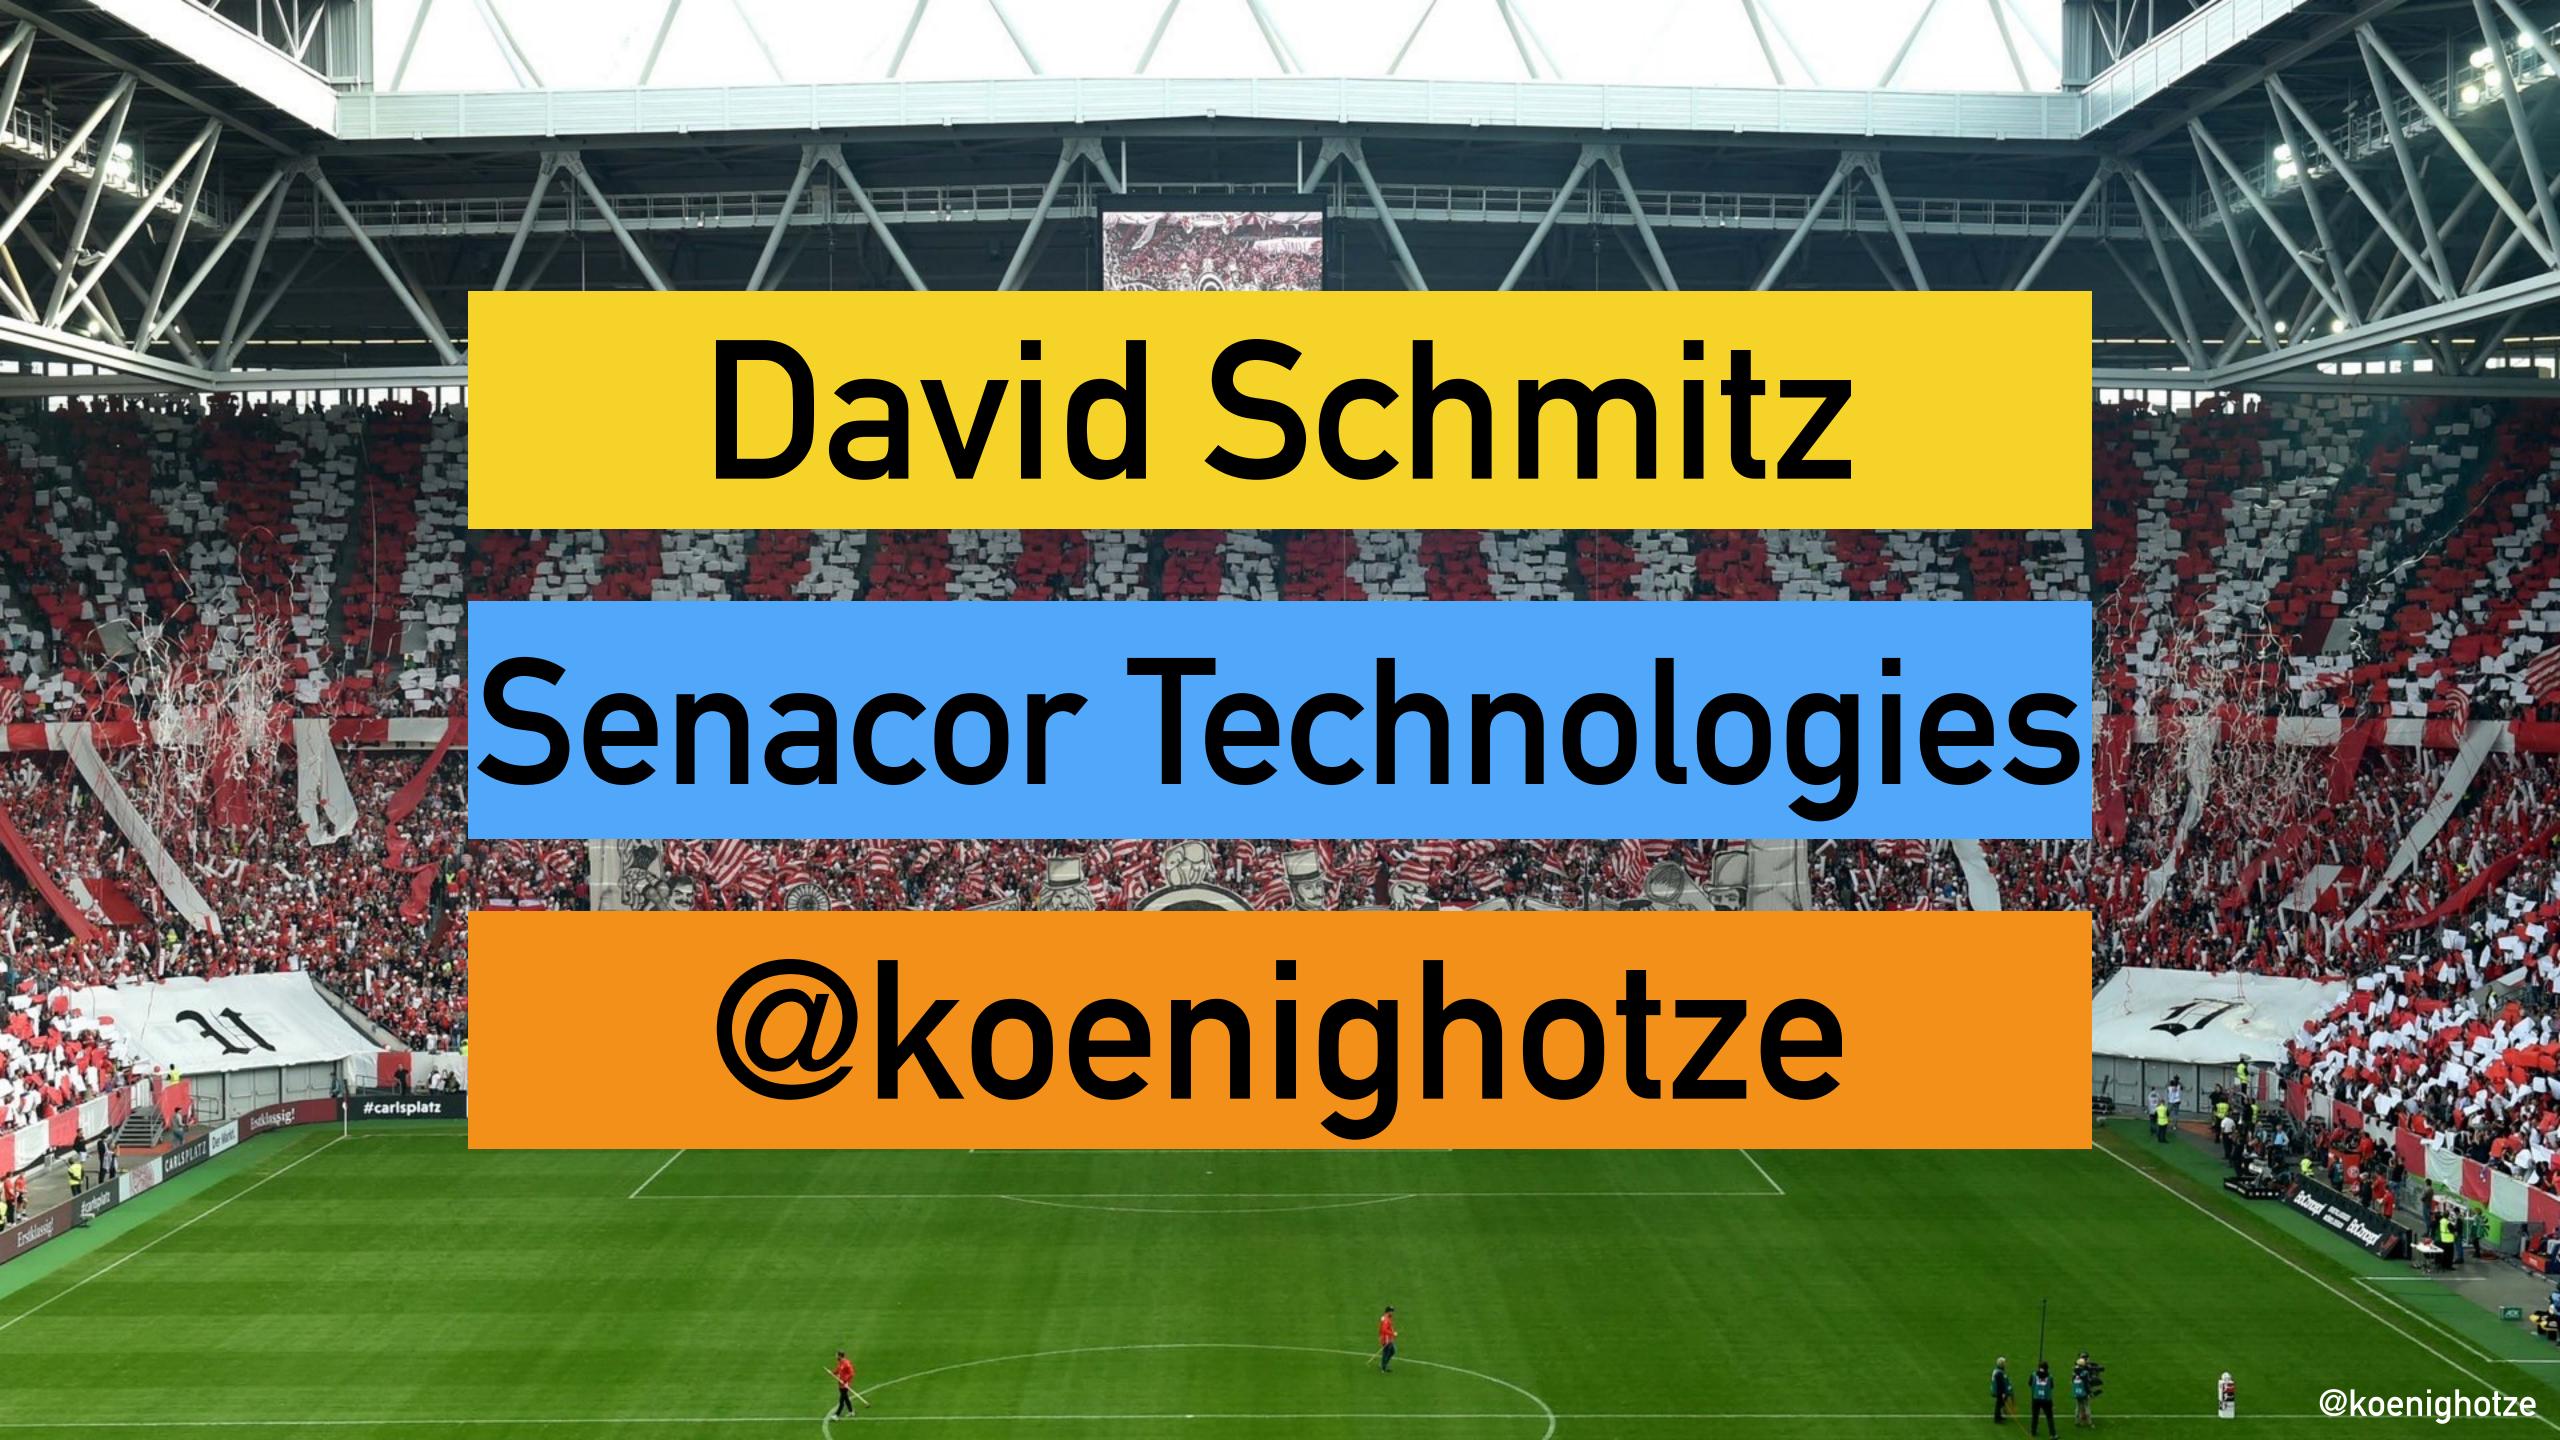




Event sourcing - You are maybe doing parts of it wrong because we made some mistakes along the way and so will you, I guess. This is difficult because there are no easy right/wrong answers; only trade-offs.





Are You

...building microservices? ...doing Domain Driven Design? ...applying event sourcing? ...using Kafka as an event store?

Typical misconceptions Patterns "we" found useful Traps to avoid Not a Kafka-rant No silver bullets





Event driven business

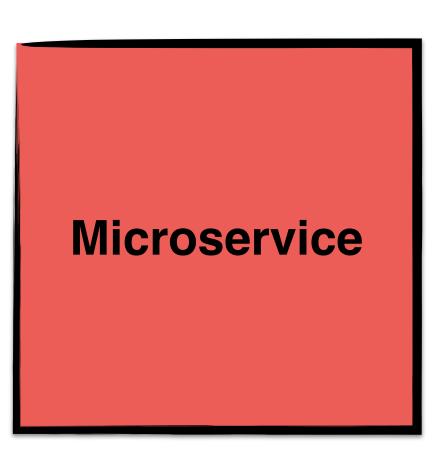






Behaviour of a domain made explicit as a first class concept

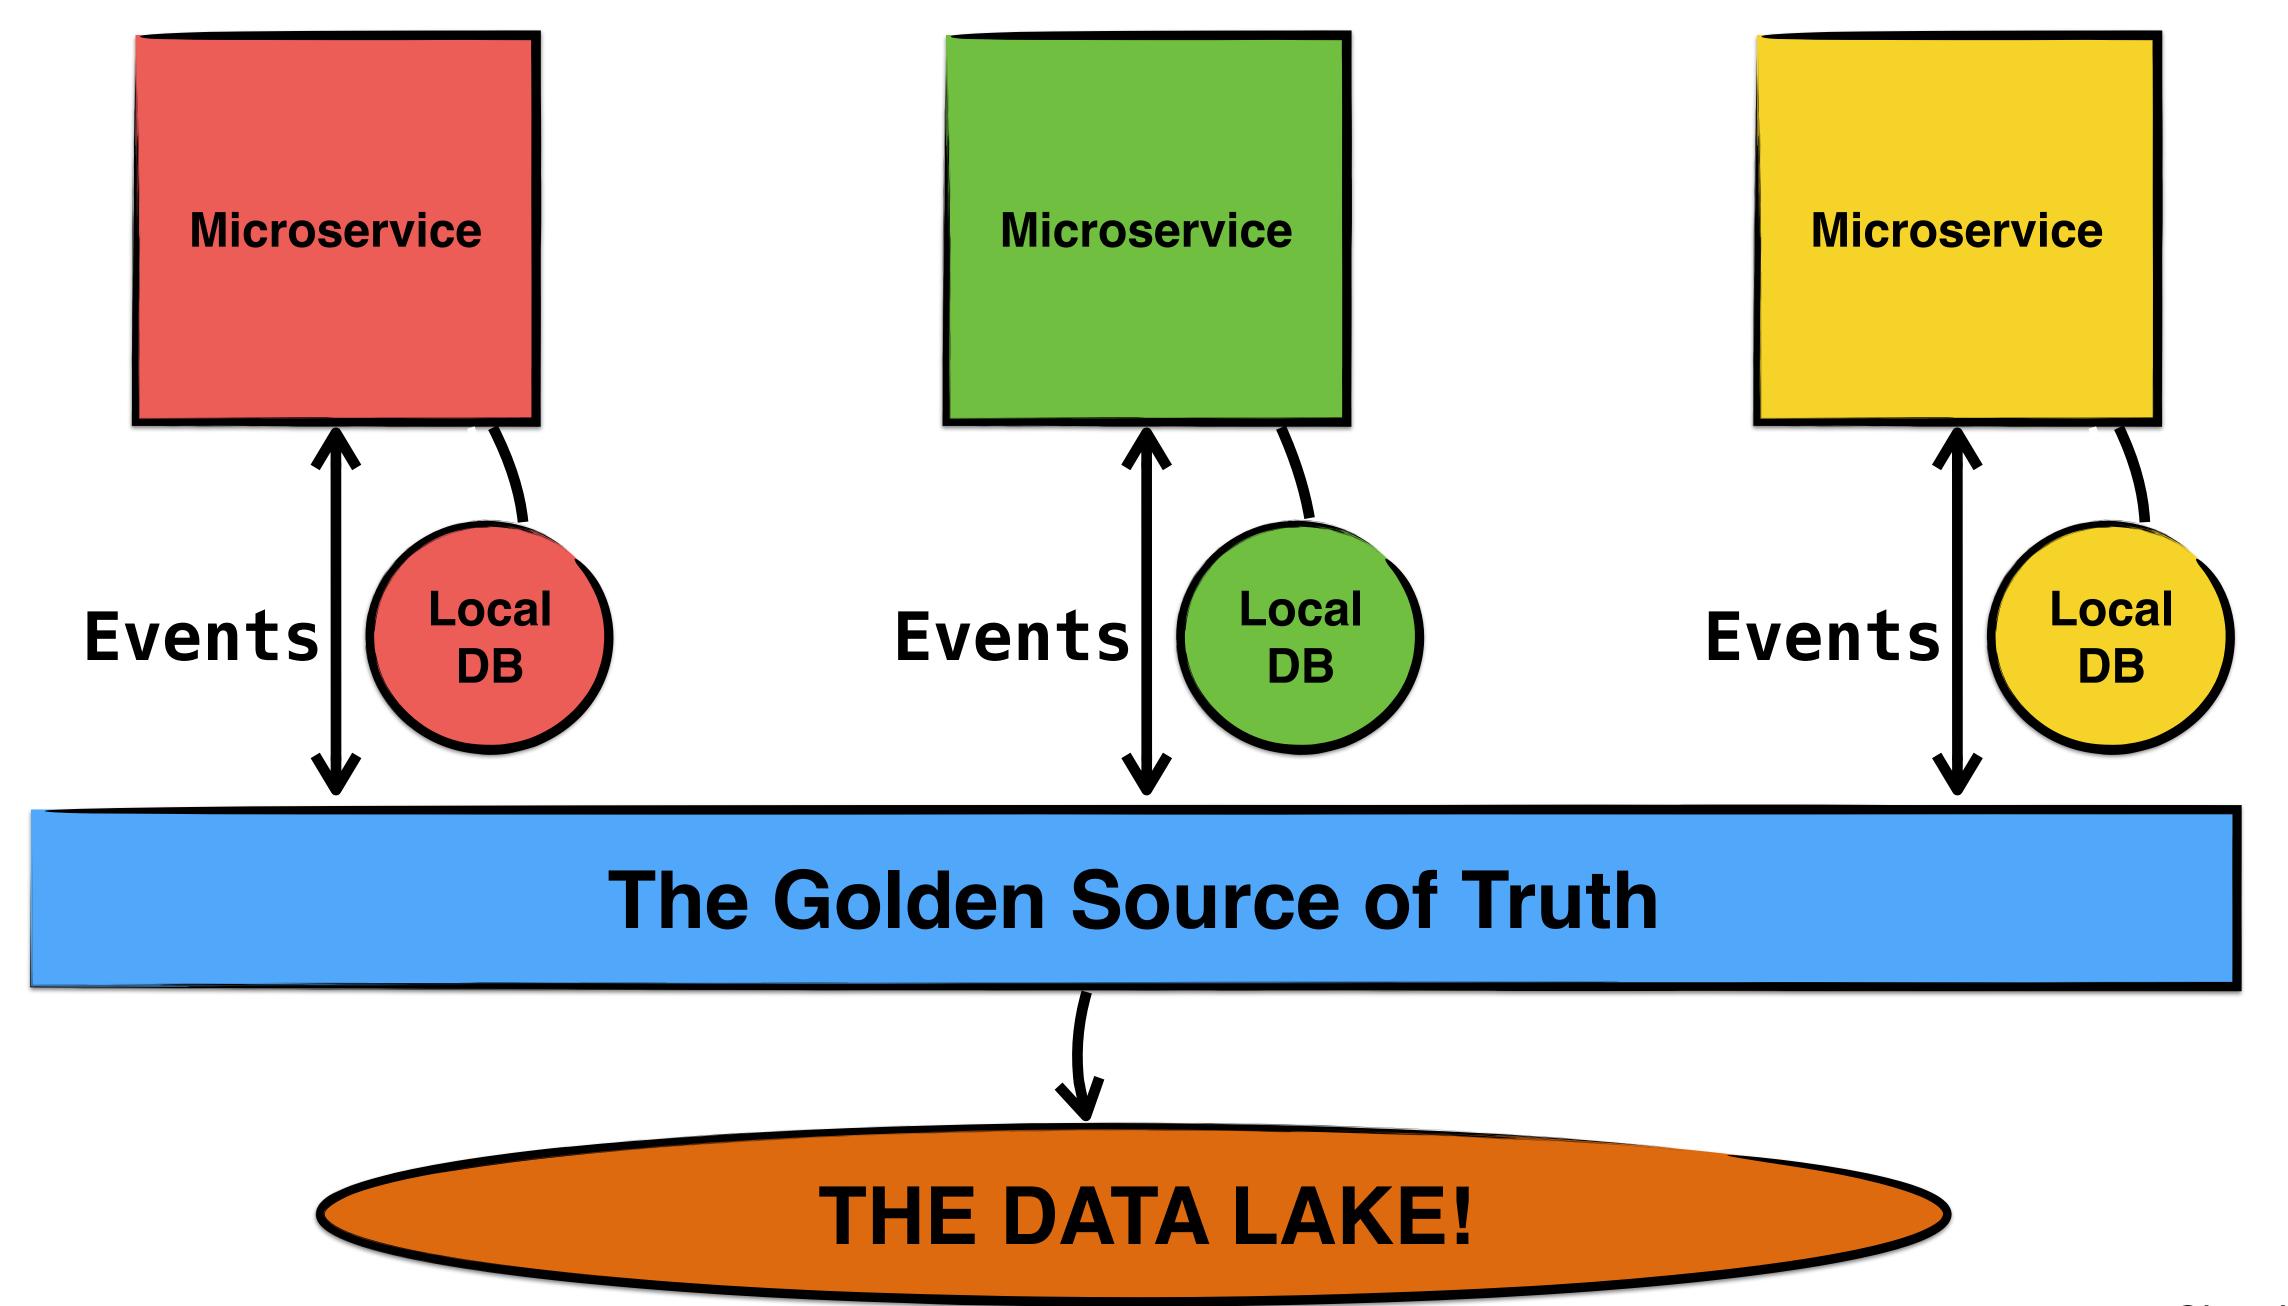
And on a technical level?



Microservice

Microservice

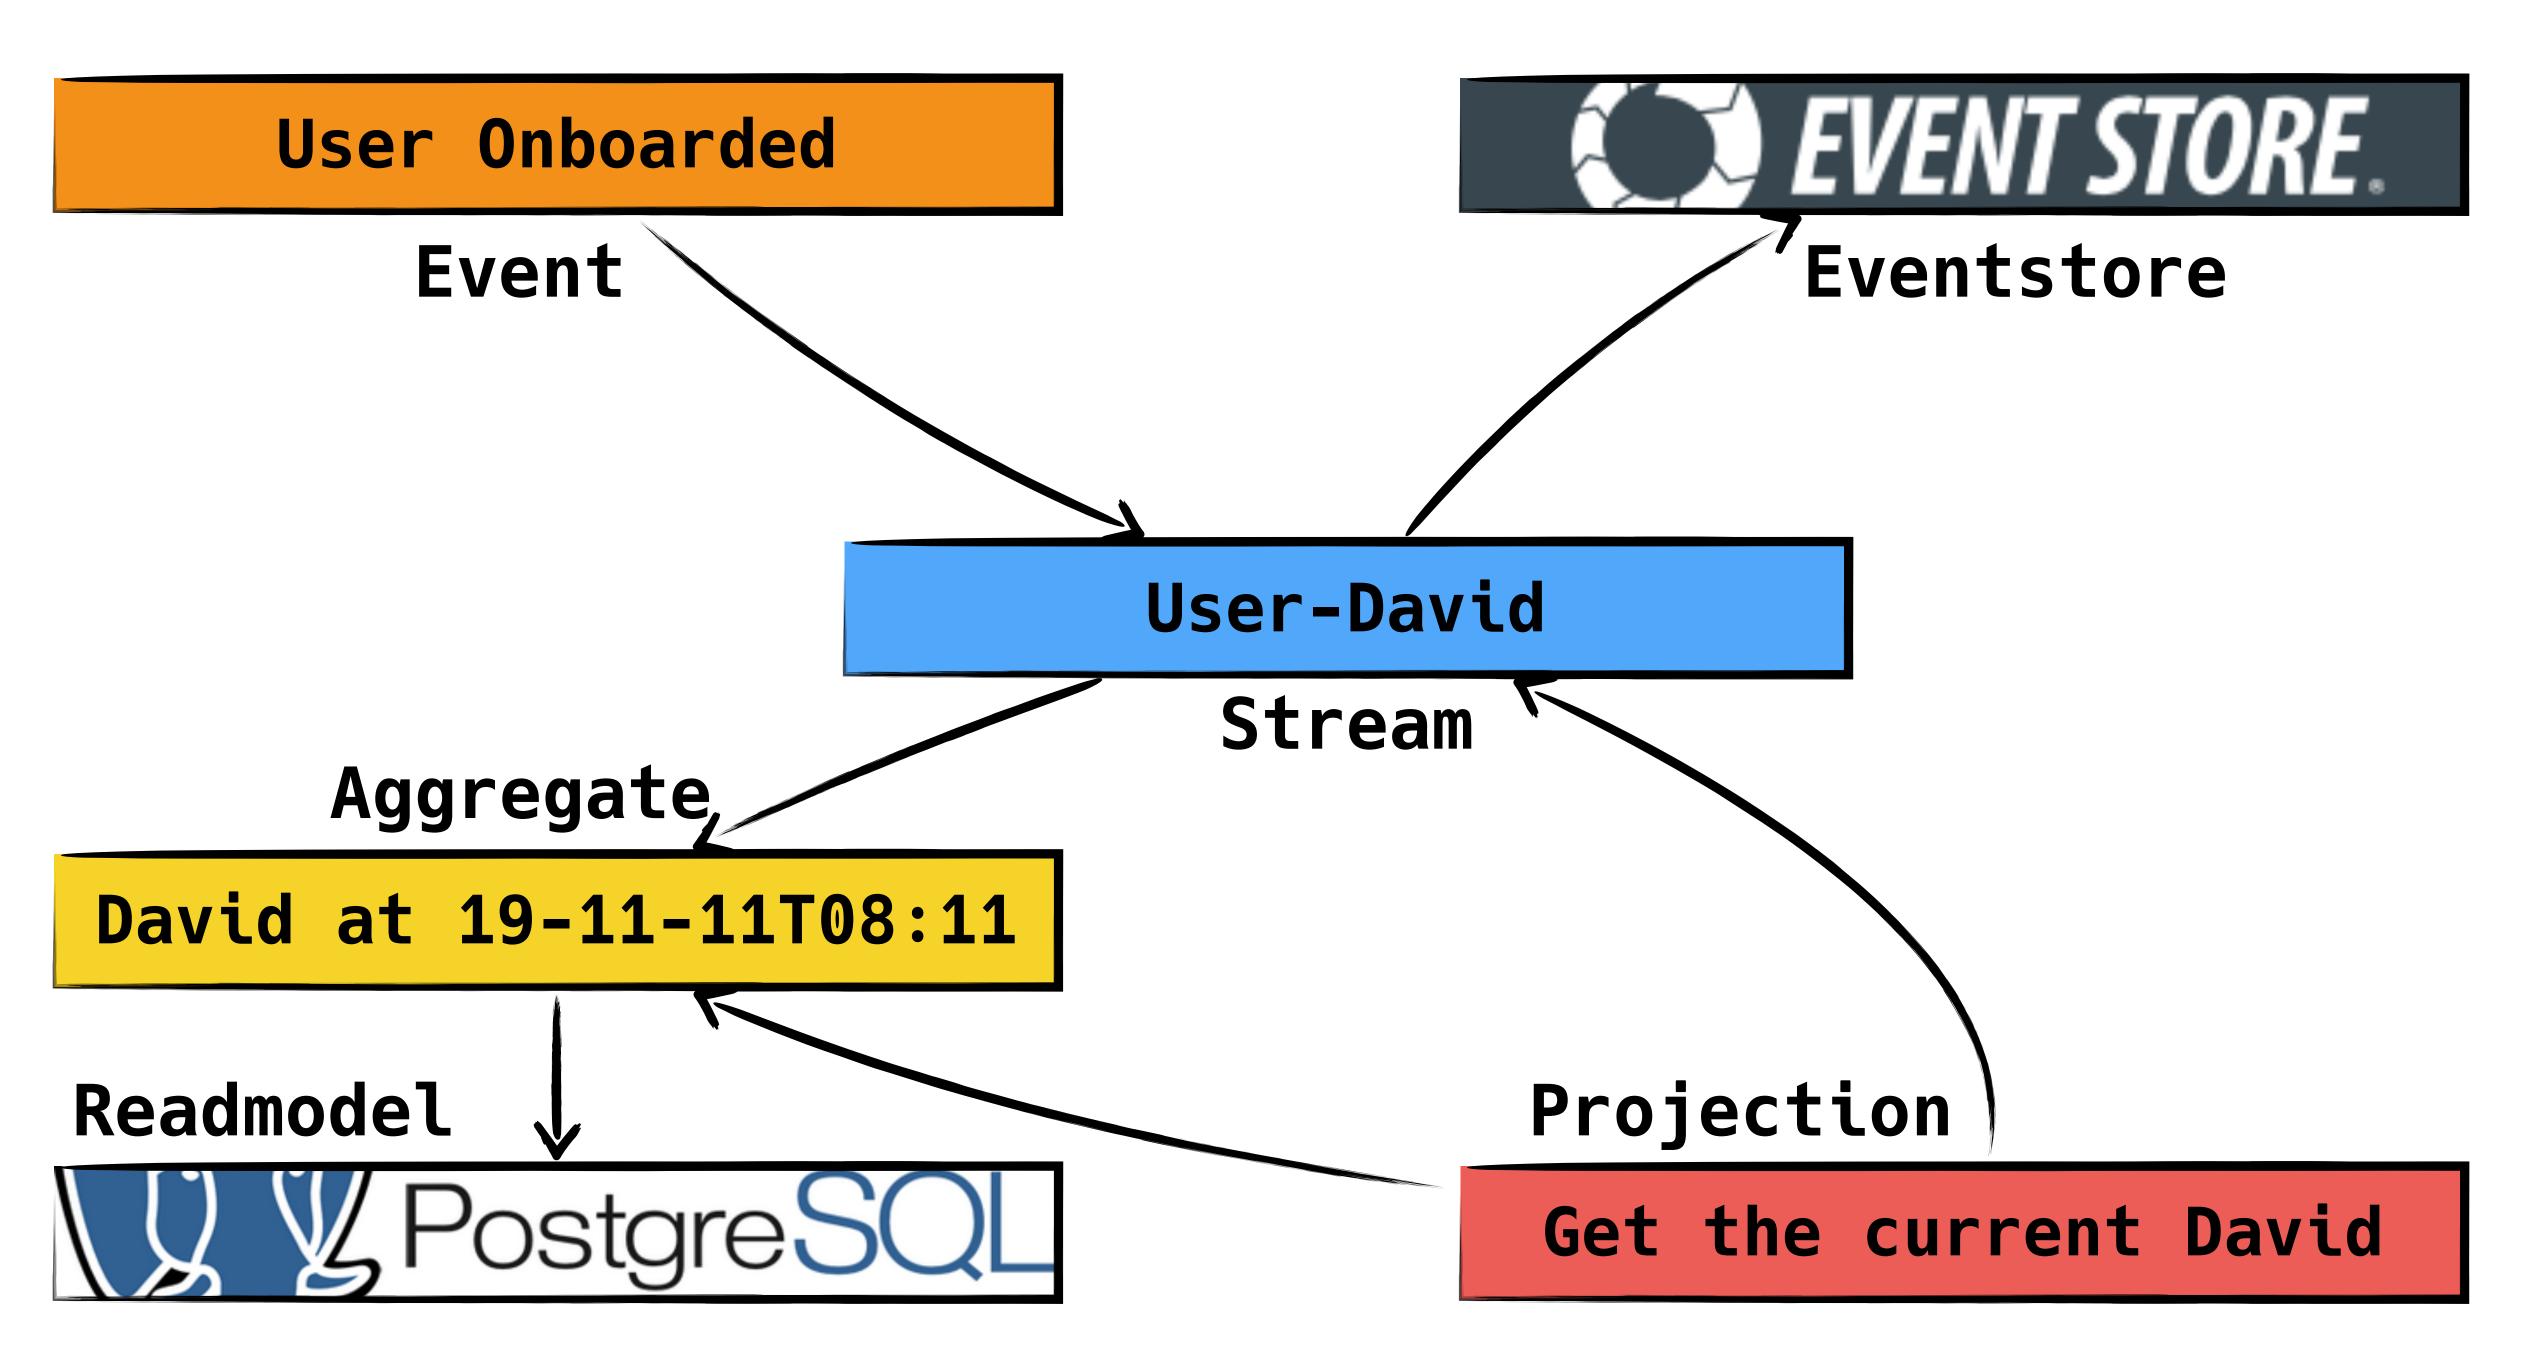
How to represent data and how to represent behaviour?

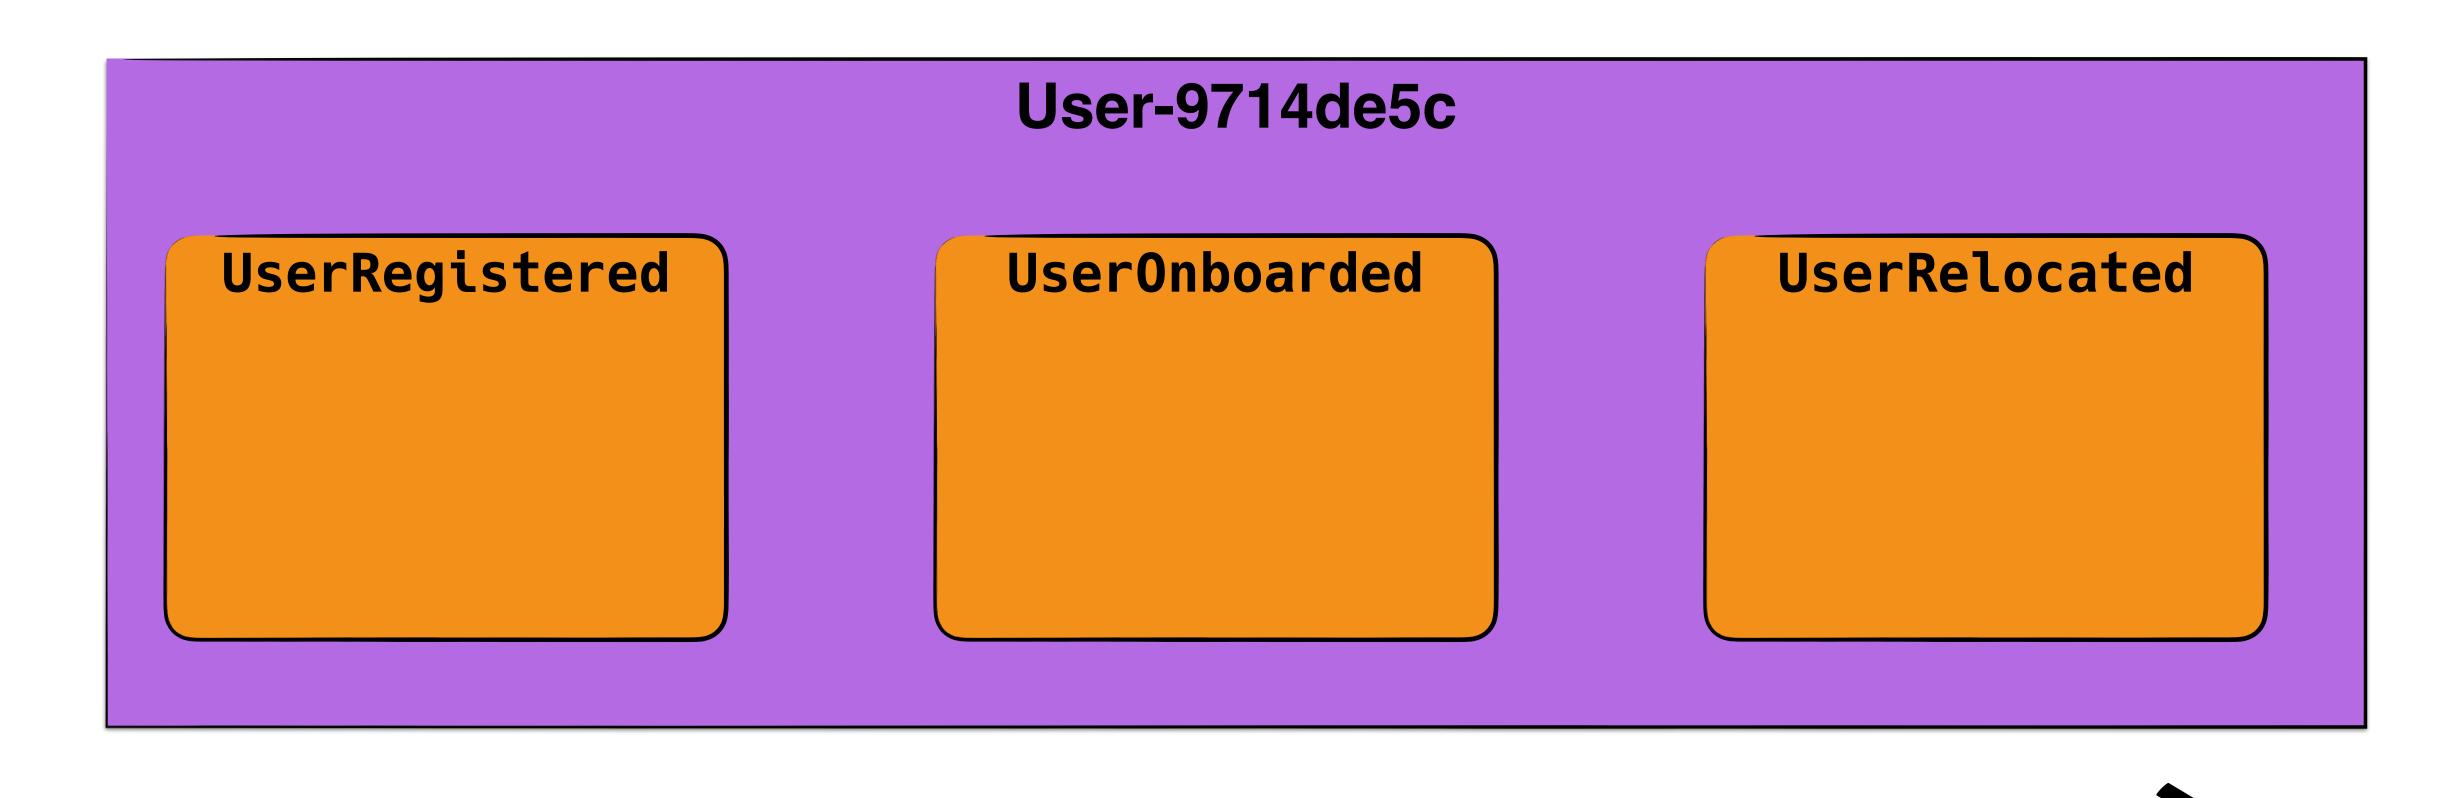




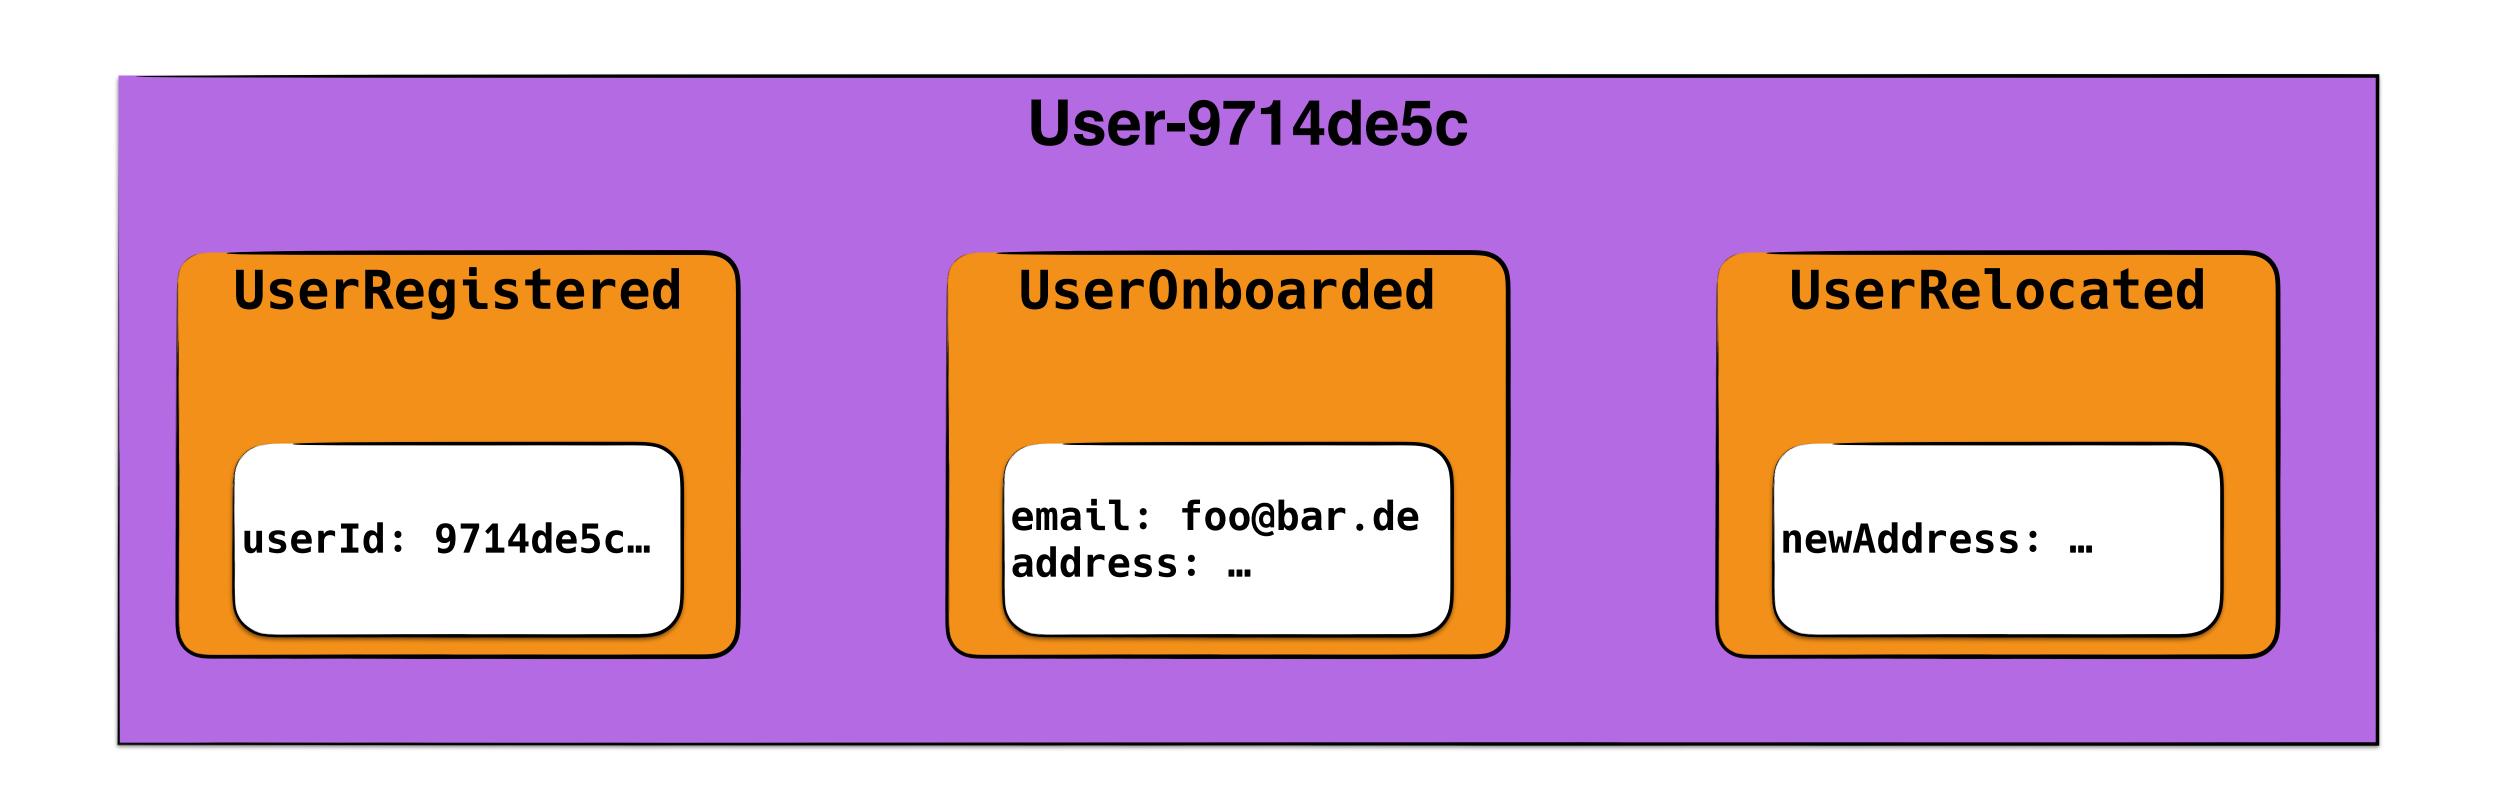


Just use magic

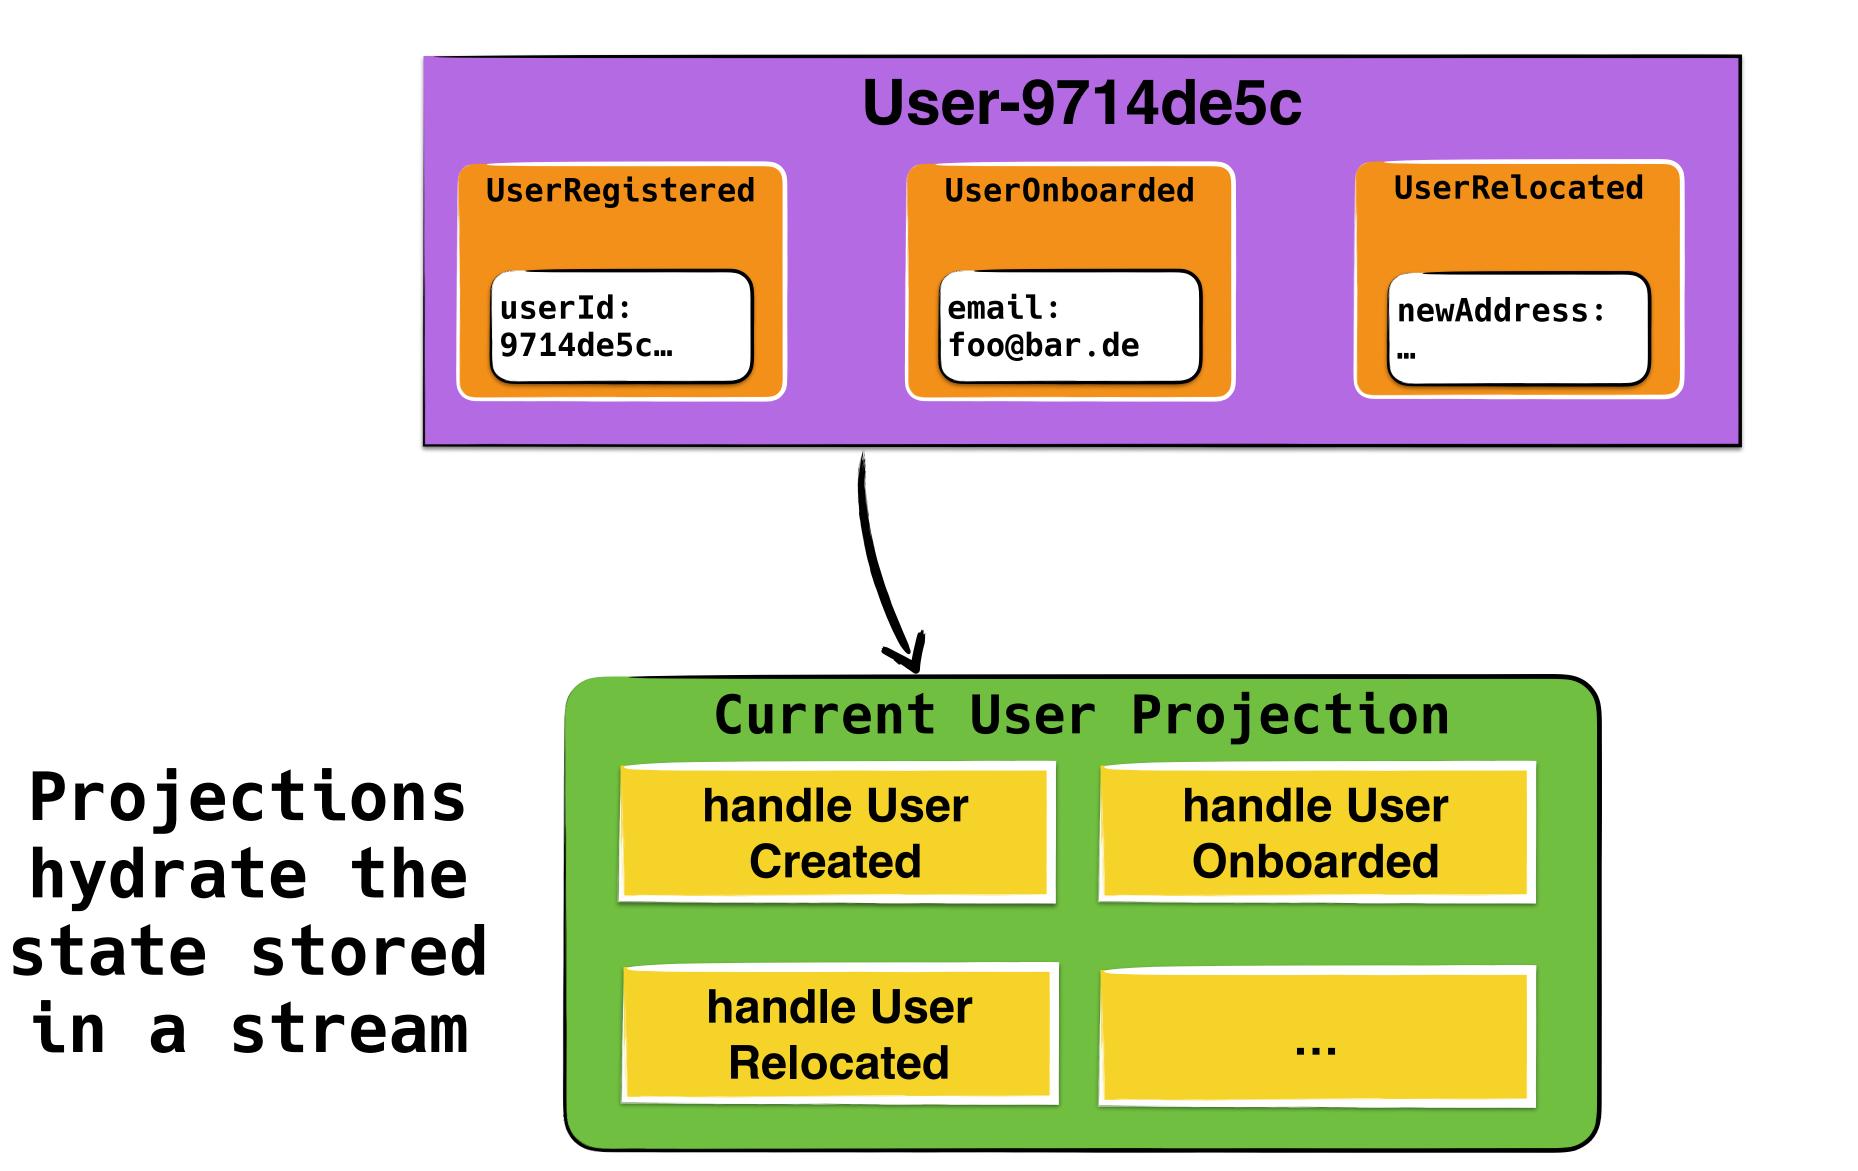


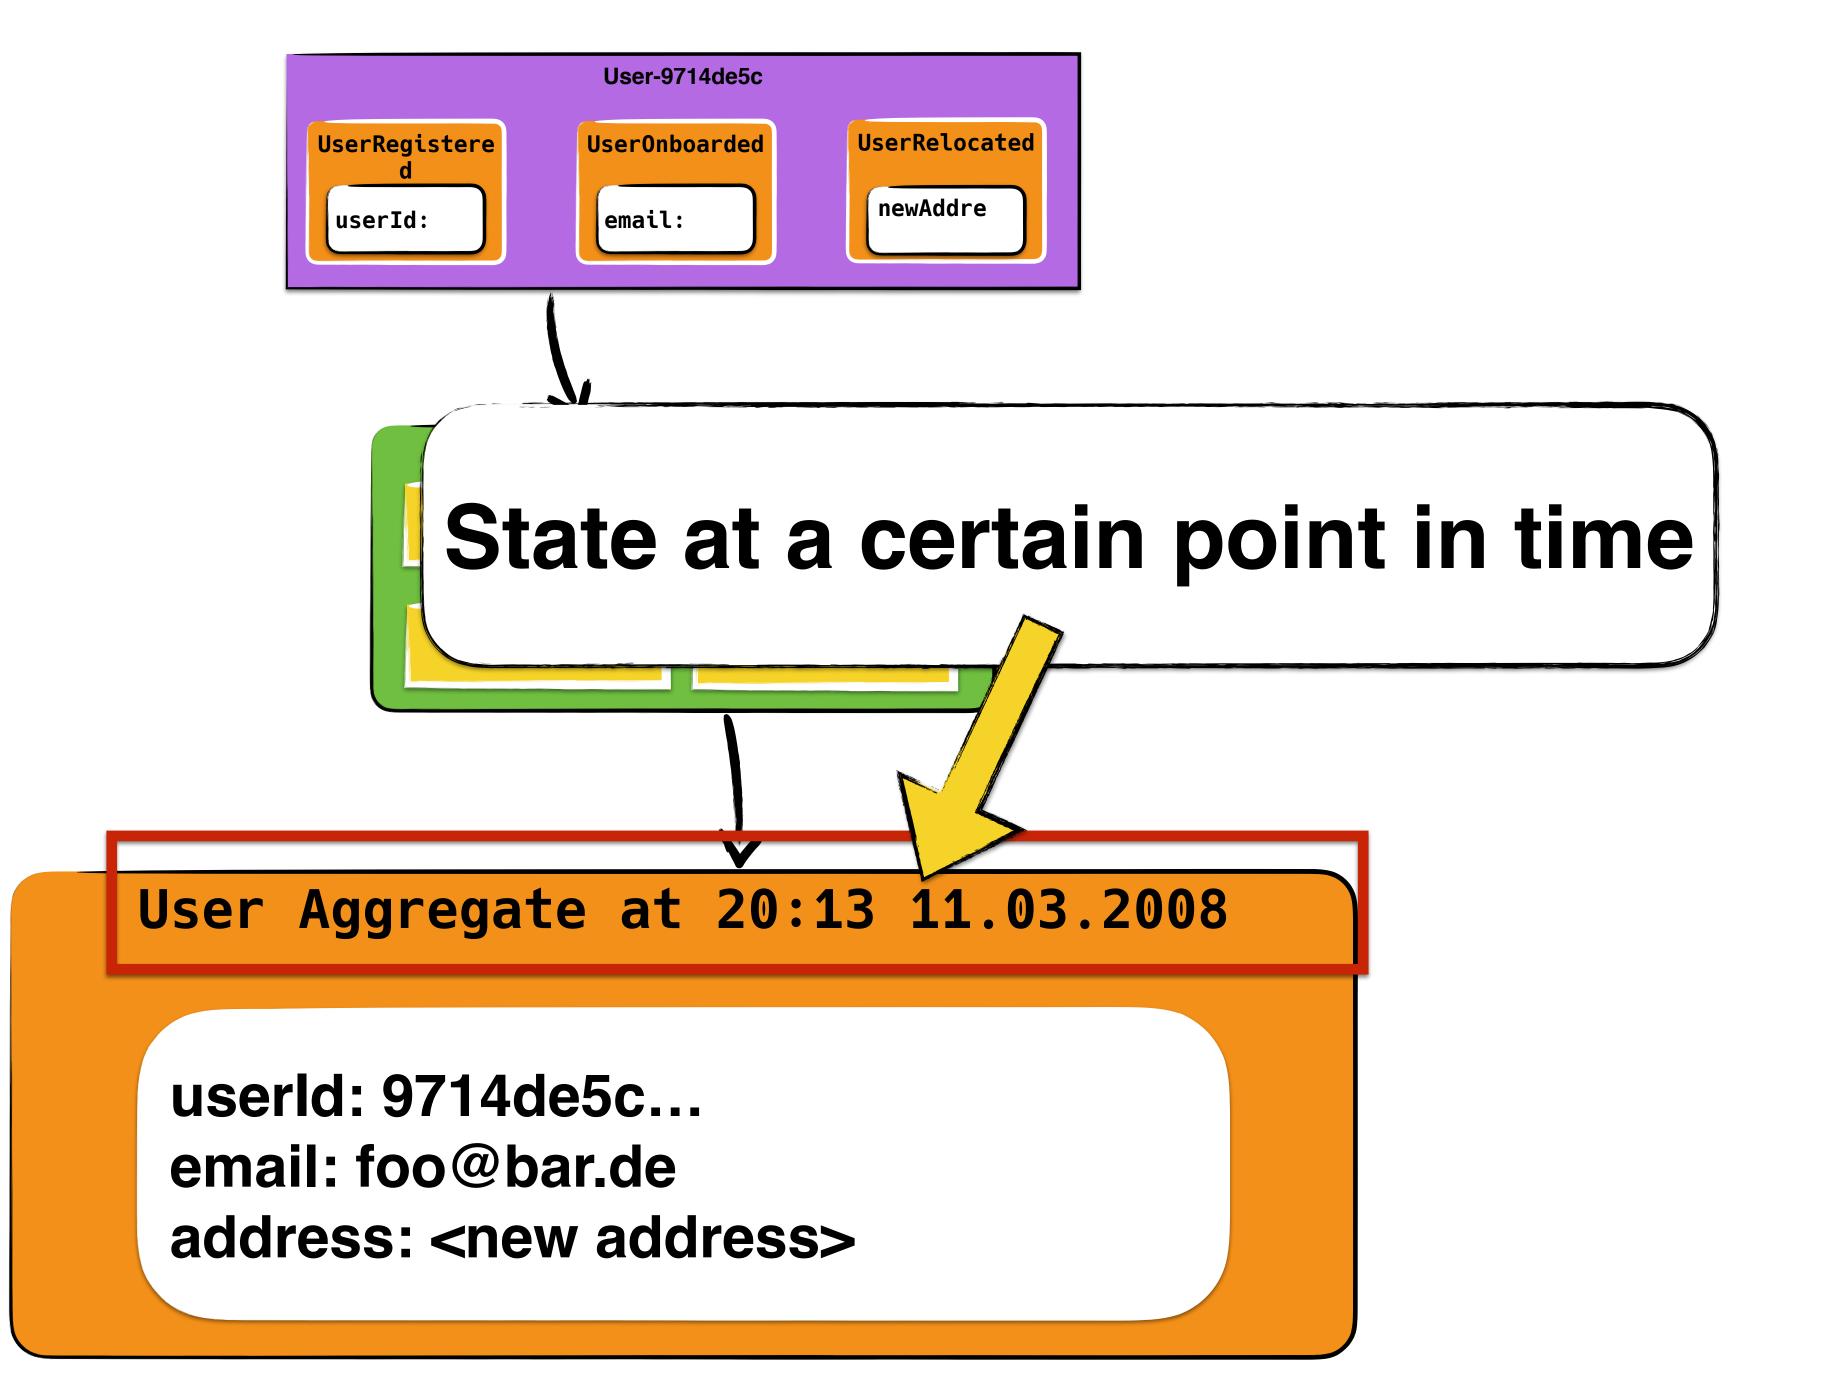


Direction of time



Direction of time





Represents
the user at
the time of
the last
event read

Read models

Transactions and concurrency

Versions, up-front-design

Dealing with errors

GDPR, compliance and eventsourcing



User-9714de5c

UserRegistered

userId: 9714de5c...

UserOnboarded

email: foo@bar.de address: ...

UserRelocated

newAddress: ...

UserRelocated newAddress: ...

User Aggregate at 20:13 11.03.2008

userld: 9714de5c...

email: foo@bar.de

address: <new address>

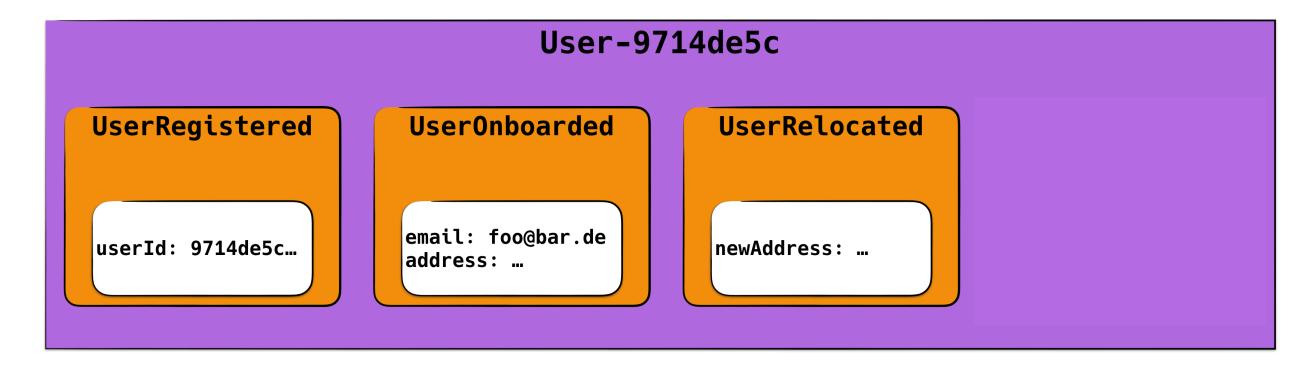
How can we store read models?

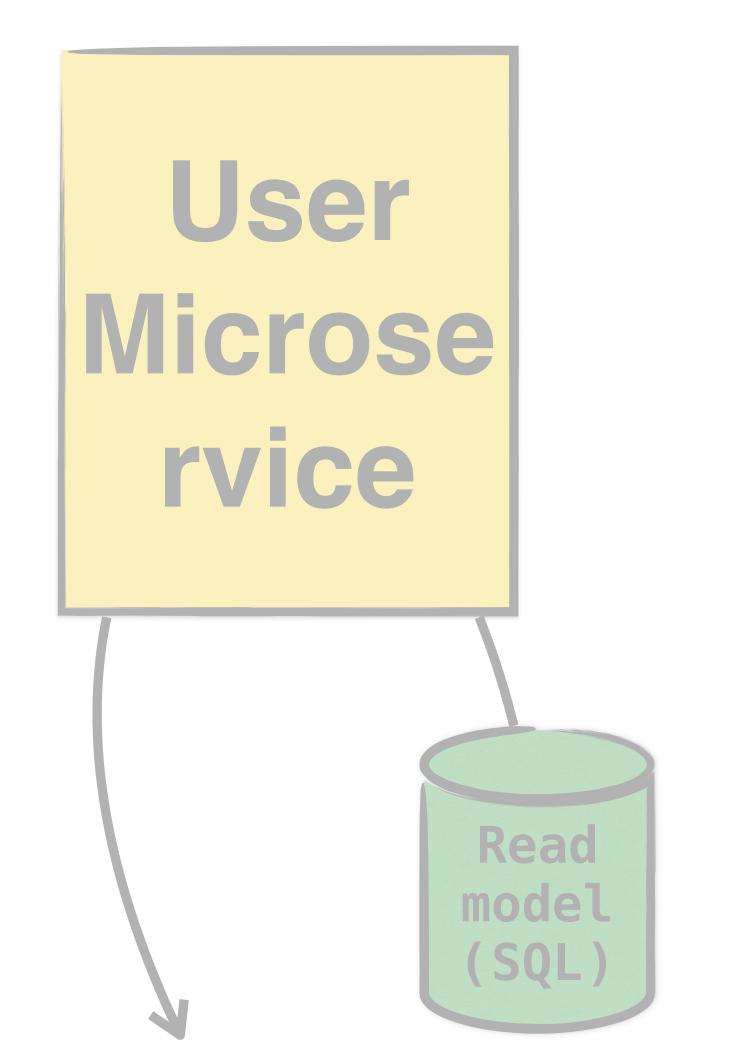
Just use a local database



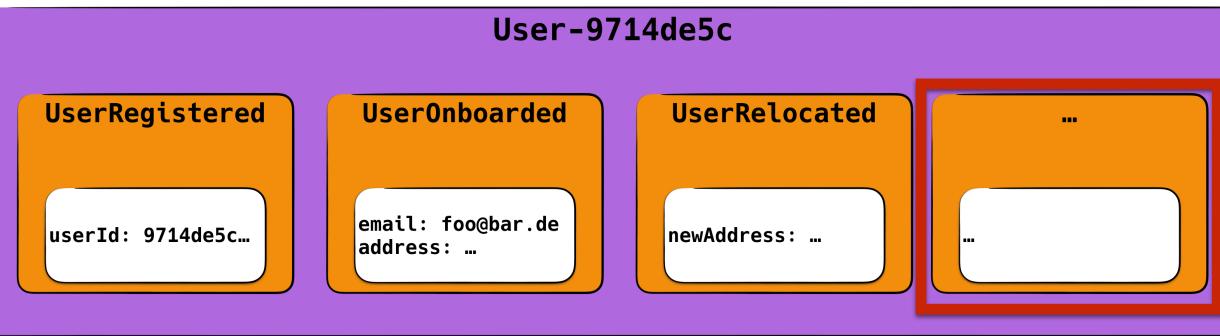


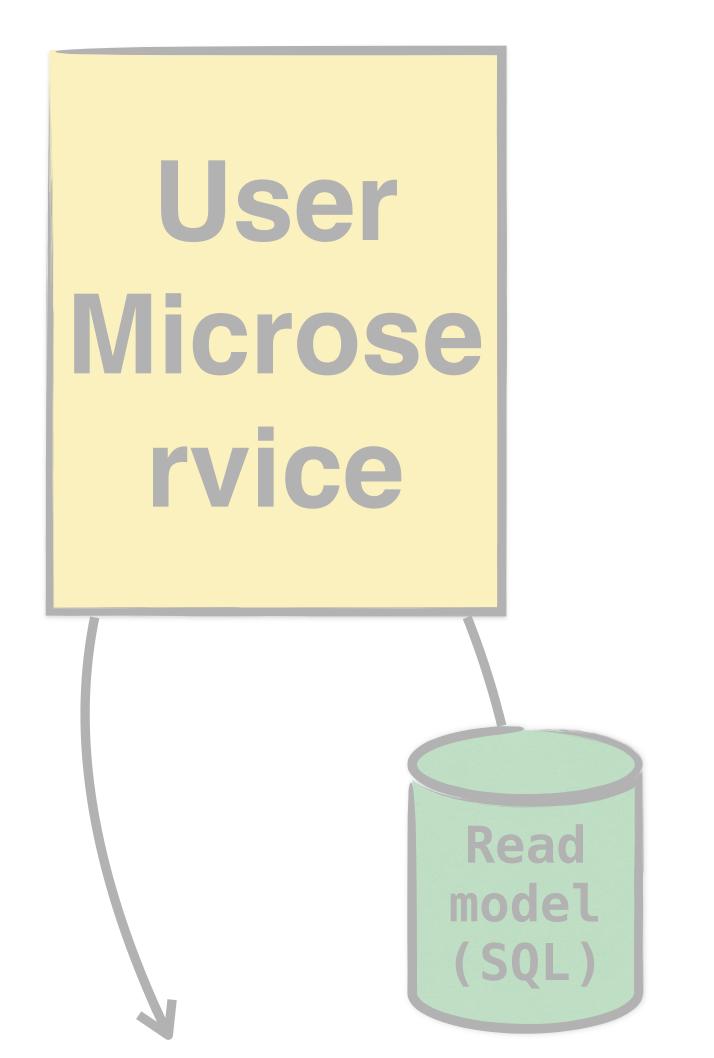
user id	last event	user name	email	street
9741	3	David	foo@bar.de	
4532	7	Martin	null	



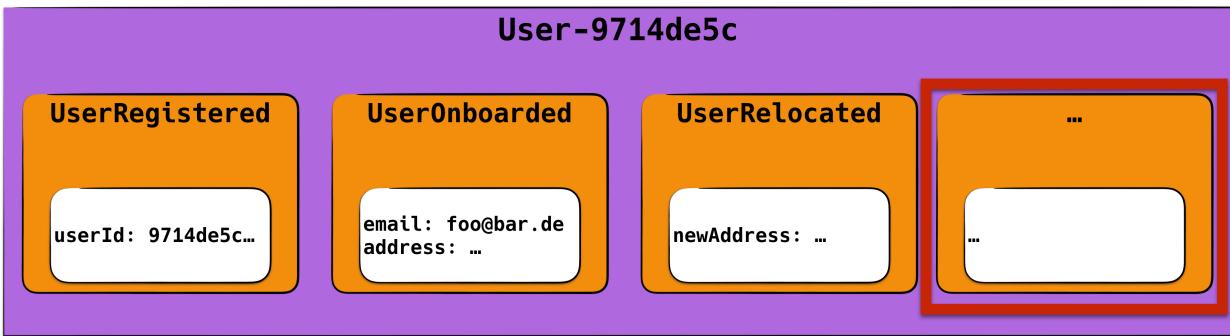


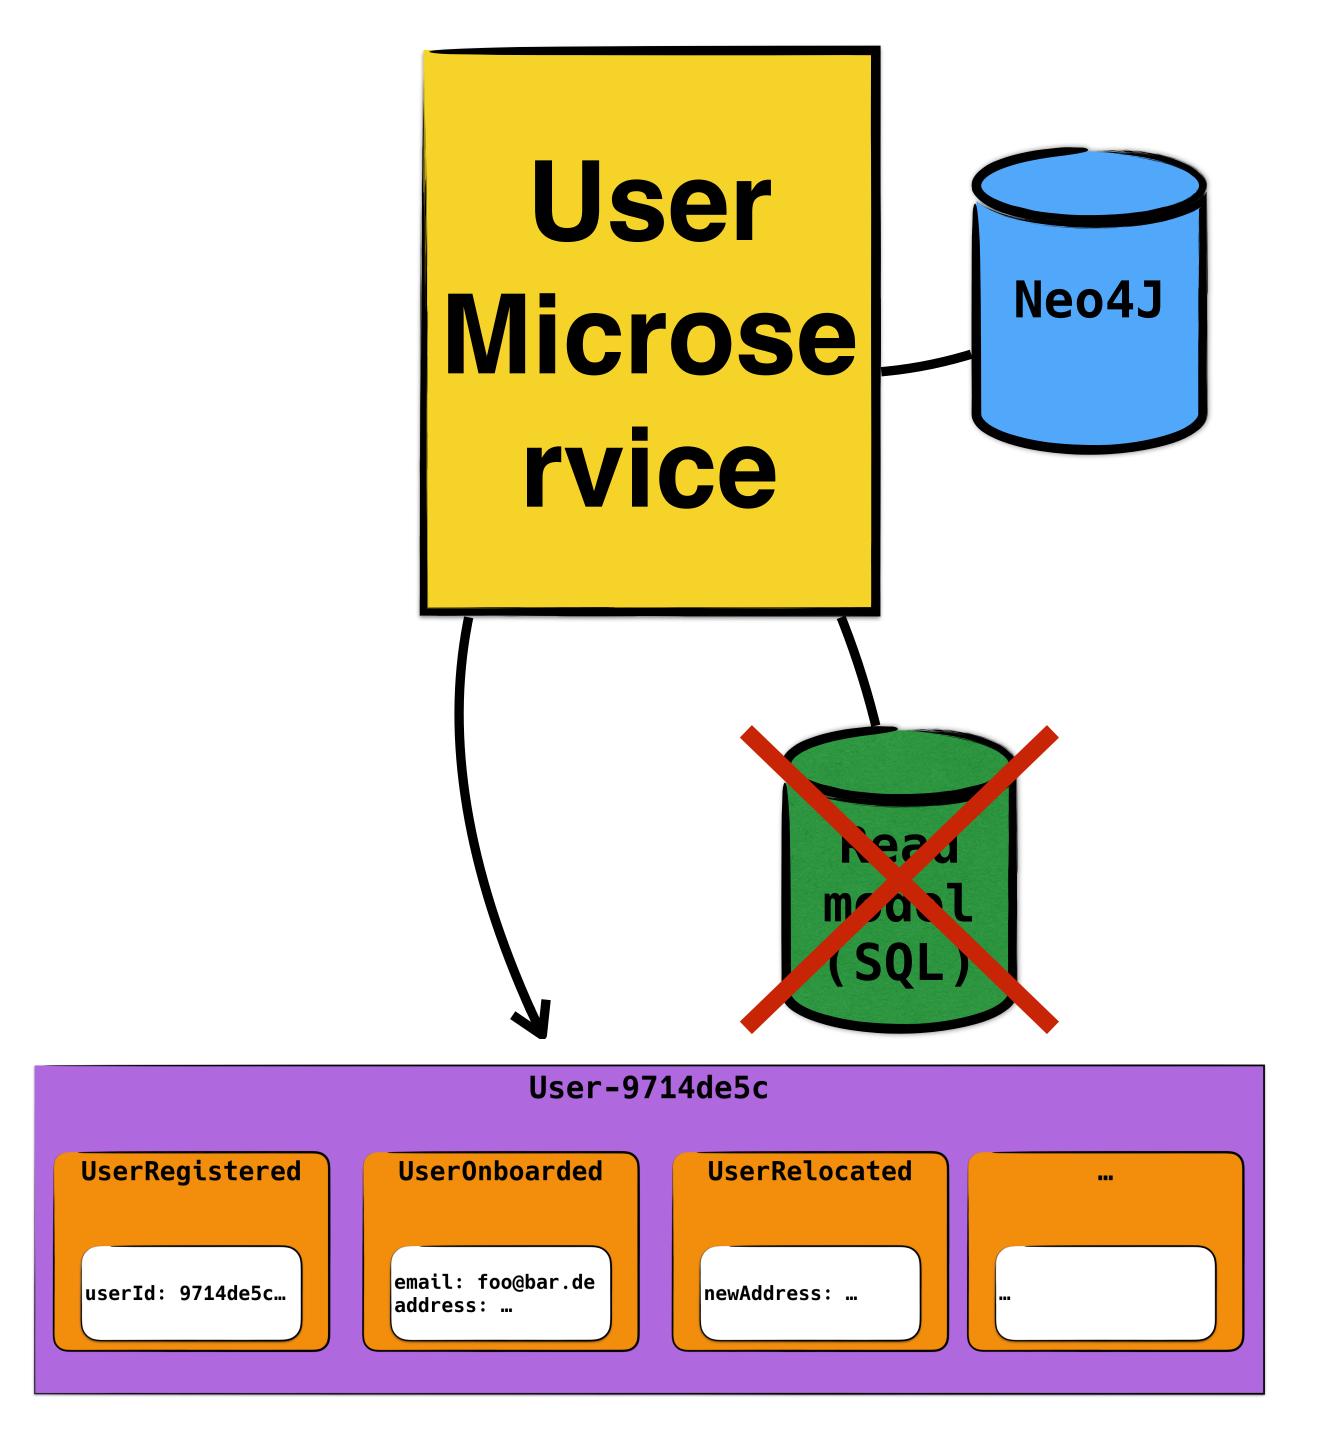
user id	last event id	user name	email	street
9741	3	David	foo@bar.de	
4532	7	Martin	null	





user id	last event id	user name	email	street
9741	4	David	qux@ba.ze	
4532	7	Martin	null	



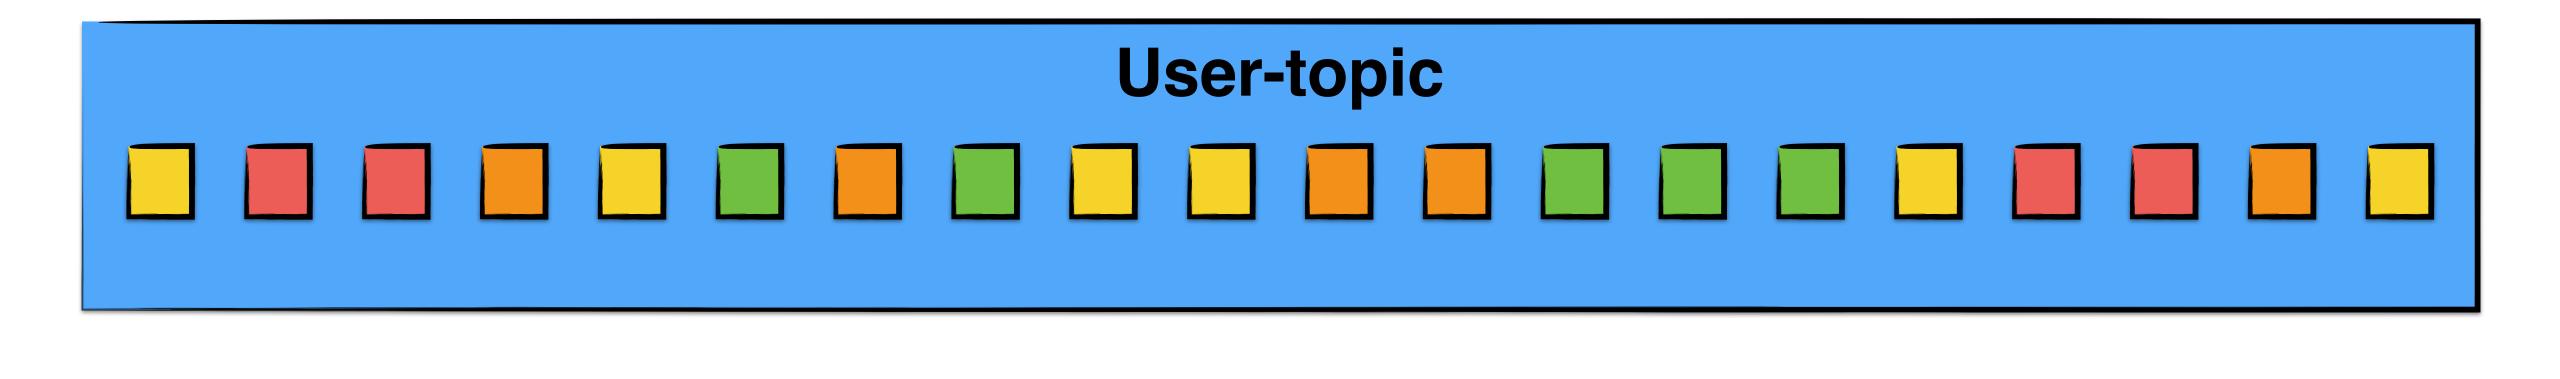


Challenges?

Eventual consistency Zero downtime replays and rebuilds Re-deliveries and effectively once Operational complexity

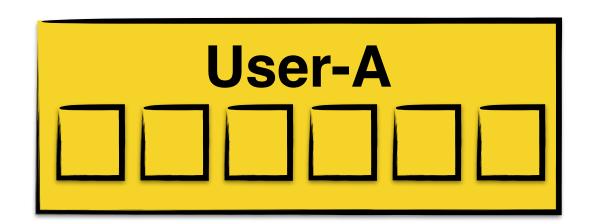
You may not need a read model

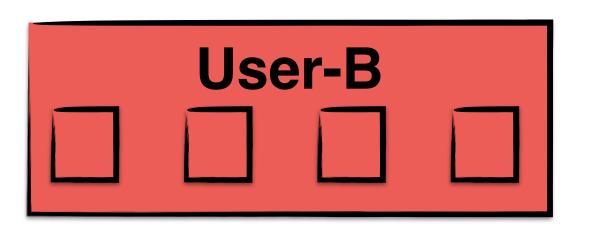
Typical strategies for storing events

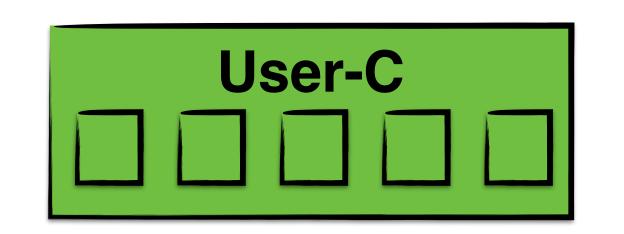


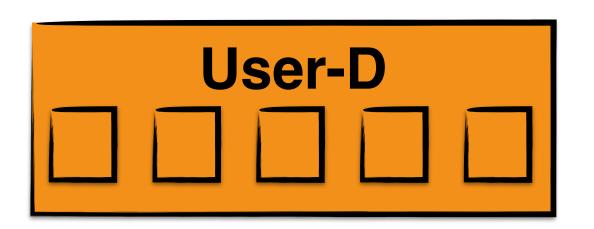
User A User B User C User D

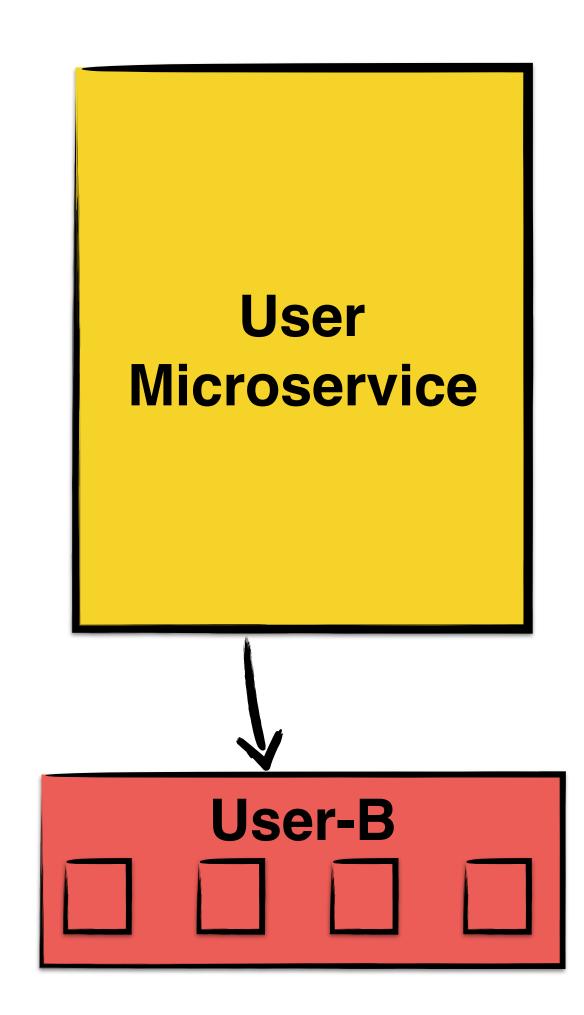
Better: One stream per aggregate











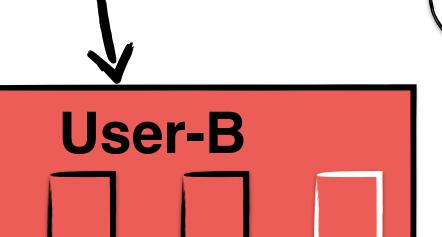
GET /users/B

```
User
Microservice
  User-B
```

```
handlers = {
  'UserCreated': (current, event) => {},
  'UserOnboarded': (current, event) => {},
  'UserDeleted': ...
aggregate = readAggregateFromStream(
  'user',
  'B',
  {},
  fromStartOfStream,
  handlers
```

GET /users/B

User Microservice

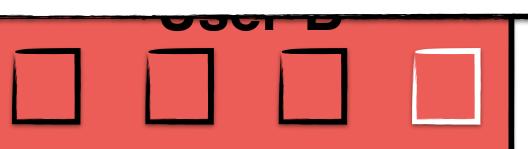


```
handlers = {
         'UserCreated': (current, event) => {},
          'UserOnboarded': (current, event) => {},
          'UserDeleted': ...
                             rgateFromStream(
    Each event type has a
corresponding handler function
         fromStartOfStream,
         handlers
```

GET /users/B

User Microservice

How shall events be handled?



```
handlers = {
  'UserCreated': (current, event) => {},
  'UserOnboarded': (current, event) => {},
  'UserDeleted': ...
aggregate = readAggregateFromStream(
  'user',
                       Which aggregate?
  'B',
  {},
  fromStartOfStream,
  handlers
```

Consistent read No operational overhead Super-simple programming logic

But, what about speed?

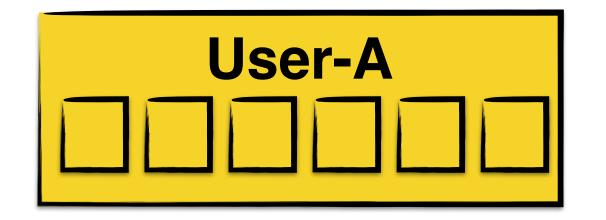
100 events: 66.047ms

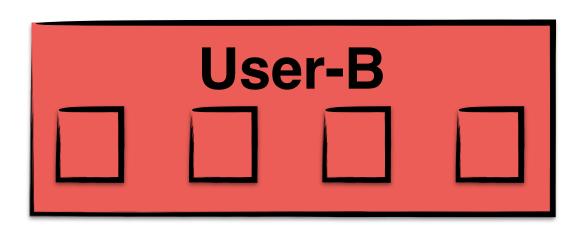
But, what about queries?

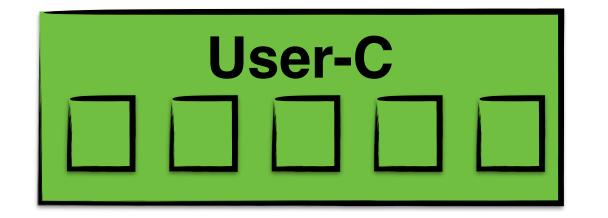
SELECT * FROM USERS WHERE AGE >= 18

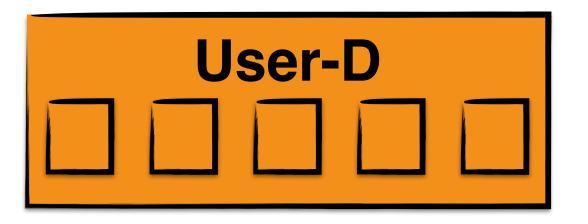
Build specific query projections

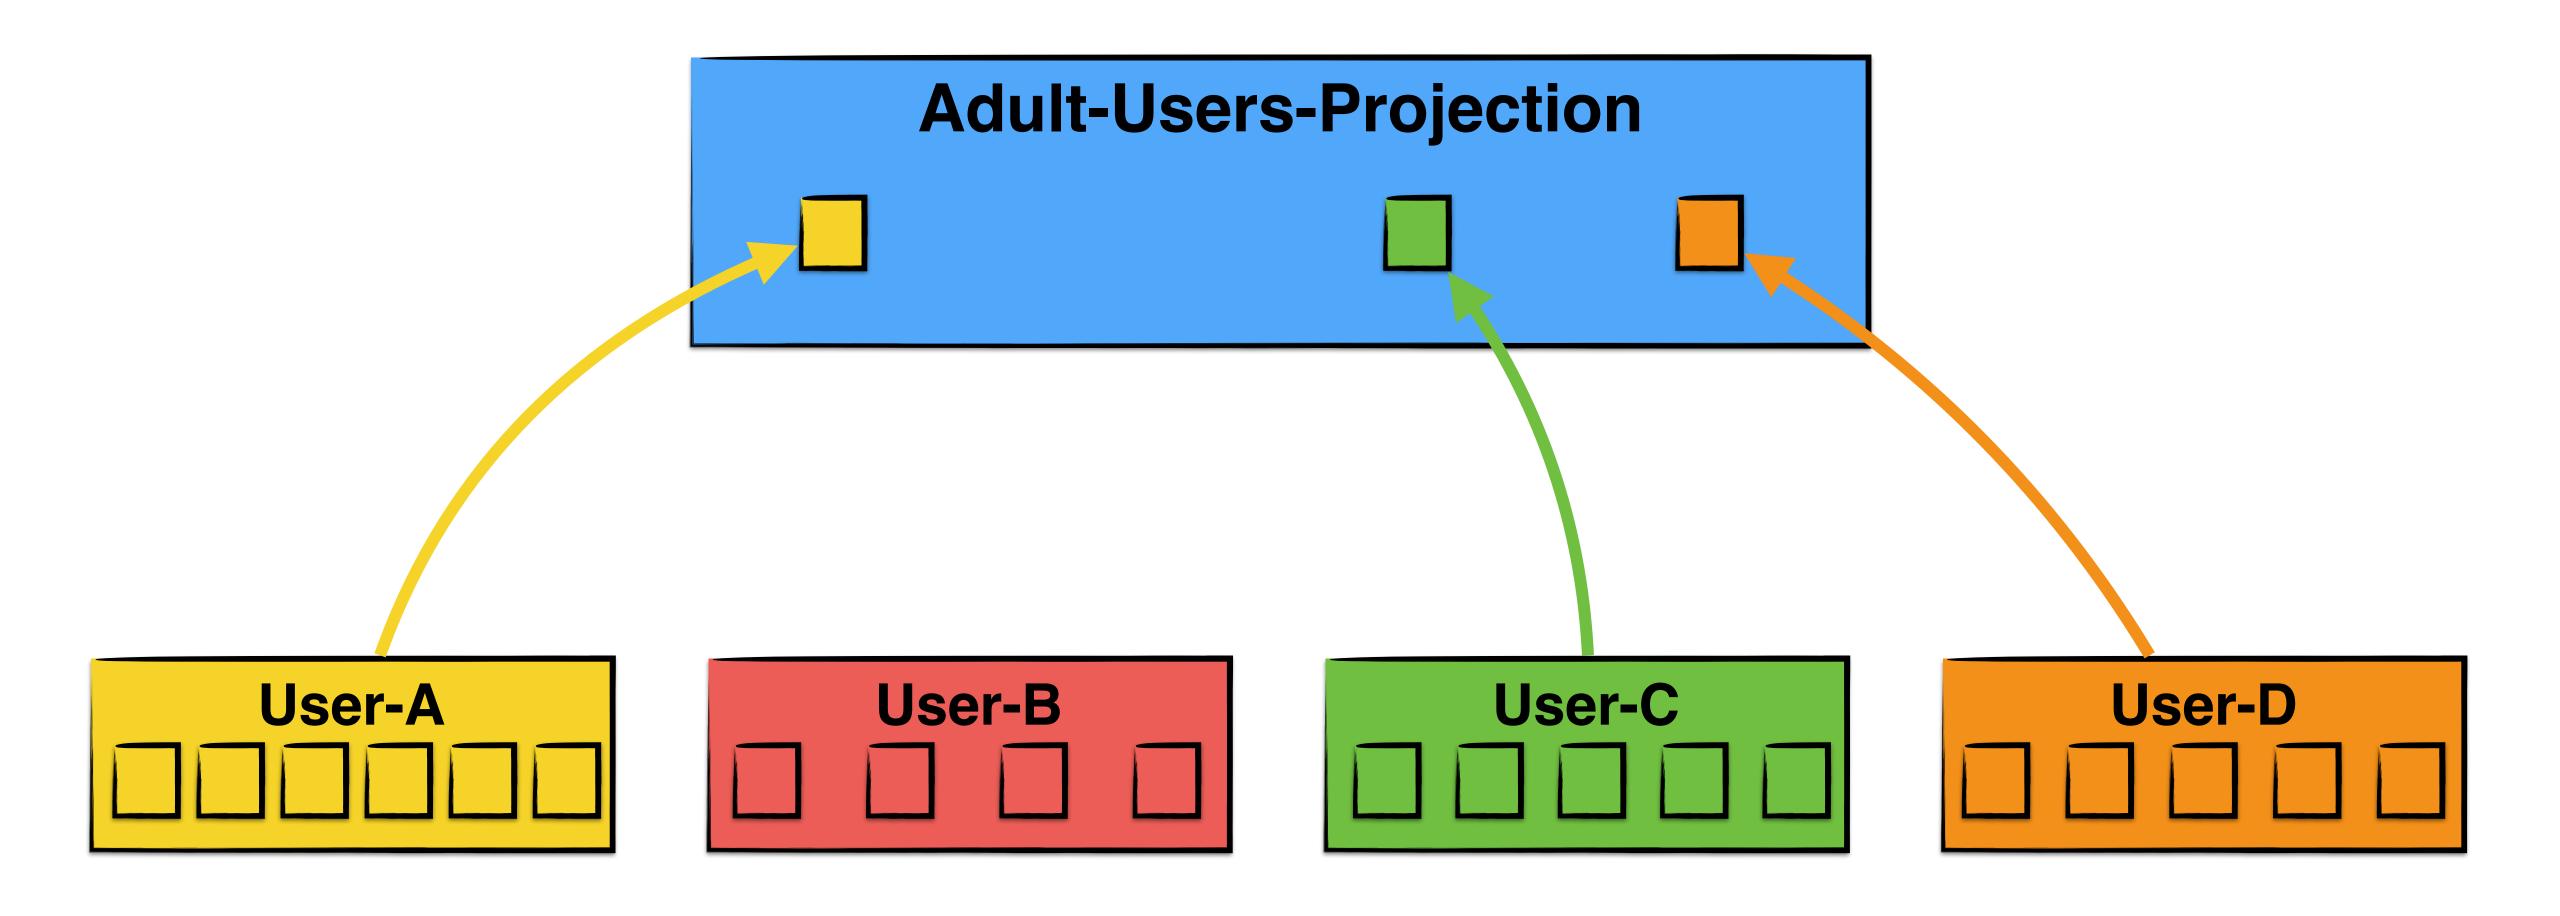
Adult-Users-Projection











Handlers are often trivial Prefer small aggregates Measure first Introduce read models only ifnedded

Transactions, concurrency and your eventstore





How can we guarantee correctness when writing?

The outcome of a business operation depends on the order of events

"Only withdraw money, if the bank account holds enough money!"

"Only withdraw money, if the bank account holds enough money!"*

*Actually, a real bank would not want such a business rule. They earn money if you overdraw your account. An overdraft fee is one of the most expensive fees banks charge. Just saying...

MoneyDeposited

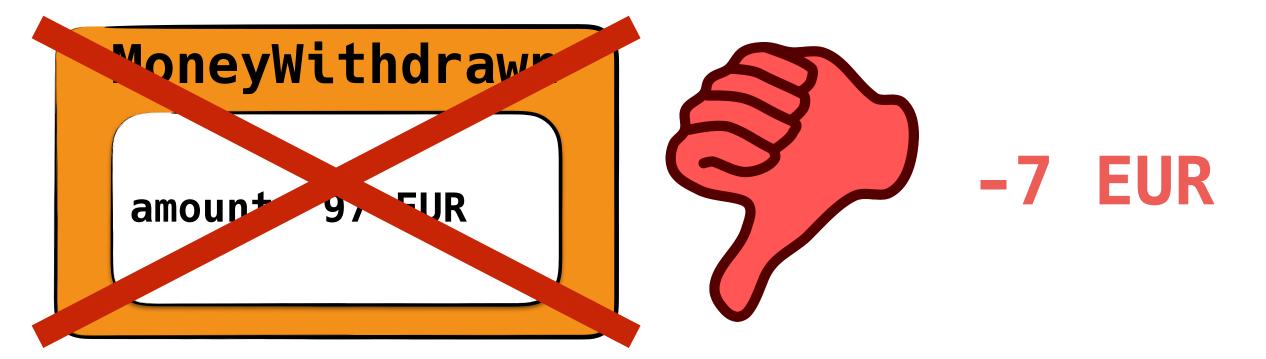
amount: 100 EUR

100 EUR

MoneyWithdrawn

amount: 10 EUR

90 EUR



MoneyDeposited

amount: 50 EUR

140 EUR

Let's shuffle

MoneyDeposited

amount: 100 EUR

100 EUR

MoneyWithdrawn

amount: 10 EUR

90 EUR

MoneyWithdrawn

amount: 97 EUR

140 EUR

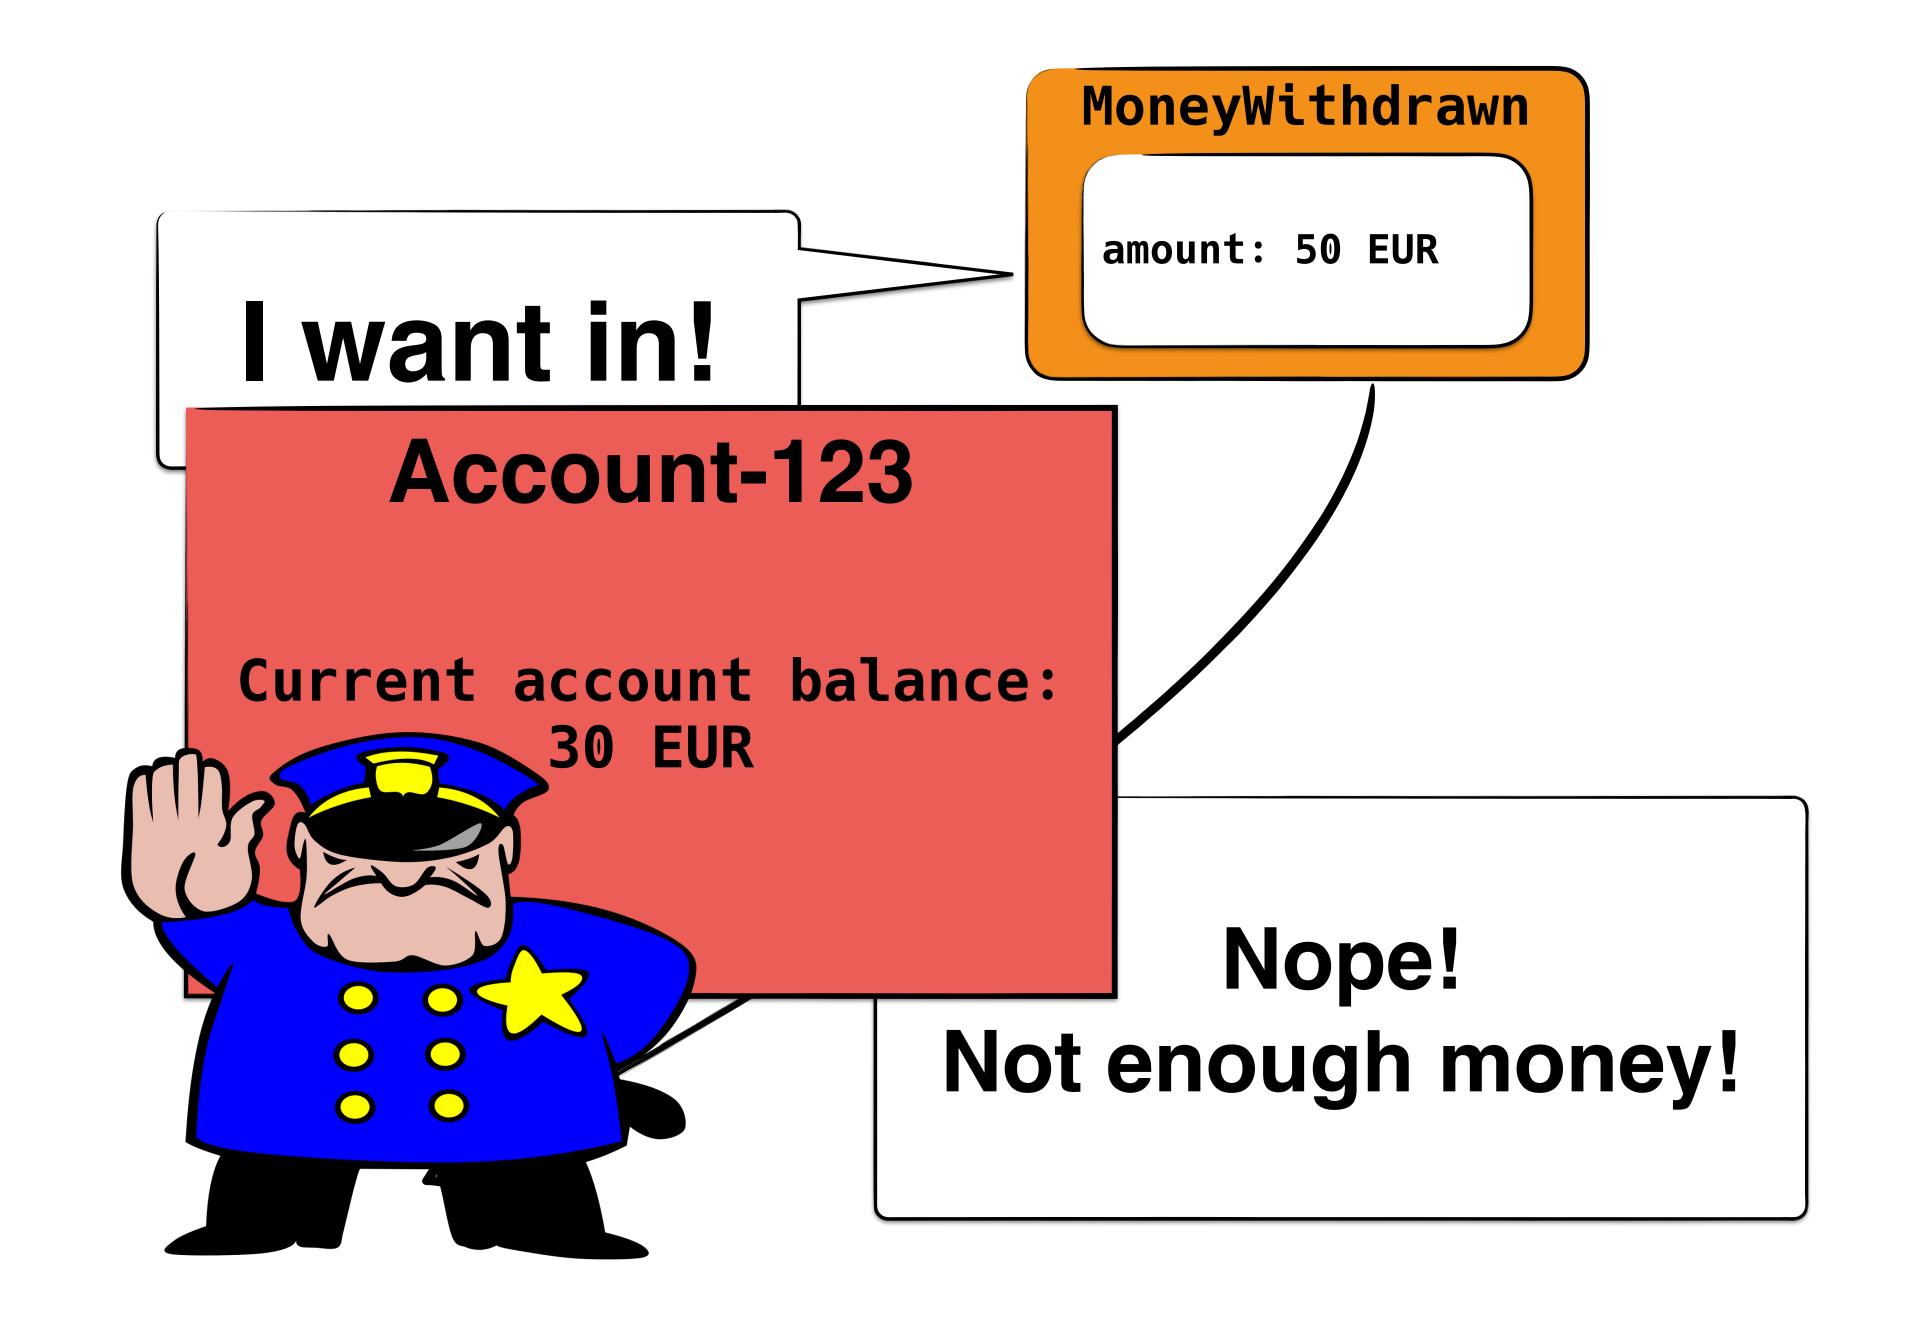


MoneyDeposited

amount: 50 EUR

43 EUR

The aggregate is responsible for enforcing business invariants



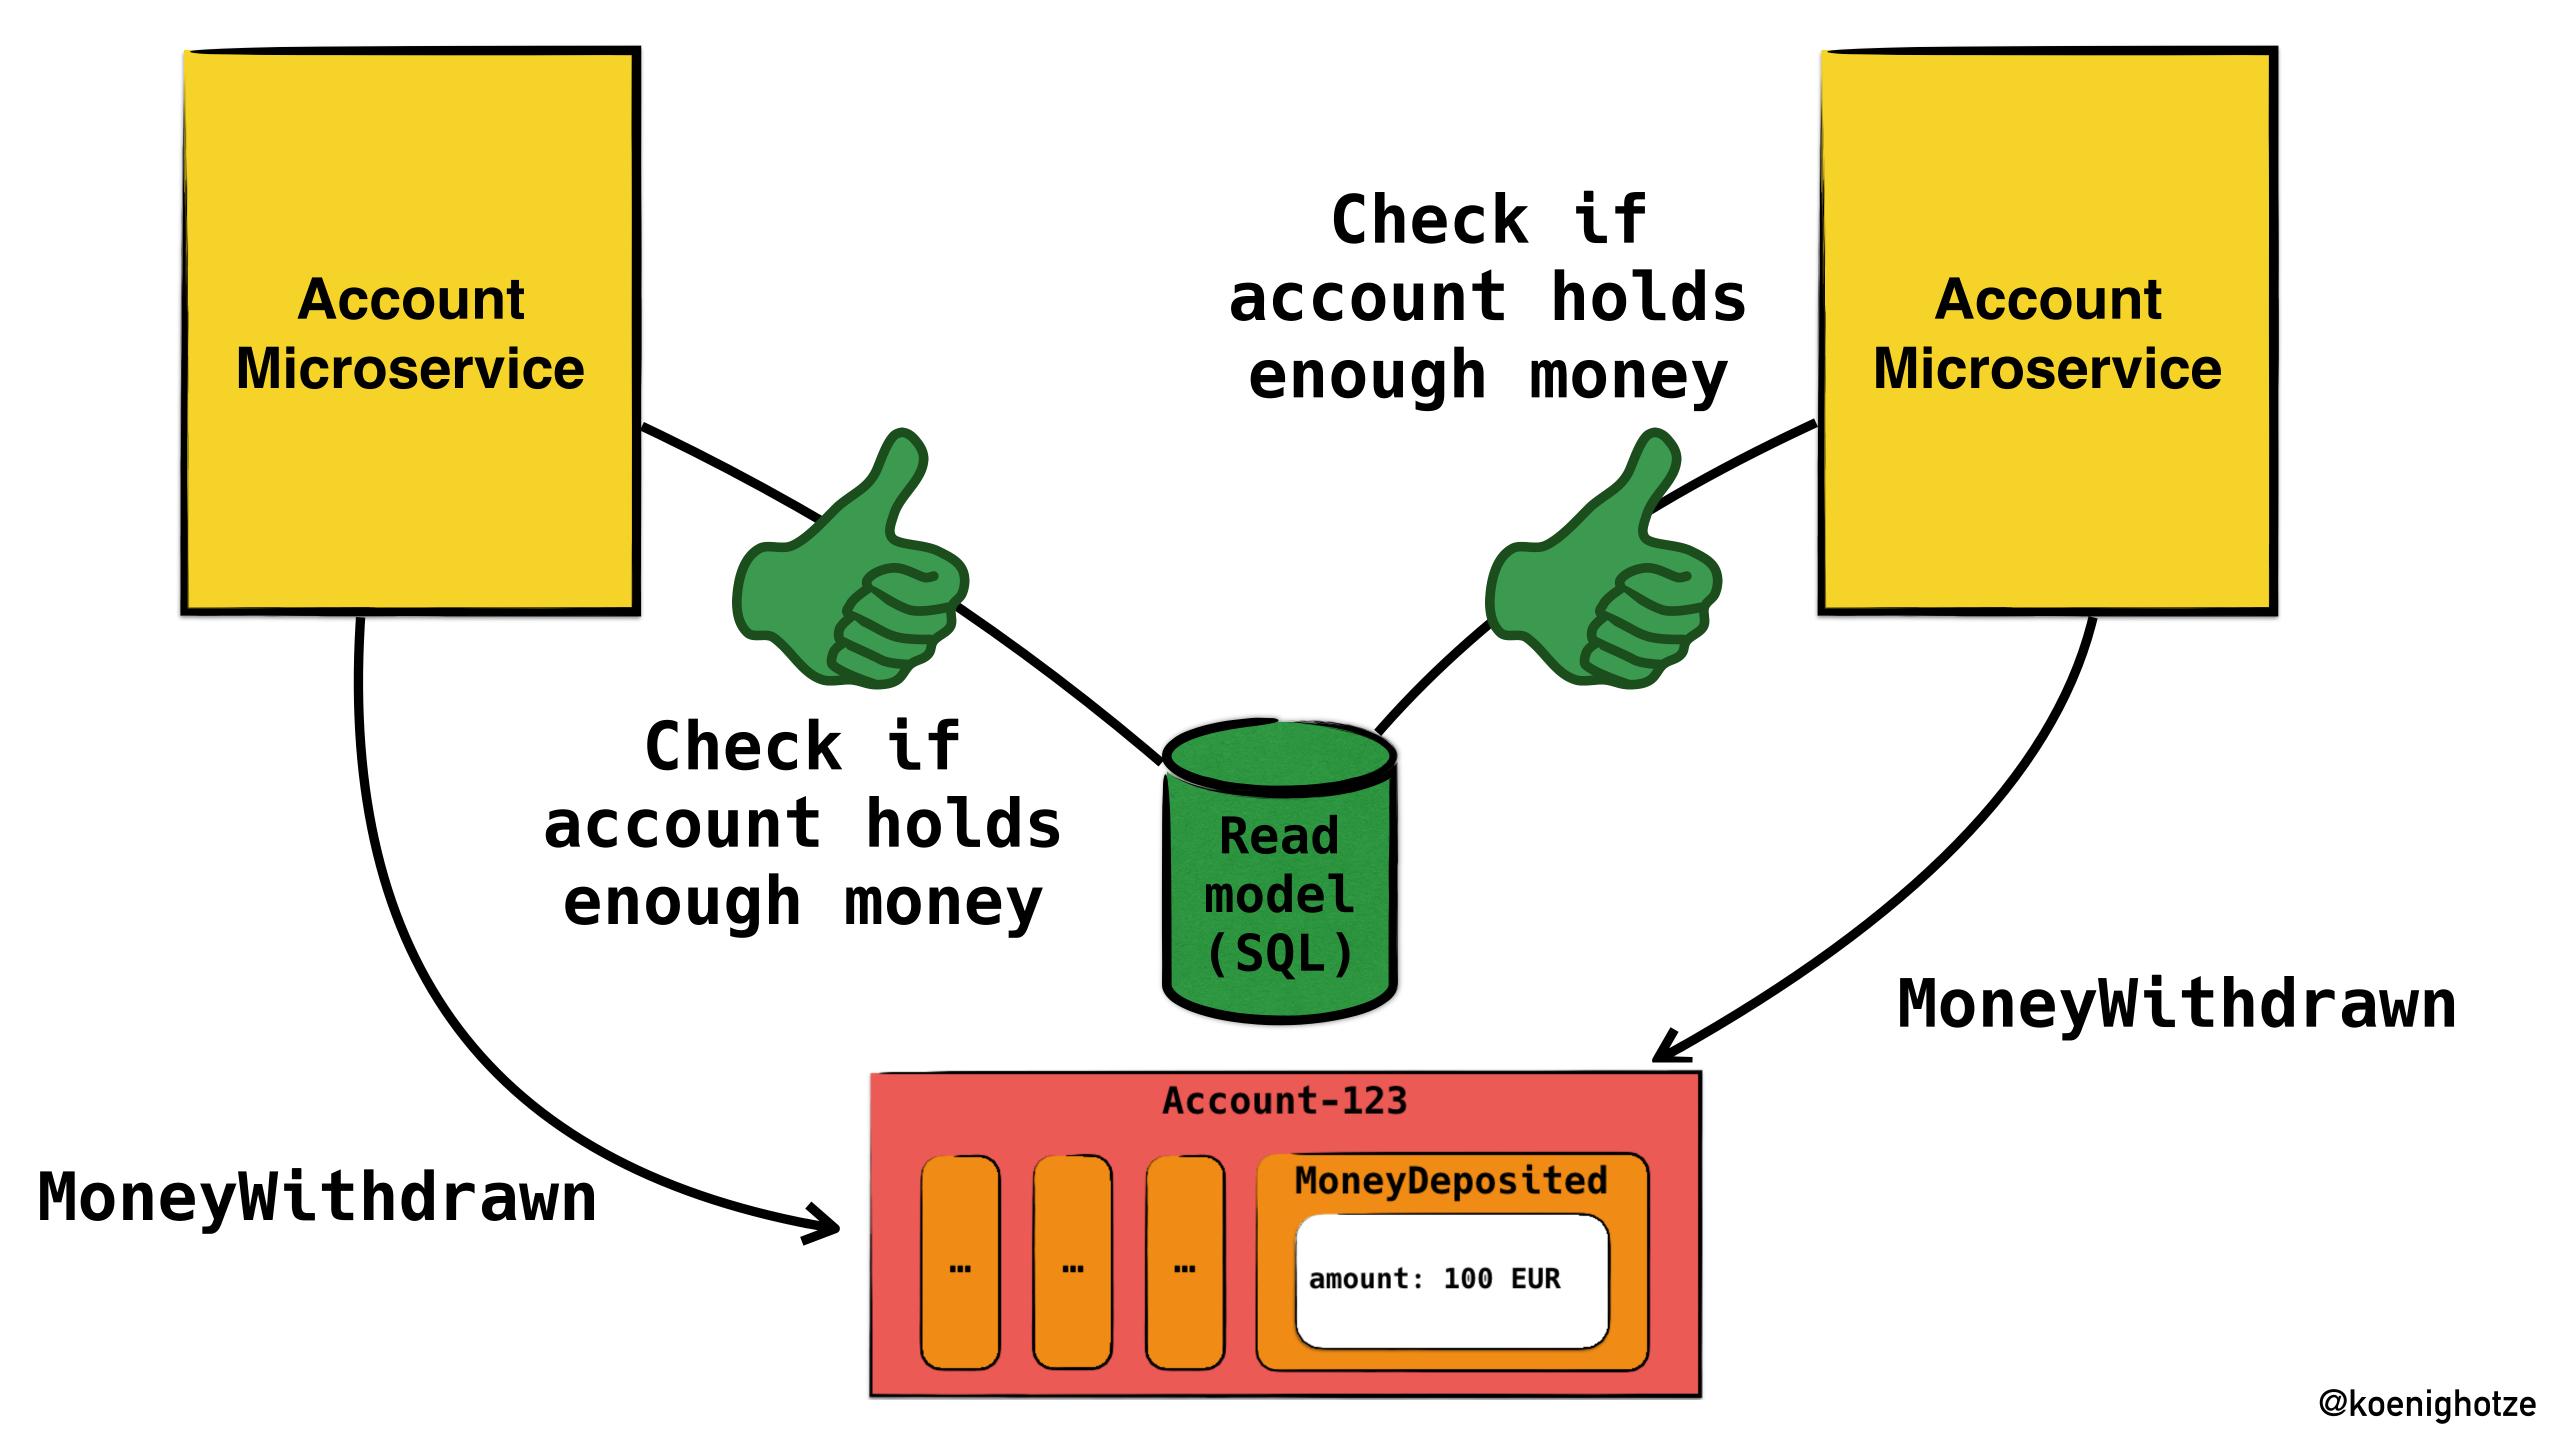
The aggregate is the "transaction boundary"

Just use a database for handling transactions



Maybe not

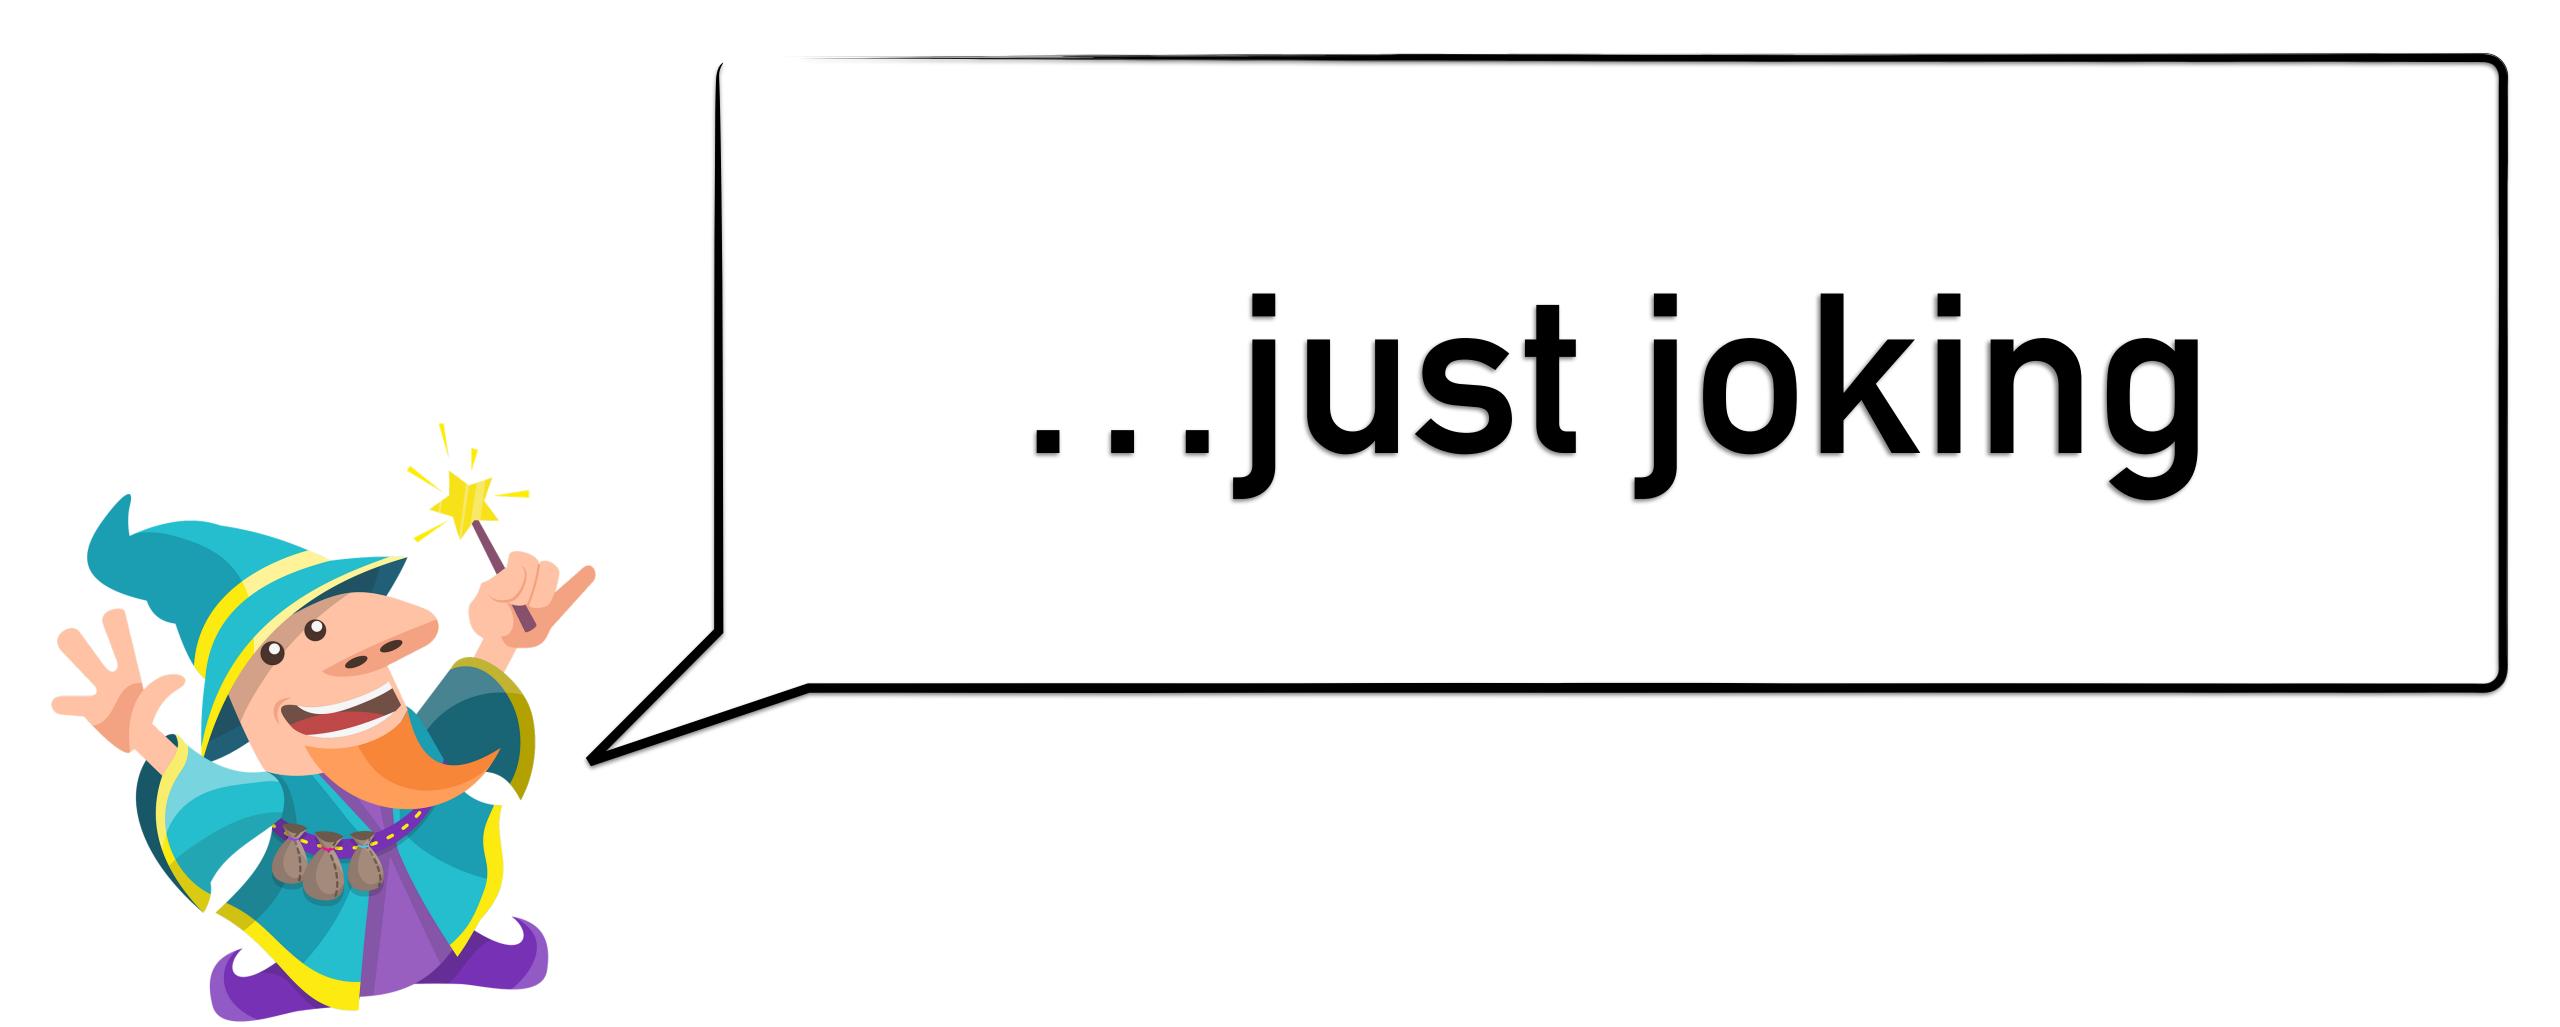
Are YOUR systems single-threaded?



Validation against a read model is prone to inconsistencies

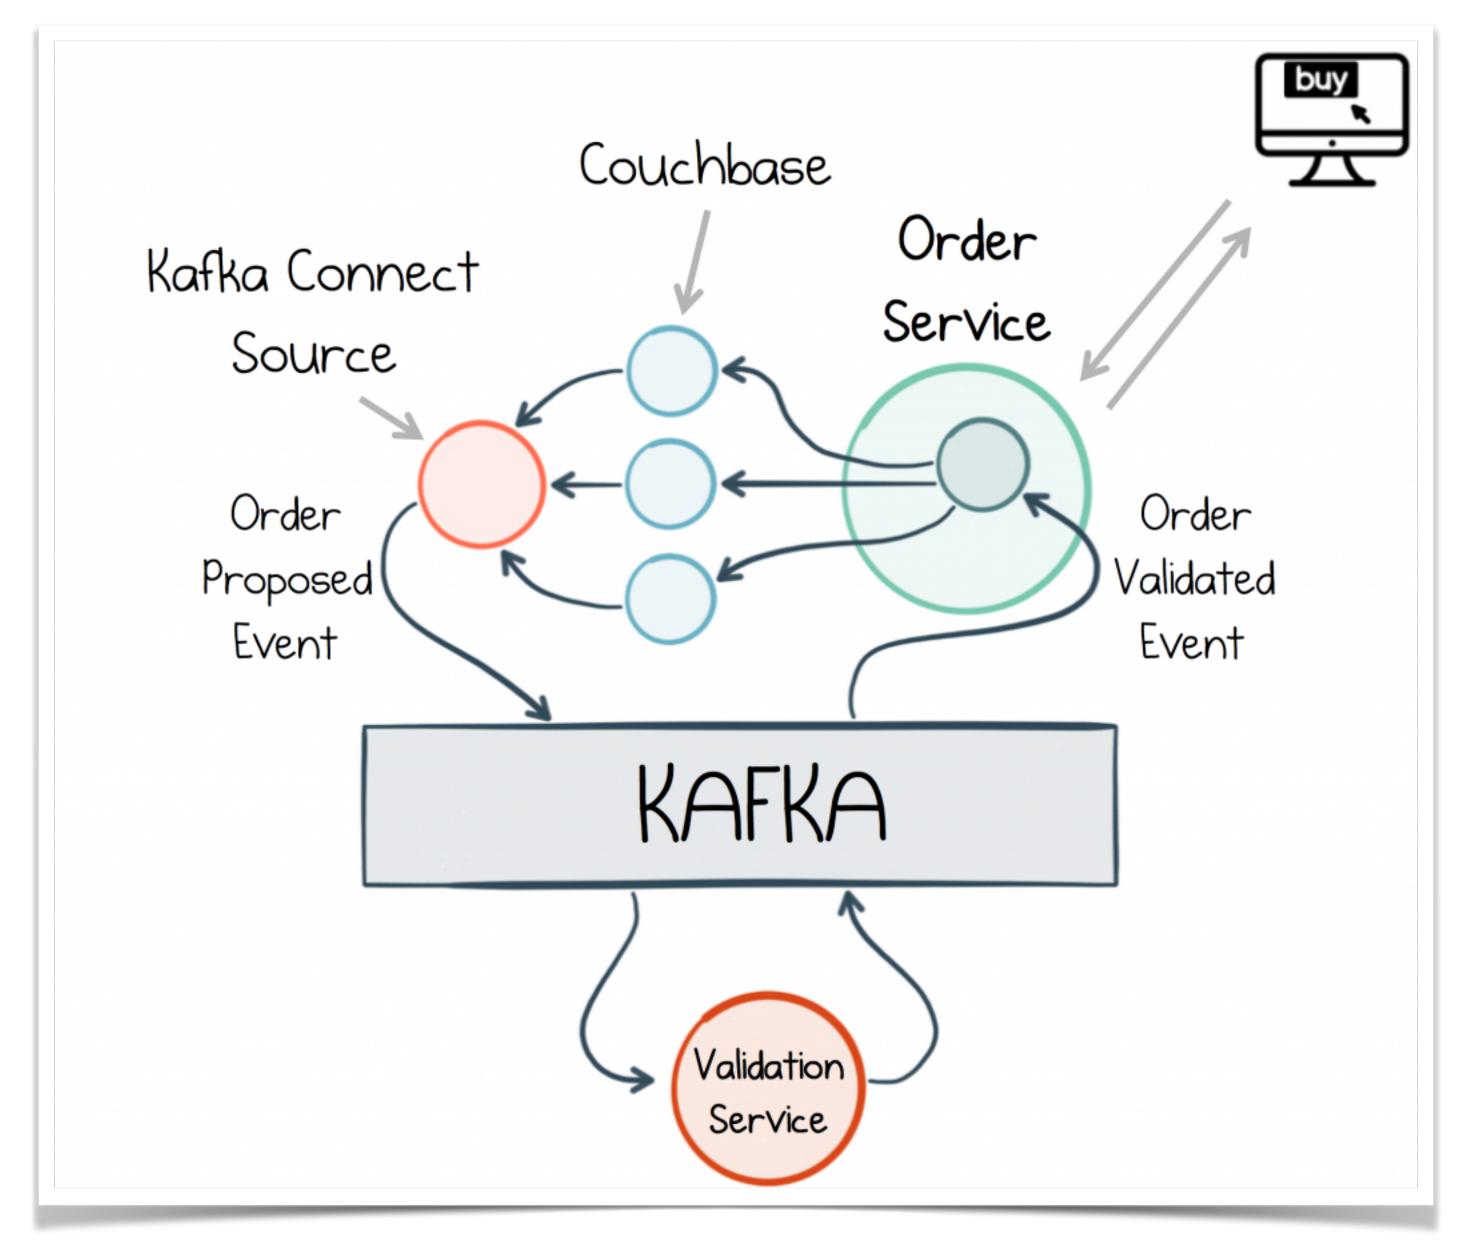


Just use distributed transactions!



Just manage transactions with a database and use single-writer





https://www.confluent.io/blog/messaging-single-source-truth/

Quick tip for finding friends in ops:

Ask them just to install and maintain production grade Kafka and Couchbase installations on AWS

Optimistic concurrency control

From Wikipedia, the free encyclopedia

Optimistic concurrency control (OCC) is a concurrency control method applied to transactional systems such as relational database management systems and software transactional memory. OCC assumes that multiple transactions can frequently complete without interfering with each other. While running, transactions use data resources without acquiring locks on those resources. Before committing, each transaction verifies that no other transaction has modified the data it has read. If the check reveals conflicting modifications, the committing transaction rolls back and can be restarted.^[1] Optimistic concurrency control was first proposed by H.T. Kung and John T. Robinson.^[2]

https://en.wikipedia.org/wiki/Optimistic_concurrency_control

Optimistic concurrency control

From Wikipedia, the free encyclopedia

OCC assumes that multiple transactions can frequently complete without interfering with each other.

transaction has modified the data it has read. If the check reveals conflicting modifications, the committing transaction rolls back and can be restarted.^[1] Optimistic concurrency control was first proposed by H.T. Kung and John T. Robinson.^[2]

https://en.wikipedia.org/wiki/Optimistic_concurrency_control

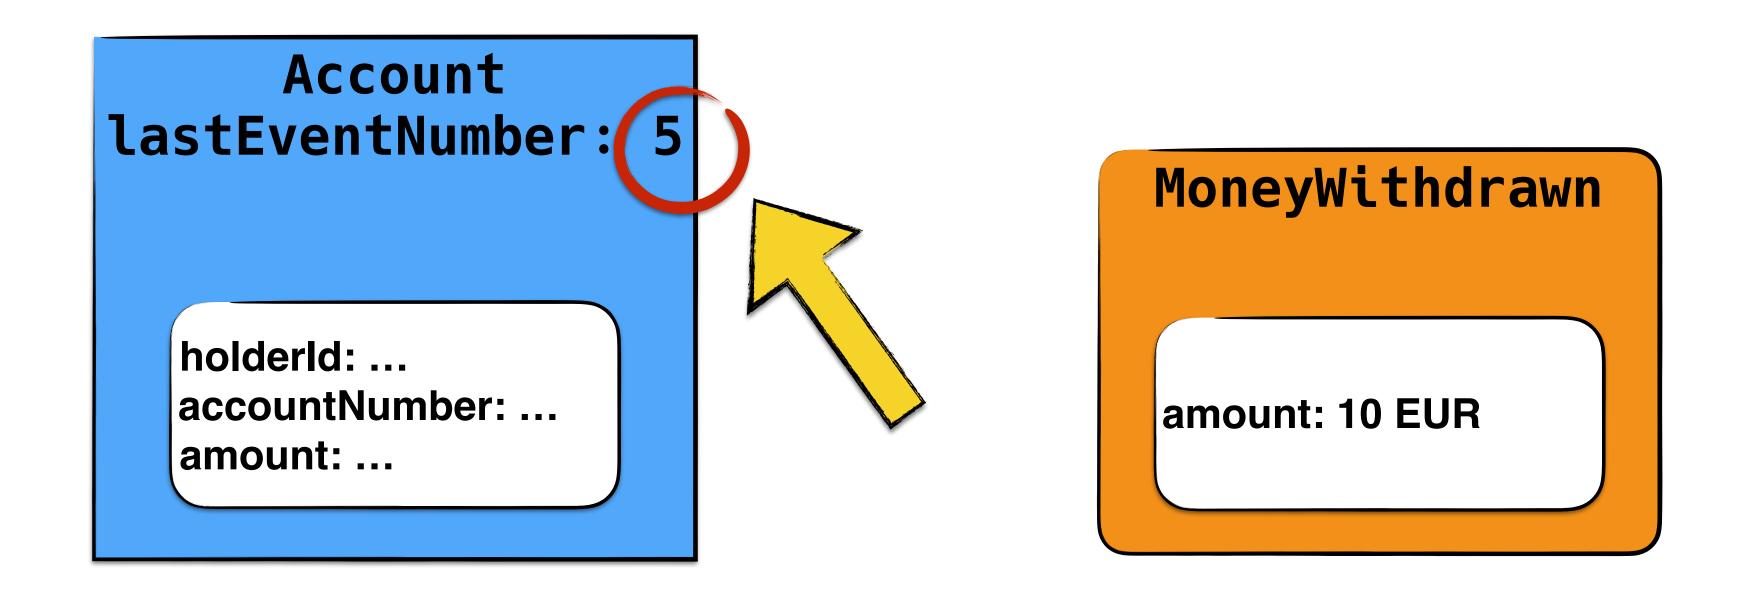
Optimistic concurrency control

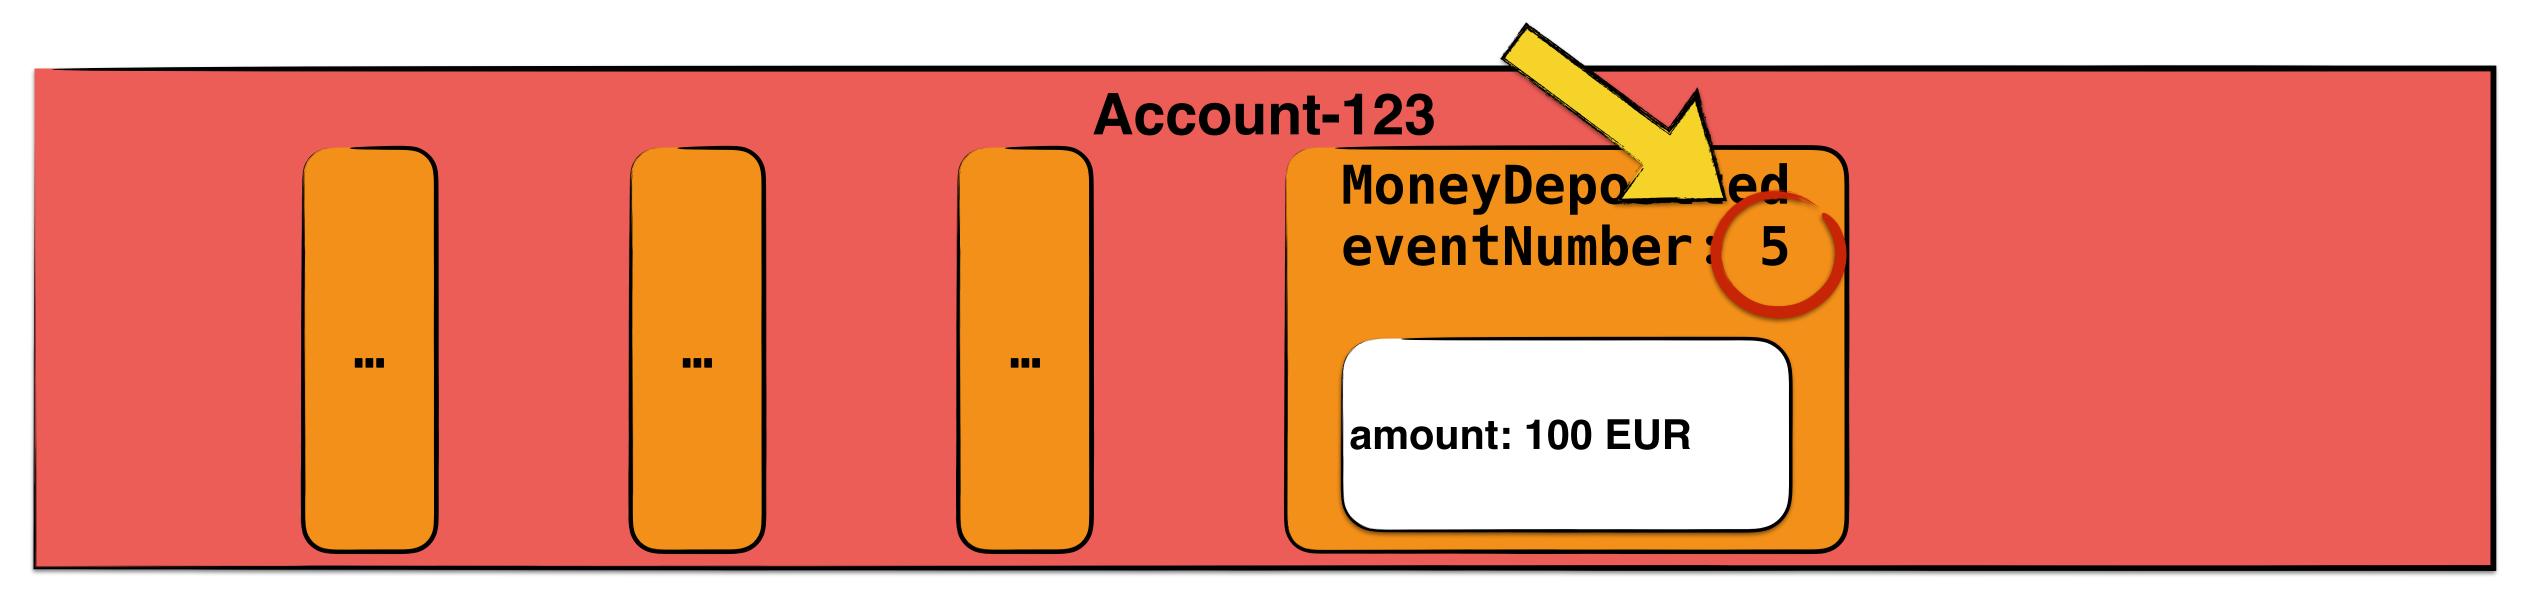
From Wikipedia, the free encyclopedia

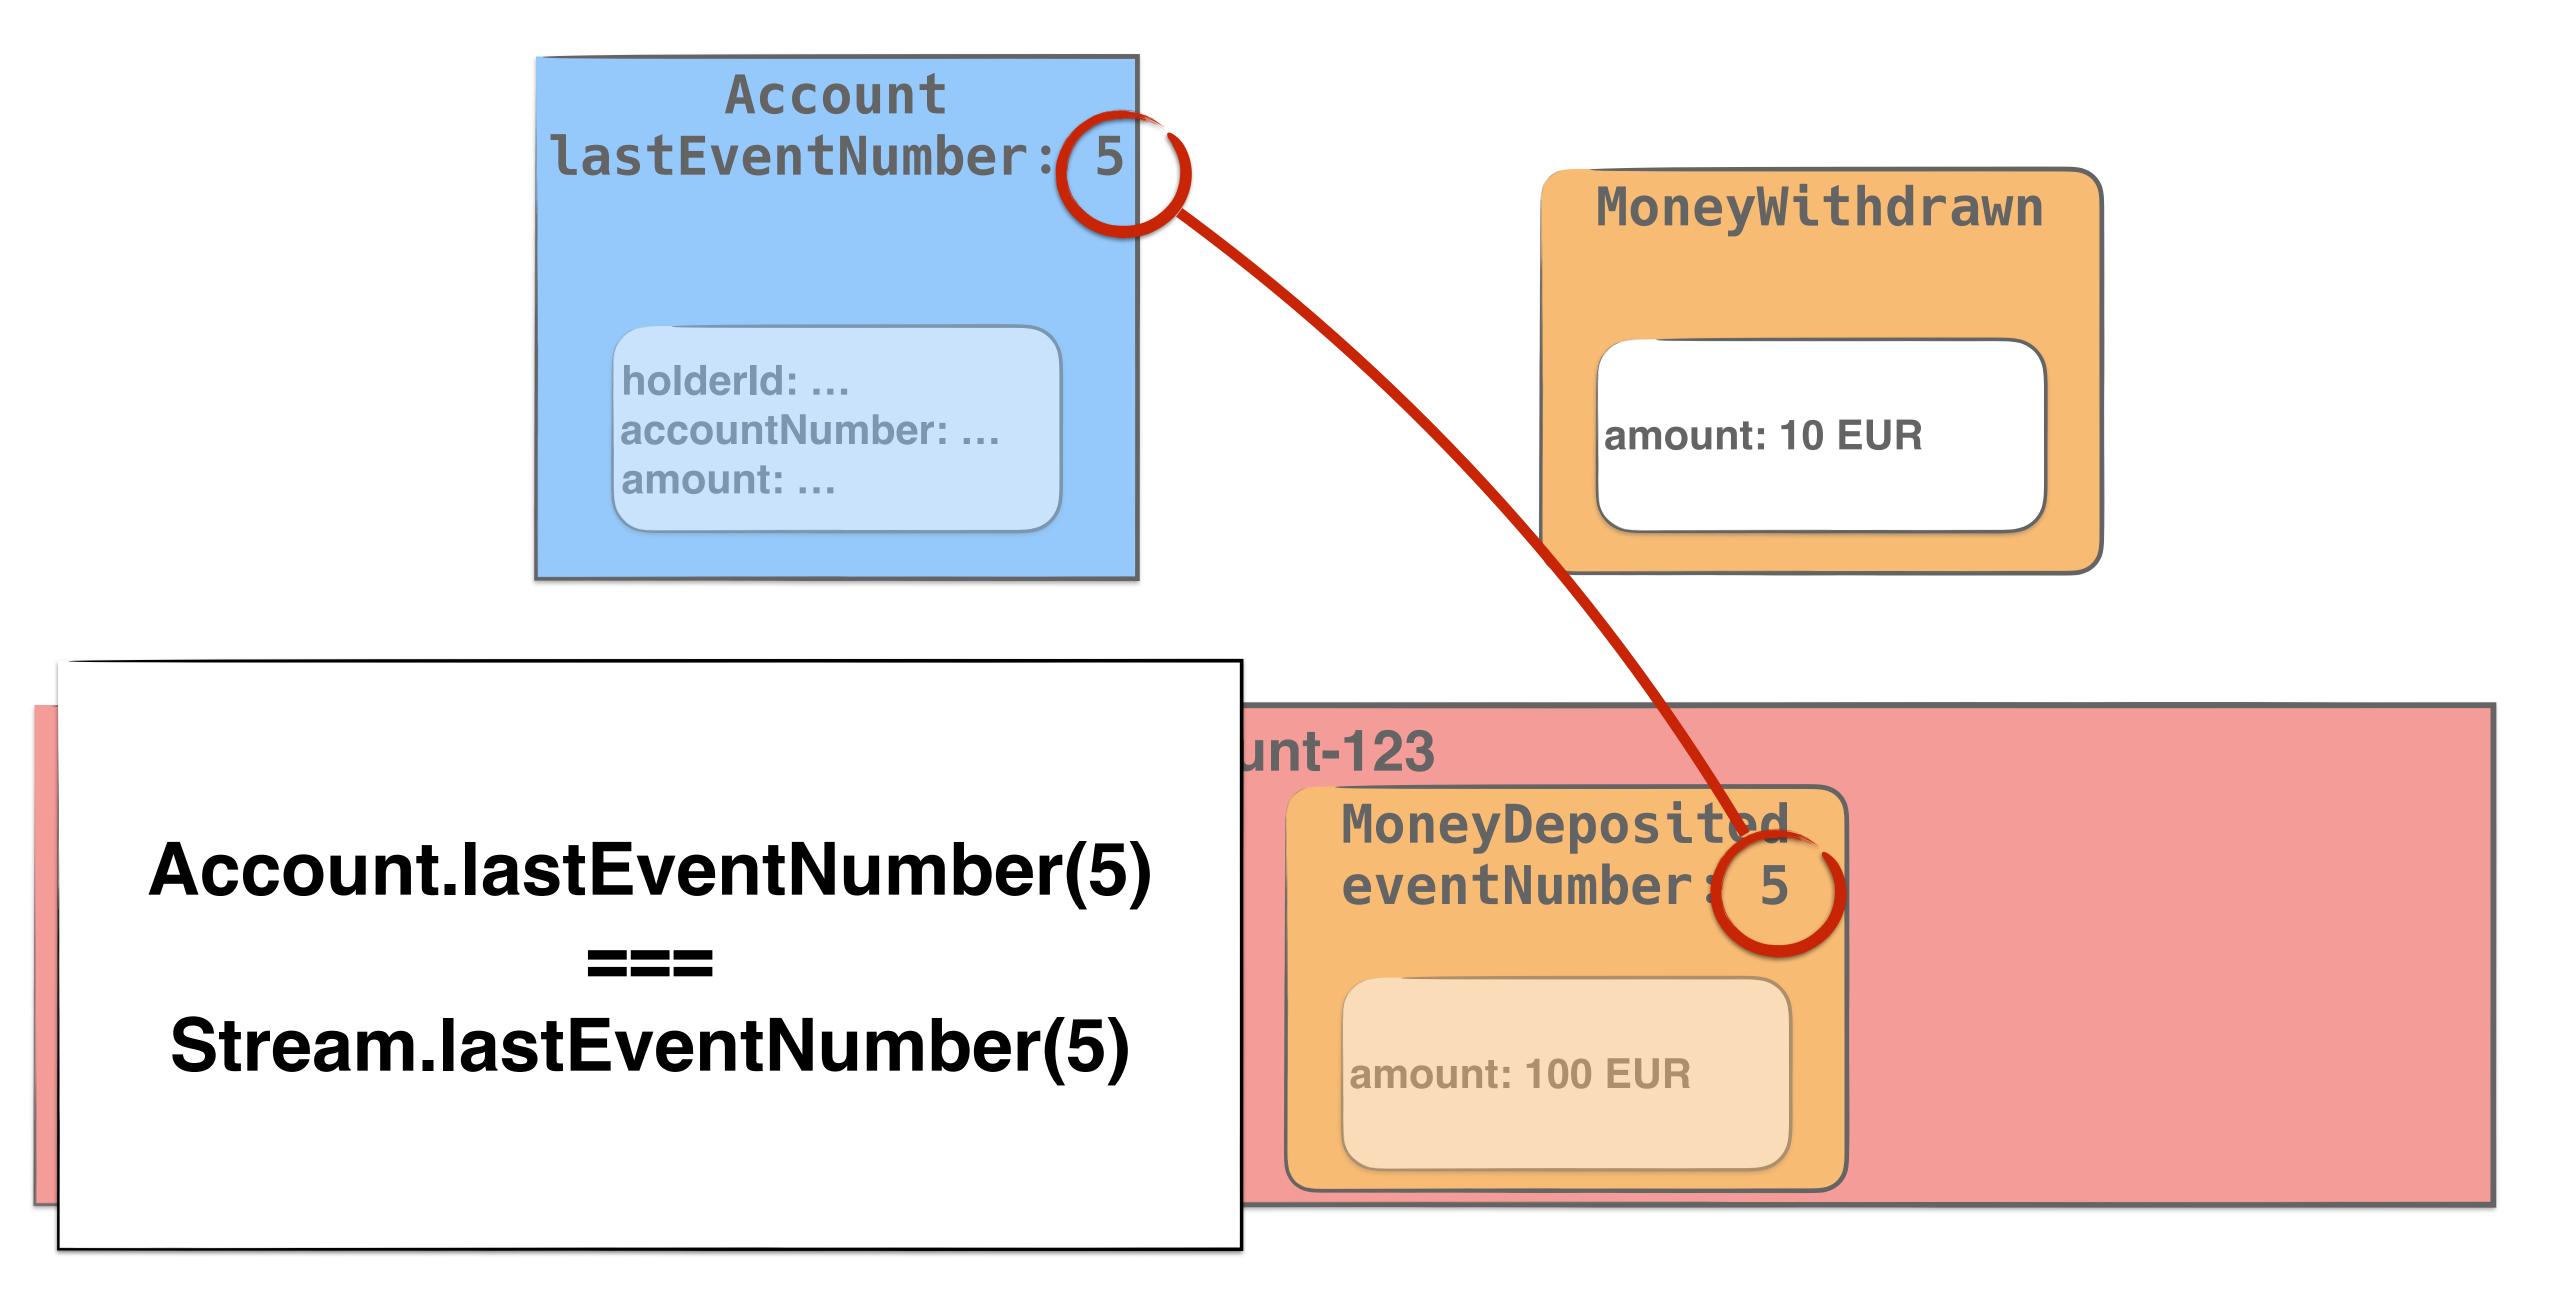
transaction has modified the data it has read. If the check reveals conflicting modifications, the committing transaction rolls back and can be restarted.^[1] Optimistic concurrency control was first proposed by H.T. Kung and John T. Robinson.^[2]

https://en.wikipedia.org/wiki/Optimistic_concurrency_control

The happy path





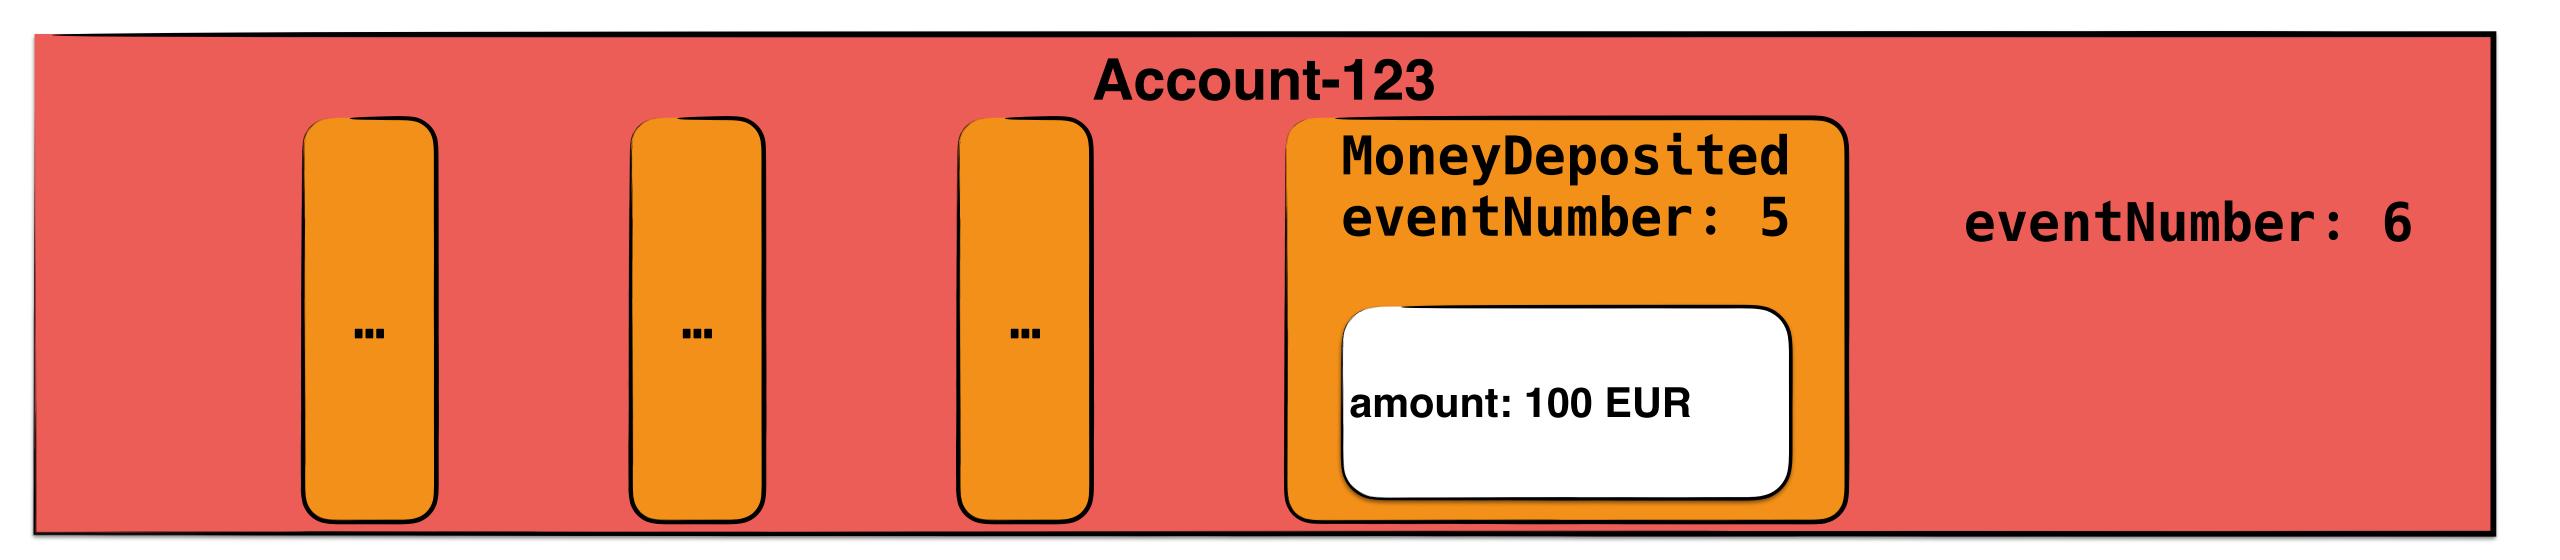


Account lastEventNumber: 5

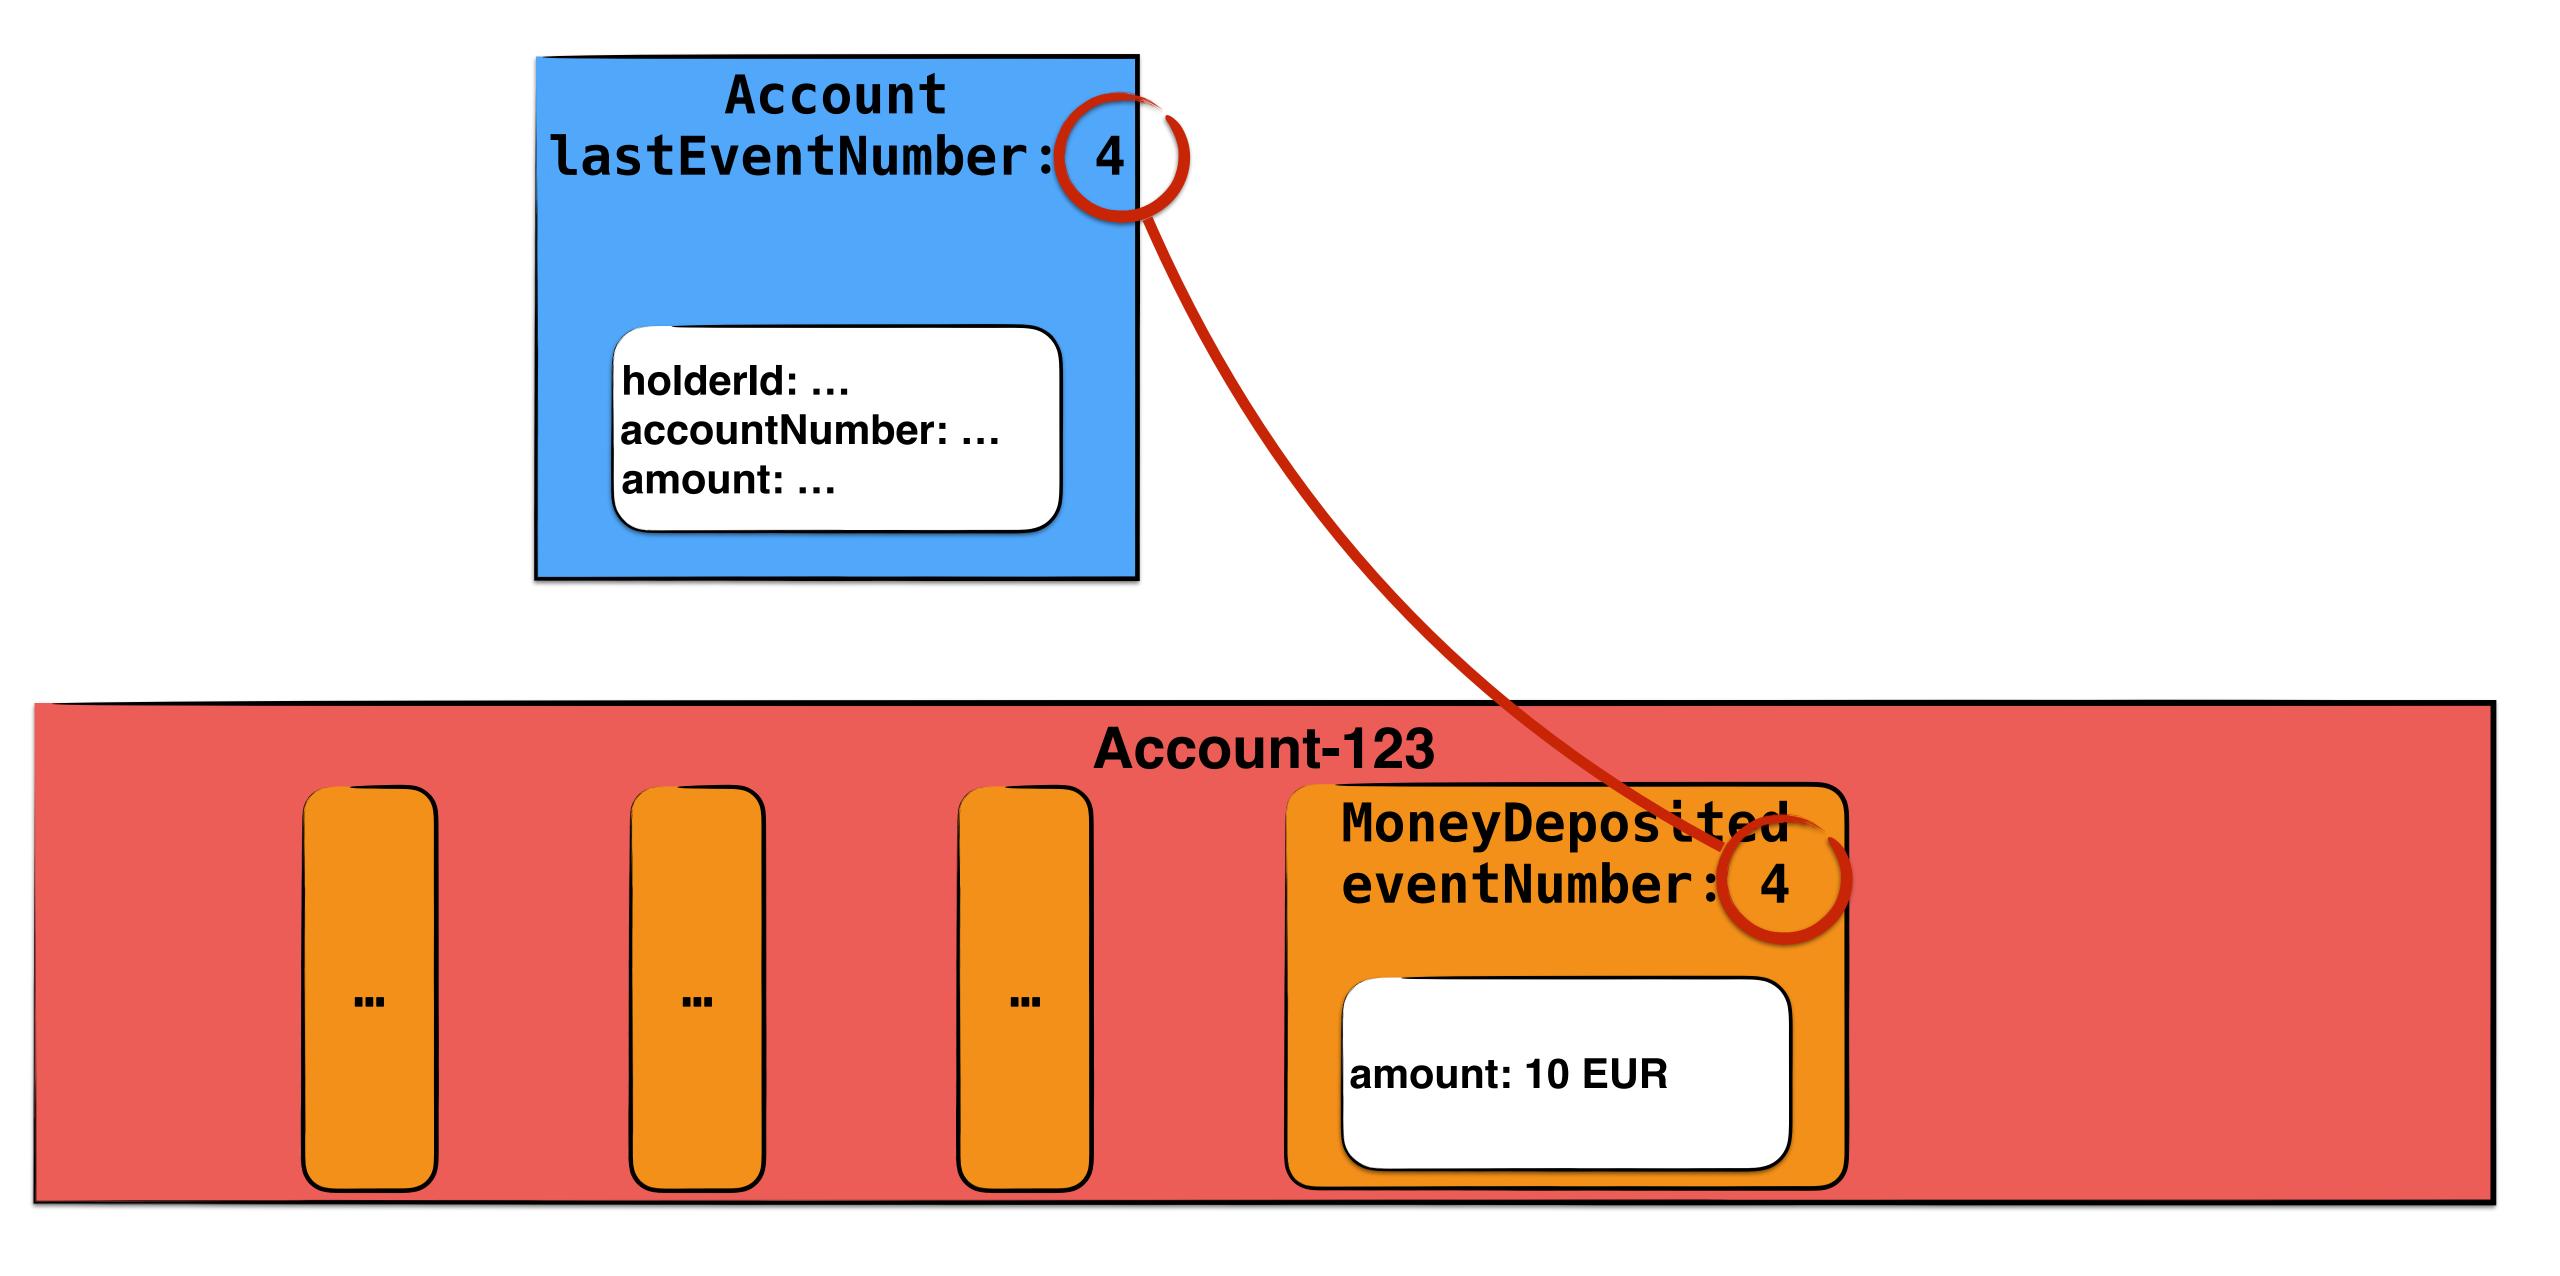
holderld: ... accountNumber: ... amount: ...

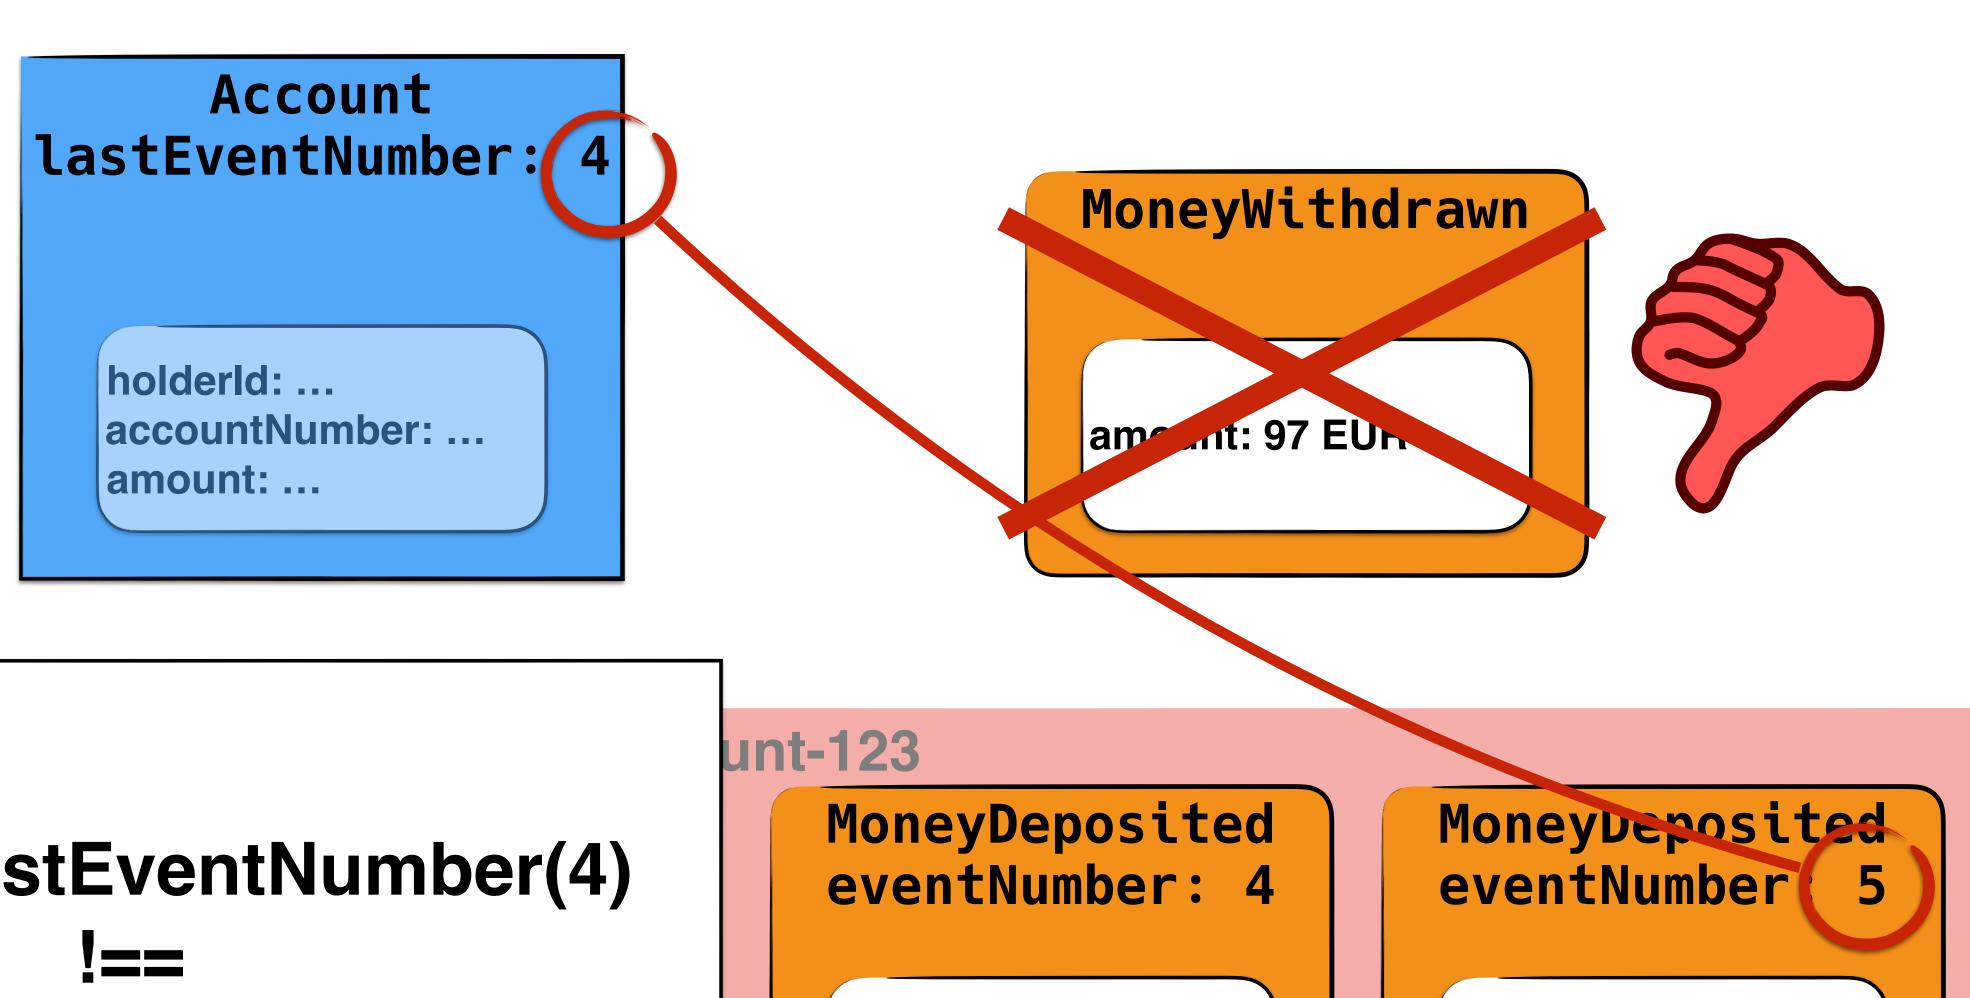
MoneyWithdrawn

amount: 10 EUR



The not-so-happy path

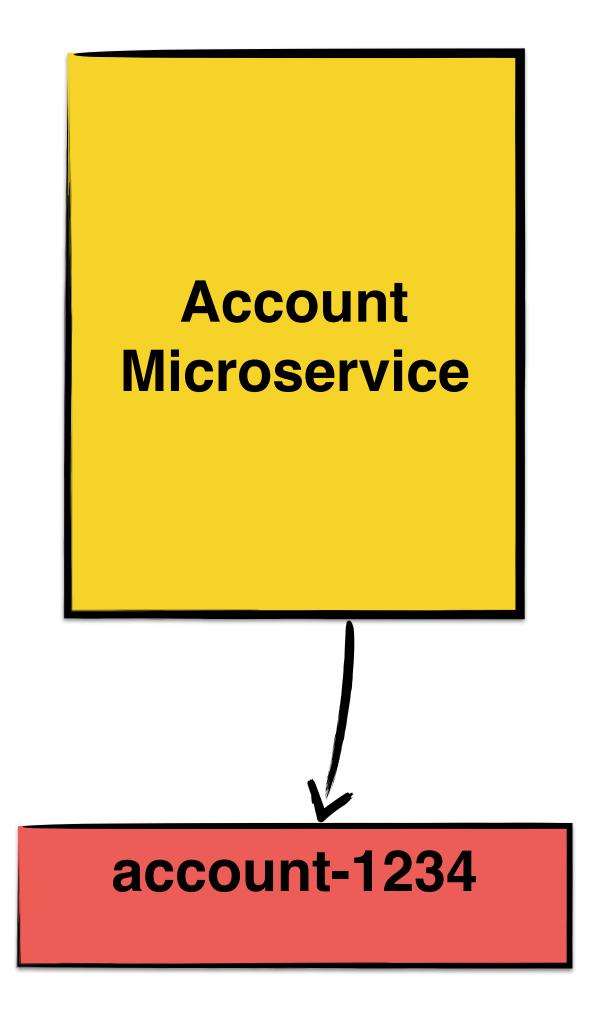




amount: 10 EUR

Account.lastEventNumber(4)
!==
Stream.lastEventNumber(5)

amount: 100 EUR



```
Account
Microservice
account-1234
```

```
{ aggregate, lastVersionNumber } = readAggregateFromStream(...)
events = executeBusinessLogic(...)
emitEvents('account', '1234', events, lastVersionNumber)
```

```
Account
                  { aggregate, lastVersionNumber } = readAggregateFromStream(...)
Microservice
                  events = executeBusinessLogic(...)
                  emitEvents('account', '1234', events, lastVersionNumber)
account-1234
```

```
Account
                  { aggregate, lastVersionNumber } = readAggregateFromStream(...)
Microservice
                  events = executeBusinessLogic(...)
                  emitEvents('account', '1234', events, lastVersionNumber)
account-1234
```

```
Account
                  { aggregate, lastVersionNumber } = readAggregateFromStream(...)
Microservice
                  events = executeBusinessLogic(...)
                  emitEvents('account', '1234', events, lastVersionNumber)
account-1234
```

```
Account
                  { aggregate, lastVersionNumber } = readAggregateFromStream(...)
Microservice
                                       essLogic(...)
                  events = executeBu
                                        '1234', evrats, lastVersionNumber
                  emitEvents('accoun
                              Optimistic
                          concurrency control
account-1234
```

And Kafka?

Allow specifying expected offset on produce

Details

Priority:

Type:

Improvement

★ Minor Resolution:

Affects Version/s: None

Component/s: producer

Labels: None

Description

I'd like to propose a change that adds a simple CAS-like mechanism to the Kafka producer. This update has a small footprint, but enables a bunch of interesting uses in stream processing or as a commit log for process state.

Status:

Fix Version/s:

OPEN

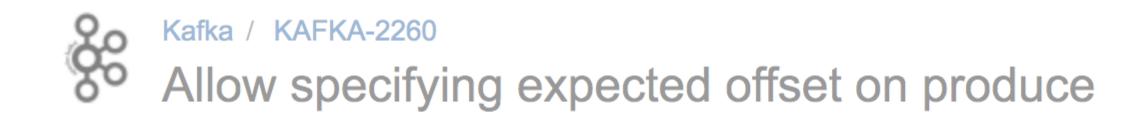
None

Unresolved

- Andy Bryant added a comment 27/Jul/18 04:14
 - Would prove very handy in event source based designs
- → Russell Ferriday added a comment 9 hours ago

This would enable full-on eventsourcing on Kafka, without having to restrict to single-thread designs. One example of a great (>250 github star) FOSS project being held back by this:

https://github.com/johnbywater/eventsourcing/issues/108



Details

Improvement Type:

Priority:

Affects Version/s: Fix Version/s: None None



added a comment - 27/Jul/18 04:14

OPEN

Unresolved



Would prove very handy in event source based designs

Status:

Resolution:

- Andy Bryant added a comment 27/Jul/18 04:14
 - Would prove very handy in event source based designs
- Russell Ferriday added a comment 9 hours ago

This would enable full-on eventsourcing on Kafka, without having to restrict to single-thread designs. One example of a great (>250 github star) FOSS project being held back by this:

https://github.com/johnbywater/eventsourcing/issues/108

Details			
Type:		Status:	OPEN
Priority:		Resolution:	Unresolved
Affects Version/s:	None	Fix Version/s:	None
Component/s:	producer		

This would enable full-on eventsourcing on Kafka, without having to restrict to single-thread designs. One example of a great (>250 github star) FOSS project being held back by this:

✓ ■ Andy Bryant added a comment - 27/Jul/18 04:14

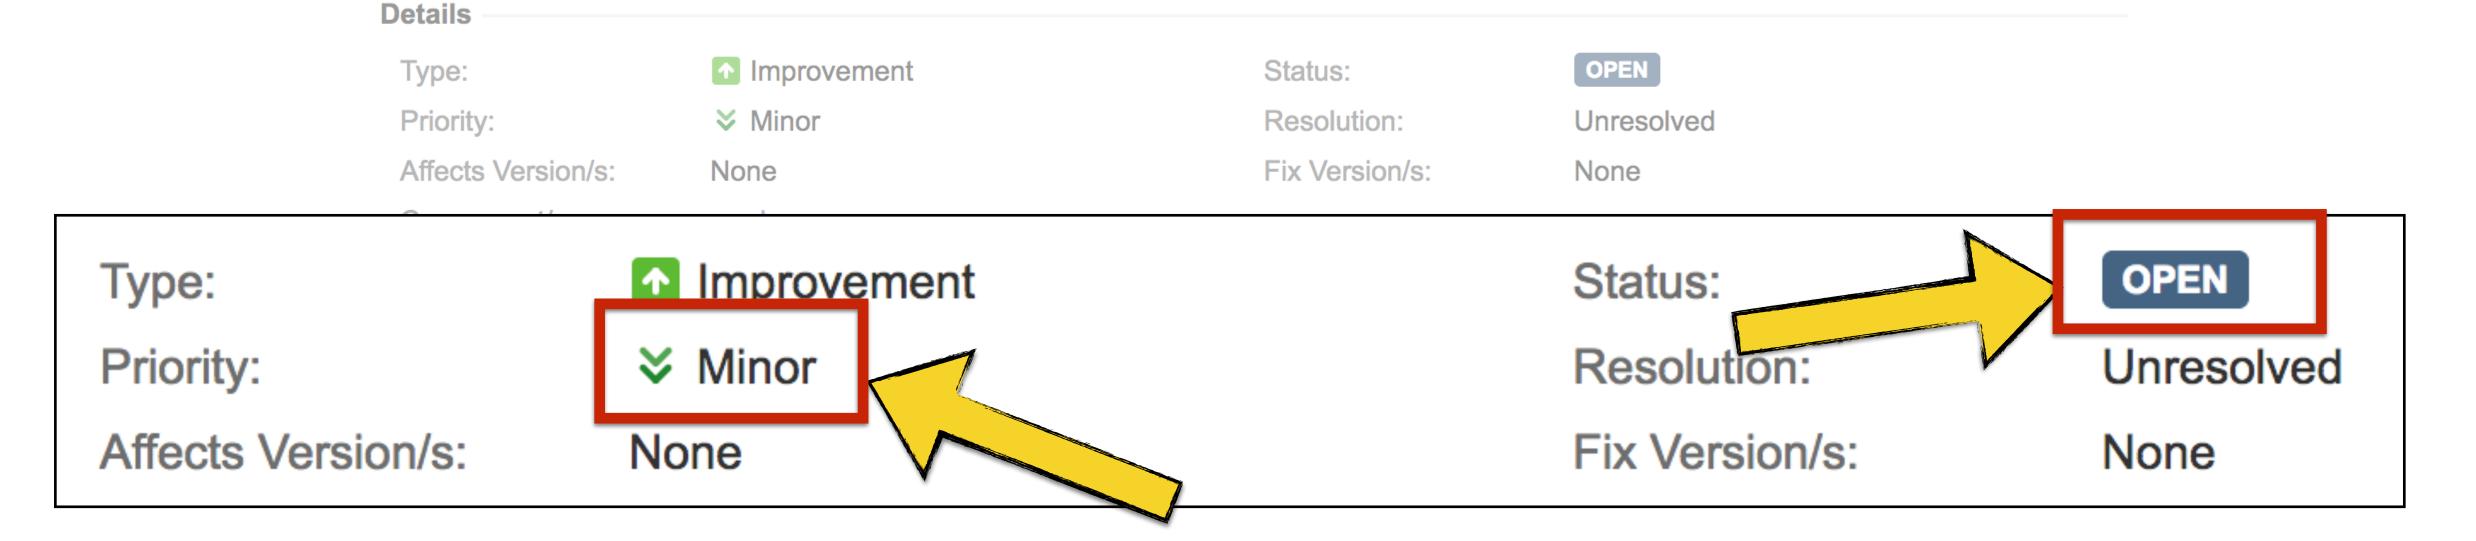
♣ Would prove very handy in event source based designs

Russell Ferriday added a comment - 9 hours ago

This would enable full-on eventsourcing on Kafka, without having to restrict to single-thread designs. One example of a great (>250 github star) FOSS project being held back by this:

https://github.com/johnbywater/eventsourcing/issues/108





- Andy Bryant added a comment 27/Jul/18 04:14
 - Would prove very handy in event source based designs
- Russell Ferriday added a comment 9 hours ago

This would enable full-on eventsourcing on Kafka, without having to restrict to single-thread designs. One example of a great (>250 github star) FOSS project being held back by this:

https://github.com/johnbywater/eventsourcing/issues/108

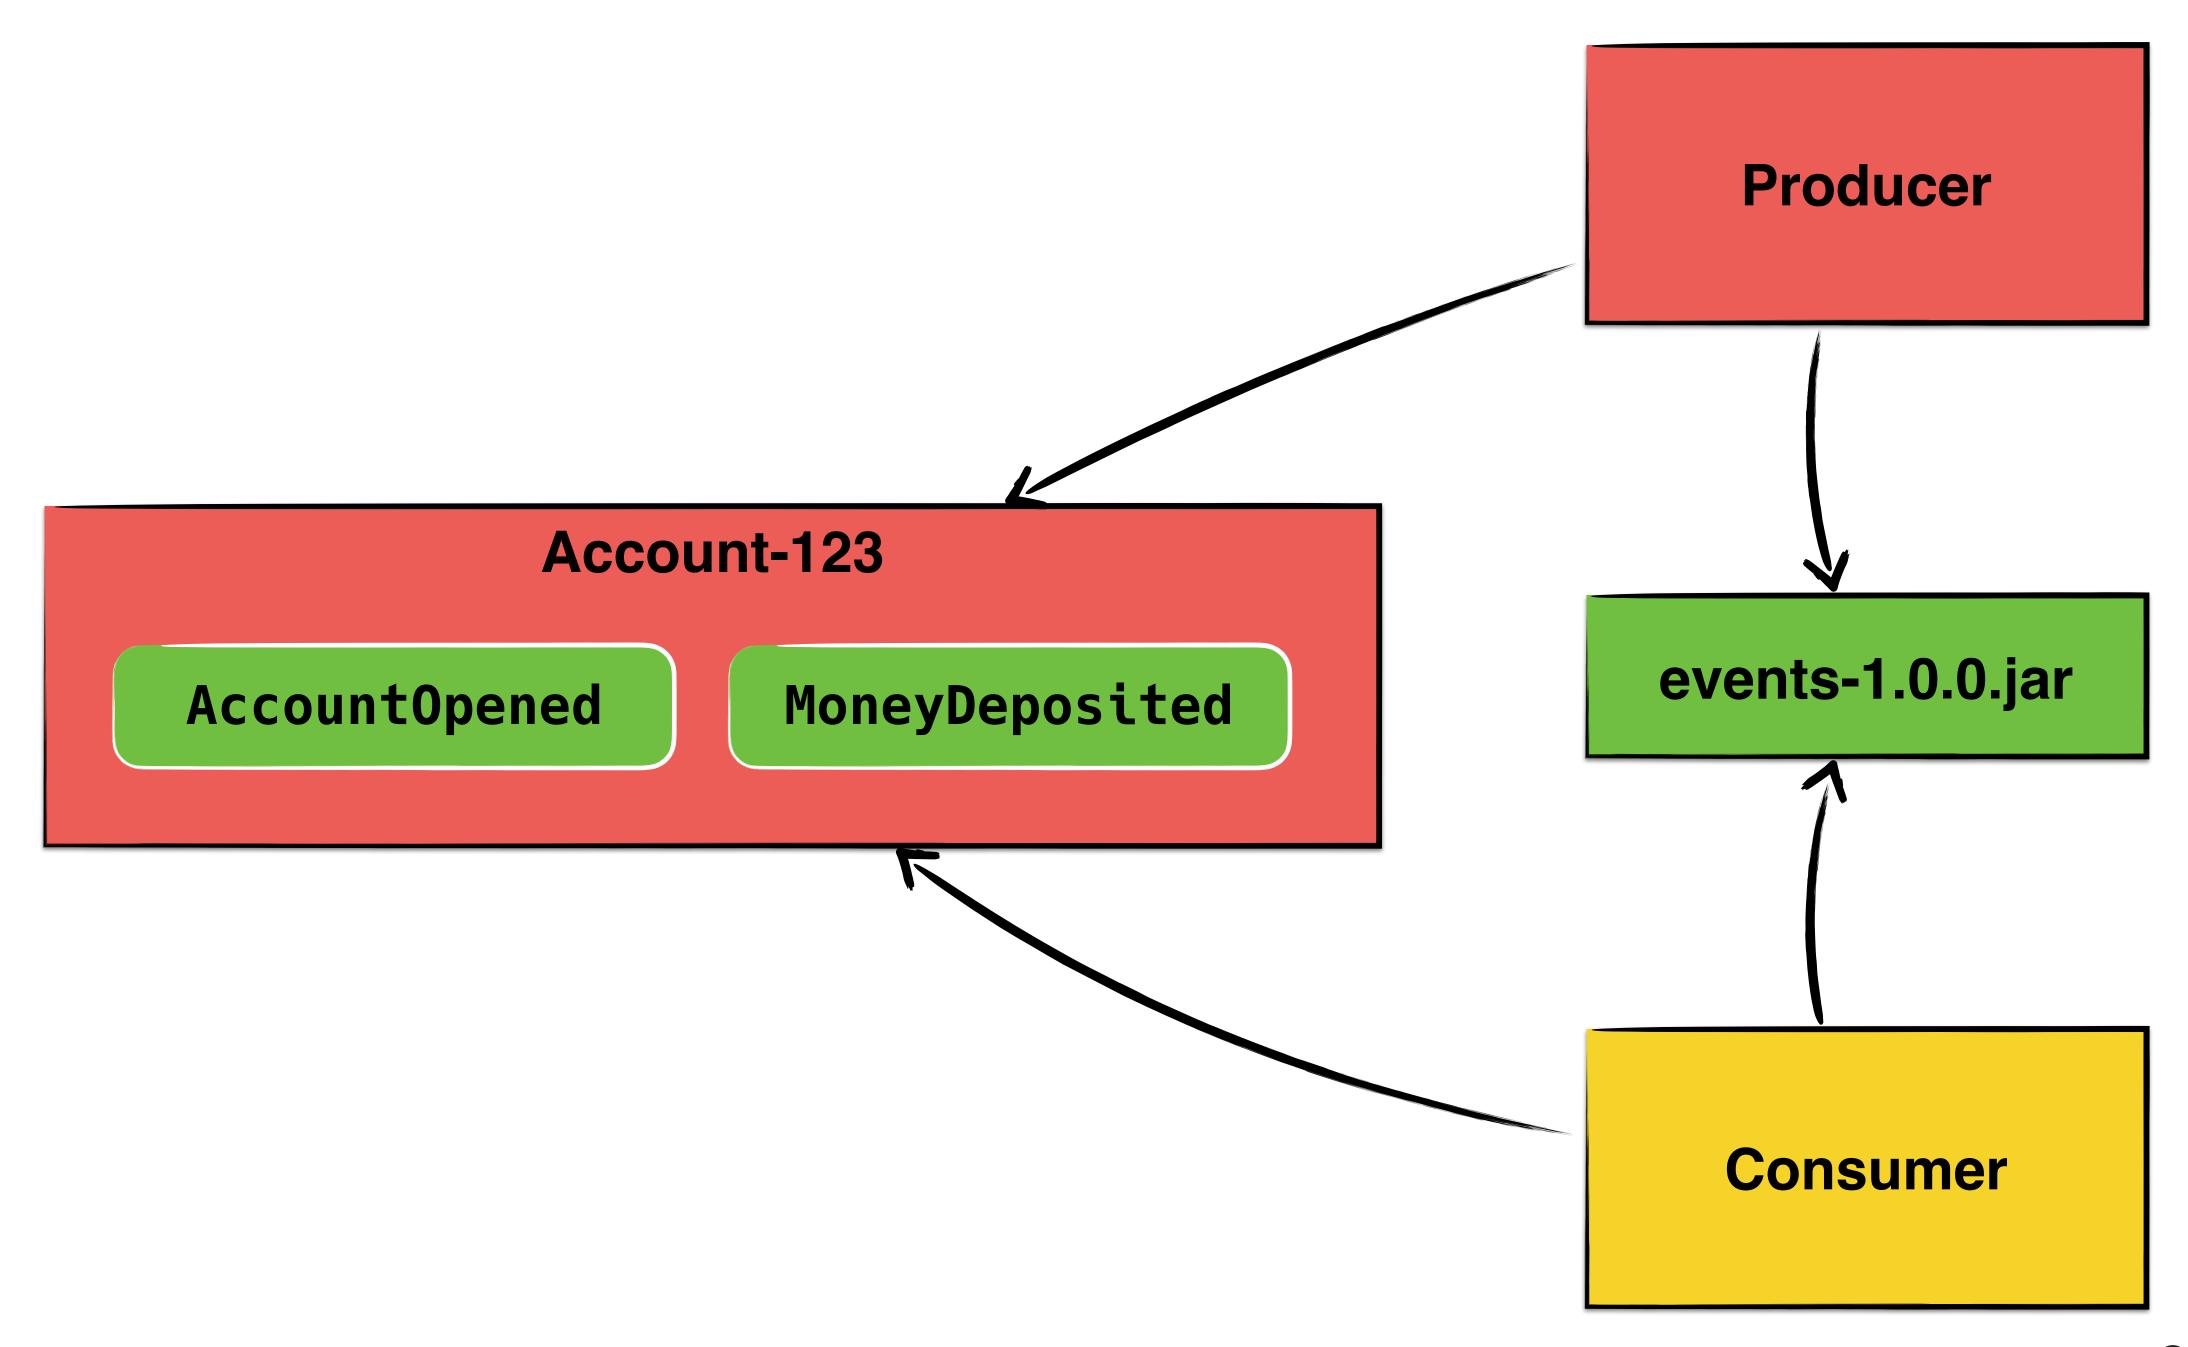
Scalability without locks Consistency Design choice Super-simple programming logic



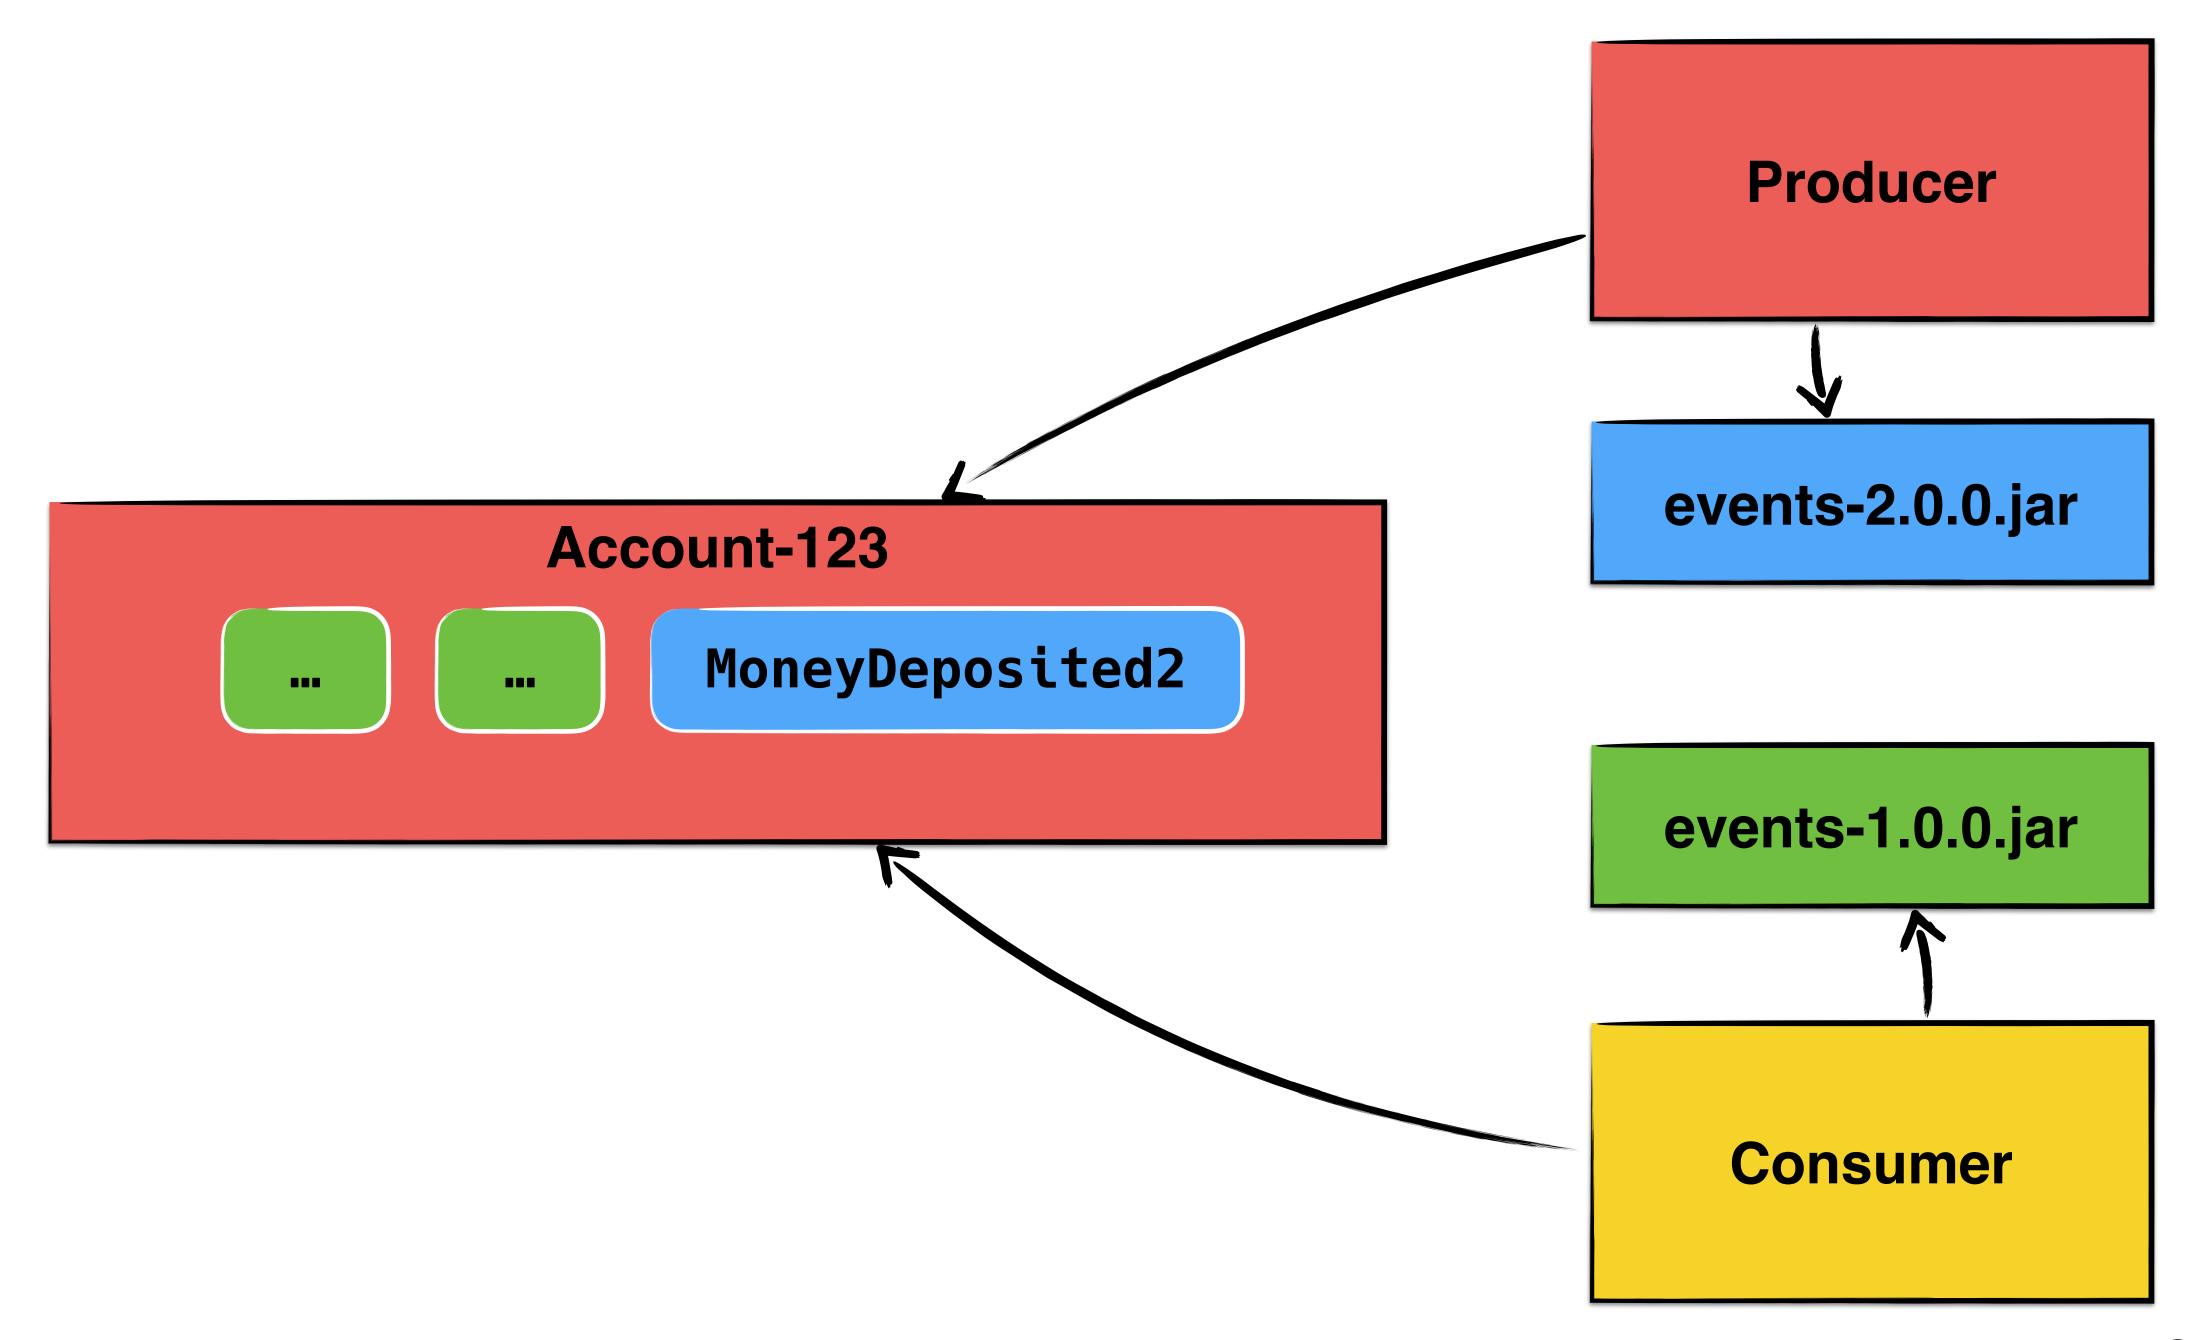
How can we deal with versions without going crazy?

Just use semantic versioning and typed events





...then change some event types

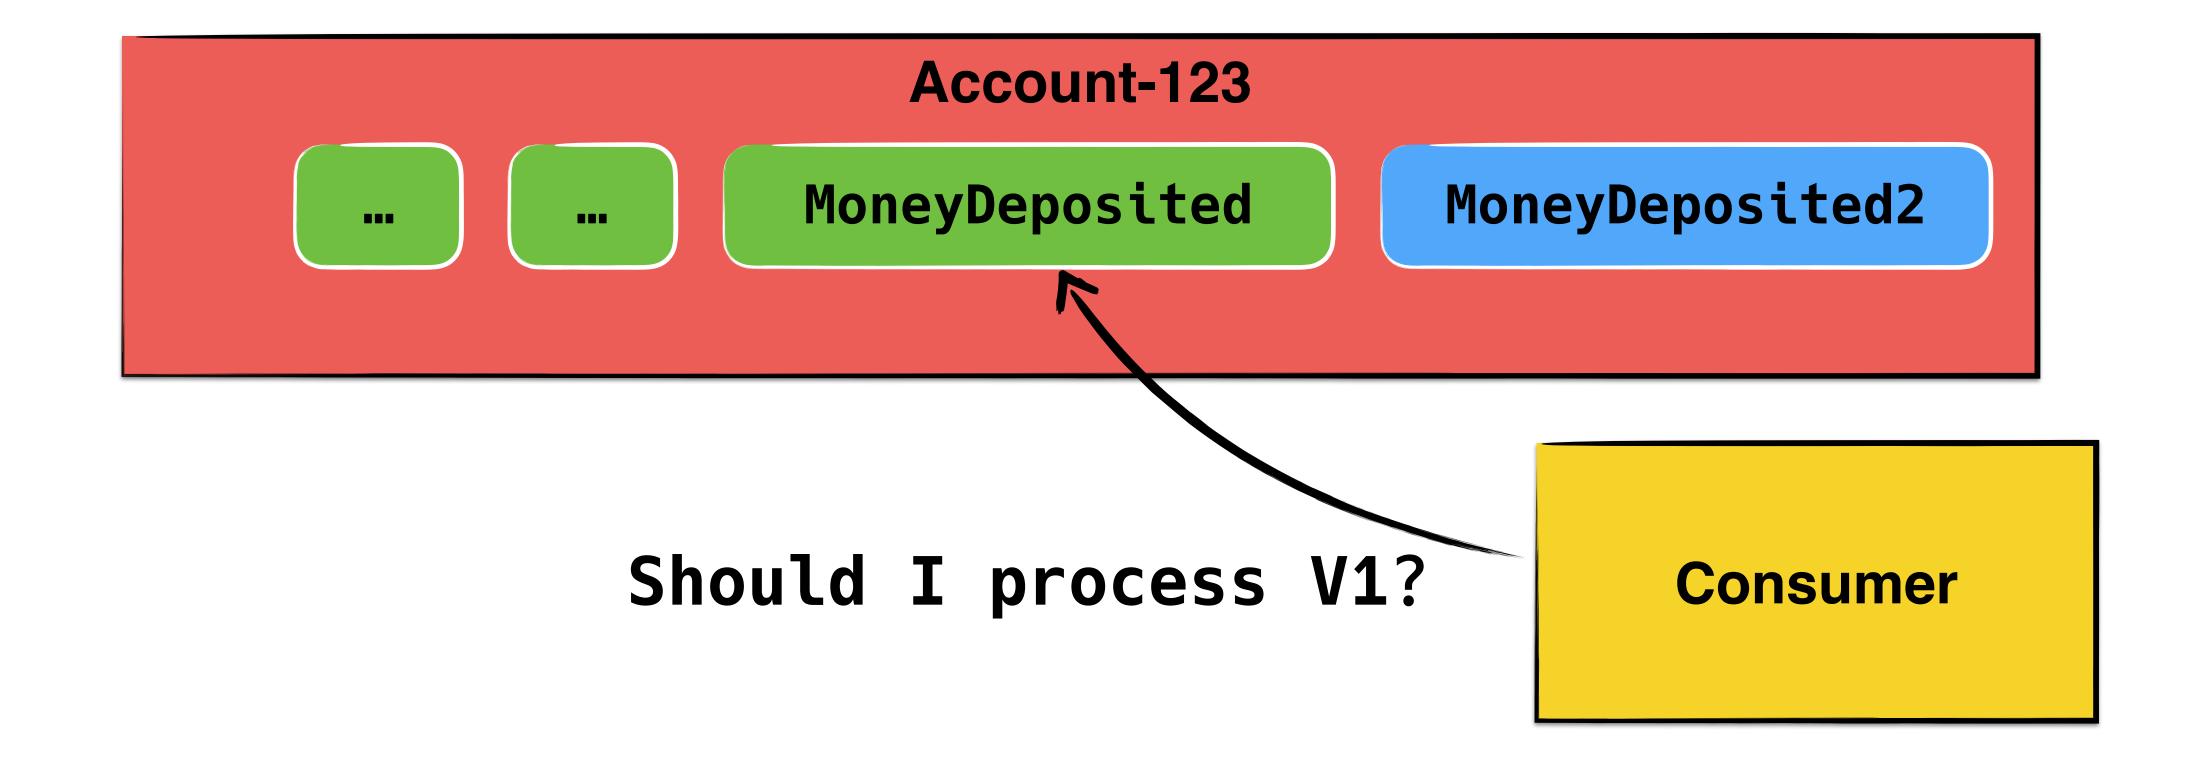


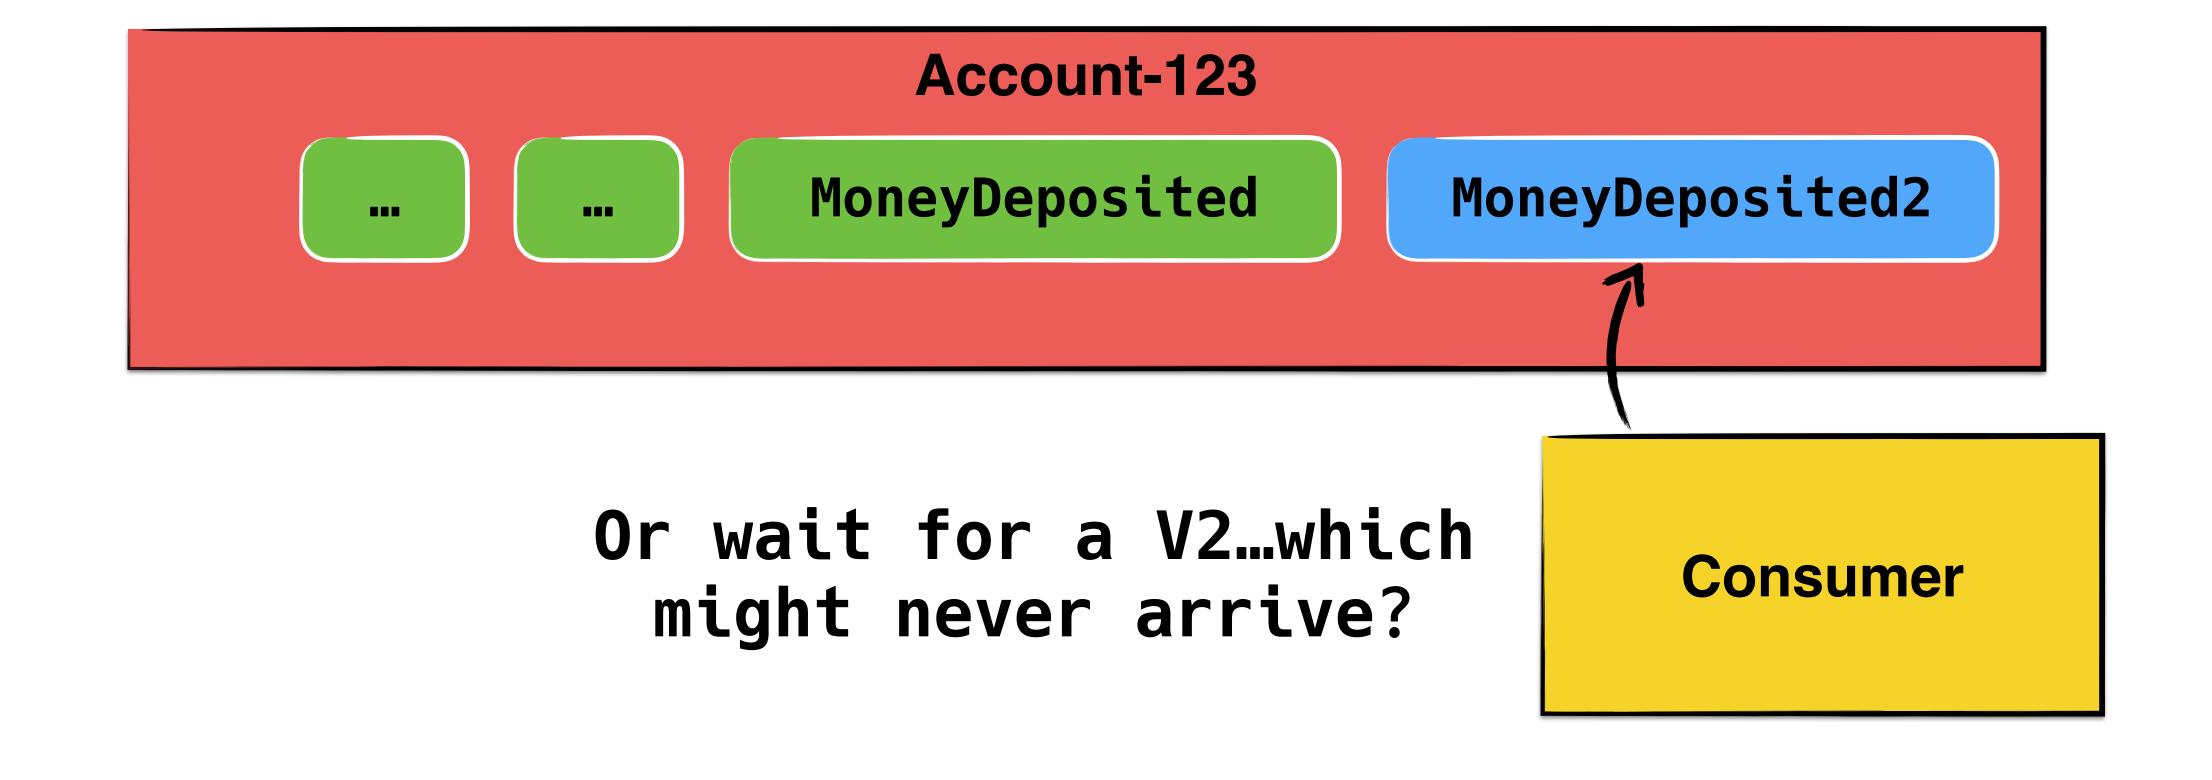
```
public class InvalidClassException extends ObjectStreamException {
    ...
}
```

Gosh...just apply "Double Write"









MoneyDeposited_v1 MoneyDeposited_v2

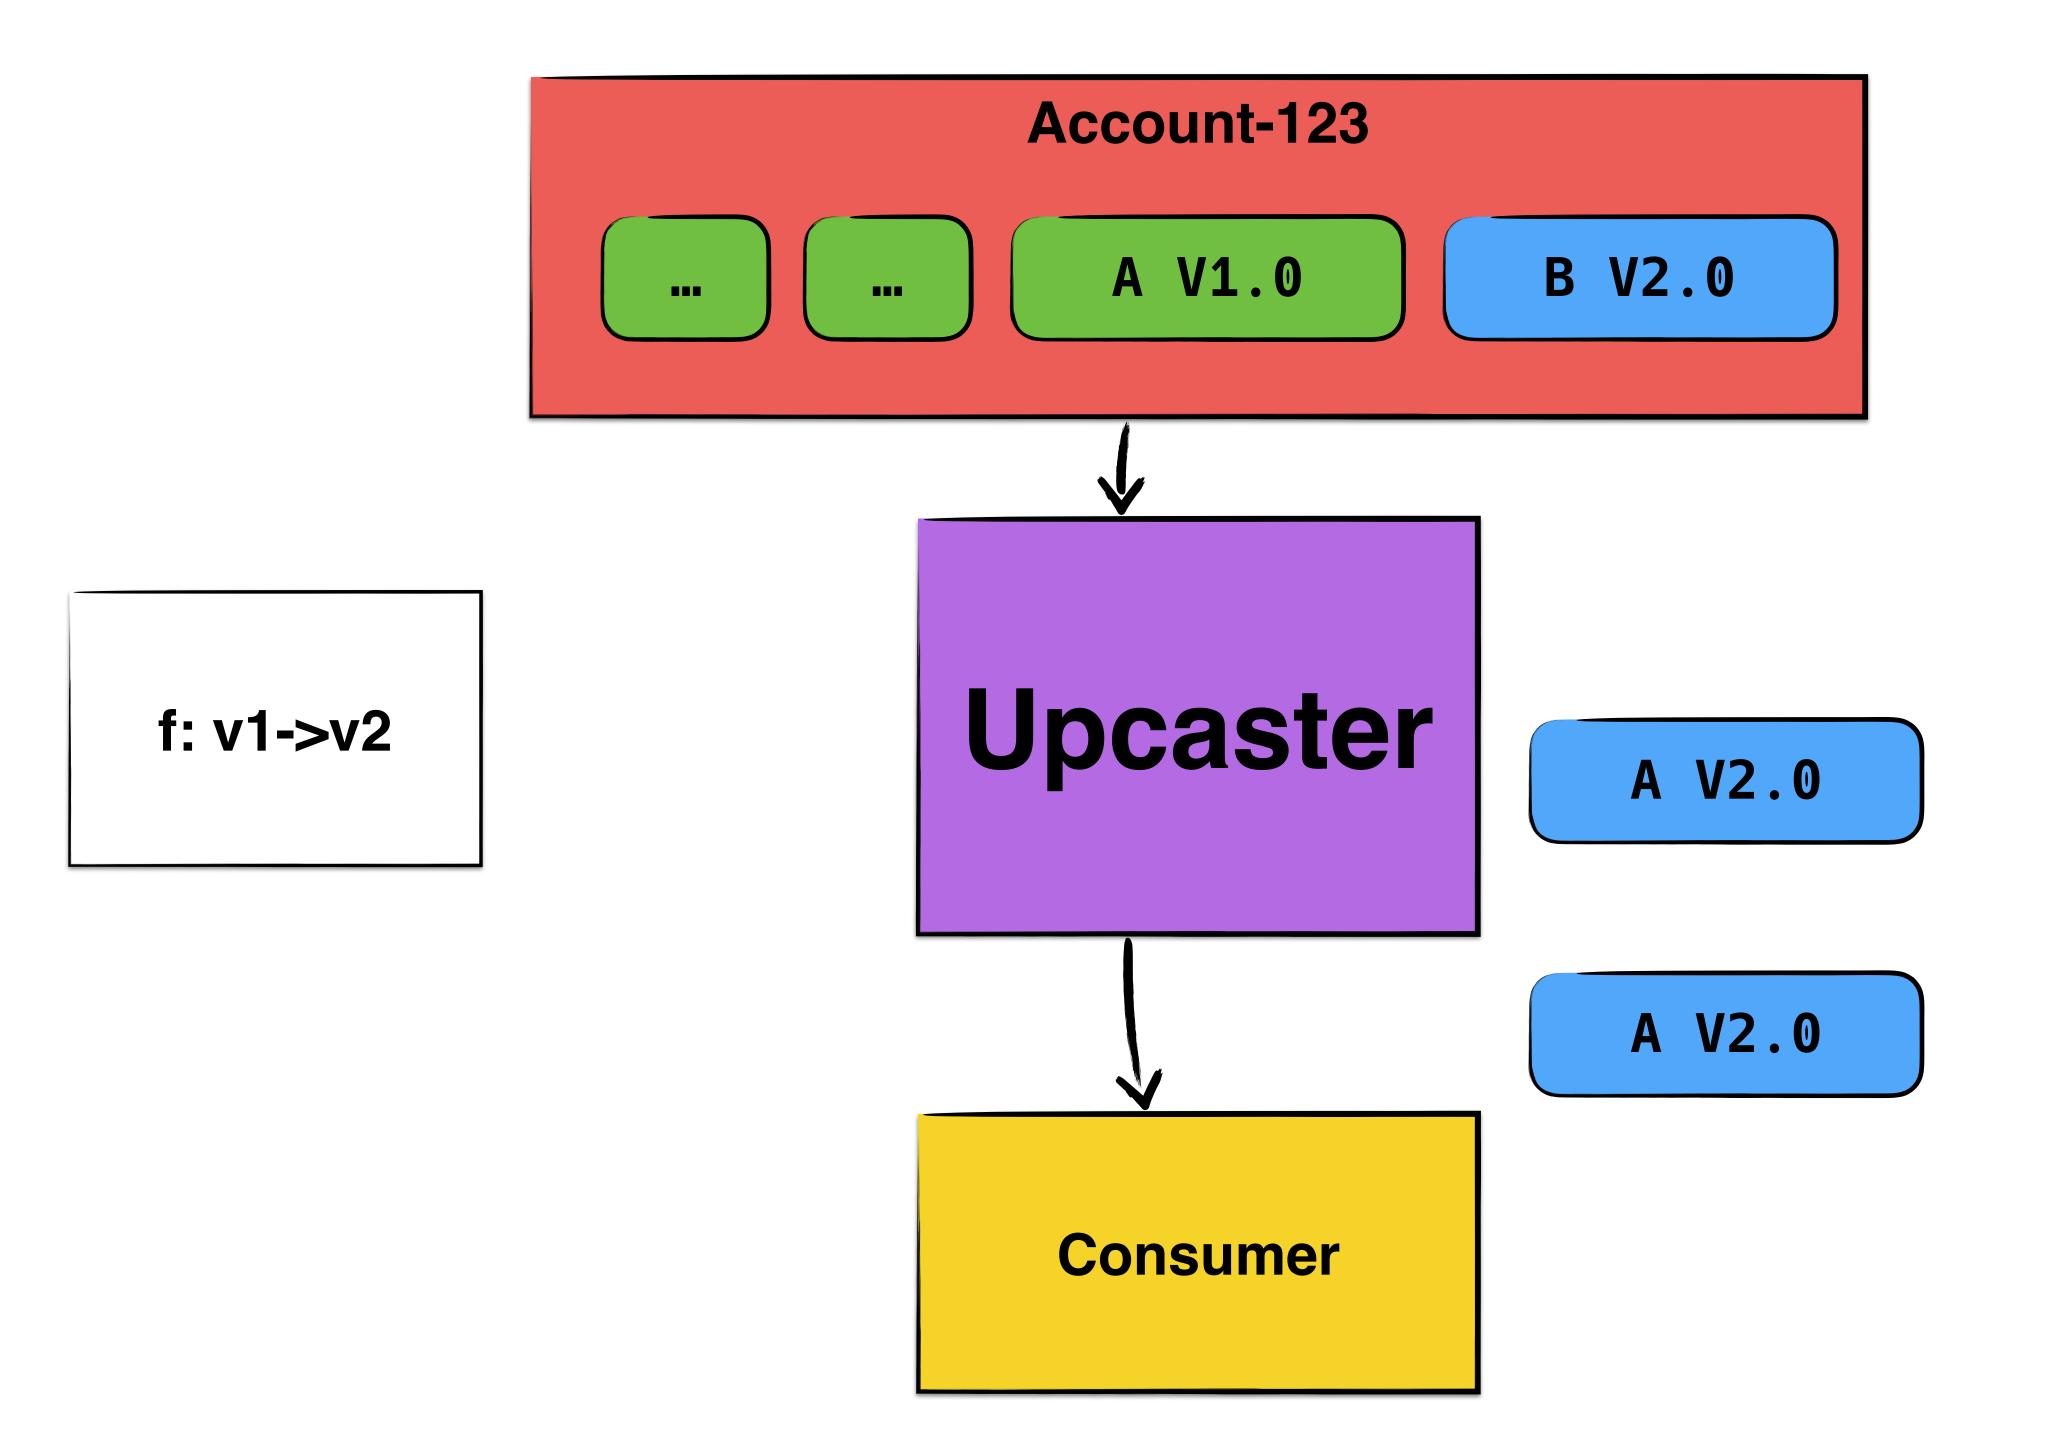
MoneyDeposited_v100

MoneyDeposited_v1Handler MoneyDeposited_v2Handler

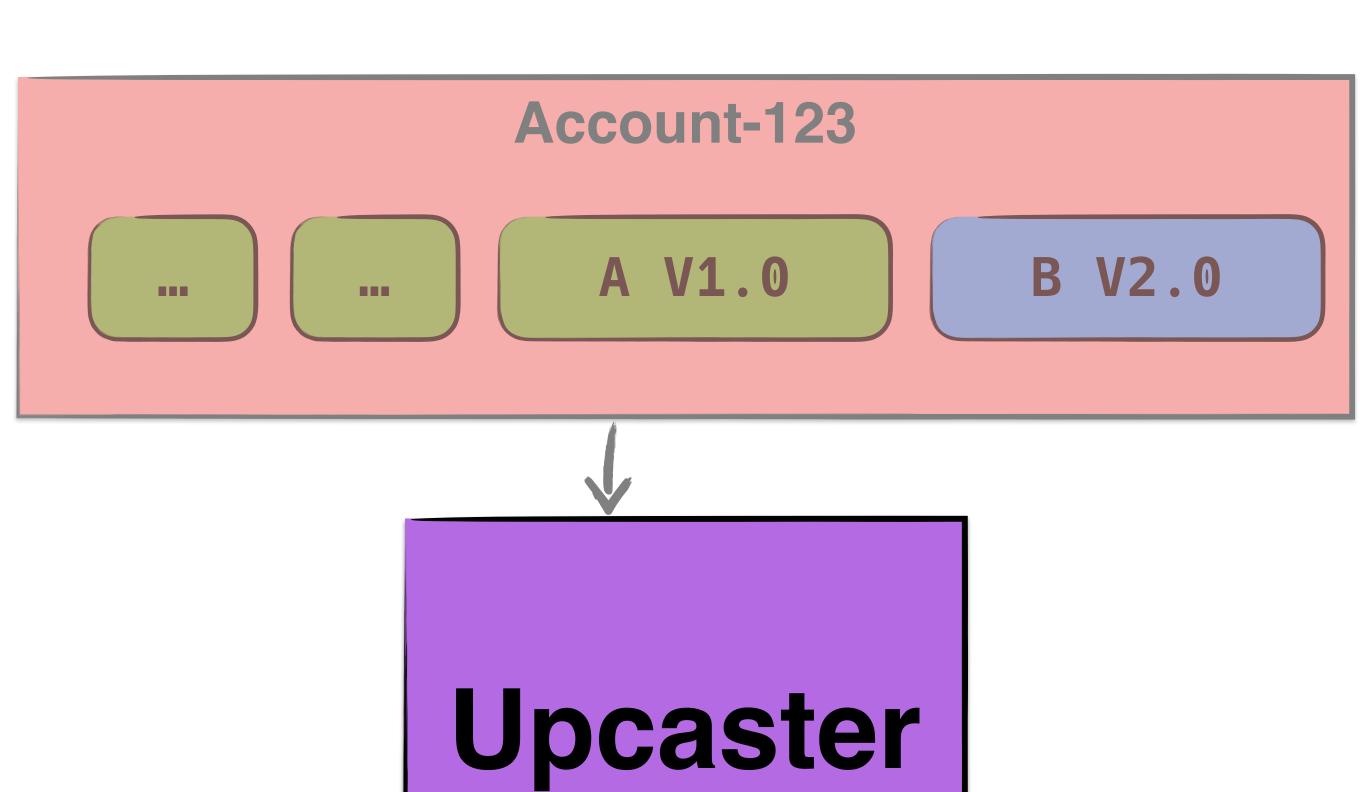
MoneyDeposited_v100Handler



Just use an Upcaster



5 months later...



f: v1->v2

f: v2->v3

f: v3->v4

f: v97->v98

f: v98->v99

f: v99->v100

Consumer

Good luck maintaining that monster

Incompatible event changes indicate a new event

Prefer simple, text-based, human readable events

Fancy speak for JSON

```
"eventType": "MoneyTransferred",
"aggregateId": "1234",
"iban": "DE12",
"accountNumber": "12312312",
"amount": 10,
"currency": "EUR"
```

And correctness?

String-ly typed events work really well

Weak schema to the rescue

```
"$schema": "http://json-schema.org/draft-07/schema",
"title": "UserCreated",
"description": "Creates a user",
"type": "object",
"properties": {
  "userId": {
    "description": "The new user's ID",
    "type": "string",
    "format": "uuid"
"additionalProperties": false,
"required": [
  "userId"
```

```
"$schema": "http://json-schema.org/draft-07/schema",
"title": "UserCreated",
"description": "Creates a user",
"type": "object",
"properties": {
 "userId": {
    "description": "The new user's ID",
    "type": "string",
    "format": "uuid"
"additionalProperties": false,
"required": [
 "userId"
```

```
"$schema": "http://json-schema.org/draft-07/schema",
"title": "UserCreated",
"description": "Creates a user",
"type": "object",
"properties": {
  "userId": {
    "description": "The new user's ID",
    "type": "string",
    "format": "uuid"
"additionalProperties": false,
"required": [
 "userId"
```

```
assertIsValid(eventData, ajv.compile(schema))
event = newEvent(
      aggregateId,
      aggregateType,
      eventData
```

```
assertIsValid(eventData, ajv.compile(schema))
event = newEvent(
     aggregateId,
      aggregateType,
      eventData
```

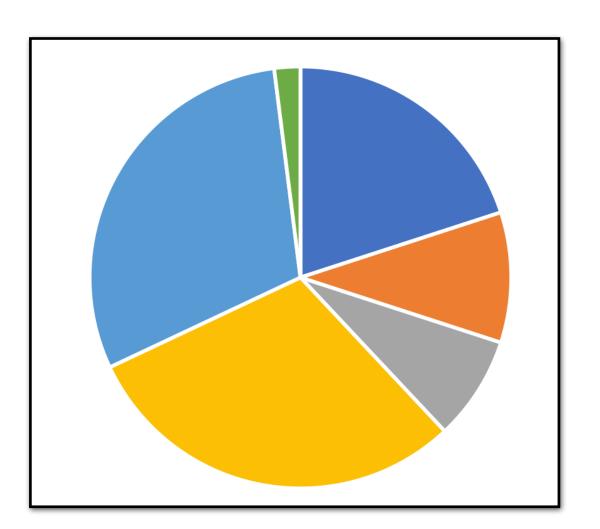
Schema-on-Read

Reusing event data?

Transactionledger Microservice



Budget Planer Microservice



Transactionledger Microservice

TransactionBooked

Budget Planer Microservice

TransactionTagged

Just copy data into different events, just so convenient



transactionId: ...

accountNumber: ...

amount: ...

currency: ...

bookingTime: ...

purpose: ...

transactionId: ...
accountNumber: ...
amount: ...
currency: ...

bookingTime: ...

purpose: ...

TransactionTagged

tagld: ...

categoryName: "..."

transactionId: ...

amount: ...

currency: ...

```
transactionId: ...
accountNumber: ...
amount: ...
currency: ...
bookingTime: ...
```

purpose: ...

TransactionTagged

```
tagld: ...
categoryName: "..."
transactionId: ...
amount: ...
currency: ...
```

But I need to display the transaction purpose, too

Budget Planer Microservice

transactionId: ...
accountNumber: ...
amount: ...
currency: ...
bookingTime: ...
purpose: ...

TransactionTagged

tagld: ...
categoryName: "..."
transactionId: ...
amount: ...
currency: ...
???

The lossy event

Only reference aggregates via their root id or event ids

eventld:...
transactionld: ...

accountNumber: ...

amount: ...

currency: ...

bookingTime: ...

purpose: ...

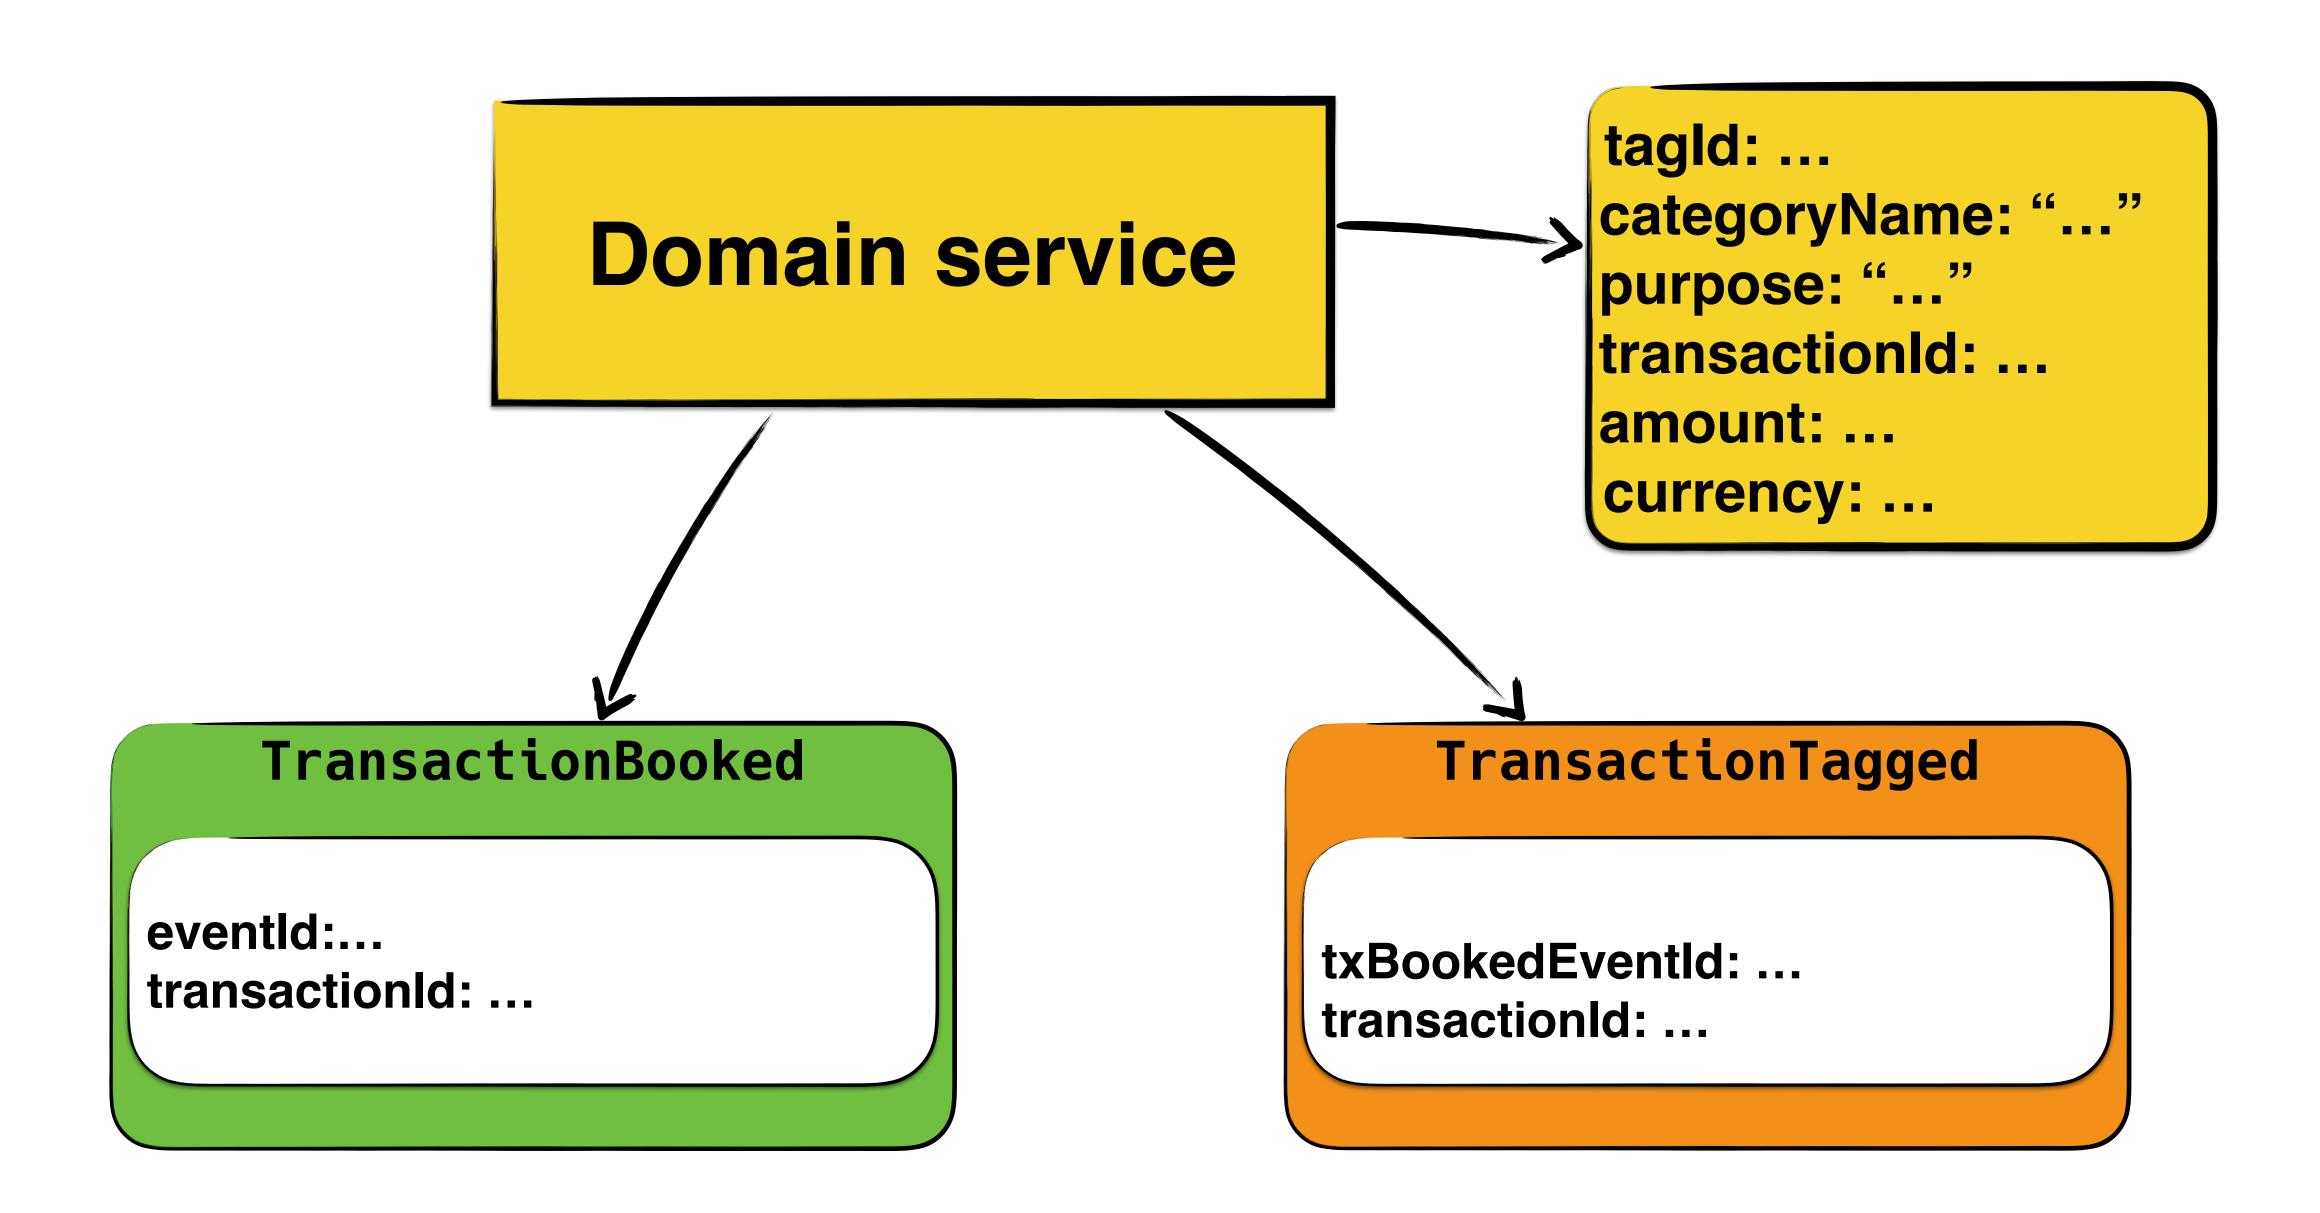
TransactionTagged

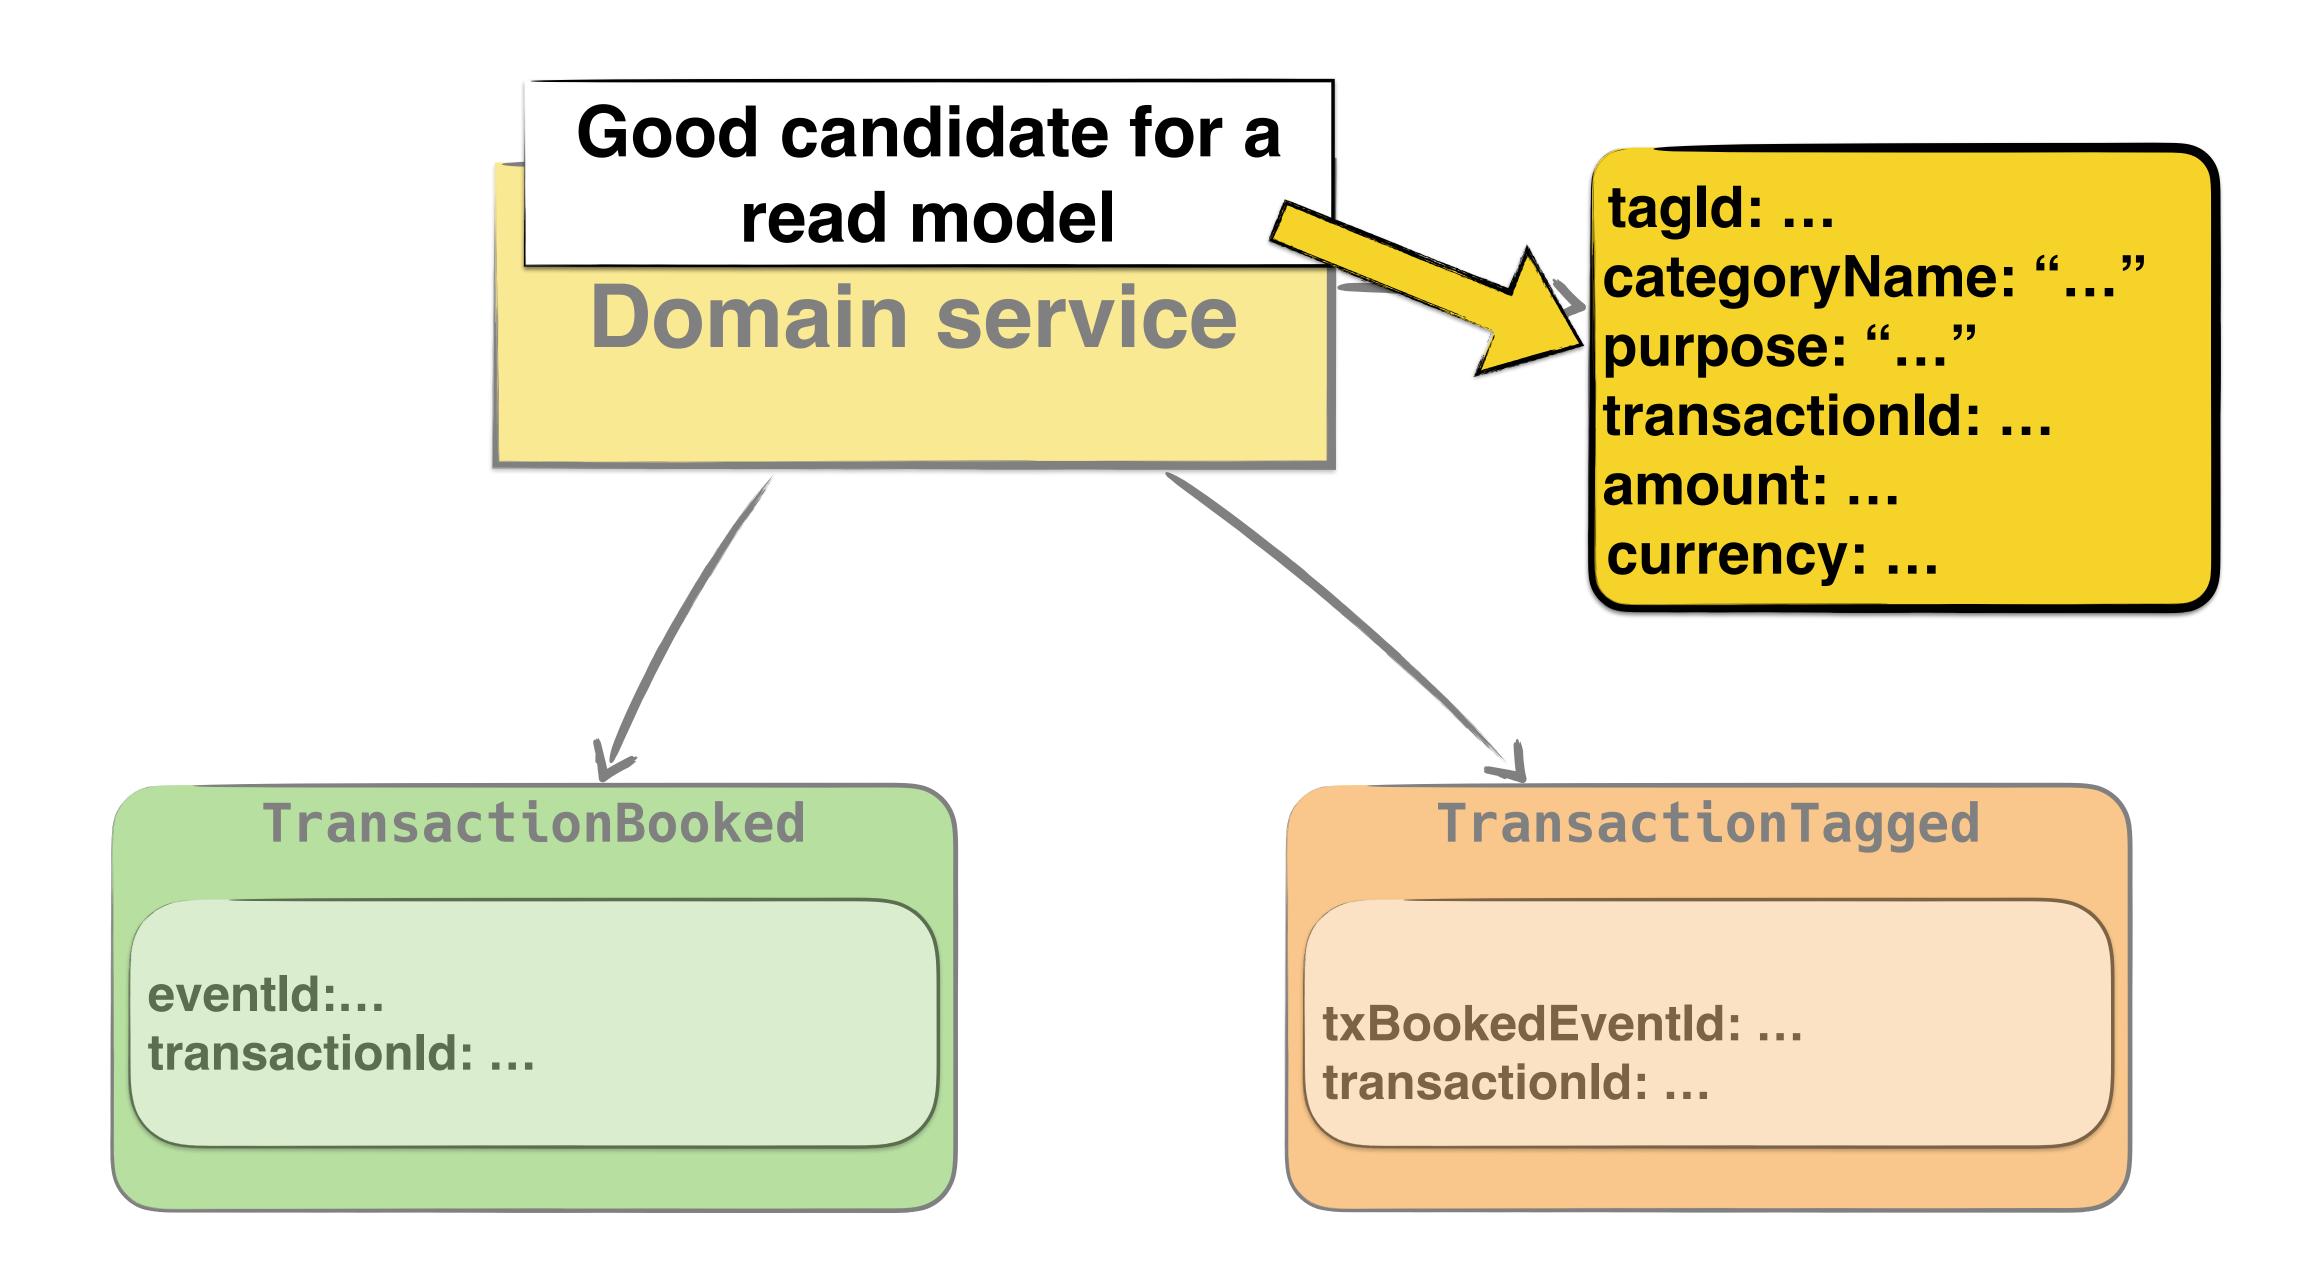
tagld: ...

categoryName: "..."

txBookedEventId: ...

transactionId: ...





Don't copy parts of an event

Prefer building use case specific projections

How do I handle large streams?

How can you handle event data over a long period of time?

You don't

Just create a snapshot of the stream



Year's end procedure

Year end – also known as an accounting reference date – is the completion of an accounting period. At this time, businesses need to carry out specific procedures to close their books.

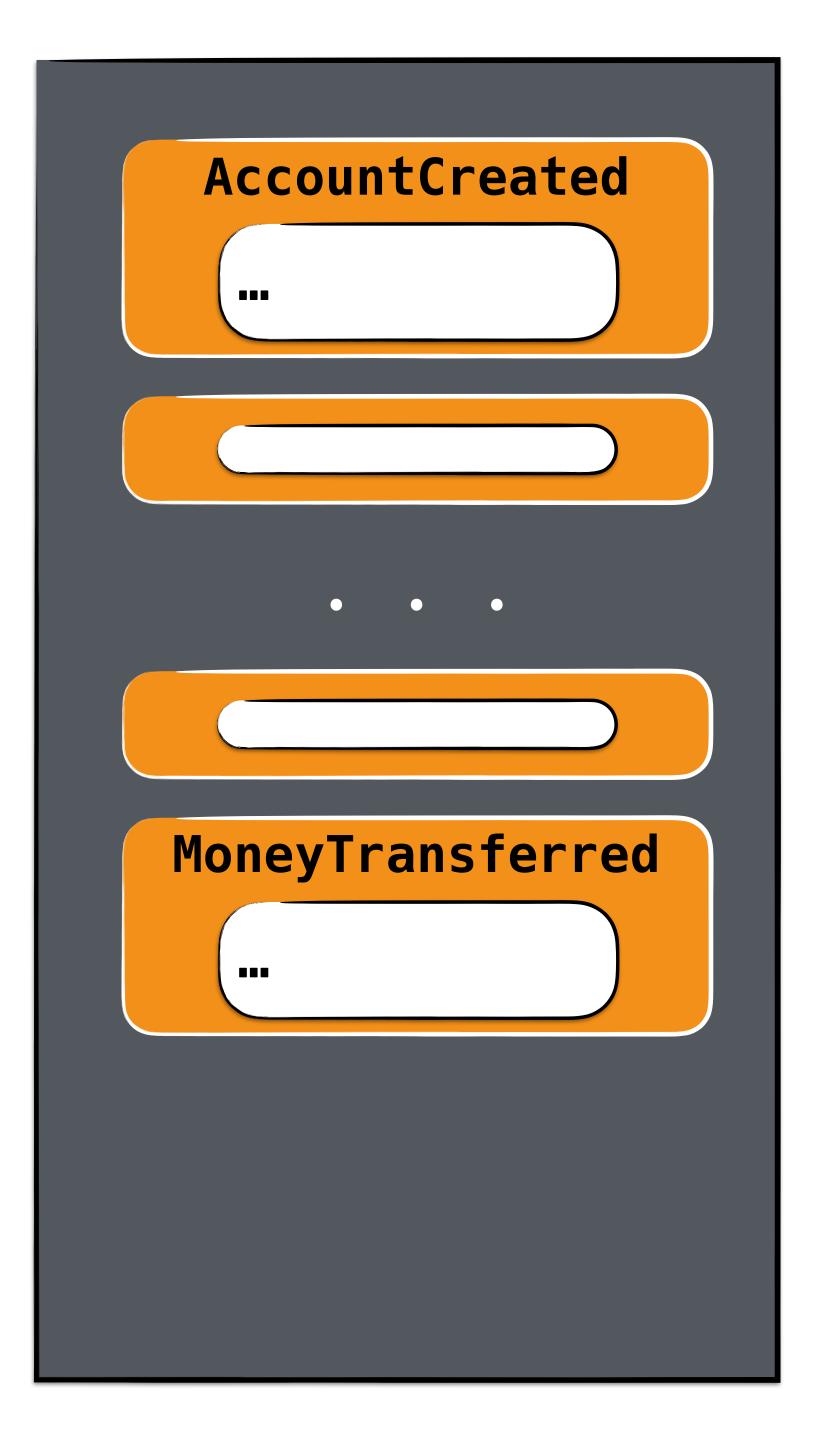
https://debitoor.com/dictionary/year-end

Year end – also known as an accounting reference date – is the completion of an accounting period. At this time, businesses need to carry out specific procedures to close their books.

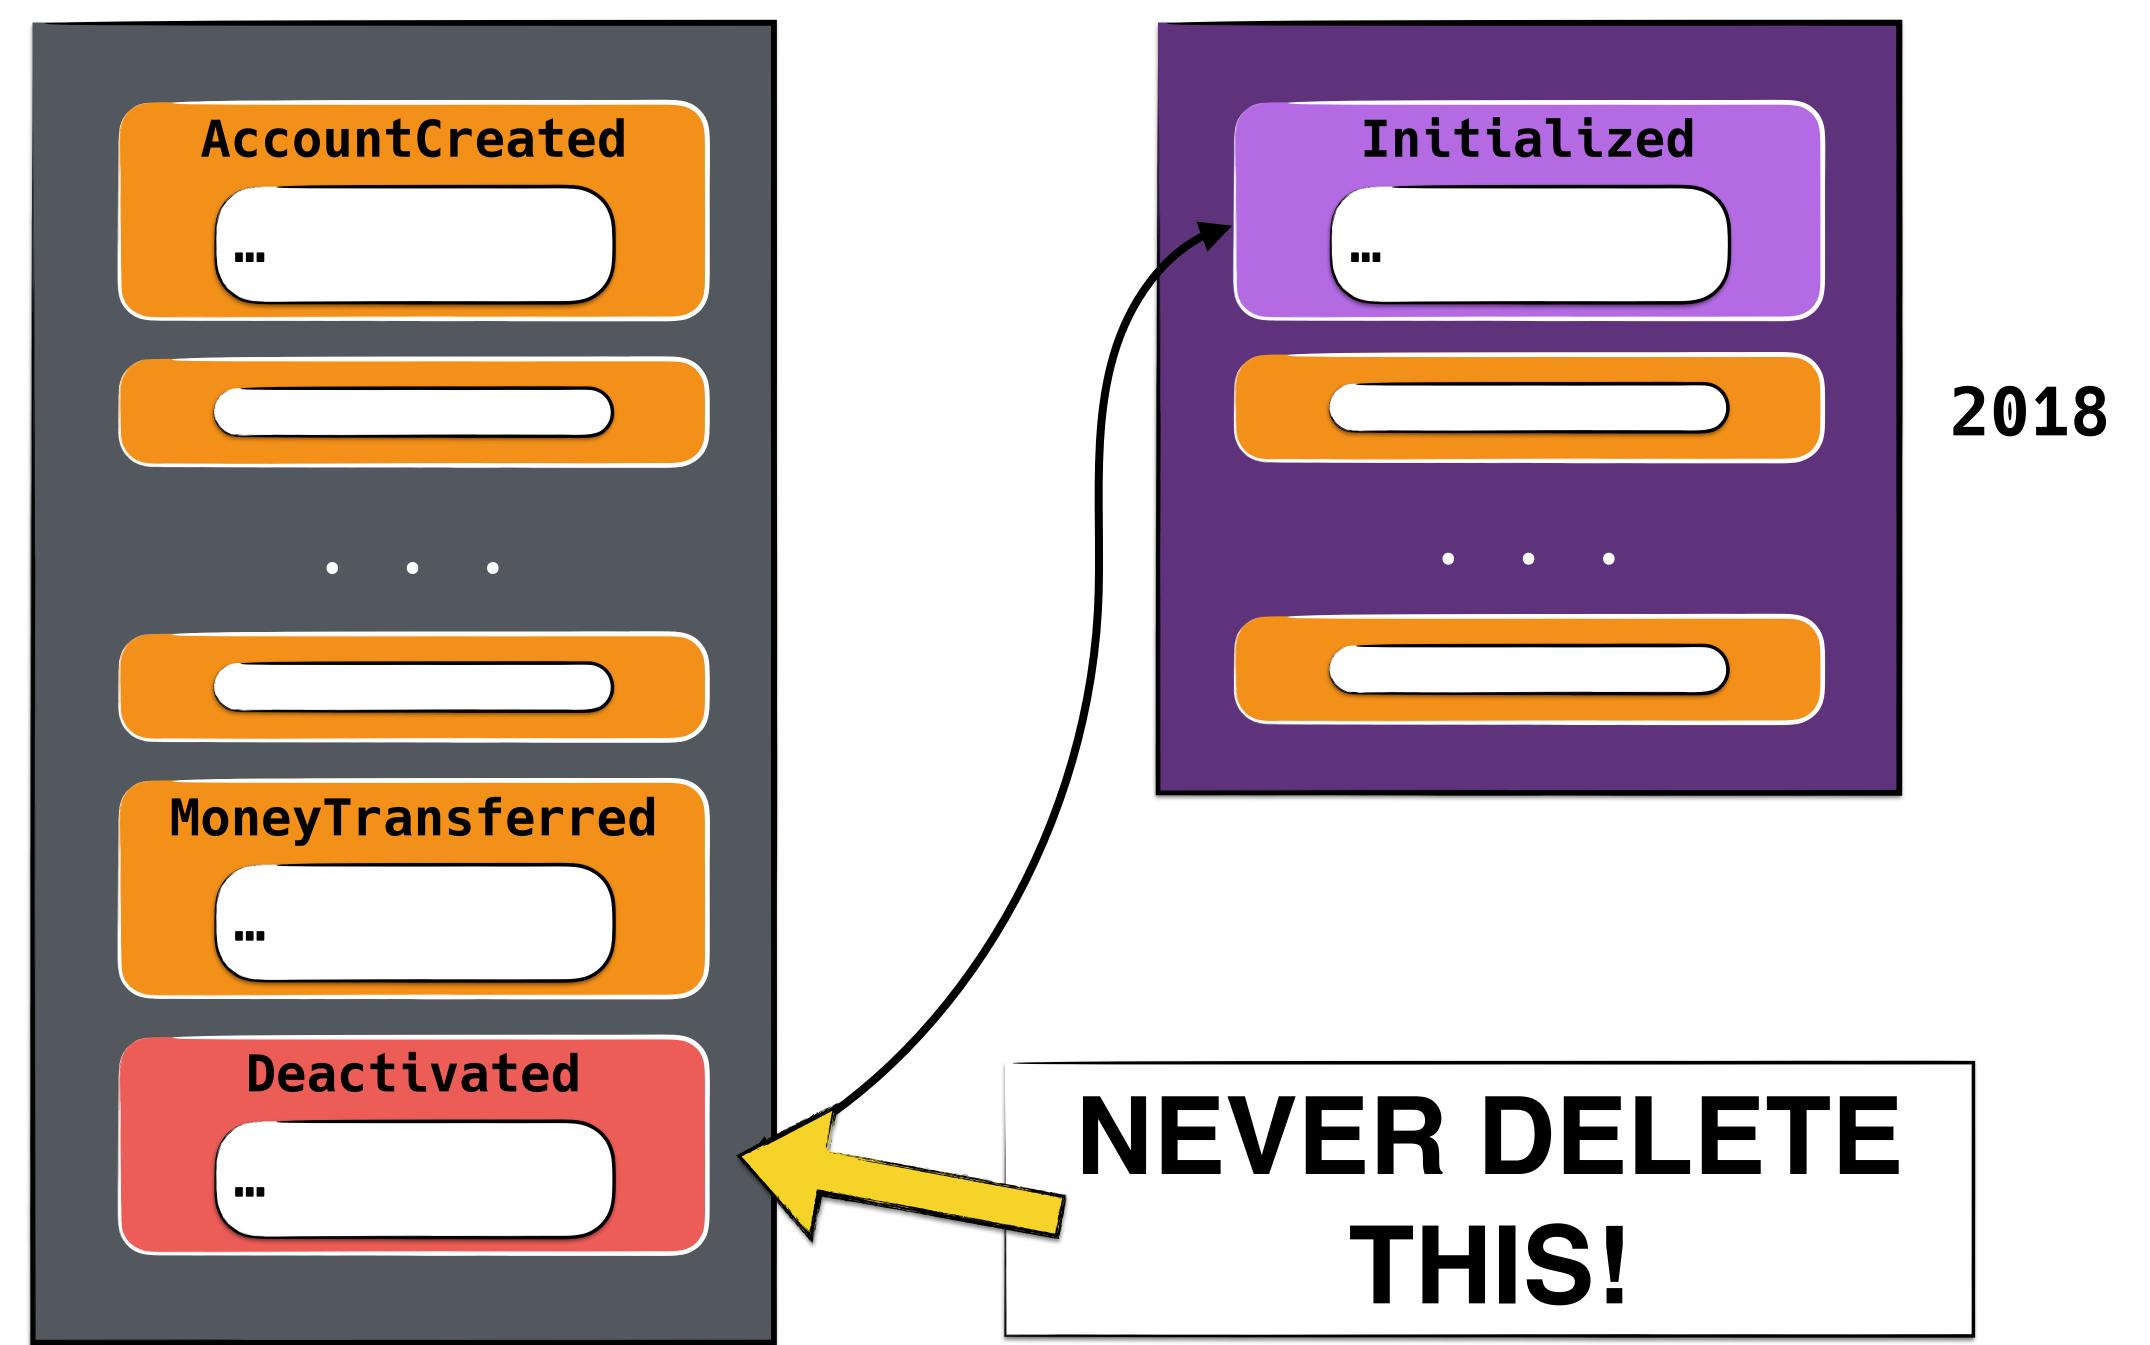
https://debitoor.com/dictionary/year-end

Copy-Transform

a.k.a. event sourcing refactoring powertool

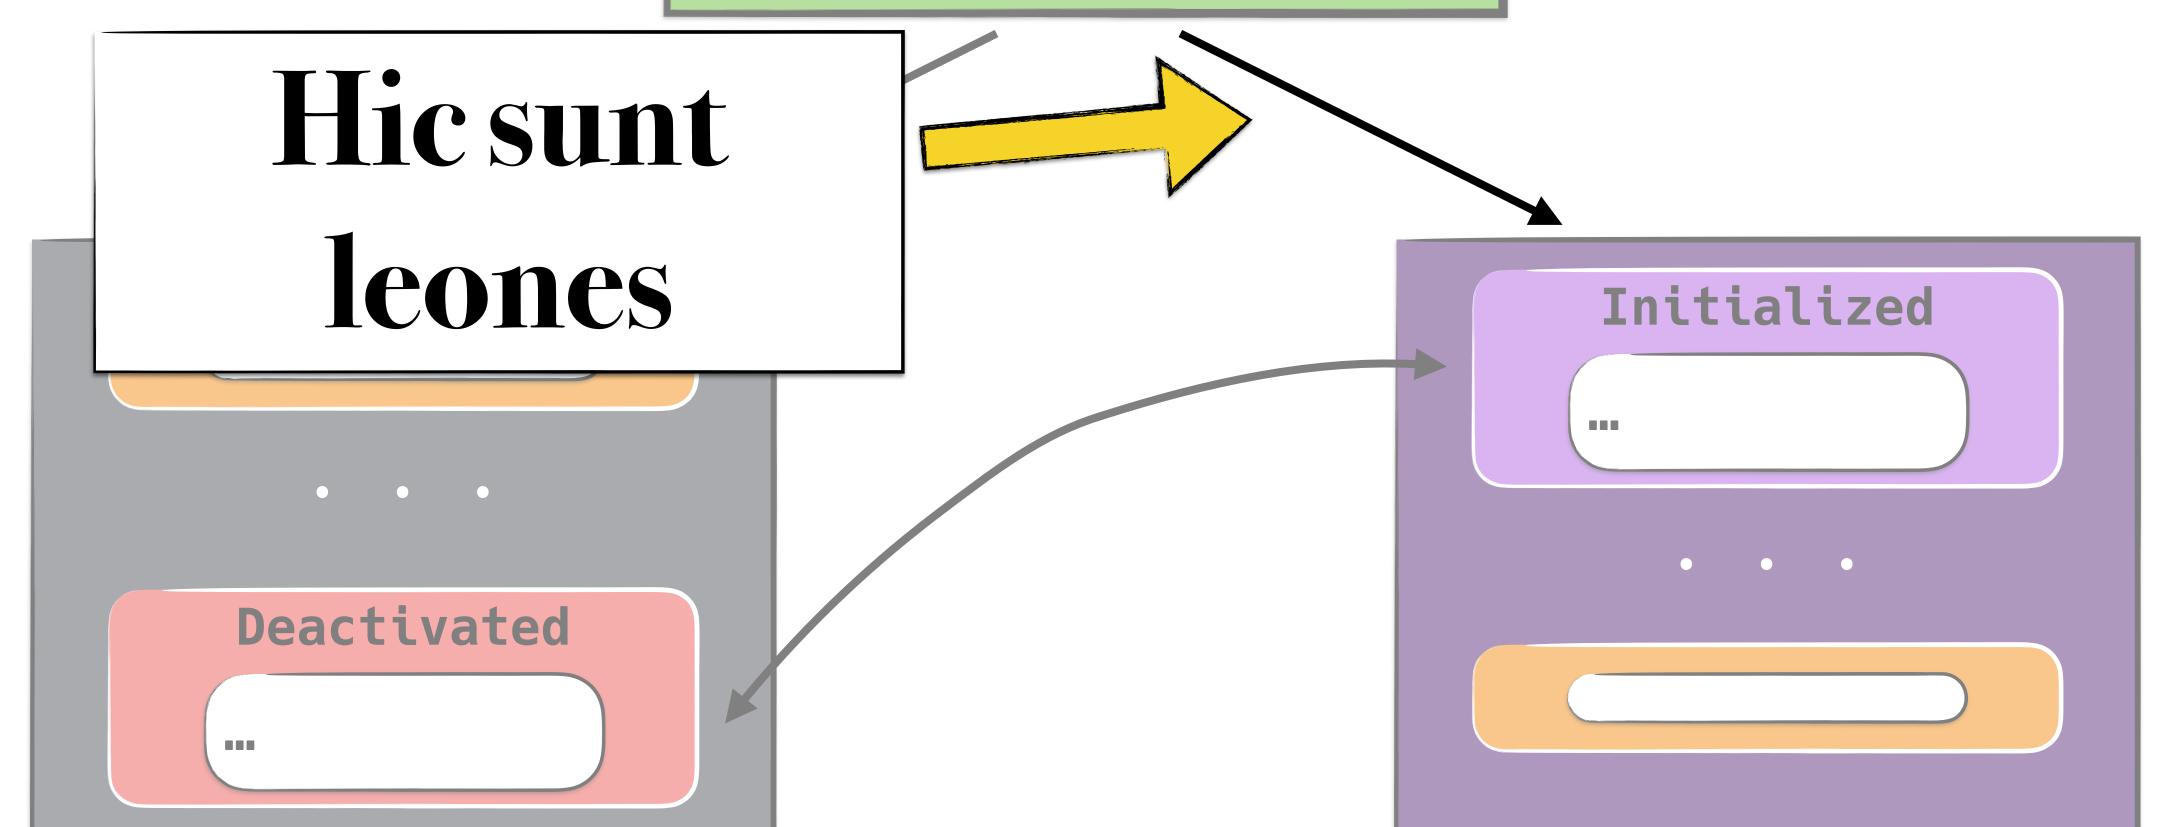


2017



2017

Transactionledger Microservice

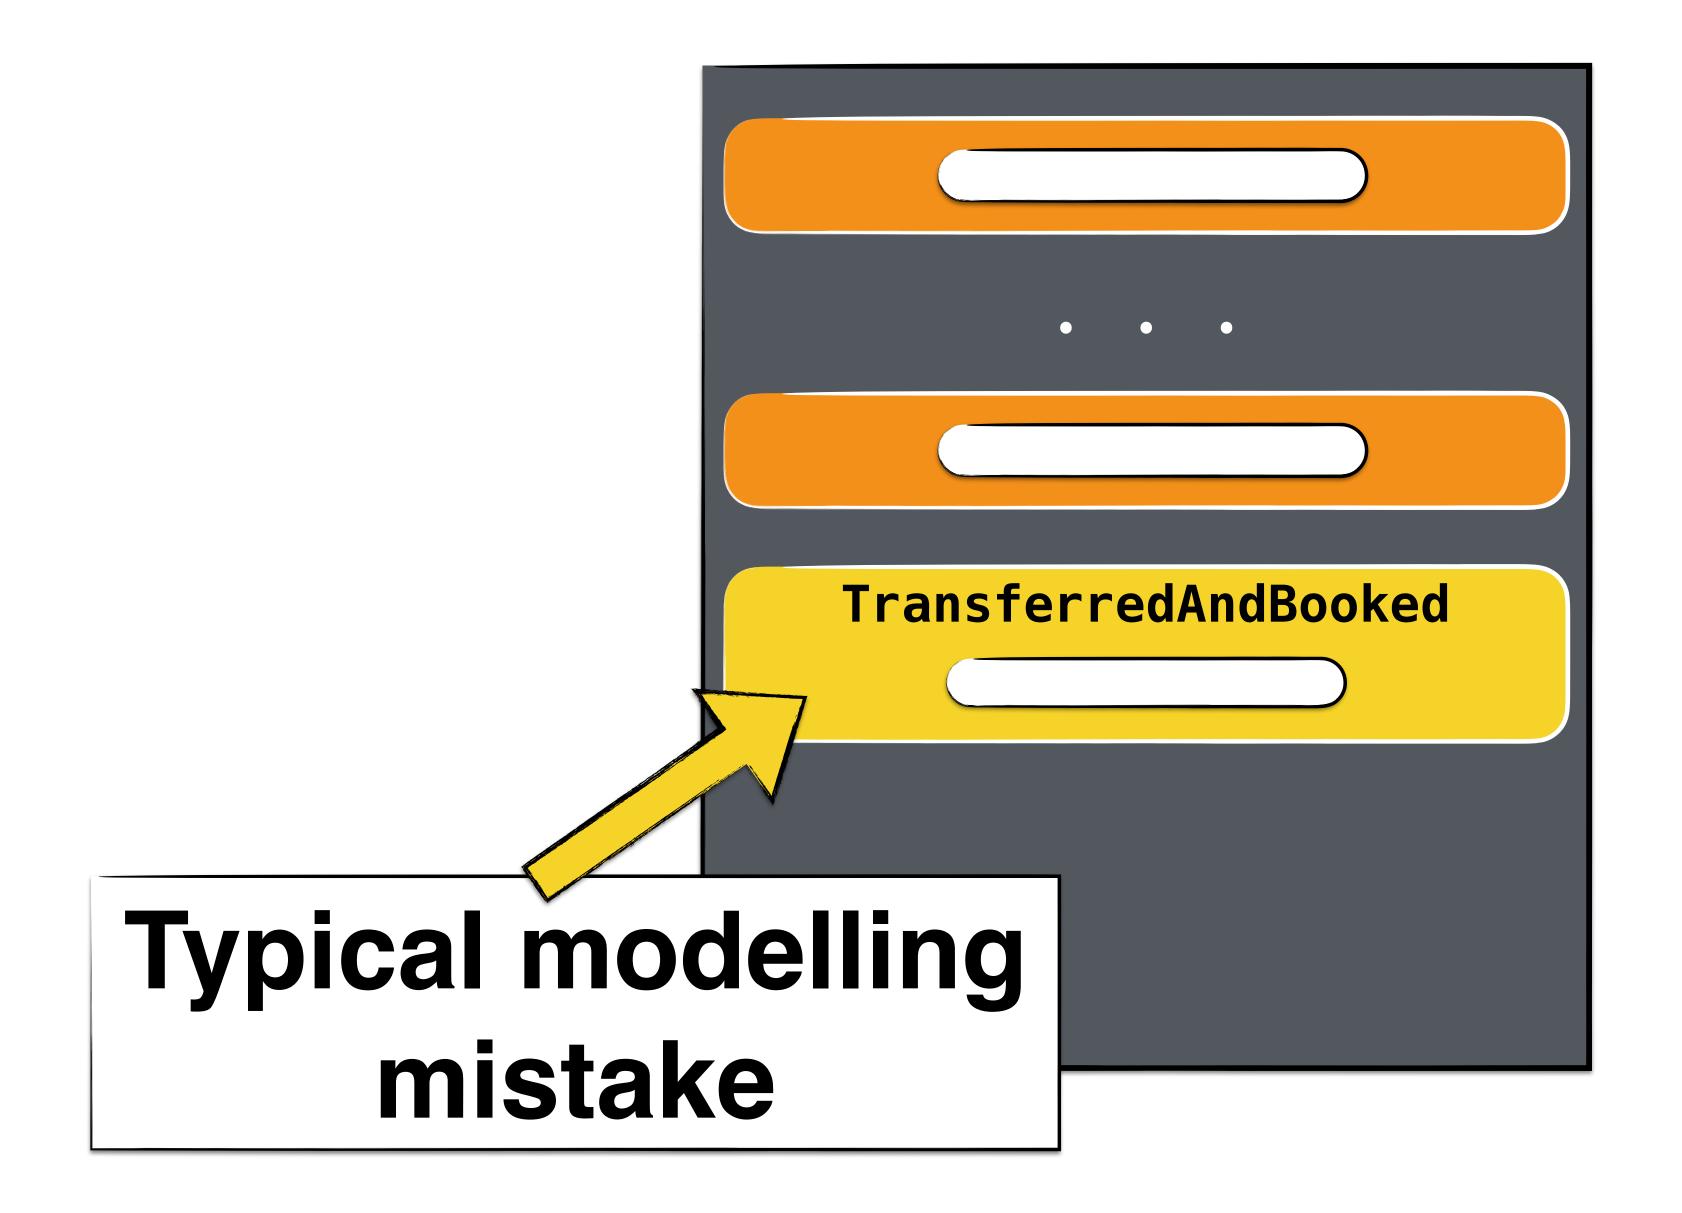


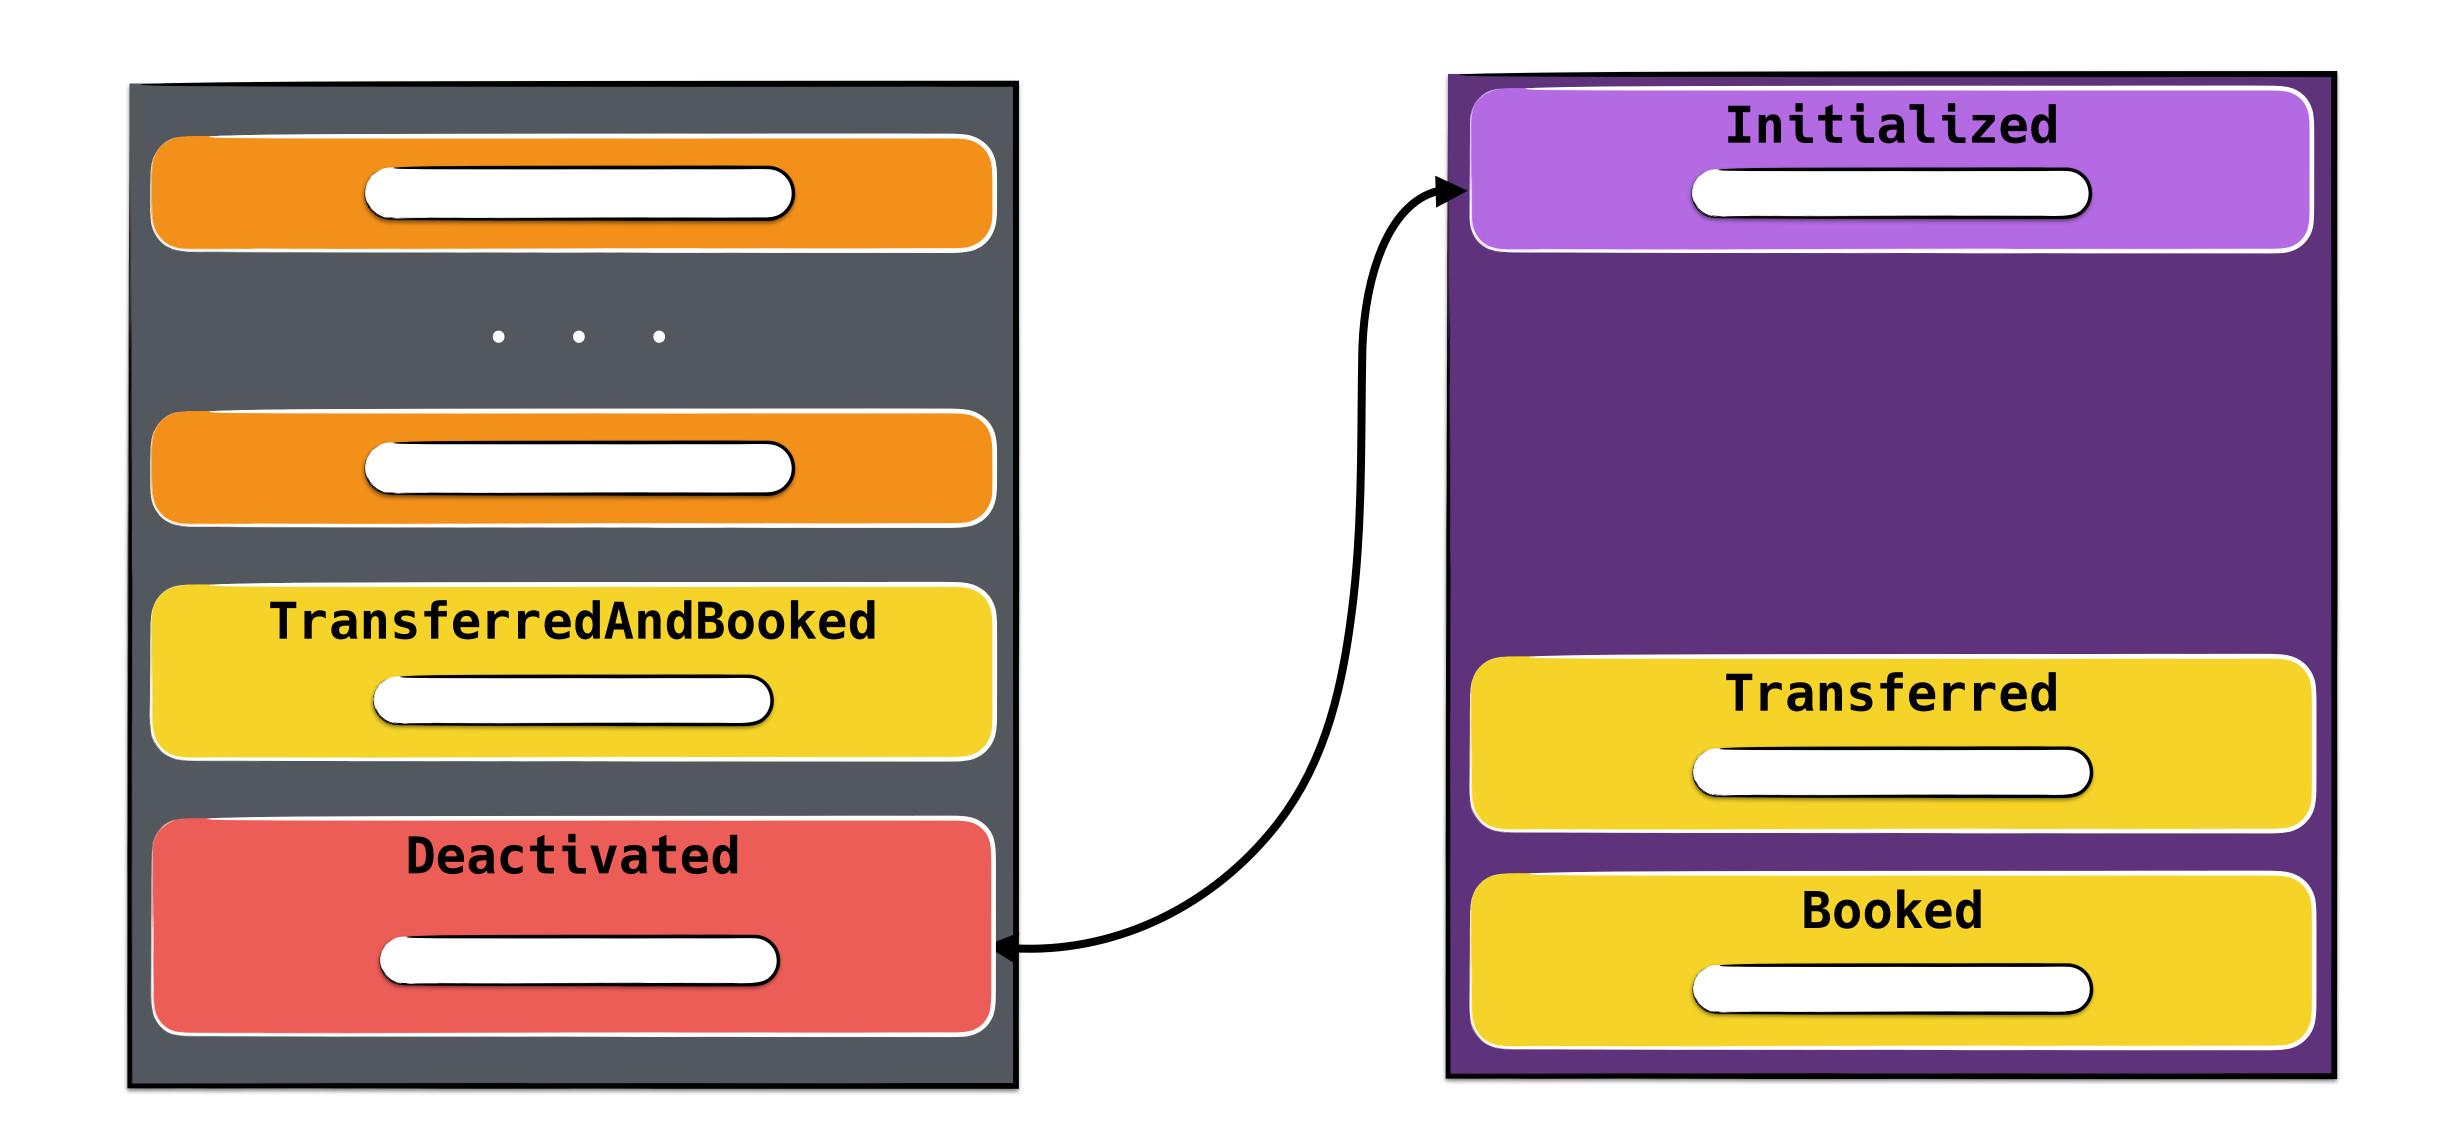
2017

2018

Same idea if you need to remodel your domain!

The FooANDBarEvent





The devil is in the detail

Dealing with errors



MoneyTransferred eventId: 231233

amount: 97 (Euro)

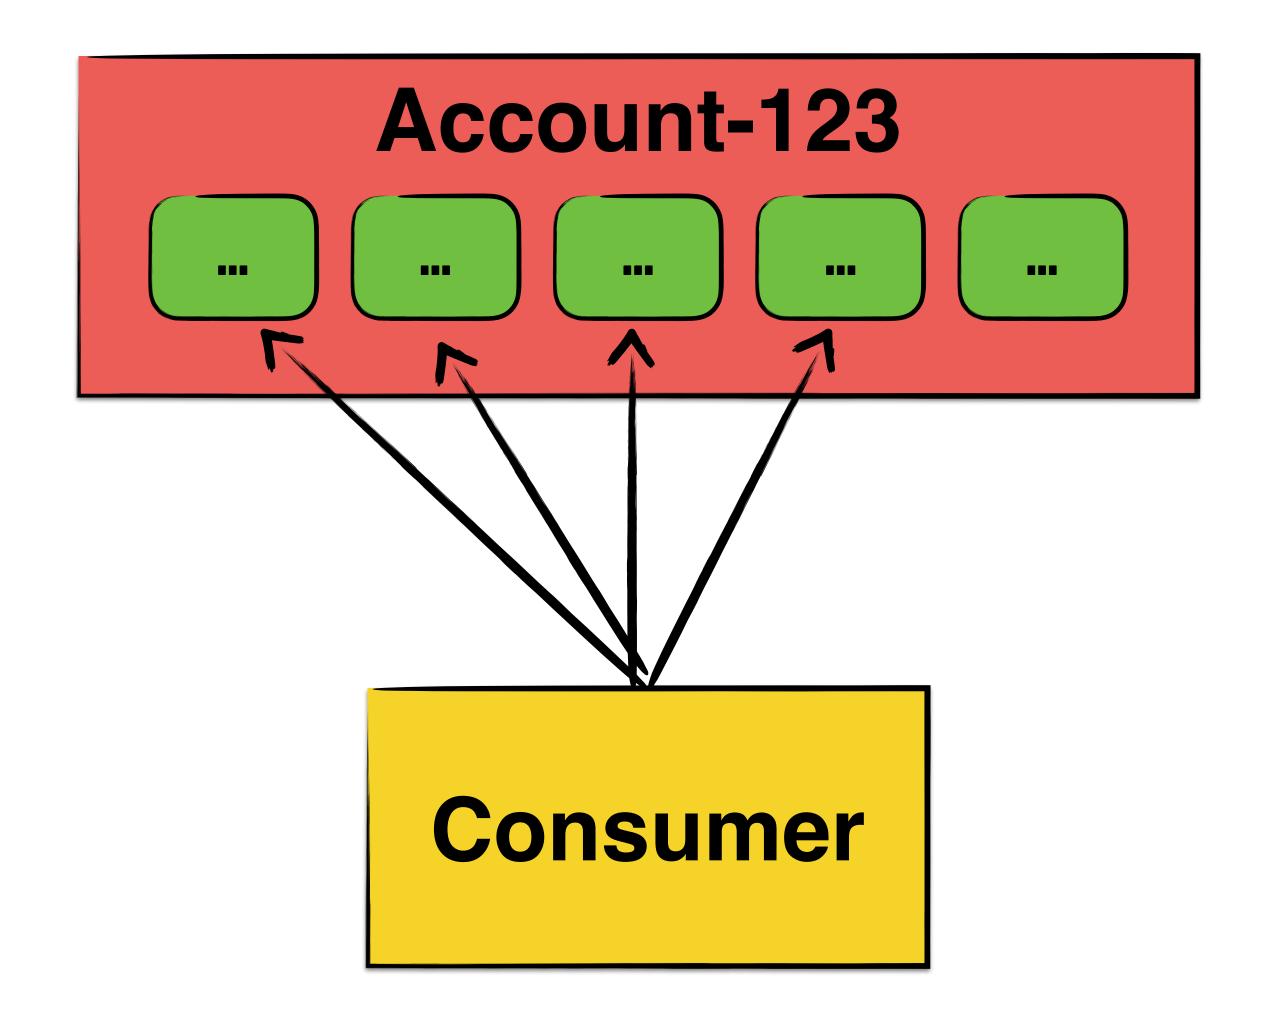
withdrawnAt: 2018-08-30T08:58:26.624



Just update the event in the eventstore!



Challenges?



Account-123 Update ... I already know that event. Why should I re-read? Consumer - (ツ)_/

Ok. Then just use compensation events



The cancelled or corrected event

Partial compensation?

MoneyTransferred eventId: 8

amount: 97 Euro

MoneyTransferAmountCorrected eventId: 9

amount: 97 EUR

reasonEventId: 8

Full compensation do as accountants do

MoneyTransferred eventId: 1

amount: 97 Euro

•••

MoneyTransferCancelled eventId: 2

reasonEventId: 1

reason: ...

MoneyTransferred eventId: 3

amount: 97 EUR

•••

A full compensation is the explicit reason for compensation







REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016

on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)





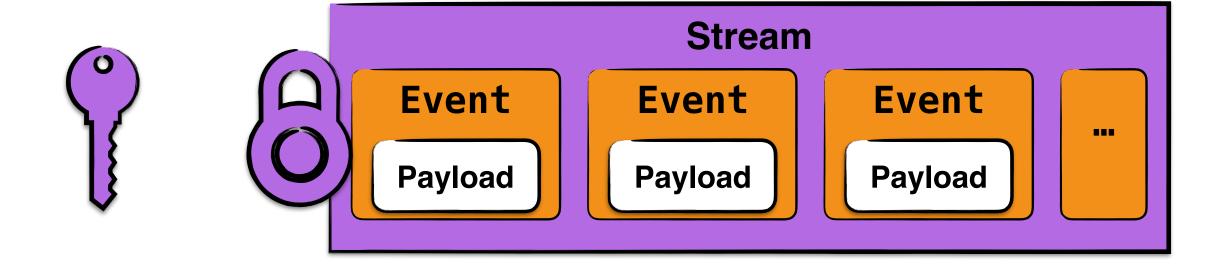
Article 17

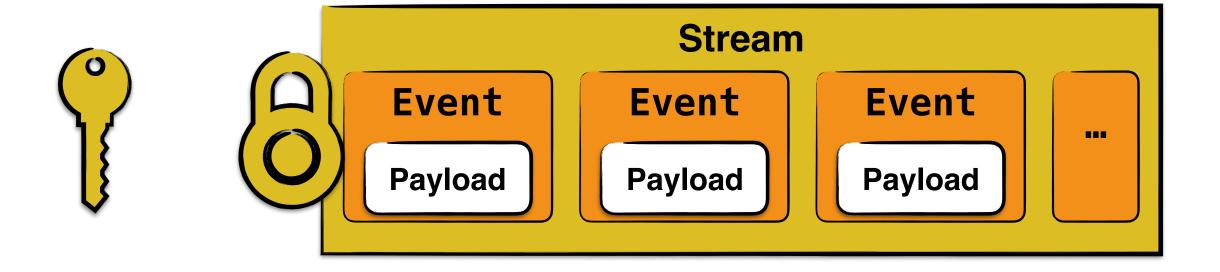
Right to erasure ('right to be forgotten')

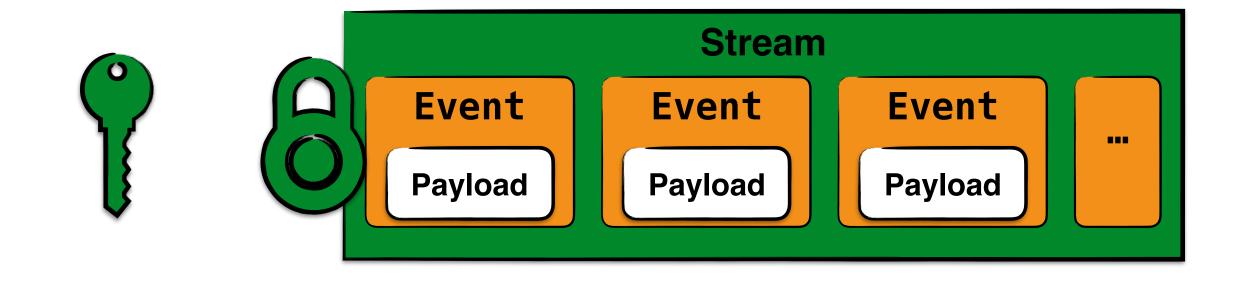
1. The data subject shall have the right to obtain from the controller the erasure of personal data concerning him or her without undue delay and the controller shall have the obligation to erase personal data without undue delay where one of the following grounds applies:







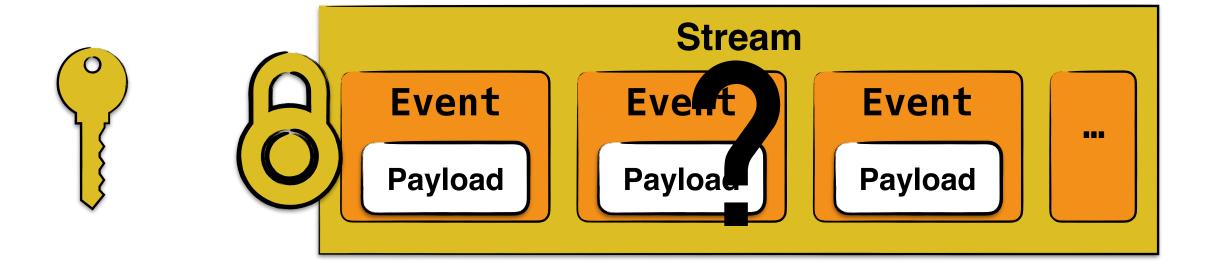


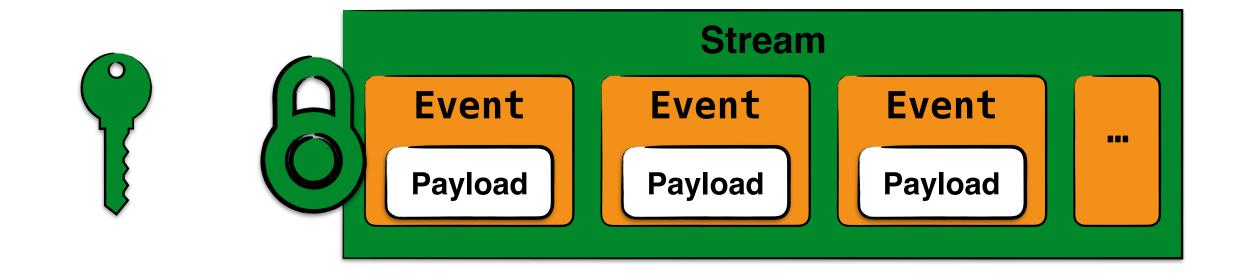


"Please delete all my data"

Deletion is effectively deleting the stream-specific key







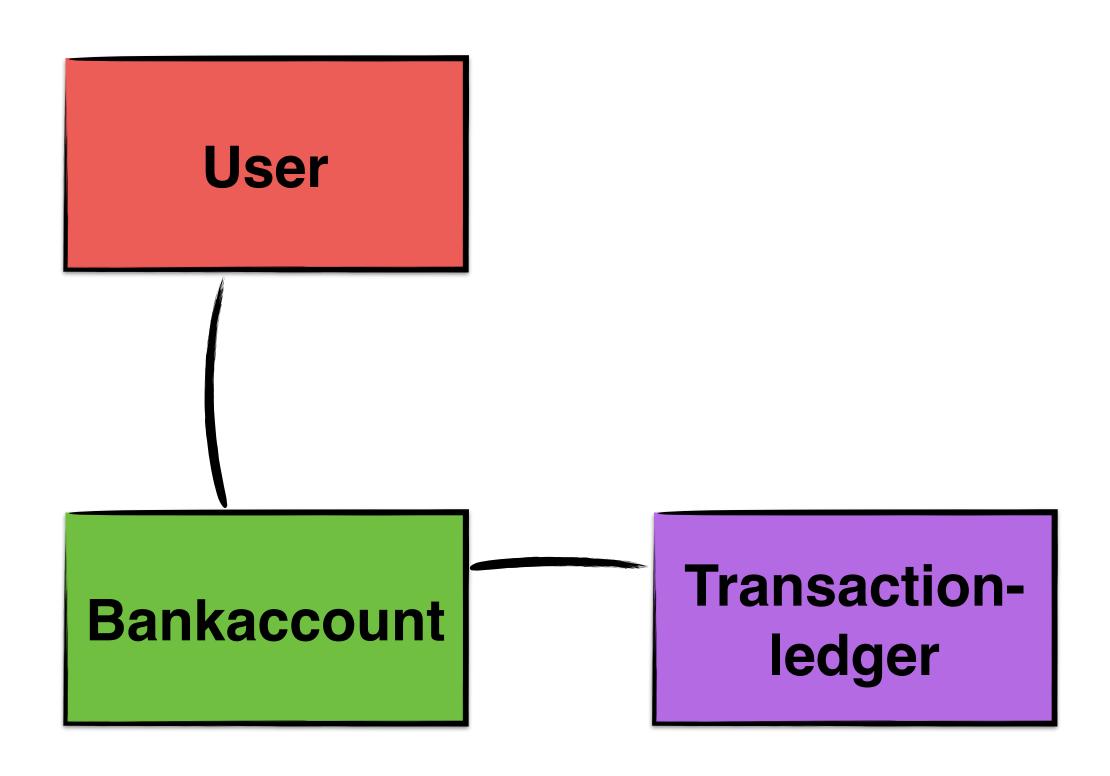
Challenges?

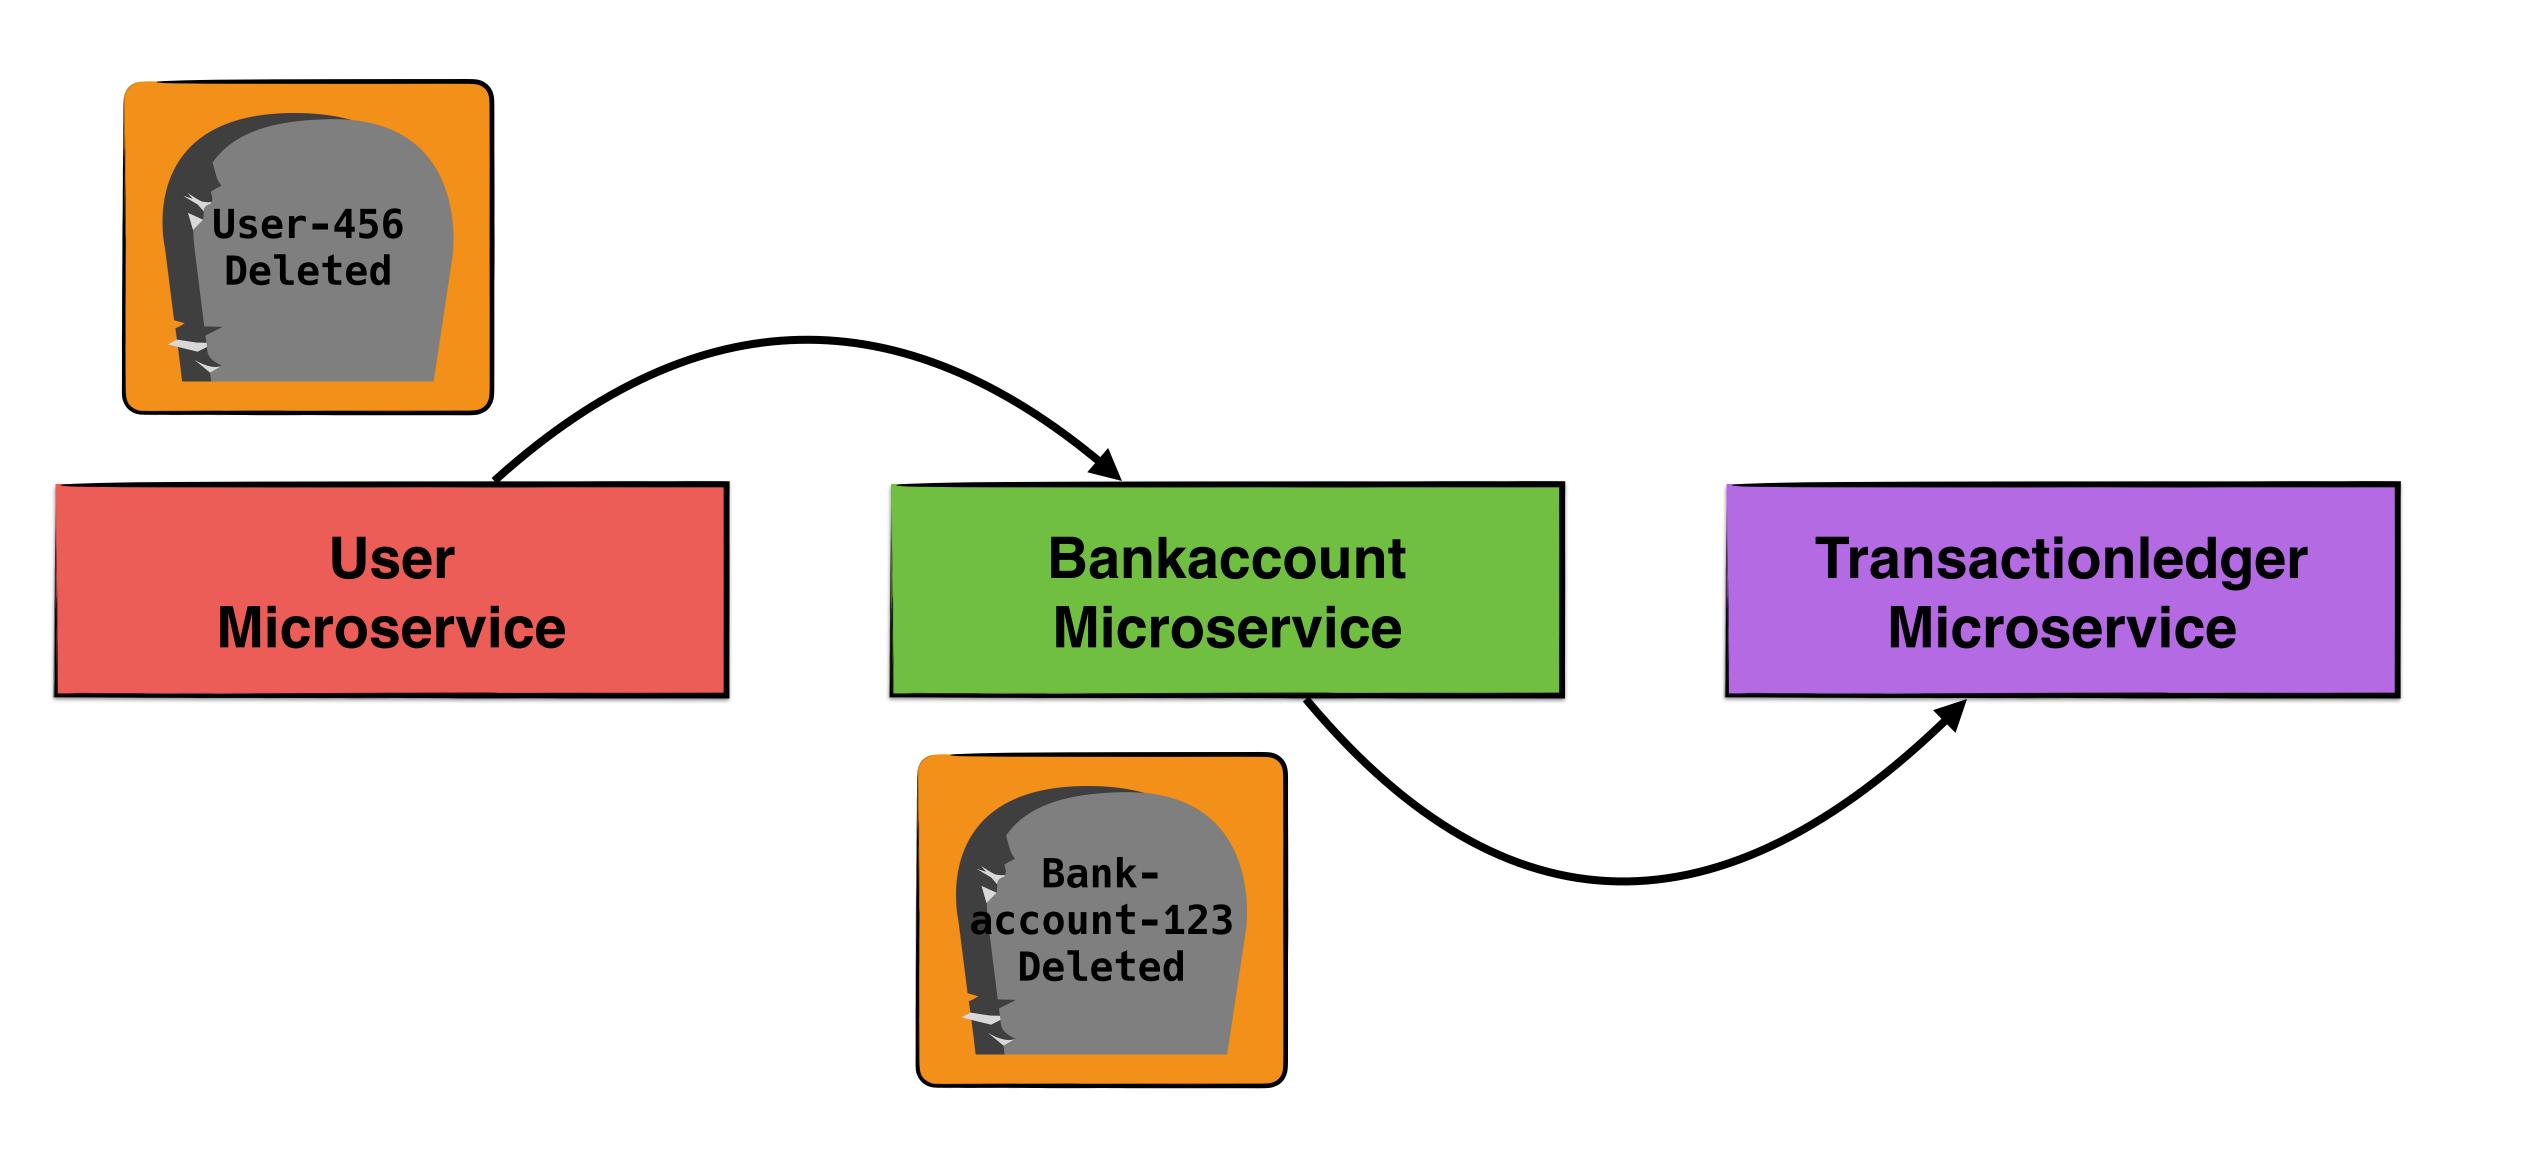
Key administration Finding what needs to be deleted Storage implications Coding complexity Dashboards, Monitoring

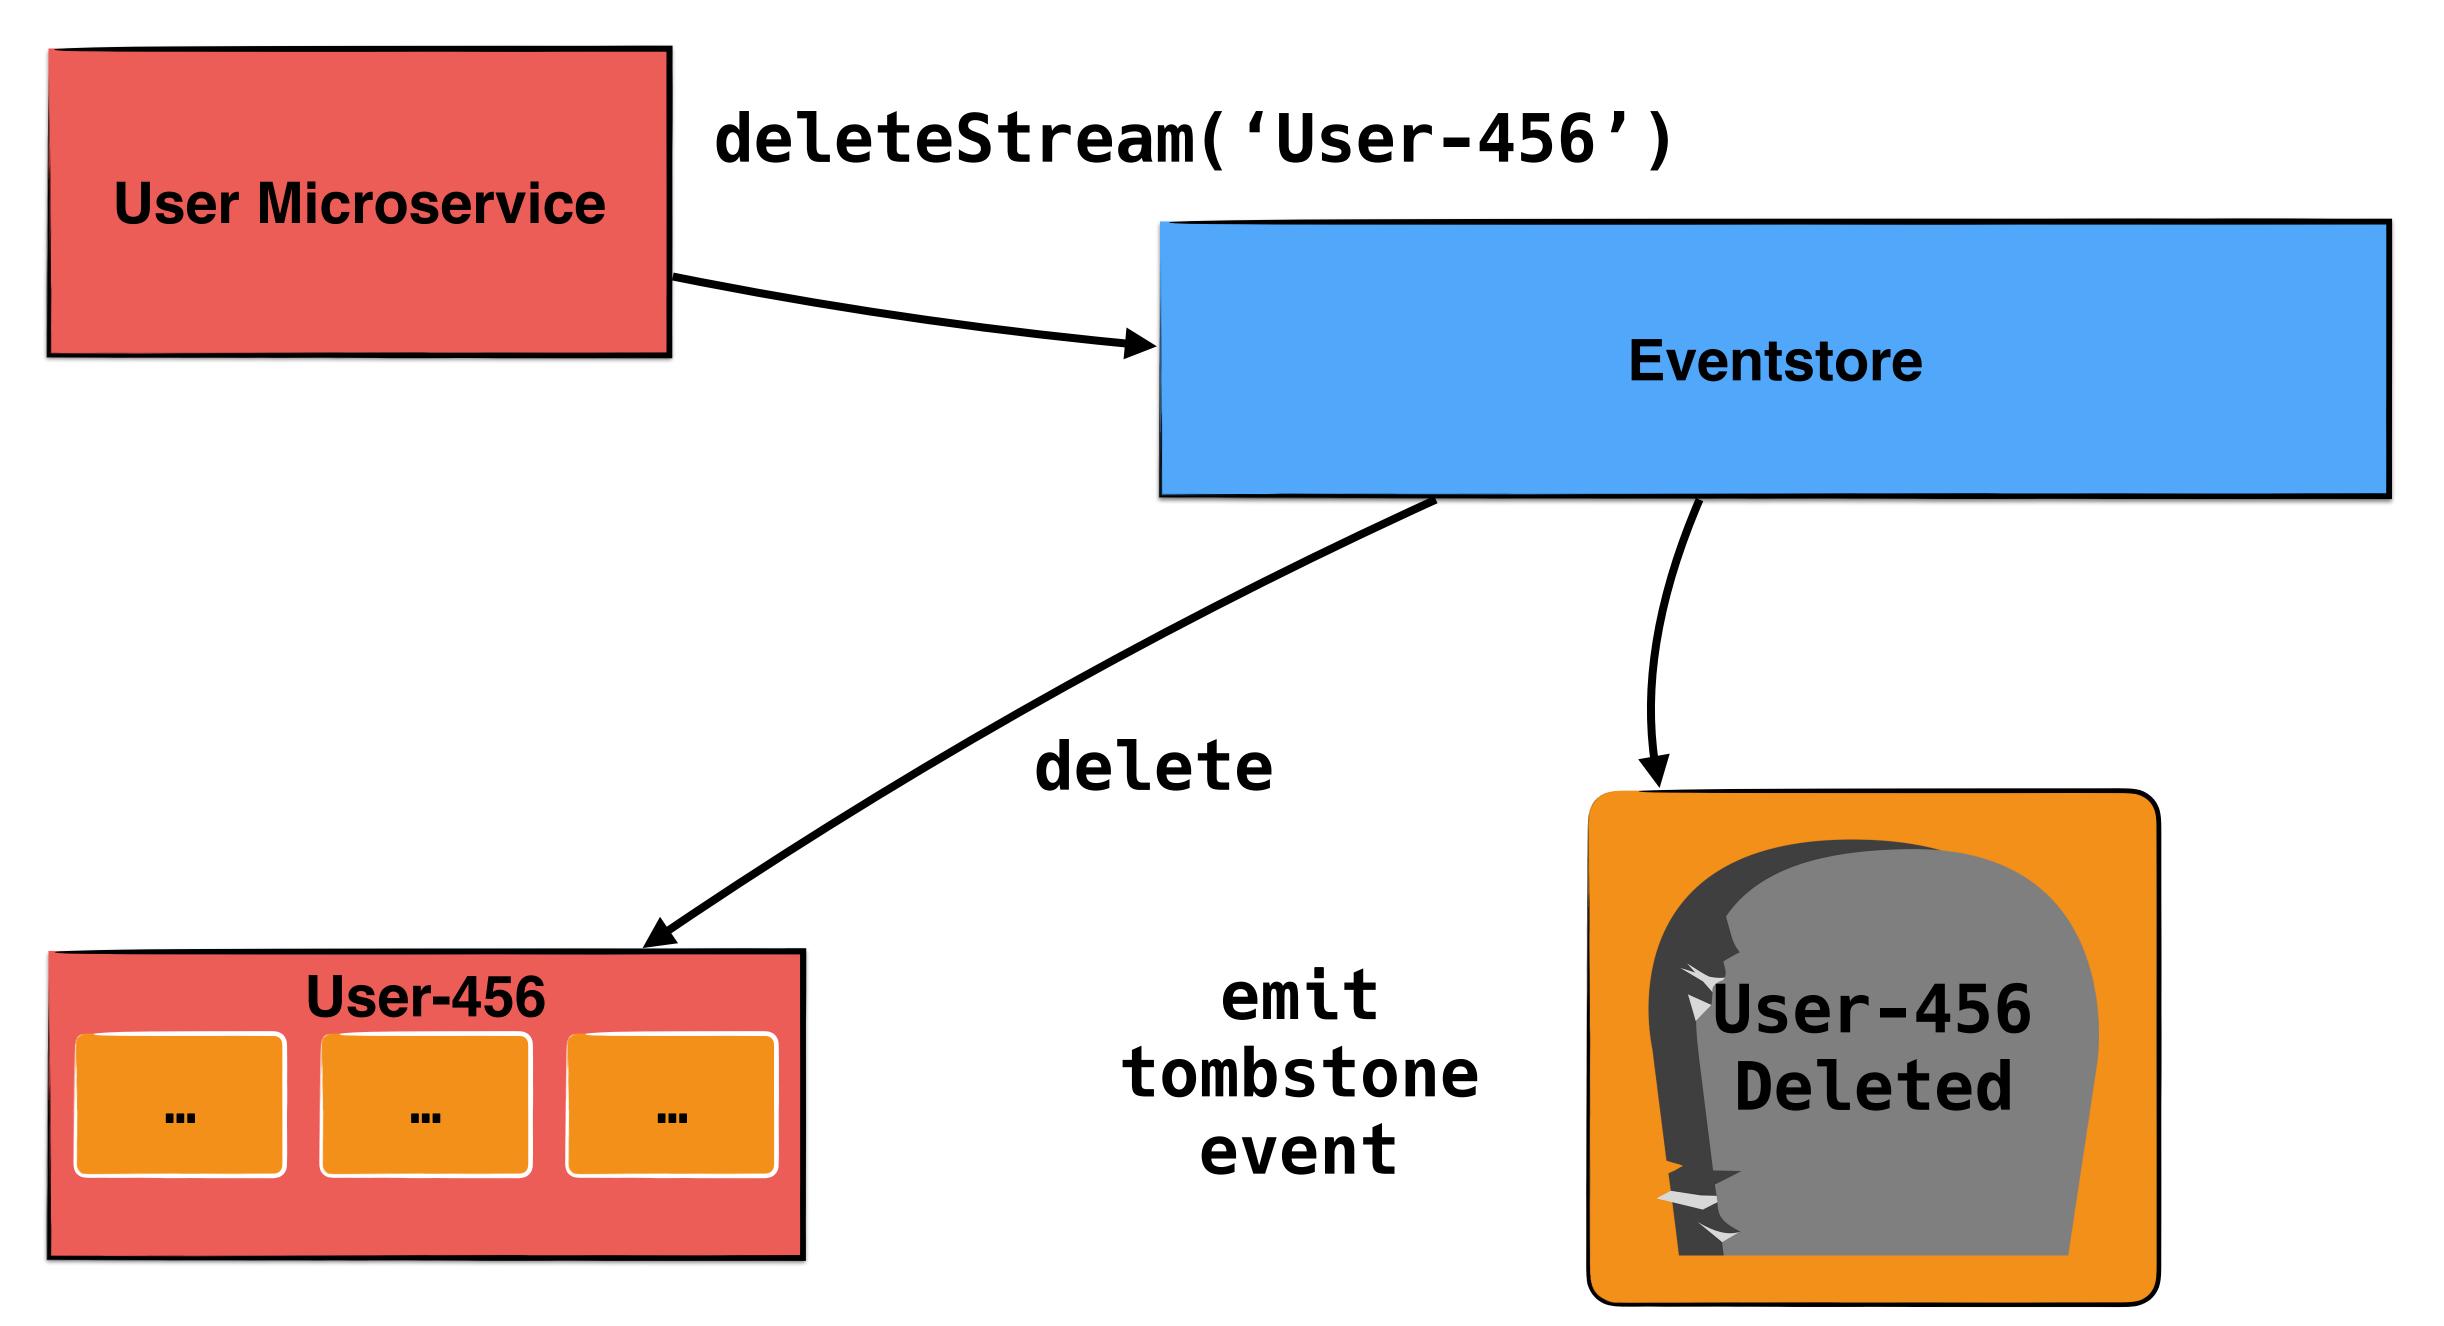
Being able to delete is awesome

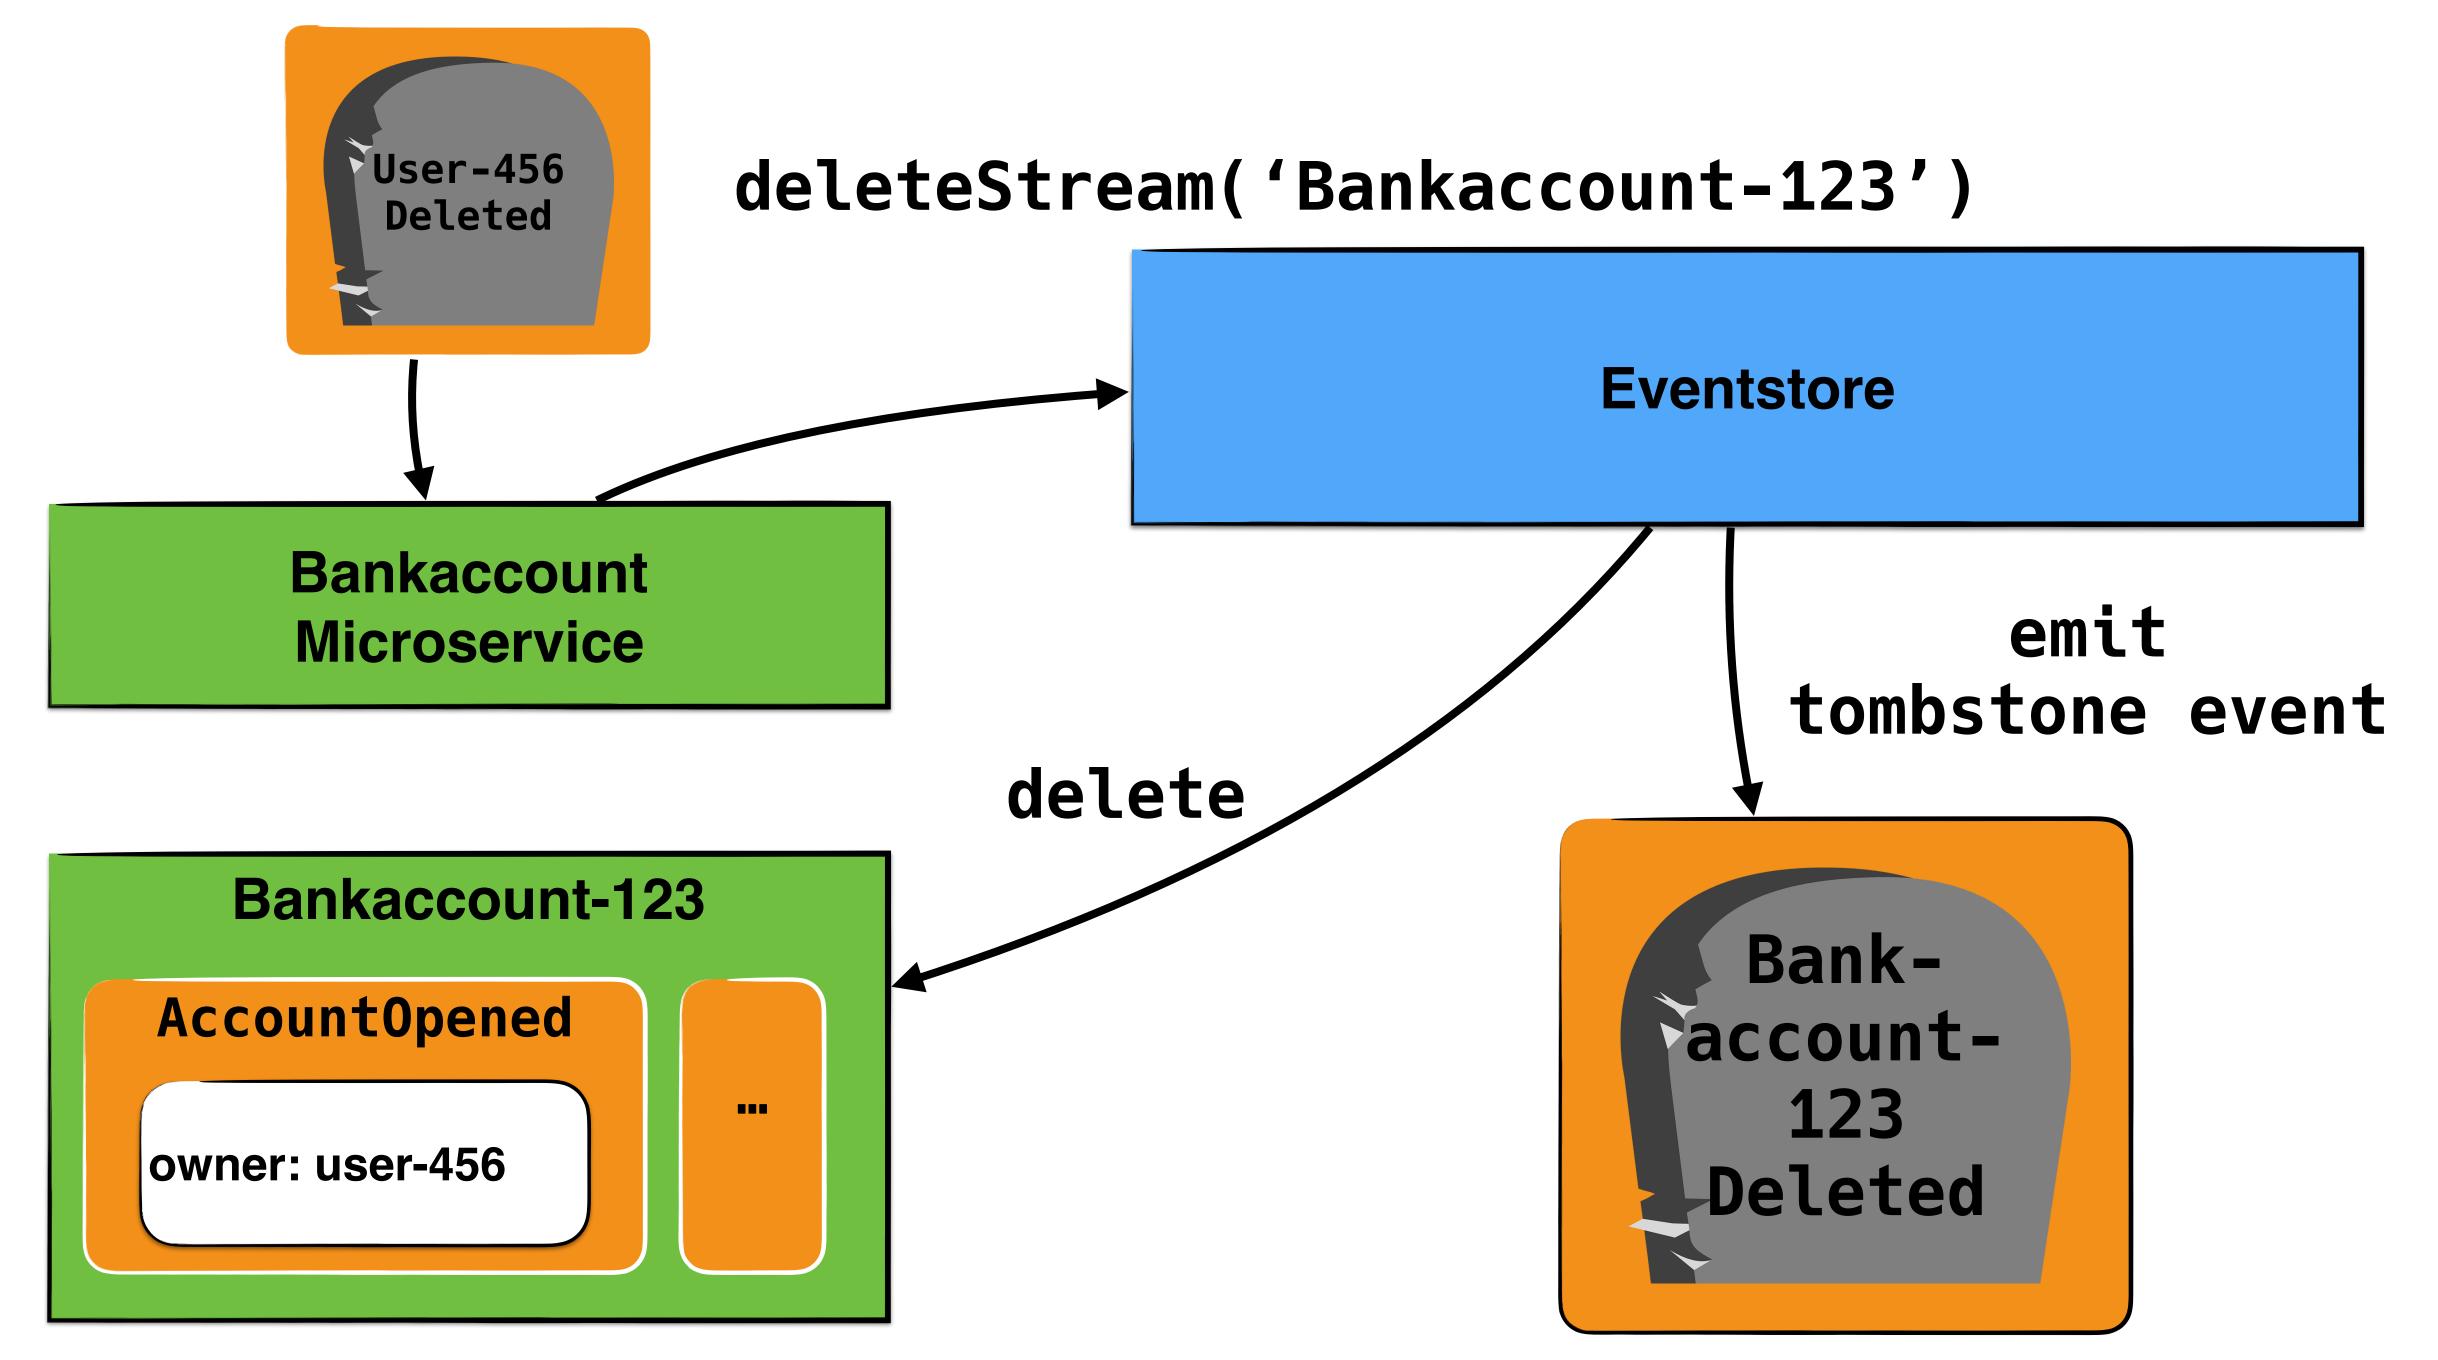
Cascading deletes with tombstones

"Please delete USER-456"











What about dependencies between aggregates?

Just anonymise the data





Recital 26 EU GDPR

(26) The principles of data protection should apply to any information concerning an identified or identifiable natural person.

Personal data which have undergone pseudonymisation, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person.





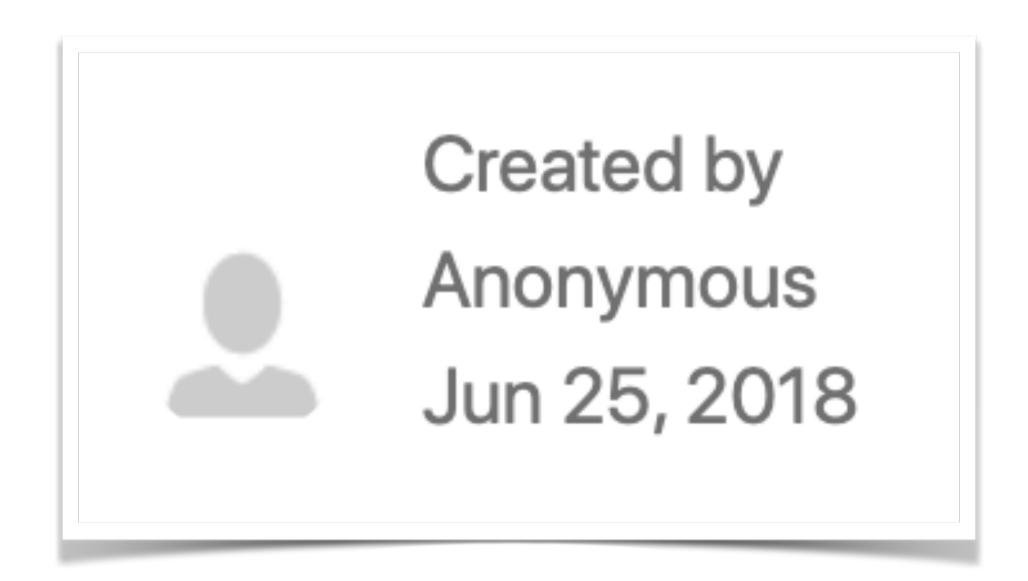
Recital 26 EU GDPR

(26) The principles of data protection should apply to any information concerning an identified or identifiable natural person.

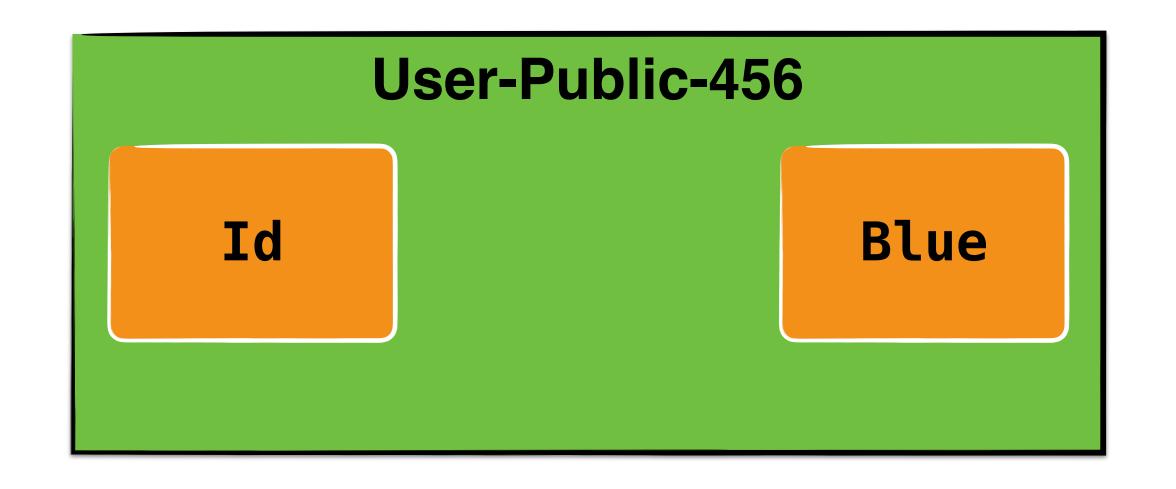
Personal data which have undergone pseudonymisation, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person.

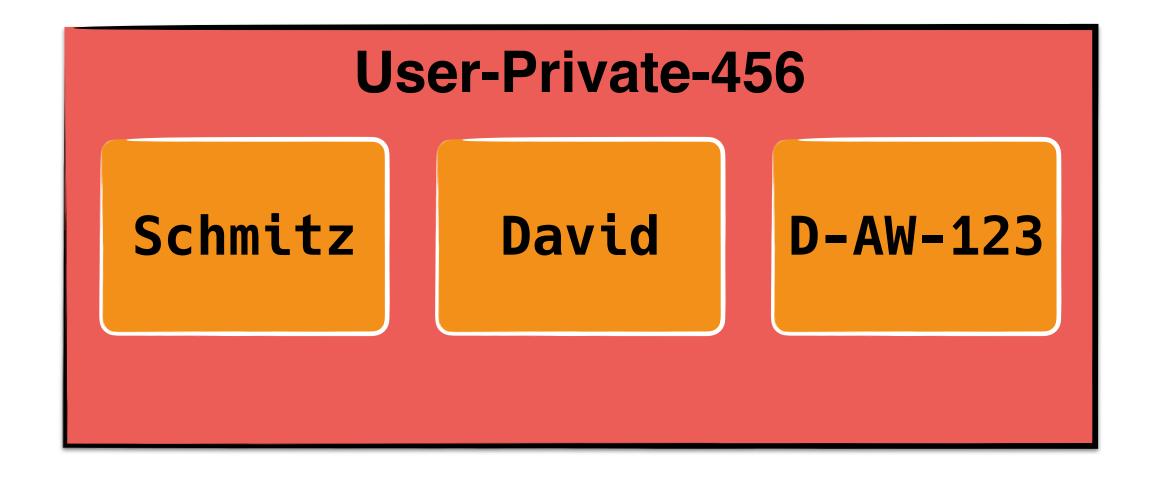


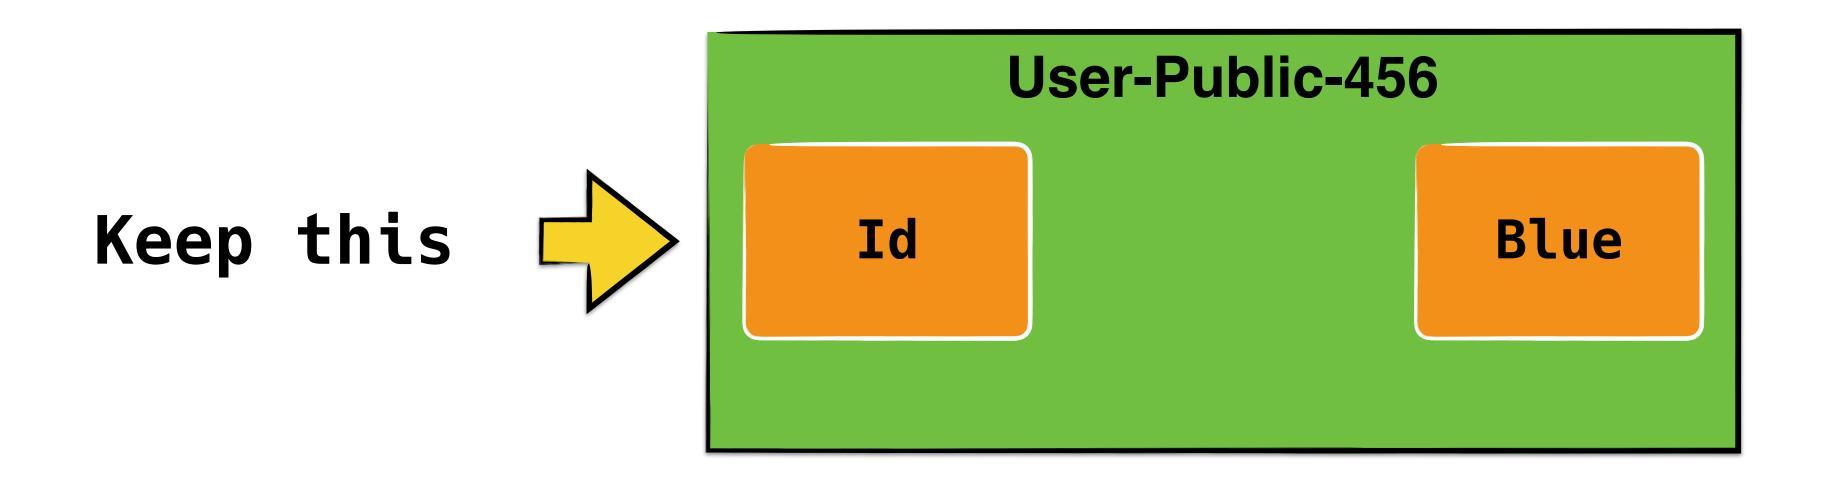
What if I delete a confluence account?

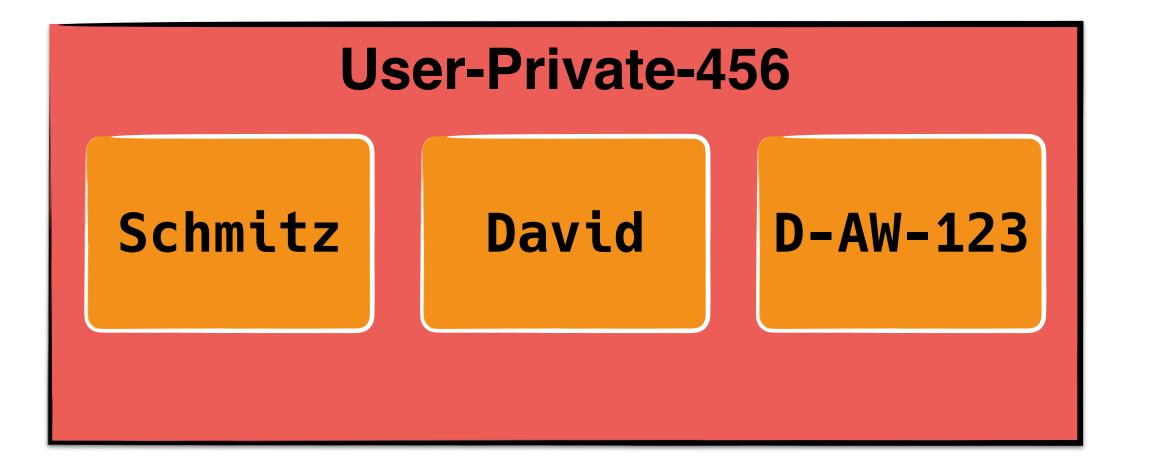


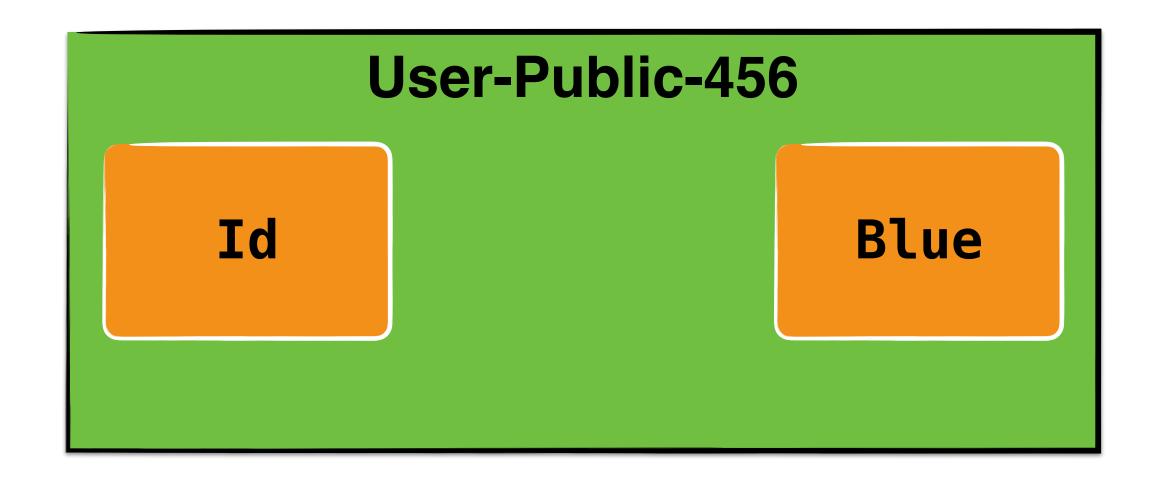
Public/private data

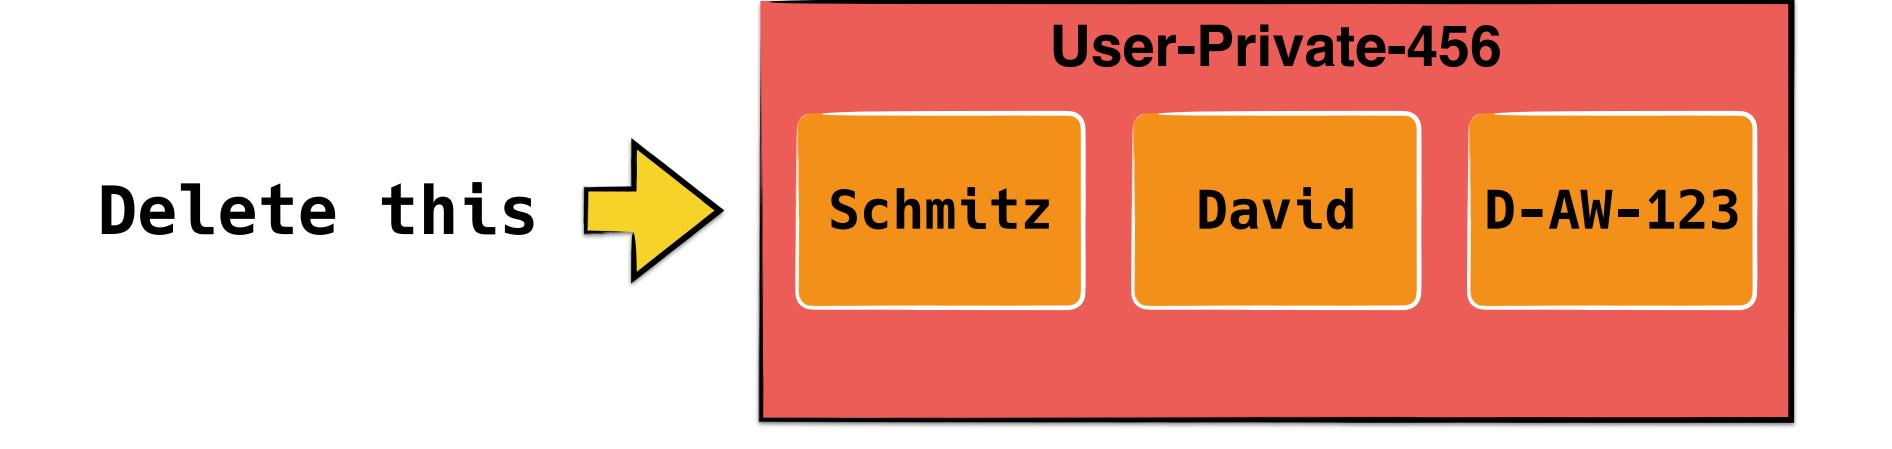












You may be able to keep references to the public data

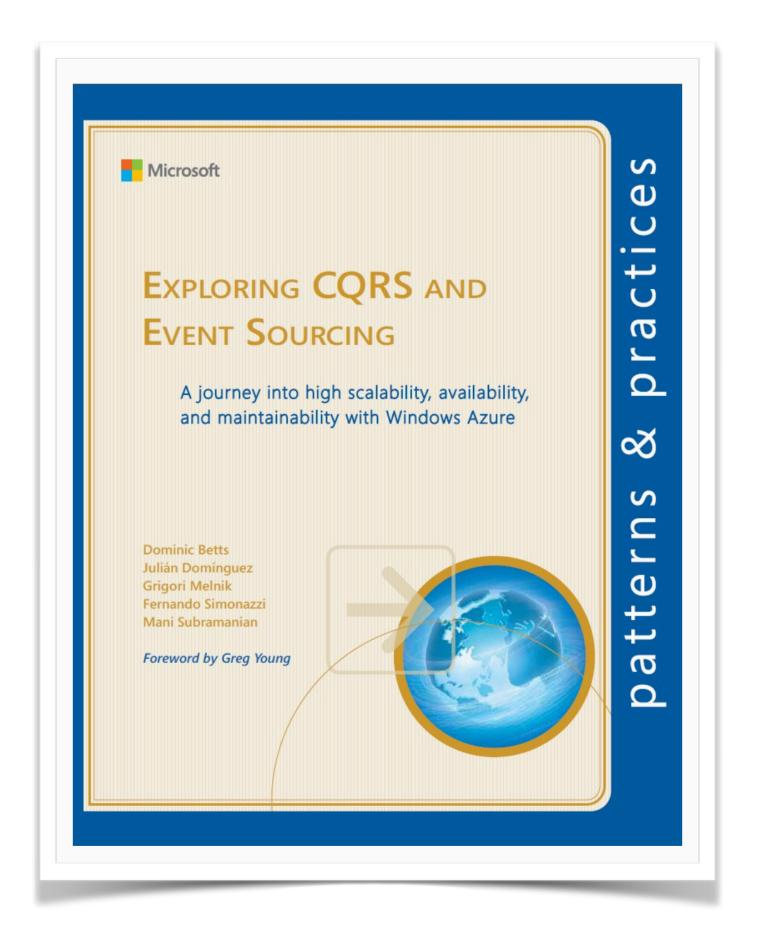
Surprise: No easy answers

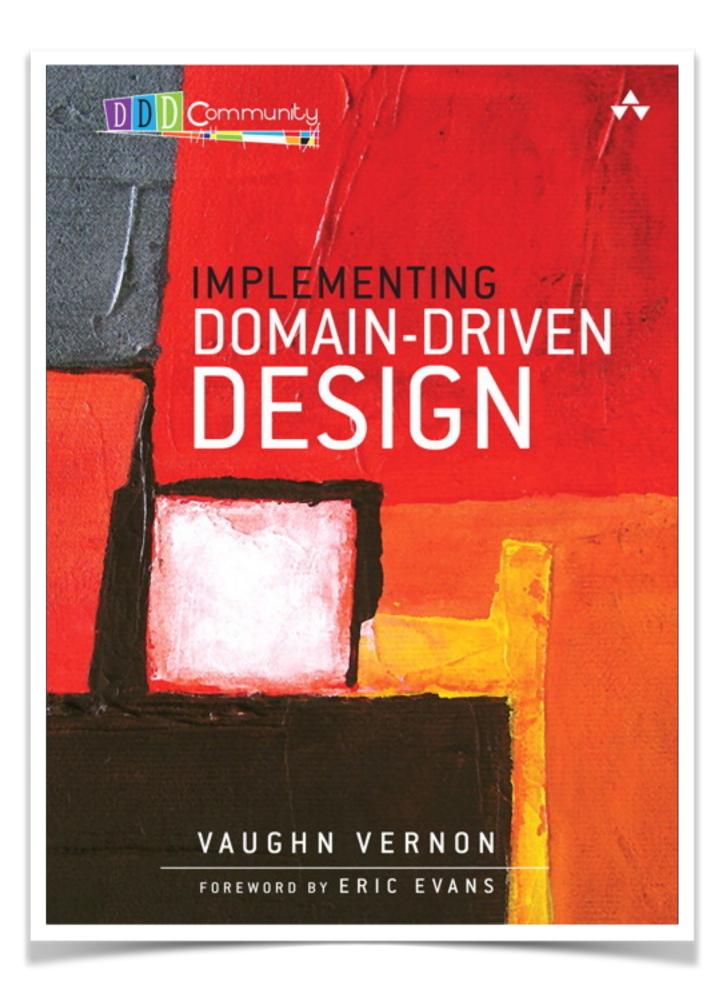
Ask your lawyer or CISO



Microservices + DDD + ES= Needs more up-front design You can refactor, you can clean up Not enough in-depth books Avoid framework lock-in Beware: "just..." or "...made easy"

Forget this talk... read some papers





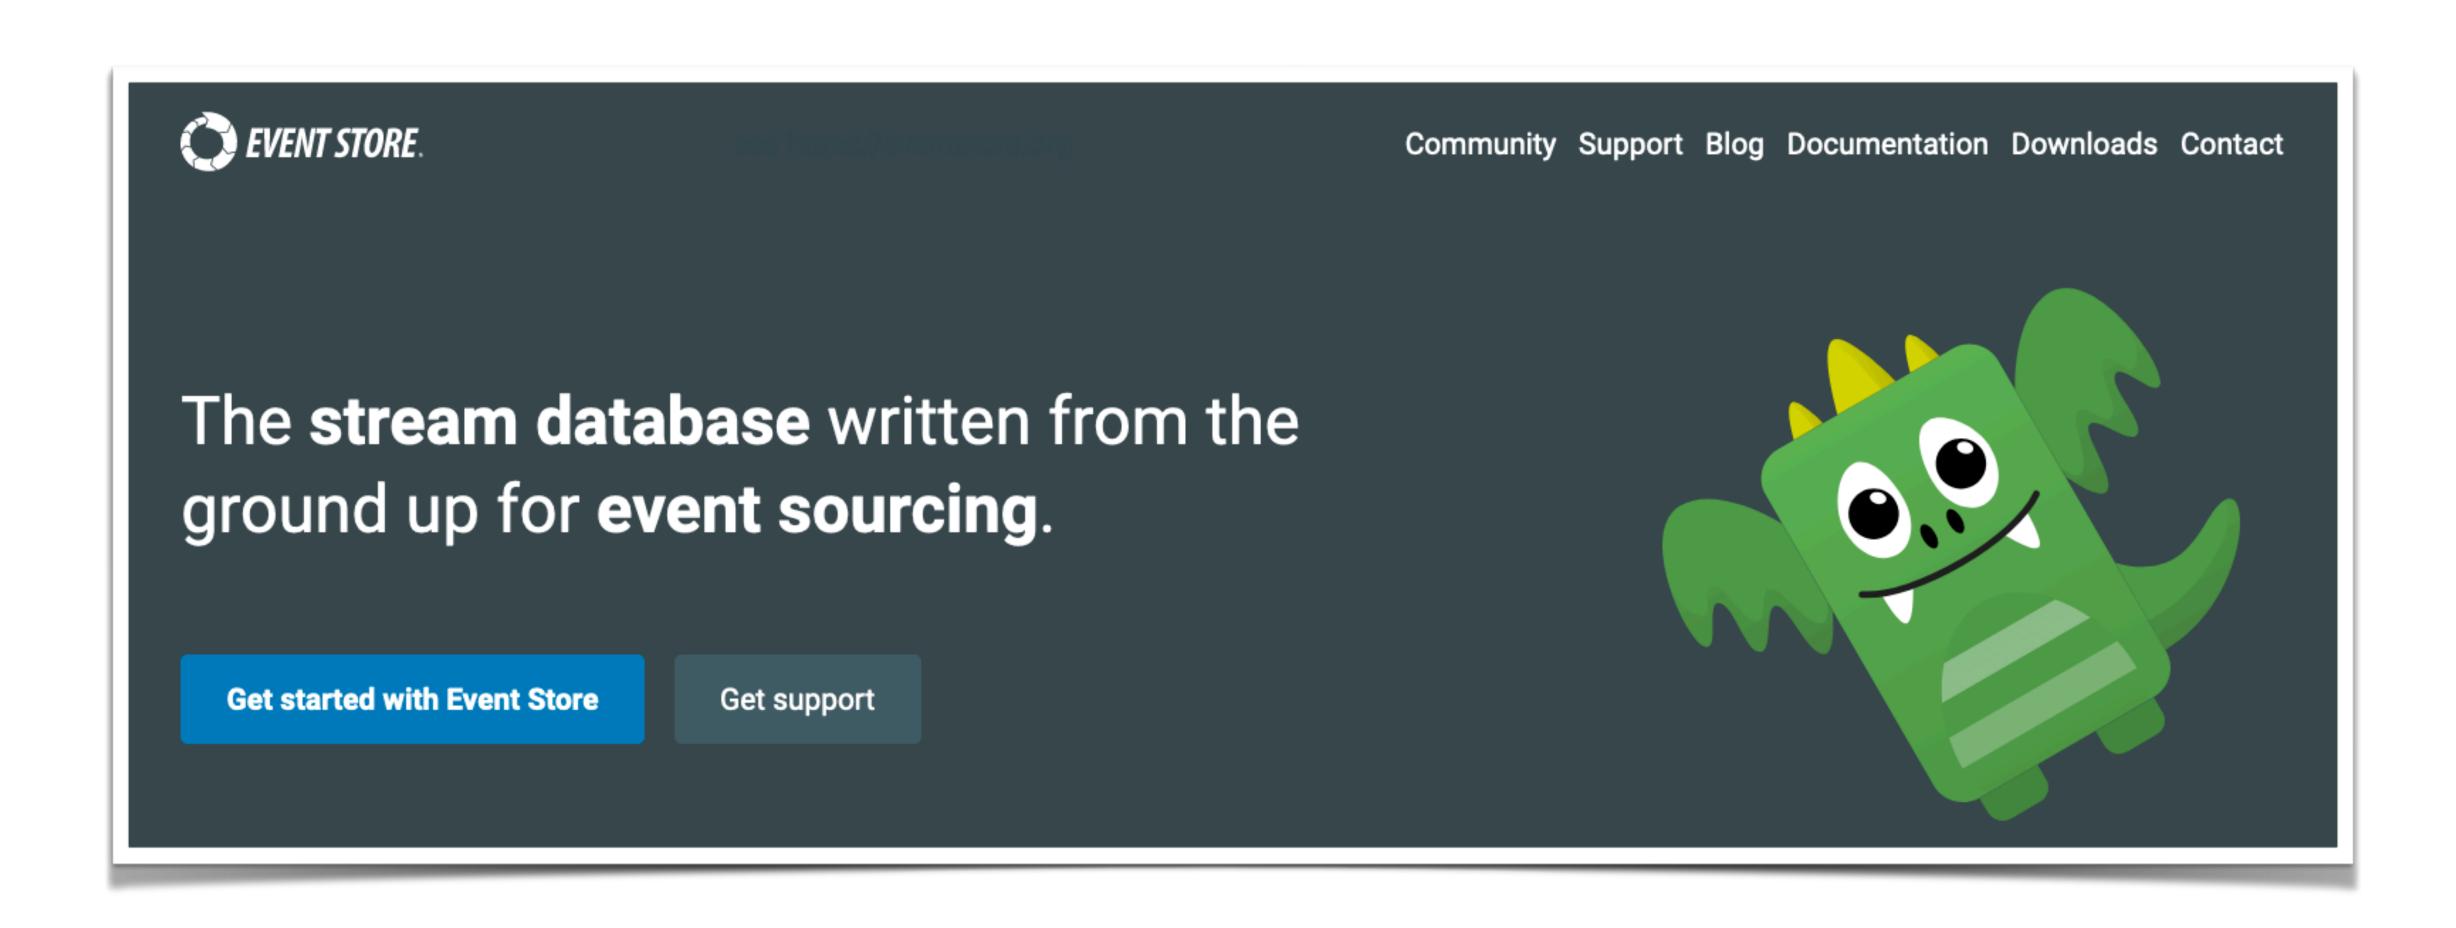
The Dark Side of Event Sourcing: Managing Data Conversion

Michiel Overeem¹, Marten Spoor¹, and Slinger Jansen²

Versioning in an Event Sourced System

Gregory Young

Choose the "right" tool?



eventstore.org

Questions and Feedback?

@Koenighotze

