

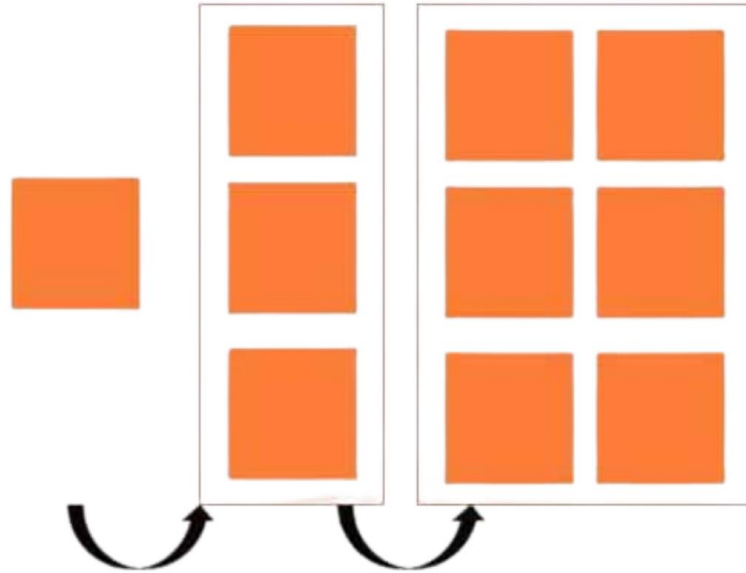
Herausforderungen des Realtime Big Data Stream Processing

Ursula Deriu / jugs.ch

Definition Big Data

Der Begriff '**Big Data**' bezeichnet Methoden zur Speicherung, Abfrage, Verarbeitung und Analyse von Daten mit Hilfe von **verteilten und horizontal skalierbaren Systemen**.

Horizontal skalierbar



Big Data RZ




Quelle Bilder: Google




Braucht Big Data auch Big Infrastructure?



Einsatzmöglichkeiten



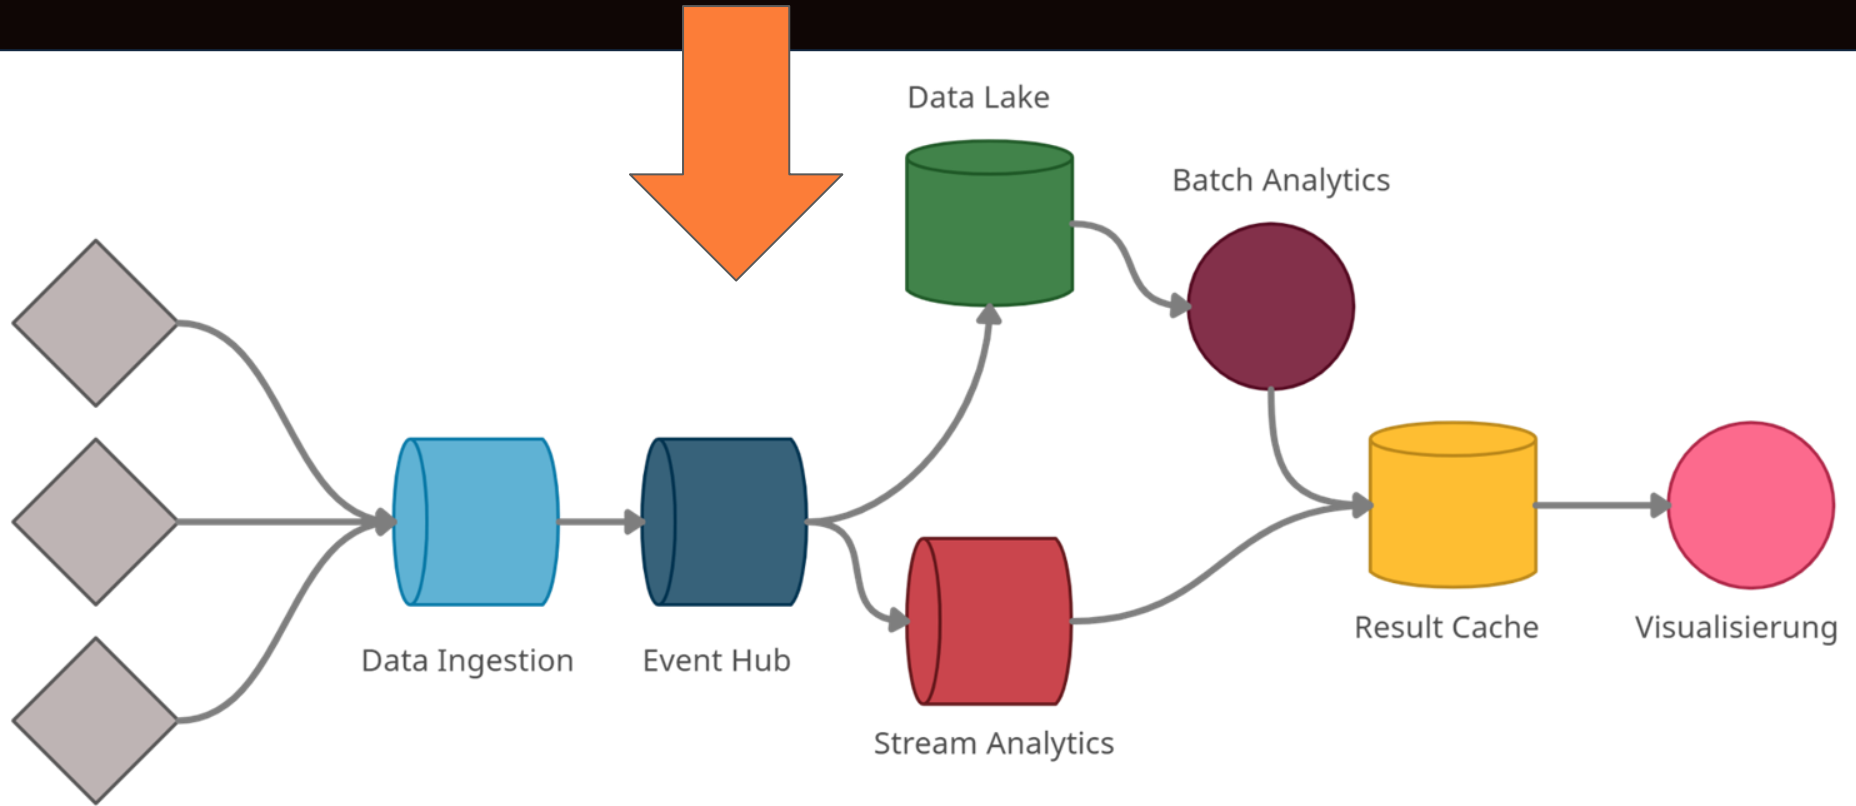
Anatomie einer Big-Data-Streaming-Pipeline



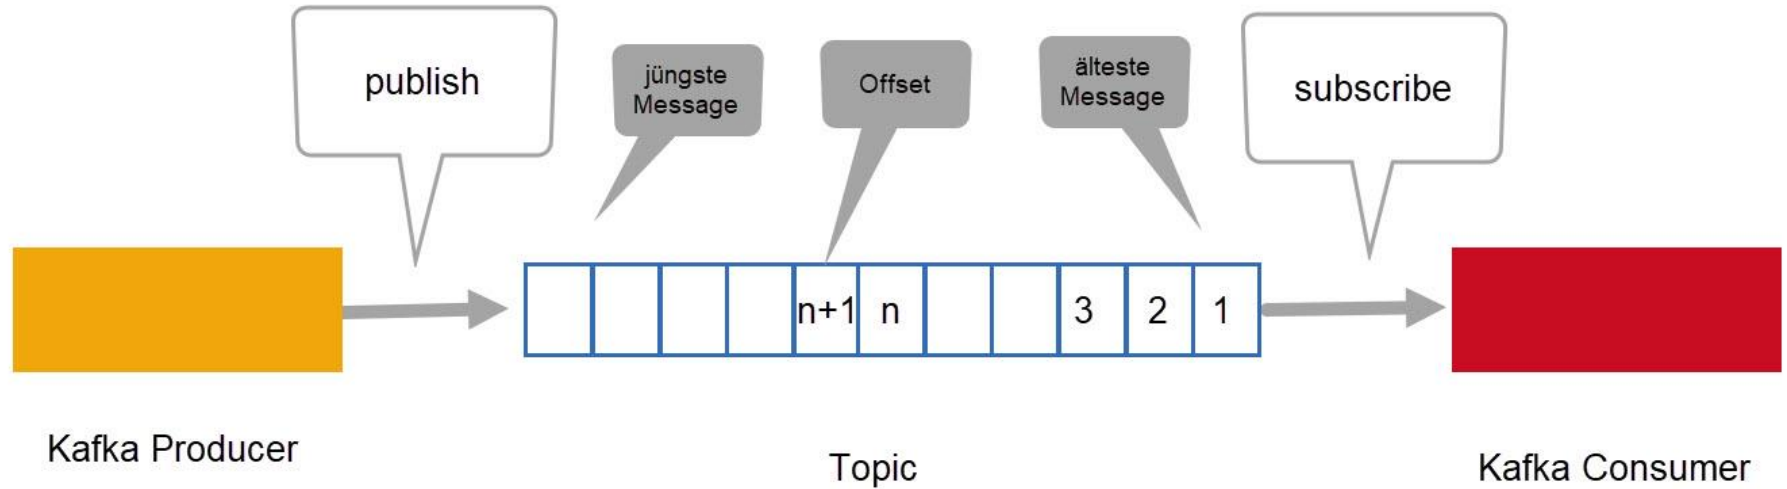
Datenquellen

Apache Kafka als Event Hub

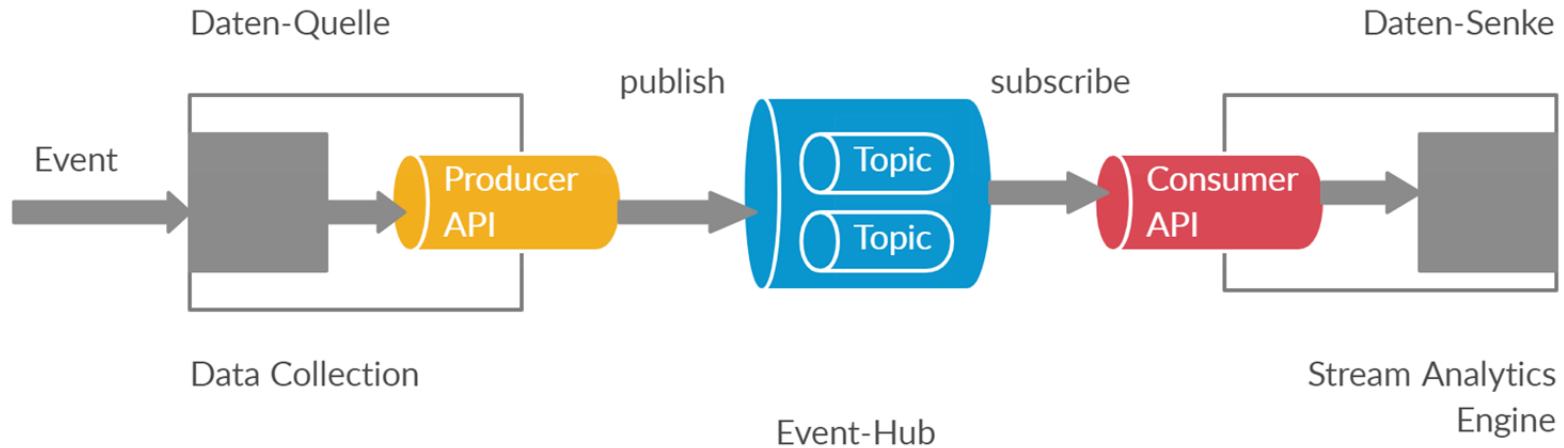
Apache Kafka als Event Hub



Warteschlange




Verteilter Event Hub



Apache Spark als Stream Analytics Engine

Apache Spark als Stream Analytics Engine




Datenquellen


Blick hinter die Kulissen von Apache Spark



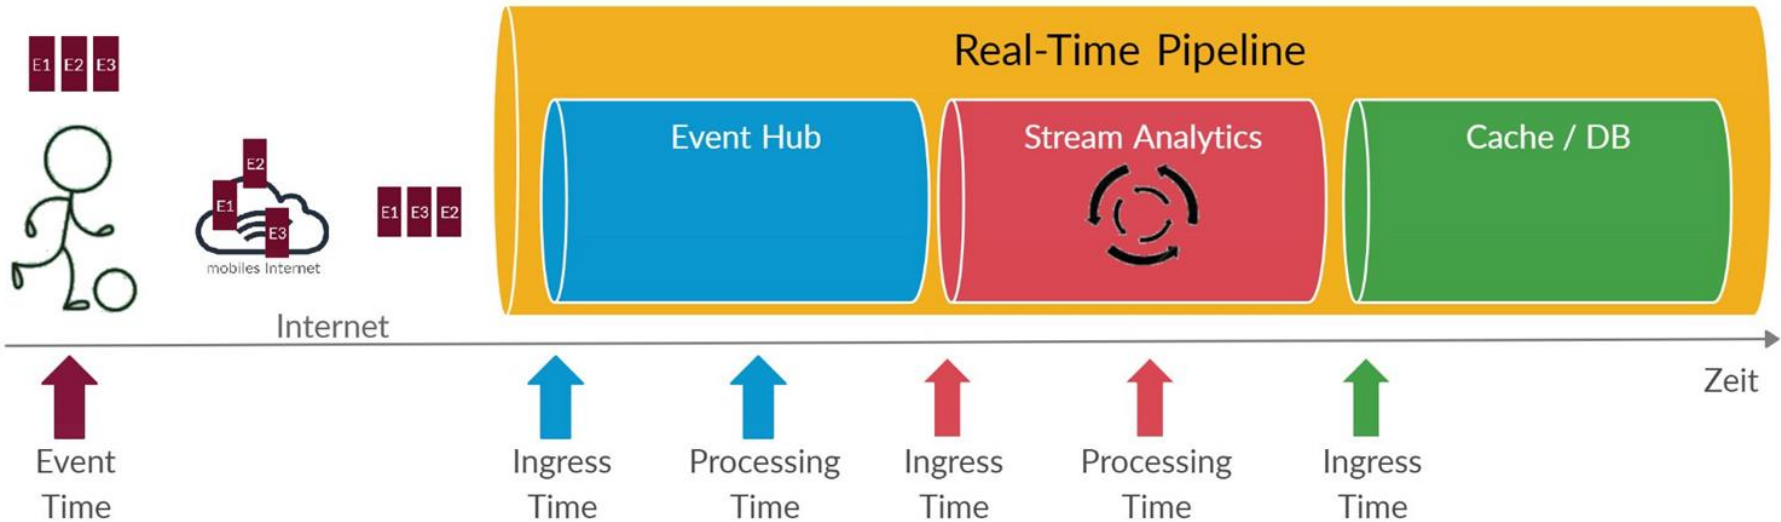
Messung




Analyse



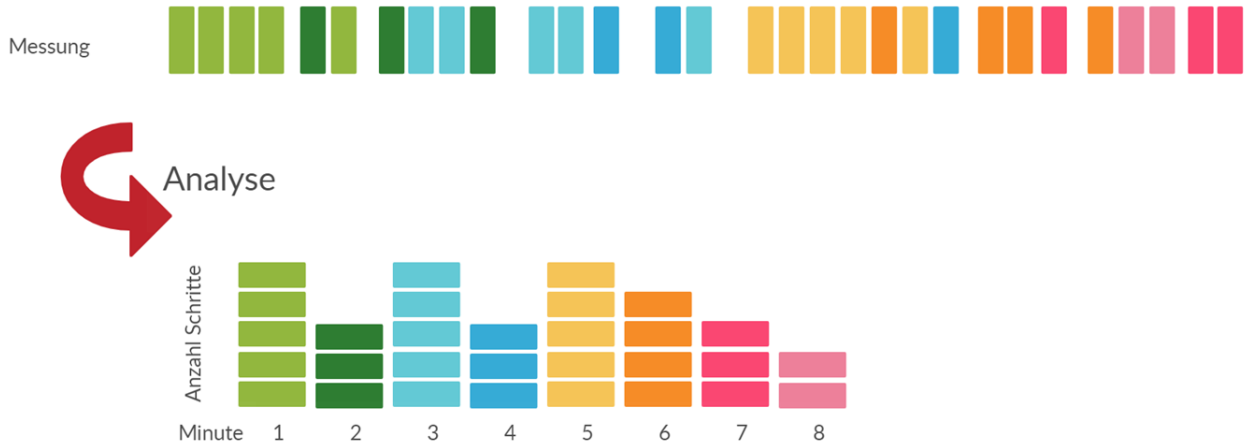
jetzt






Ziel: Exactly Once


```
schritteProMinute = schrittstream \
  .withWatermark("timestamp", "2 minutes") \
  .groupBy(
    window (schrittstream.timestamp, "1 minute"),
    schrittstream.spieler) \
  .count()
```




Data Stream Processing




Elementweise



Aggregation




Join



Machine Learning

Aggregation



Windows - Zeitfenster



Fixed




Sliding




Sessions

Beispiel: Herzüberwachung

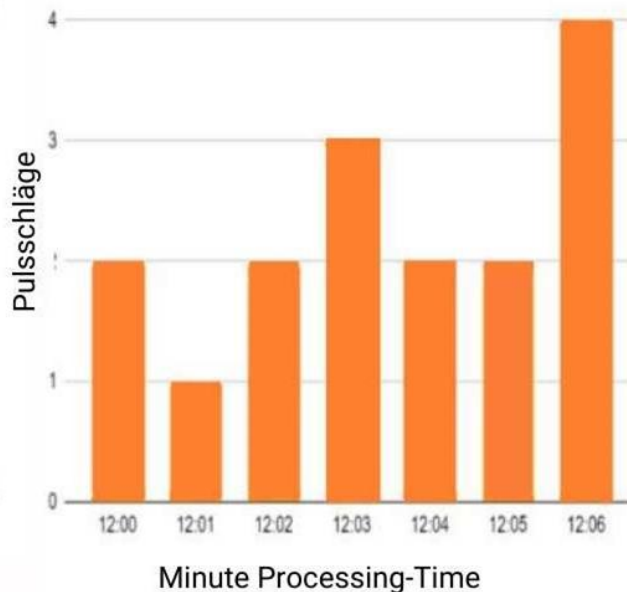
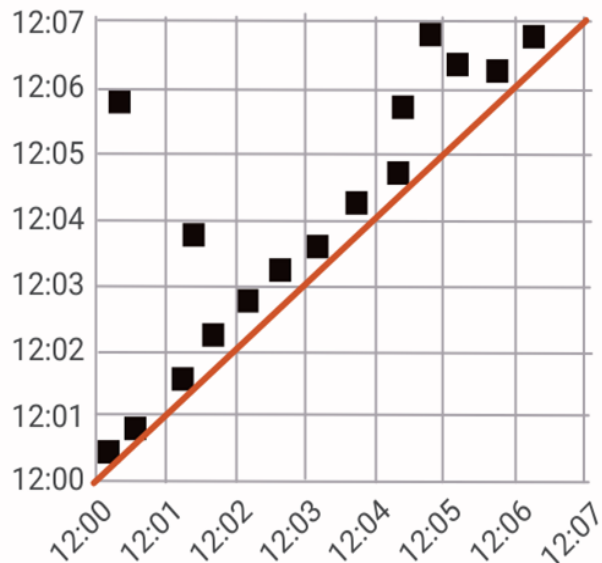


(c) Tirsus / Ursula Deriu

Windows - nach Processing Time


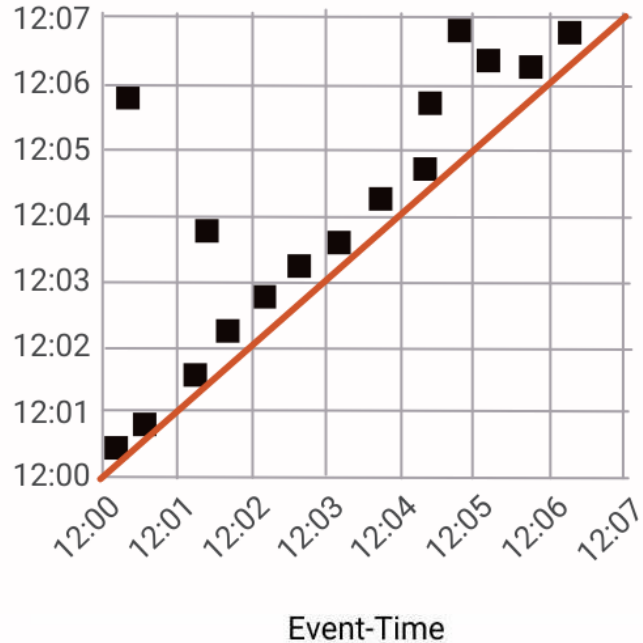


Processing-Time - "einfach erfüllbar"


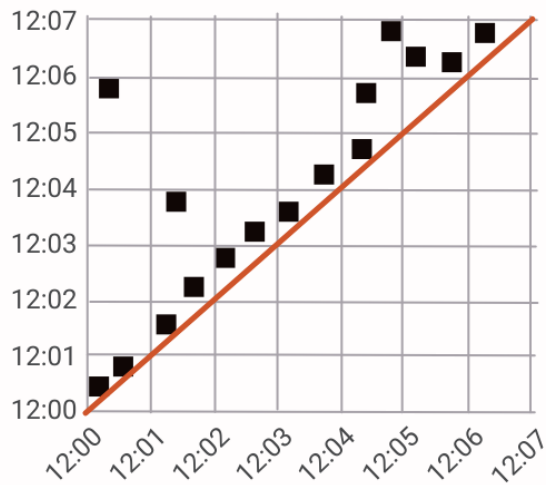


(c) Tirsus / Ursula Deriu

Erwartung: Event-Time



Event-Time vs. Processing-Time



Triggers und Watermarks



Animation hier:

<https://tirsus.com/stream-processing-zeit/>

Dive Deeper

Exactly Once Semantik

Exactly Once Jedes Event wird genau ein Mal verarbeitet.

At Most Once Wir nehmen in Kauf, dass im Fehlerfall Events verloren gehen.

At Least Once Wir nehmen in Kauf, dass im Fehlerfall Events mehrfach verarbeitet werden.

Exactly Once -> Transaktionen

Source Code mit DataFrames/DataSets

```
val lines = spark
    .readStream.format("kafka")
    .option("kafka.bootstrap.servers", "kafka-2:9092")
    .option("subscribe", "test-002")
    .load()
val values = lines
    .selectExpr("CAST(value AS STRING)")
val words = values
    .as[String]
    .flatMap(_.split(" "))
val query = words
    .writeStream
    .format("memory")
    .queryName("kafka2")
    .outputMode("update")
    .start()
```

Logischer Plan

readStream

transform

aggregate

write

Physicher Plan

Source

Optimierte
Operationen

Sink

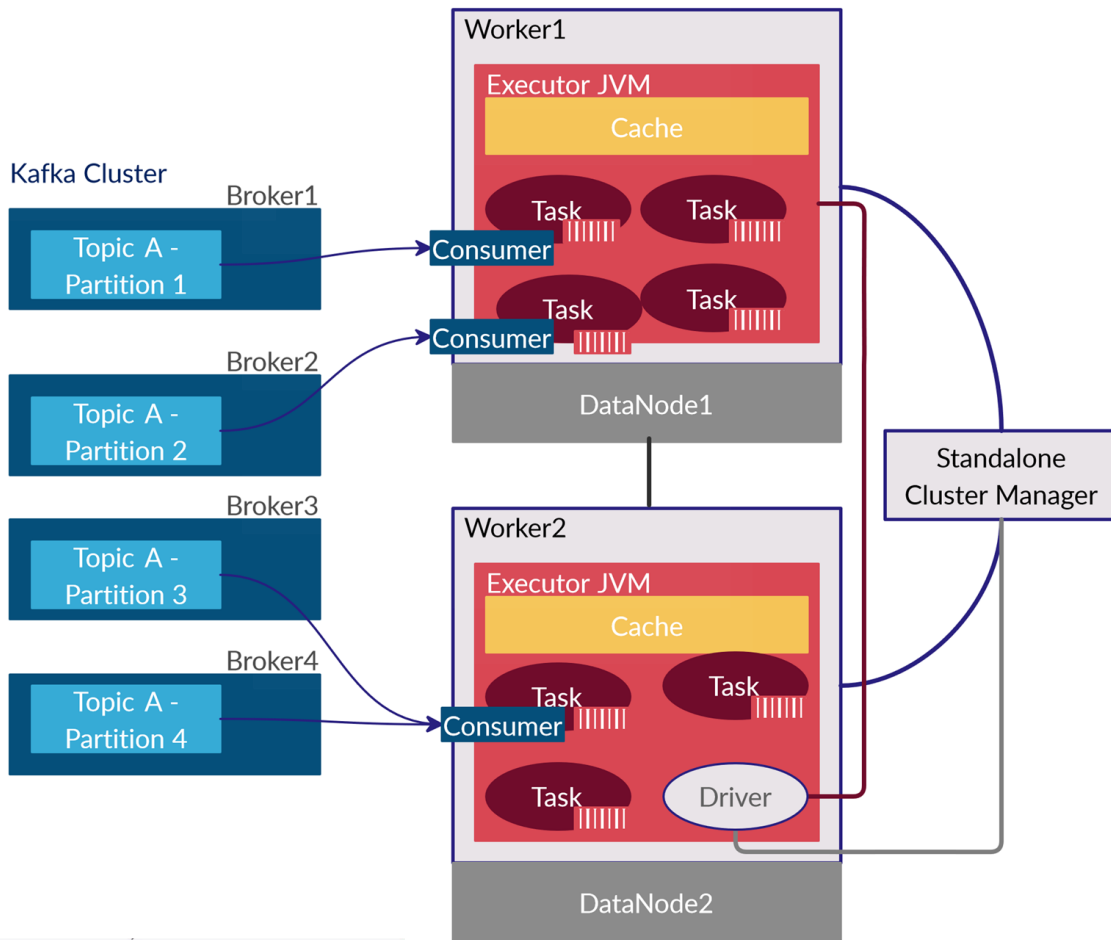
Folge inkrementeller Ausführungspläne

t=1

t=2


t=3

Zeit



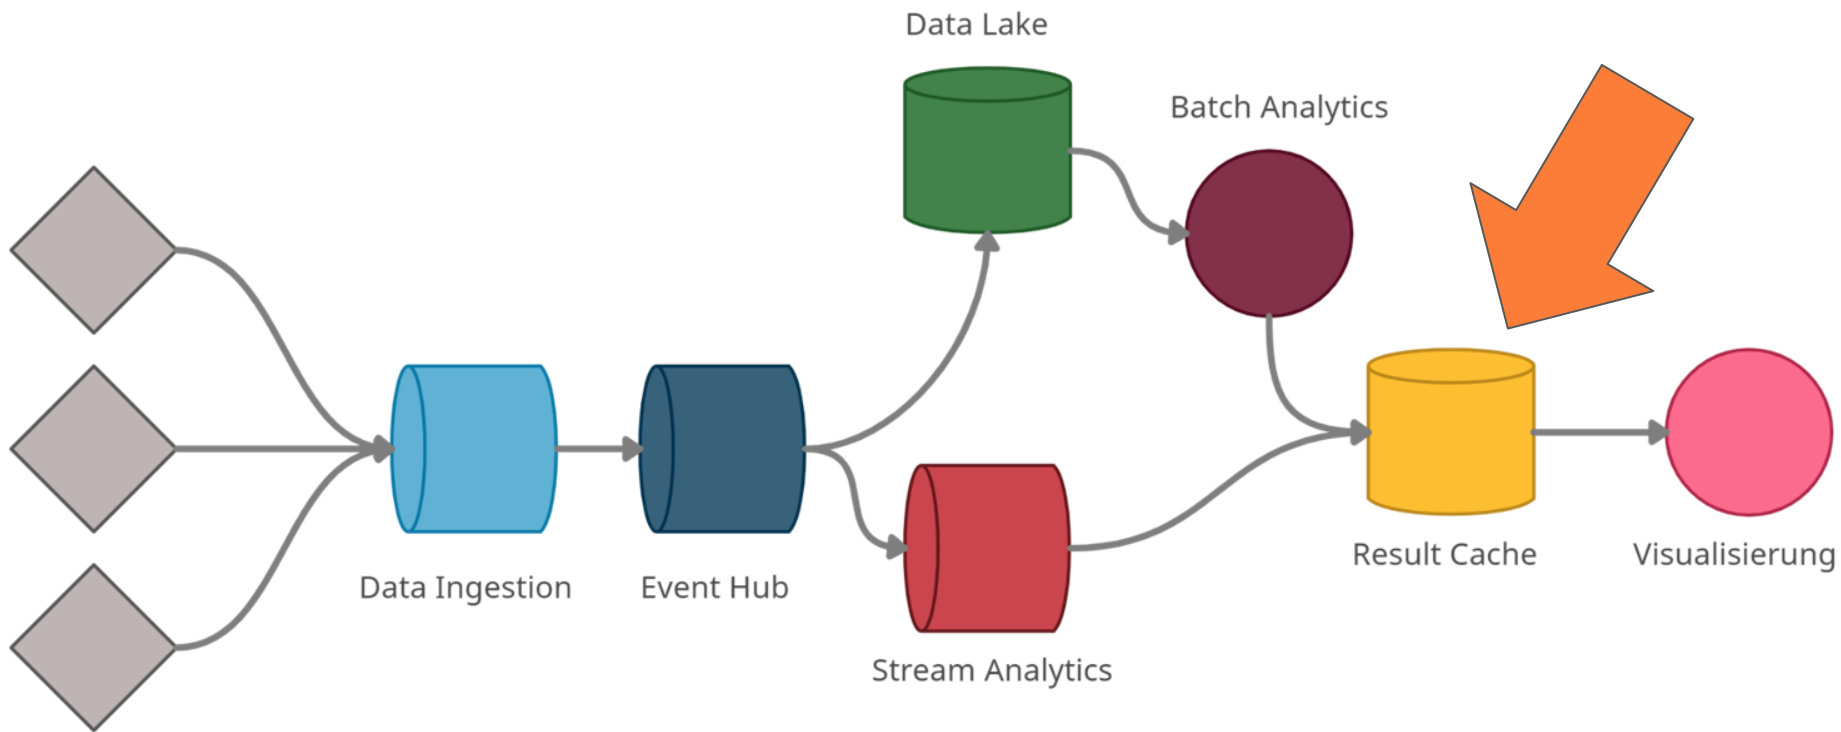
Spark liest von Kafka

Workerausfall?






Der Result Cache

Der Result Cache



Datenquellen

Idempotente Schreiboperationen



Trade-Offs der Real-Time Analyse großer Datenströme



Free E-Mail-Training



<https://tirsus.com/pipeline/>

Ursula Deriu
ursula.deriu@tirsus.com

 <https://www.linkedin.com/in/ursuladeriu>
tirsus.com/big-data

Q & A