

BIG DATA TECHNOLOGY WARSAW SUMMIT 2017
FEBRUARY 9, 2017

REAL-TIME DATA PROCESSING AT RTB HOUSE

ARCHITECTURE & LESSONS LEARNED

BARTOSZ ŁOŚ

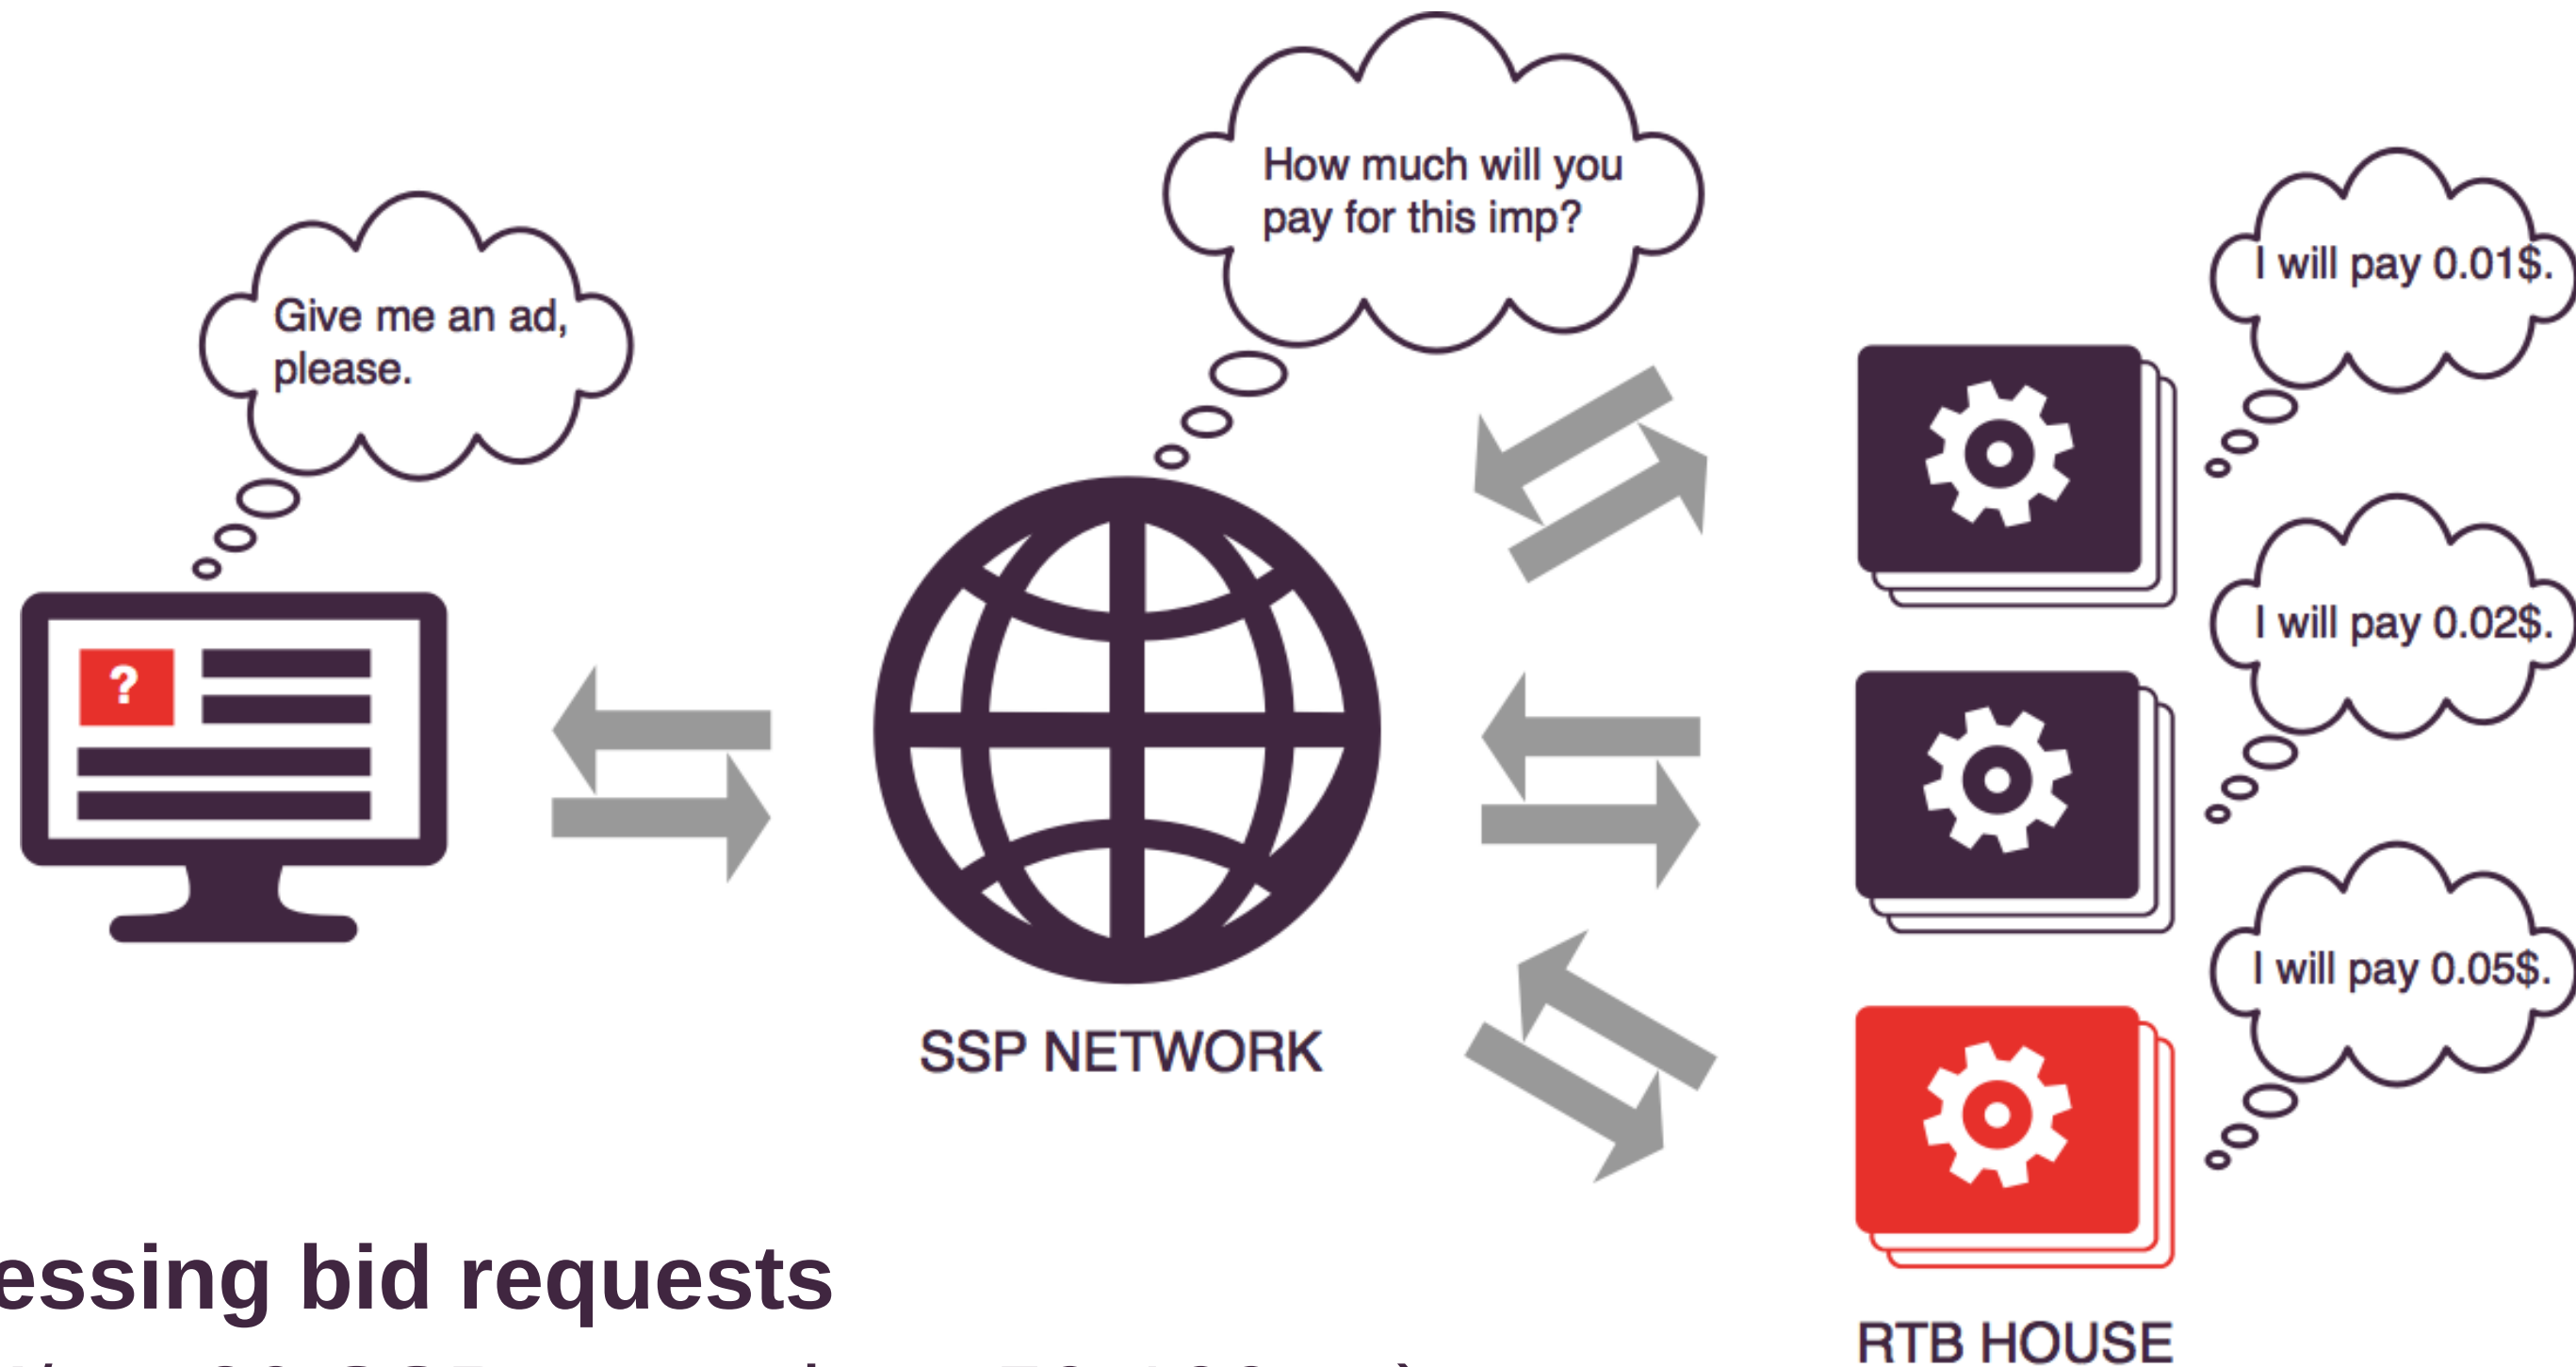
RTB HOUSE =



Agenda:

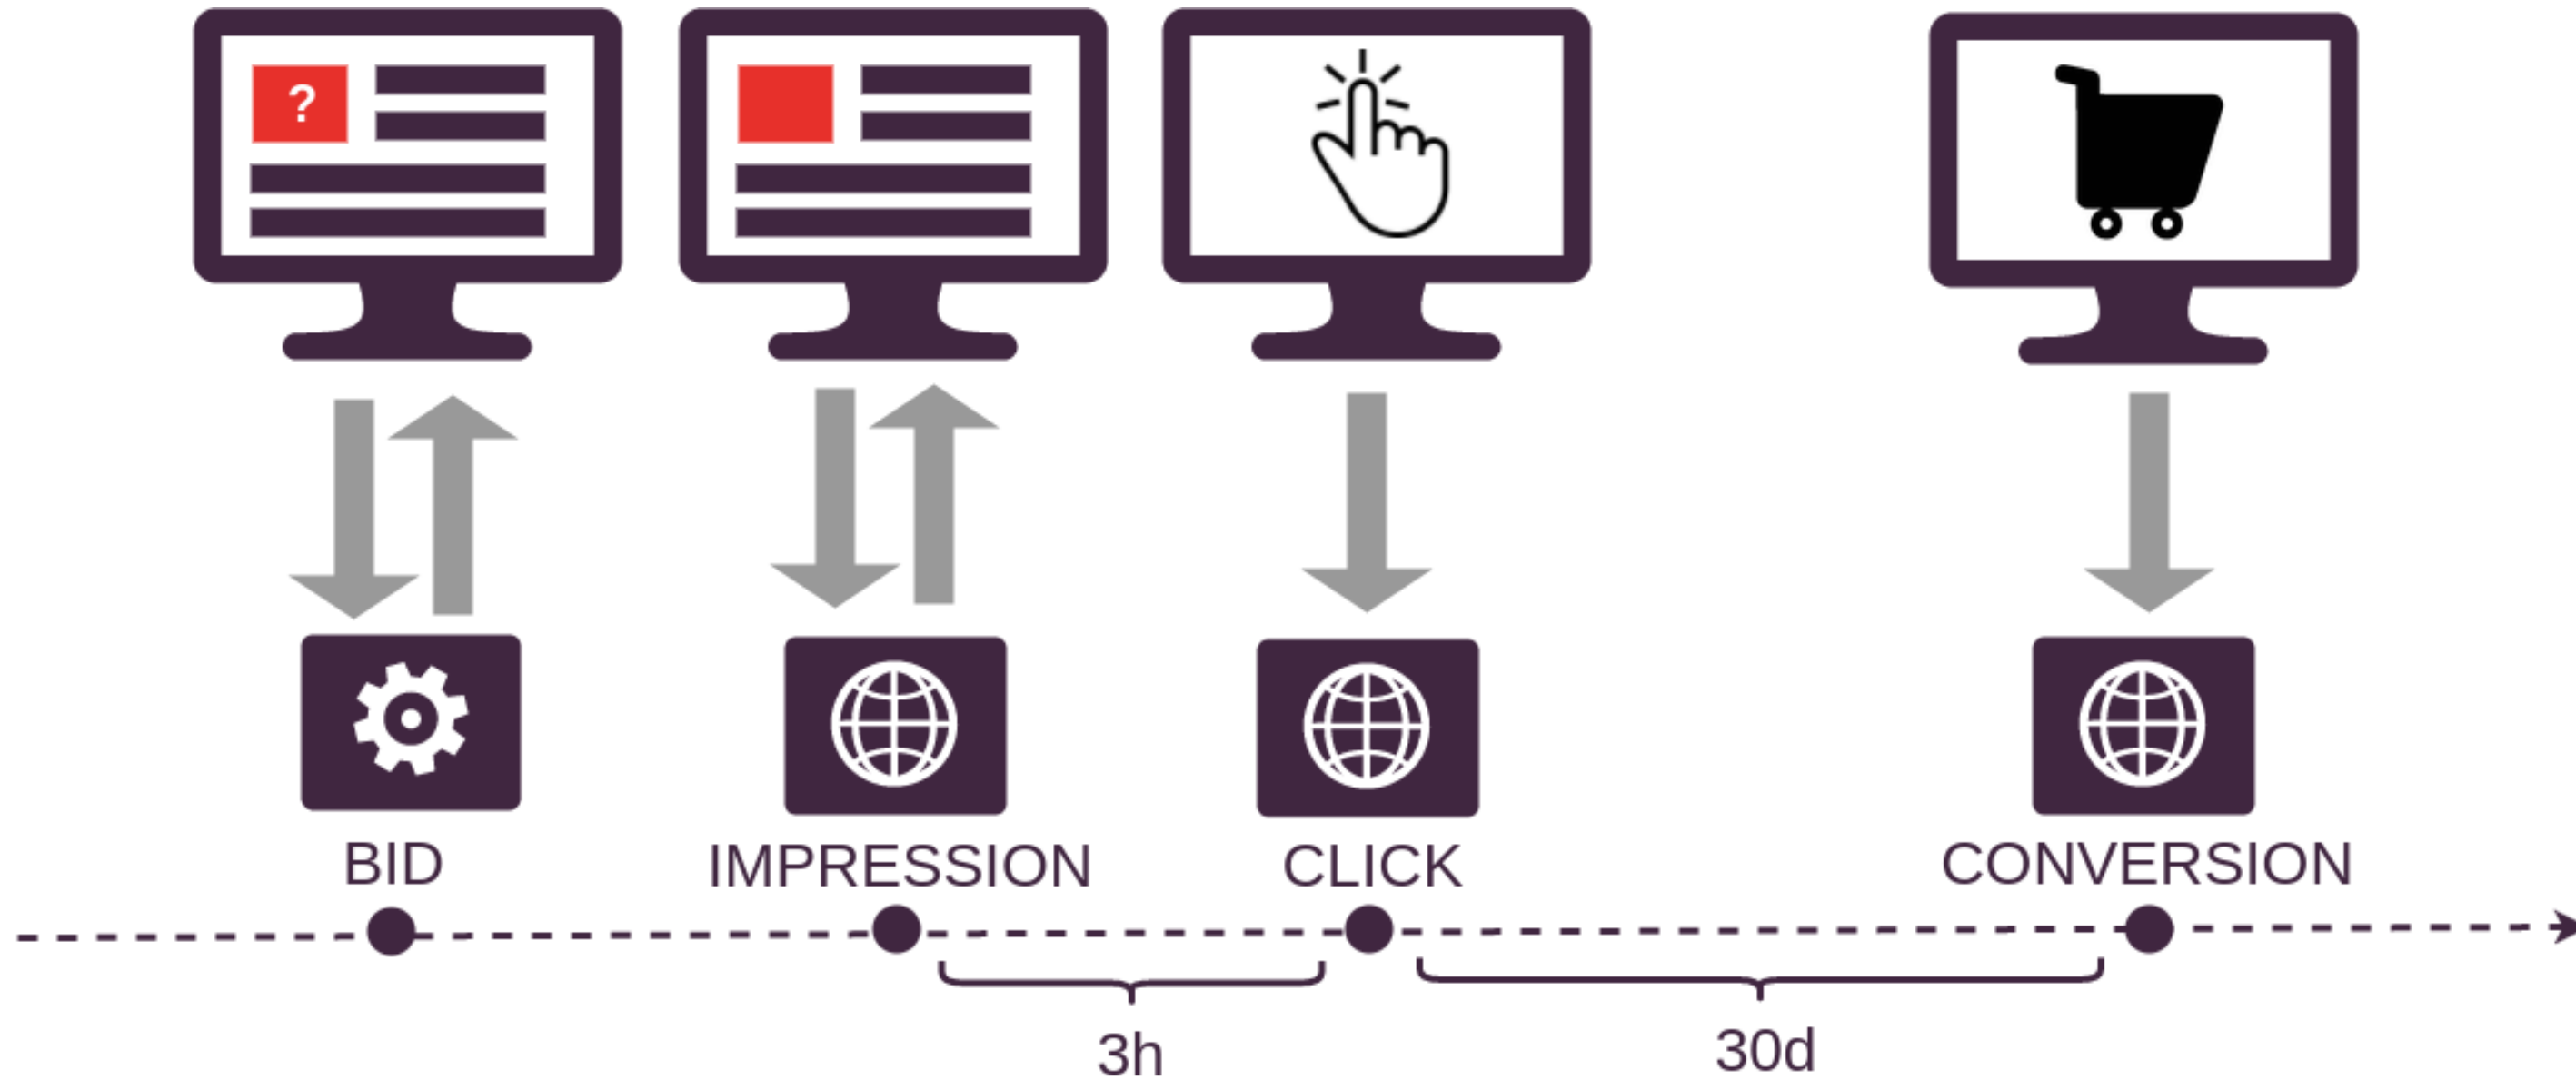
- real-time bidding
- the first iteration: mutable structures
- the second iteration: data-flow
- the third iteration: immutable streams of events

REAL-TIME BIDDING



Processing bid requests

(350K/s, ~30 SSP networks, <50-100ms)



Impressions:

~ 150M events / day
~ 4TB data / day

Clicks:

~ 1M events / day
~ 35GB data / day

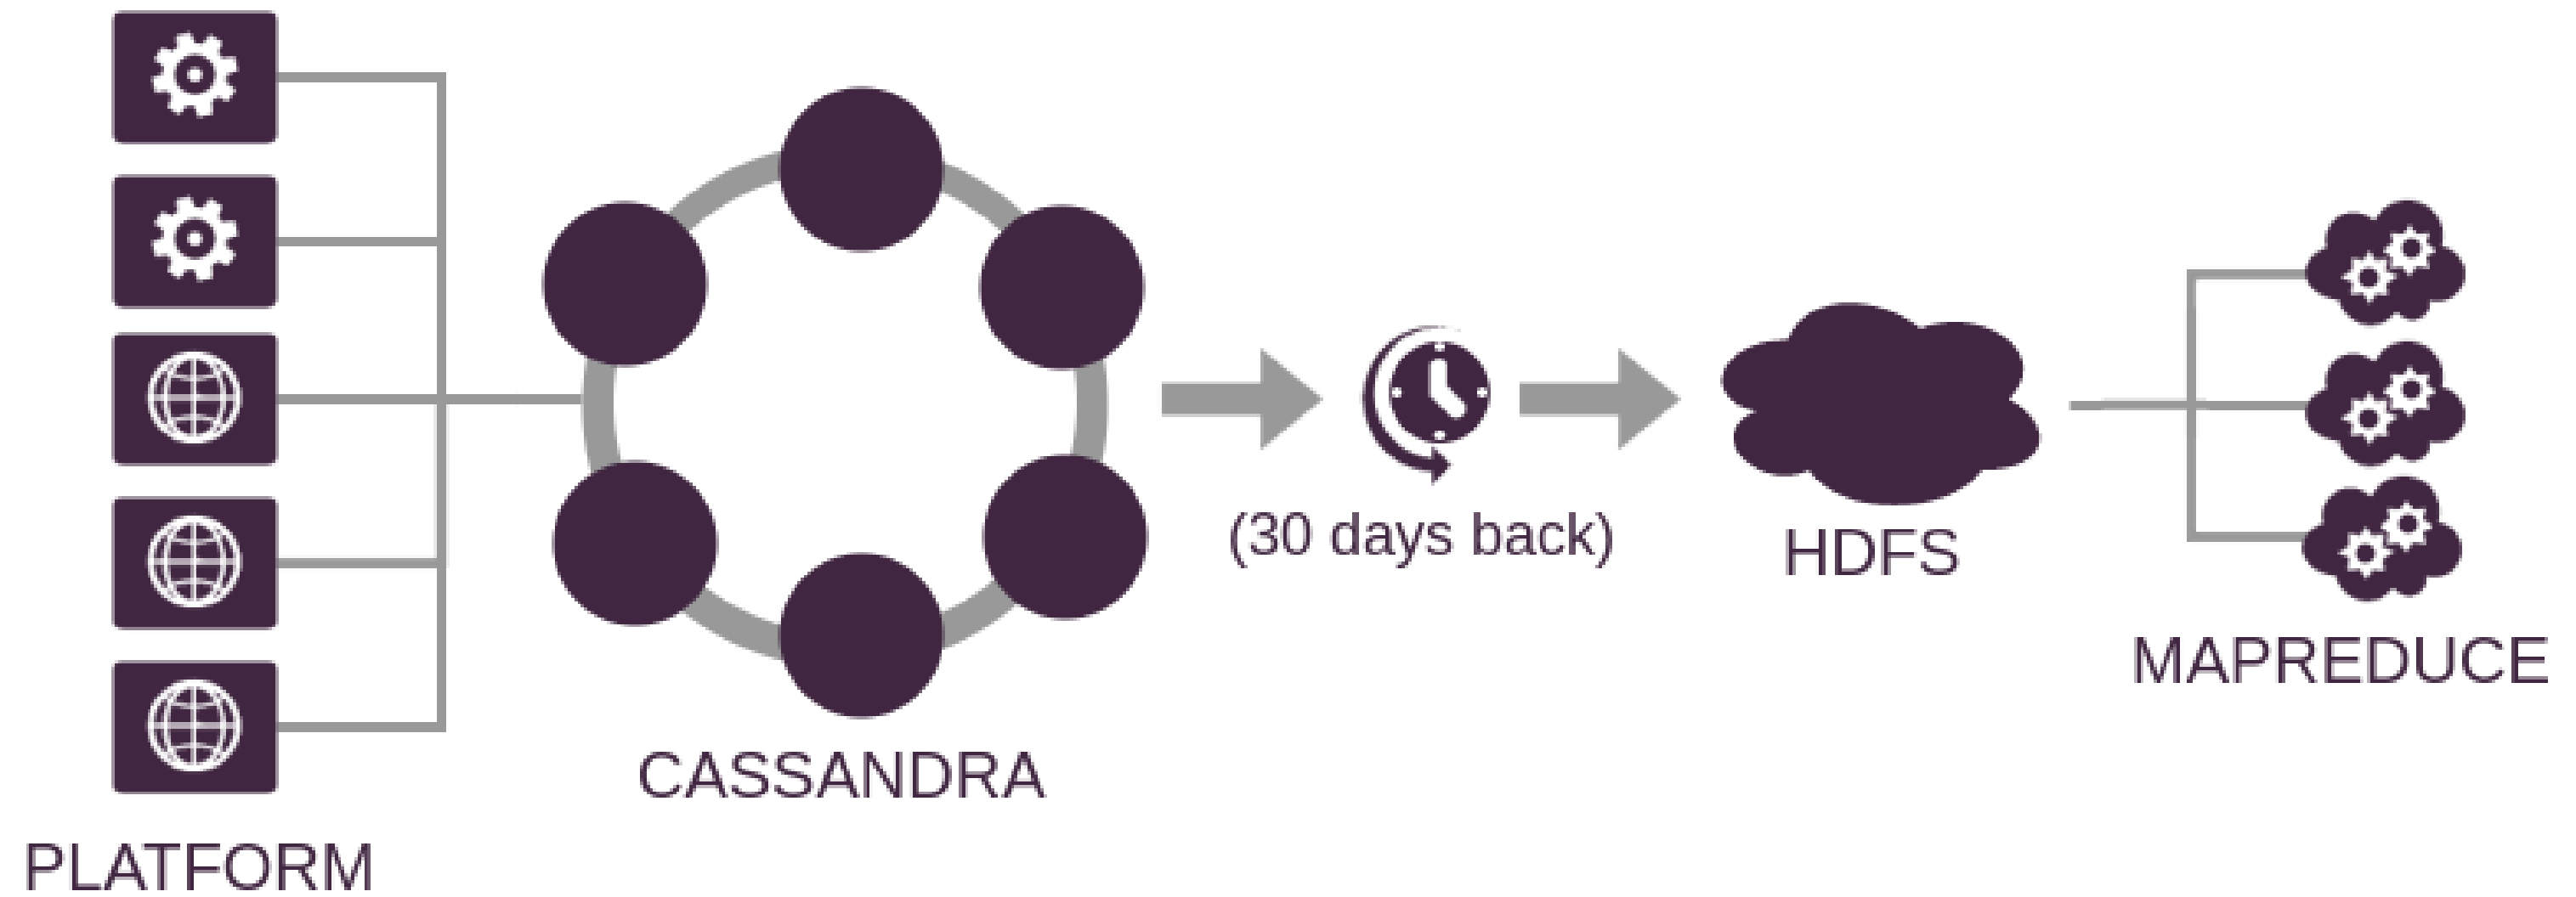
Conversions:

~ 450K events / day
~ 25GB data / day

THE FIRST ITERATION

THE 1ST ITERATION: MUTABLE IMPRESSIONS

07/23



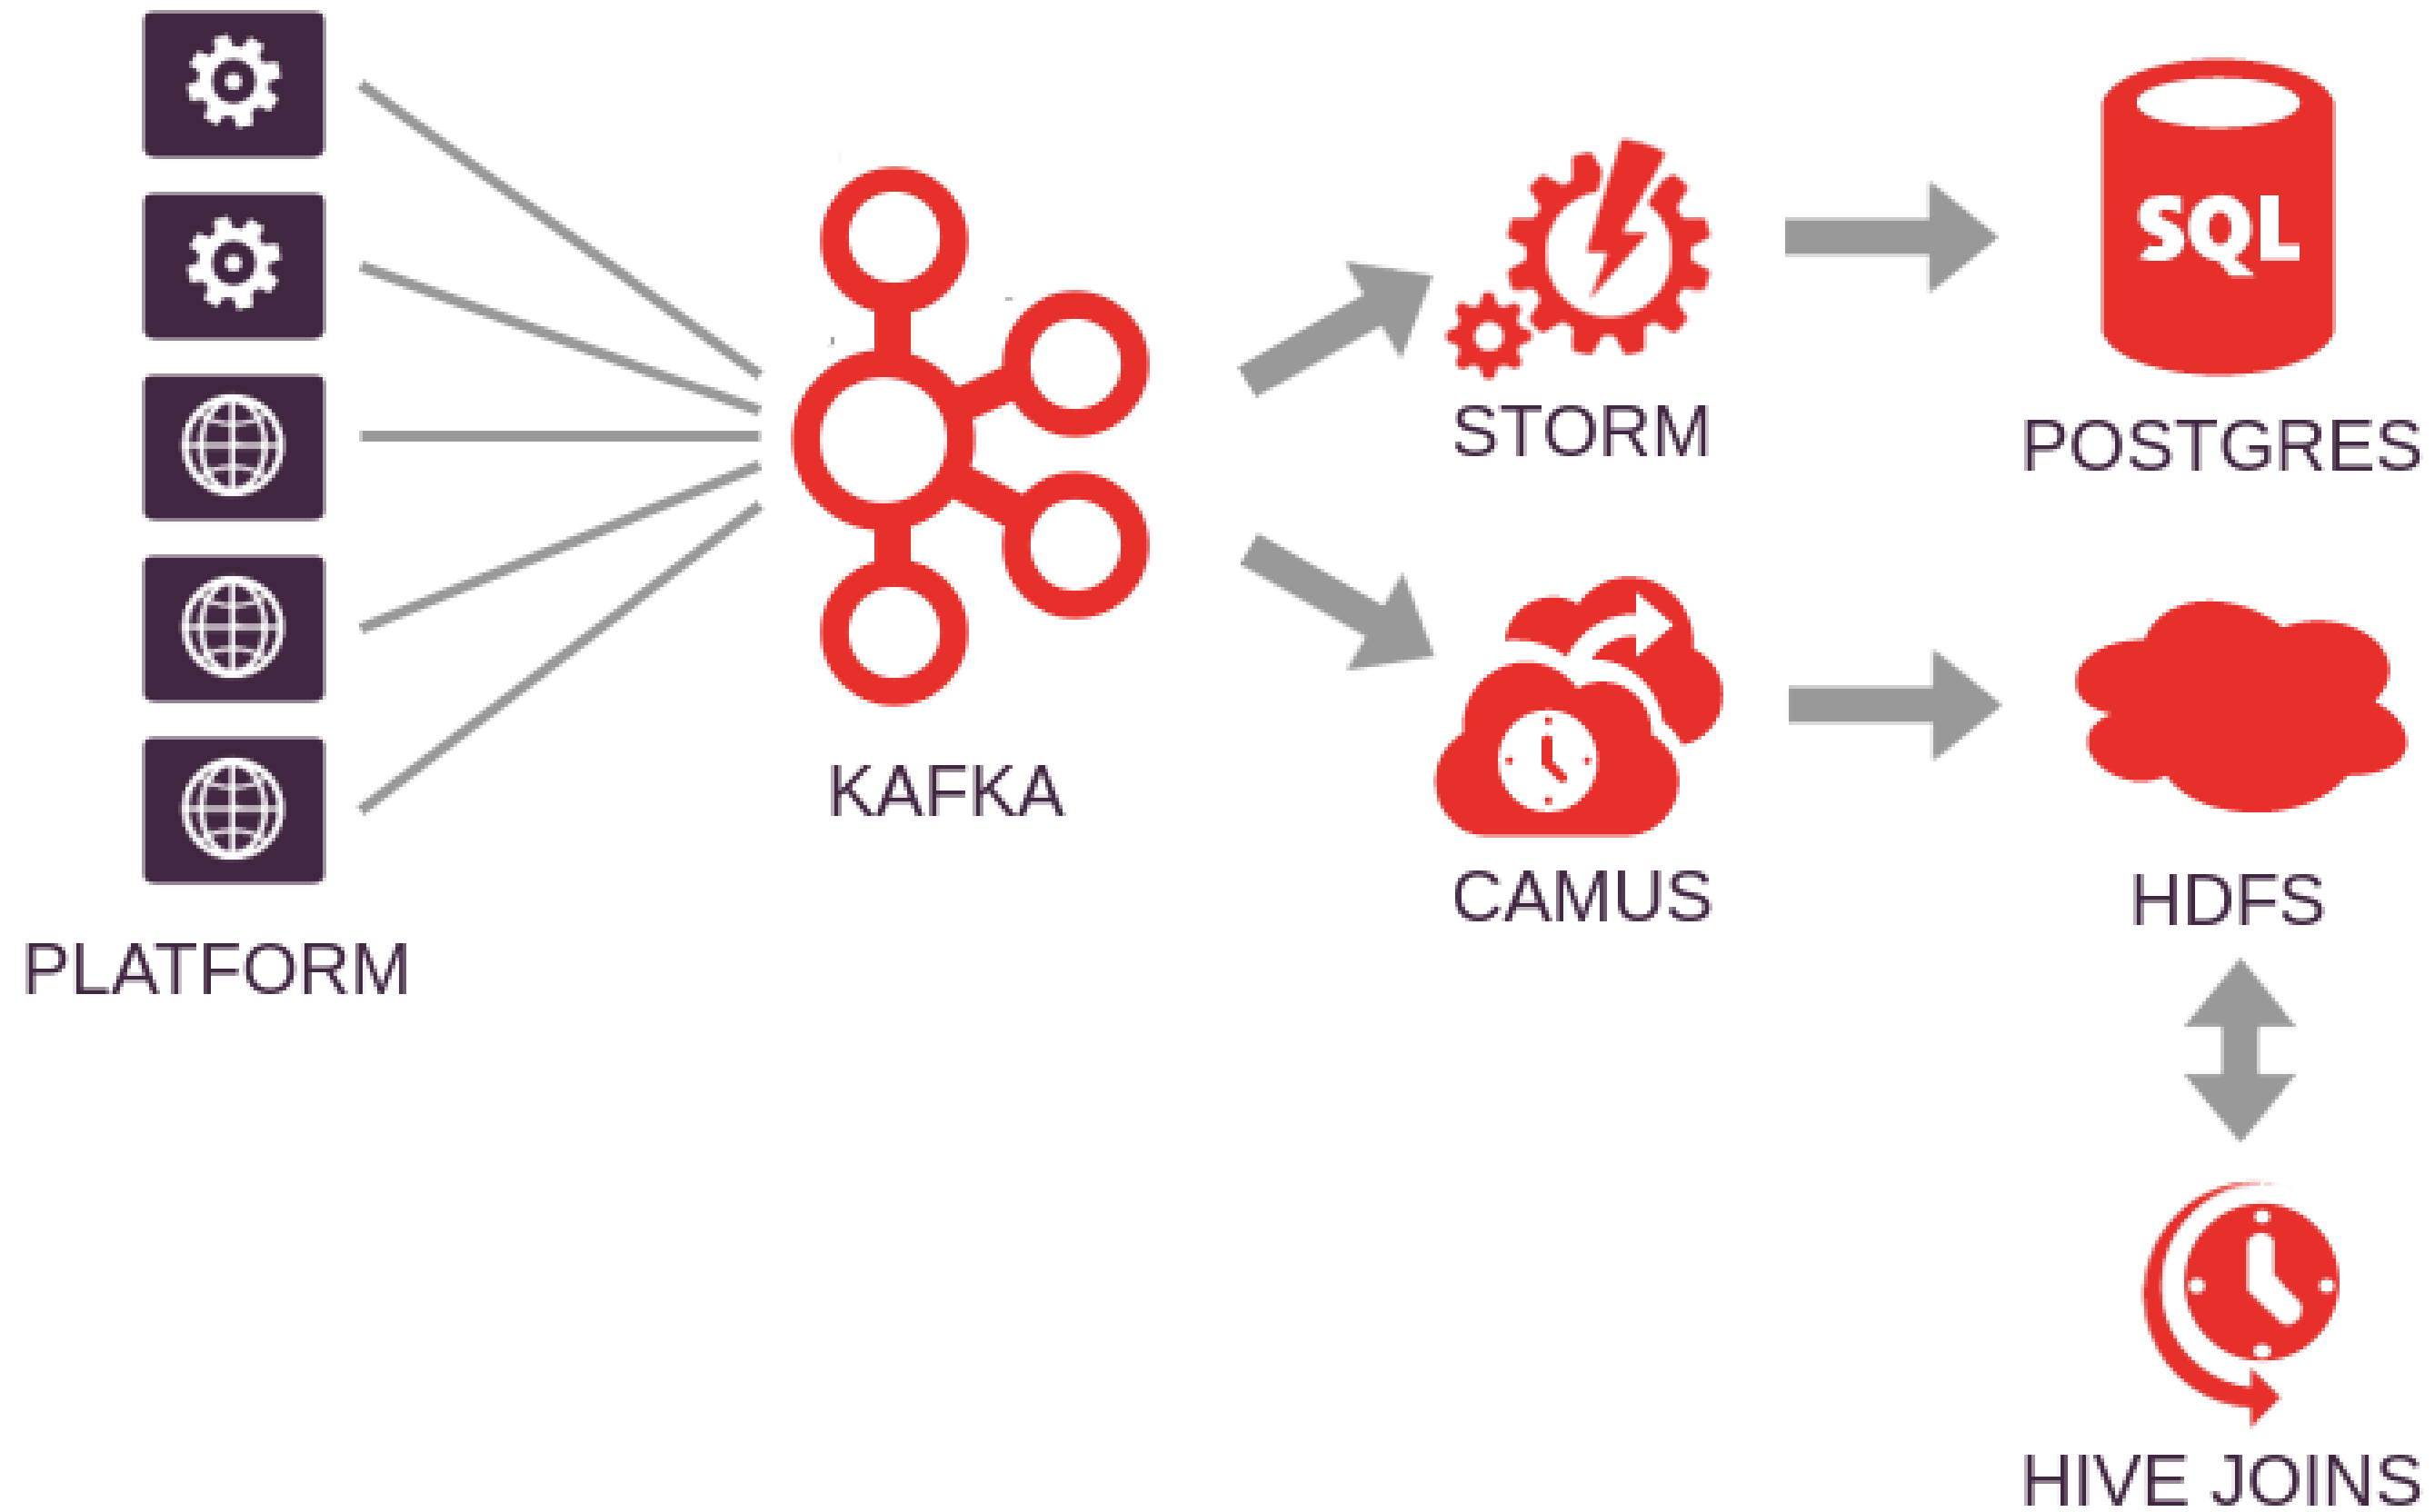
Issues:

- long, overloading data migrations (30 days back)
- complex servlets' logic, inability to reprocess
- inflexible, various schemas
- single-DC
- inconsistencies

THE SECOND ITERATION: DATA-FLOW

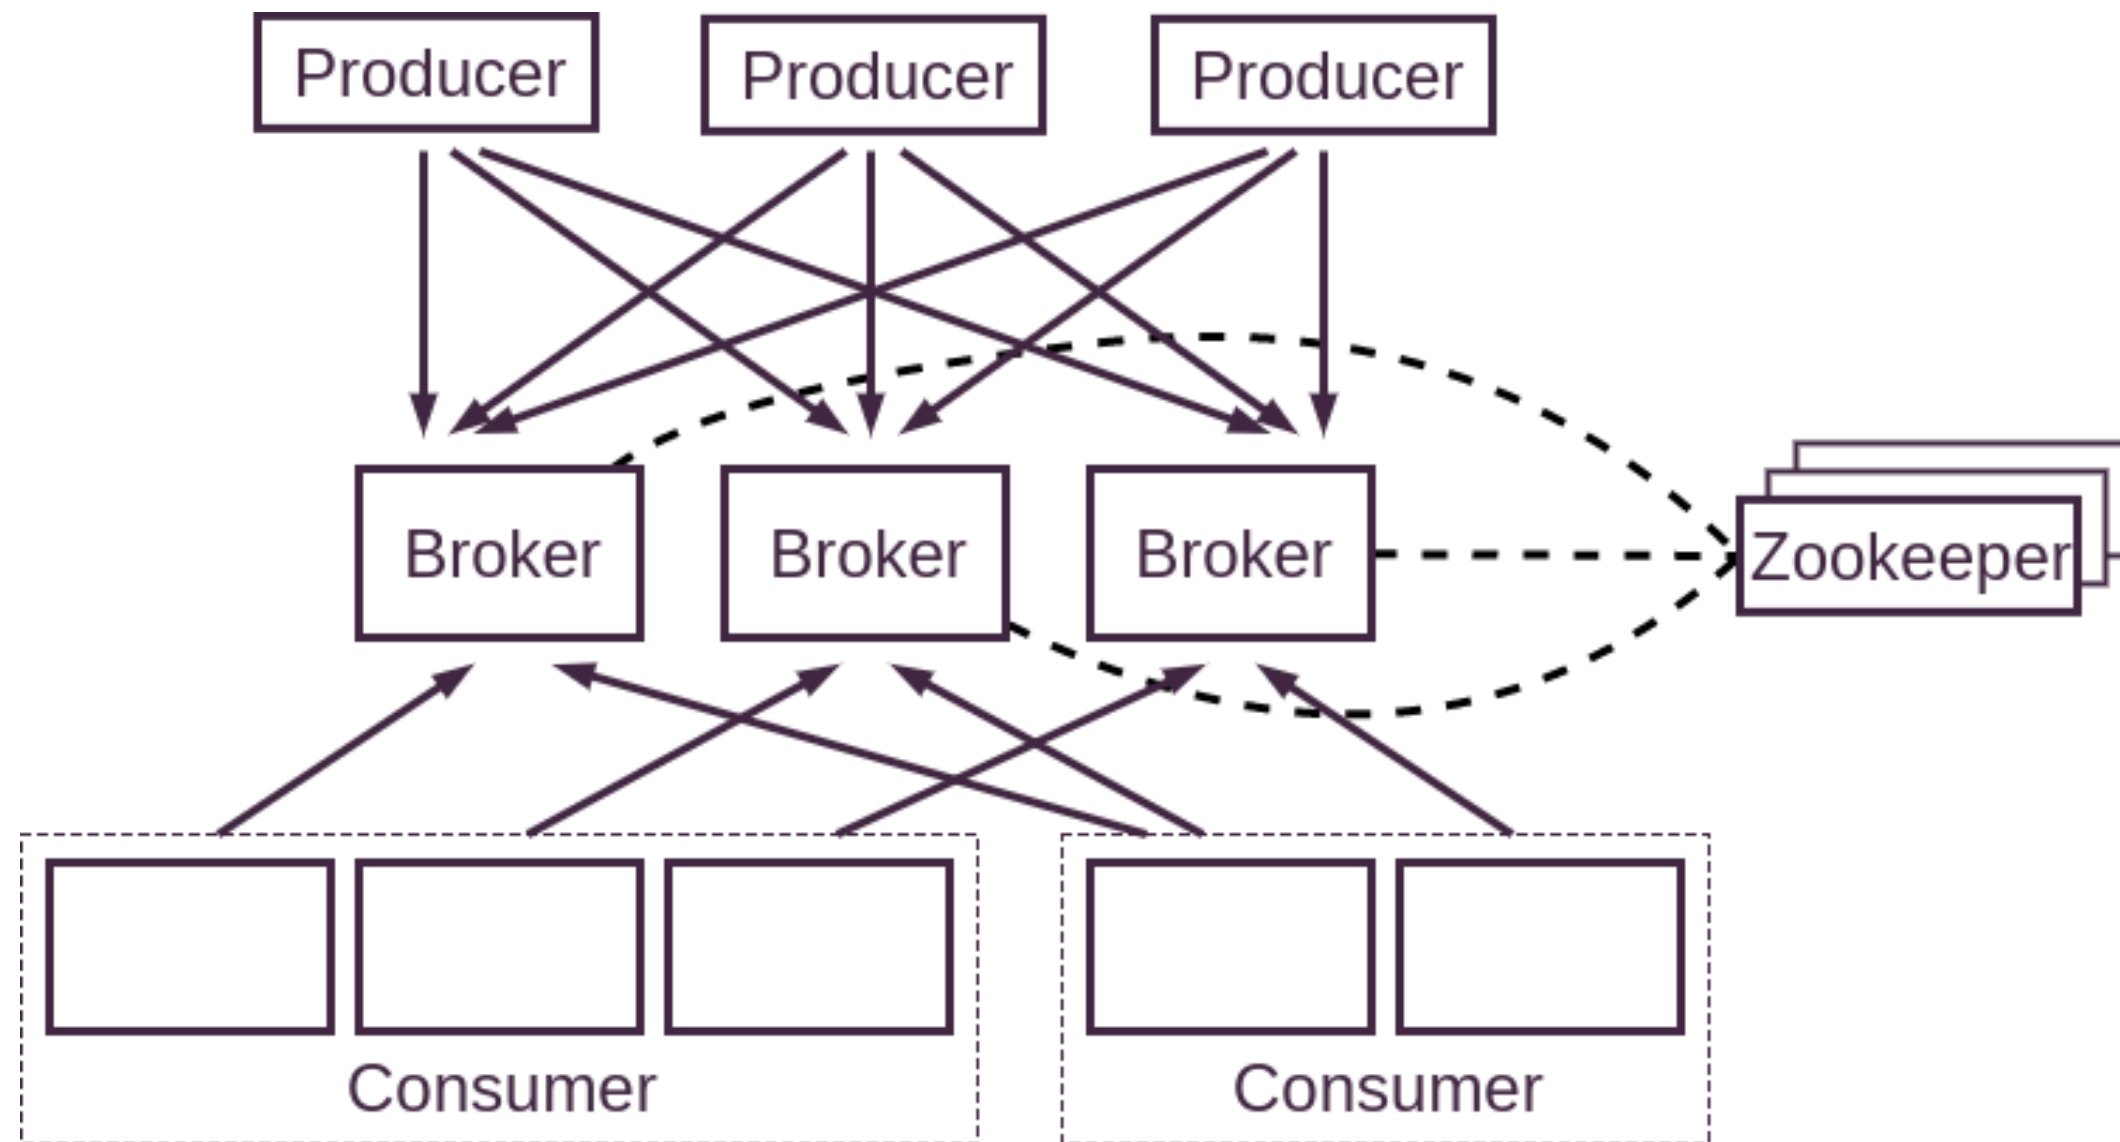
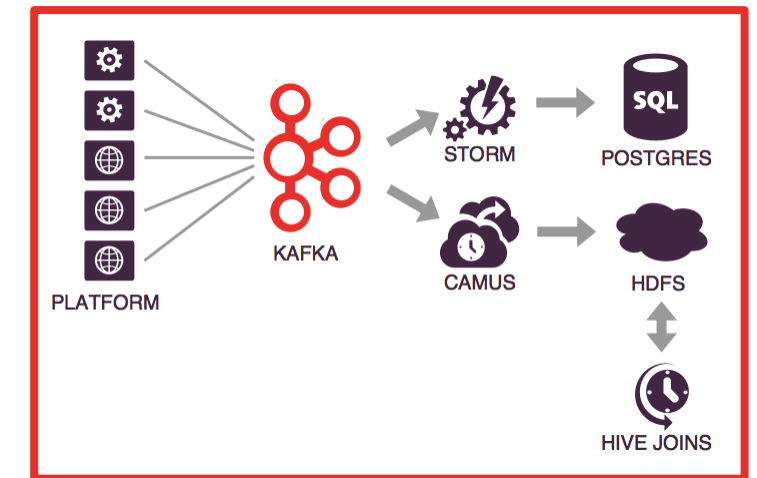
THE 2ND ITERATION: THE 1ST DATA-FLOW ARCHITECTURE

10/23



THE 2ND ITERATION: DISTRIBUTED LOG

11/23

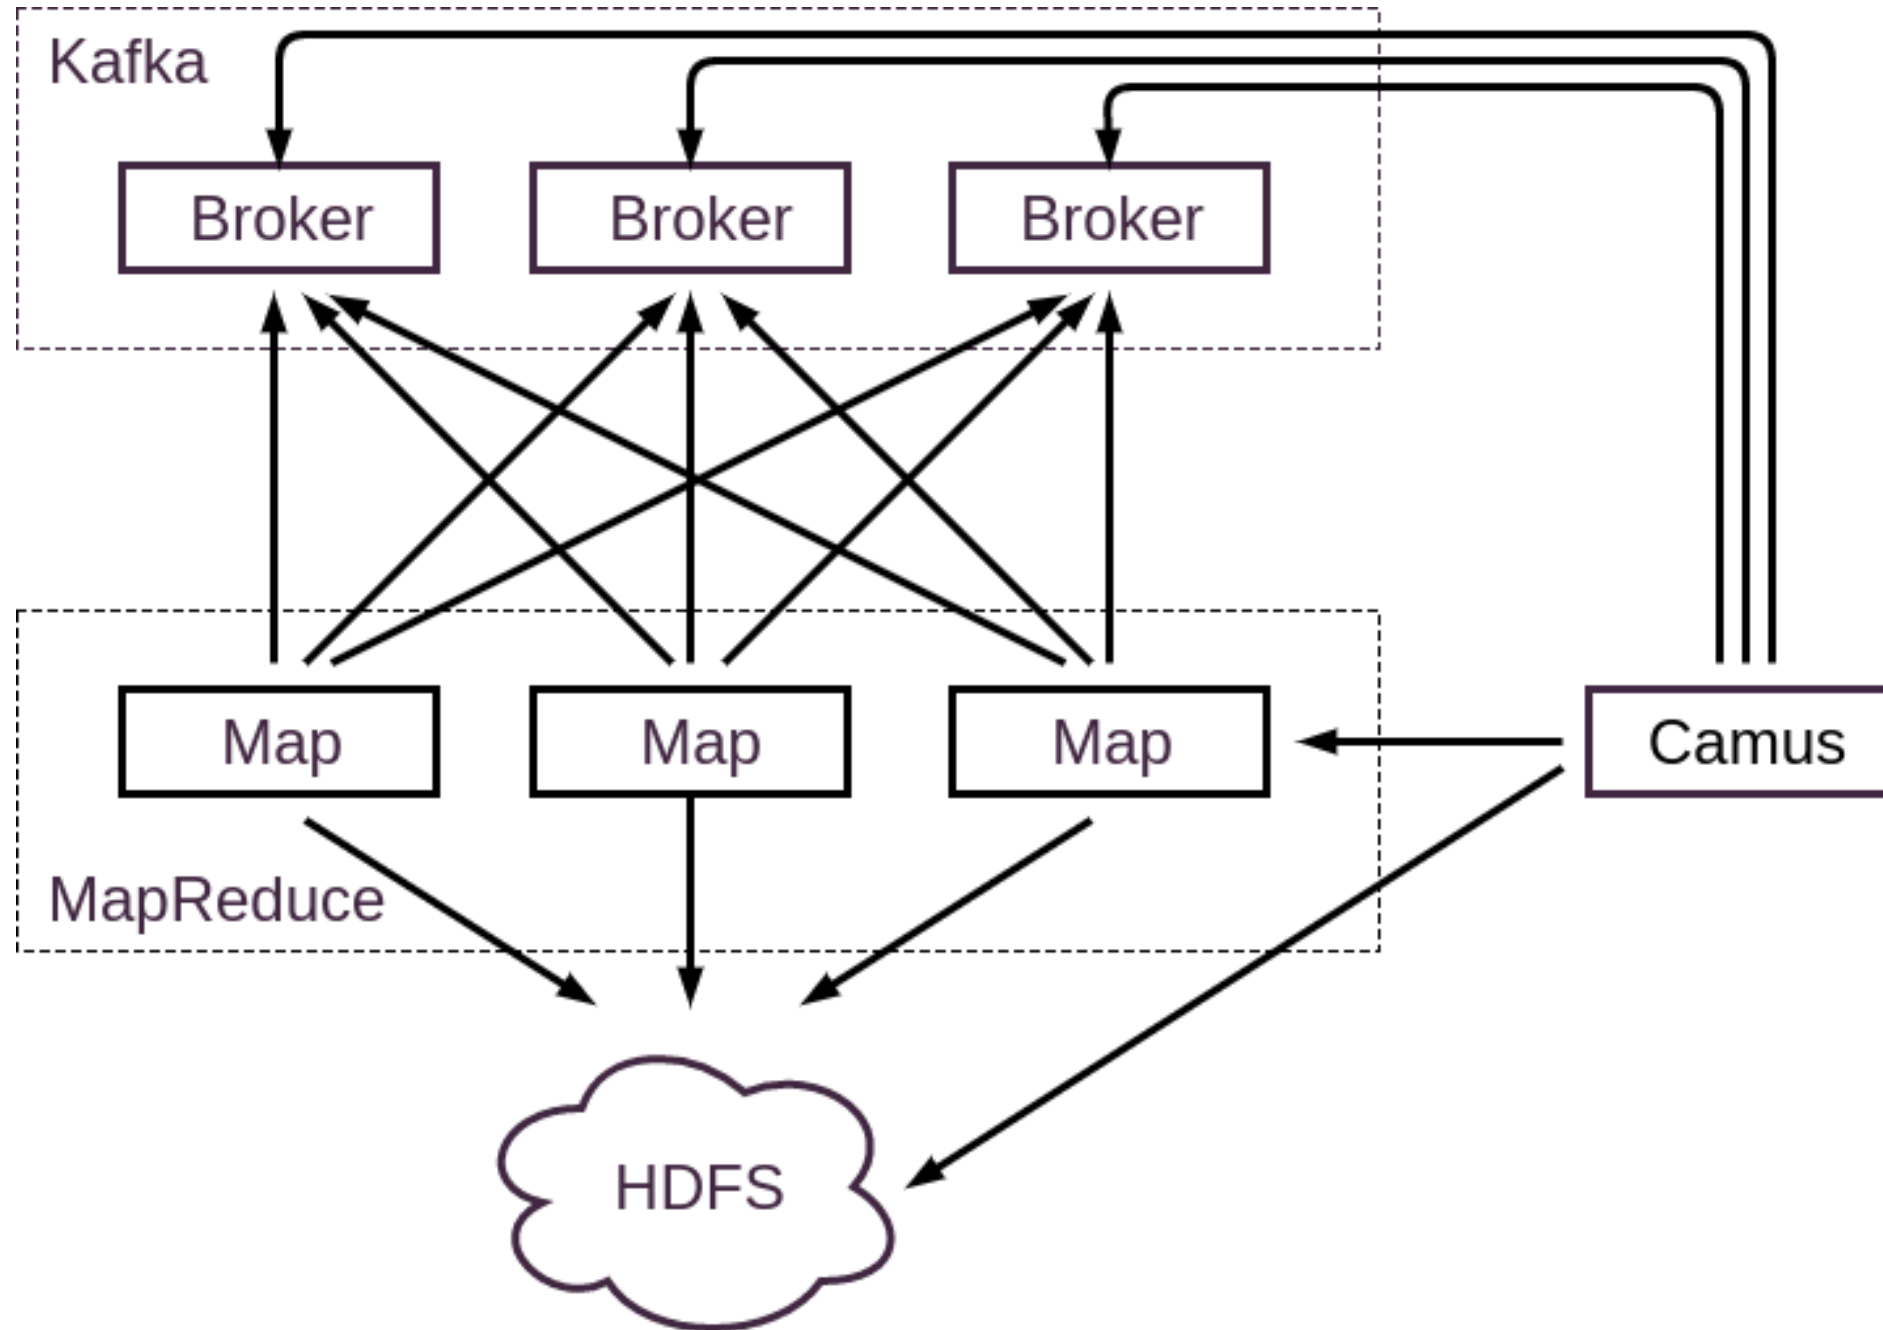
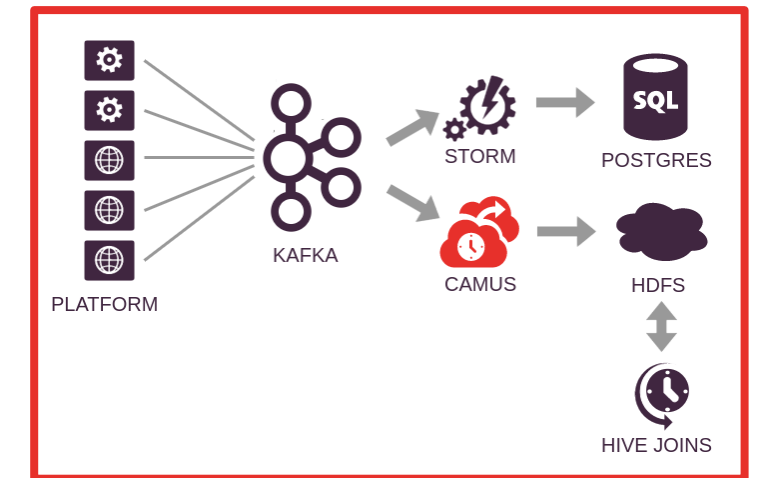


Why Apache Kafka:

- distributed log
- topics partitioning
- partition replication
- log retention
- stateless
- efficient data consuming

THE 2ND ITERATION: BATCH LOADING

12/23



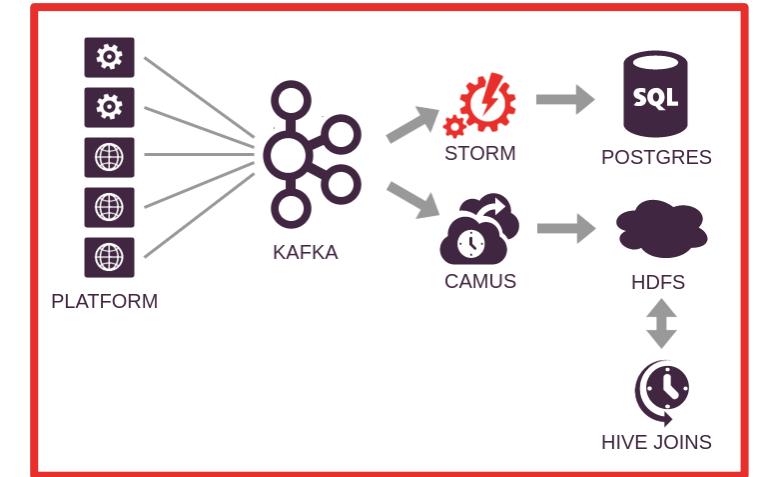
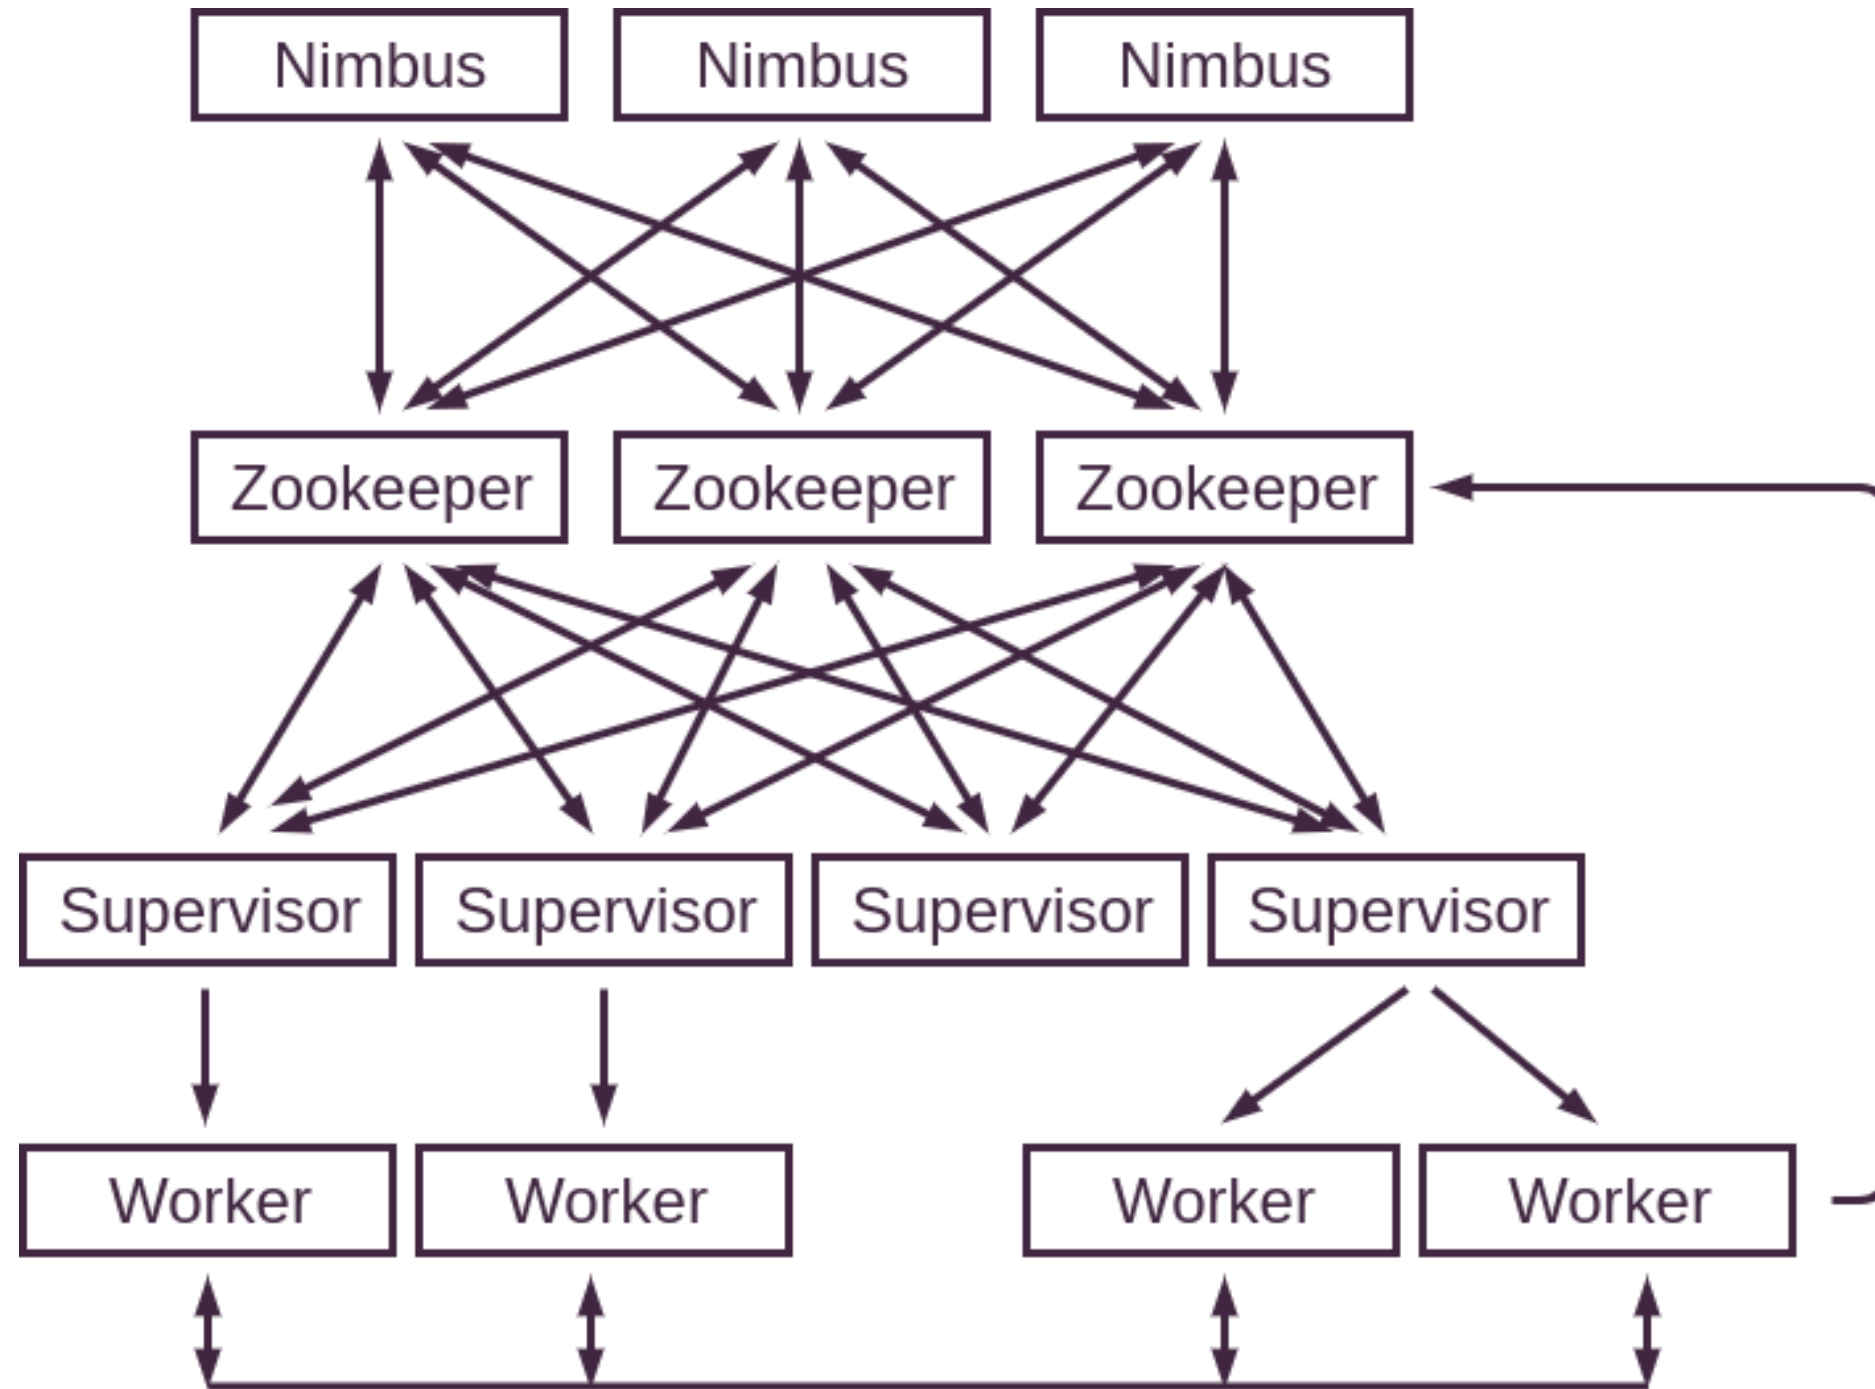
Why Apache Camus:

- "Kafka to HDFS" pipeline
- map-reduce jobs, batches
- storing offsets in log files
- data partitioning

Why Apache Avro:

- data serialization framework
- rich data structures
- self-describing container files
- reader & writer schemas
- binary data format
- schema registry

THE 2ND ITERATION: ACCURATE STATISTICS



Why Apache Storm:

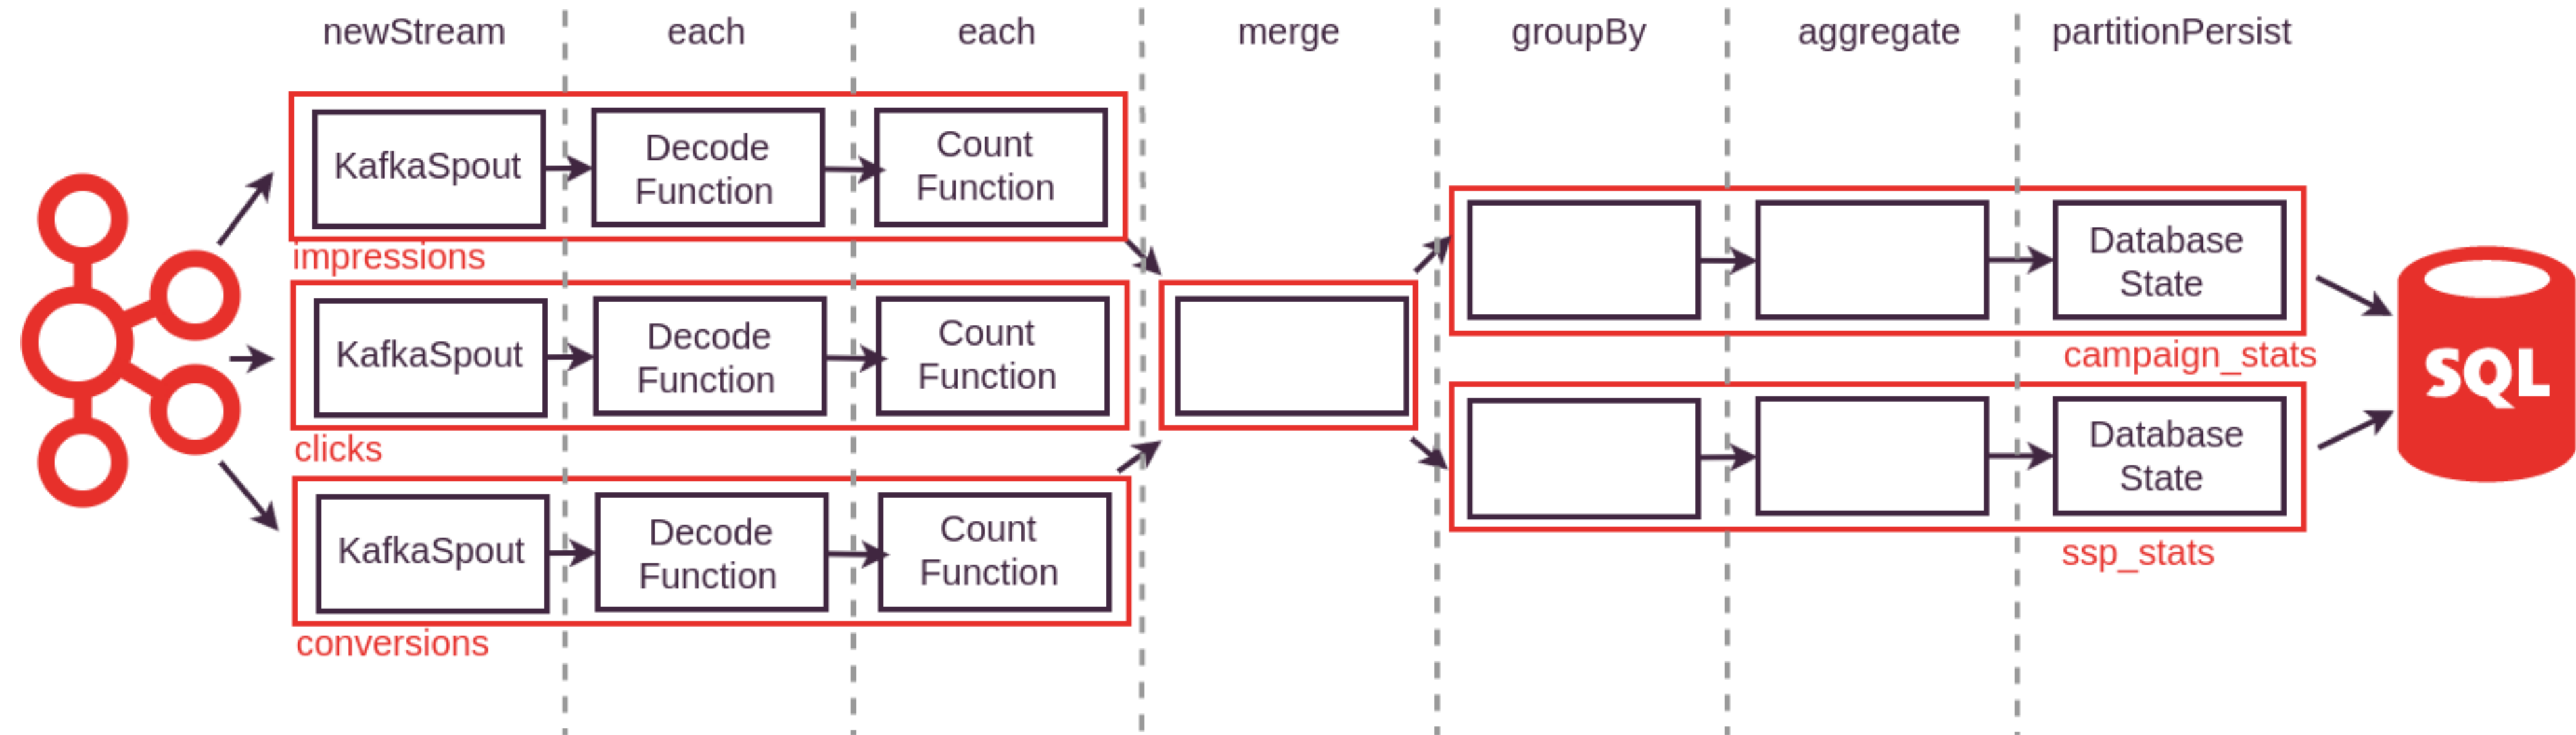
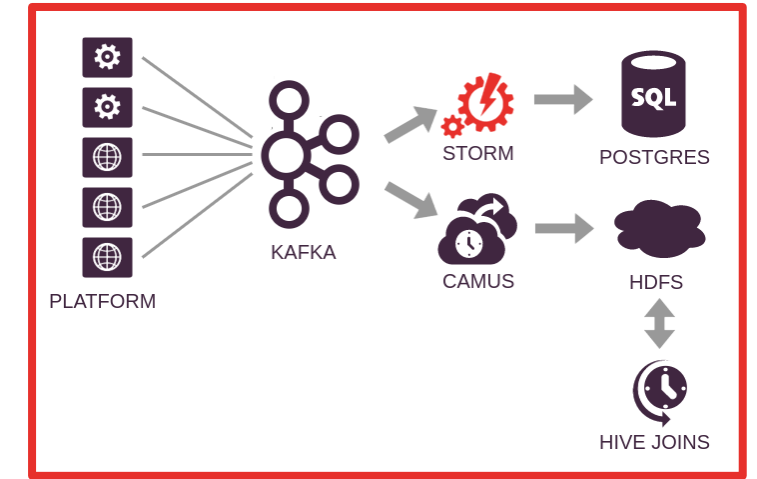
- real-time processing
- streams of tuples, topologies
- fault-tolerance

Why Trident:

- transactions, exactly-once processing
- microbatches (latency & throughput)

THE 2ND ITERATION: STATS-COUNTER TOPOLOGY

15/23



Hybrid architecture:

- aggregates (real-time)
- raw events (2-hour batches)
- joined events (end-of-day batch jobs)

Other issues:

- Hive joins
- mutable events
- servlets' complex logic

THE THIRD ITERATION: NEW APPROACH

New approach:

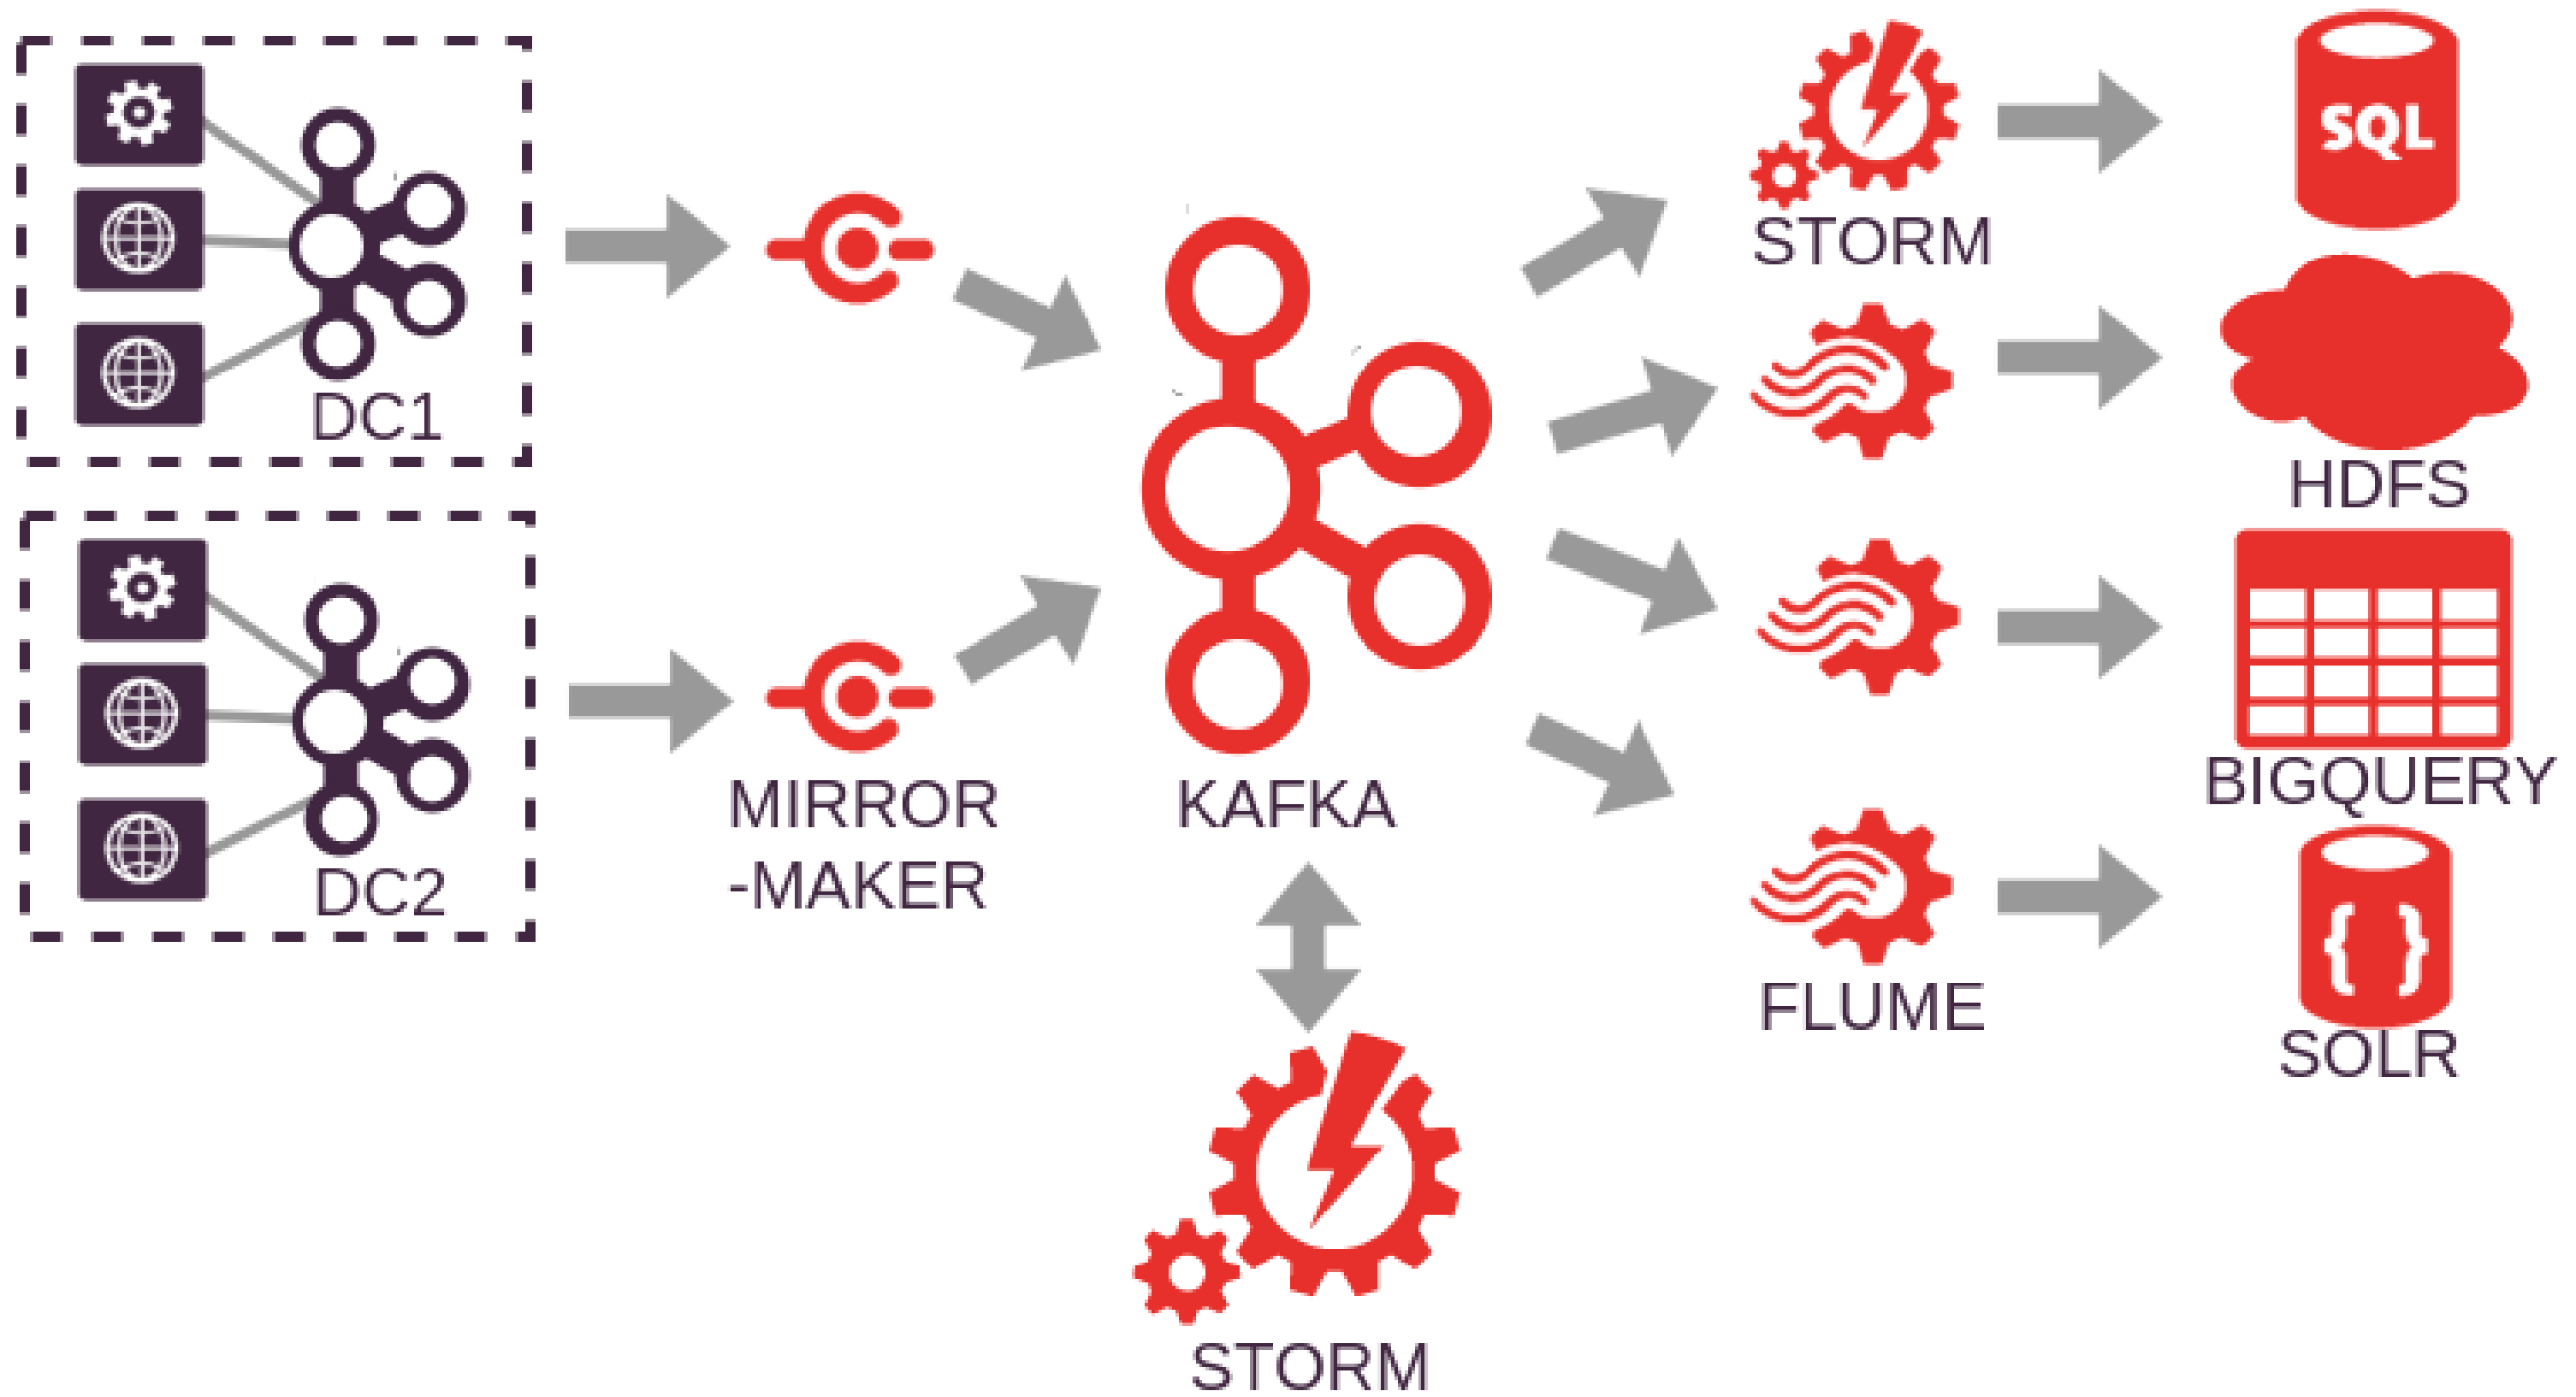
- real-time processing
- publishing light events
- immutable streams of events

```
{ "IMPRESSION":  
  "URL",  
  "TIME",  
  "CREATIVE",  
  ...  
  "CLICKS",  
  "CONVERSIONS"  
}
```

```
{ "CLICK":  
  "TIME",  
  "IMPRESSION_ID",  
  ...  
  "IMPRESSION"  
}
```

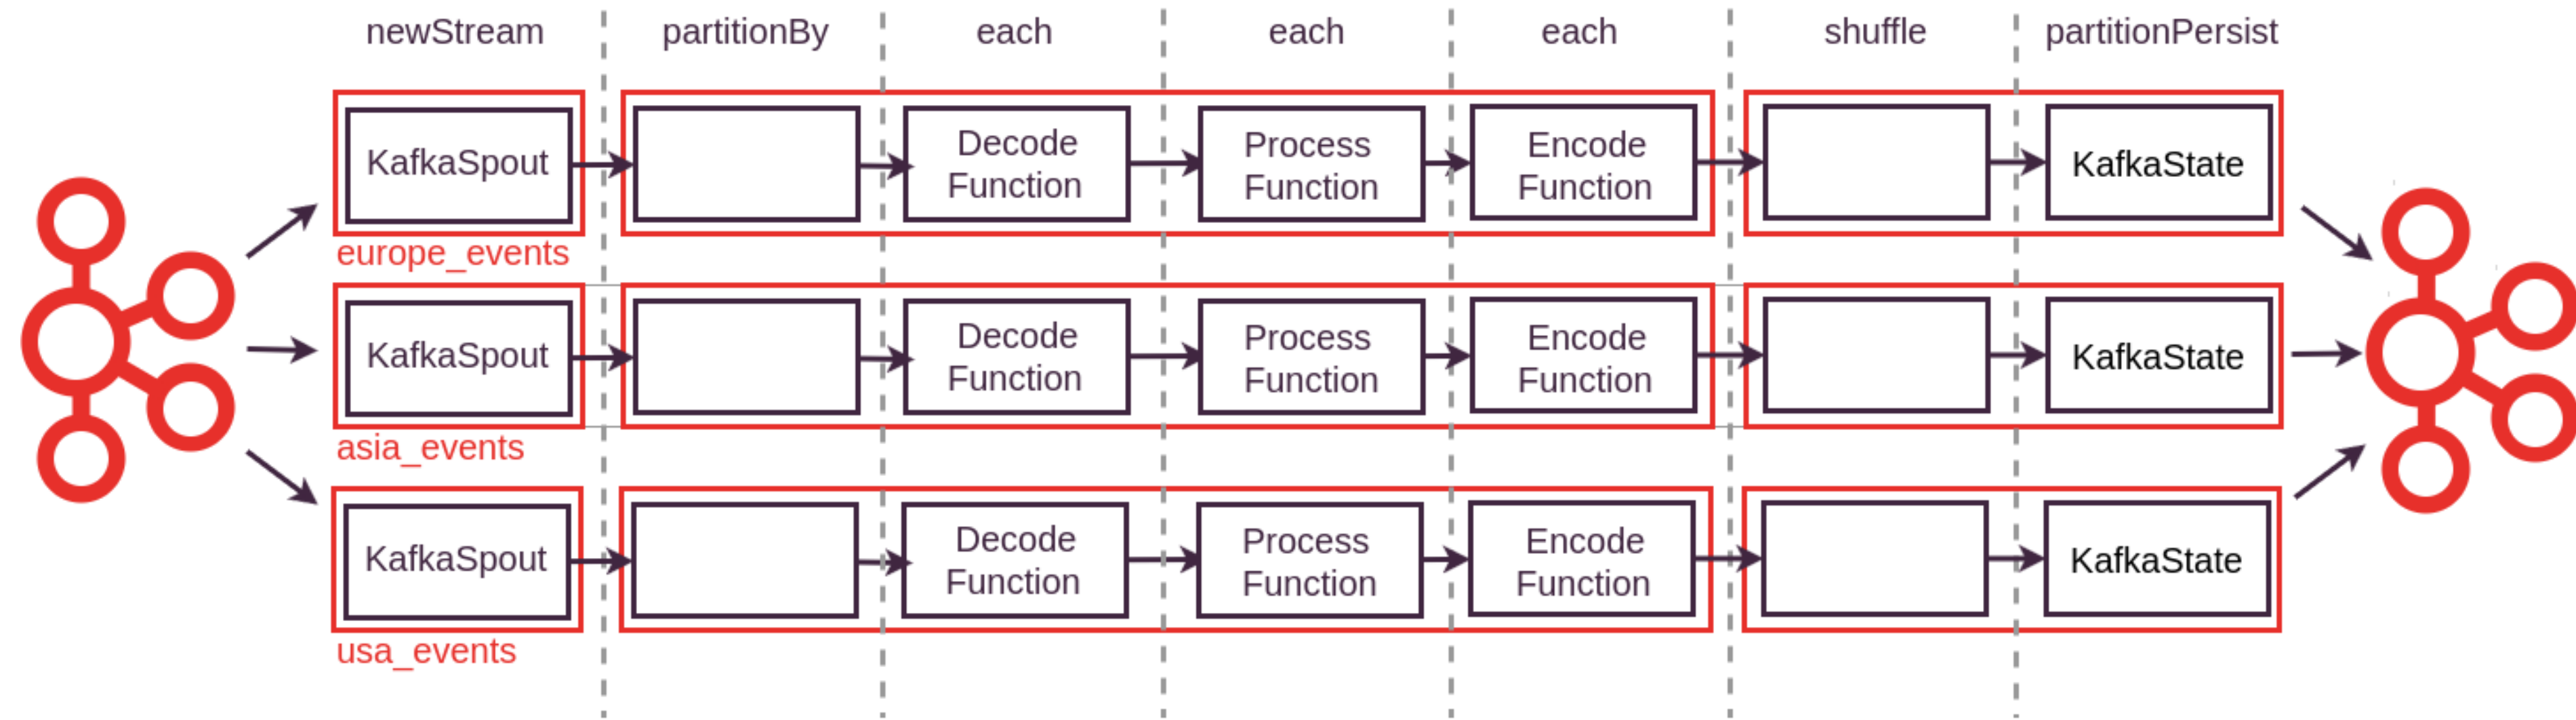

```
{ "CONVERSION":  
  "TIME",  
  "CLICK_ID",  
  ...  
  "CLICK"  
}
```

THE 3RD ITERATION: HIGH-LEVEL ARCHITECTURE

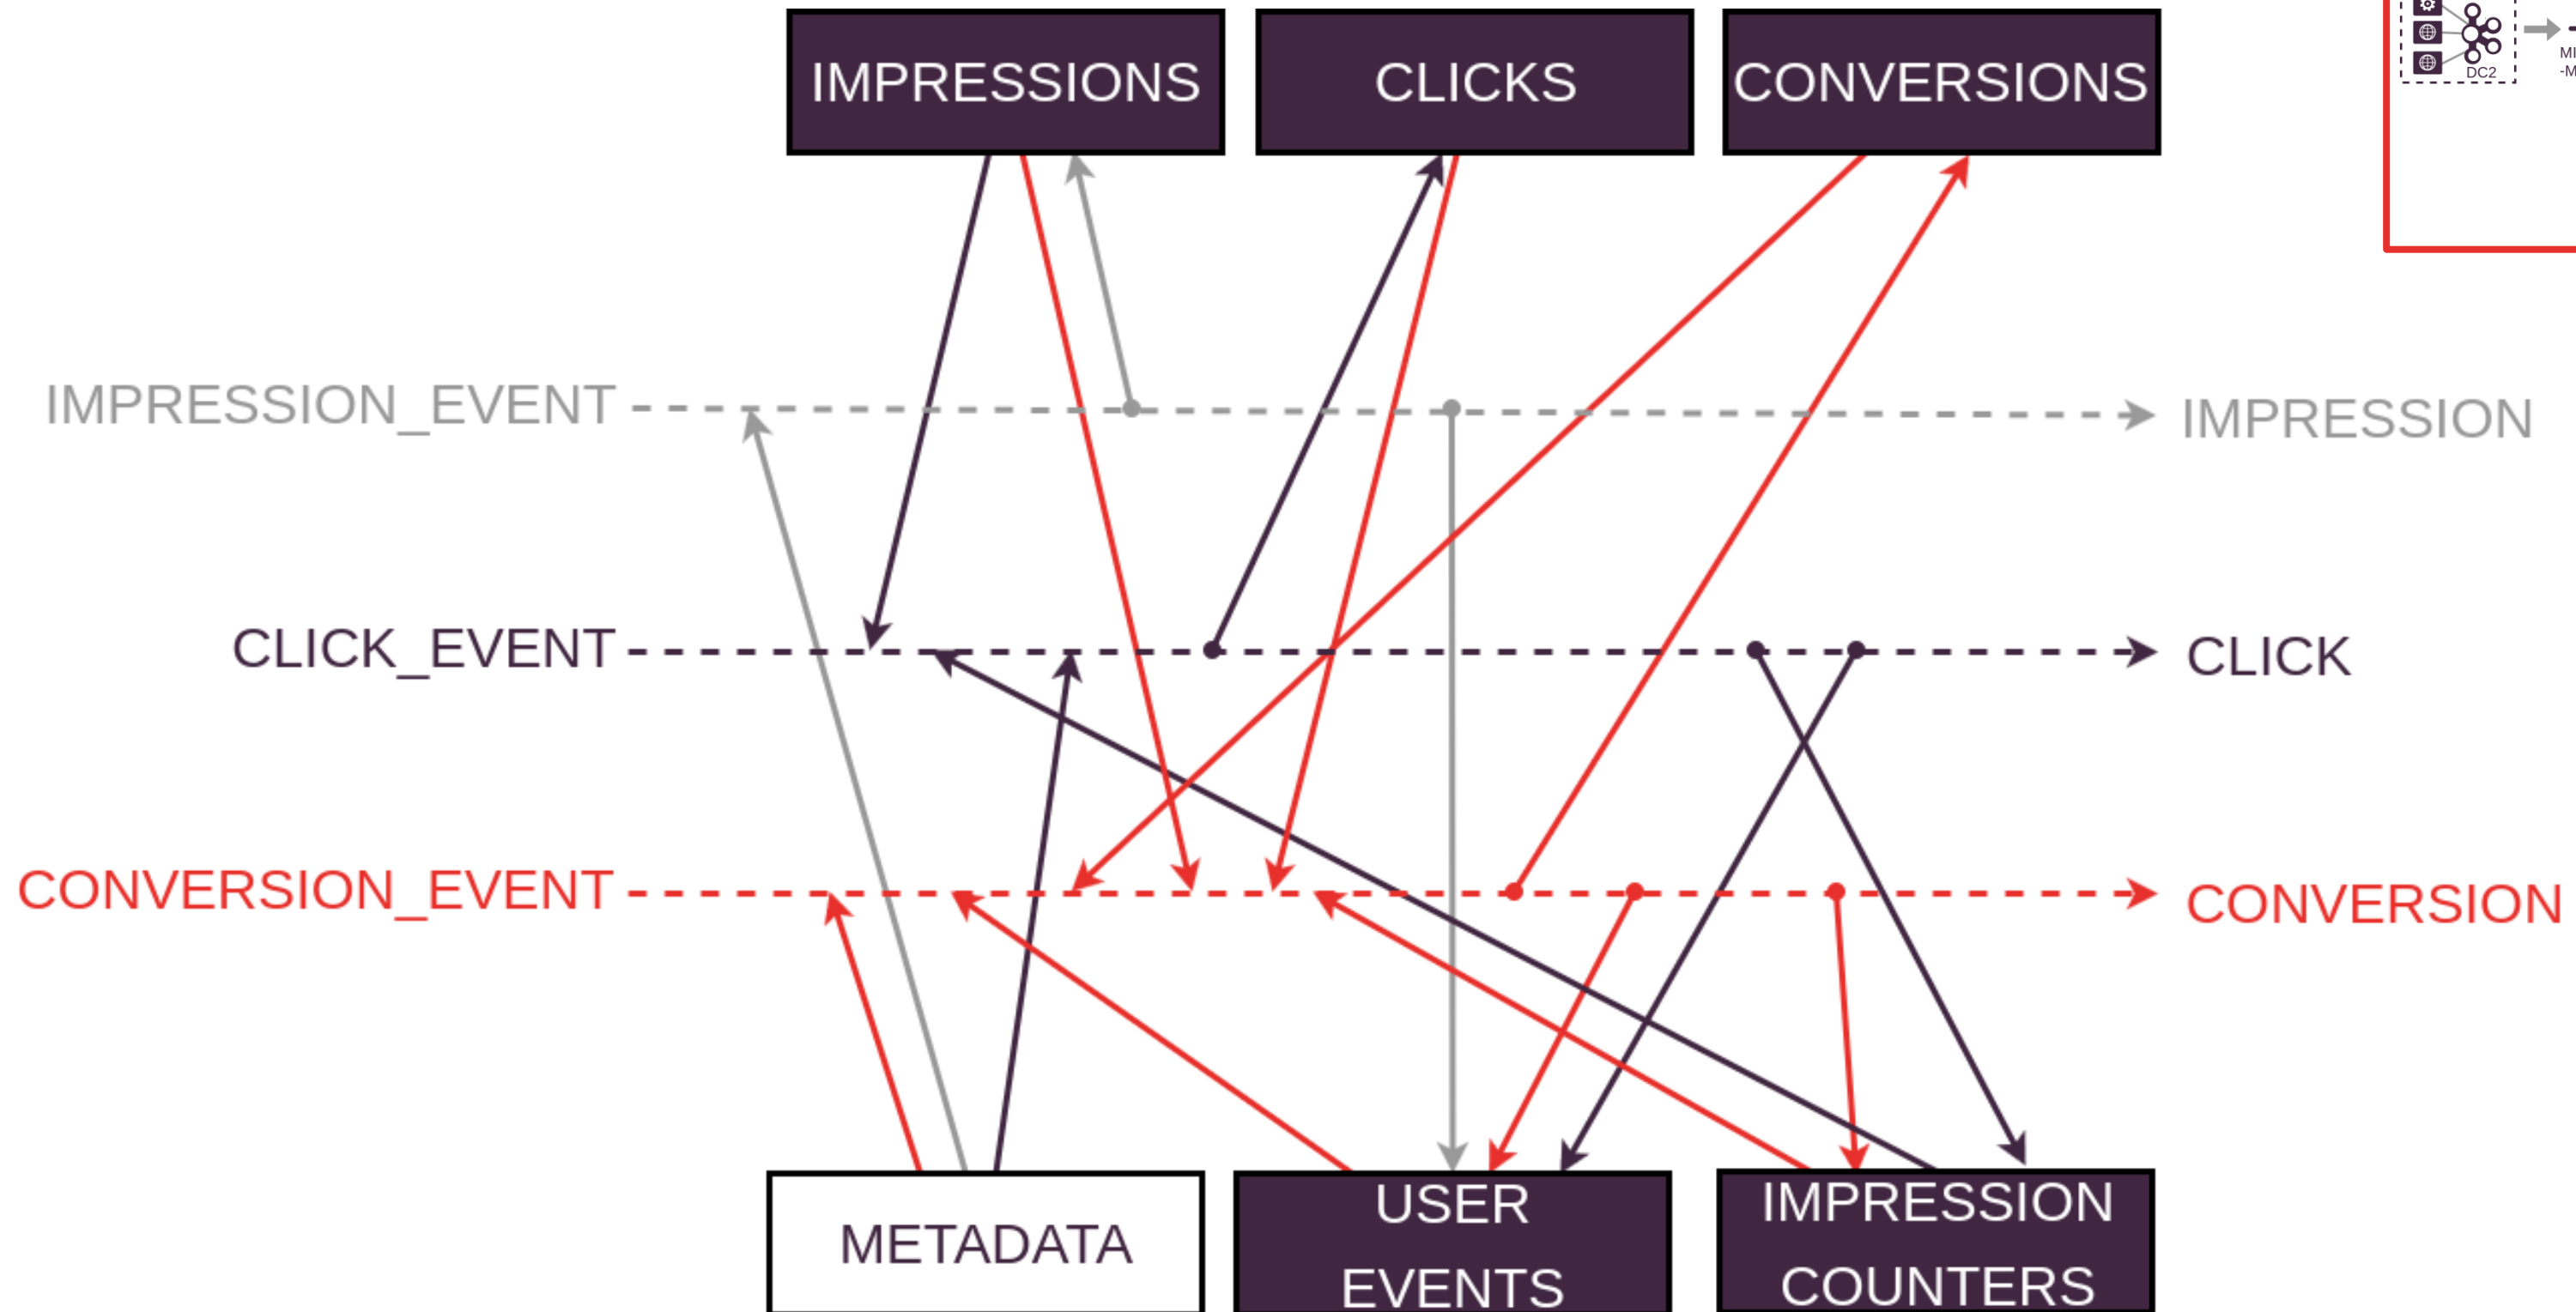



THE 3RD ITERATION: DATA-FLOW TOPOLOGY

20/23



THE 3RD ITERATION: **EVENTS MERGE**



What we have achieved:

- multi-DC architecture
- HDFS & BigQuery streaming
- platform monitoring
- much more stable platform
- higher quality of data processing
- better data-flow monitoring, deployment & maintenance



THANK YOU
FOR YOUR ATTENTION