

# Not Just Events: Developing Asynchronous Microservices

Chris Richardson

Founder of Eventuate.io

Founder of the original CloudFoundry.com

Author of POJOs in Action and Microservices Patterns

 @crichardson

chris@chrisrichardson.net

<http://learn.microservices.io>

**GOTO**  
CHICAGO 2019



# Presentation goal

Implementing transactions and  
queries in a microservice  
architecture using  
asynchronous messaging



# Presentation goal

Microservices > REST

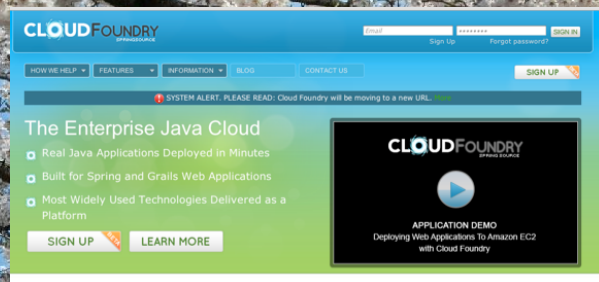


Microservices > Events

Microservices  $\neq$  Event Sourcing

Apache Kafka  $\neq$  Event Store



# About Chris





# About Chris

Consultant and trainer  
focussed on helping  
organizations adopt the  
microservice architecture  
(<http://www.chrisrichardson.net/>)



# About Chris


Founder of a startup that is creating  
an open-source/SaaS platform  
that simplifies the development of  
transactional microservices

(<http://eventuate.io>)





# About Chris

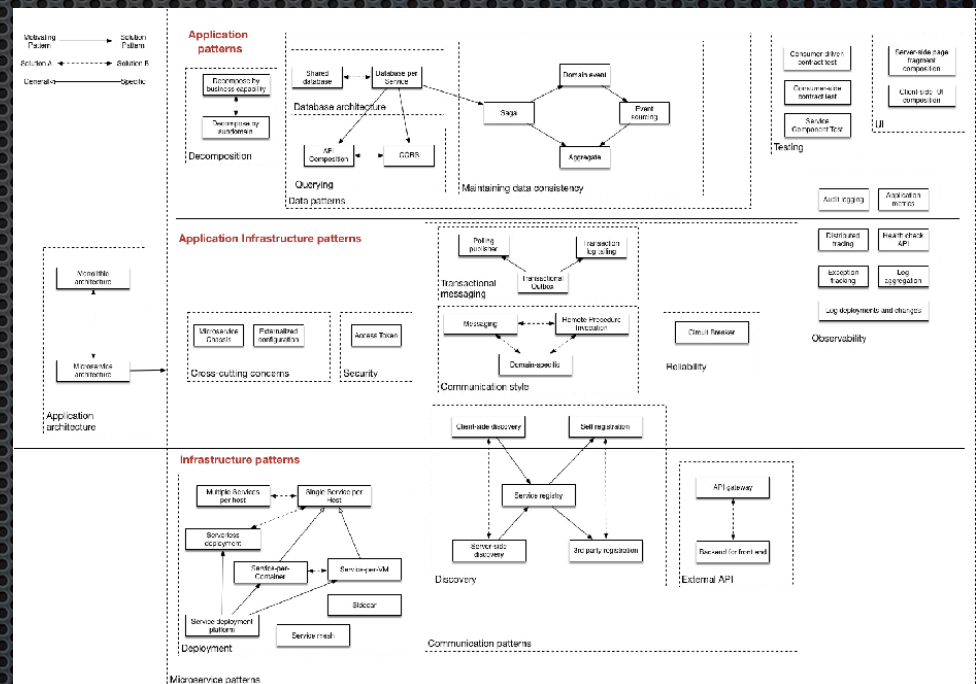


Get a 40% discount:  
<http://bit.ly/gotochgo2019-microservices>



# About Chris: microservices.io

- ✦ Microservices pattern language
- ✦ Articles
- ✦ Code examples
- ✦ Microservices Assessment Platform - <http://microservices.io/platform>





# Agenda

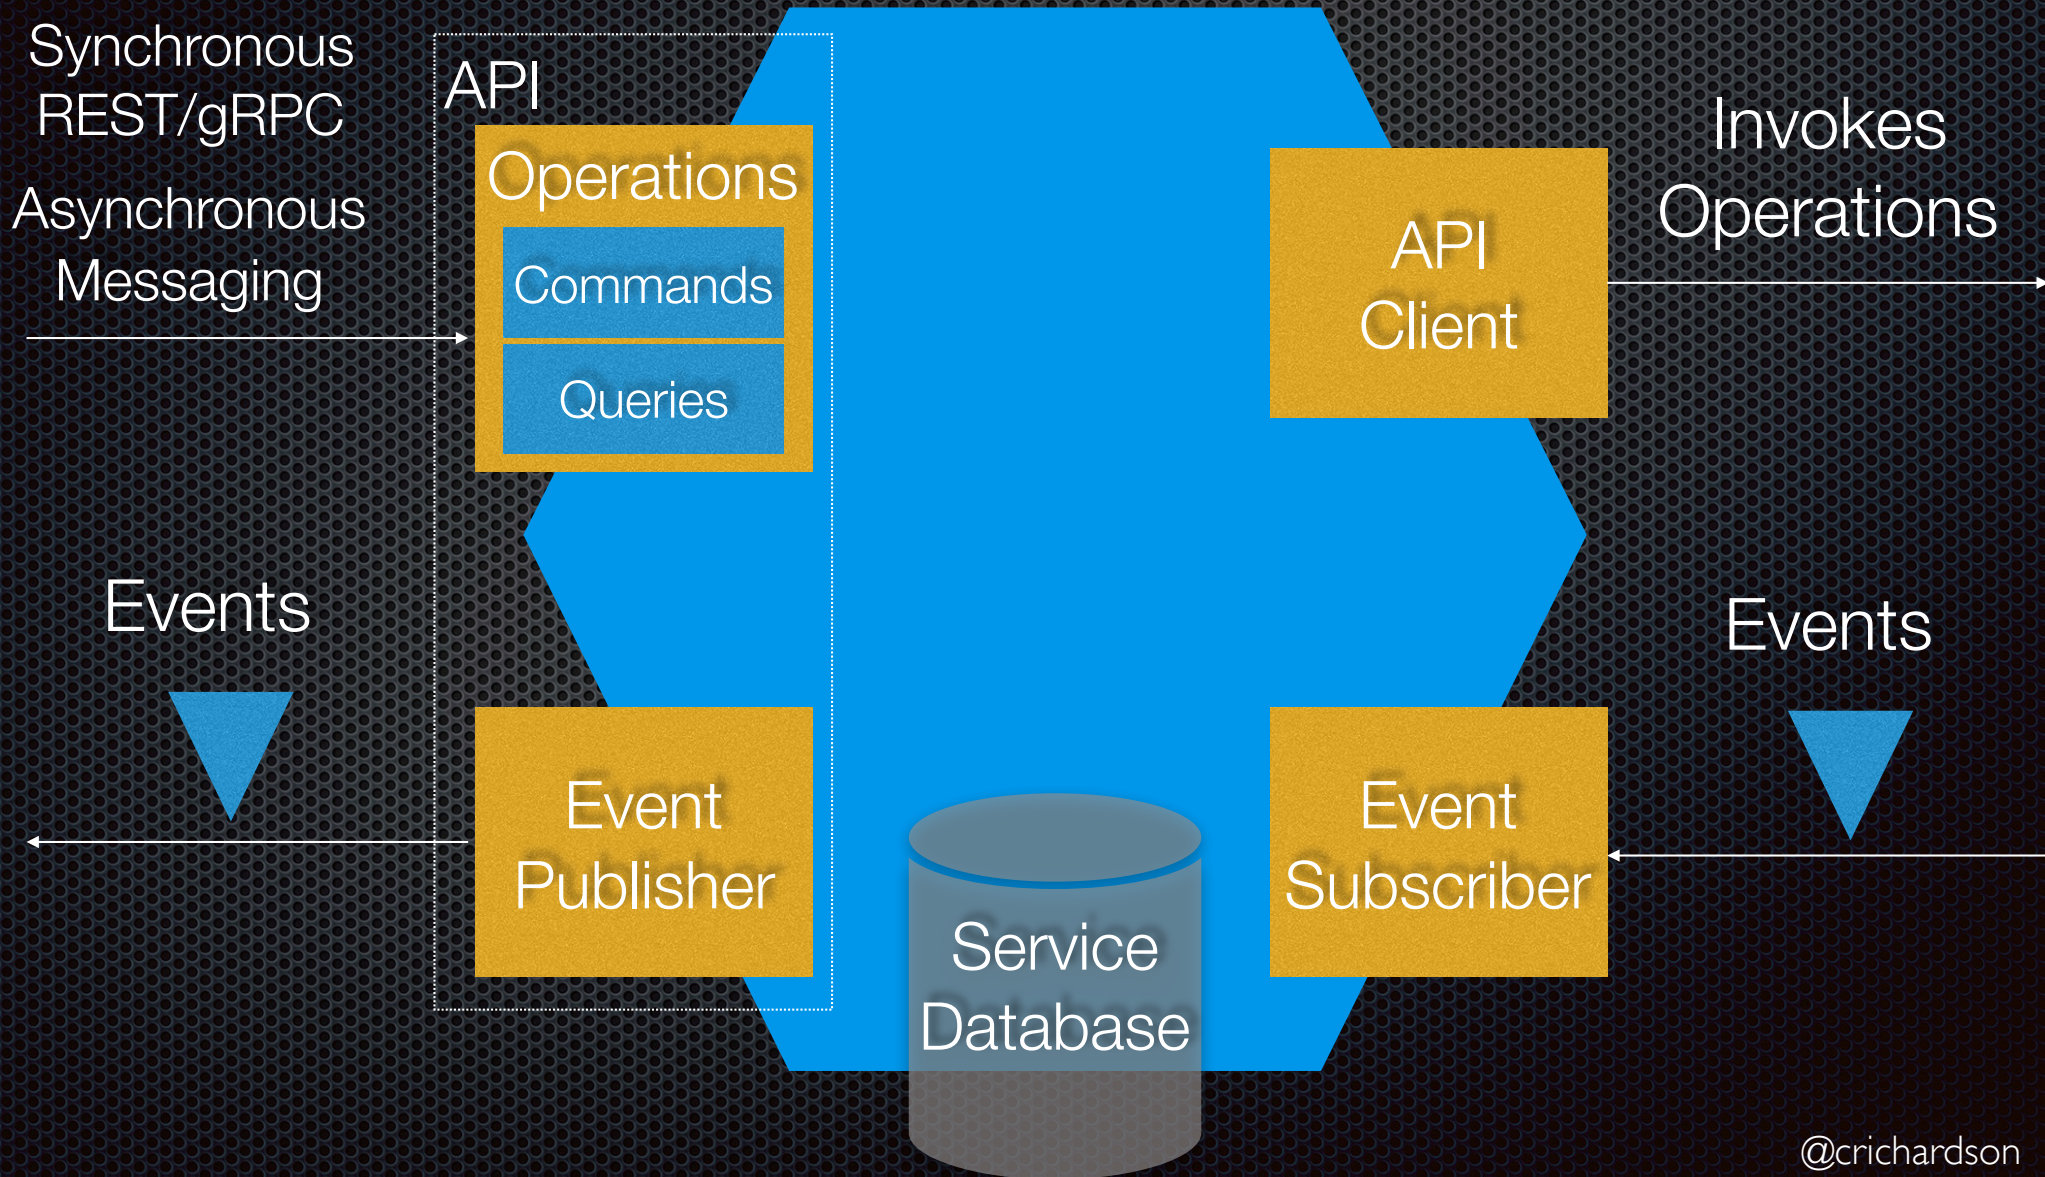
- ✦ Transactions, queries and microservices
- ✦ Managing transactions with sagas
- ✦ Implementing queries with CQRS
- ✦ Implementing transactional messaging



The microservice architecture  
**structures**  
an application as a  
**set of loosely coupled**  
**services**

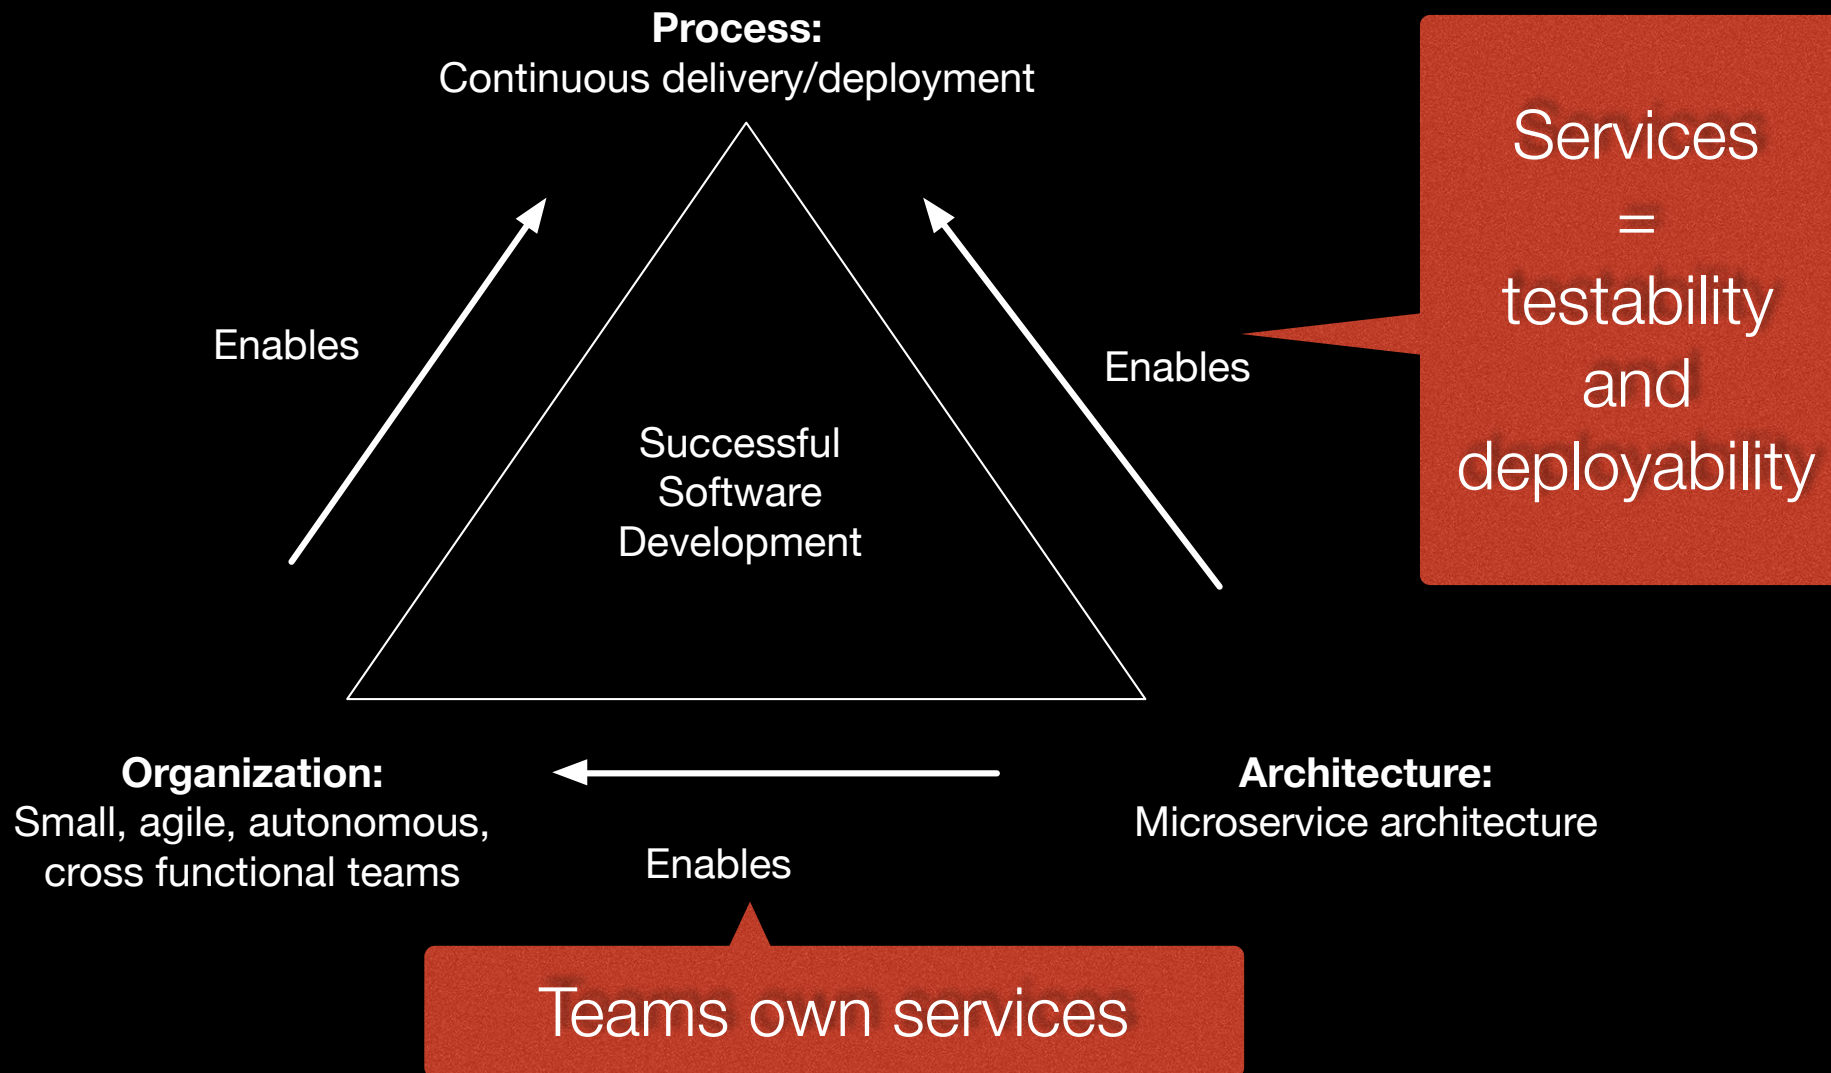


# Service = independently deployable component





# Microservices enable continuous delivery/deployment





# Let's imagine that you are building an online store API

`createCustomer(creditLimit)`

`createOrder(customerId, orderTotal)`

`findOrdersForCustomer(customerId)`

`findRecentCustomers()`

●  
REST API

Customer  
Management

Order  
Management

...



Retrieve data from both services 🤔

createCustomer()  
createOrder()  
findOrdersForCustomer()  
findRecentCustomers()

REST API

Must reserve customer's credit 🤔

API Gateway

createCustomer()

REST API

Customer Service

createOrder()

REST API

Order Service

Essential for loose coupling

Customer Database

Customer table

availableCredit

...

Order Database

Order table

orderTotal

...



# No ACID transactions that span services

Distributed transactions

BEGIN TRANSACTION

...

SELECT ORDER\_TOTAL  
FROM **ORDERS** WHERE CUSTOMER\_ID = ?

Private to the  
Order Service

...

SELECT CREDIT\_LIMIT  
FROM **CUSTOMERS** WHERE CUSTOMER\_ID = ?

...

INSERT INTO ORDERS ...

Private to the  
Customer Service

...

COMMIT TRANSACTION



# Querying across services is not straightforward

Private to Customer Service

Private to Order Service

```
SELECT *  
FROM CUSTOMER c, ORDER o  
WHERE  
    c.id = o.ID  
    AND c.id = ?
```

Find customer and their orders



# Agenda

- ✦ Transactions, queries and microservices
- ✦ Managing transactions with sagas
- ✦ Implementing queries with CQRS
- ✦ Implementing transactional messaging



From a 1987 paper


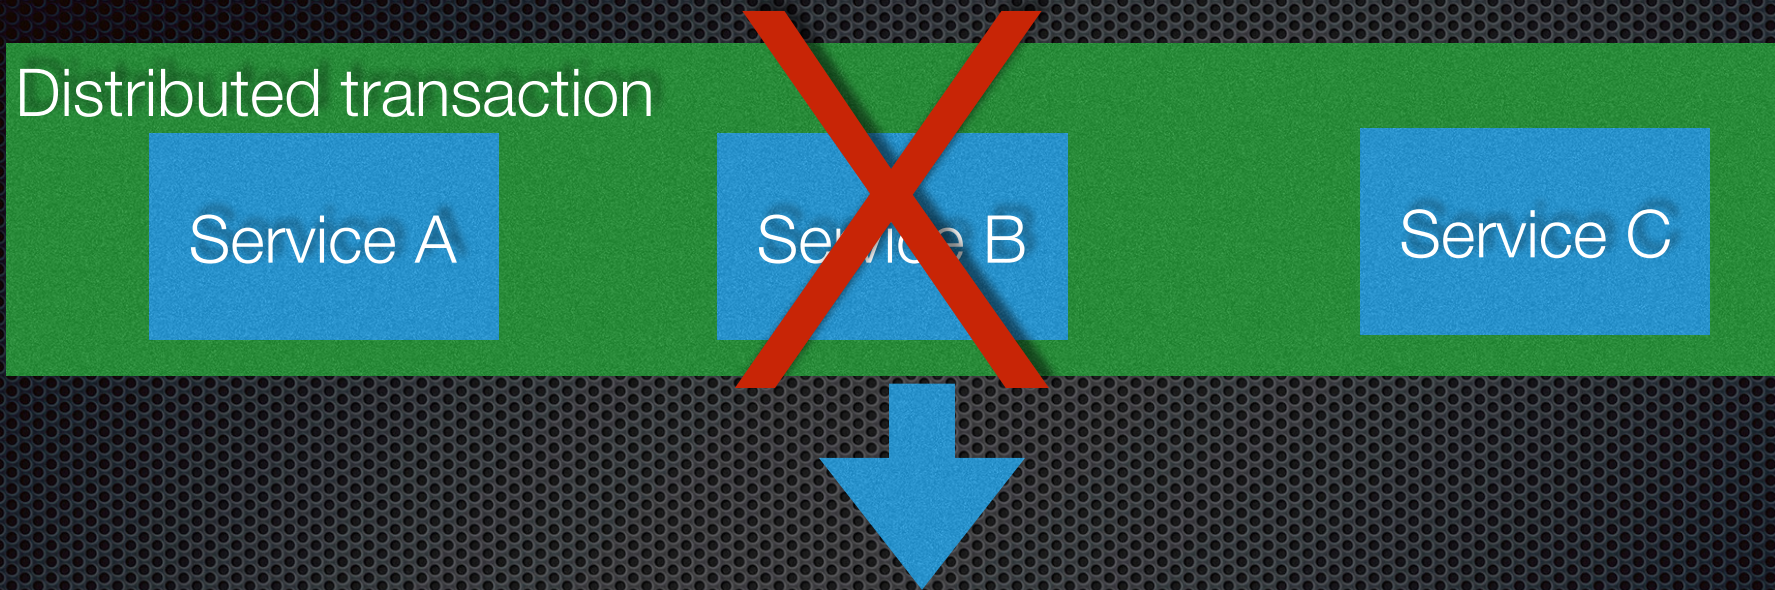
## **SAGAS**

*Hector Garcia-Molina*  
*Kenneth Salem*

**Department of Computer Science**  
**Princeton University**  
**Princeton, N J 08544**



# Use Sagas instead of 2PC





# Create Order Saga

createOrder()

Initiates saga

Order Service

Local transaction

createOrder()

Order

state=PENDING

Customer Service

Local transaction

reserveCredit()

Customer


Order Service

Local transaction

approve  
order()

Order

state=APPROVED






# Saga design challenges

- ✦ API design
  - ✦ Synchronous REST API initiates asynchronous saga
  - ✦ When to send back a response?
- ✦ Rollback  $\Rightarrow$  compensating transactions
- ✦ Sagas are ACD - No I
  - ✦ Sagas are interleaved  $\Rightarrow$  anomalies, such as lost updates
  - ✦ Must use countermeasures




# How do the saga participants communicate?

- Synchronous communication, e.g. REST = temporal coupling
- Client and server need to be both available
- Customer Service fails  $\Rightarrow$  retry provided it's idempotent
- Order Service fails  $\Rightarrow$  Oops





# Collaboration using asynchronous, broker-based messaging



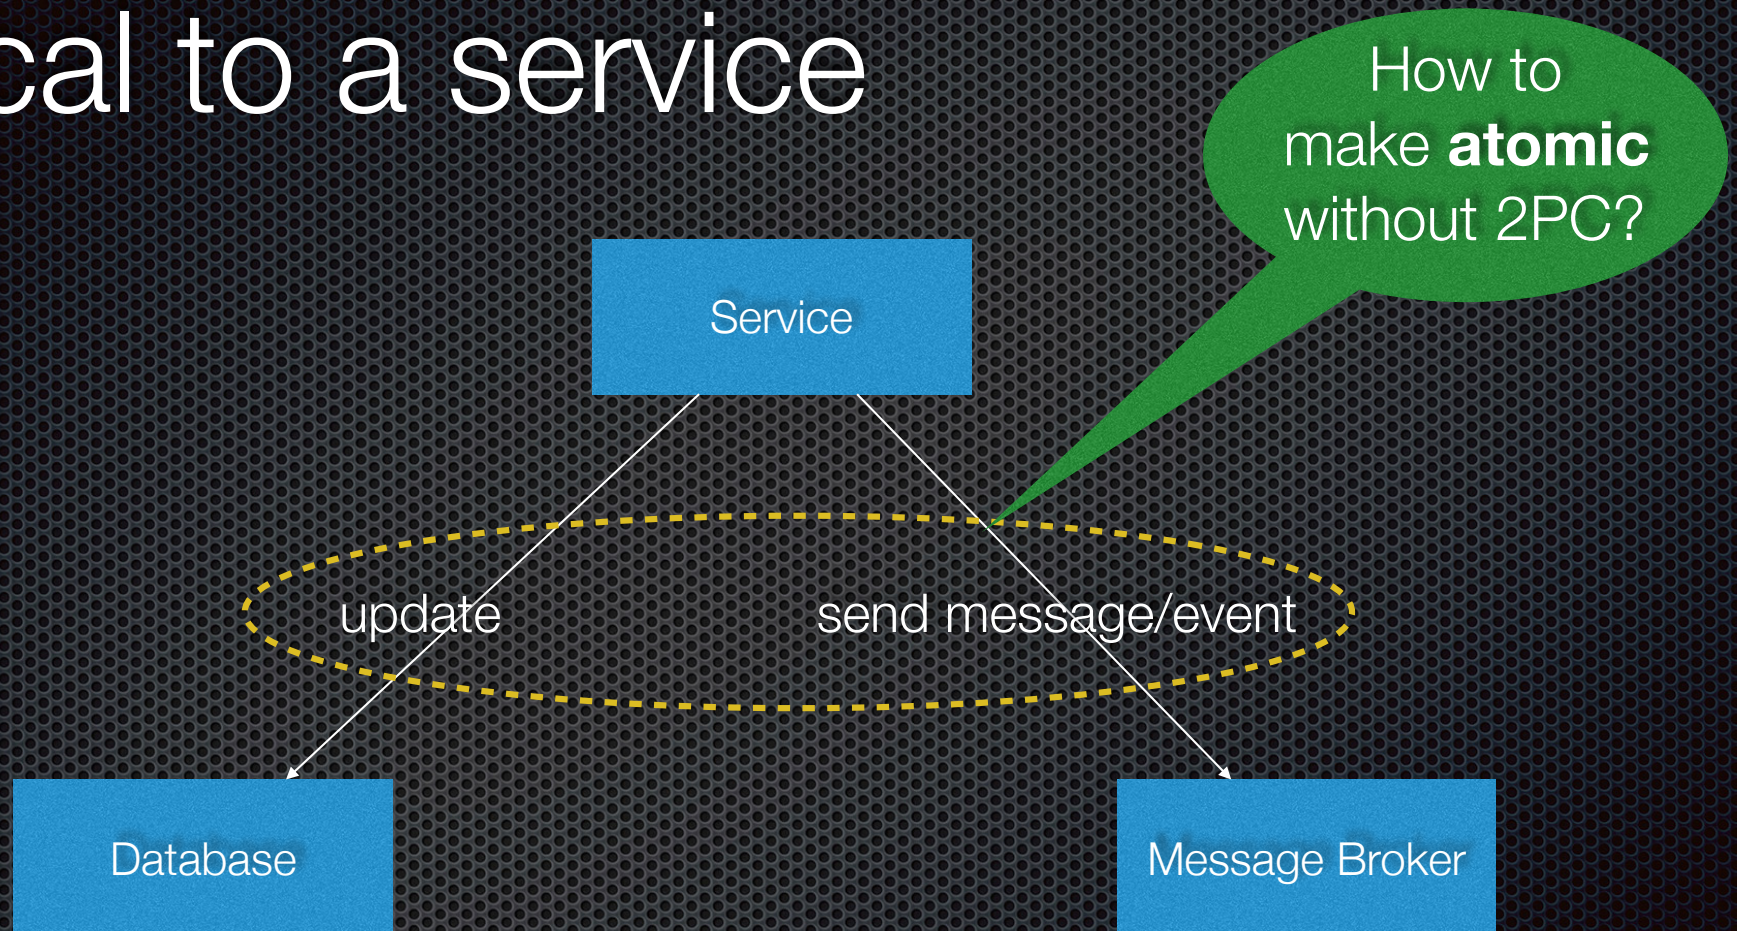


# About the message broker

- At least once delivery
  - Ensures a saga completes when its participants are temporarily unavailable
- Ordered delivery
- Mechanism for scaling consumers that preserves ordering e.g.
  - Apache Kafka consumer group
  - ActiveMQ message group
  - ...



# Saga step = a transaction local to a service





# How to sequence the saga transactions?

- After the completion of transaction  $T_i$  “something” must decide what step to execute next
- Success: which  $T_{i+1}$  - branching
- Failure:  $C(i - 1)$



Choreography: **distributed** decision making

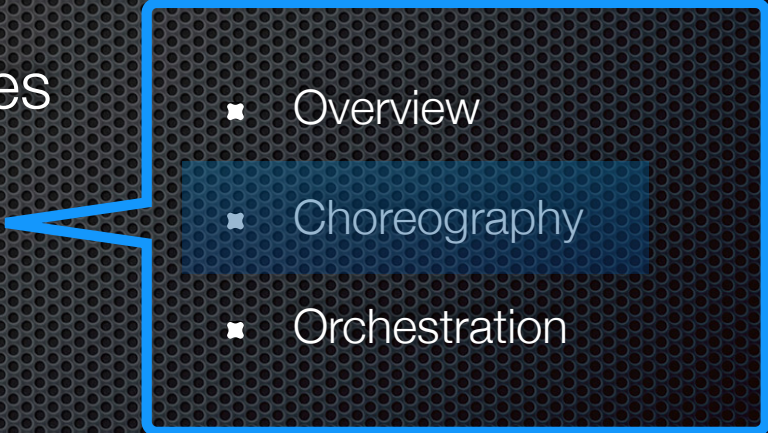
VS.

Orchestration: **centralized** decision making



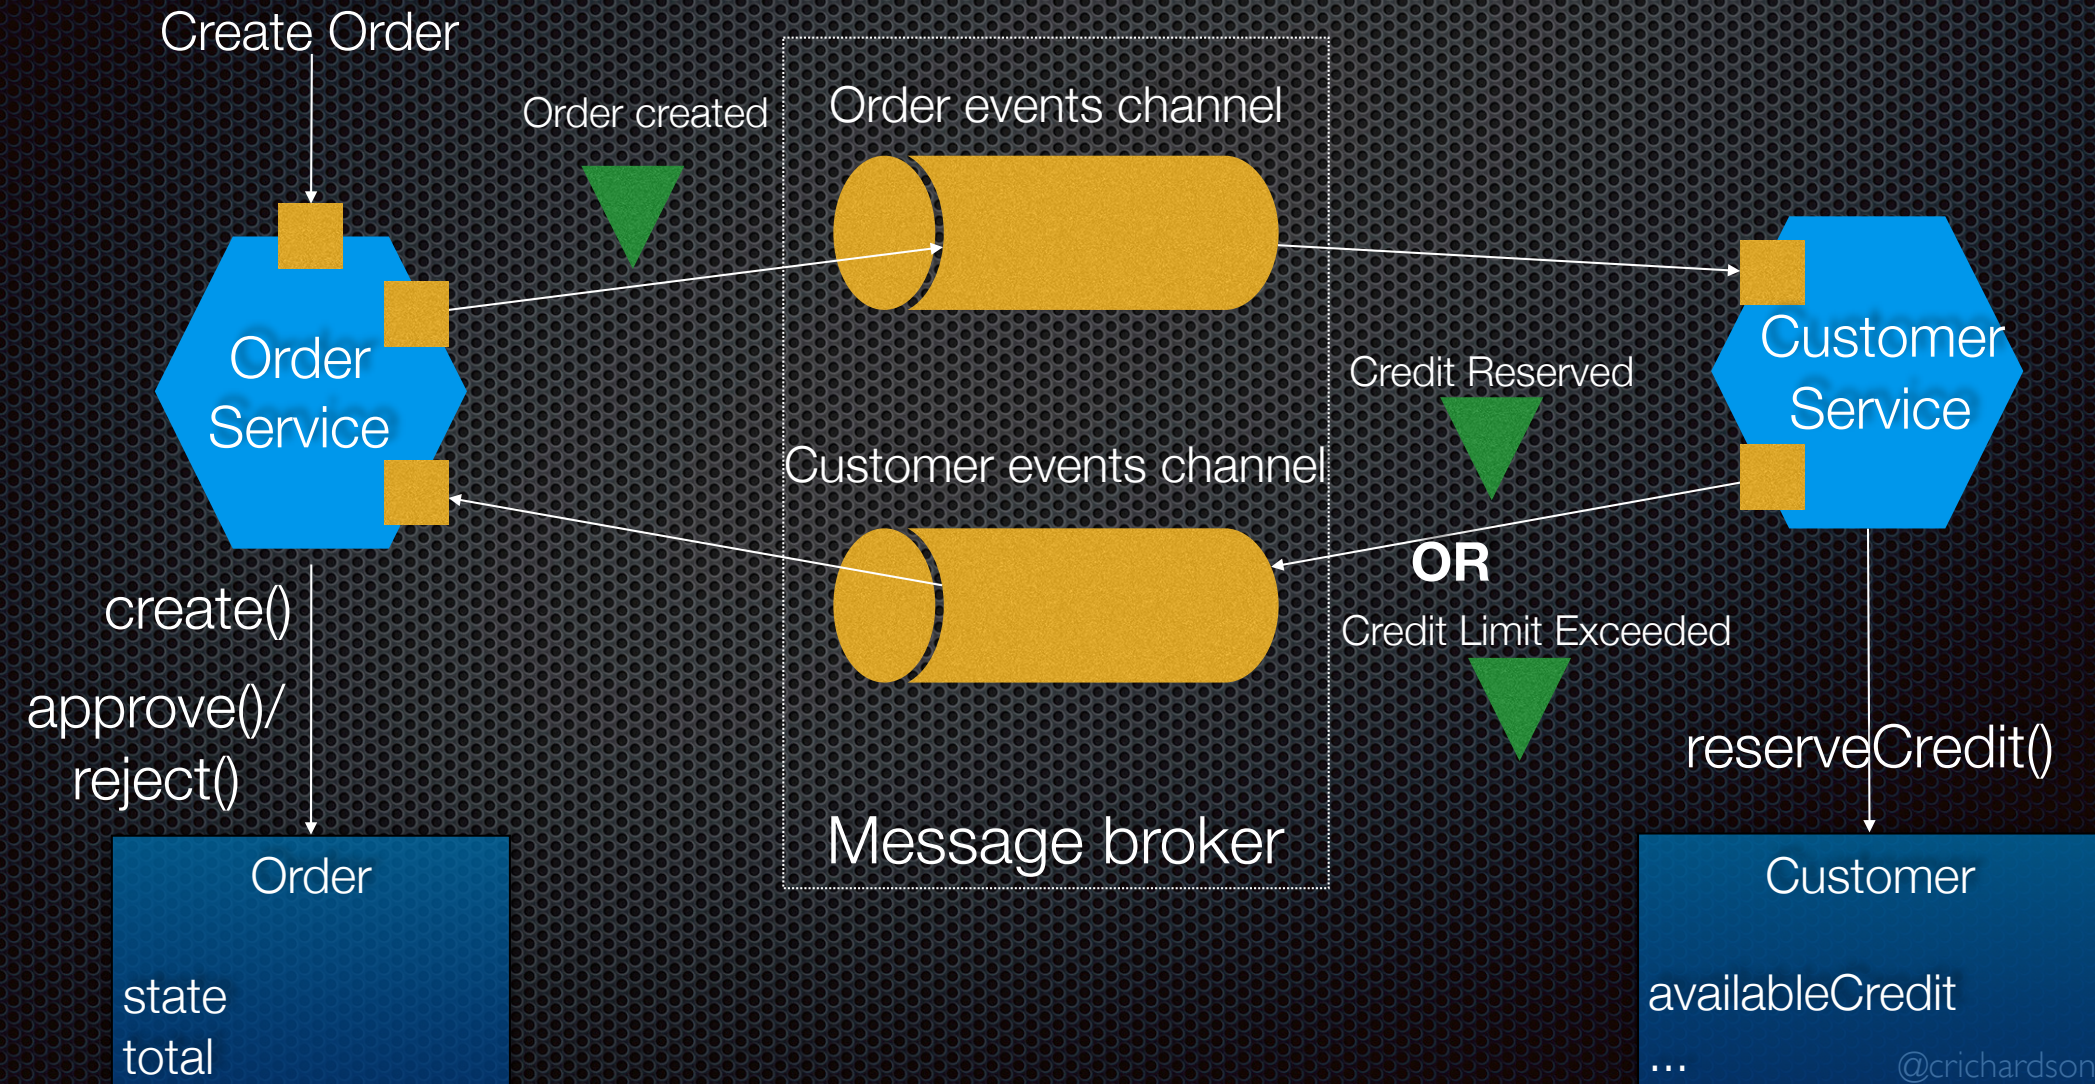
# Agenda

- Transactions, queries and microservices
- Managing transactions with sagas
- Implementing queries with CQRS
- Implementing transactional messaging

- 
- Overview
  - Choreography
  - Orchestration



# Choreography-based Create Order Saga





# Benefits and drawbacks of choreography

## Benefits

- Simple, especially when using event sourcing
- Participants are loosely coupled

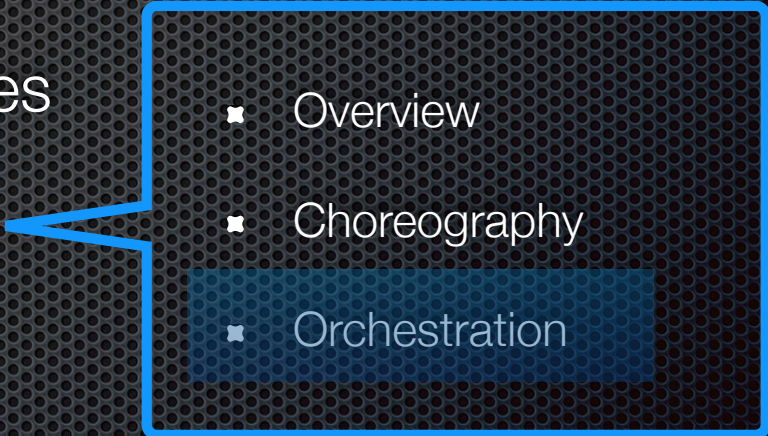
## Drawbacks

- Decentralized implementation - potentially difficult to understand
- Cyclic dependencies - services listen to each other's events, e.g. Customer Service **must know** about all Order events that affect credit
- Overloads domain objects, e.g. Order and Customer **know** too much
- Events = indirect way to make something happen



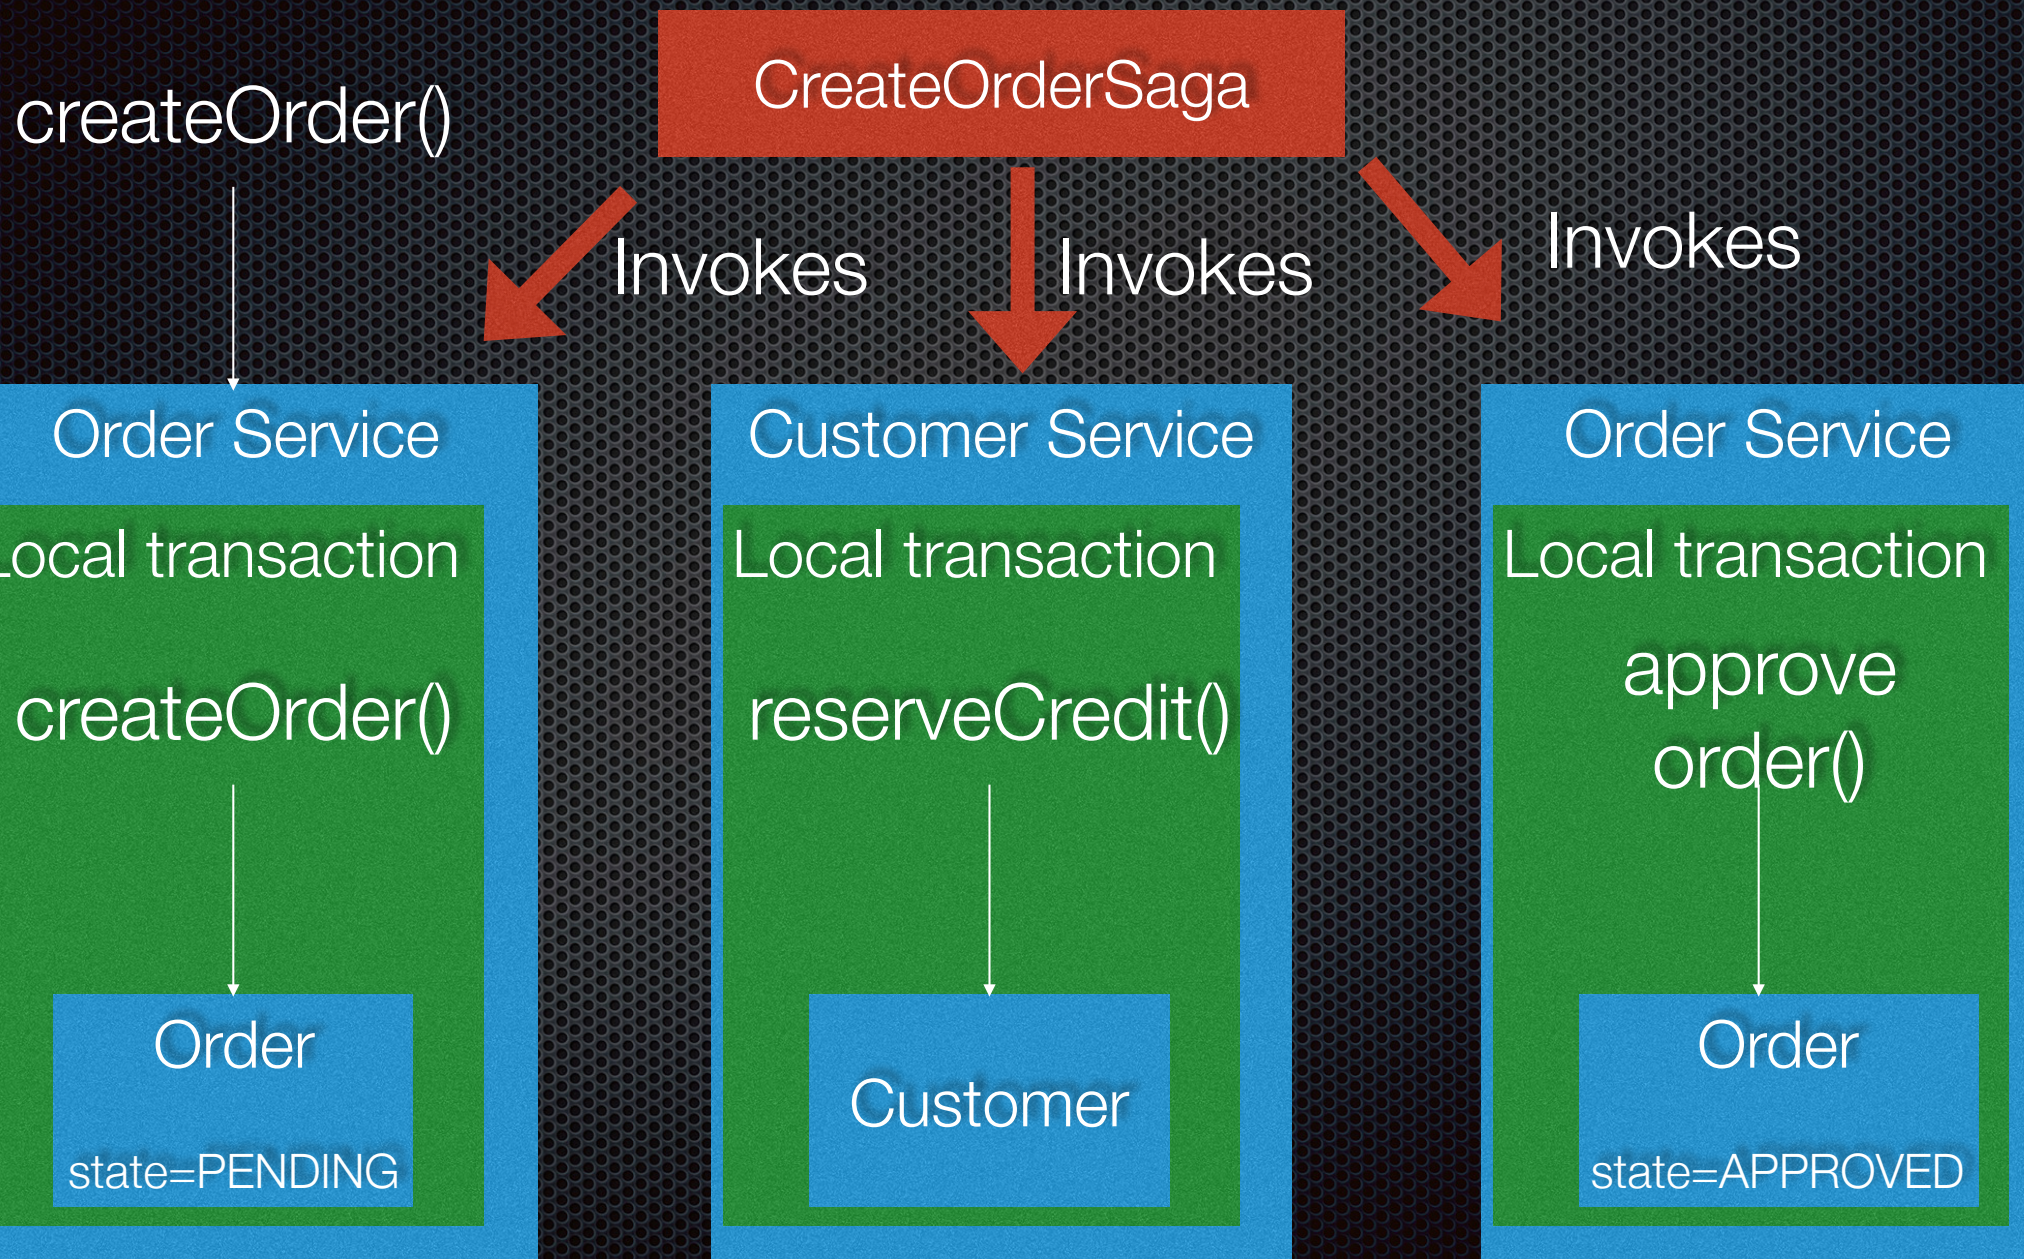
# Agenda

- Transactions, queries and microservices
- Managing transactions with sagas
- Implementing queries with CQRS
- Implementing transactional messaging

- 
- Overview
  - Choreography
  - Orchestration



# Orchestration-based coordination using command messages





A saga (orchestrator)  
is a **persistent object**  
that  
implements a state machine  
and  
**invokes** the participants

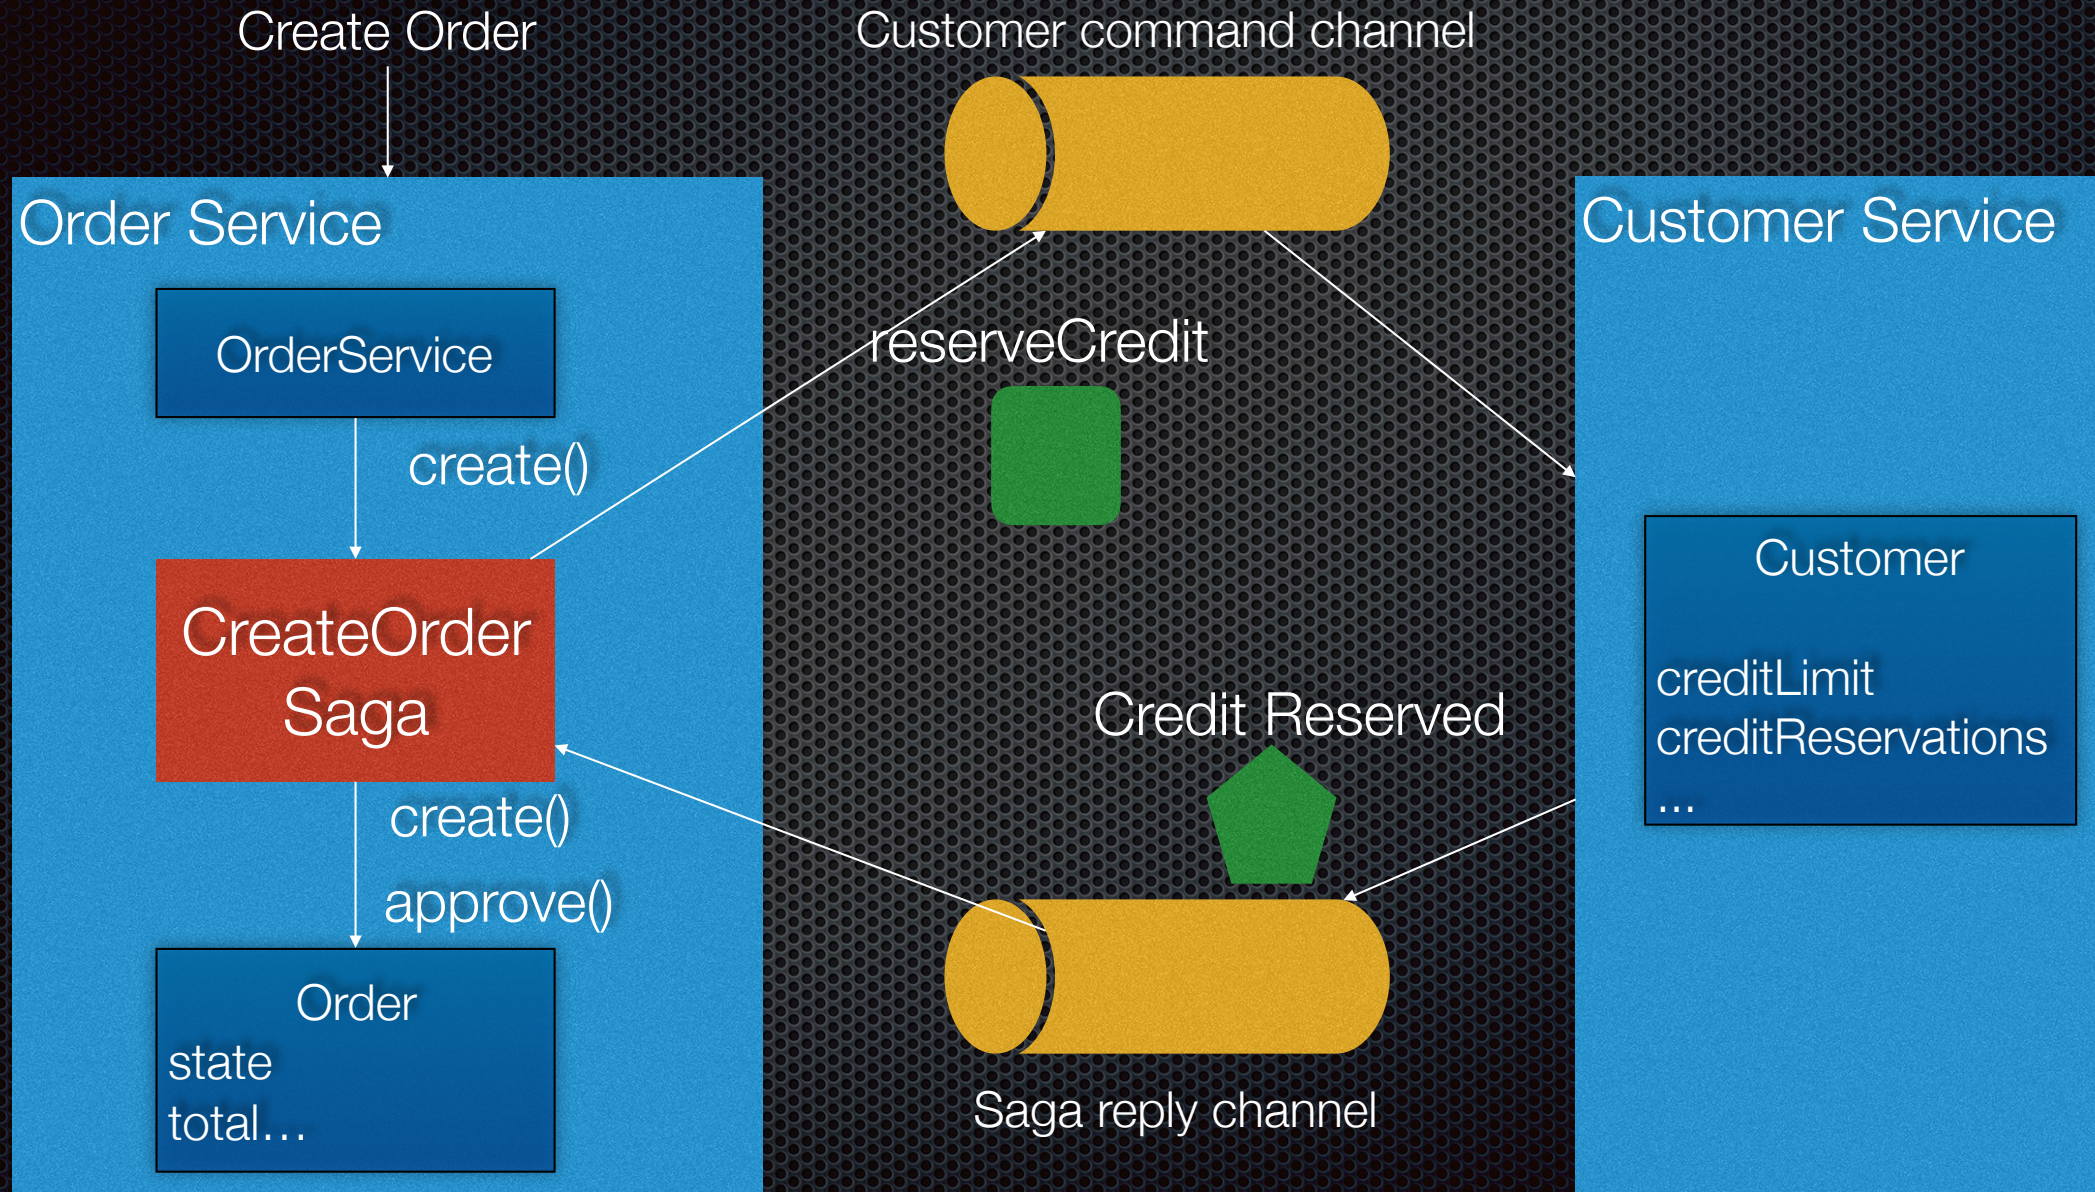


# Saga orchestrator behavior

- On create:
  - Invokes a saga participant
  - Persists state in database
  - Wait for a reply
- On reply:
  - Load state from database
  - Determine which saga participant to invoke next
  - Invokes saga participant
  - Updates its state
  - Persists updated state
  - Wait for a reply
  - ...



# CreateOrderSaga orchestrator





# Benefits and drawbacks of orchestration

## Benefits

- ✦ Centralized coordination logic is easier to understand
- ✦ Reduced coupling, e.g. Customer Service knows less. Simply has API for managing available credit.
- ✦ Reduces cyclic dependencies

## Drawbacks

- ✦ Risk of smart sagas directing dumb services



# Agenda

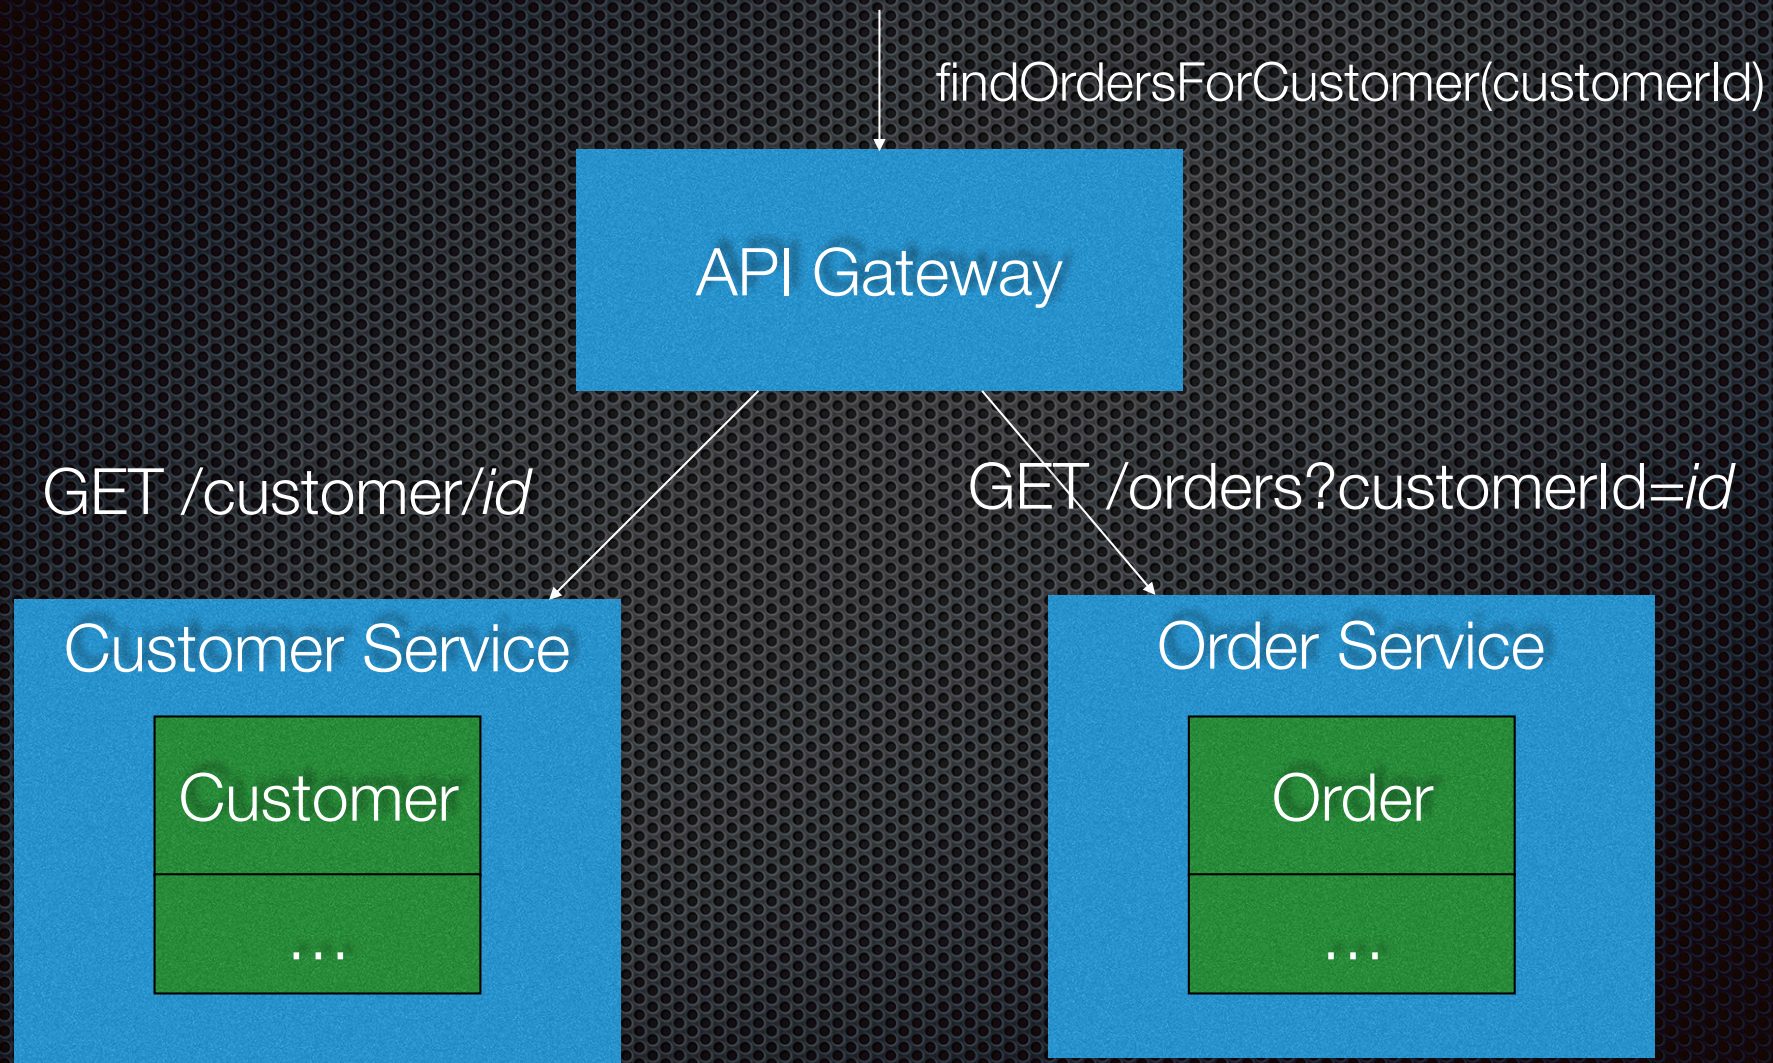
- ✦ Transactions, queries and microservices
- ✦ Managing transactions with sagas
- ✦ Implementing queries with CQRS
- ✦ Implementing transactional messaging



Queries often retrieve data  
owned by multiple services



# API Composition pattern





# Find recent, valuable customers

Customer  
Service

Order Service

```
SELECT *  
FROM CUSTOMER c, ORDER o  
WHERE  
    c.id = o.ID  
    AND o.ORDER_TOTAL > 100000  
    AND o.STATE = 'SHIPPED'  
    AND c.CREATION_DATE > ?
```

Not efficiently implemented using API Composition

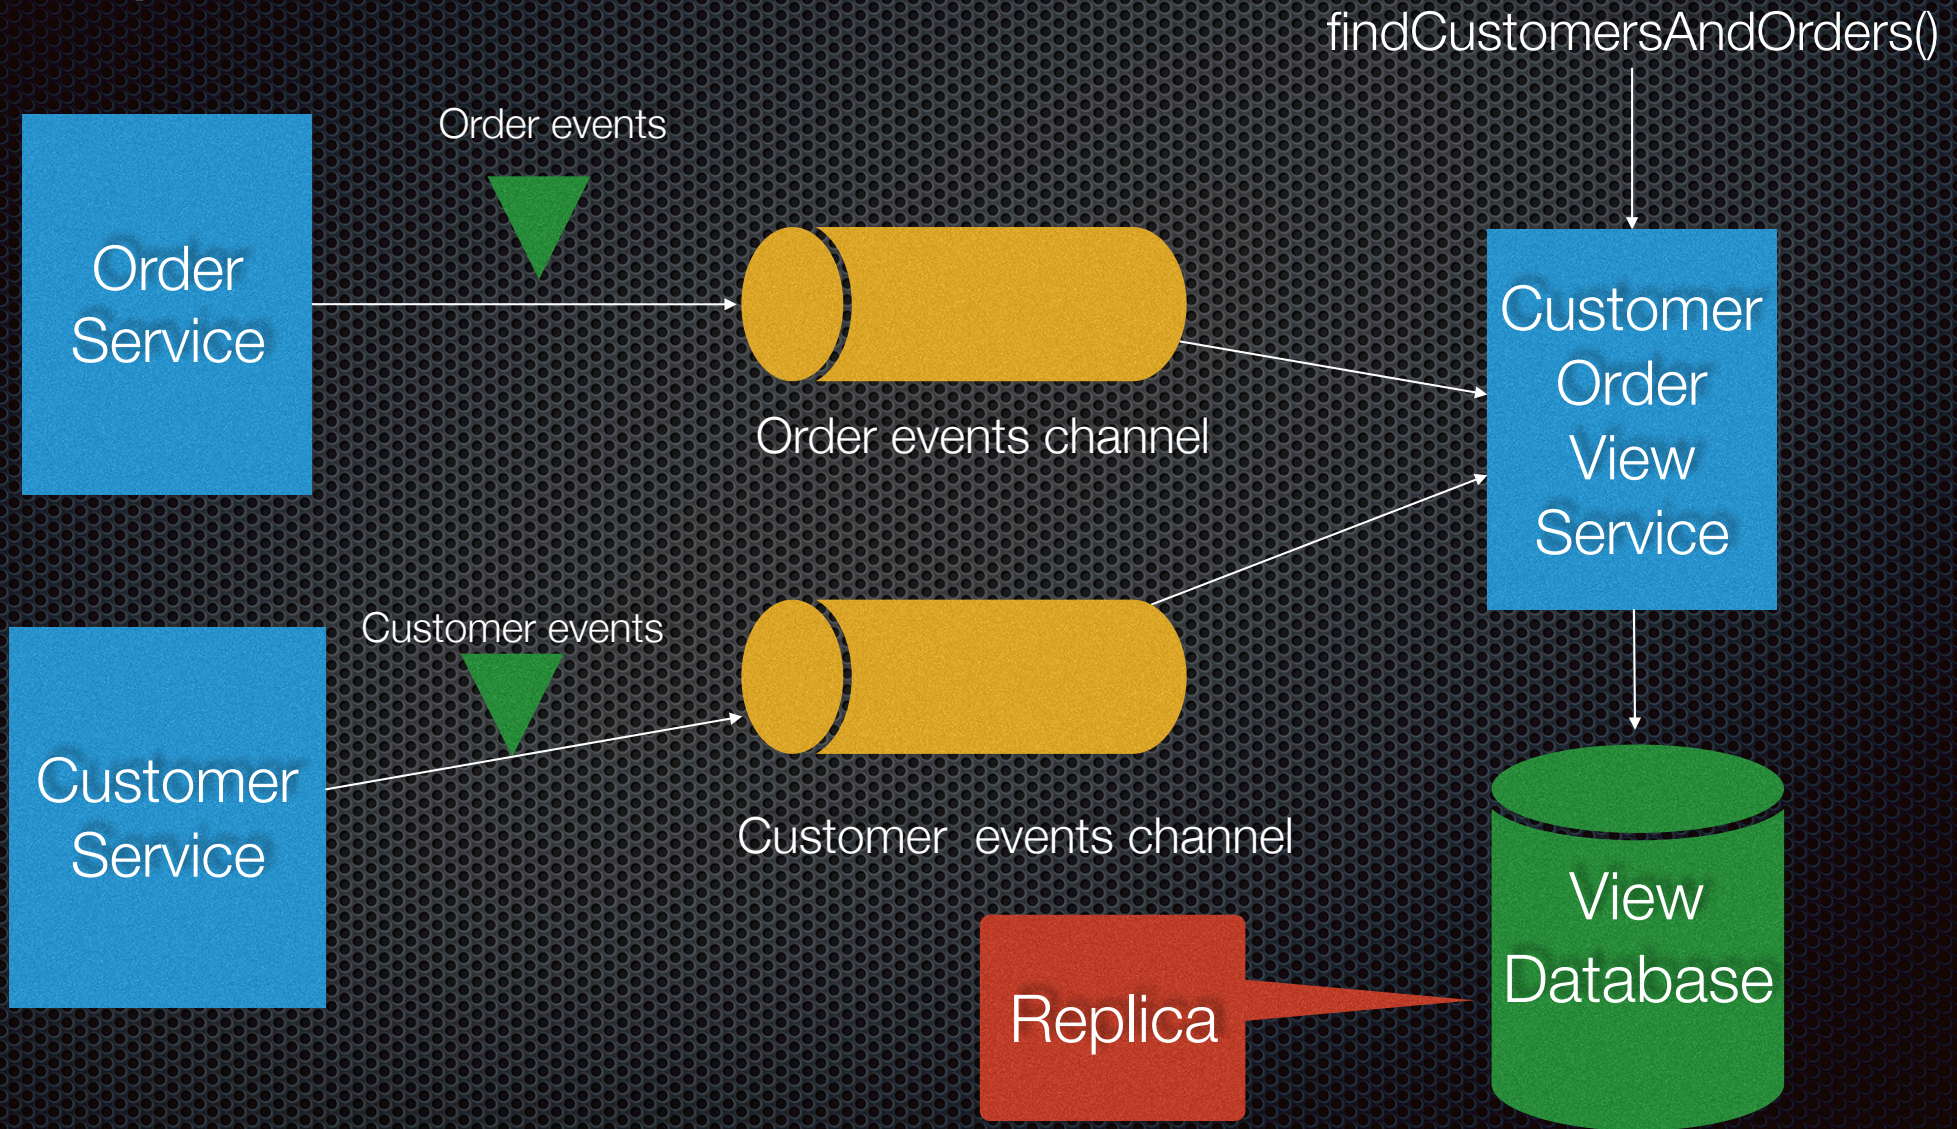


# API Composition would be inefficient

- ✦ 1 + N strategy:
  - ✦ Fetch recent customers
  - ✦ Iterate through customers fetching their shipped orders
  - ✦ Lots of round trips  $\Rightarrow$  high-latency
- ✦ Alternative strategy:
  - ✦ Fetch recent customers
  - ✦ Fetch recent orders
  - ✦ Join
  - ✦ 2 roundtrips but potentially large datasets  $\Rightarrow$  inefficient



# Using events to update a queryable replica = CQRS



<https://microservices.io/patterns/data/cqrs.html>



# Persisting a customer and order history in MongoDB

```
{
  "_id" : "0000014f9a45004b 0a00270000000000",
  "name" : "Fred",
  "creditLimit" : {
    "amount" : "2000"
  },
  "orders" : {
    "0000014f9a450063 0a00270000000000" : {
      "state" : "APPROVED",
      "orderId" : "0000014f9a450063 0a00270000000000",
      "orderTotal" : {
        "amount" : "1234"
      }
    },
    "0000014f9a450063 0a00270000000001" : {
      "state" : "REJECTED",
      "orderId" : "0000014f9a450063 0a00270000000001",
      "orderTotal" : {
        "amount" : "3000"
      }
    }
  }
}
```

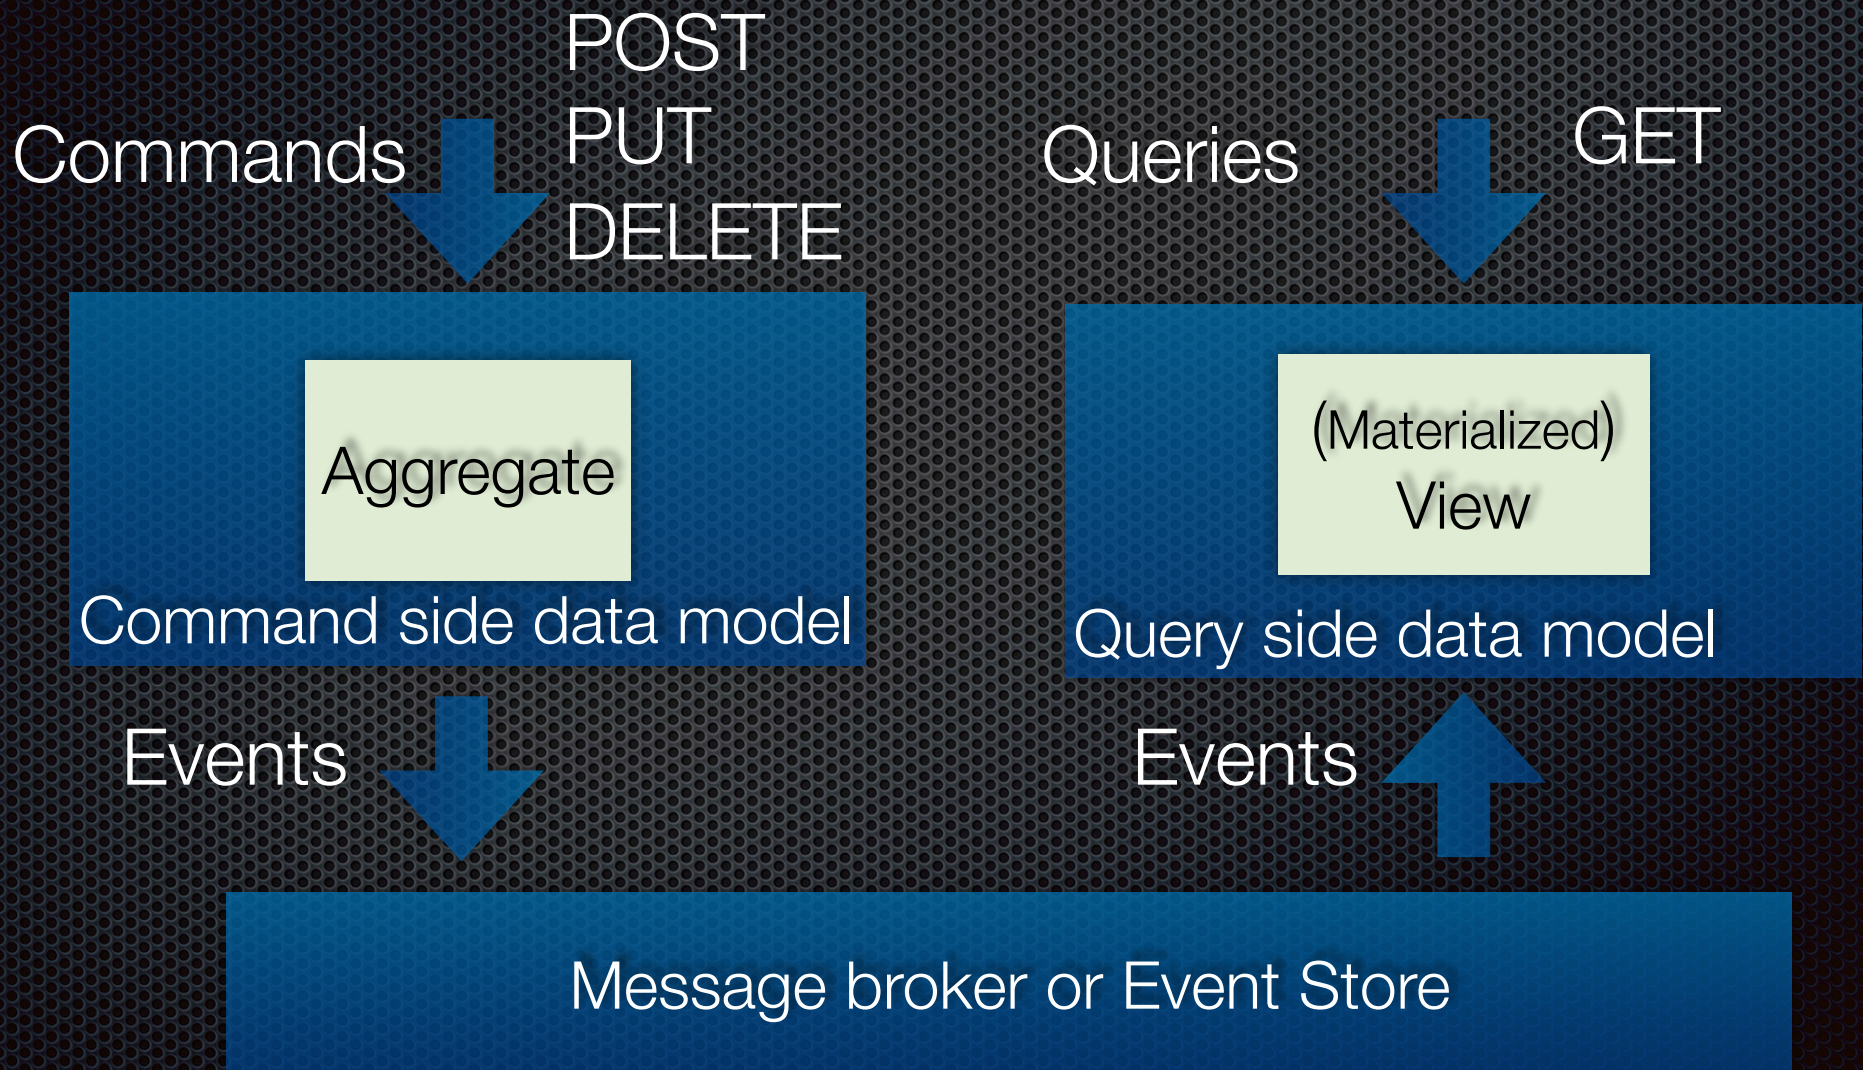
Customer  
information

Order  
information

Denormalized = efficient lookup



# Command Query Responsibility Segregation (CQRS)





# Queries $\Rightarrow$ database (type)

POST  
PUT  
DELETE

GET /customers/id

GET /orders?text=xyz

GET ...



Aggregate

MongoDB

ElasticSearch

Neo4j

Command side

Query side

Query side

Query side

Events



Event Store/Message Broker



# CQRS views are disposable

- ✦ Rebuild when needed from *source of truth*
- ✦ Using event sourcing
  - ✦ (Conceptually) replay all events from beginning of time
- ✦ Using traditional persistence
  - ✦ “ETL” from *source of truth* databases



# Handling replication lag

- ✦ Lag between updating command side and CQRS view
- ✦ Risk of showing stale data to user
- ✦ Either:
  - ✦ Update UI/client-side model without querying
  - ✦ Use updated aggregate version to “wait” for query view to be updated




# Agenda

- ✦ Transactions, queries and microservices
- ✦ Managing transactions with sagas
- ✦ Implementing queries with CQRS
- ✦ Implementing transactional messaging



# Messaging must be transactional

How to make **atomic** without 2PC?






# Publish to message broker first?

- Guarantees atomicity

**BUT**

- Service can't read its own writes
- Difficult to write business logic







Option: Event sourcing  
=  
Event centric approach to  
business logic and persistence

<http://eventuate.io/>

@crichardson



# Event sourcing: persists an object as a sequence of events



Event Store

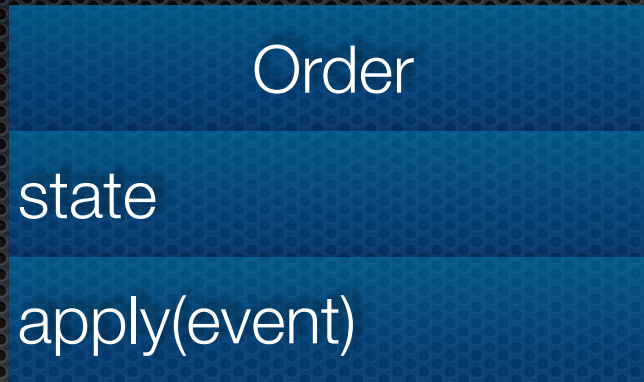
Event table

Entity id	Entity type	Event id	Event type	Event data
101	Order	901	OrderCreated	...
101	Order	902	OrderApproved	...
101	Order	903	OrderShipped	...




# Replay events to recreate in memory state

Instantiate with  
default  
constructor



Event store =  
database



Load events by **ID** and call apply()

Event Store

Event table

Entity id	Entity type	Event id	Event type	Event data
101	Order	901	OrderCreated	...
101	Order	902	OrderApproved	...
101	Order	903	OrderShipped	...

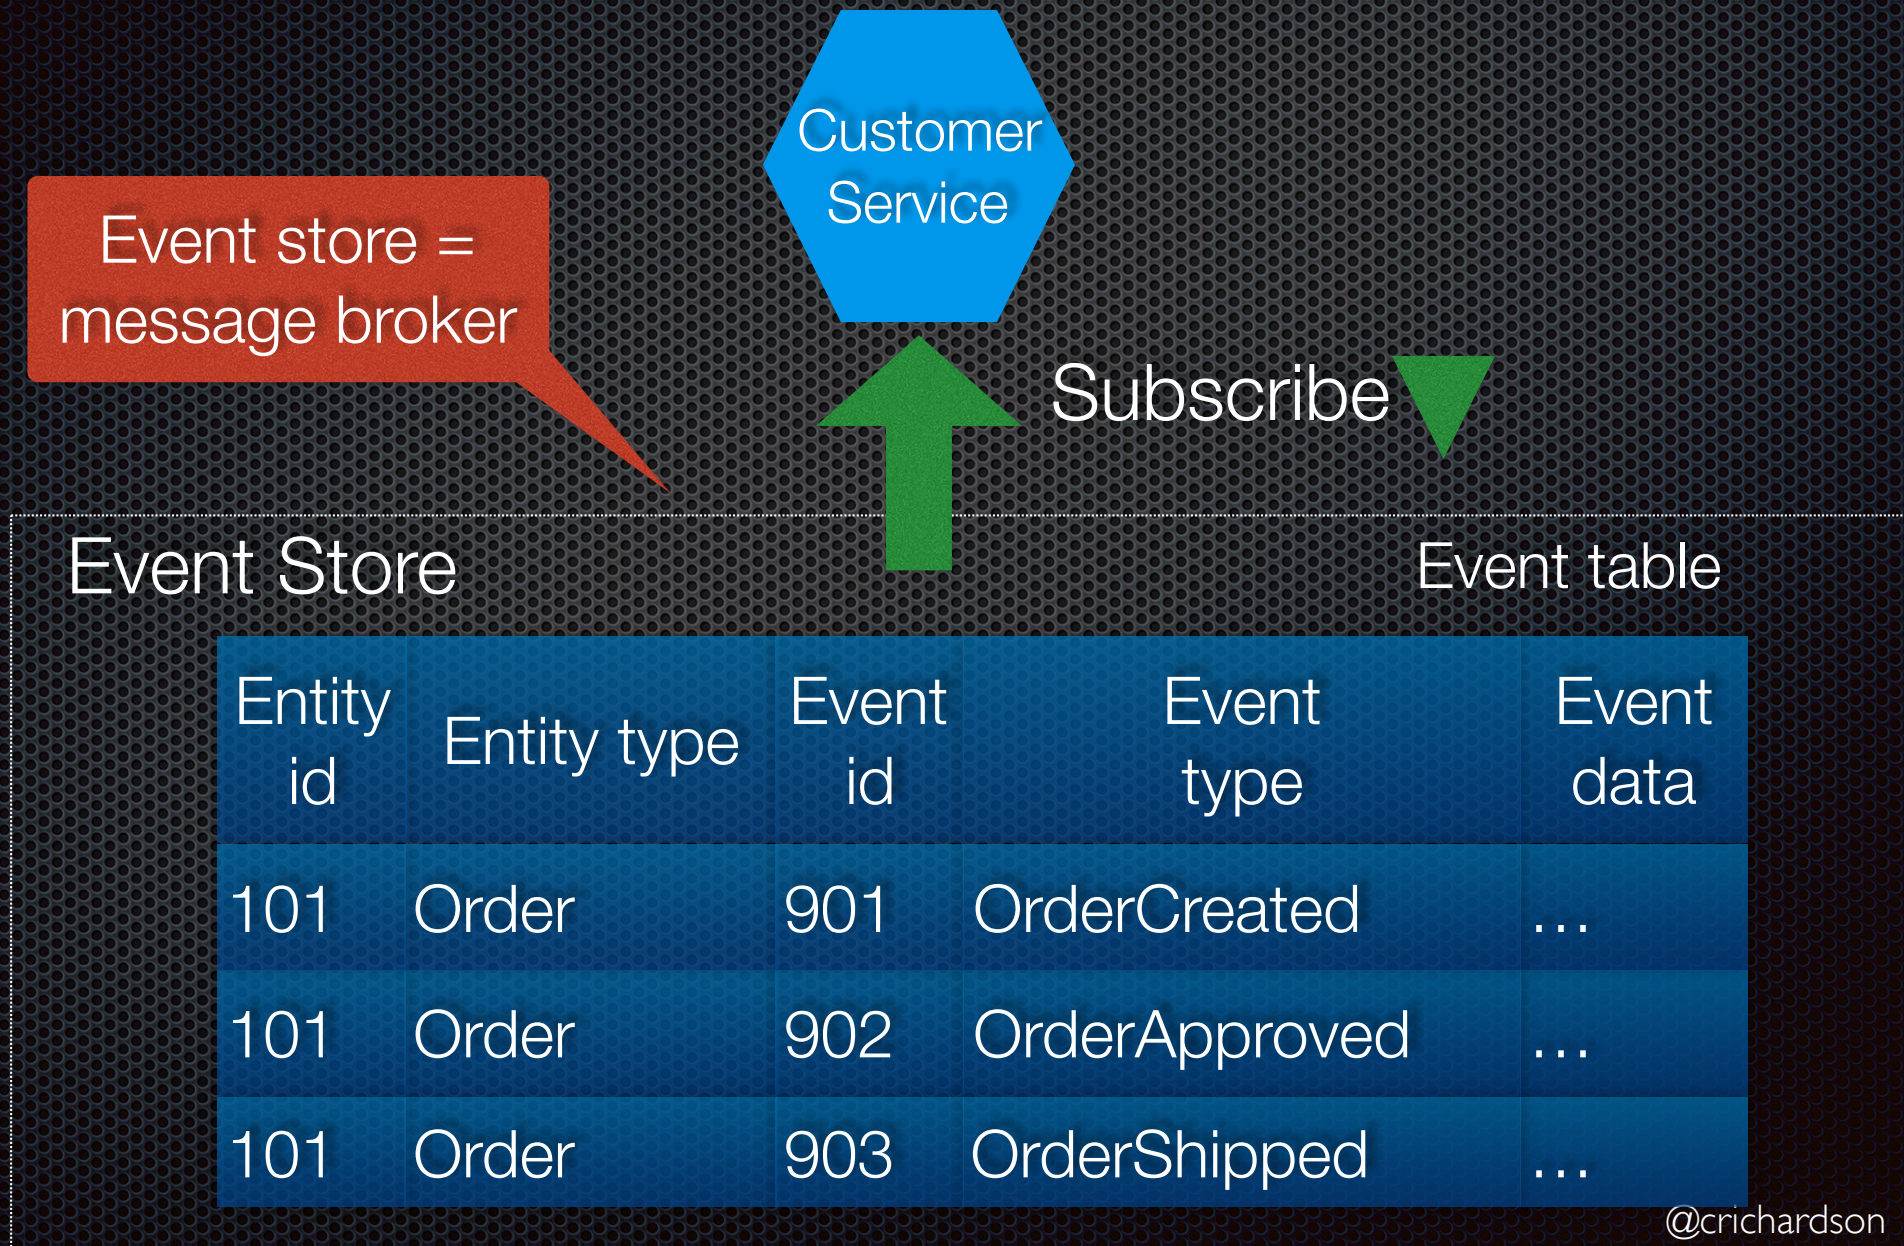


FYI:

Apache Kafka != event store



# Event sourcing guarantees: state change $\Rightarrow$ event is published





# Other benefits of event sourcing

Preserves history of domain objects

Supports temporal queries

Simplifies retroactive correction

Built-in auditing



# Drawbacks of event sourcing

Unfamiliar programming model

Evolving the schema of long-lived events

Event store only supports PK-based access  
⇒ requires CQRS ⇒ less consistent reads



# Drawbacks of event sourcing

Good fit for choreography-based sagas  
BUT orchestration is more challenging



Use event handler to translate event  
into command/reply message



Option:

Traditional persistence  
(JPA, MyBatis,...)

+

Transactional outbox pattern



# Spring Data for JPA example

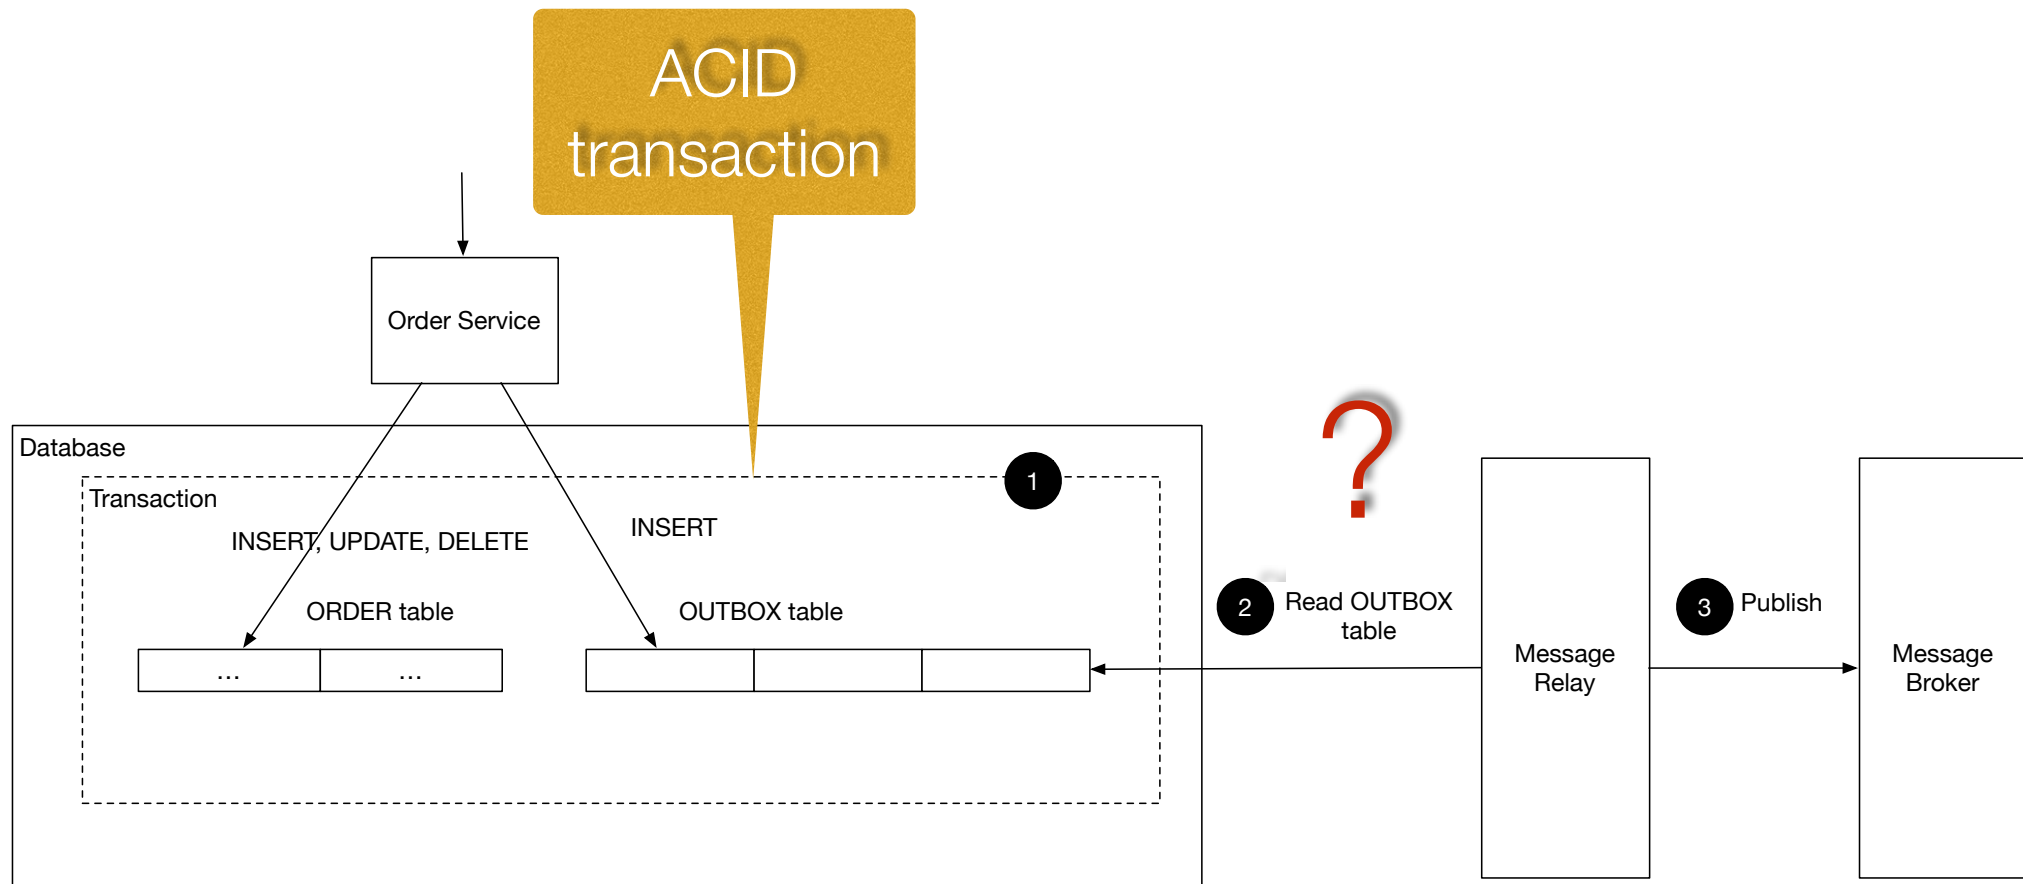
```
public class OrderService {  
  
    @Autowired  
    private DomainEventPublisher domainEventPublisher;  
  
    @Autowired  
    private OrderRepository orderRepository;  
  
    @Transactional  
    public Order createOrder(OrderDetails orderDetails) {  
        Order order = Order.createOrder(orderDetails);  
        orderRepository.save(order);  
        domainEventPublisher.publish(Order.class,  
            order.getId(),  
            singletonList(new OrderCreatedEvent(order.getId(), orderDetails)));  
        return order;  
    }  
}
```

Save order

Publish event



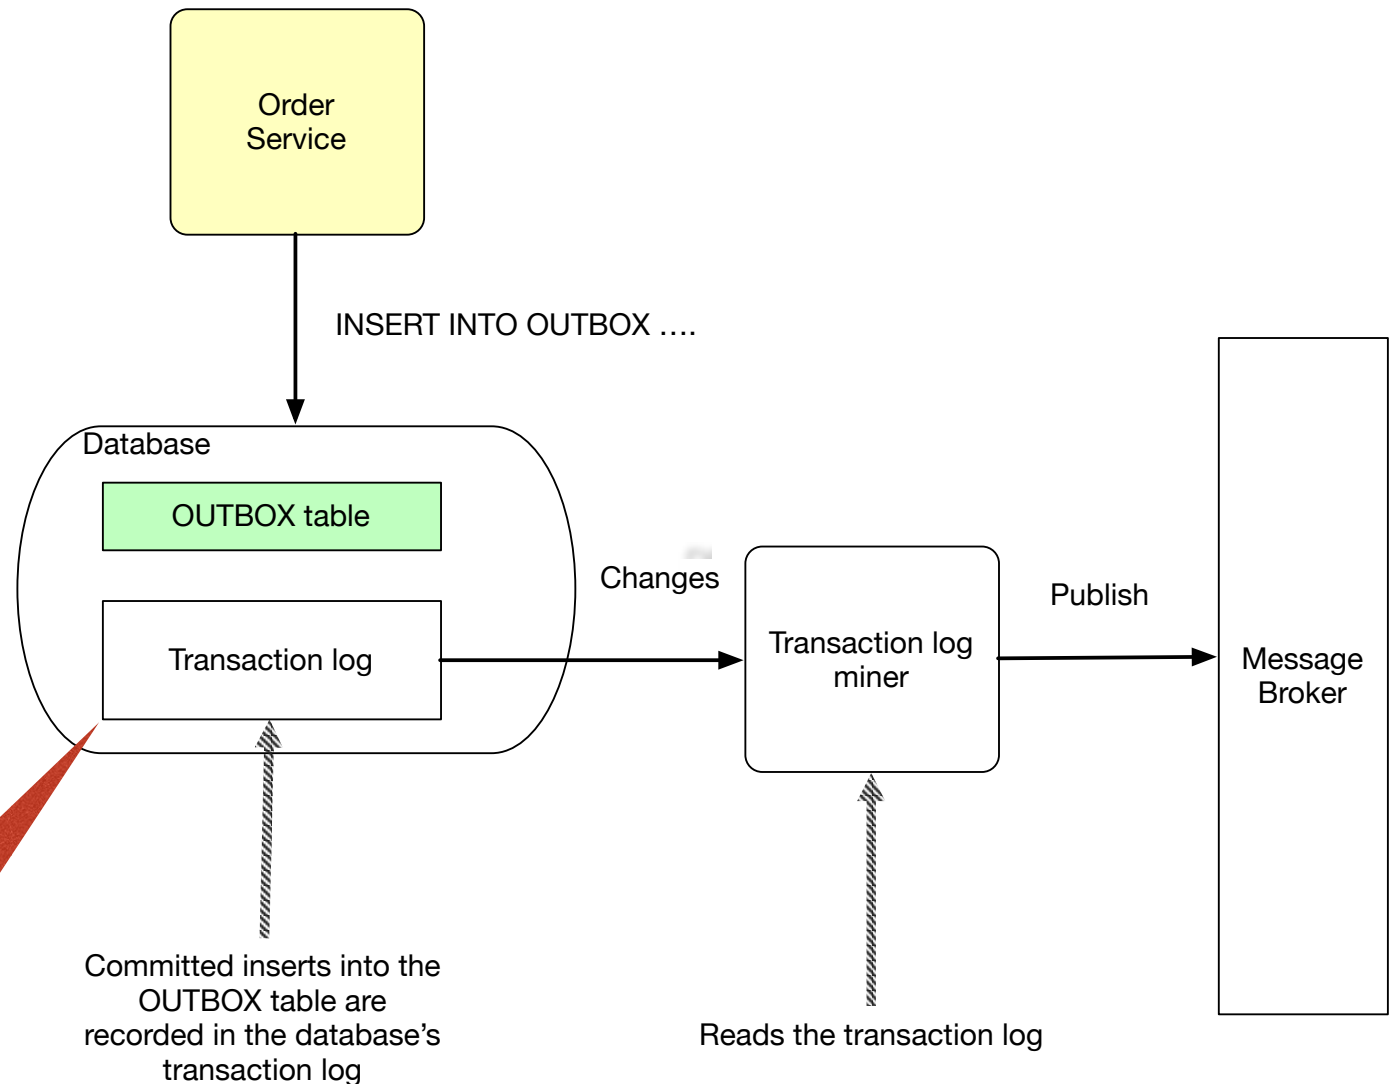
# Transactional Outbox pattern



- ❑ <https://microservices.io/patterns/data/transactional-outbox.html>
- ❑ <https://eventuate.io/>



# Transaction log tailing



MySQL binlog

Postgres WAL

AWS DynamoDB table streams

MongoDB change streams



# Polling the message table

- ✦ Simple
- ✦ Works for all databases
- ✦ BUT what about polling frequency

SELECT \* FROM MESSAGE...  
UPDATE MESSAGE






# Summary

- Use asynchronous messaging to solve distributed data management problems
- Services publish events to implement
  - choreography-based sagas
  - queries using CQRS views
- Services send command/reply messages to implement orchestration-based sagas
- Services must atomically update state and send messages
  - Event sourcing
  - Transactional outbox



 @crichardson chris@chrisrichardson.net



Questions?

Get a 40% discount

<http://bit.ly/gotochgo2019-microservices>