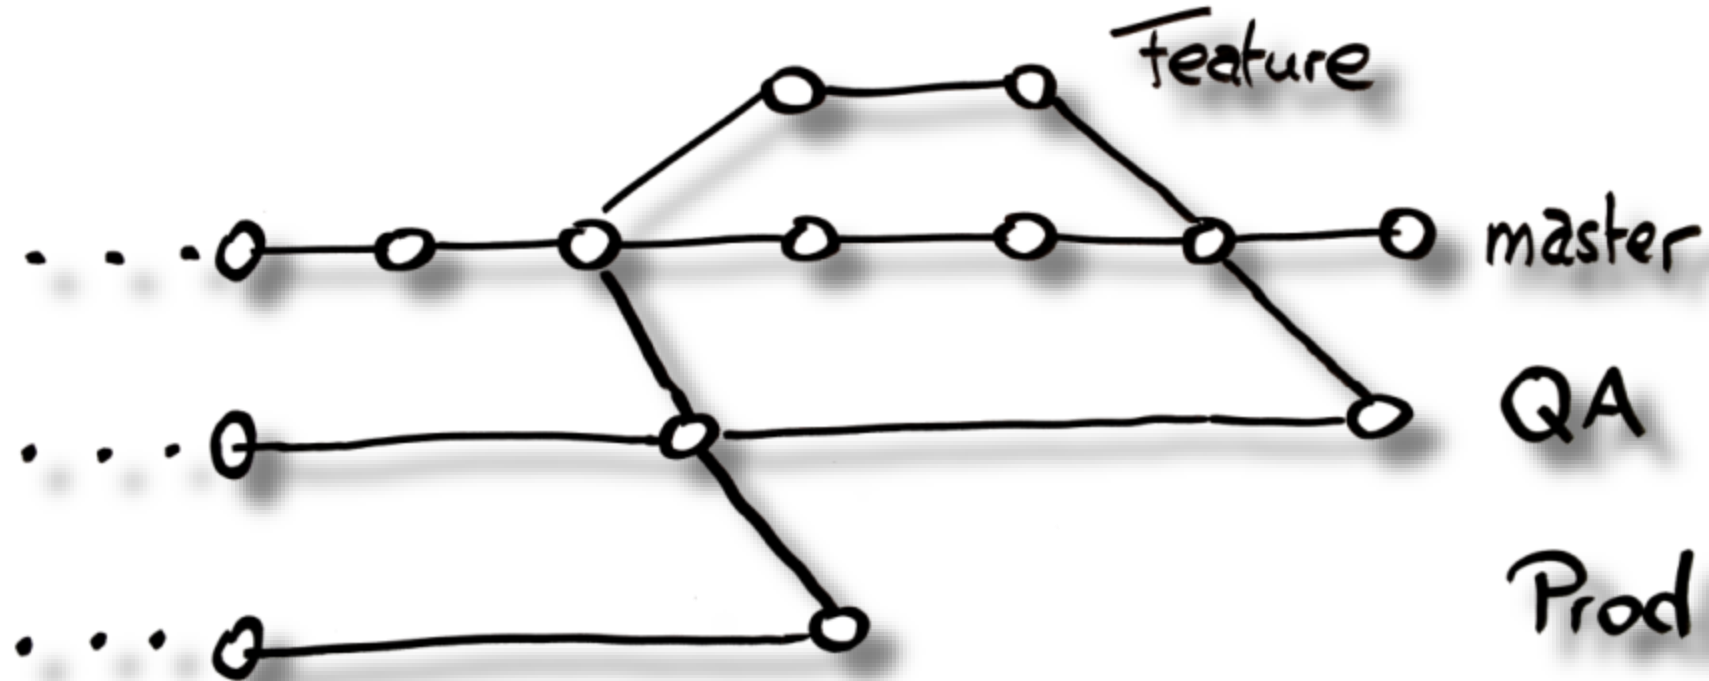
Even if it worked OK, it may be suboptimal. DBAs are paid that much for a reason.

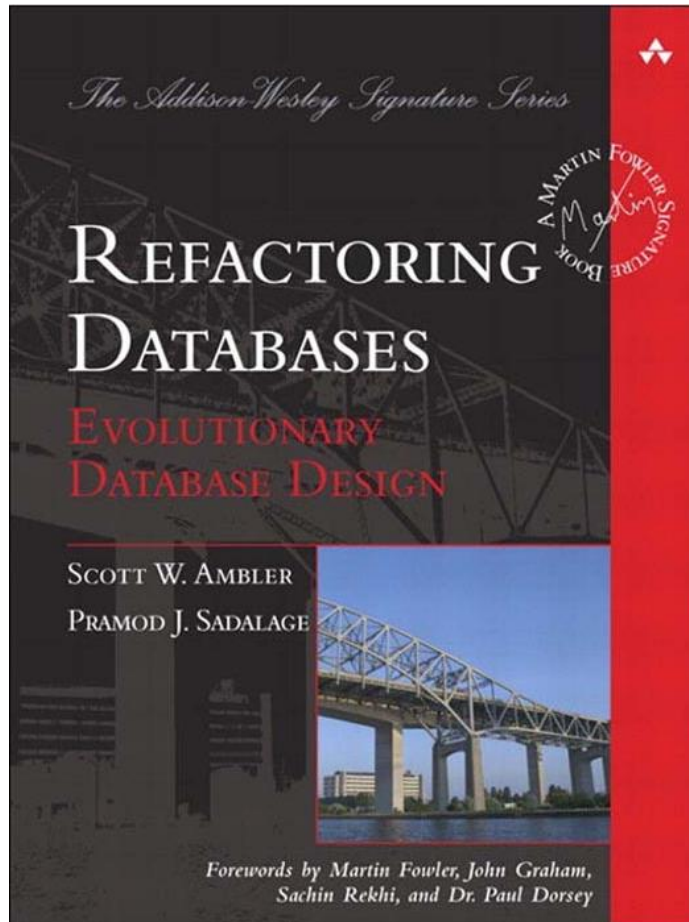
„Vor dem Update von **Version 1.0.2** auf **1.1** musst du erst noch **die neuen SQL-Skript** einspielen.“

- 090123012_HSQL_DB_create_constraints_DDL.sql
- 090123025_HSQL_DB_KARF.ReportingUnit_insert_test_data.sql
- 090123036_HSQL_DB_KARF.Currency_insert_test_data.sql
- 090123037_HSQL_DB_KARF.LocalBusinessUnit_insert_test_data.sql
- 091300022_HSQL_DB_Parts.PartNumberGroup_insert_test_data.sql
- 091300023_HSQL_DB_Parts.GlobalPartNumber_insert_test_data.sql
- 091300027_HSQL_DB_KARF.AssignmentRules_insert_test_data.sql
- 091652021_HSQL_DB_SourceSystem.SourceSystem_insert_test_data.sql
- 100432010_HSQL_DB_drop_schemata_tables_DDL.sql
- 100432024_HSQL_DB_Parts.LocalPartNumber_insert_test_data.sql
- 100432026_HSQL_DB_SourceSystem.SourceSystem_ReportingUnit_insert_test_data.sql
- 100614044_HSQL_DB_KARF.CountryClass_insert_test_data.sql
- 114446035_HSQL_DB_KARF.Country_insert_test_data.sql
- 122135032_HSQL_DB_KARF.BusinessUnit_insert_test_data.sql
- 122135033_HSQL_DB_KARF.ProductGroup_insert_test_data.sql
- 122323034_HSQL_DB_KARF.ActivityType_insert_test_data.sql
- 122323038_HSQL_DB_KARF.LocalProductGroupUnit_insert_test_data.sql
- 135817051_HSQL_DB_MDF.MDFCodeLocation_insert_test_data.sql
- 142308031_HSQL_DB_KARF.Division_insert_test_data.sql
- 154725011_HSQL_DB_create_schemata_tables_DDL.sql



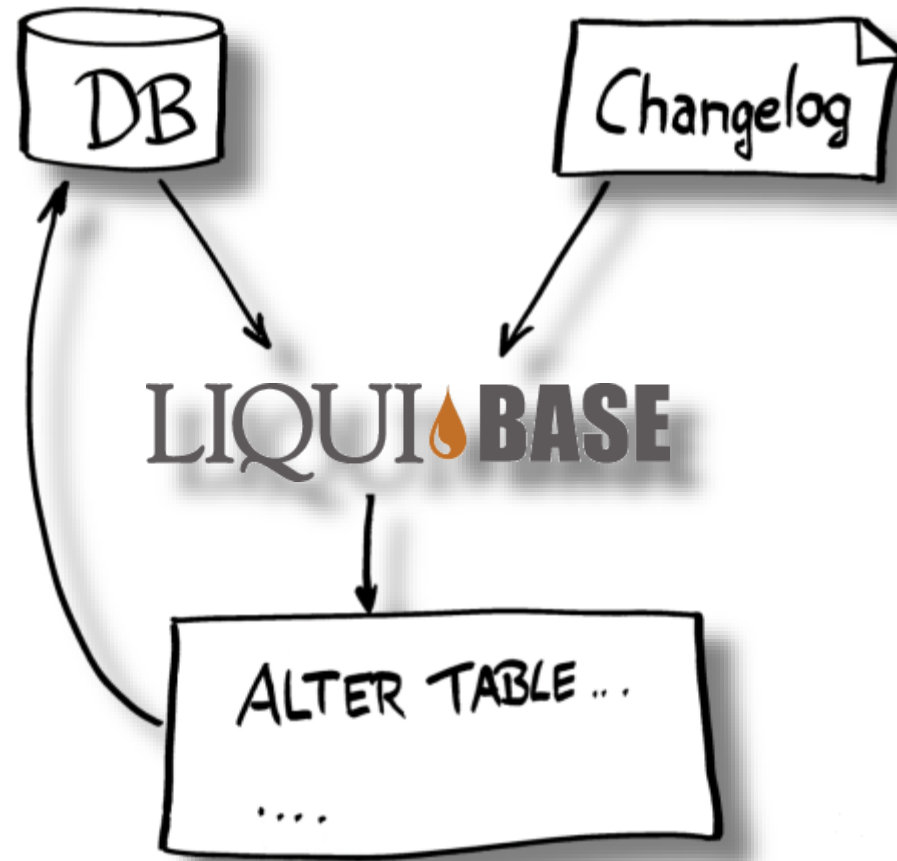
Eigentlich noch komplizierter



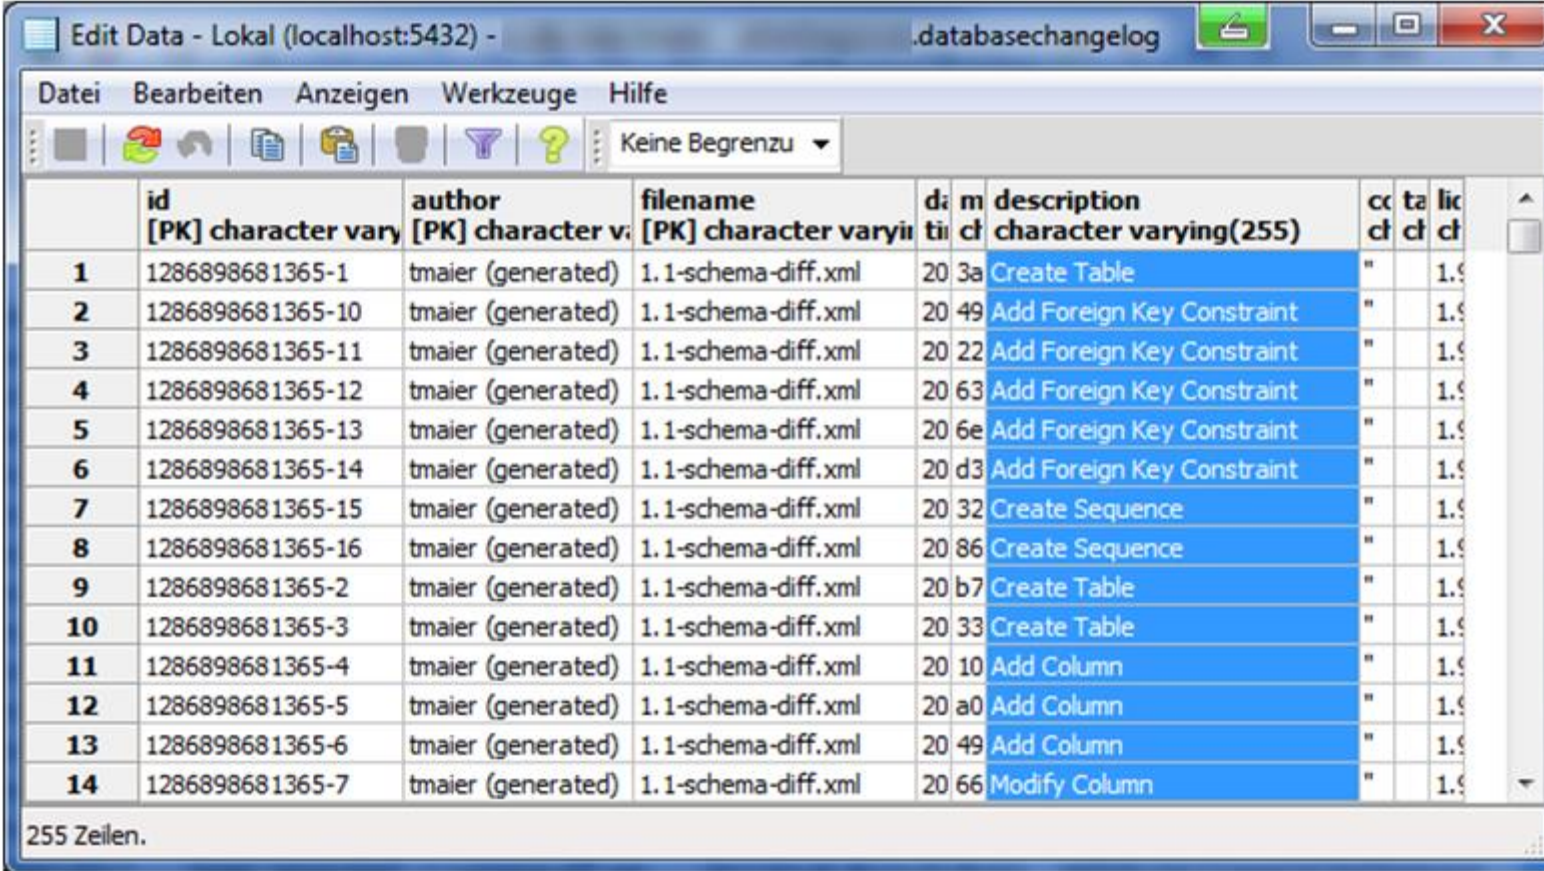


LIQUIBASE






```
<changeSet author="tmaier" id="1495491474">  
  <createTable schemaName="public" tableName="person">  
    <column name="address" type="varchar(255)"/>  
  </createTable>  
</changeSet>
```



Edit Data - Lokal (localhost:5432) - .databasechangelog

Datei Bearbeiten Anzeigen Werkzeuge Hilfe

Keine Begrenzu

	id [PK] character vary	author [PK] character v	filename [PK] character varyin	d: m	description character varying(255)	cc	ta	lic
1	1286898681365-1	tmaier (generated)	1.1-schema-diff.xml	20 3a	Create Table	"		1.9
2	1286898681365-10	tmaier (generated)	1.1-schema-diff.xml	20 49	Add Foreign Key Constraint	"		1.9
3	1286898681365-11	tmaier (generated)	1.1-schema-diff.xml	20 22	Add Foreign Key Constraint	"		1.9
4	1286898681365-12	tmaier (generated)	1.1-schema-diff.xml	20 63	Add Foreign Key Constraint	"		1.9
5	1286898681365-13	tmaier (generated)	1.1-schema-diff.xml	20 6e	Add Foreign Key Constraint	"		1.9
6	1286898681365-14	tmaier (generated)	1.1-schema-diff.xml	20 d3	Add Foreign Key Constraint	"		1.9
7	1286898681365-15	tmaier (generated)	1.1-schema-diff.xml	20 32	Create Sequence	"		1.9
8	1286898681365-16	tmaier (generated)	1.1-schema-diff.xml	20 86	Create Sequence	"		1.9
9	1286898681365-2	tmaier (generated)	1.1-schema-diff.xml	20 b7	Create Table	"		1.9
10	1286898681365-3	tmaier (generated)	1.1-schema-diff.xml	20 33	Create Table	"		1.9
11	1286898681365-4	tmaier (generated)	1.1-schema-diff.xml	20 10	Add Column	"		1.9
12	1286898681365-5	tmaier (generated)	1.1-schema-diff.xml	20 a0	Add Column	"		1.9
13	1286898681365-6	tmaier (generated)	1.1-schema-diff.xml	20 49	Add Column	"		1.9
14	1286898681365-7	tmaier (generated)	1.1-schema-diff.xml	20 66	Modify Column	"		1.9

255 Zeilen.

***Liquibase aktualisiert
beim Starten der Anwendung
automatisch die Datenbank.***

BEISPIEL

John Doe

Ersetze die Spalte **Name**
durch die Spalte
Vorname

John

ID	name	firstname
1	John Doe	
2	Dieter Mai	



ID	name	firstname
...
3	Mia Boe	Mia

ID	name	firstname
1	John Doe	John
2	Dieter Mai	Dieter
3	Mia Boe	Mia



ID	name	firstname
...

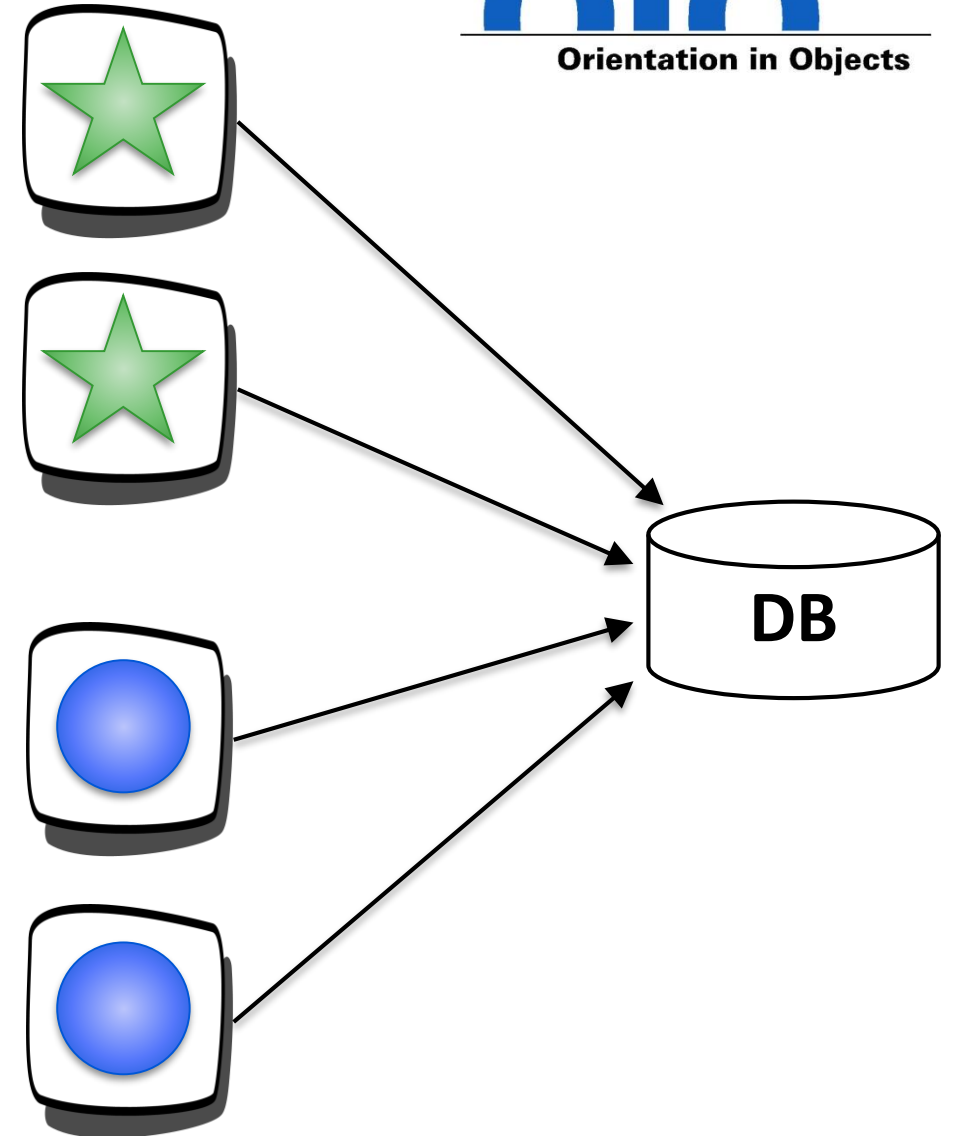
read

ID	name	firstname
...



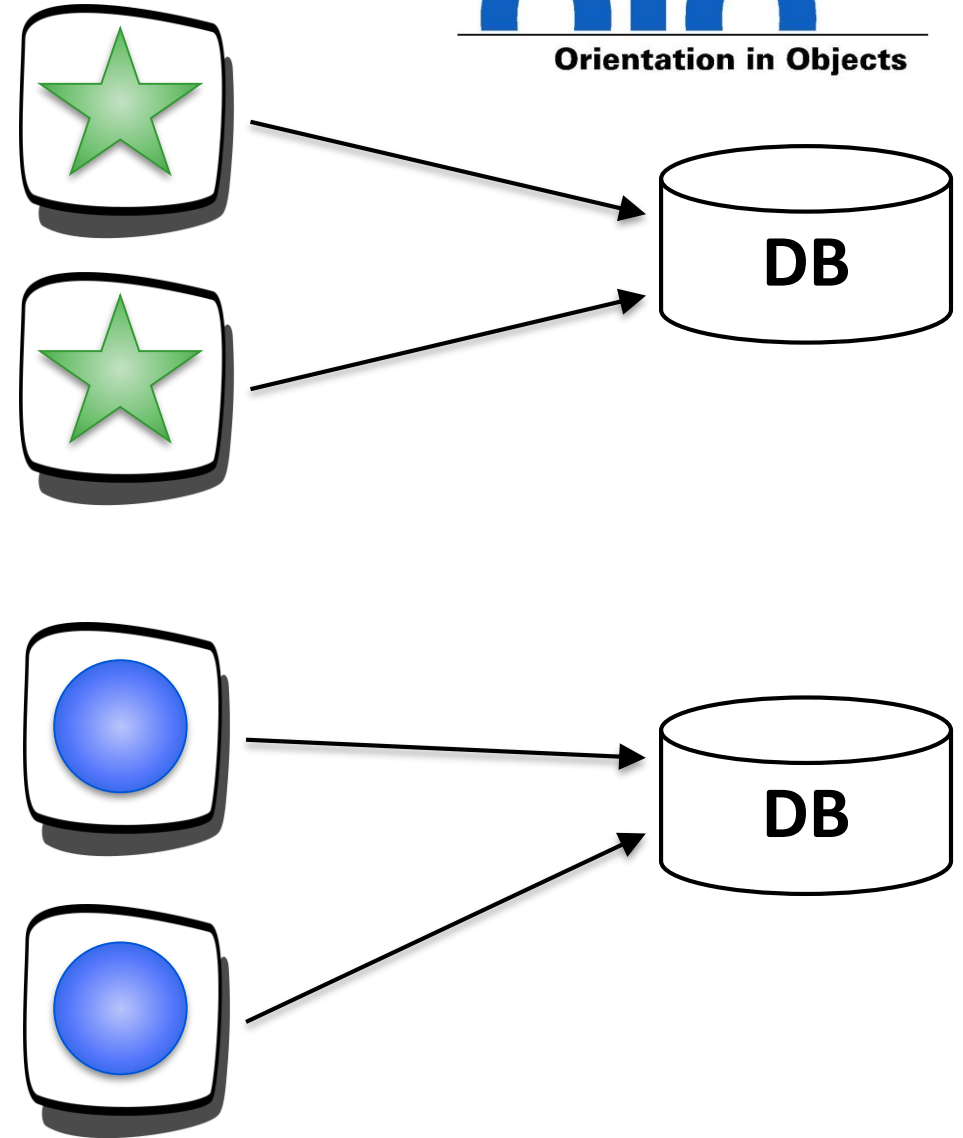
Wer ändert wann die DB?

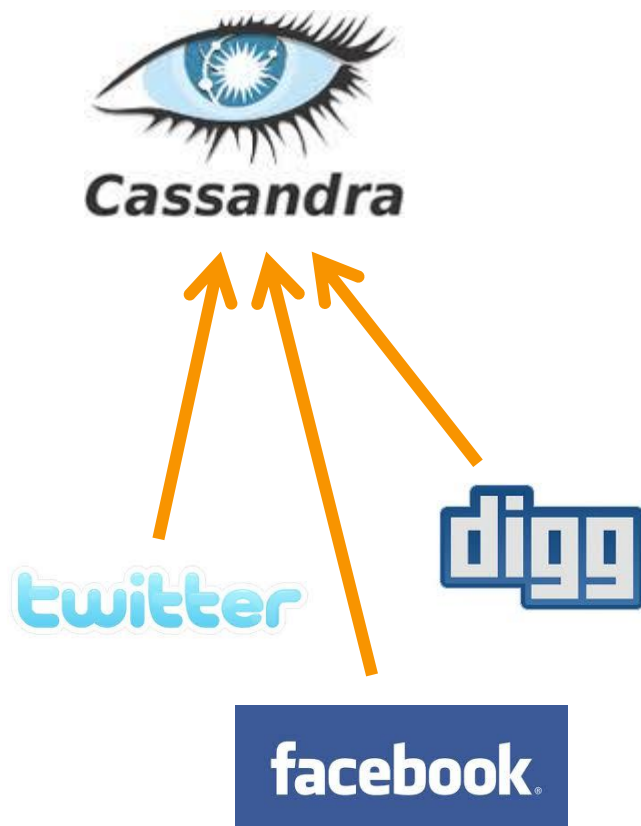
It depends ;-)



Zu viele Köche

***Ab einer gewissen
Größe nur noch so!***





*Warum nehmen wir eigentlich keine
schemalose NoSQL DB?*

Relationale Datenbanken skalieren nicht

Warum?

Normalisierung => Joins

Forderung nach Konsistenz => Transaktionen

NoSQL

=

Relationale Datenbank

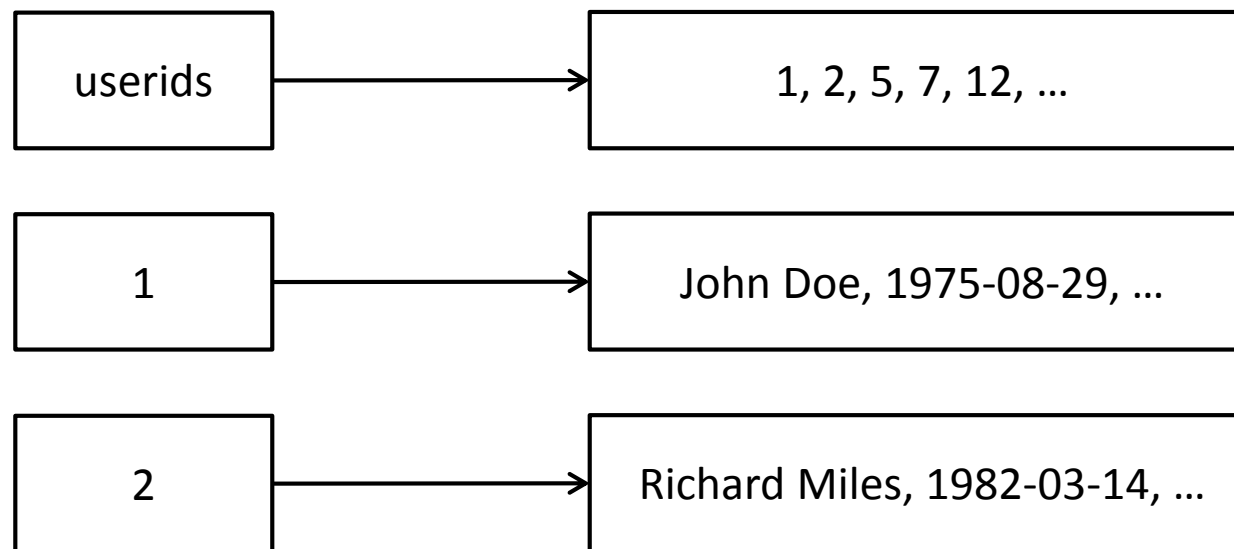
*- Transaktionen - Normalisierung - Joins - Konsistenz - **hartes Schema***

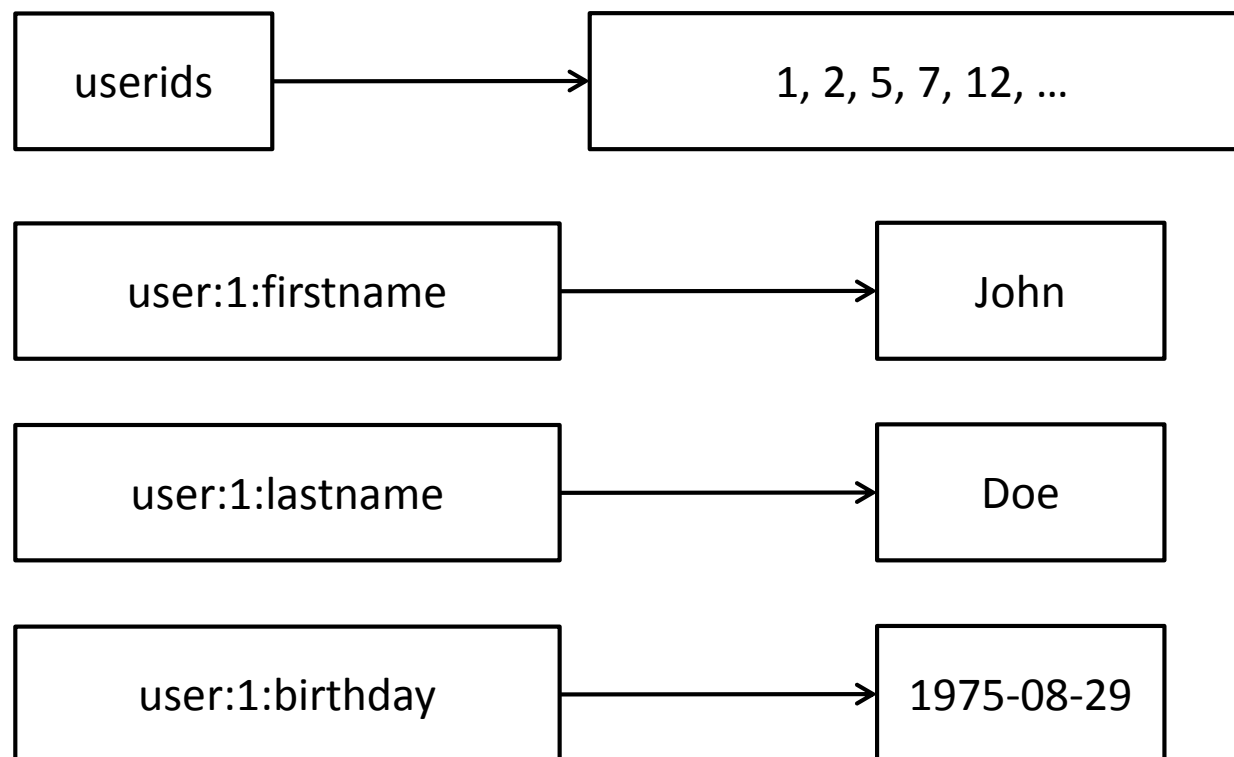
(+) Skalierung

(+) Performance

Datenmodellierung NoSQL

***Konkret für den Anwendungsfall
Denormalisierung und Duplizierung sind normal***

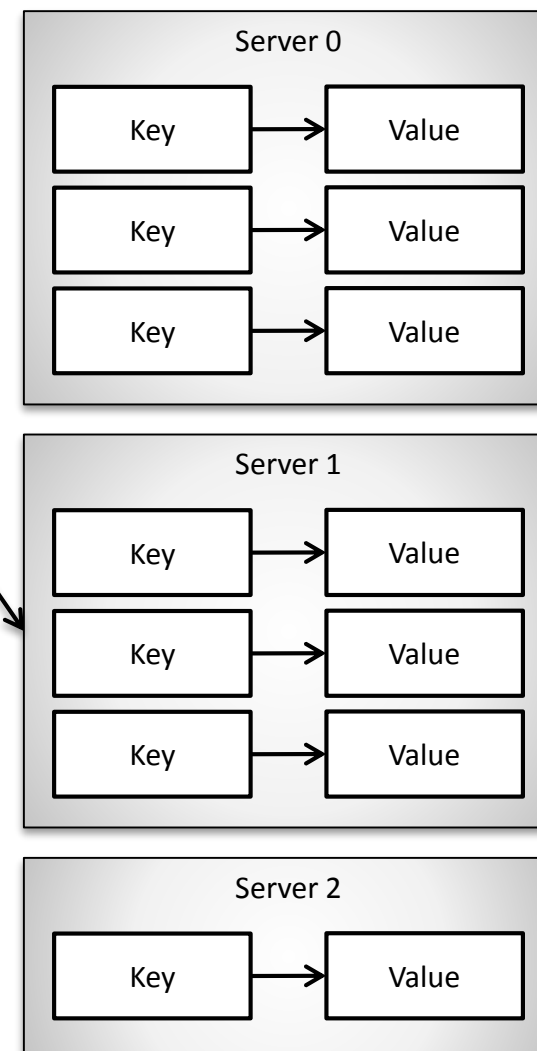


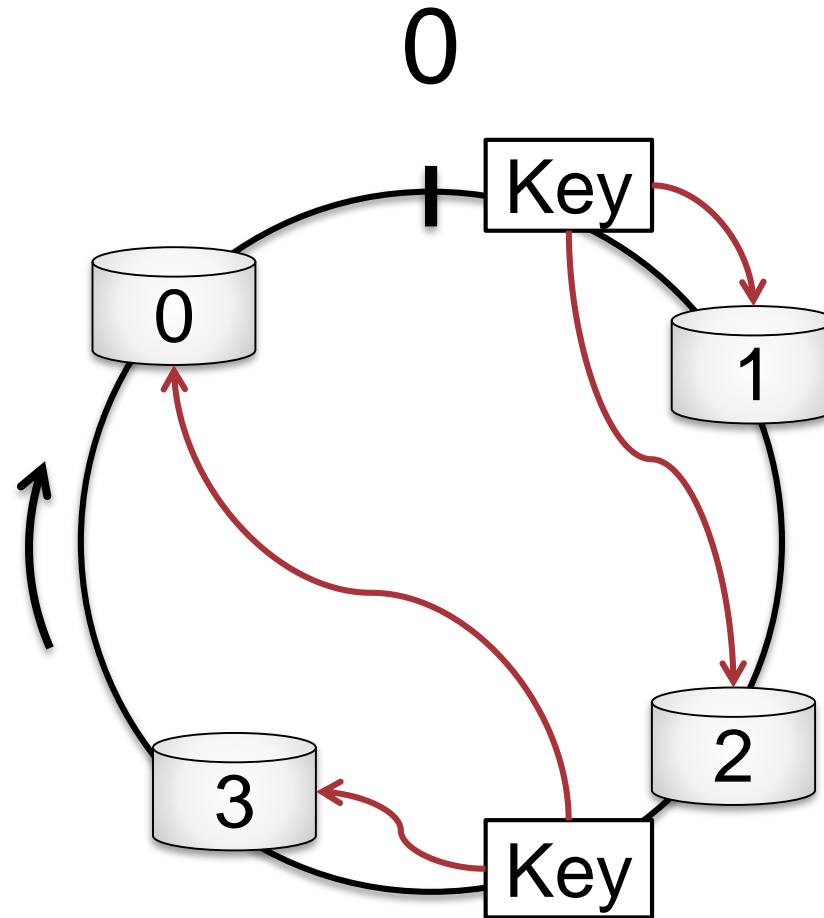


user:tmaier:name

Key-Hash: 39E5DC60AAD55349
Modulo Anzahl der Server (3): 1

→ Key auf Server 1 speichern

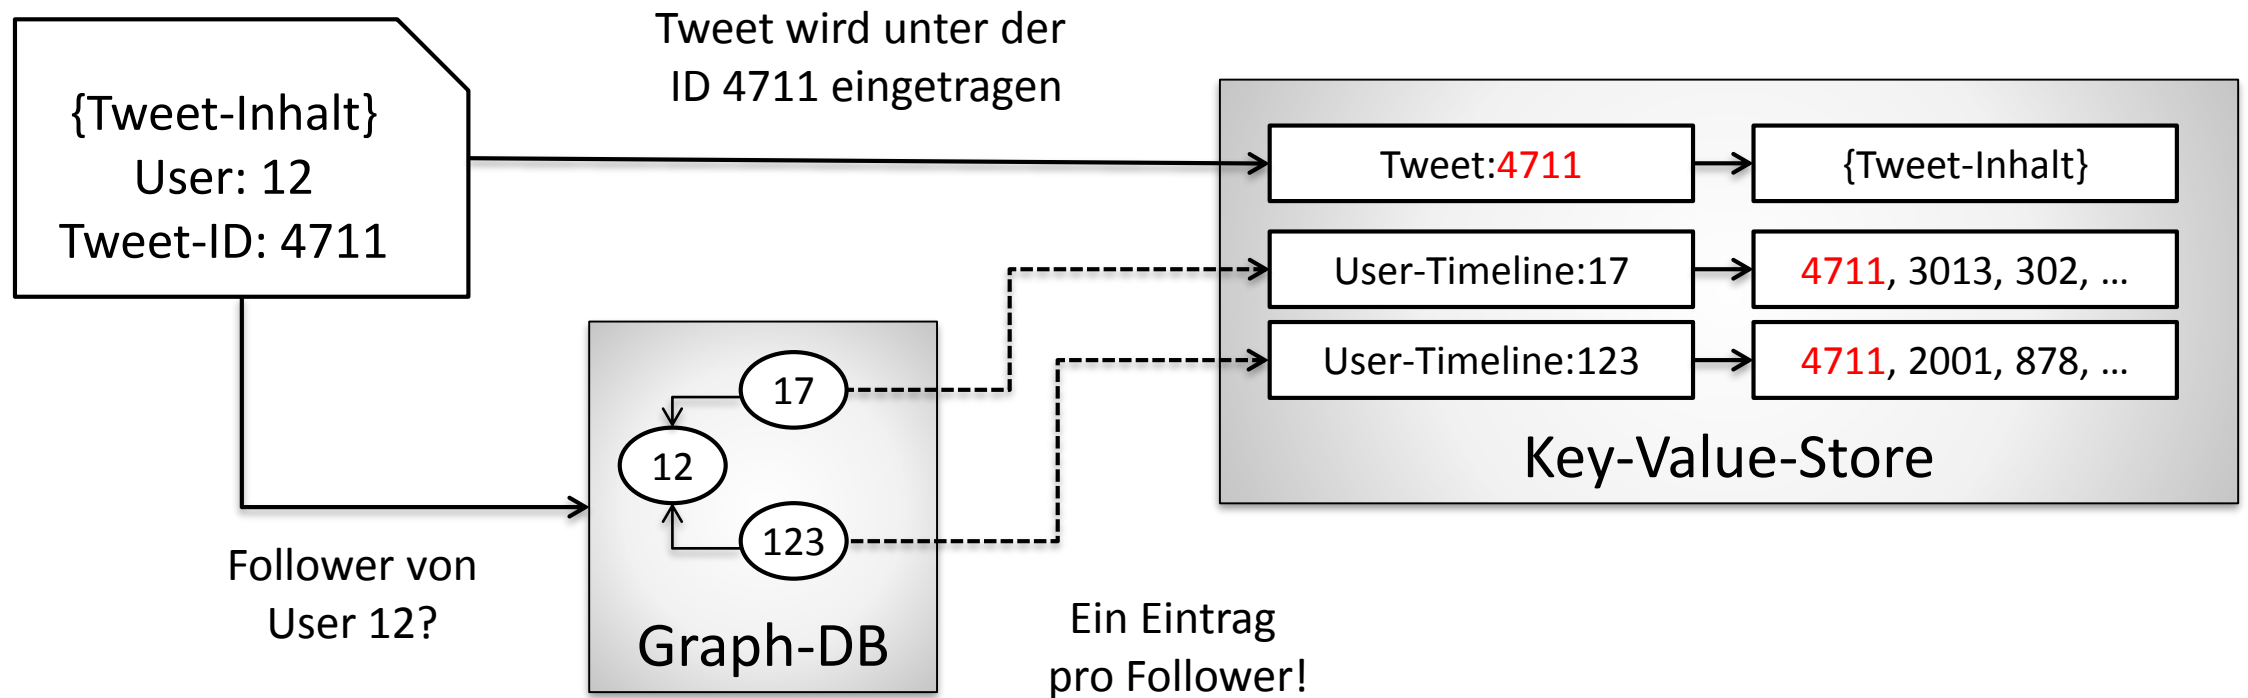




Twitter

*Viel Aktion beim Schreiben,
wenig beim Lesen*

300.000 reads/sec
6000 writes/sec



NoSQL?

Falls die **Anforderungen** passen!

Anwendung kümmert sich um das „**Schema**“



Fragen ?

Orientation in Objects GmbH

Weinheimer Str. 68
68309 Mannheim

www.oio.de
info@oio.de



Vielen Dank für ihre Aufmerksamkeit !

Orientation in Objects GmbH

Weinheimer Str. 68
68309 Mannheim

www.oio.de
info@oio.de