



# Architecting a Modern Financial Institution


# CREDIT CARD

September 2014



# GROWING QUICKLY IN A COMPLEX DOMAIN

# of clients (M)  
Credit Card



**10.5M**

Unique applications

**20**

Deploys per day

**2.6M**

Customers

**120**

Microservices

**262M**

Purchases

**105**

Engineers

**198**

Countries

# IMMUTABLE THEMES FROM OUR STACK



CLOJURE

LISP hosted on the JVM

Functional (opinionated), immutable data structures

Simple, concise, fast, concurrent

Tight REPL feedback cycles

Gradual typing (schemas)



DATOMIC



KAFKA



CLOUD

# IMMUTABLE THEMES FROM OUR STACK



CLOJURE



for your data

Accumulate-only

Reified ACID transactions, preserve what changed when

Query using data structures (Datalog)

Cloud native with integrated caching and scalable reads



DATOMIC



KAFKA



CLOUD

# IMMUTABLE THEMES FROM OUR STACK



CLOJURE

Immutable, persistent, partitioned log

Logical decoupling between services



DATOMIC

Temporal decoupling, useful for asymmetric workloads



KAFKA

Fault isolation and recovery (circuit breakers, dead letters)

Financial batch jobs expressed as a streams of messages



CLOUD

# IMMUTABLE THEMES FROM OUR STACK



CLOJURE

Infra as code (AWS)

Immutable upon provisioning (Docker)

Blue-Green deploys at service and company level

Kubernetes for speed and scalability



DATOMIC



KAFKA



CLOUD

Nubank HQ  
São Paulo, Brazil

# FUNCTIONAL BENEFITS

## HIRING

POSITIVE SELF SELECTION  
1-MONTH RAMP

## COMPLEXITY

SMALL, PURE FUNCTIONS  
STRAIGHTFORWARD TO UNTANGLE

## CONSISTENCY

COMPOSING A SMALL NUMBER OF IDIOMATIC  
LANGUAGE FEATURES



# CREDIT CARD ARCHITECTURE

Greenfield MVP





# BANK ACCOUNT

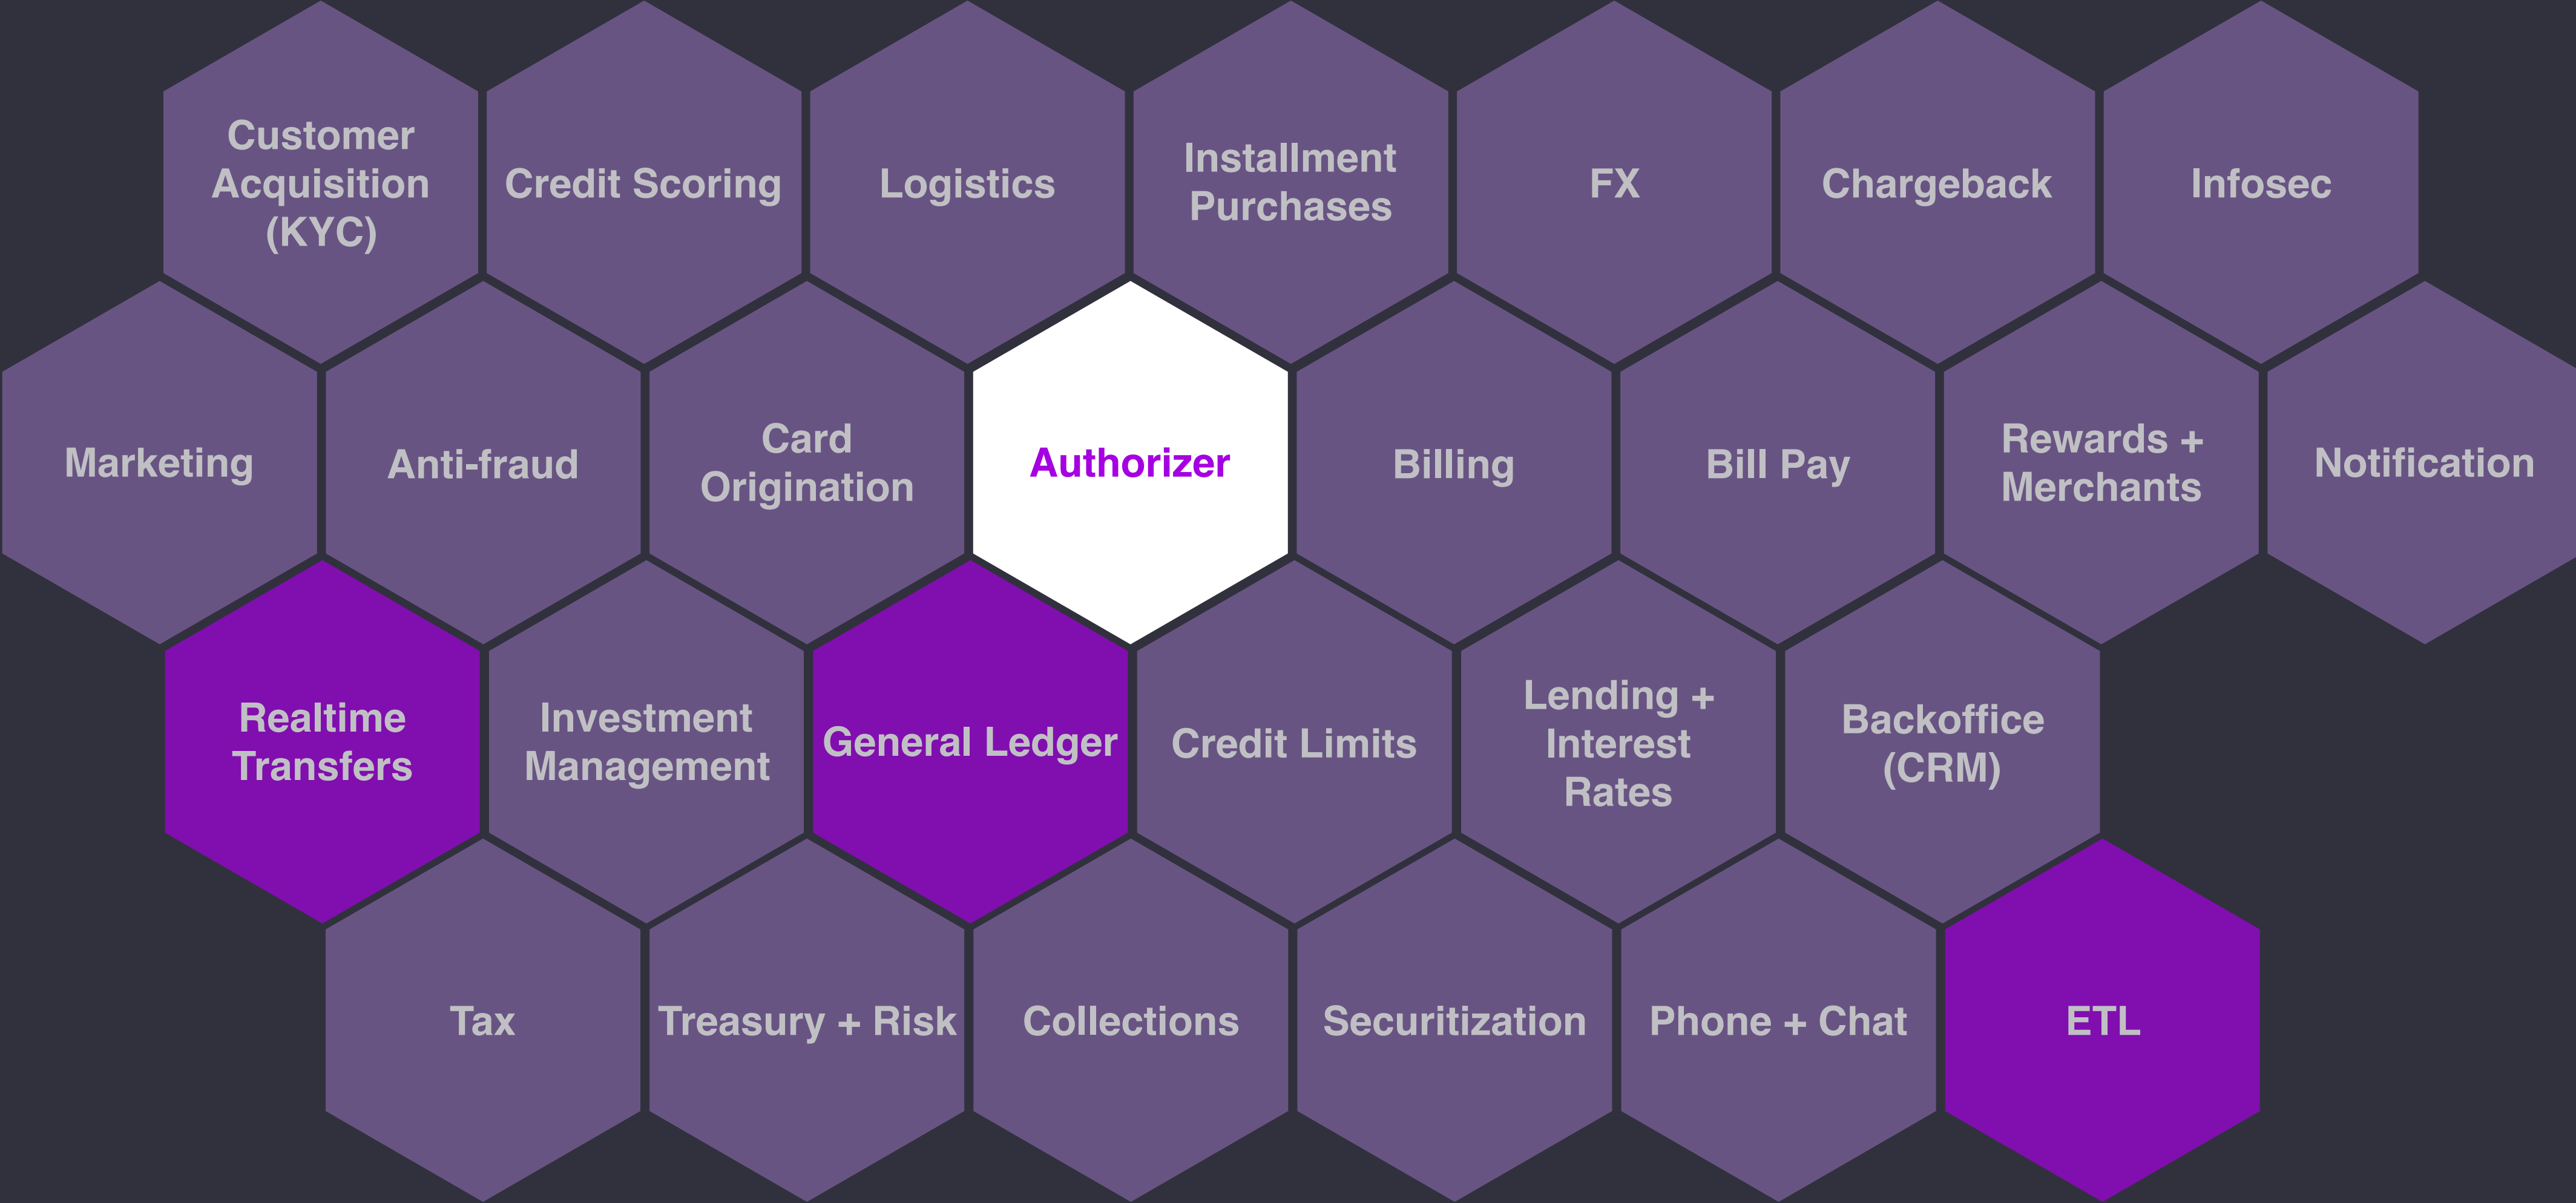
October 2017

# CORE BANKING + CREDIT CARD ARCHITECTURE



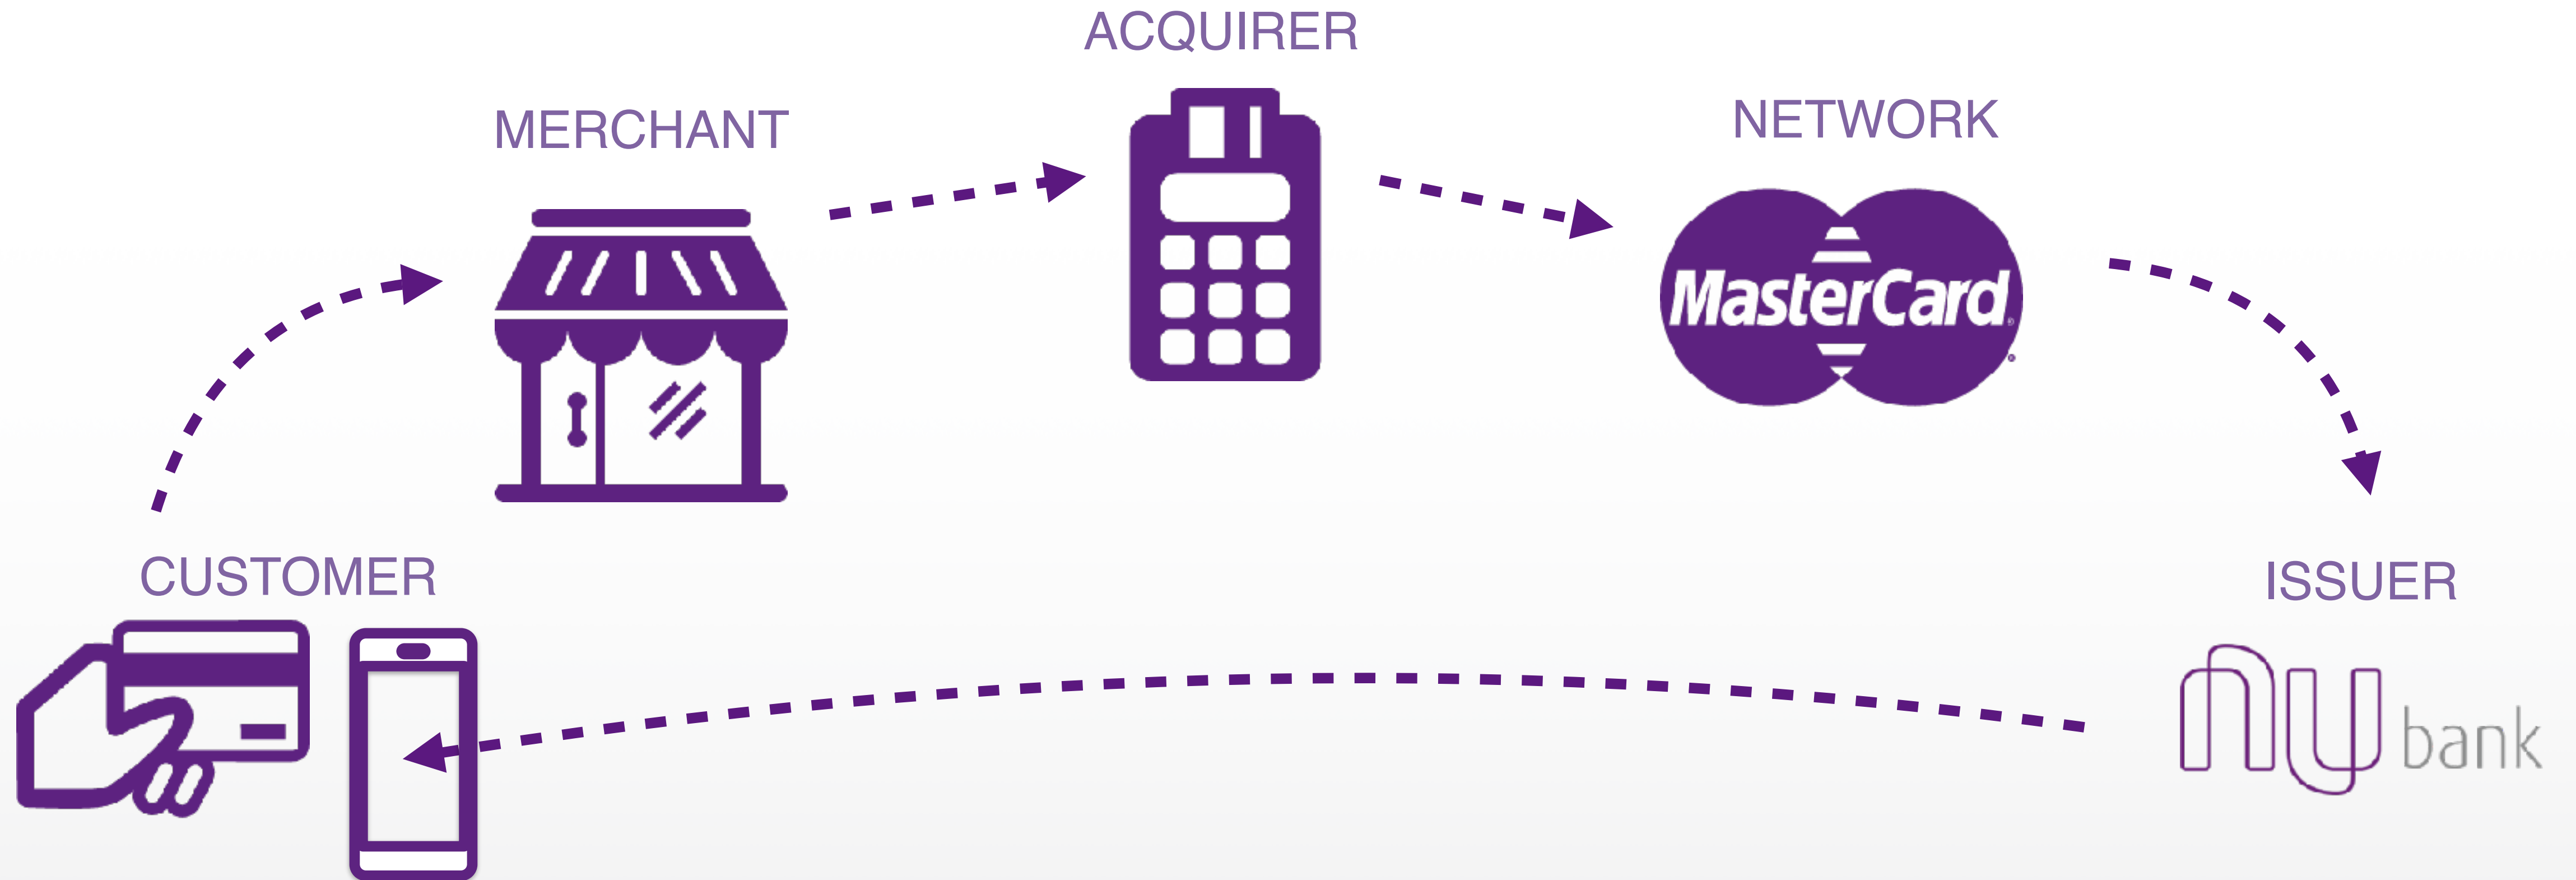
INFRASTRUCTURE

# PURCHASE AUTHORIZATION



INFRASTRUCTURE

# PURCHASE AUTHORIZATION VALUE CHAIN



# ISSUER AUTHORIZATION

NETWORK




ISSUER



# ISSUER AUTHORIZATION

MASTERCARD INTERFACE DEVICE




**1**

**Establish a connection**

**2**


**Receive authorization requests**

AUTHORIZER




# ISSUER AUTHORIZATION: ISO-8583

MASTERCARD INTERFACE DEVICE



**ISO-8583 Binary Message**

AUTHORIZER



HARDWARE SECURITY MODULE



# SCODEC BINARY PARSER FOR ISO-8583

```
object PANMappingFileD {  
  import scala.language.reflectiveCalls  
  
  val codec: Codec[SE33Subfield] = discriminated[SE33Subfield].by(intPadded(2))  
    .typecase(1, llvar(str).as[AccountNumberIndicator])  
    .typecase(2, llvar(intString(intPadded(2))).as[AccountNumber])  
    .typecase(3, llvar(yearMonth).as[ExpirationDate])  
    .typecase(4, llvar(str).as[ProductCode])  
    .typecase(5, llvar(intPadded(2)).as[TokenAssuranceLevel])  
    .typecase(6, llvar(intString(intPadded(2))).as[TokenRequestorID])  
    .typecase(7, llvar(intString(intPadded(2))).as[PANAccountRange])  
  
}
```


# ISSUER AUTHORIZATION: REQUIREMENTS

BRAND INTERFACE DEVICE



ISO-8583 Binary Message

AUTHORIZER

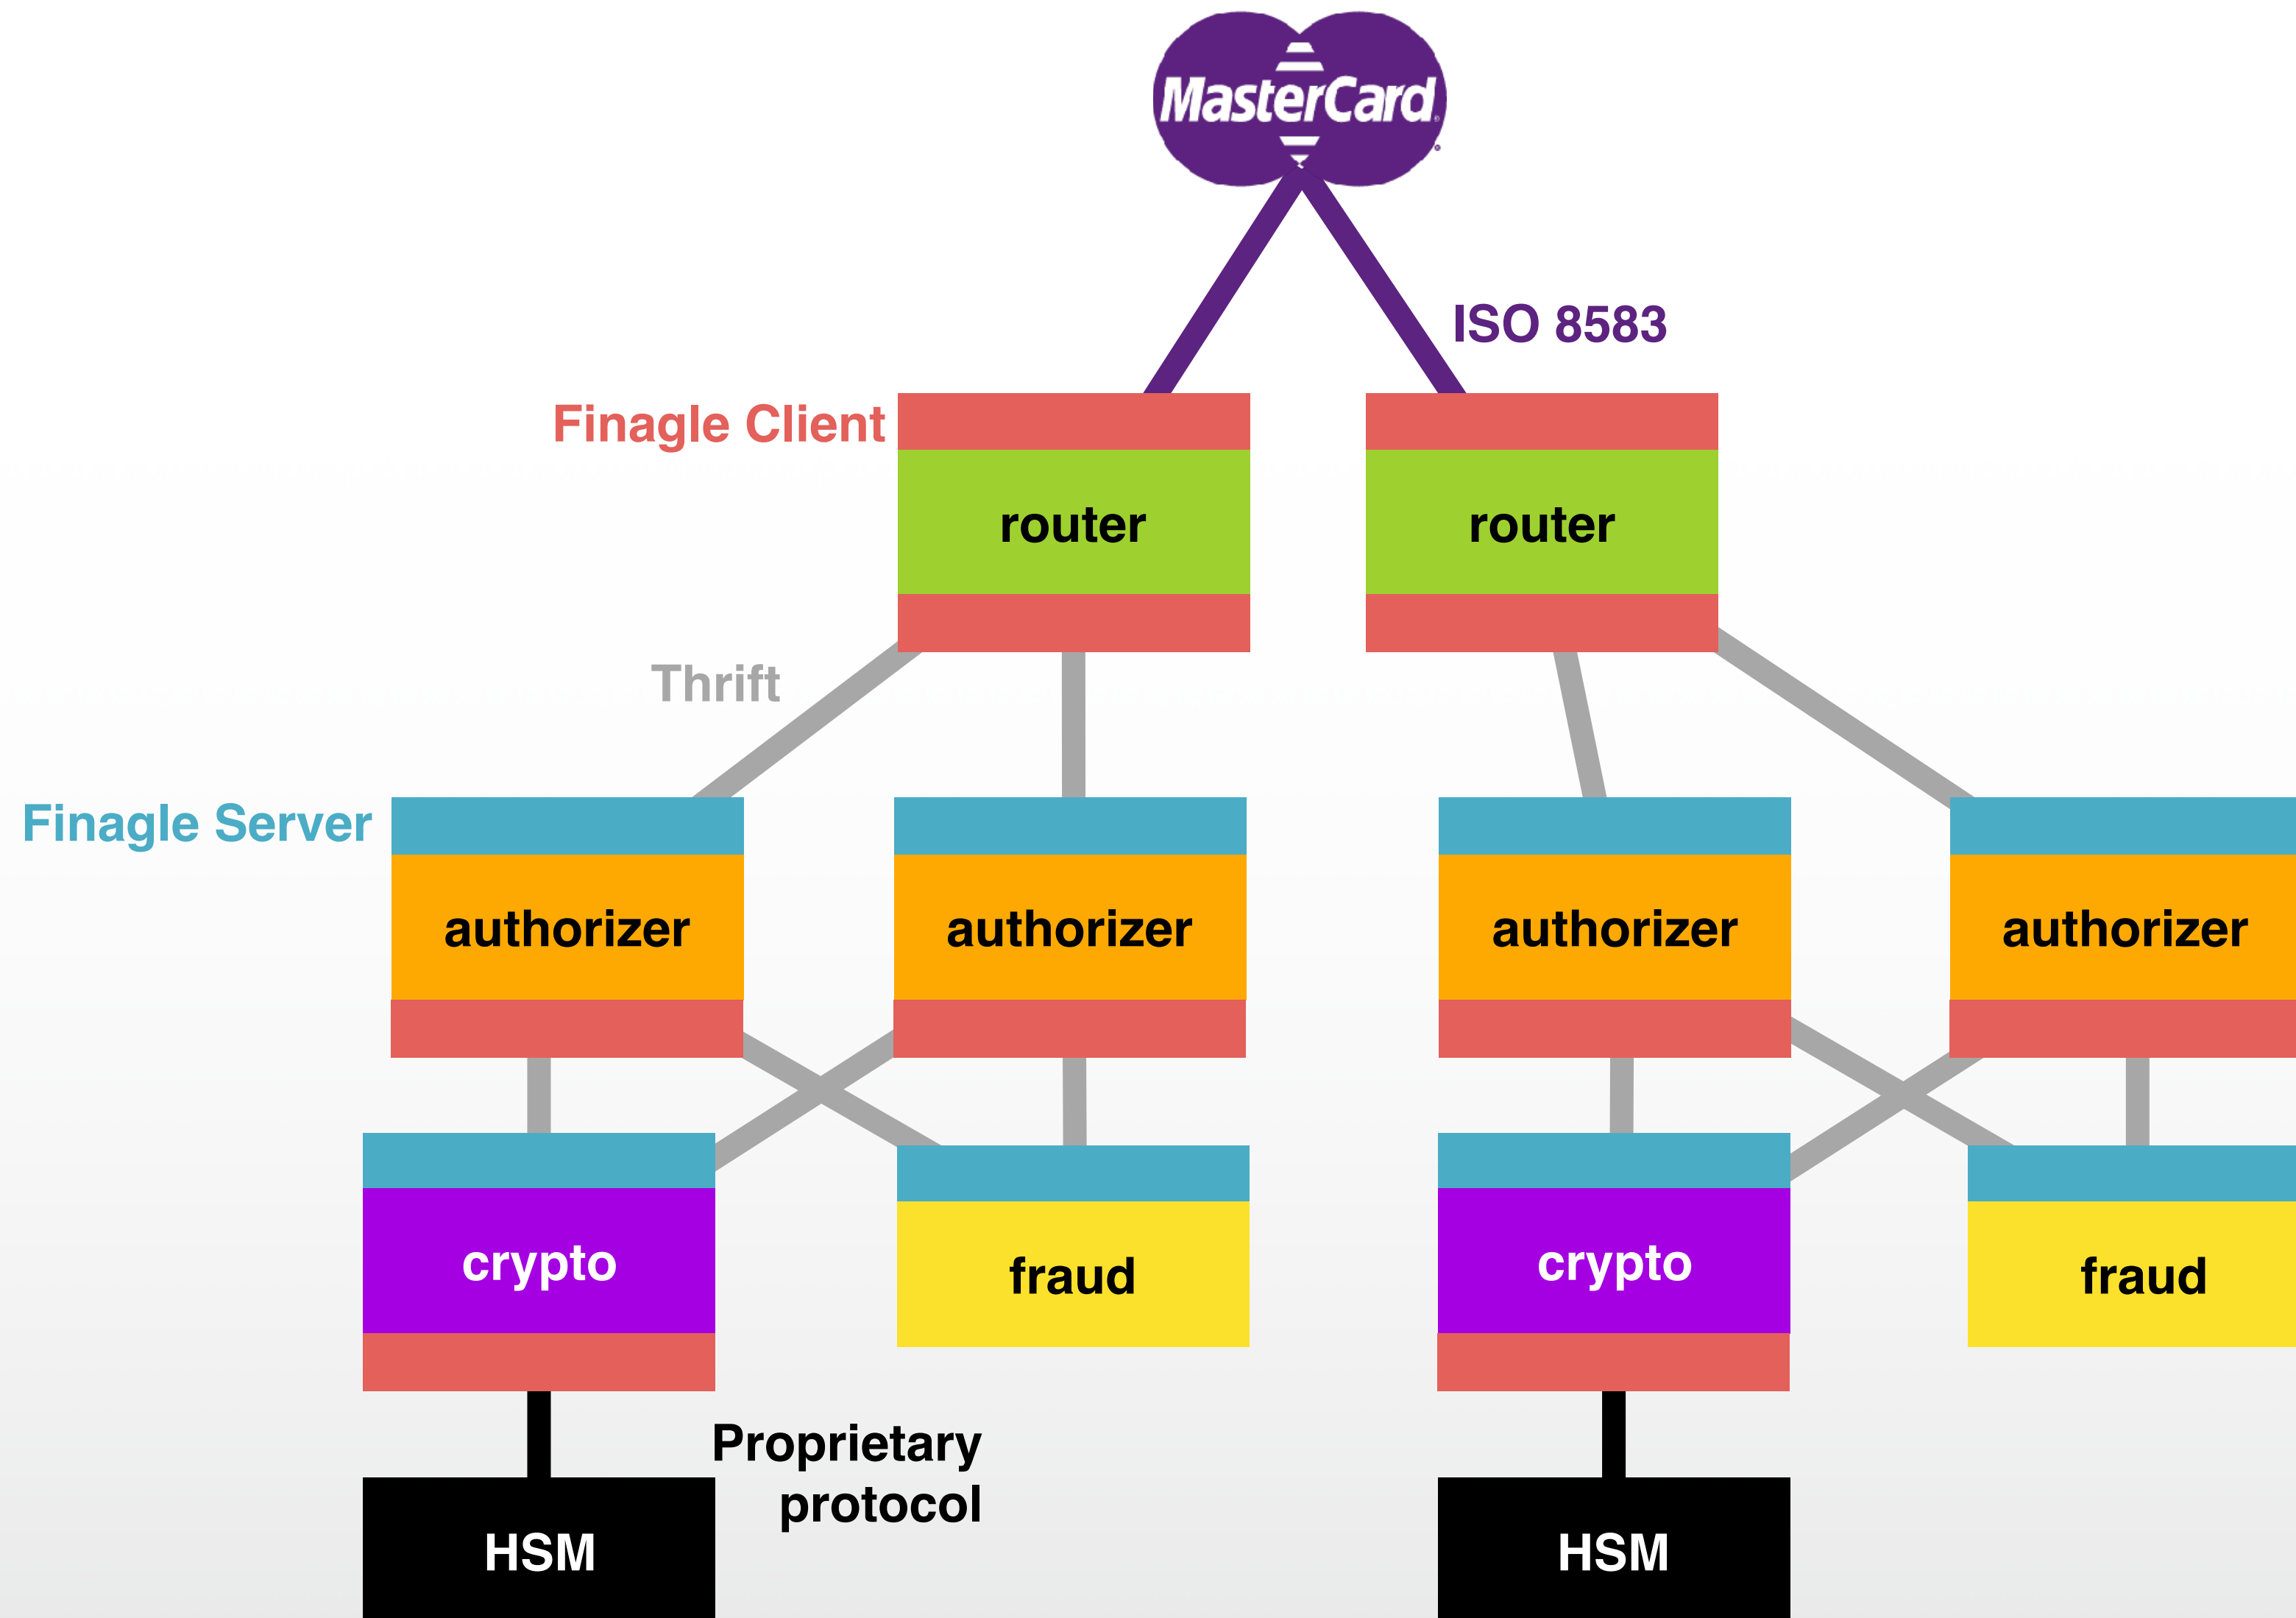


HARDWARE SECURITY MODULE (HSM)

**1.Highly Available**



**2.Physical Infrastructure**

# AUTHORIZER SERVICE LAYOUT

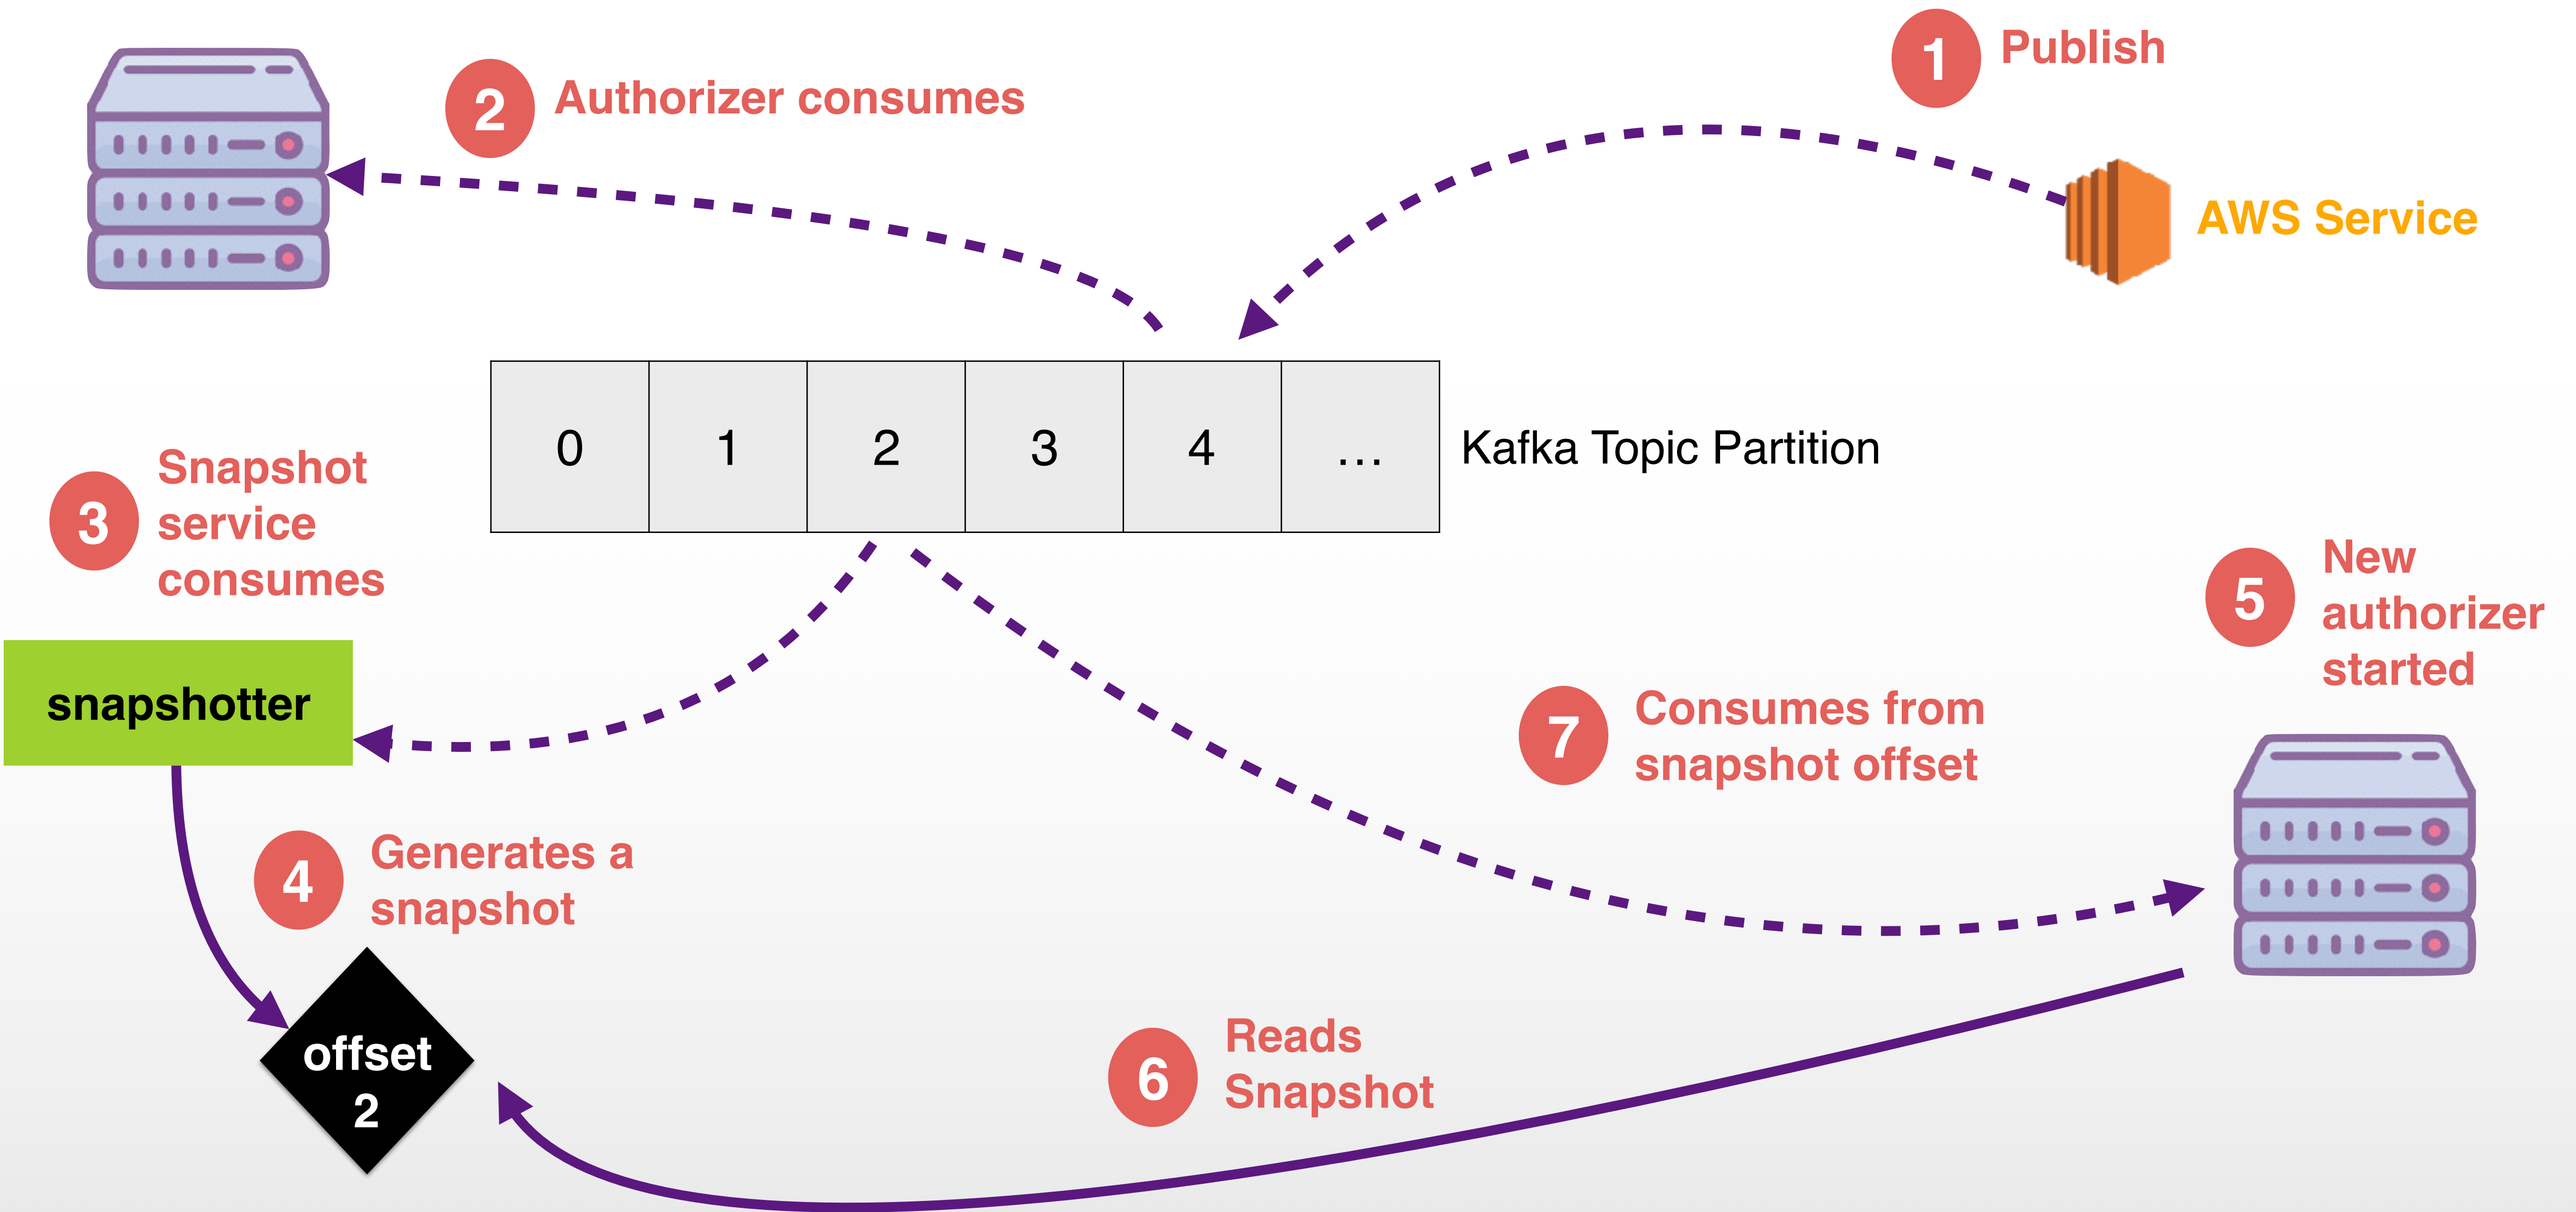


- Small set of highly available services
- Co-located with the MasterCard devices in the same datacenters
- Isolated: transaction authorization hot path does not need communication with the cloud
- Active-active disaster recovery (not shown)

# KAFKA AS THE BRIDGE BETWEEN ENVIRONMENTS

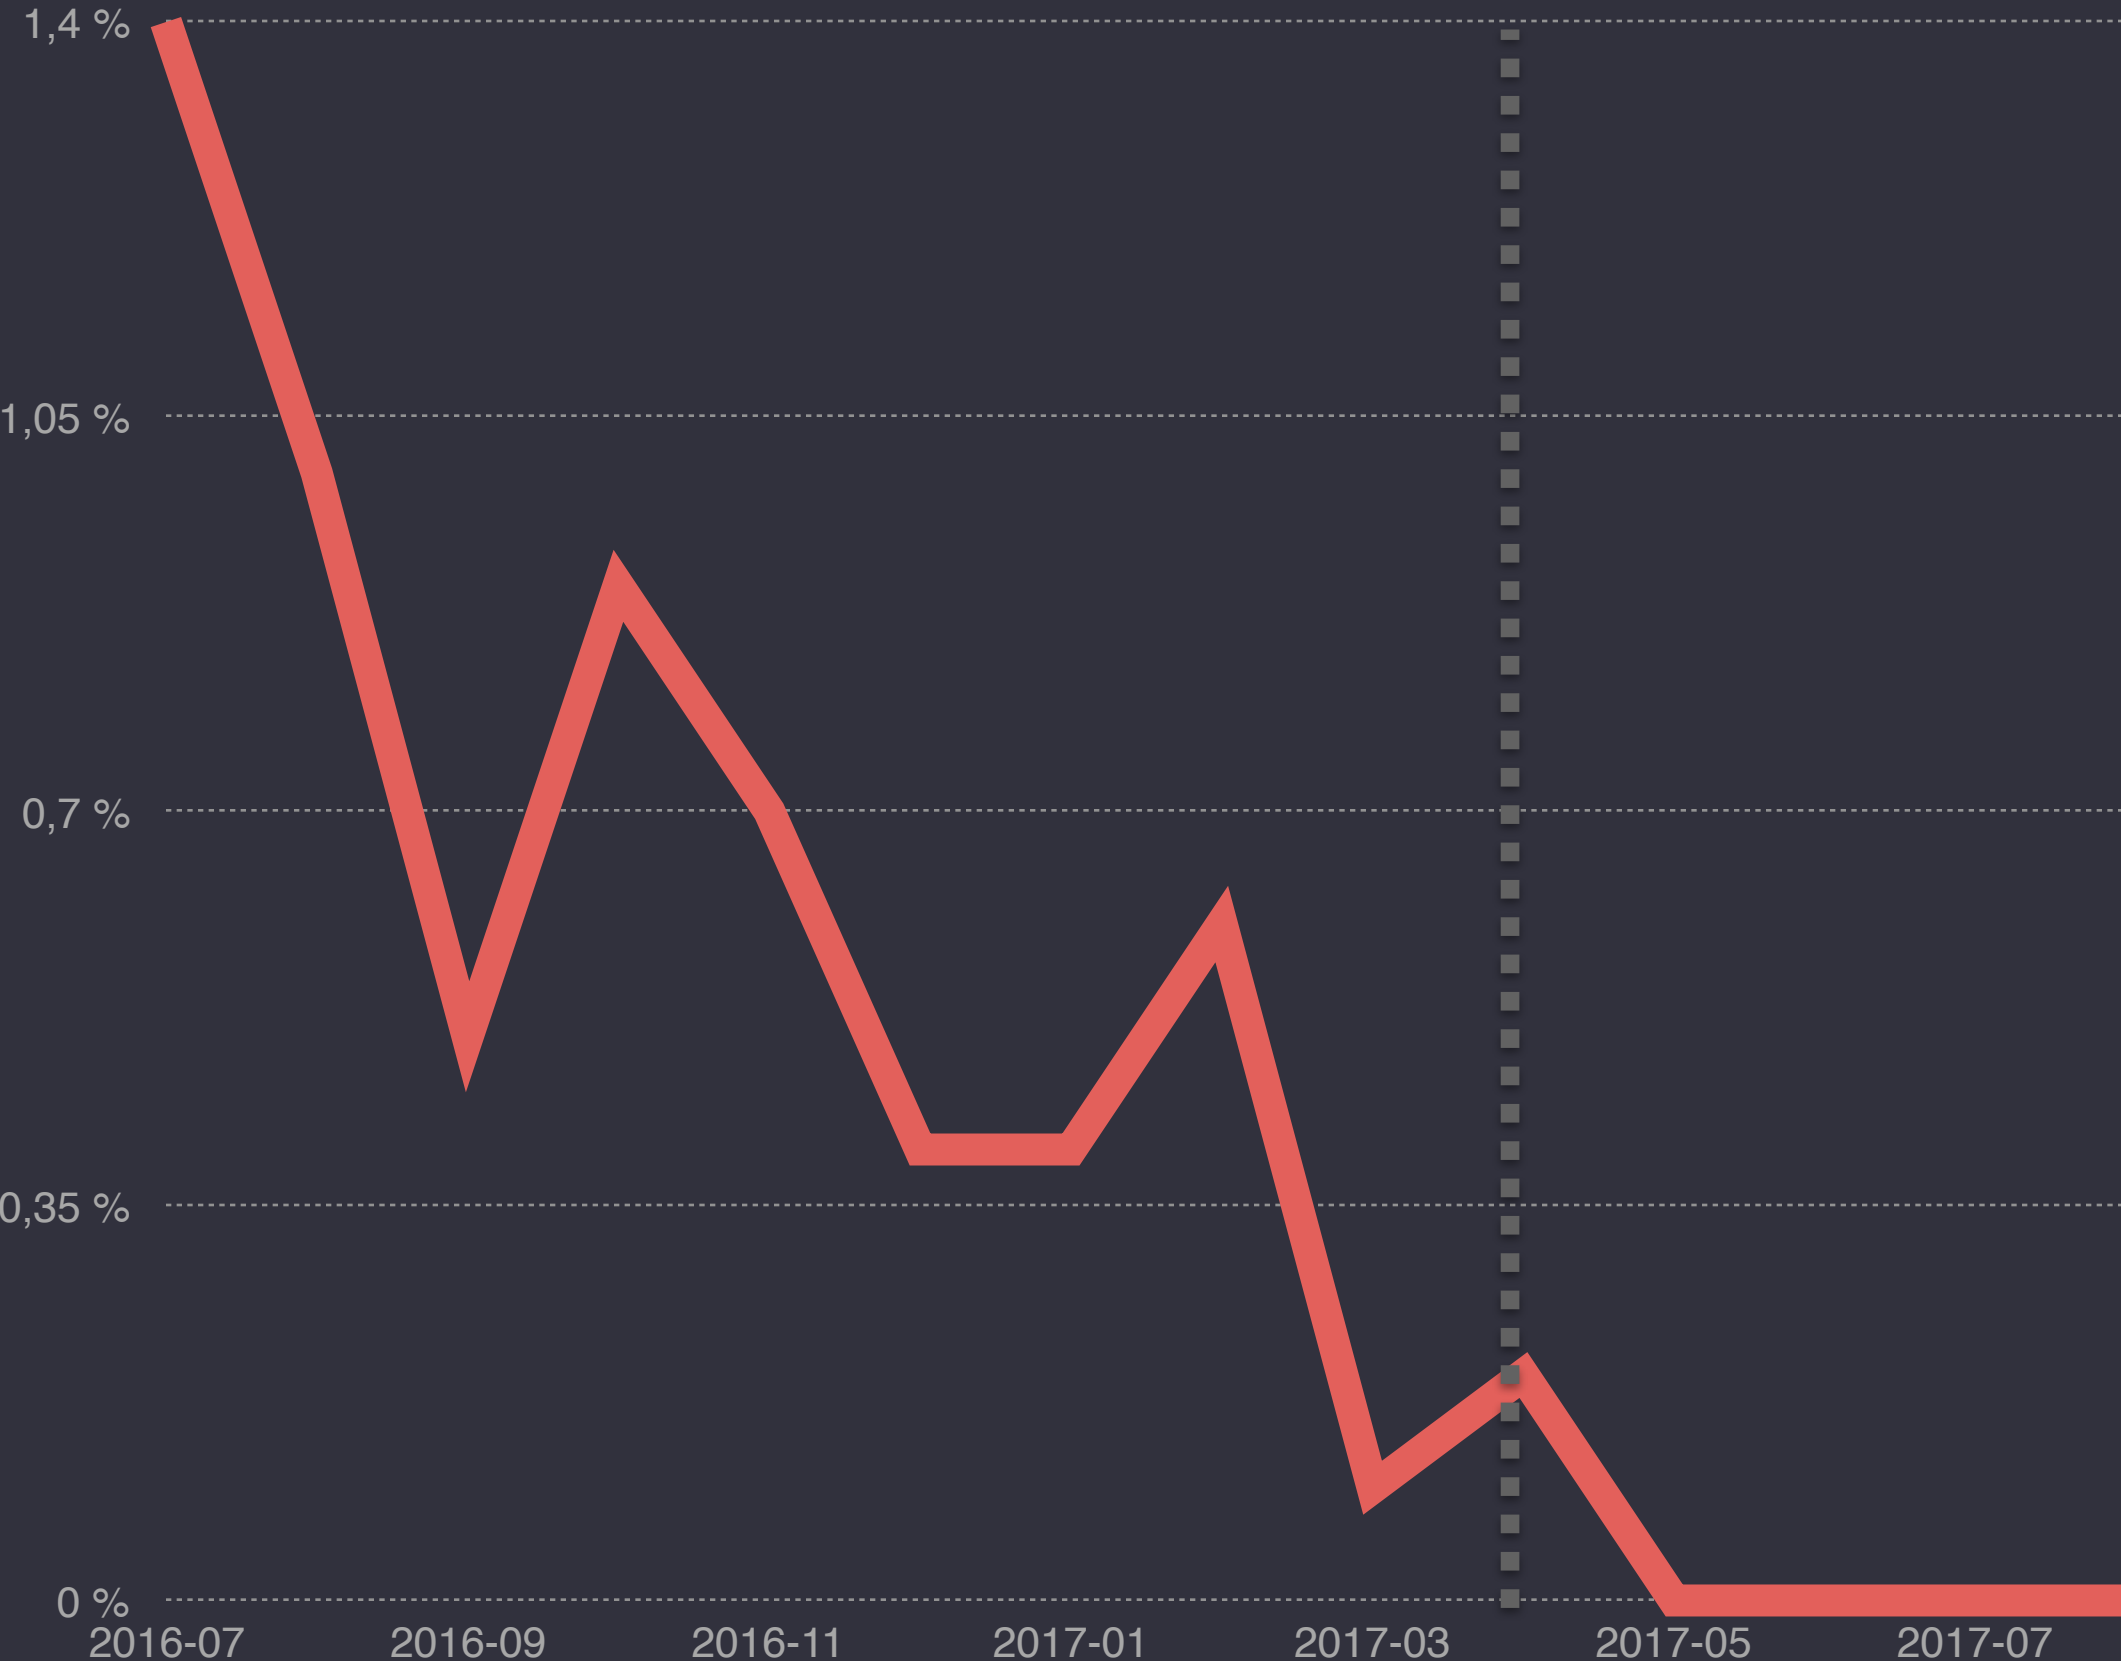


# KAFKA-BASED LOG/SNAPSHOT




# DRAMATIC IMPROVEMENTS IN RELIABILITY AND FRAUD

**% stand-in**



cutover

cutover



**% fraud capture**

**% fraud precision**

# DOUBLE ENTRY ACCOUNTING



INFRASTRUCTURE

# BUSINESS LOGIC DEPENDS ON DATA ACROSS MANY SERVICES



Purchases



Payments






Chargebacks



Interest



Currencies




Should we...

authorize a purchase?  
**Double Entry**  
block a card?  
charge interest?



# DOUBLE ENTRY: THE MODEL




$$\text{BALANCE} = \Sigma \text{\$}$$

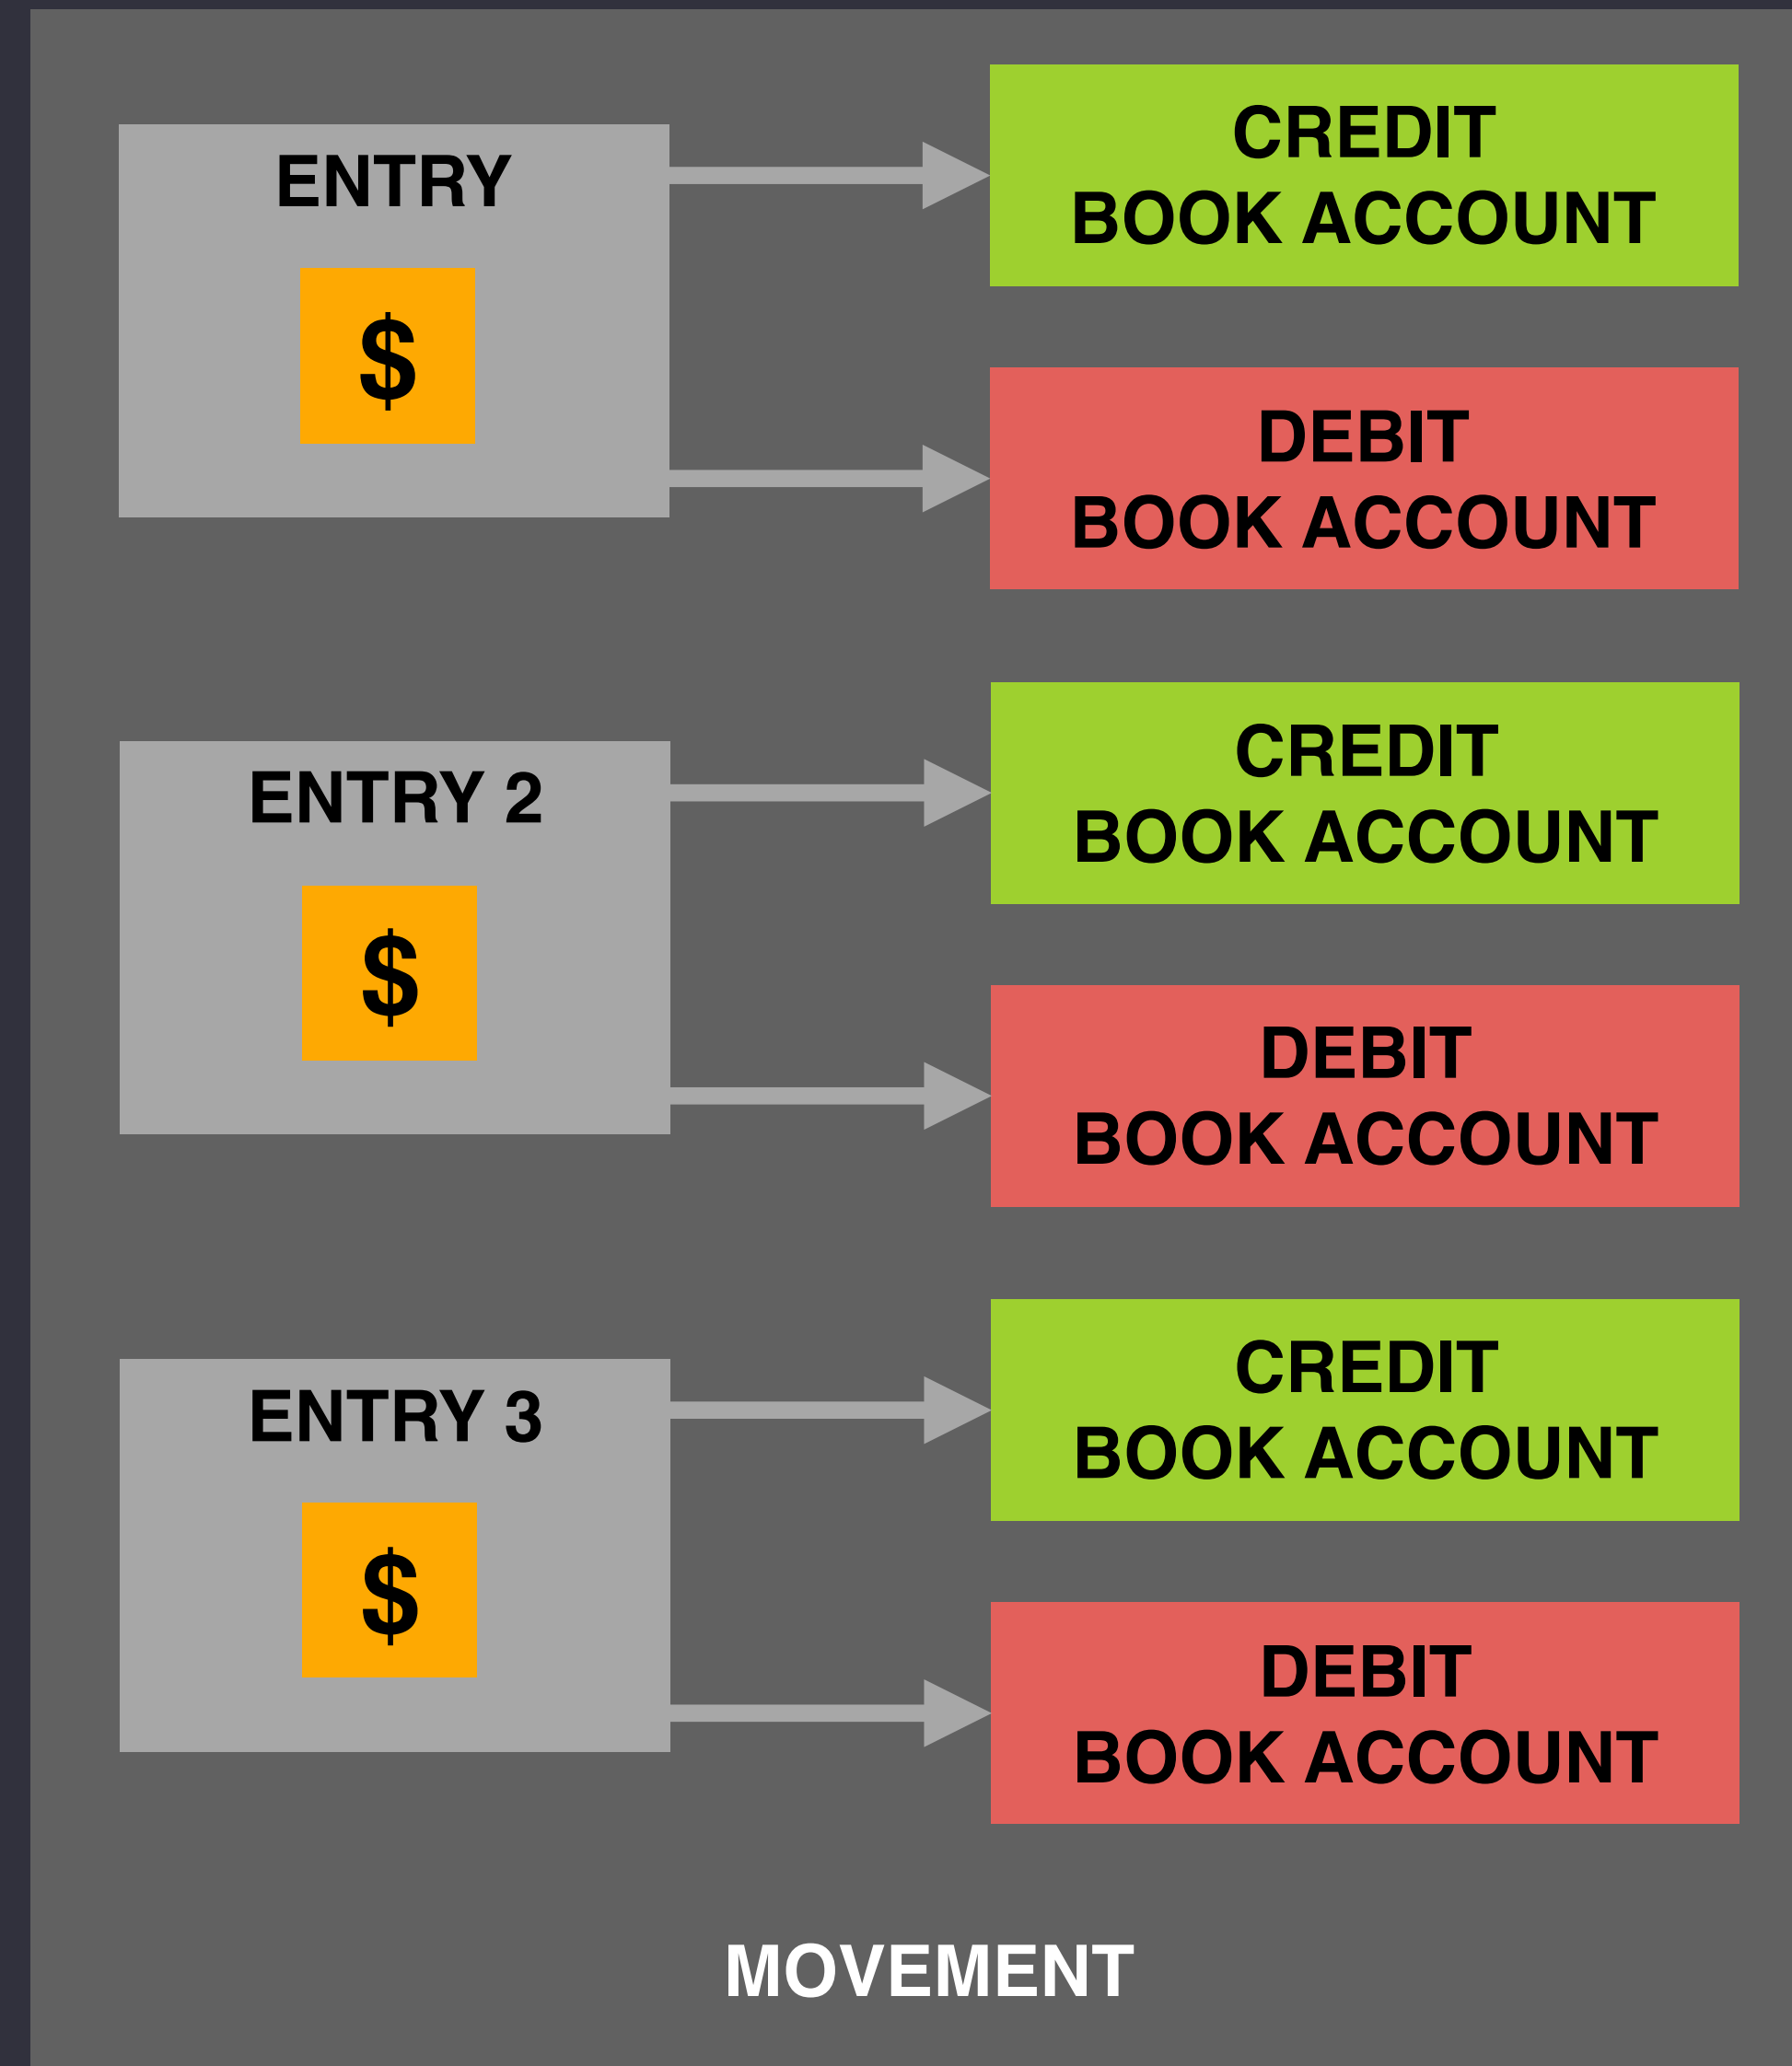
The sum of all credits and debits for one book-account is its balance

A customer's balance sheet is a cumulative function of their entire history

# DOUBLE ENTRY: THE RULEBOOK



NEW-PURCHASE  
NEW-PAYMENT  
...



# DOUBLE ENTRY: EXAMPLE MOVEMENT

```
(def unsettled-purchase  
  [
```

```
    {:entry/debit-account  :book-account-type.asset/unsettled  
     :entry/credit-account :book-account-type.liability/unsettled-counterparty  
     :entry/amount        #'transaction-amount  
     :entry/post-date     #'produced-date}
```

```
    {:entry/debit-account  :book-account-type.liability/current-limit-counterparty  
     :entry/credit-account :book-account-type.asset/current-limit  
     :entry/amount        #'transaction-amount  
     :entry/post-date     #'produced-date}
```

```
  ])
```

# DOUBLE ENTRY: CHALLENGES

ordering matters (i.e. movements are not commutative)

late arriving events (e.g. *a payment was made 3 days ago*)

fixing invariants

write throughput

# DOUBLE ENTRY: GENERATIVE TESTING OF INVARIANT

```
(def loss-property
  (prop/for-all [adjs (gen/vector (gen/one-of [gen-adjustment gen-payment gen-tx]) 1 10)
                initial-state (gen/such-that (comp not #{:late :pre-loss} :state) rbh/initial-state-gen)
                loss-event (gen/tuple (gen/no-shrink (gen/elements #{:pre-loss
                                                                    :credit-loss
                                                                    :id-fraud-loss
                                                                    :fraudster}))
                                      (tg/make-generator LocalDateTime)
                                      (tg/make-generator LocalDate))]
                (check-properties adjs initial-state loss-event)))
```

# DOUBLE ENTRY: CHALLENGES

ordering actually matters (i.e. movements are not commutative)

late arriving events (e.g. *a payment was made 3 days ago*)

fixing invariants

**write throughput**


# SHARDED, FAULT TOLERANT INFRASTRUCTURE



**INFRASTRUCTURE**

# SCALING BOTTLENECKS

# of clients (M)  
Credit Card



1. database throughput limits required throttling writes
2. batch job latency impacting customer experience



# SCALING PLAN

Need to partition the workload

Customer data is spread across services

Interactions between customers are minimal

Safe to partition the user base

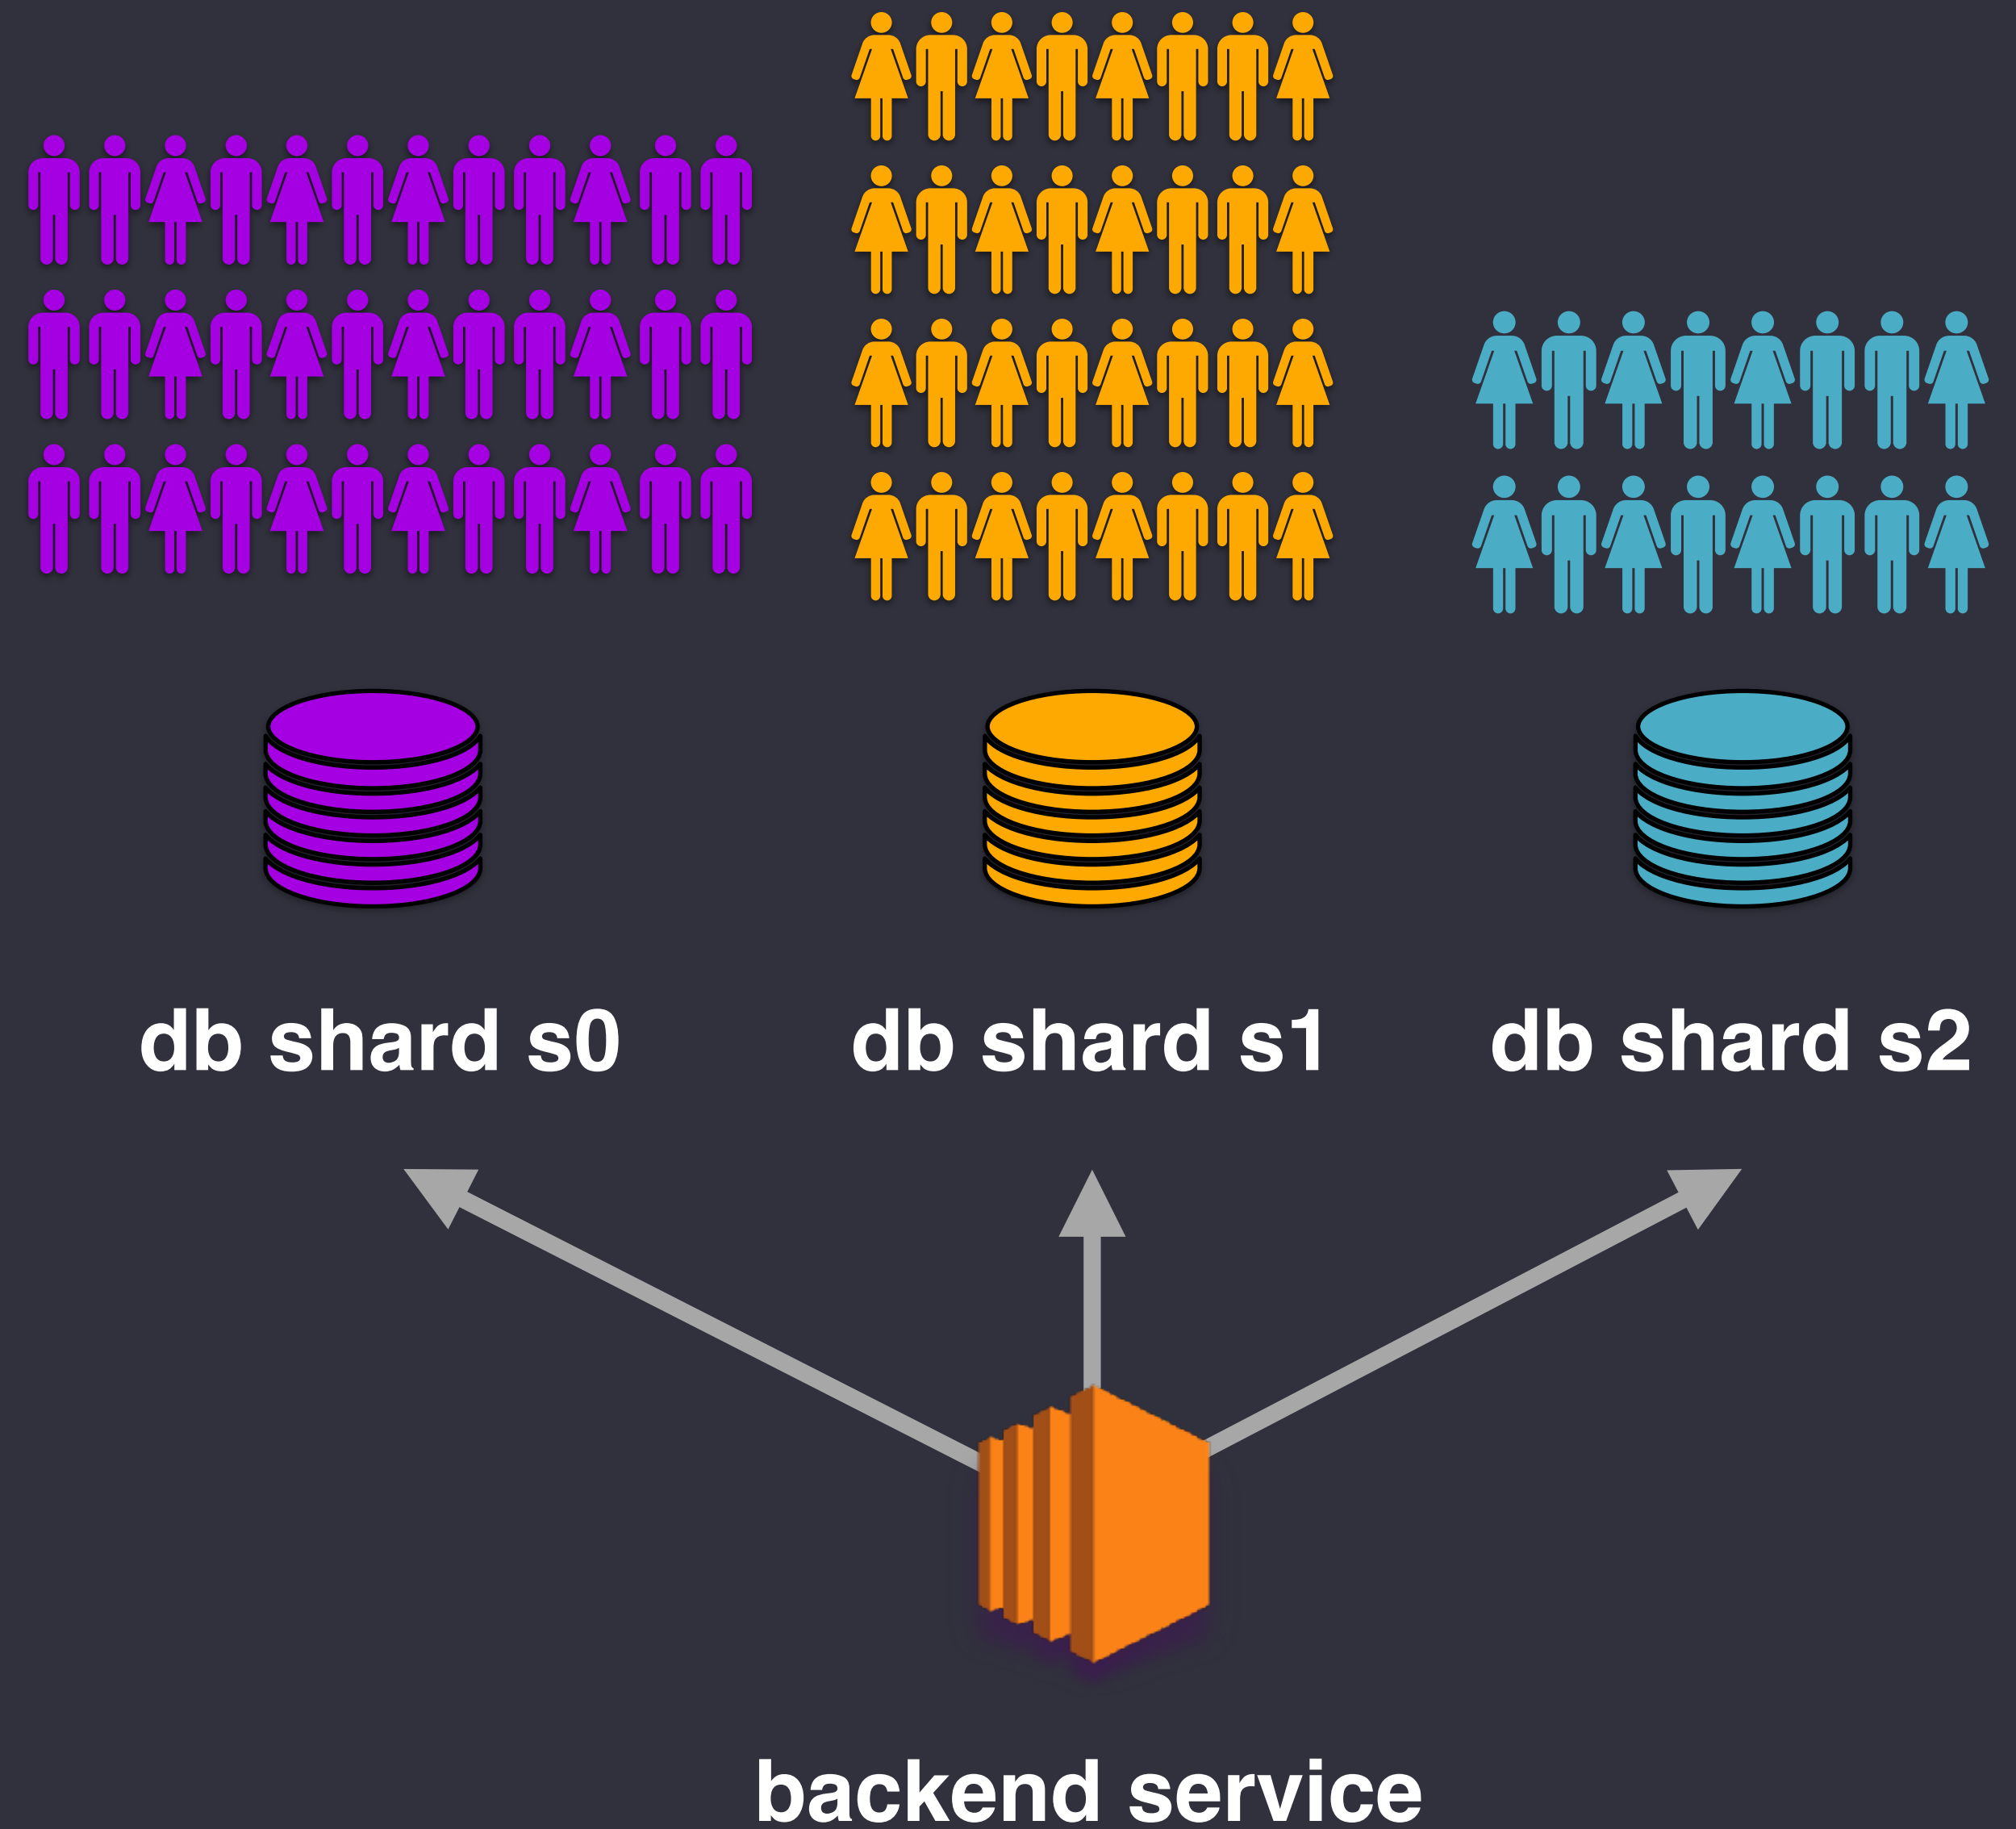


# OPTION #1: PARTITION SERVICE DATABASES

Database writes were the worst bottleneck

Option: horizontally partition each database

Change every service to route queries and writes to the appropriate shard

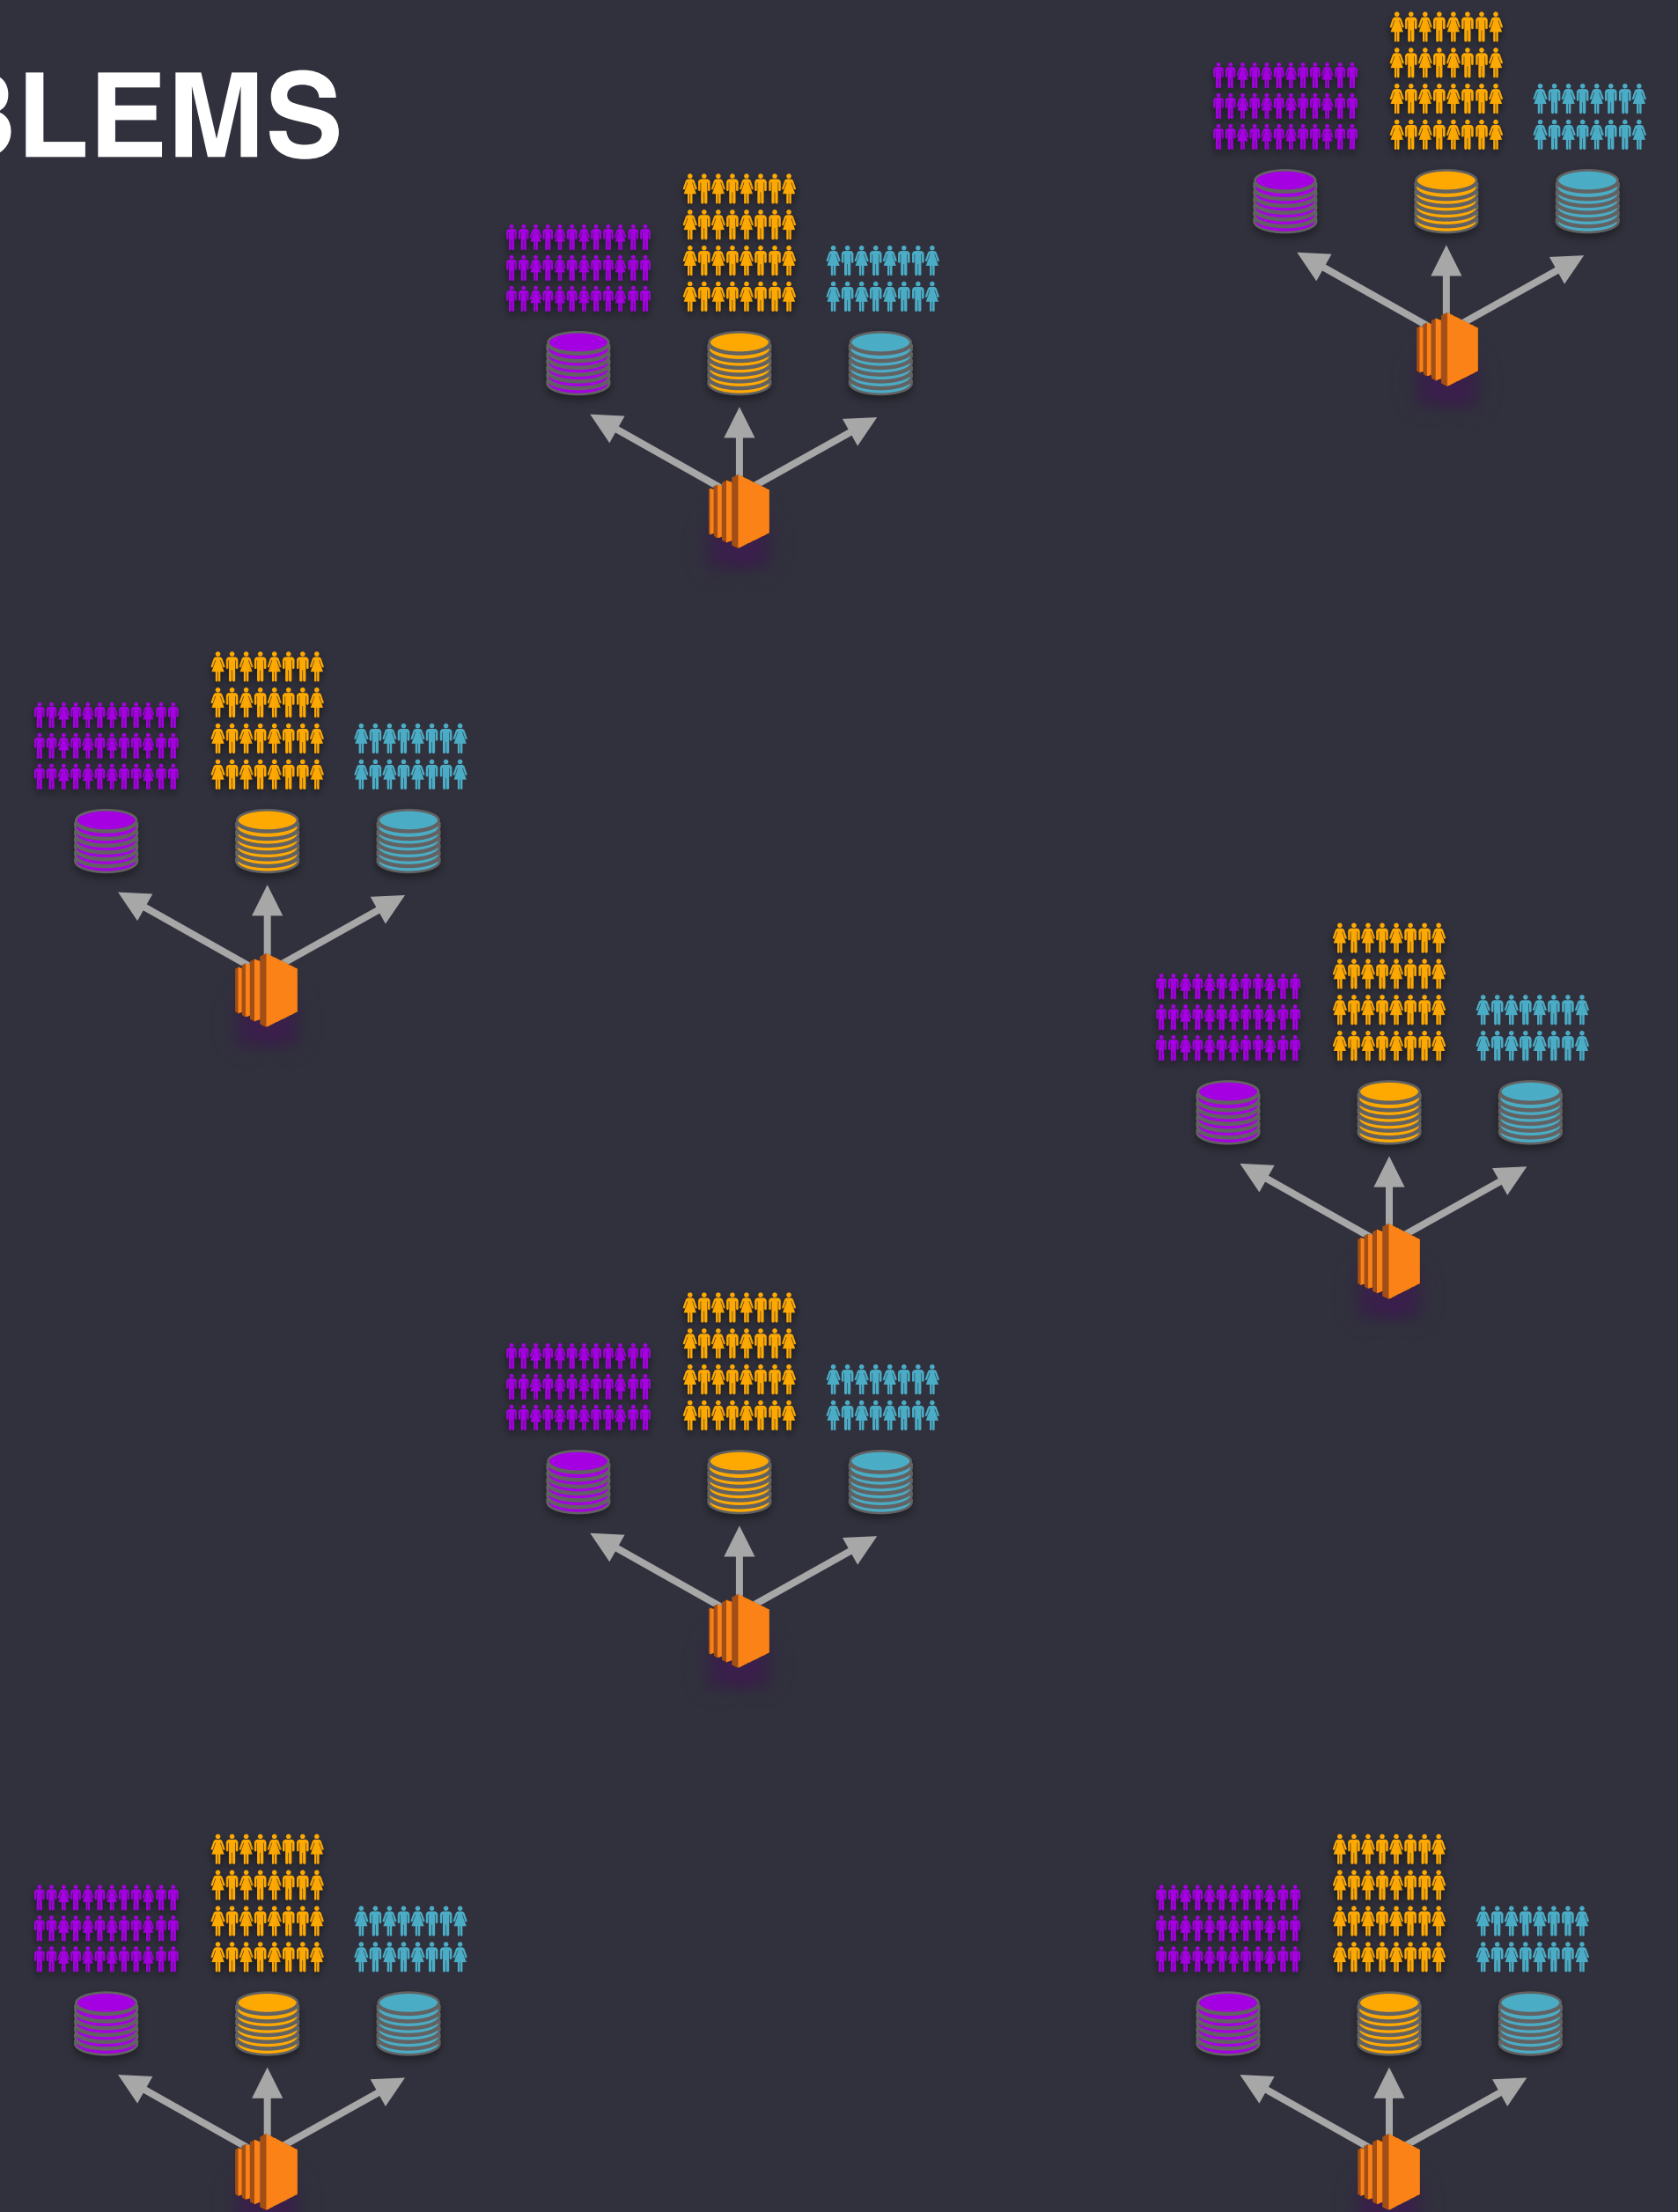


# OPTION #1: PROBLEMS



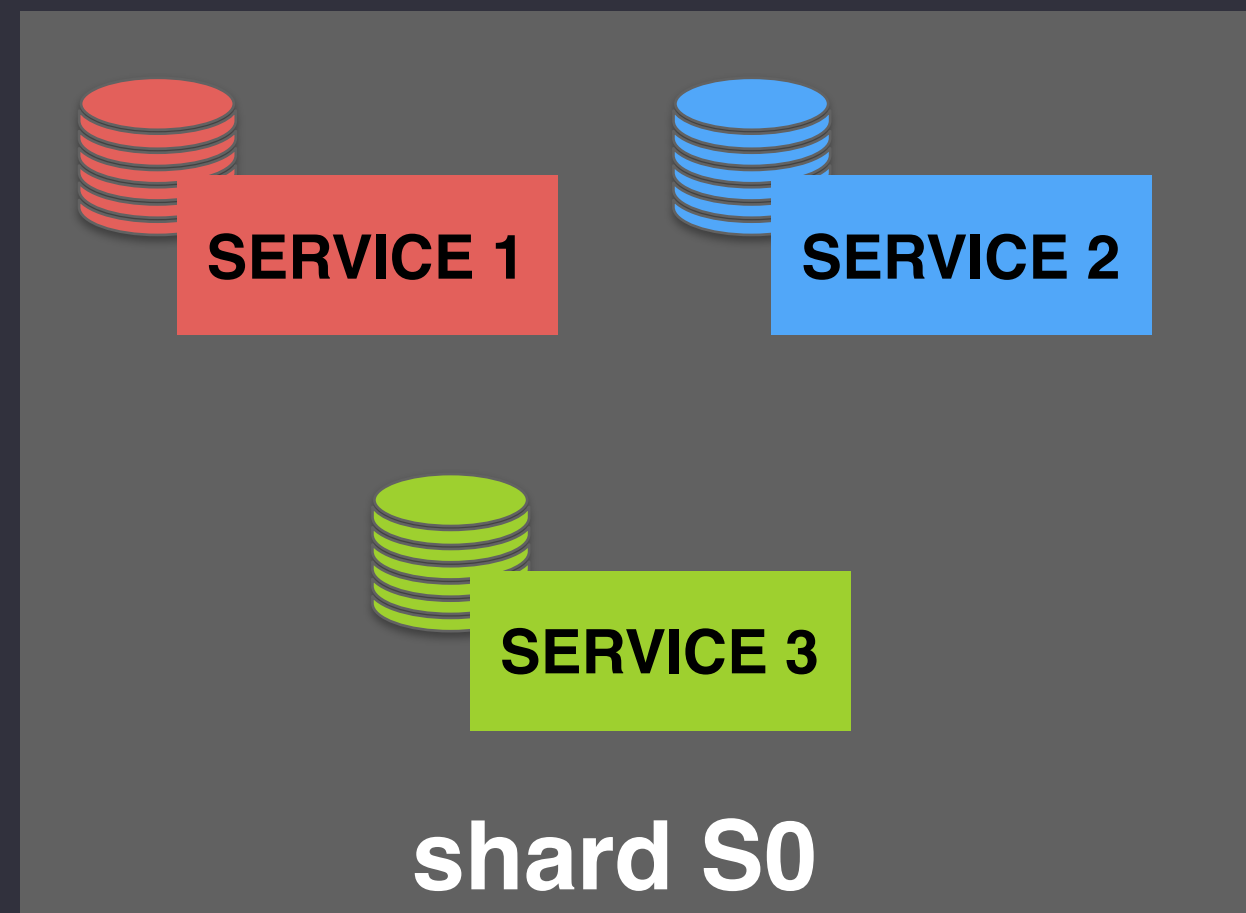
Enormous effort to change **every** service

Doesn't address **non-db** bottlenecks

Risks intermingling data infrastructure code with business logic

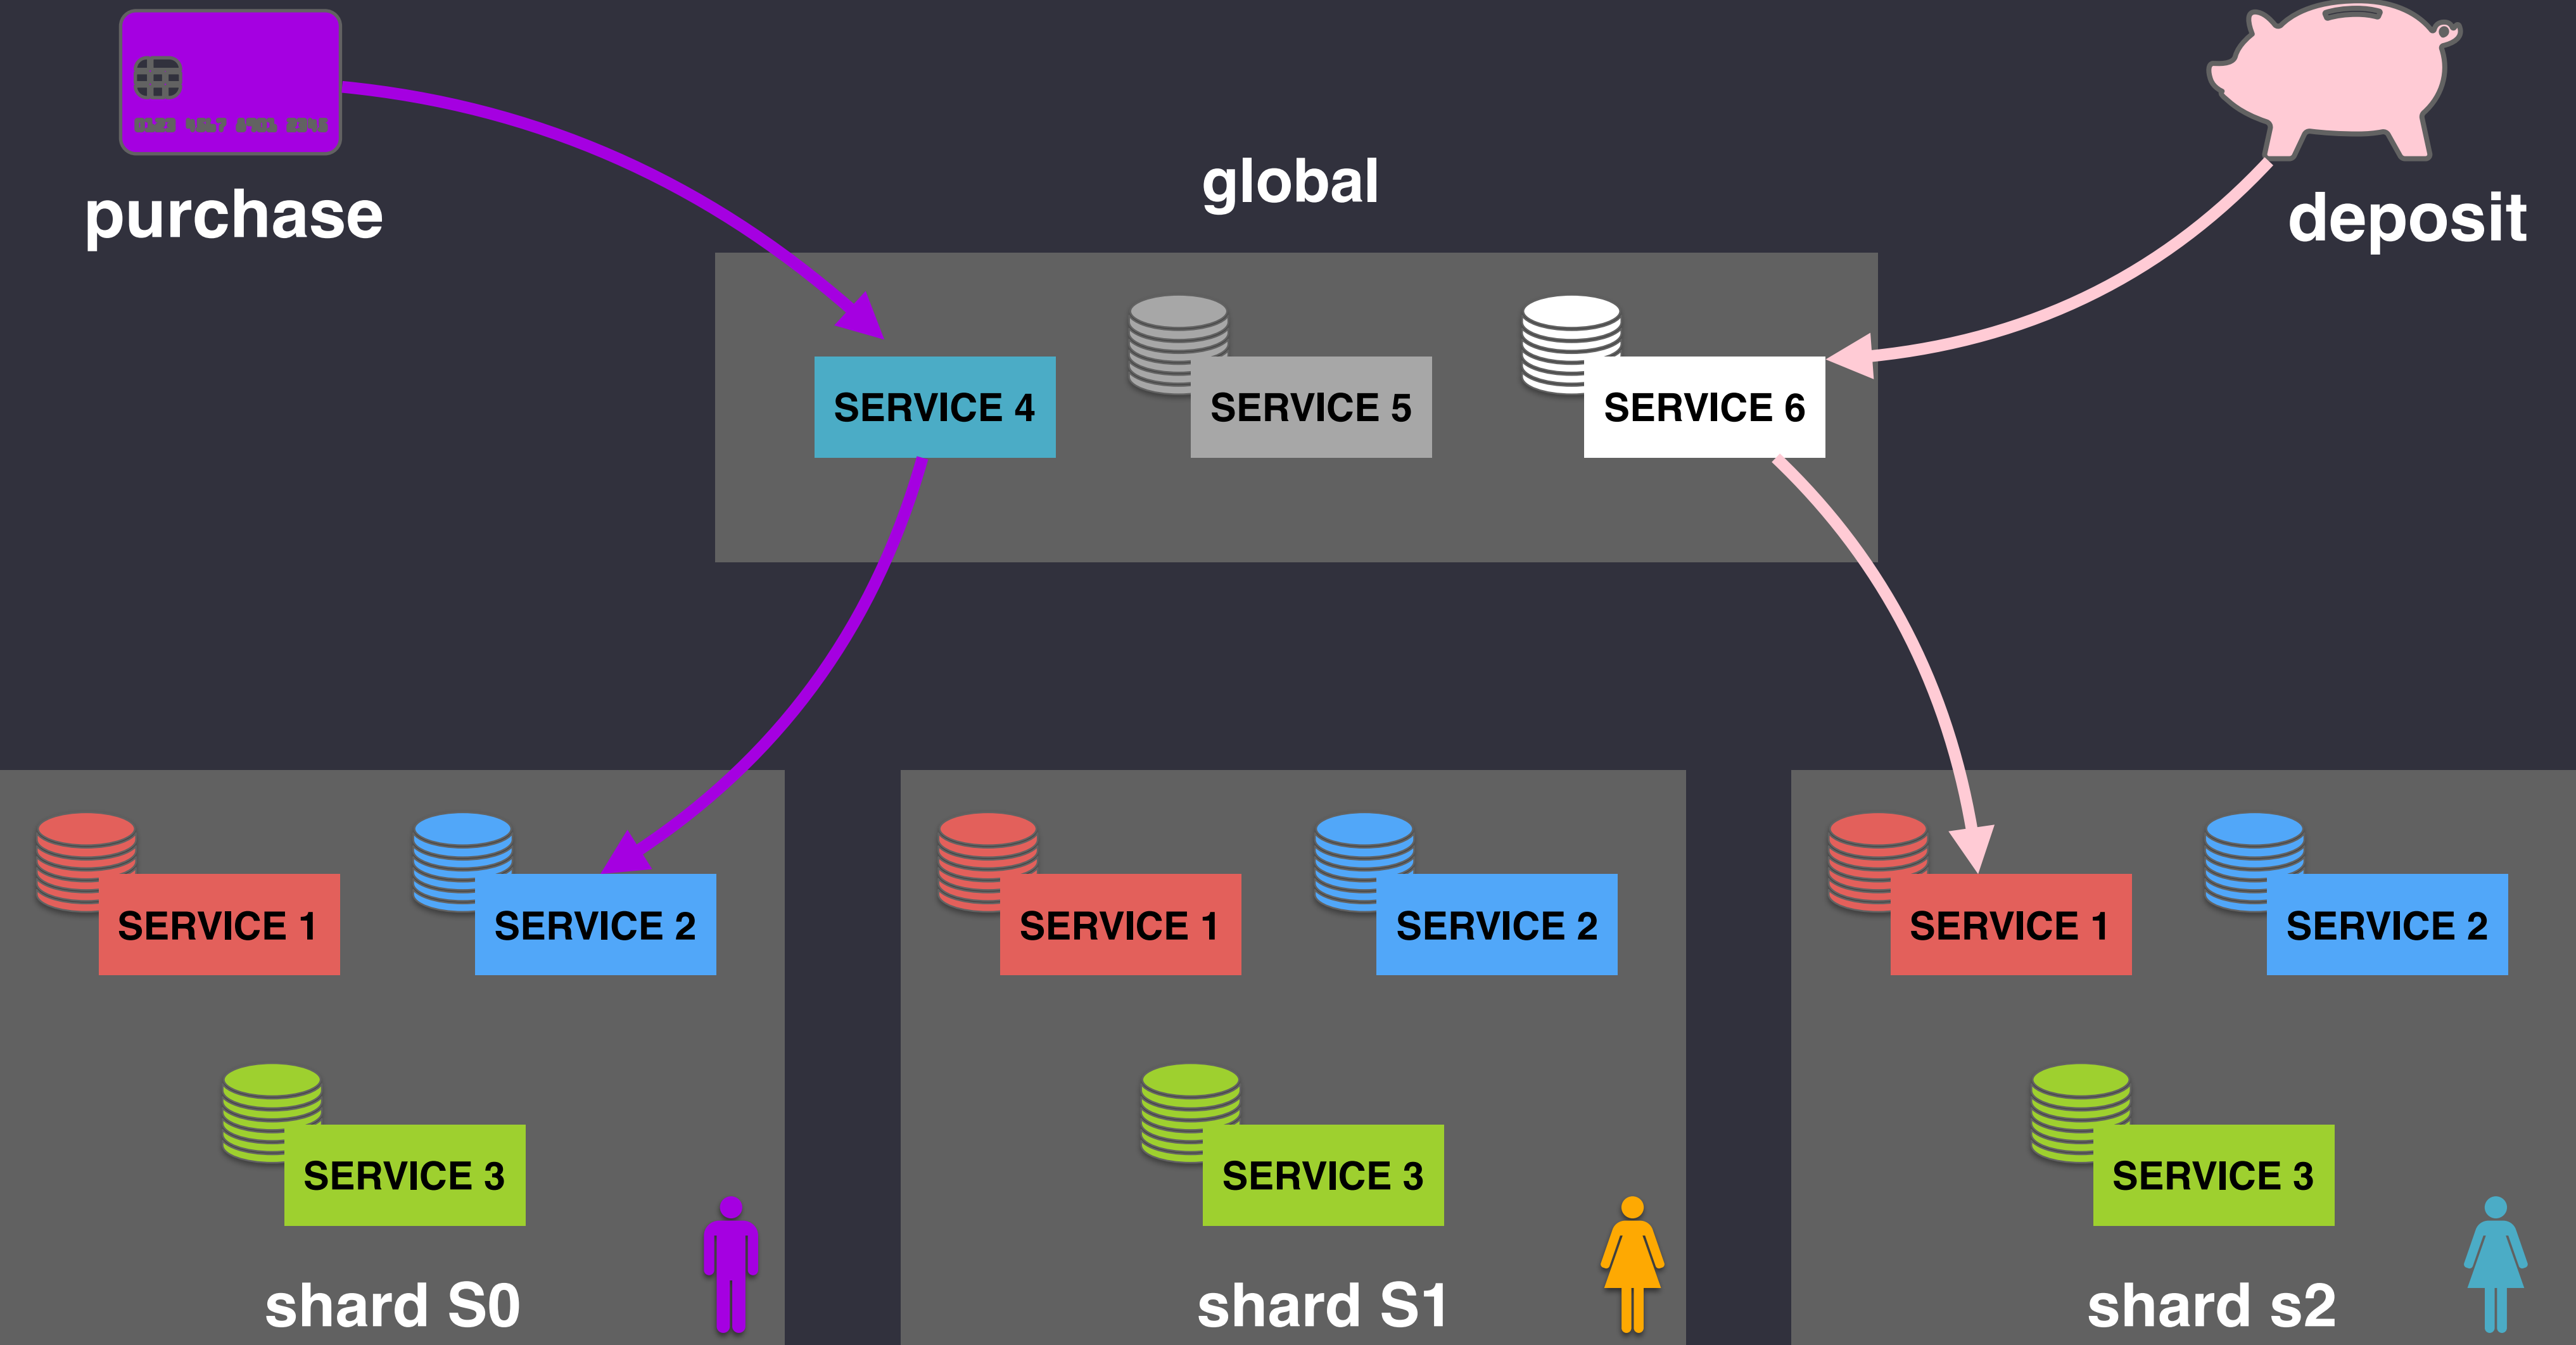


# OPTION #2: SCALABILITY UNITS

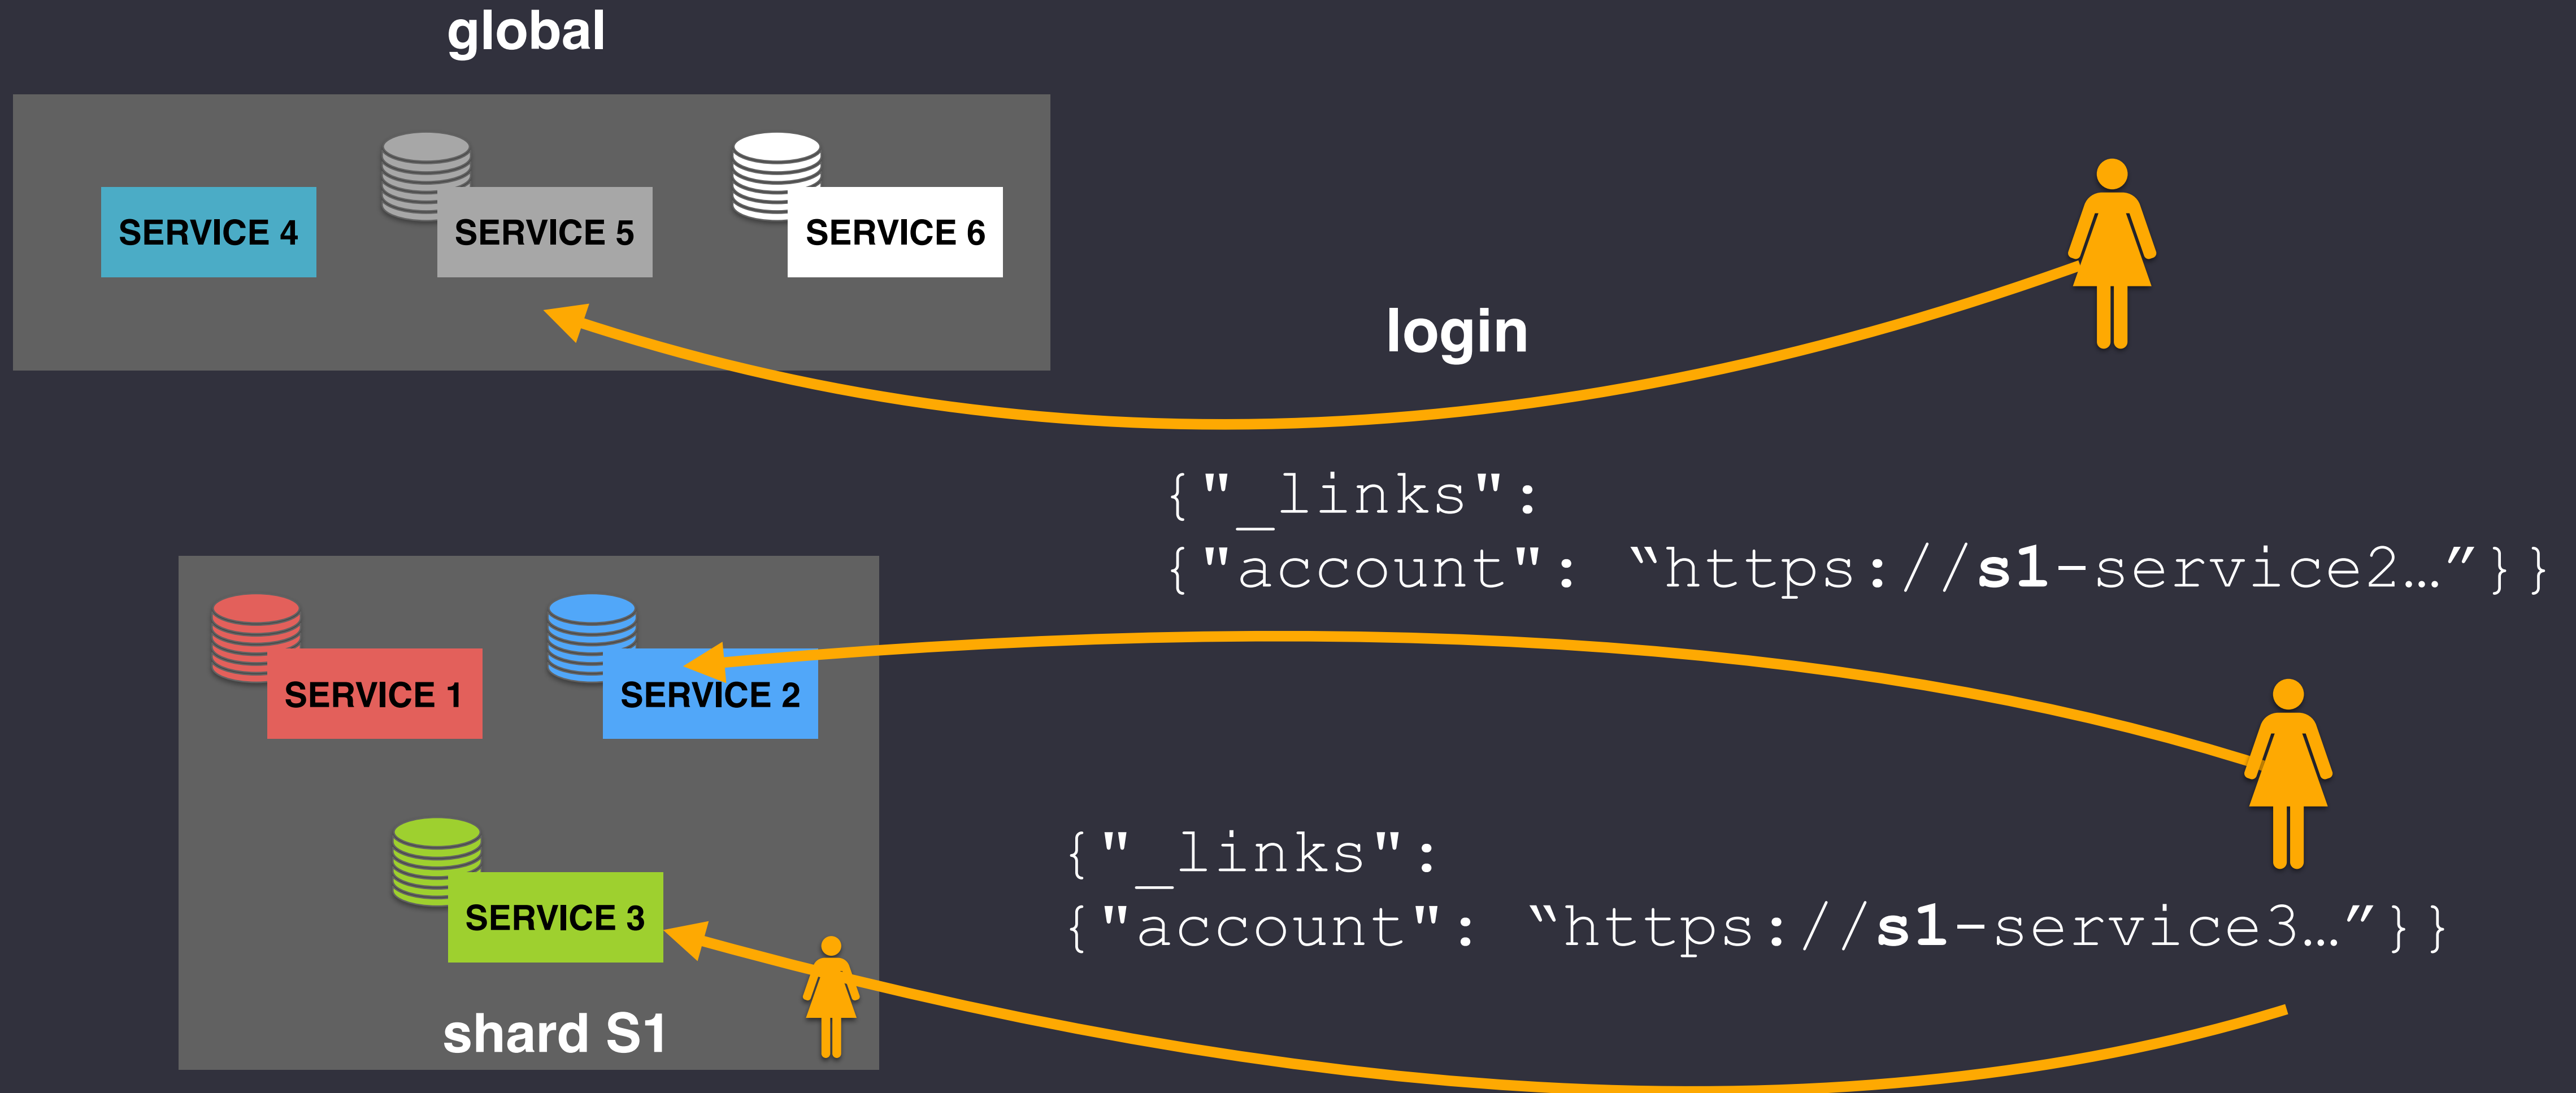


...

# OPTION #2: SCALABILITY UNITS + GLOBAL ROUTING



# OPTION #2: HYPERMEDIA FOR INTERACTIONS



# SCALING LESSONS LEARNED

## SCALABILITY UNITS WORK

works in practice, but difficult to move incrementally in that direction

## AUTOMATED IMMUTABLE INFRA

made this process much more tractable

## START EARLY

sharding was a complex project  
exponential growth defies intuition: use real growth models for planning

## BEWARE HOTSPOTS

business logic may create hot spots  
reactivated old prospects overcrowded s0

## MESSAGING AND HYPERMEDIA

provide critical flexibility for shard routing

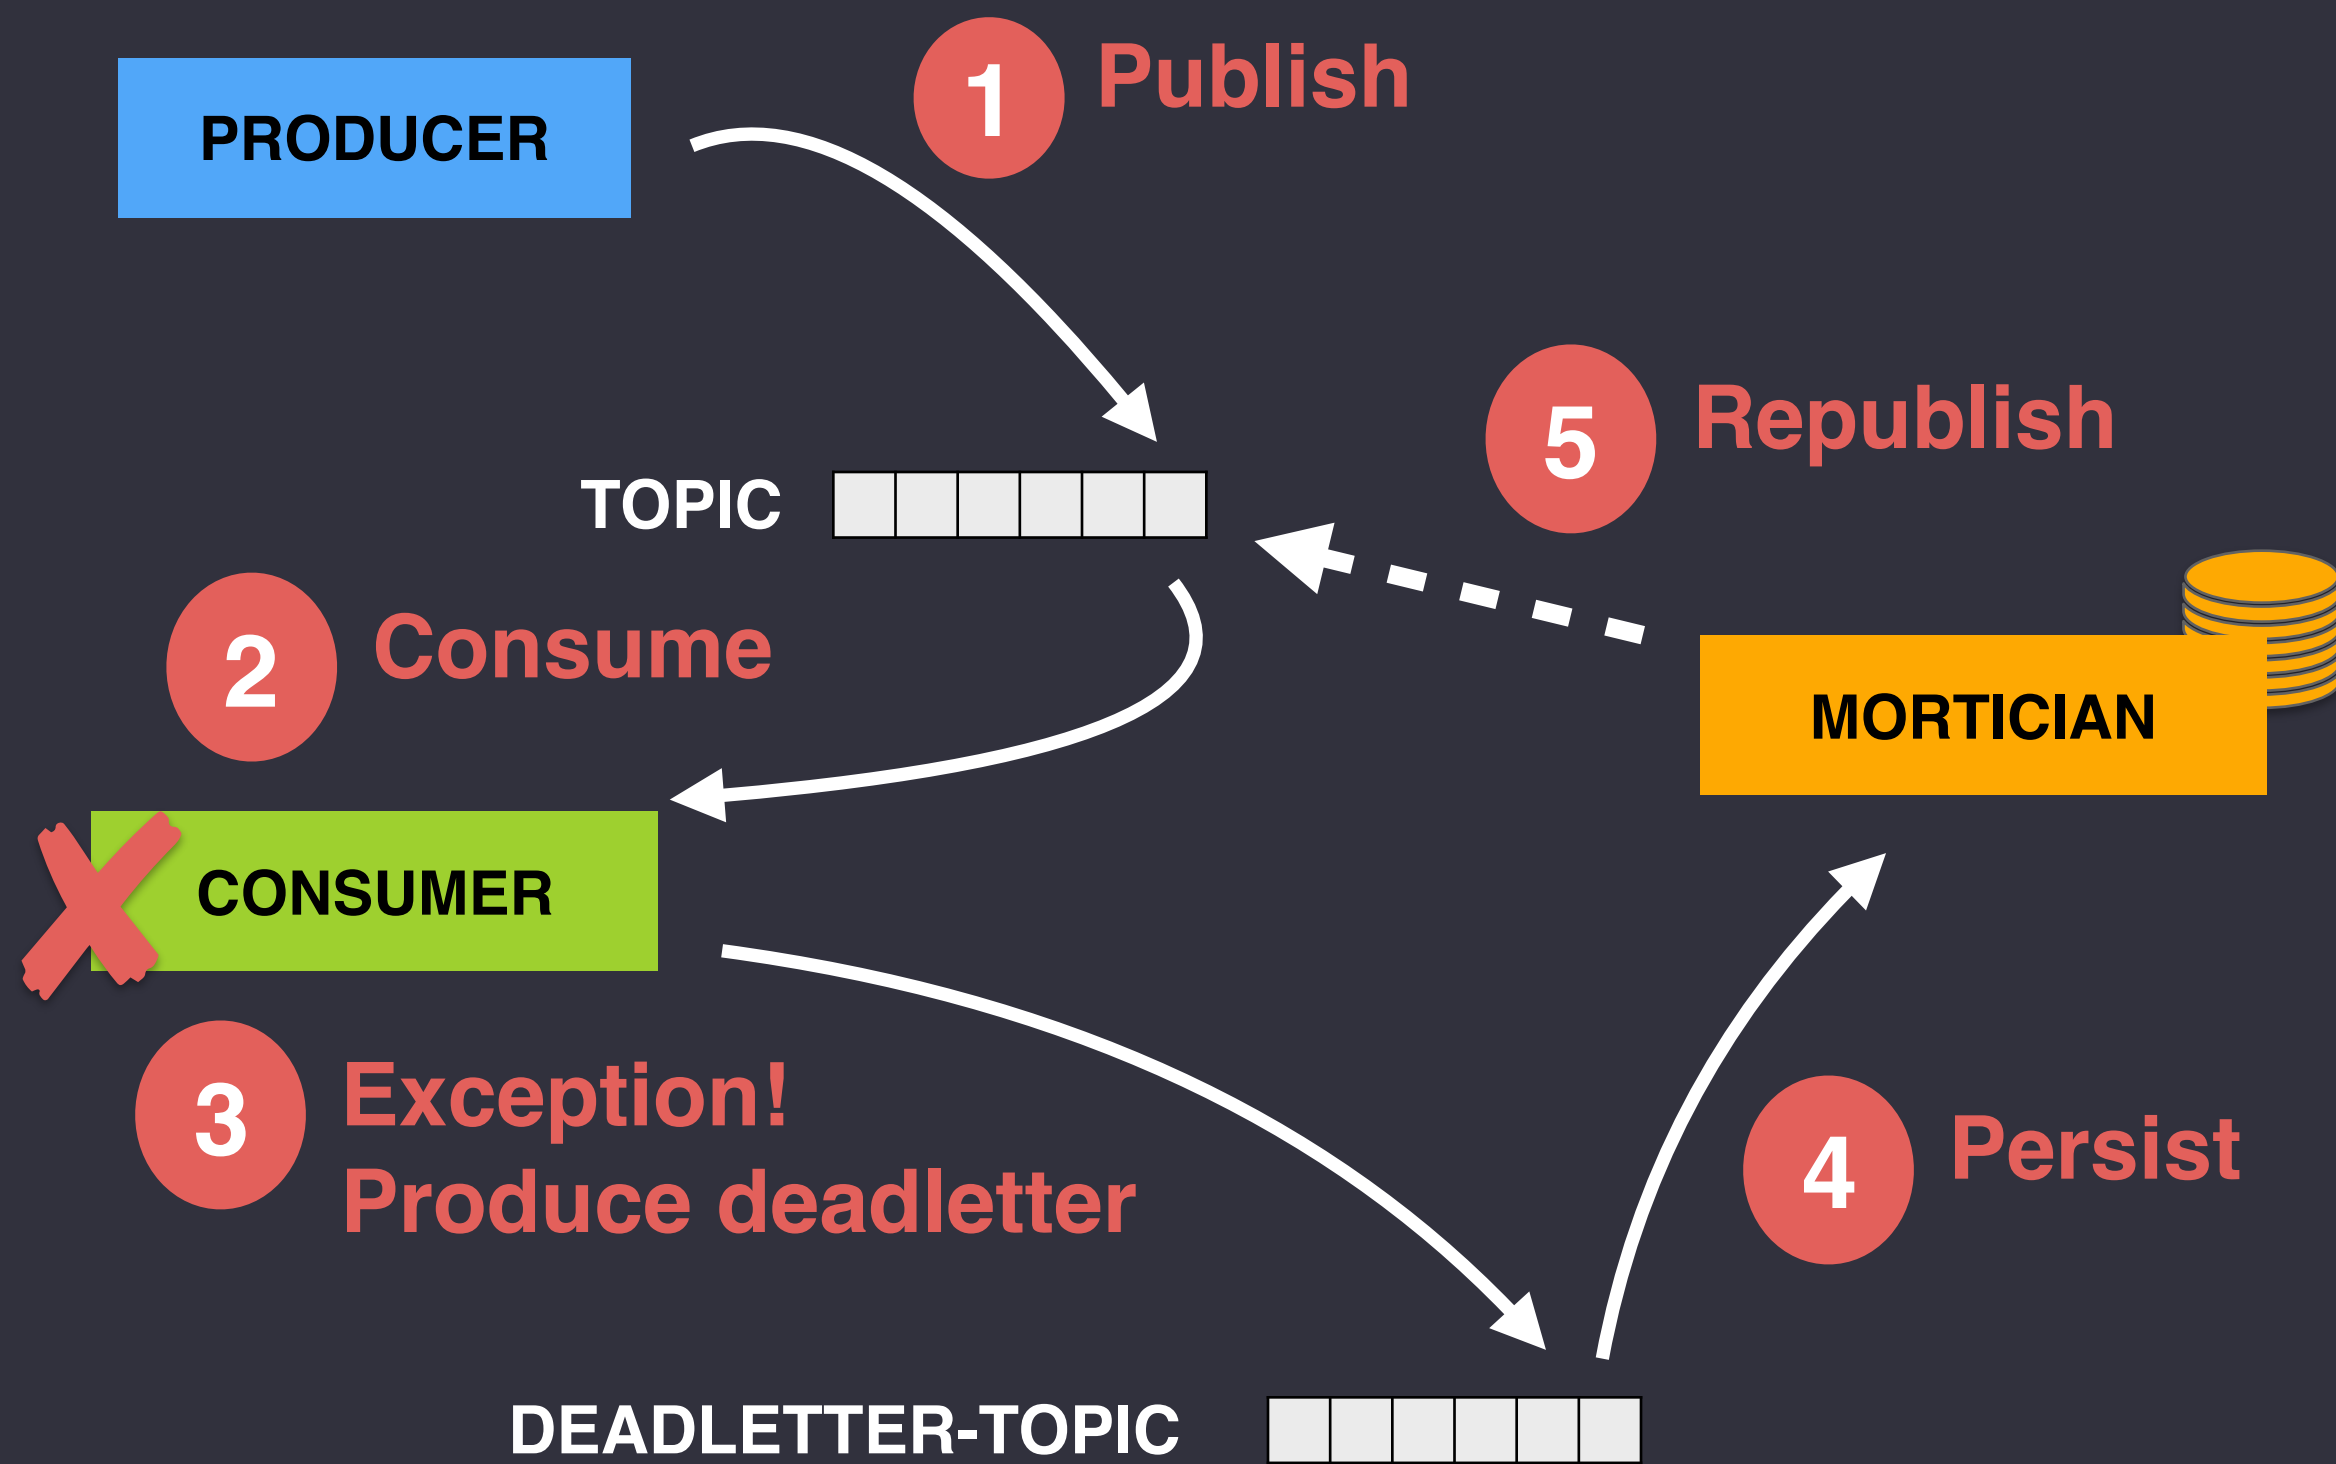
## SPLITTING EXISTING DATA

it's devilishly difficult (we avoided it, mostly)

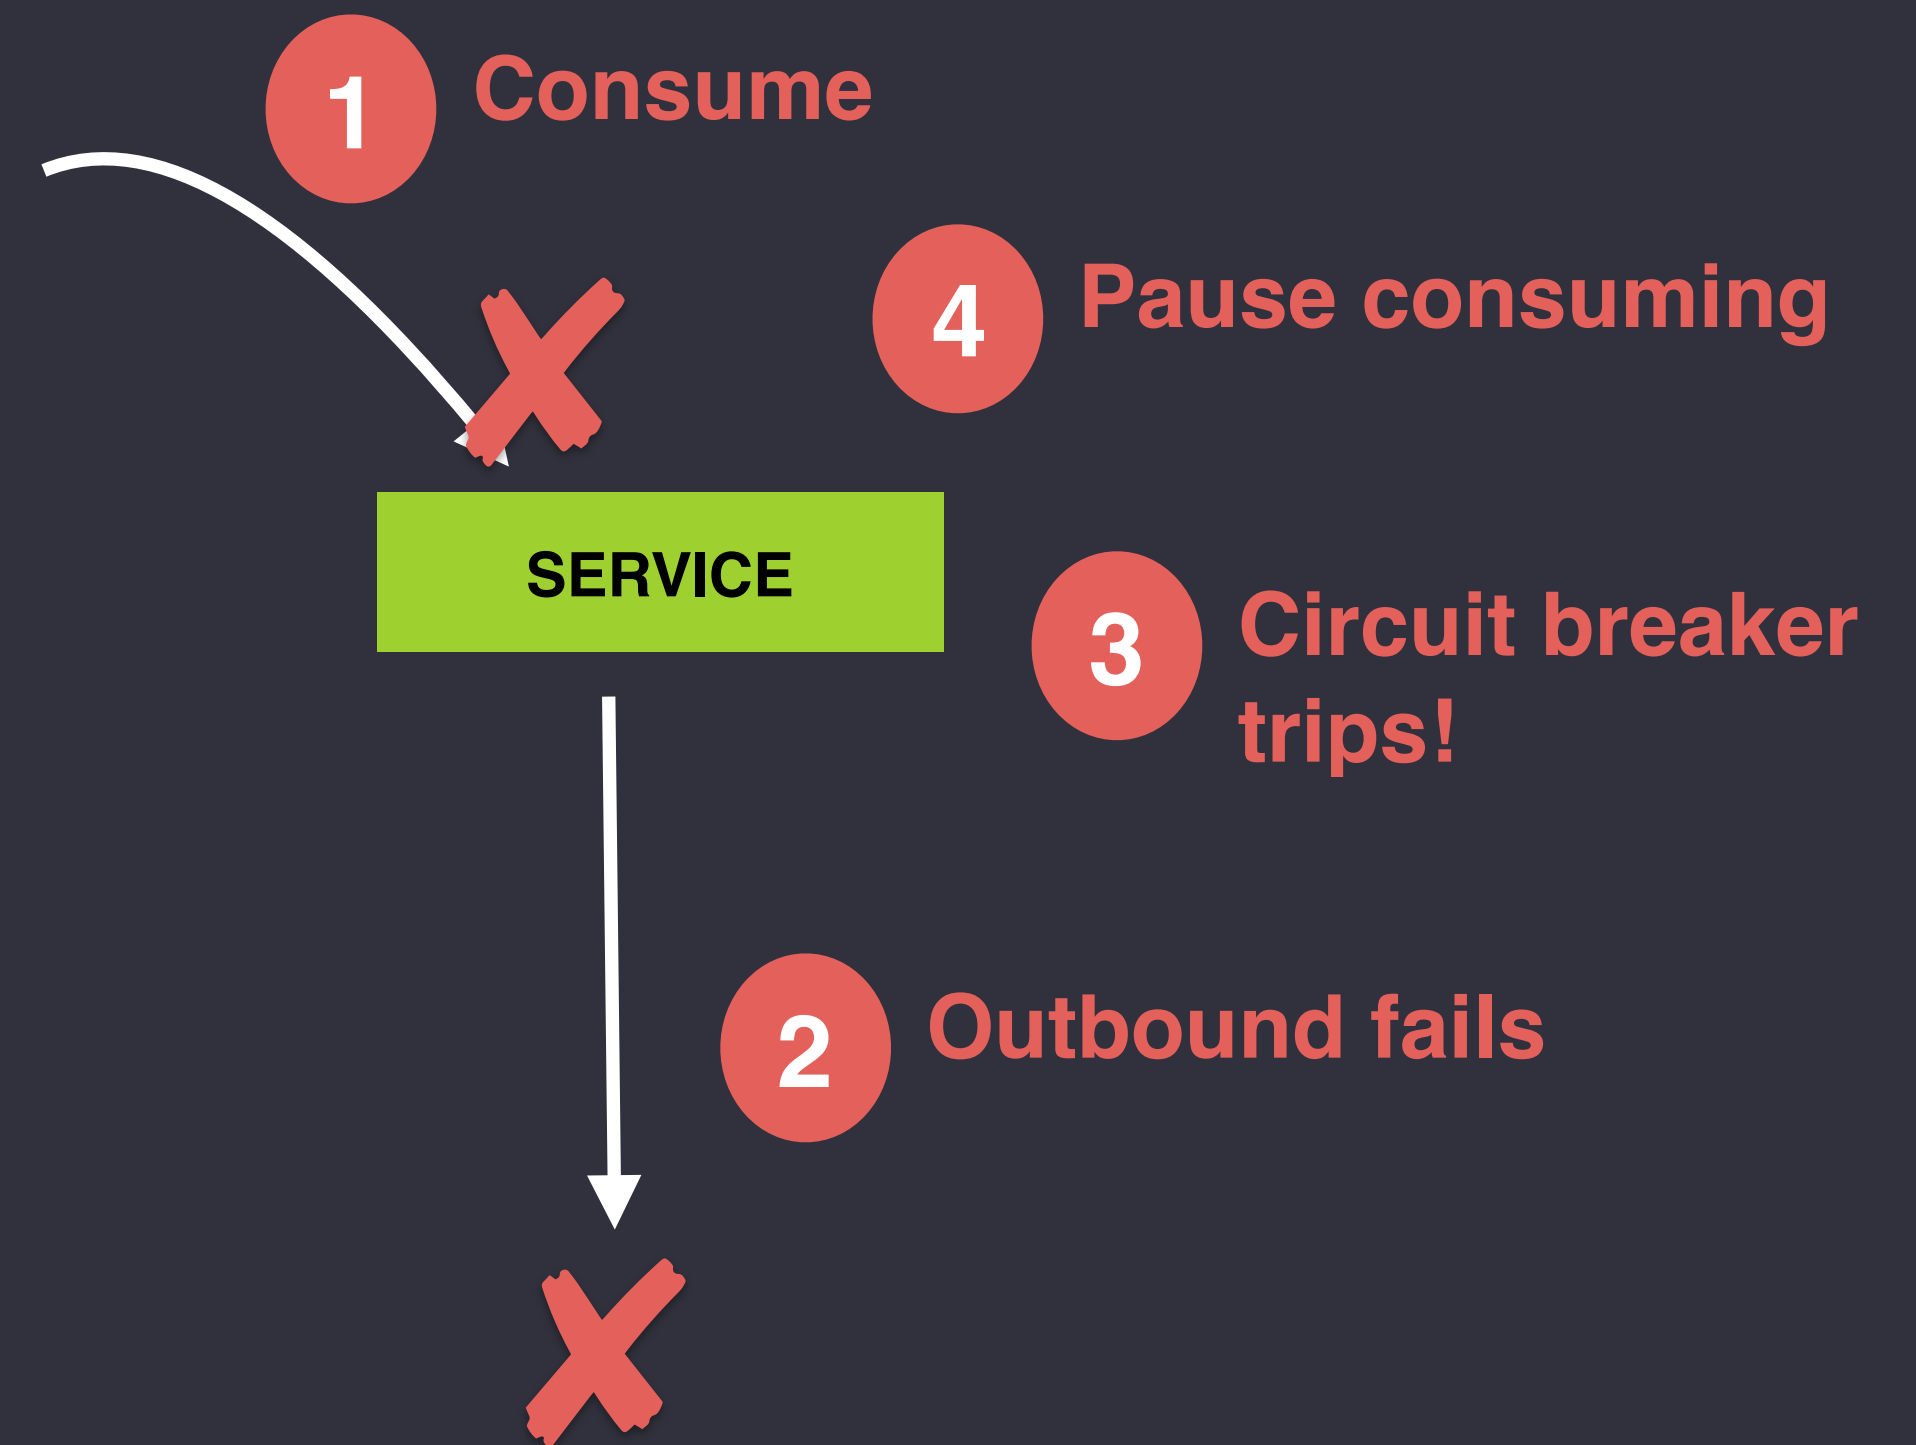
# FAULT TOLERANCE PATTERNS

Simple patterns for fault isolation and recovery

## DEADLETTERS



## CIRCUIT BREAKERS



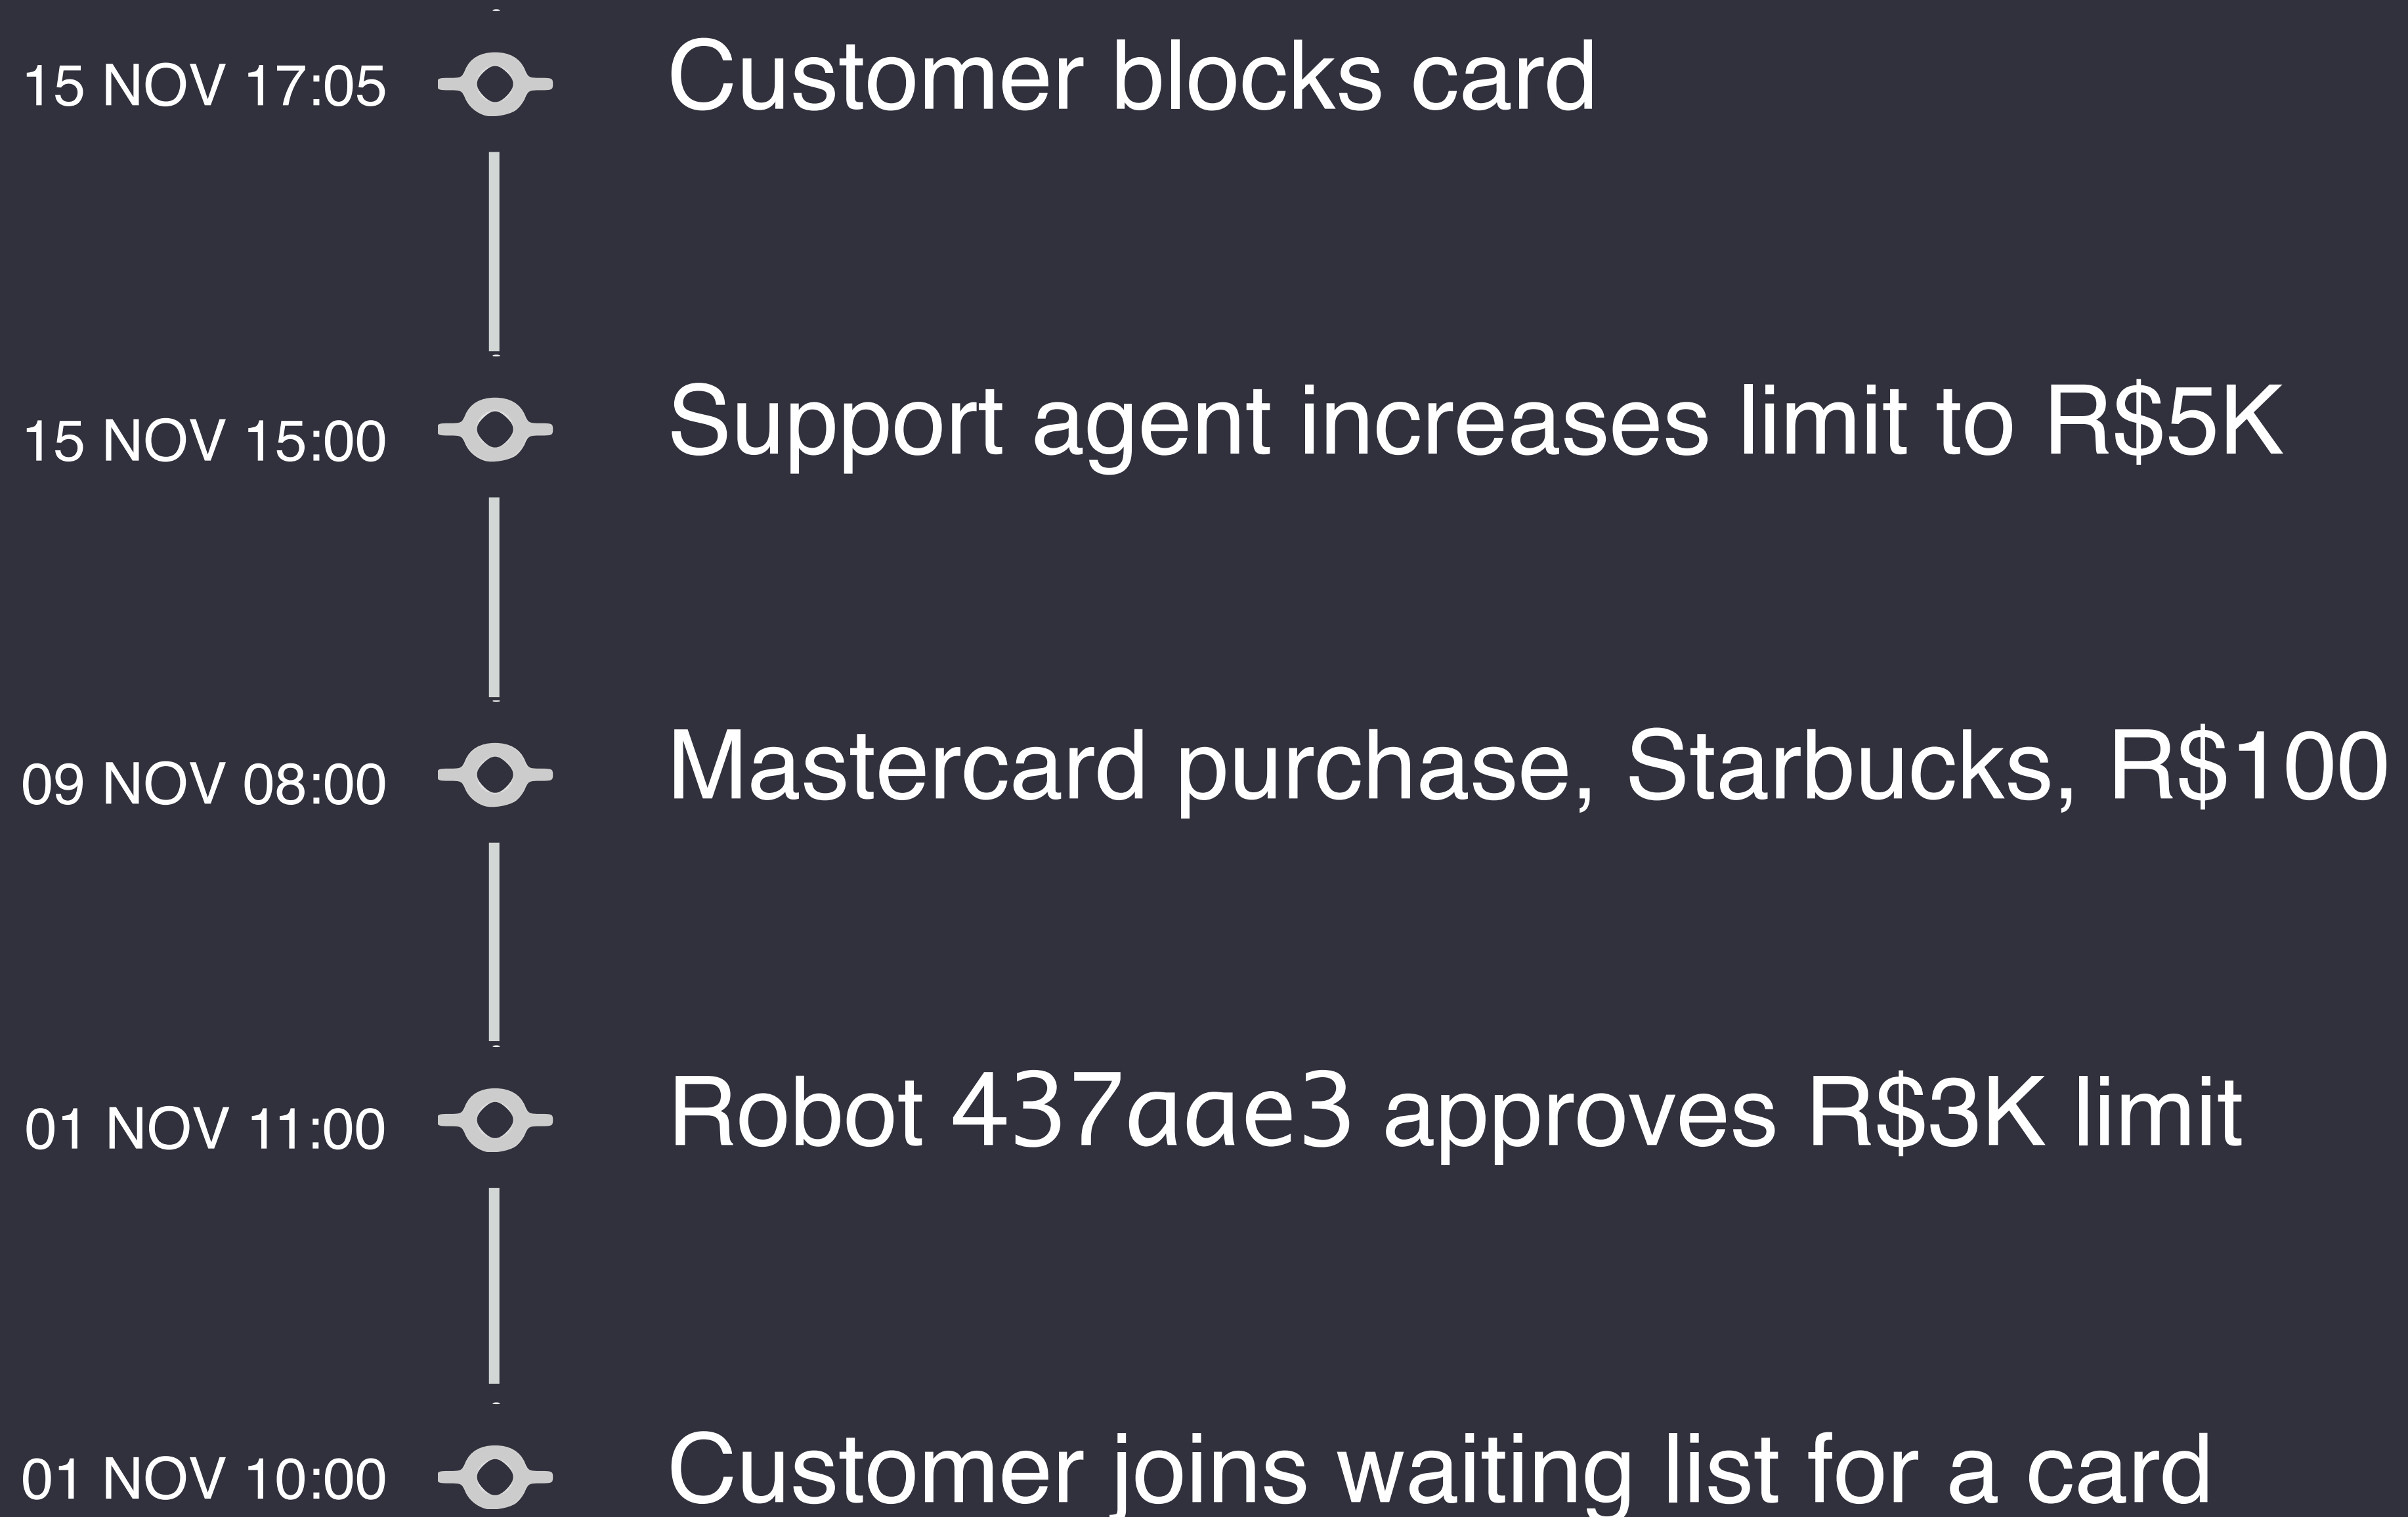


# ETL + THE ANALYTICAL ENVIRONMENT

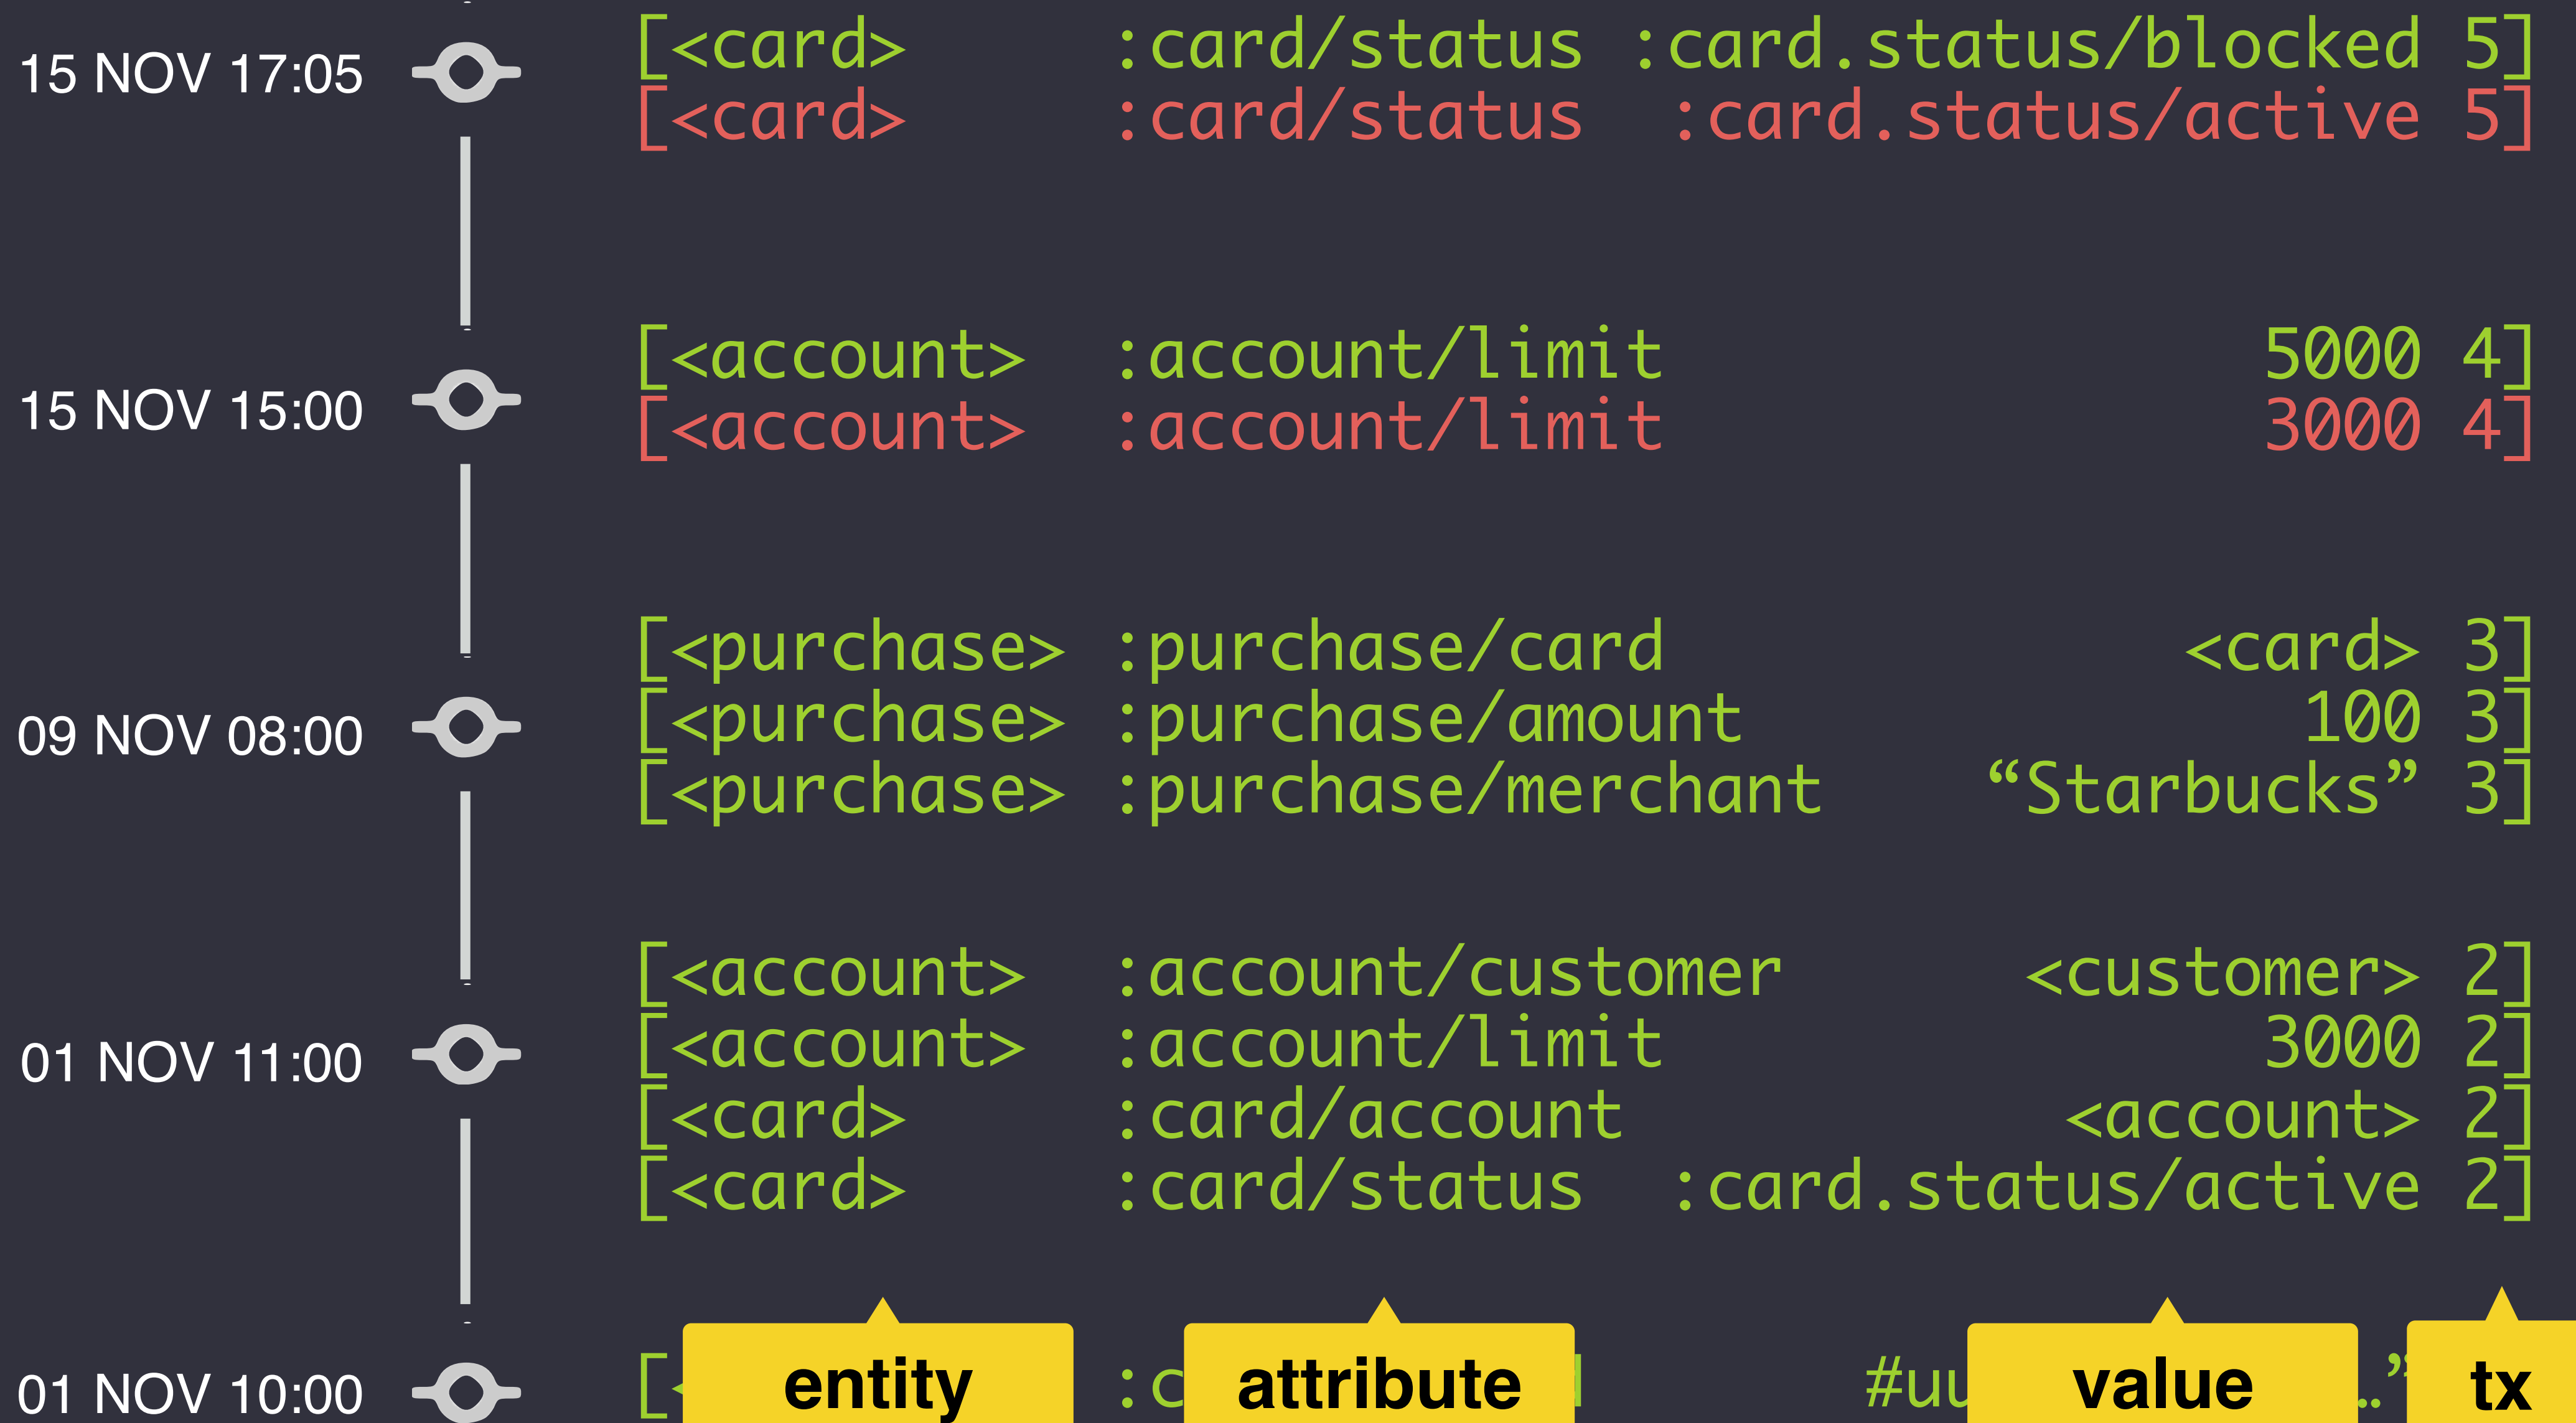


INFRASTRUCTURE

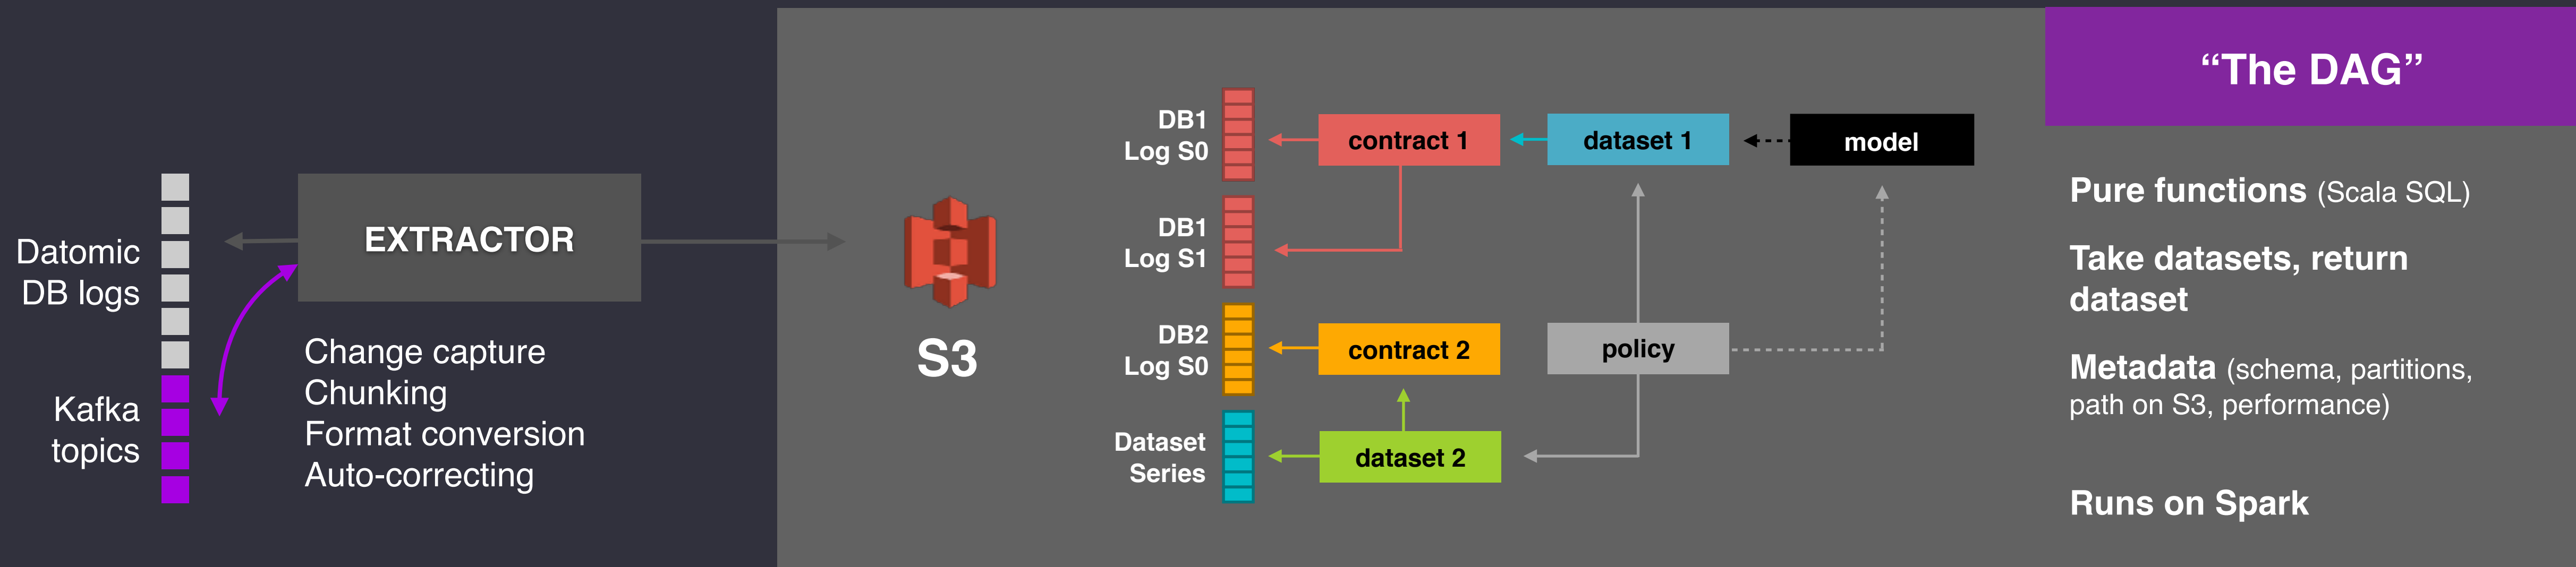
# DATOMIC PRIMER: EVENTS OVER TIME



# DATOMIC PRIMER: FACTS OVER TIME



# EXTRACT, TRANSFORM, LOAD




Datomic and Kafka log extraction feeding our data lake (S3) in real time

Analytical schemas (“contracts”) generated from Datomic entities

Shards recombined into a logical table-per-entity incrementally

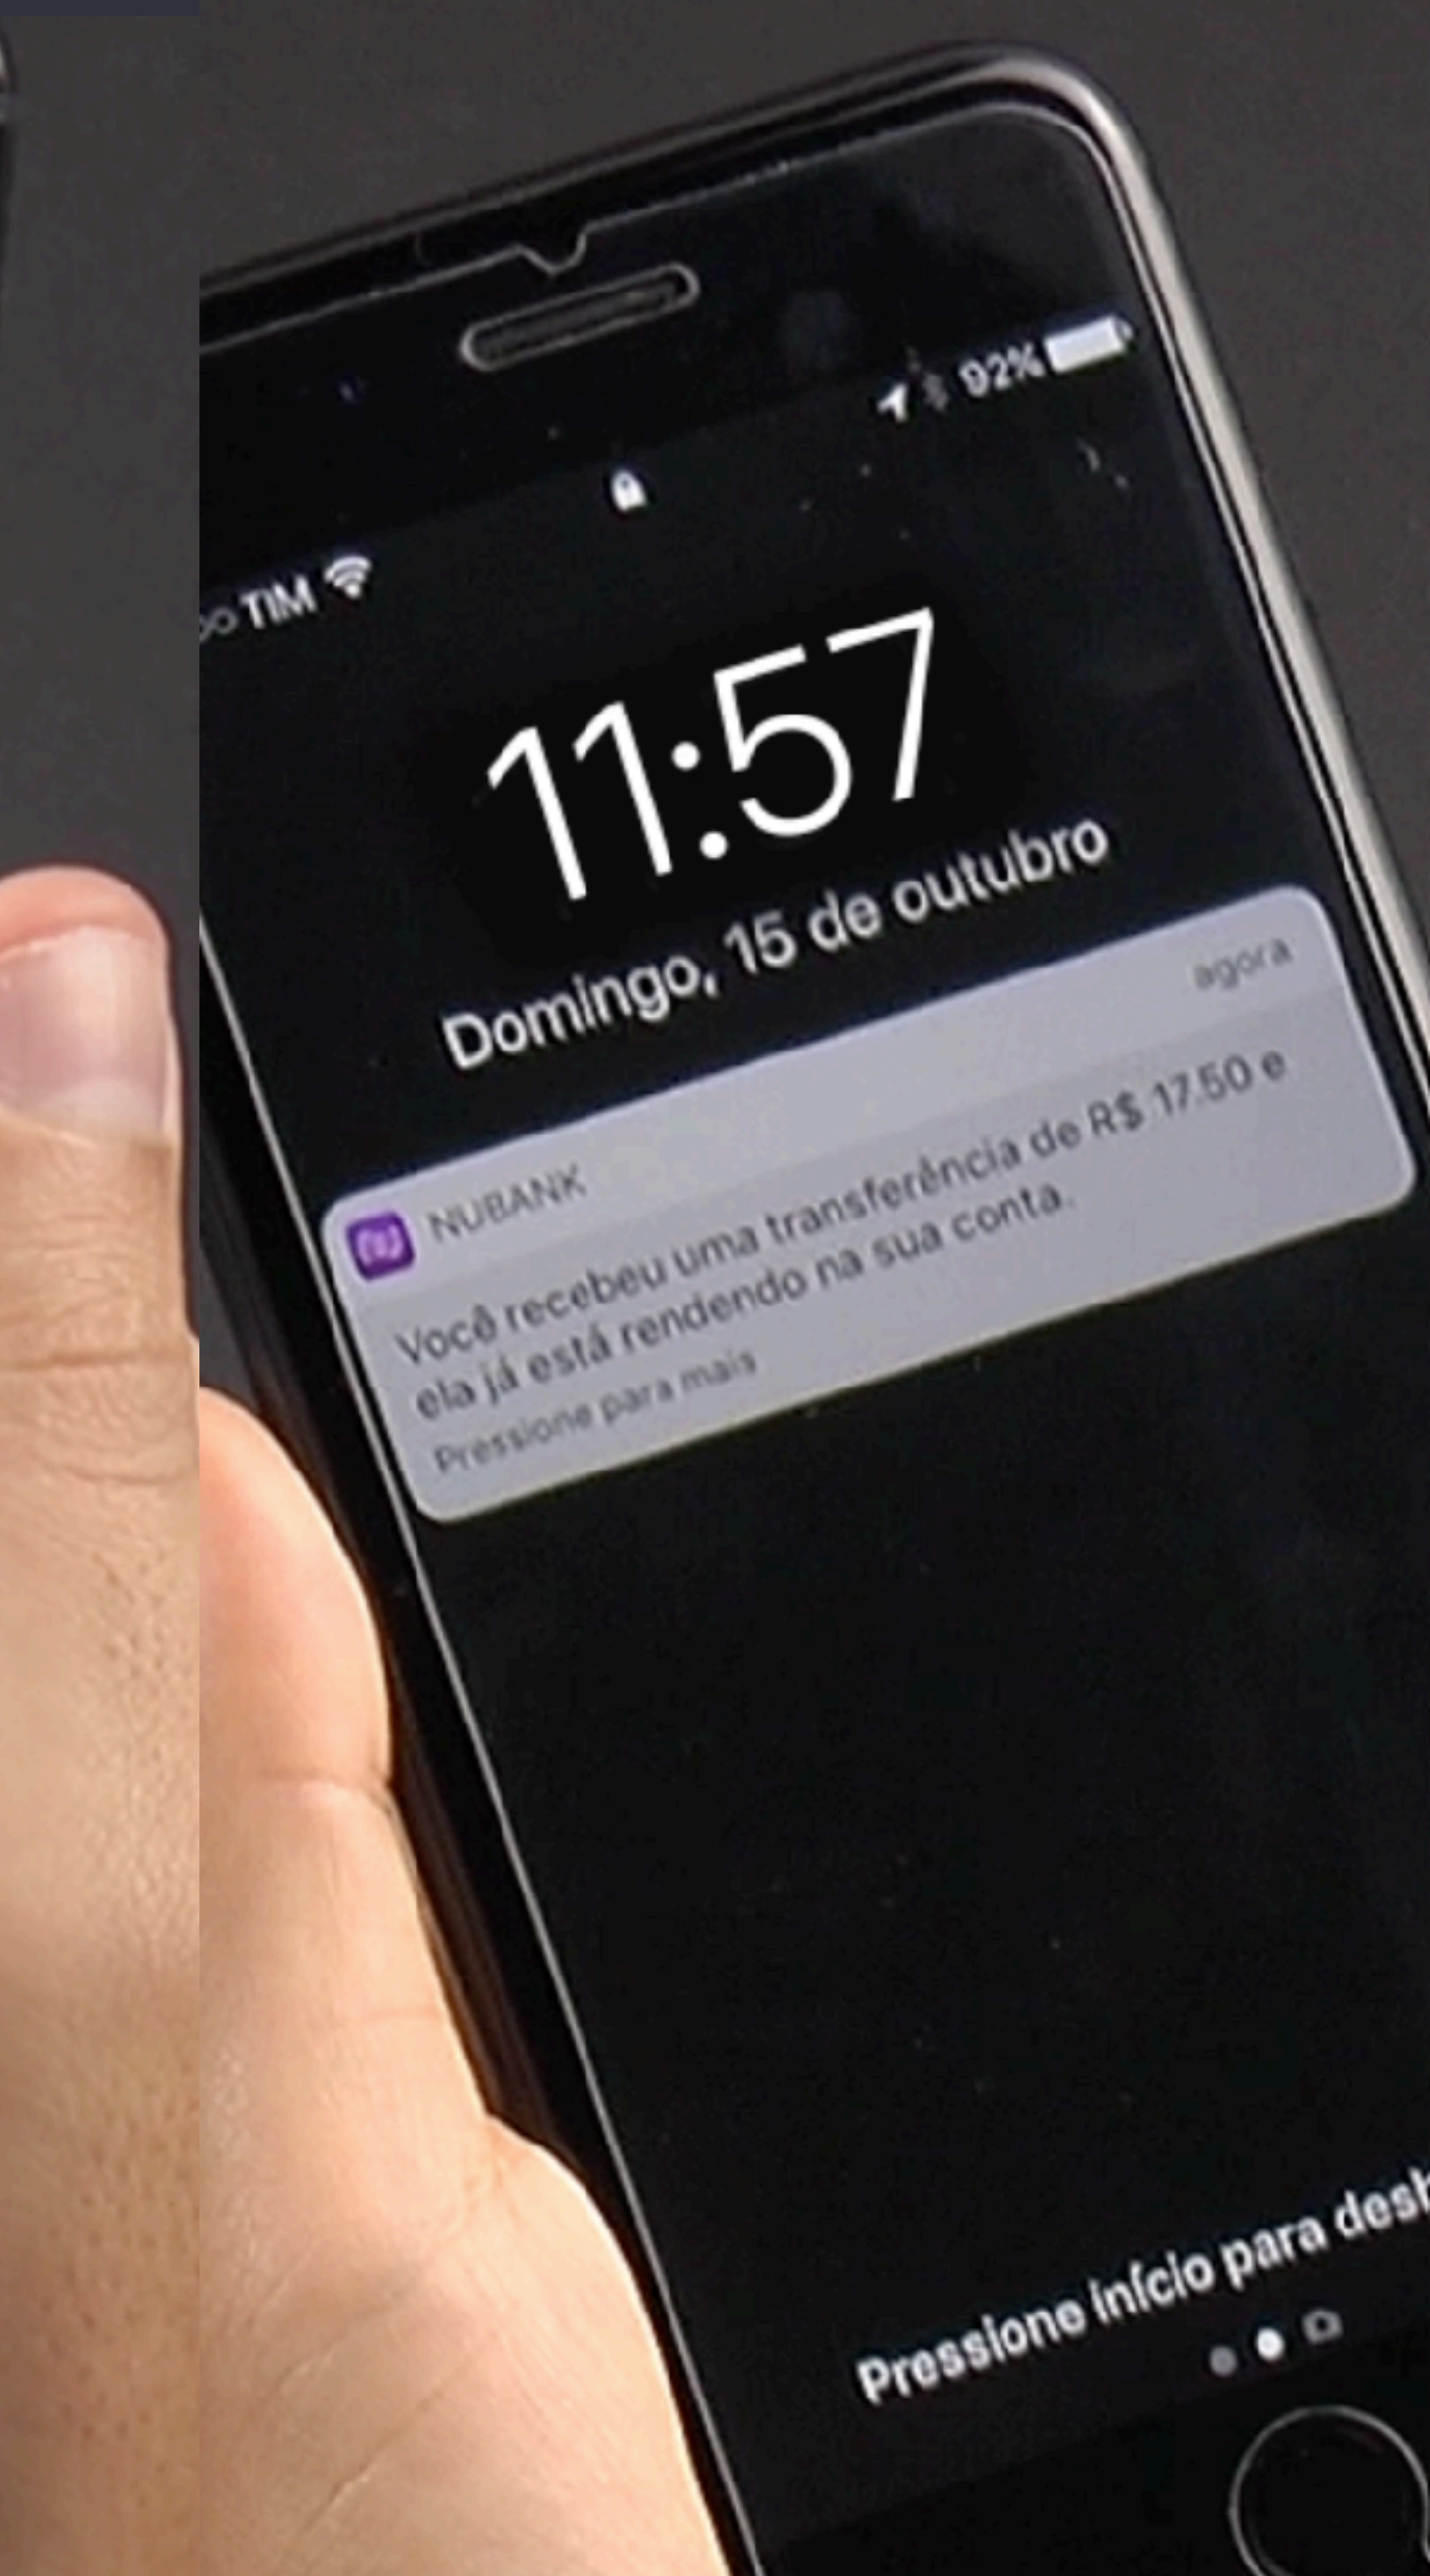
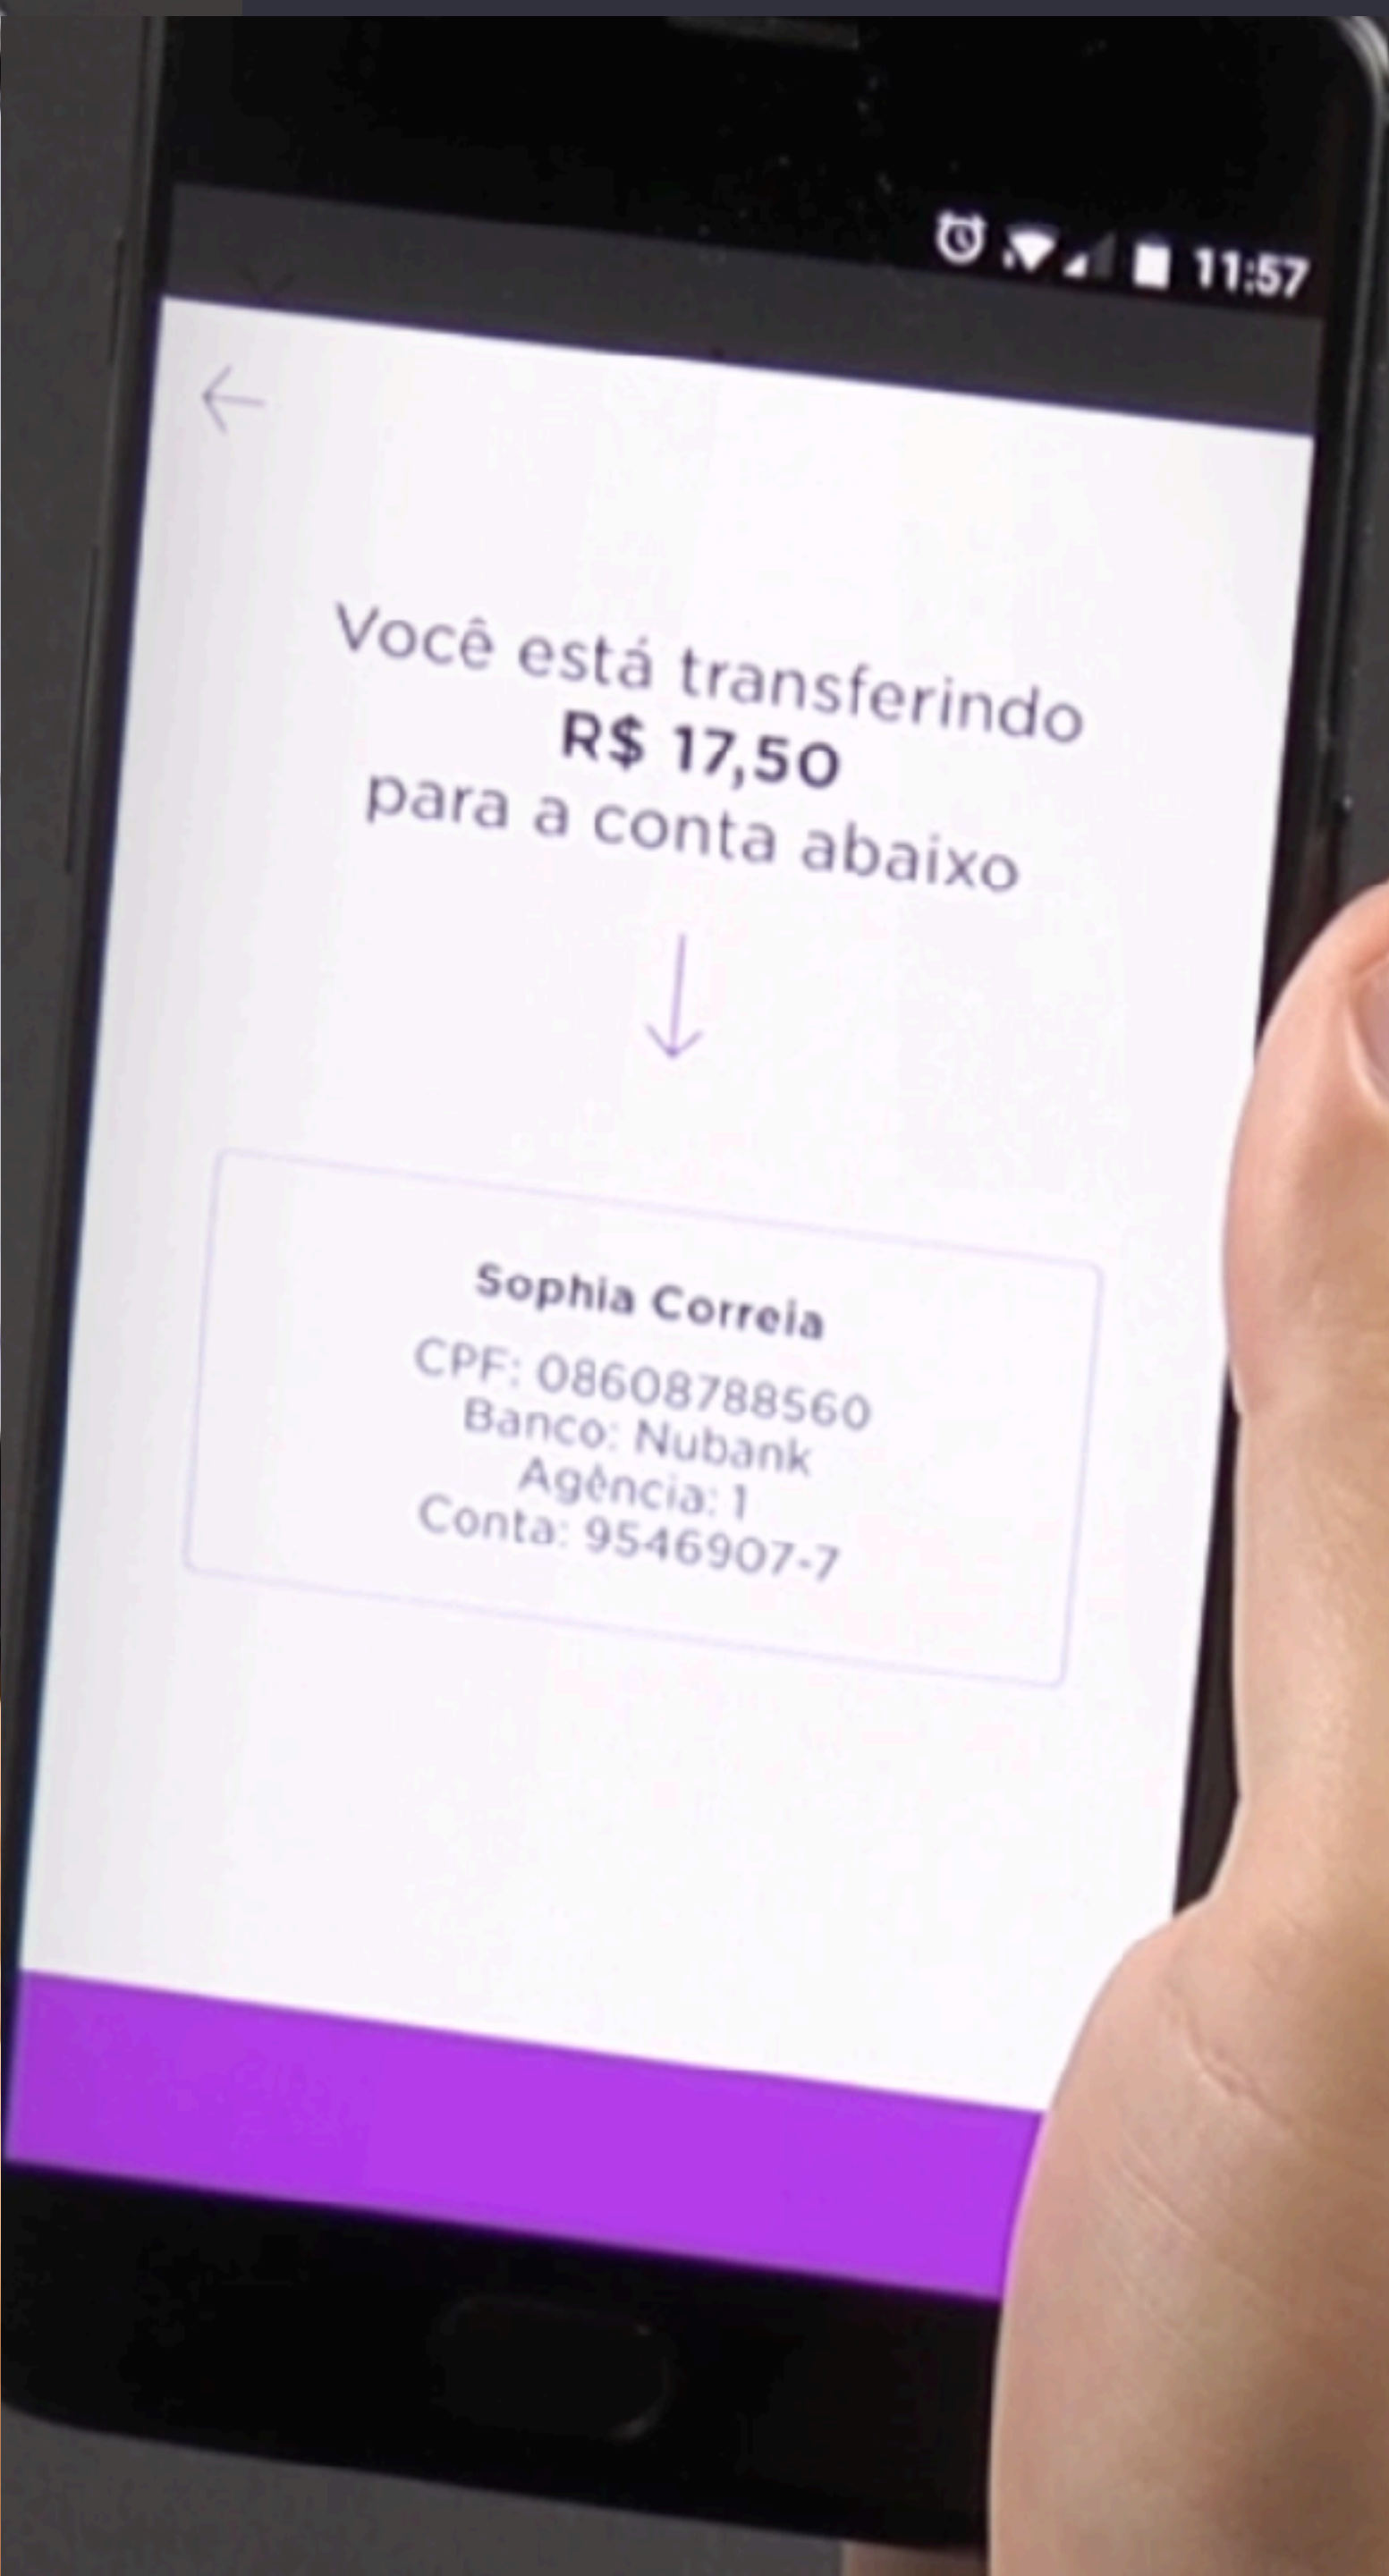
# ETL EXAMPLE: CONTRIBUTION MARGIN



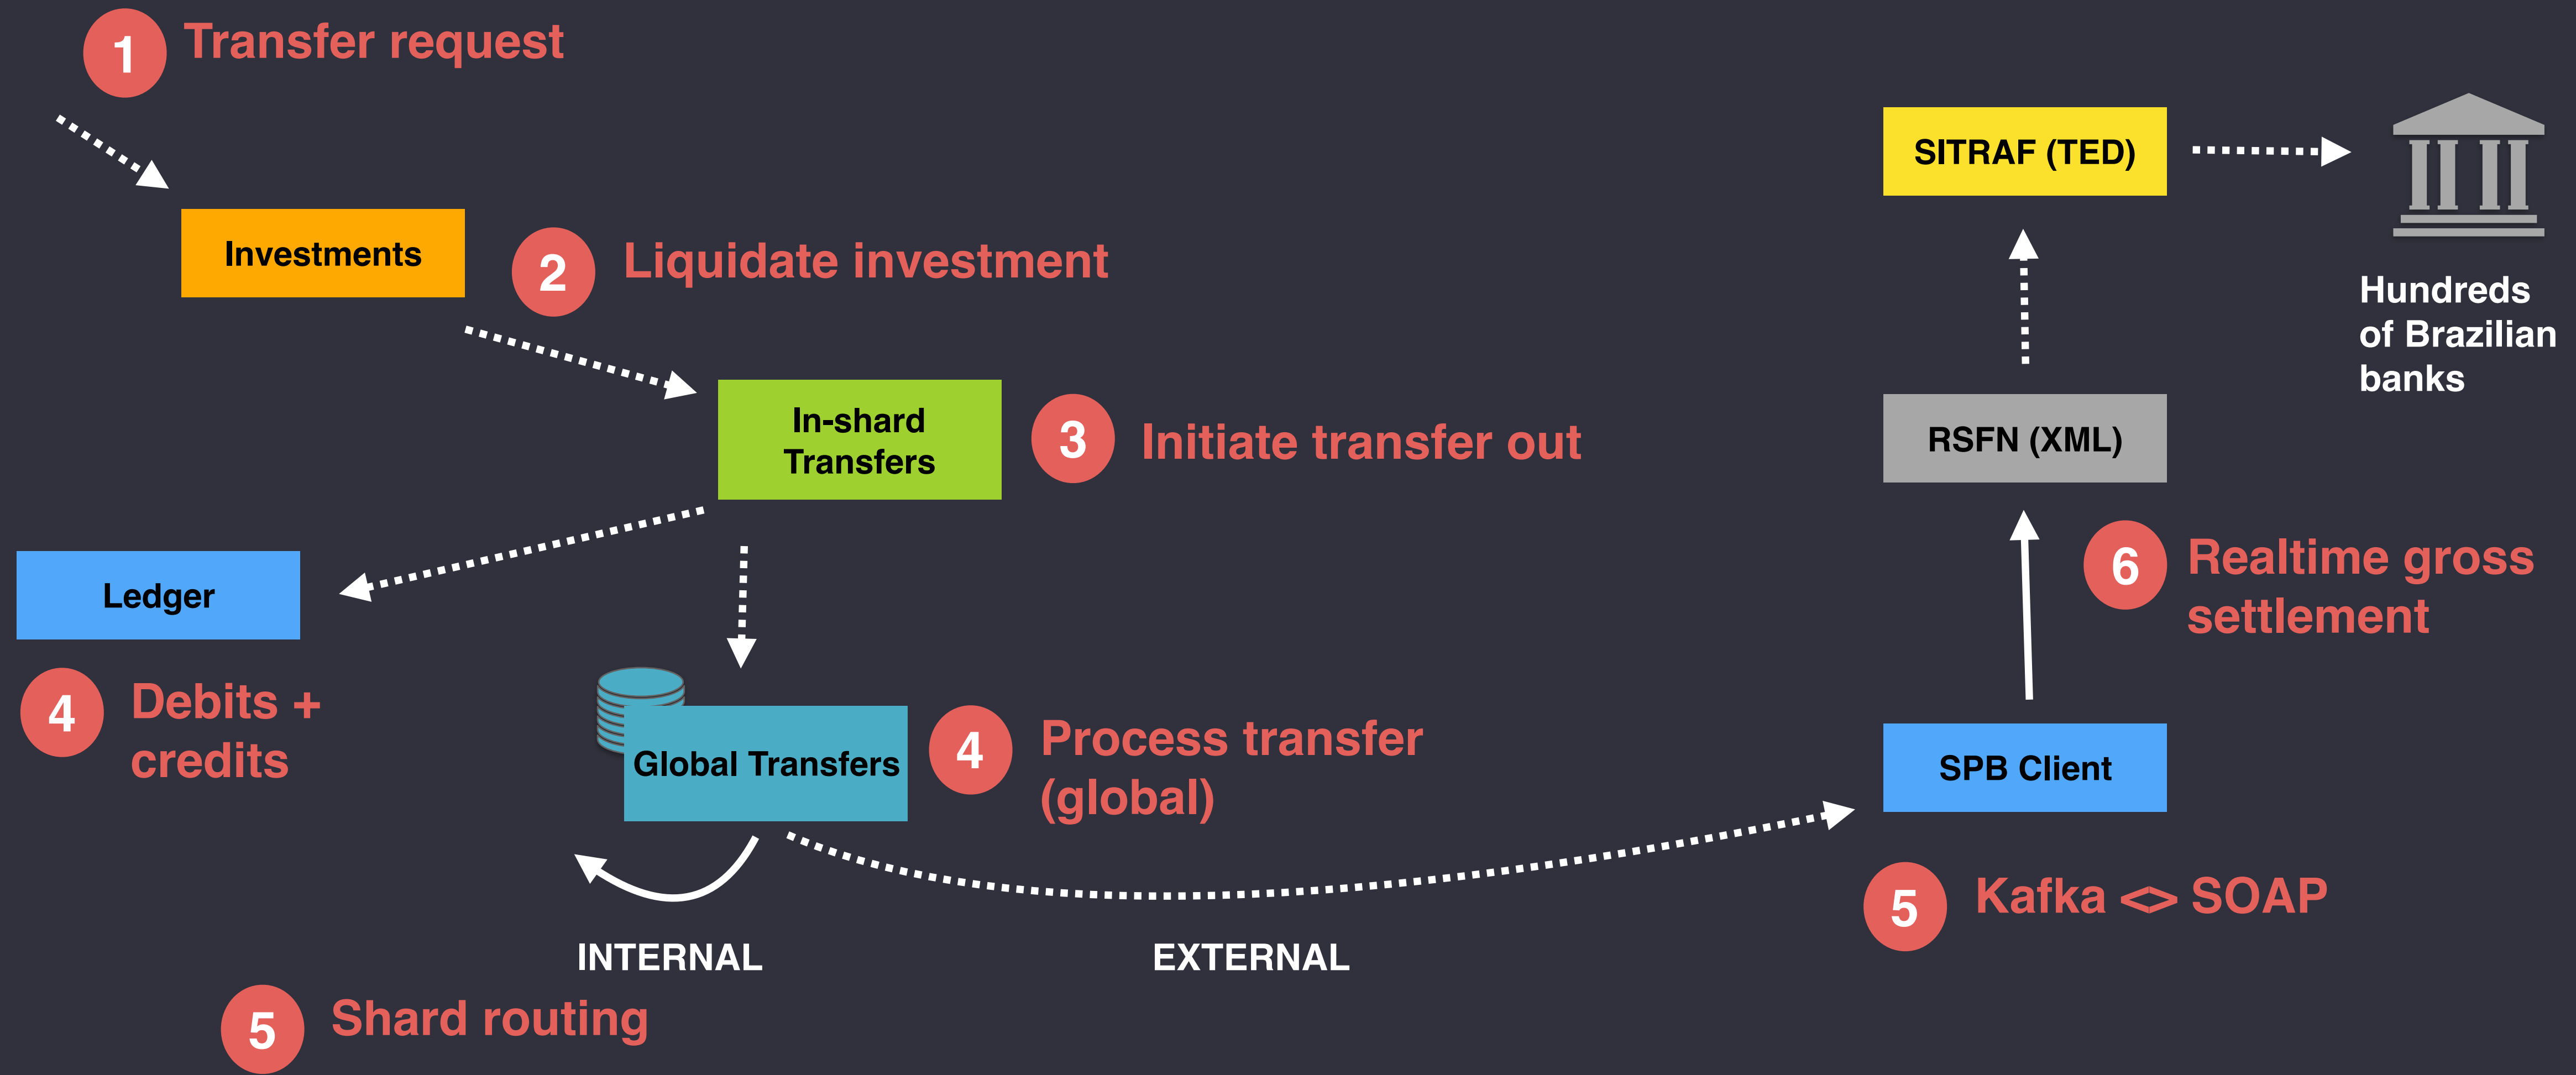
# REALTIME TRANSFERS



INFRASTRUCTURE



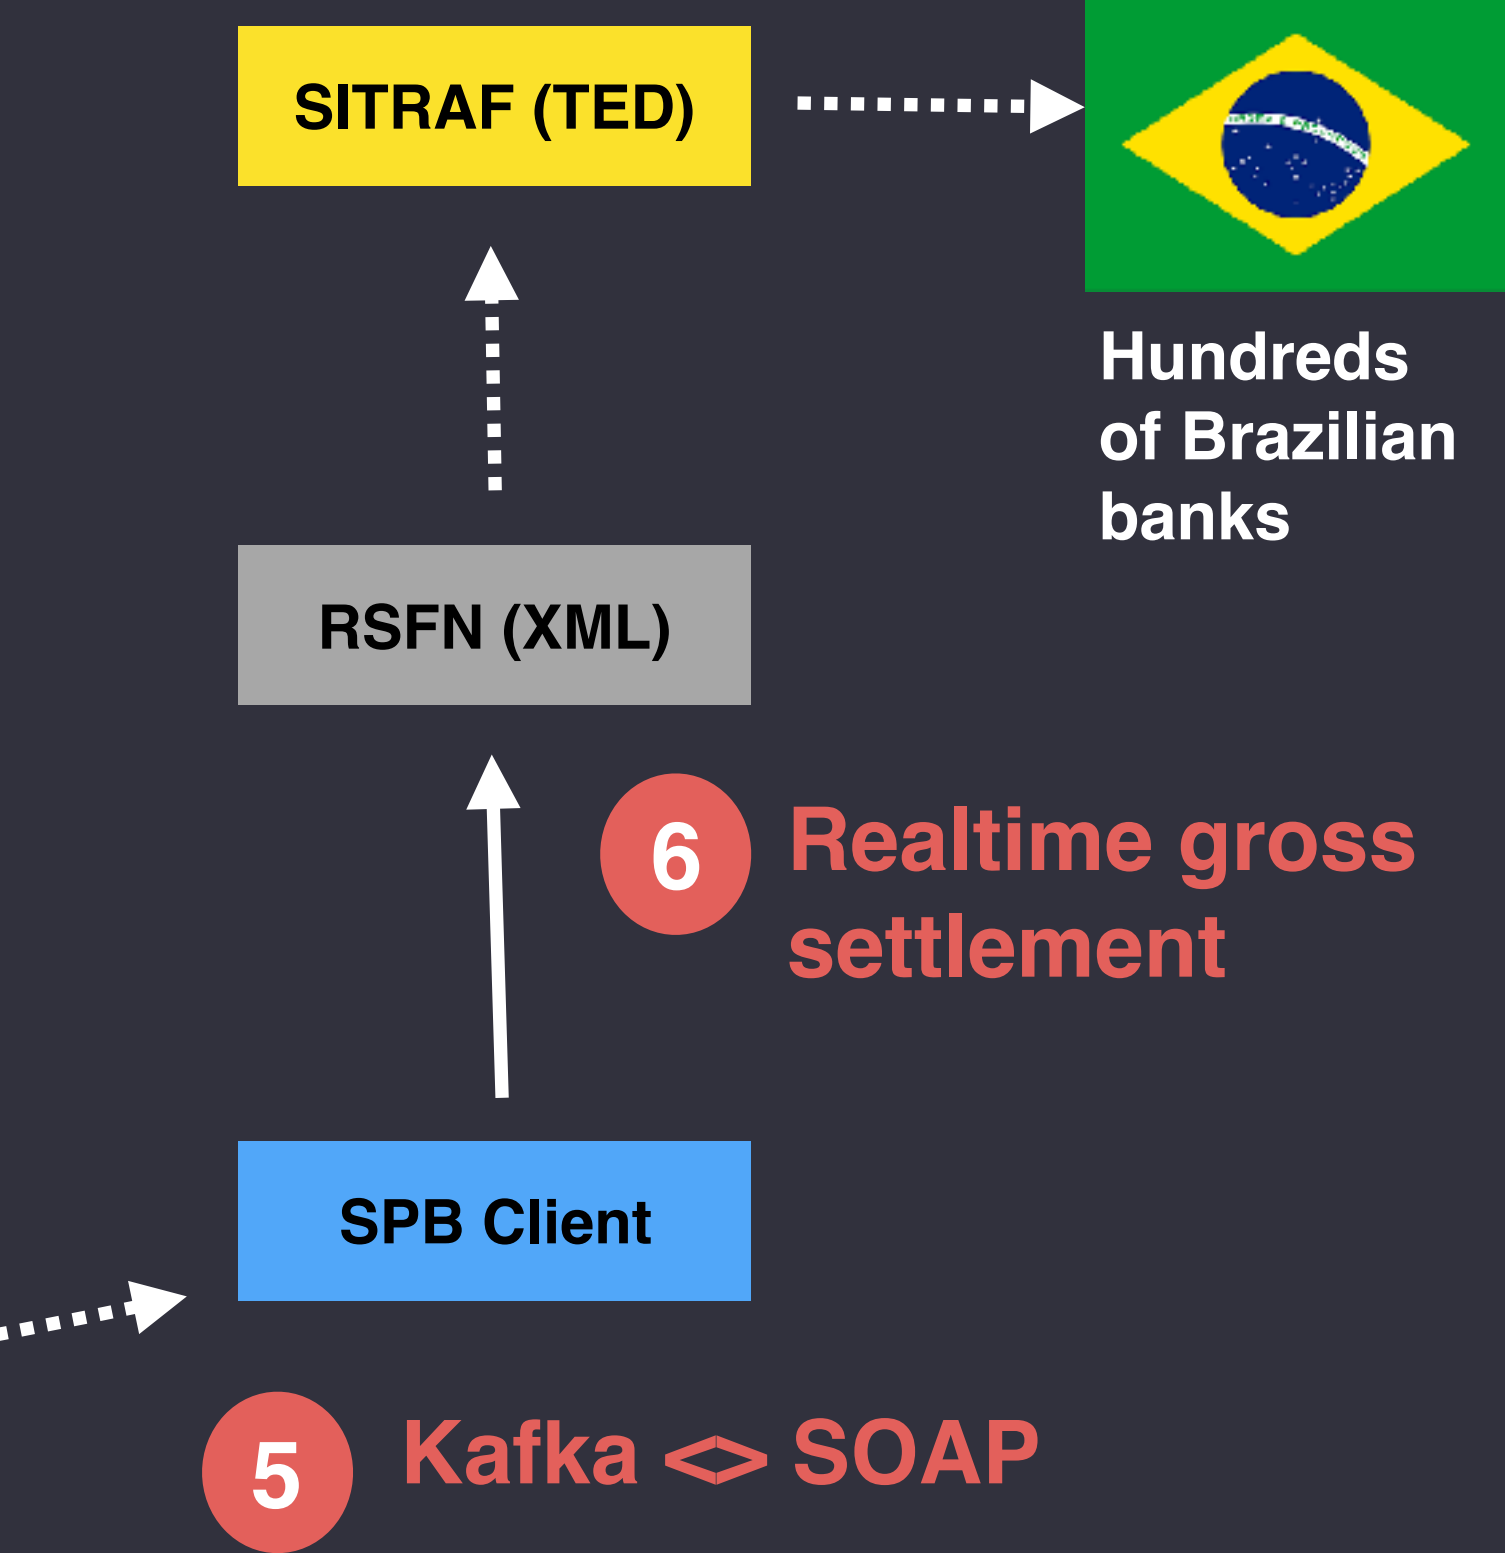
# REALTIME MONEY TRANSFER





# BRAZILIAN PAYMENTS SYSTEM

Hub and spoke model for national payments



Real time gross, irrevocable and unconditional settlement of unlimited amounts

~R\$1 trillion (US\$300B) transferred per day

06:30 - 18:30 business days

Proprietary XML protocol, IBM MQ Series messaging

# DOMAIN MODEL SUMMARY



INFRASTRUCTURE

# We're hiring

<https://nubank.workable.com>

São Paulo, Brazil

# We're hiring


<https://nubank.workable.com>

Berlin, Germany



ny bank

**THANK YOU!**





**BACKUP**

# KEY SECURITY DECISIONS

## EXTERNAL


- **Client authentication** (mutual TLS)
  - authorizing new device with reputation score
- **Immutable infrastructure**
  - Short-lived instances
  - No mutations
  - Bootstrap service identity from instance profiles using IAM
- **Uniformity** of service architecture enables rapid patching

## INTERNAL

- **Auto-revoke** of access scopes
  - Operational scopes are short lived
  - Customer contact enables access
  - Employee access bootstrapped from Google OAuth, 2FA + Yubikeys required
- **Realtime monitoring** of security events
  - Cloudtrail, Slack, Lambdas for fine-grained operational access control
- Internal **red team** / incident response team

# GROWING ORGANICALLY THROUGH REFERRALS

Each customer we book leads to 3-4 new leads



Sample Cohorts - Days after release

— 2016-02    — 2016-06    — 2016-09    — 2016-12    — 2017-03    — 2017-06    — 2017-09