Scalable Data Ingestion for Stream Processing and Beyond

Gabriel Antoniu

Joint work with Ovidiu Marcu, Alexandru Costan, Maria S. Pérez

9th JLESC workshop, UTK, Knoxville, April 15, 2019

Inria inventors for the digital world

POLITÉCNICA

INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES RENNES

UMR IRISA

From Big Data to Fast Data

Volume

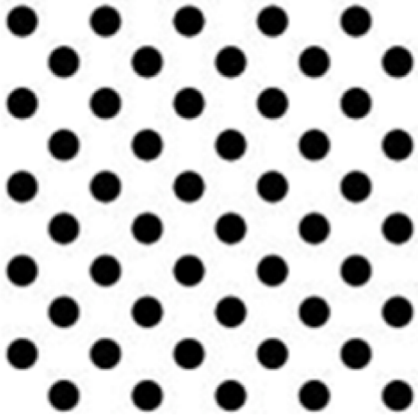

Data at rest

Stationary
Static

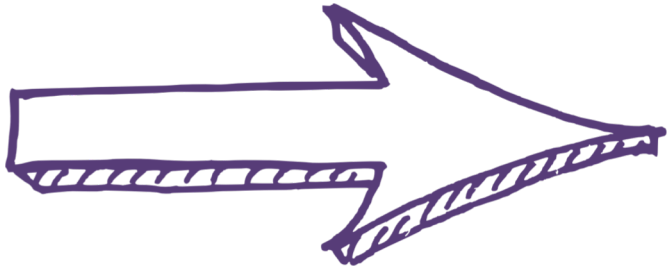

Velocity

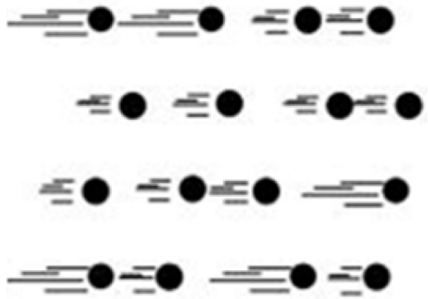

Data in motion

Fluid
Dynamic

	Batch	Streaming
		
Correctness	Exact results	Approximate results
Latency	High-latency	Low-latency
Cost	Stateless	Stateful

State of the art until recently: Lambda Architectures

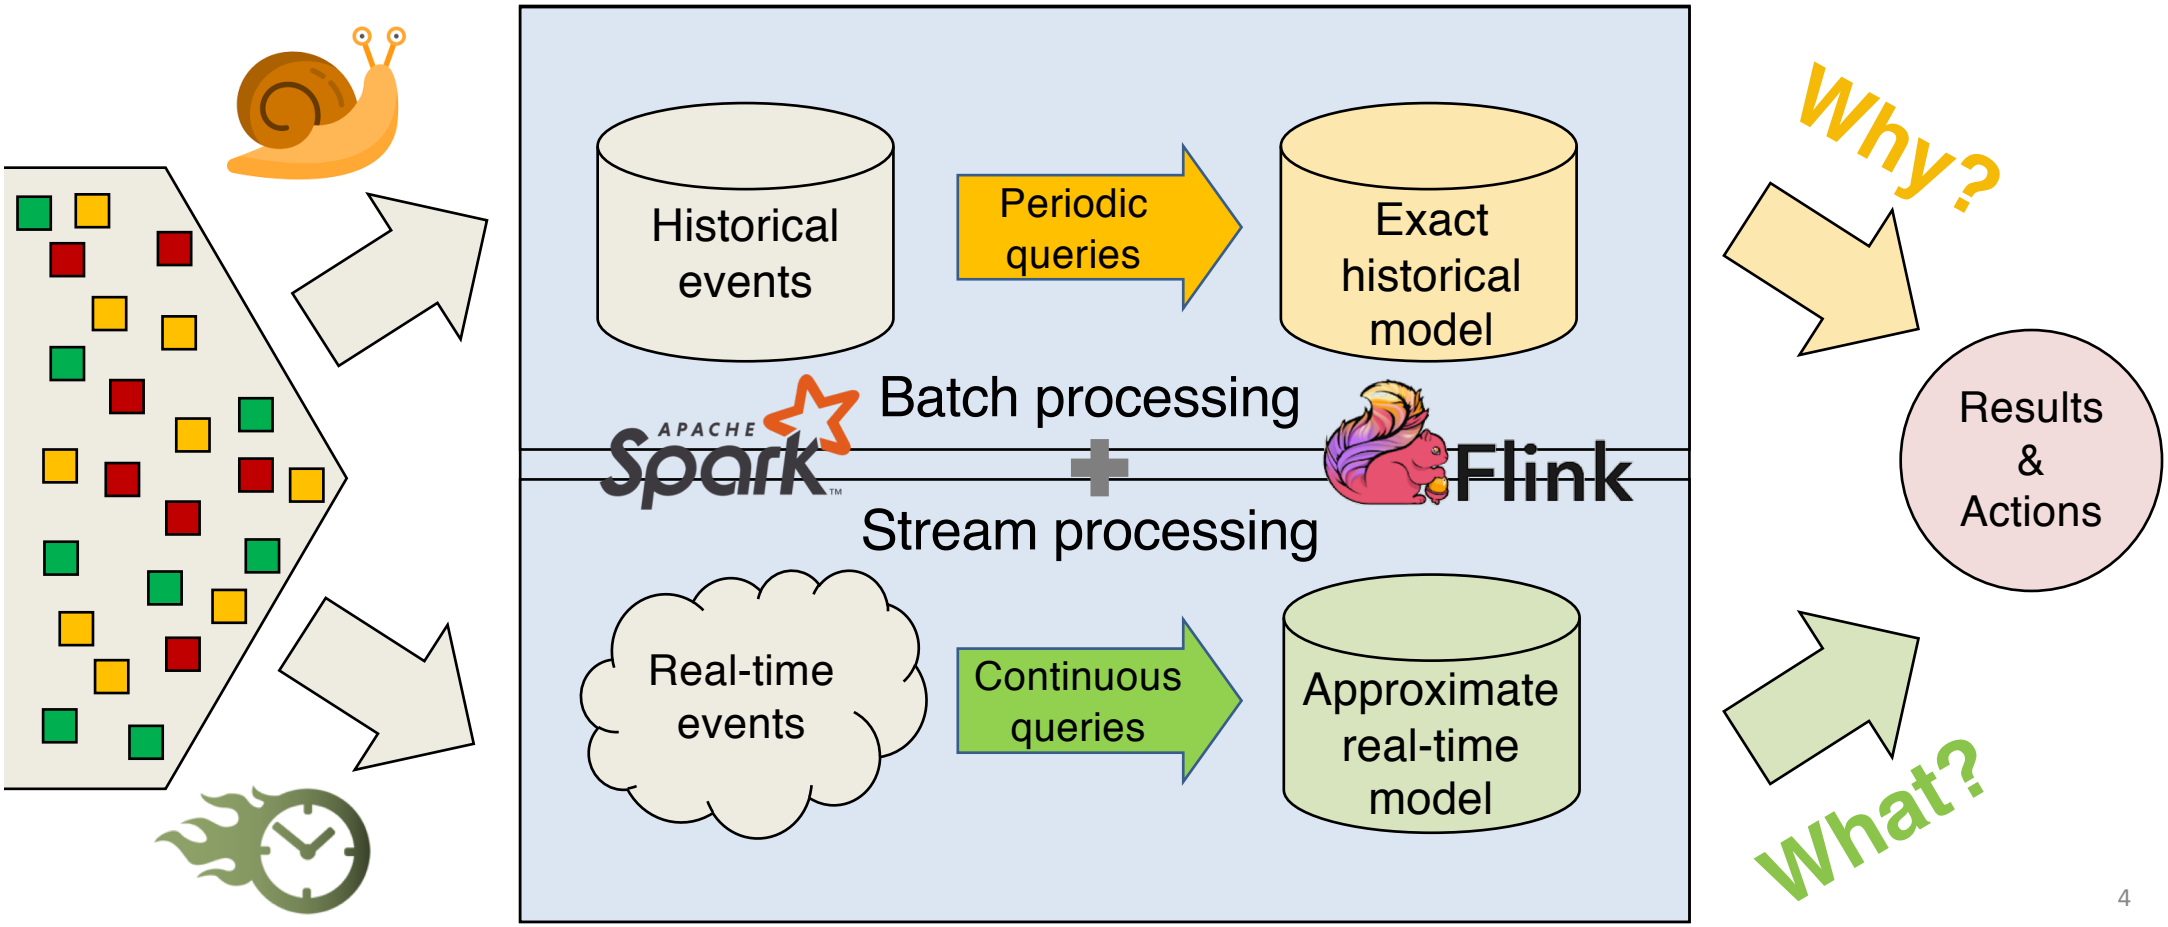

The streaming pipeline: latency happens

Cloud

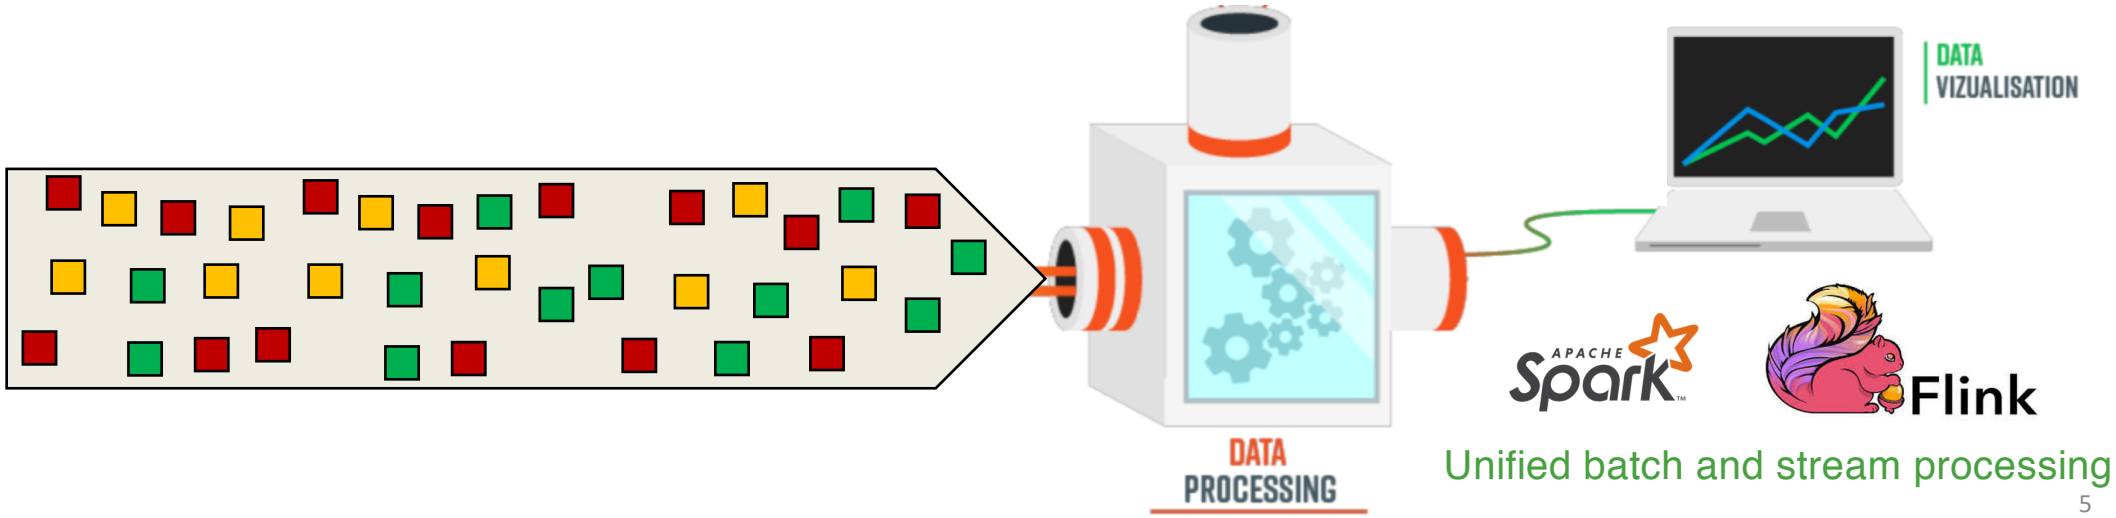

What is ingestion ?

- Collect data from various sources
  → *producers*
- Deliver them for processing / storage
  → *consumers*
- Optionally: buffer, log, pre-process

Ingestion determines the processing performance

State of the art: Apache Kafka

Spotify

50 nodes, average 200K events/s

LinkedIn

400 nodes, peak 3.2M events/s

Limitations

- Scalability
- Data duplication

The KerA approach to ingestion

- Scalability → **Dynamic partitioning**
  - Enables seamless elasticity

- Data duplication → **Unified ingestion and storage**
  - Support for both
    - Streams (unbounded data)
    - Objects (bounded data)

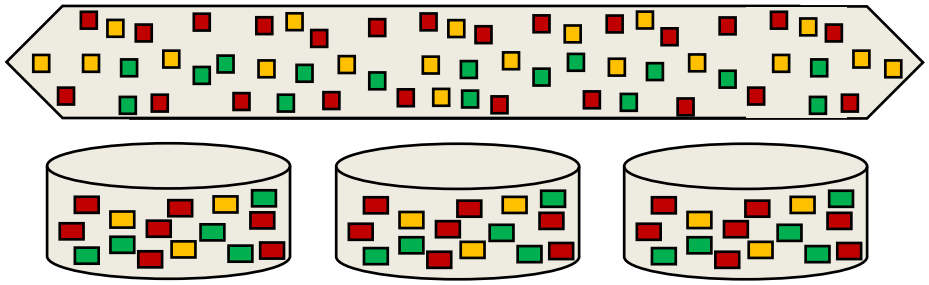

Zoom: scalability

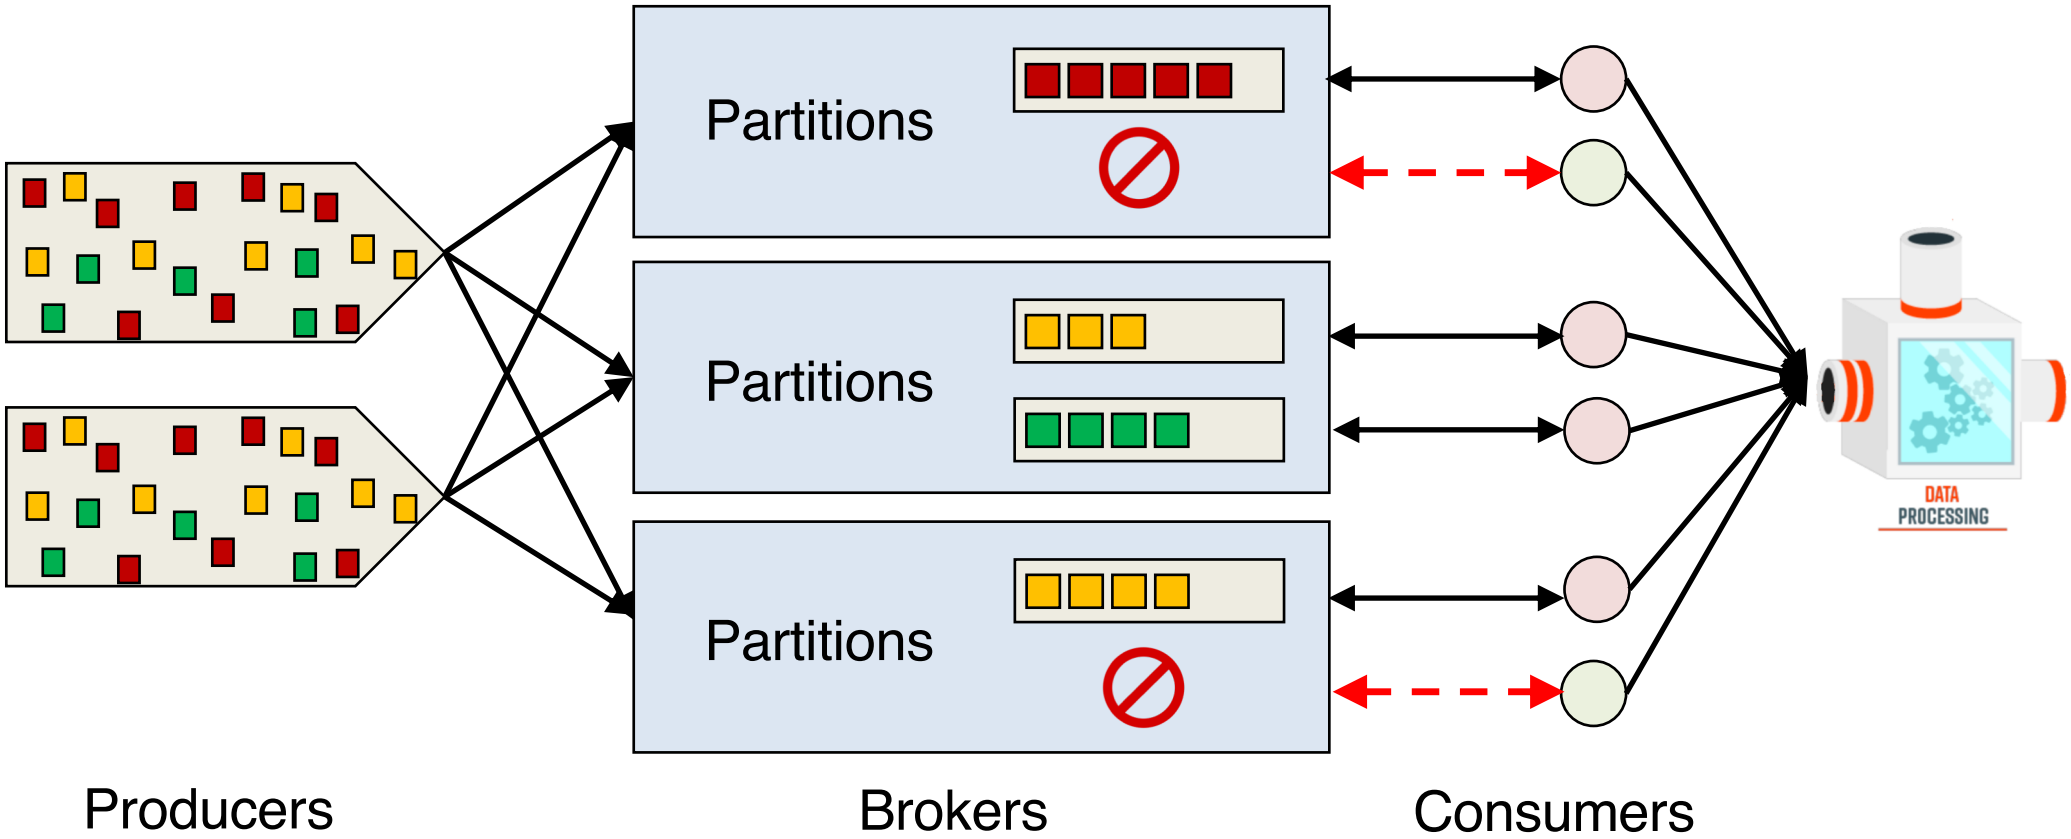

Each partition is statically associated with one consumer: limited scalability

KerA: dynamic partitioning

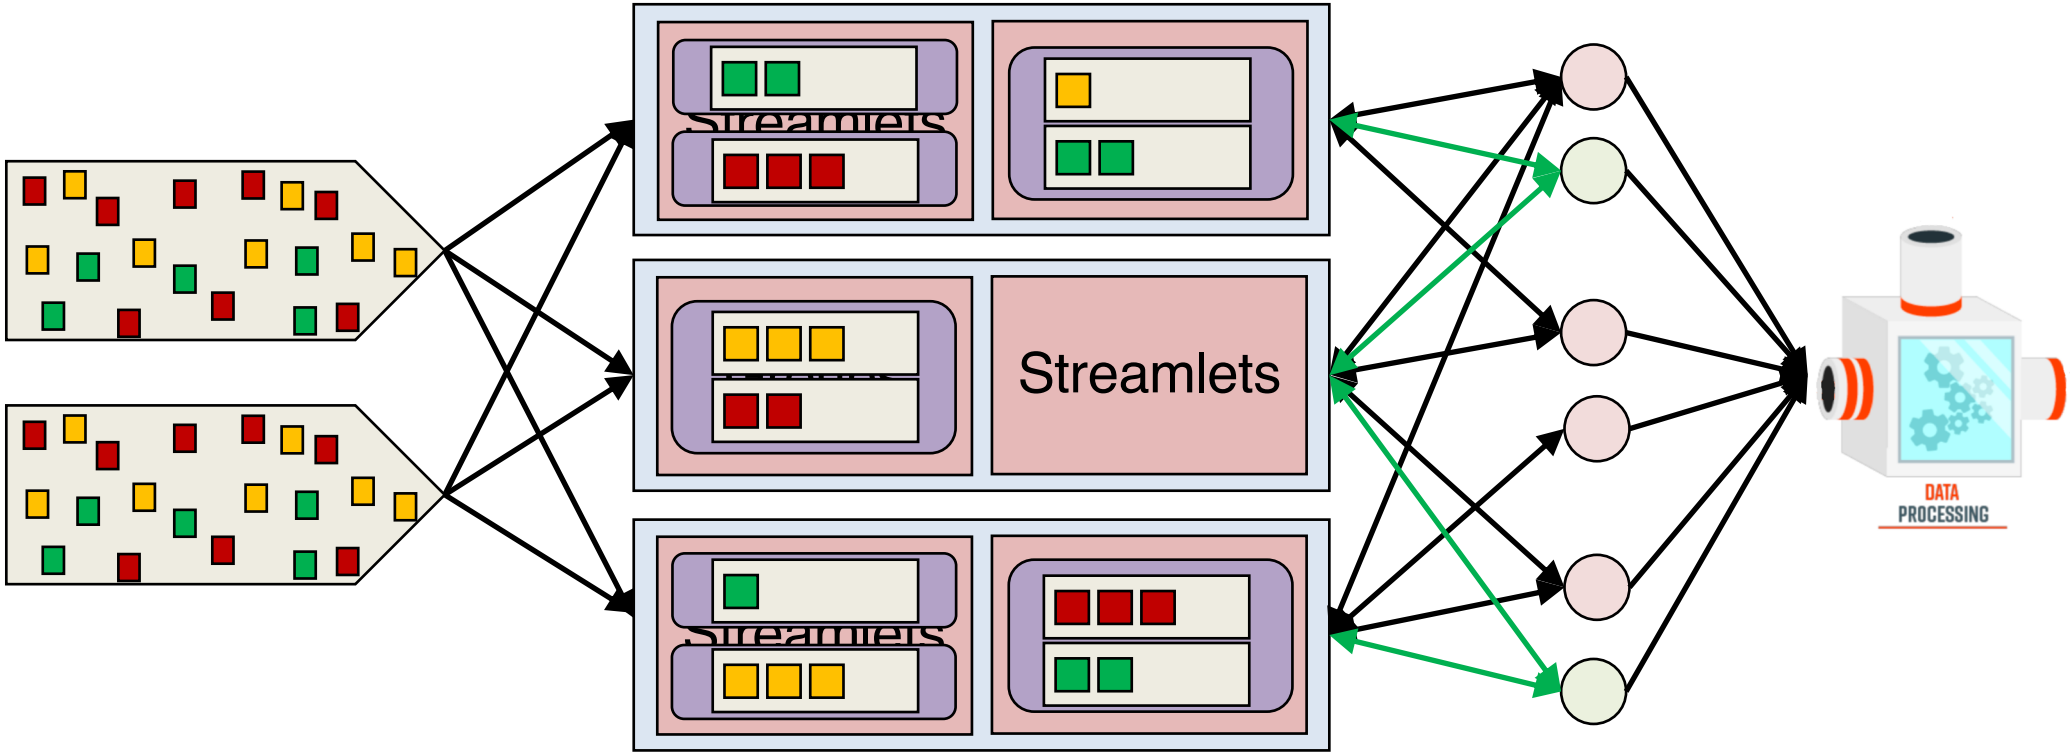

- **Streamlets:** logical stream containers; #streamlets > #brokers
- **Groups:** created and processed dynamically; maximum #active groups per broker

Zoom: data duplication

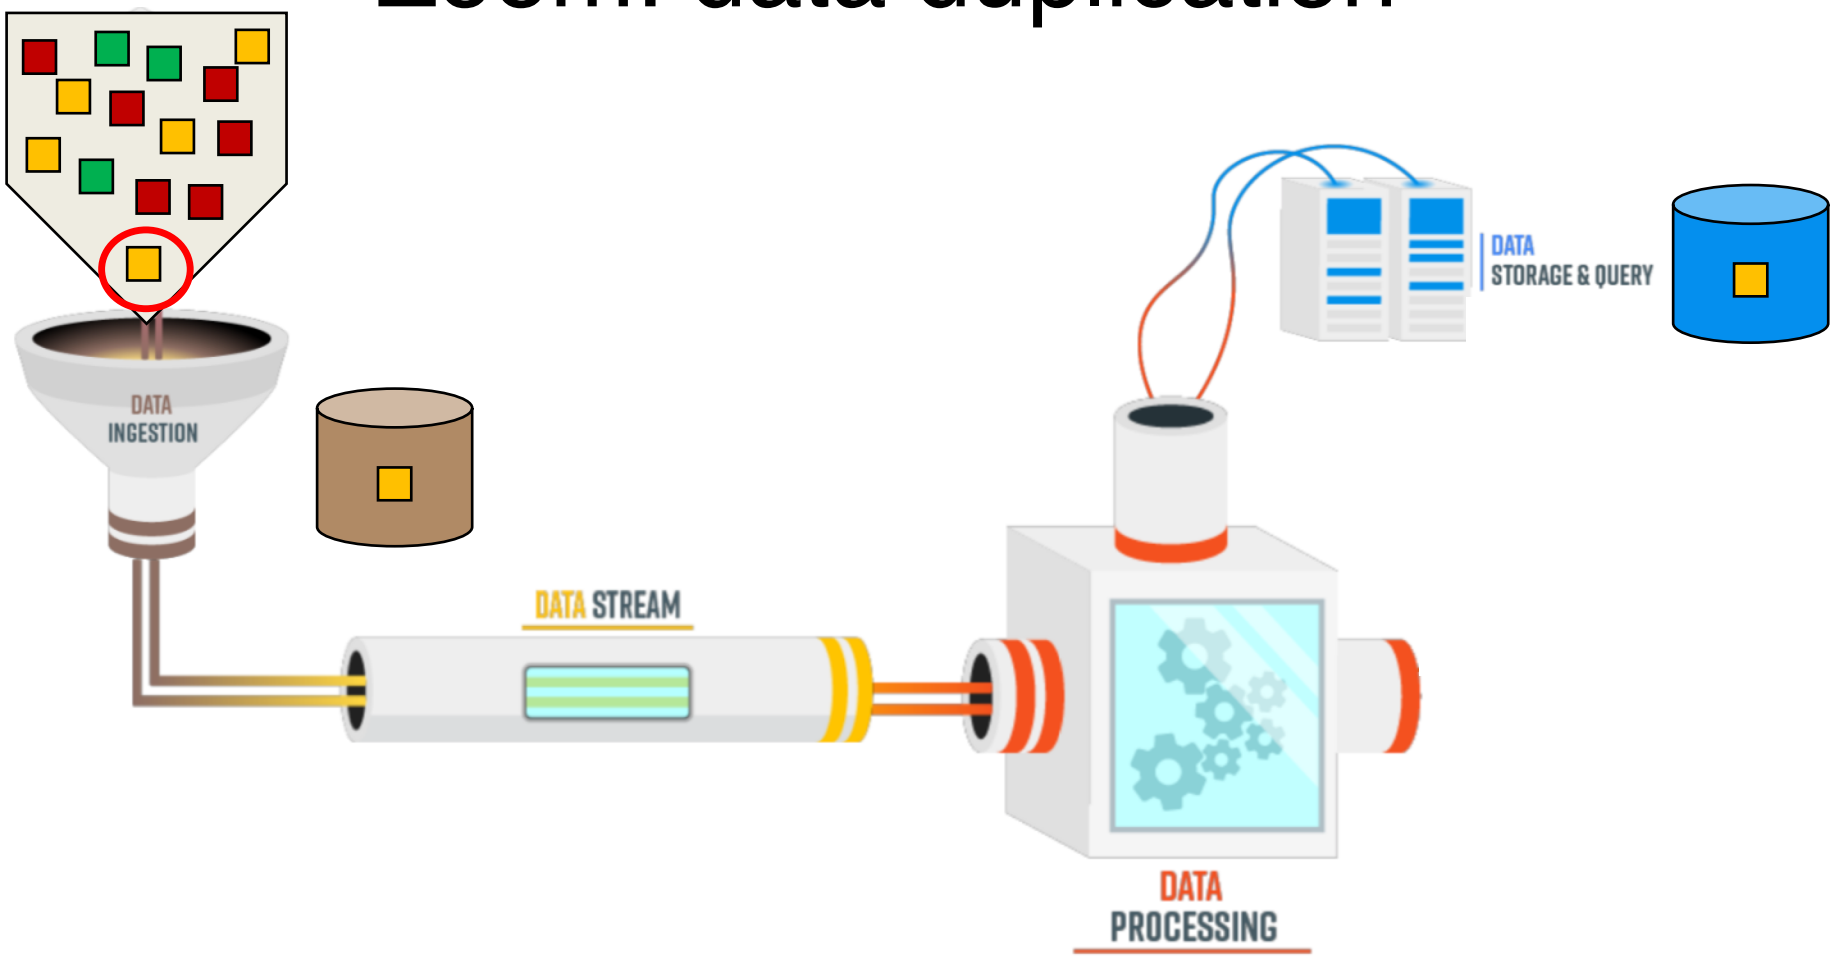

Increased network and storage overheads

KerA: unified ingestion and storage

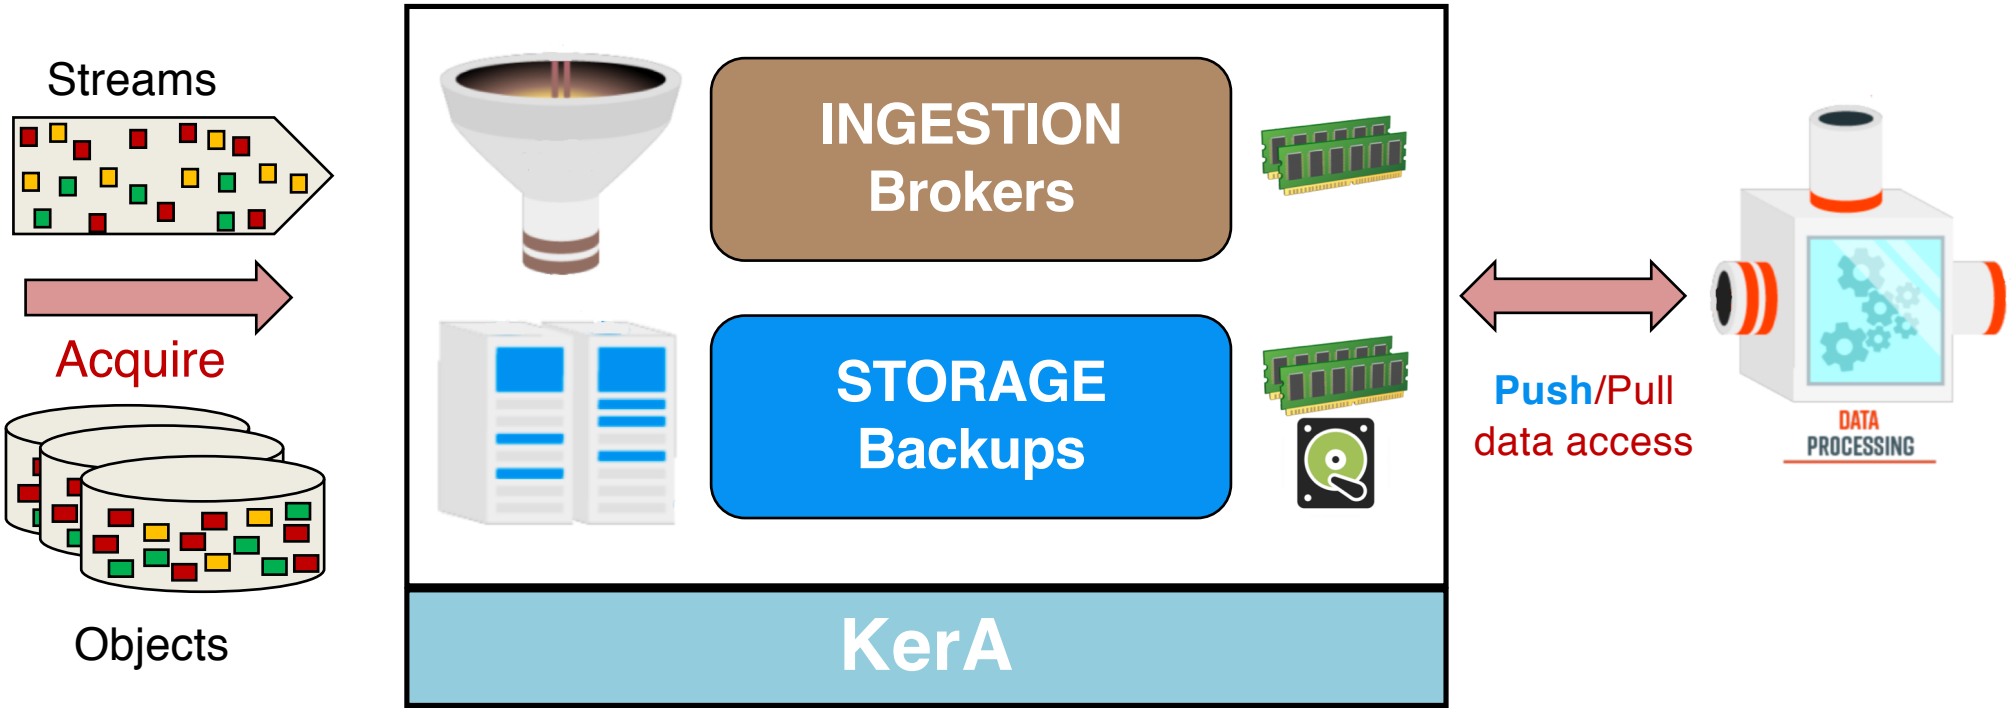

Move less data, process them faster

Common data model for streams and objects

Evaluating scalability

Vertical

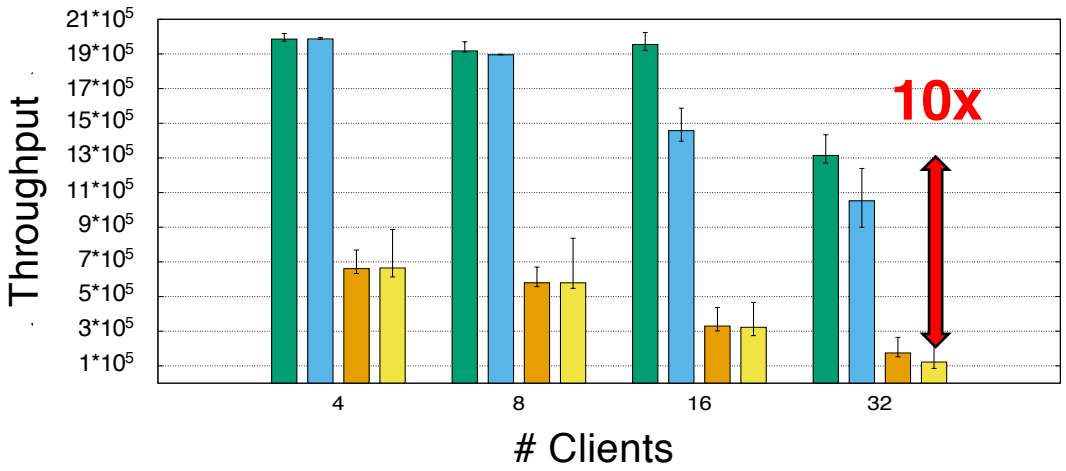

4 brokers, 32 partitions,
128KB request size, 100B records

Horizontal

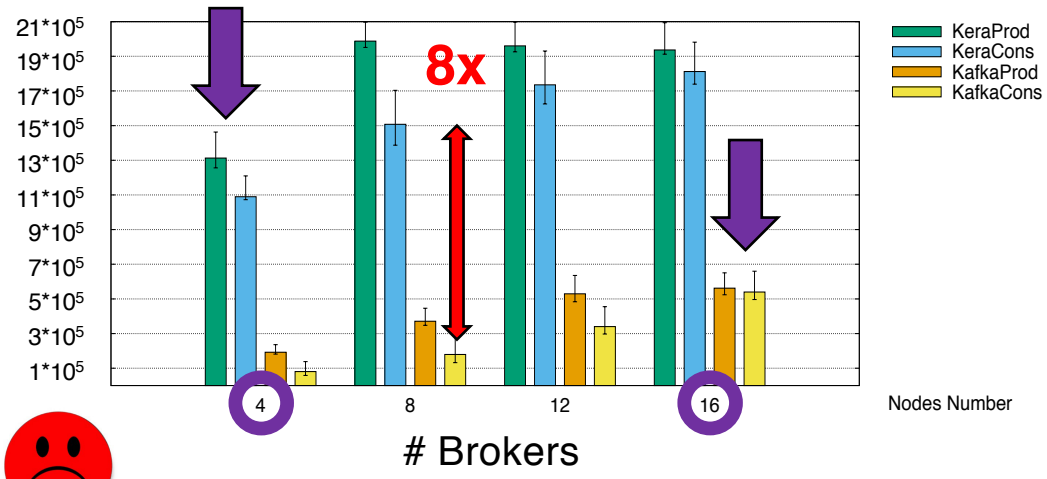

64 clients, 32 partitions,
1MB request size, 100B records

2x better throughput
with 75% less resources

Our vision: hybrid analytics architecture

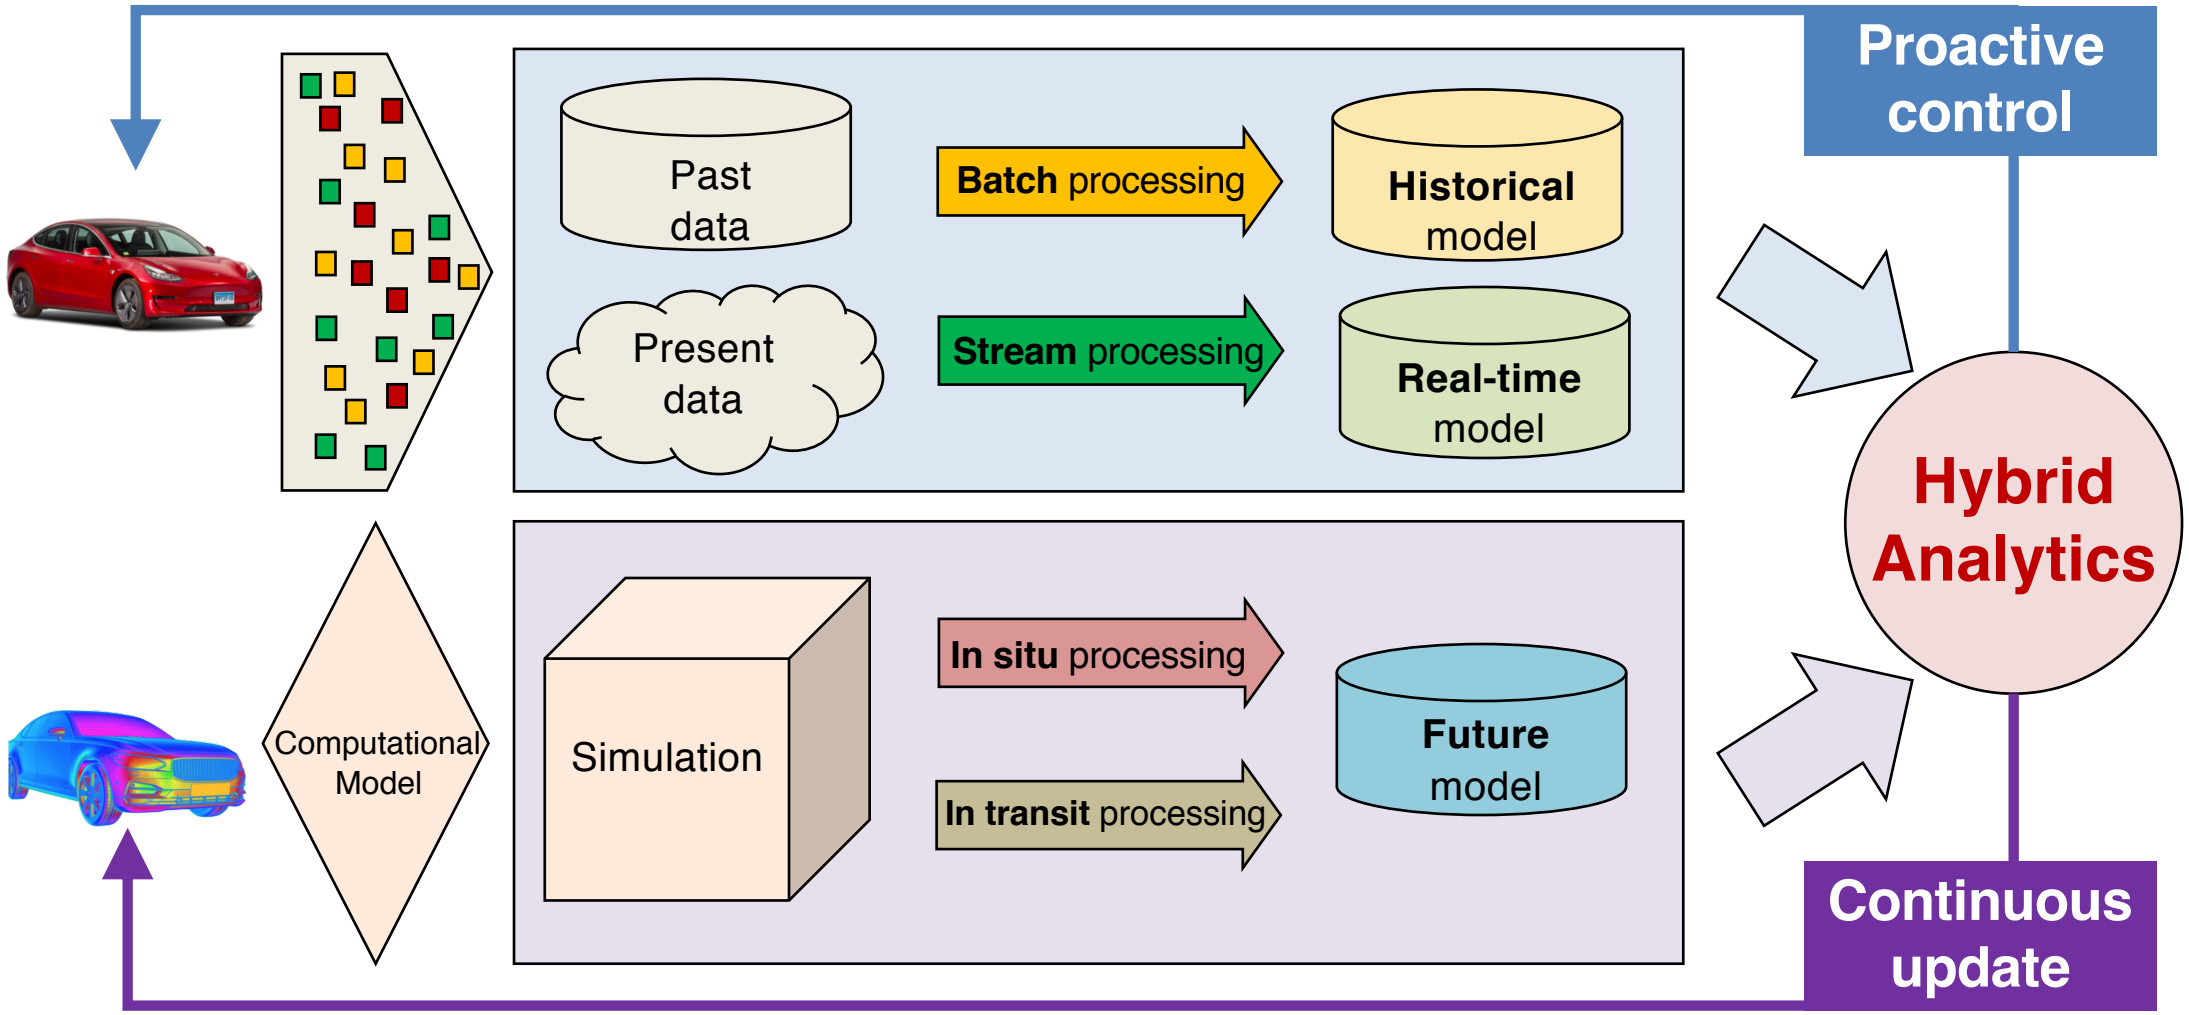

Hybrid analytics: processing architecture

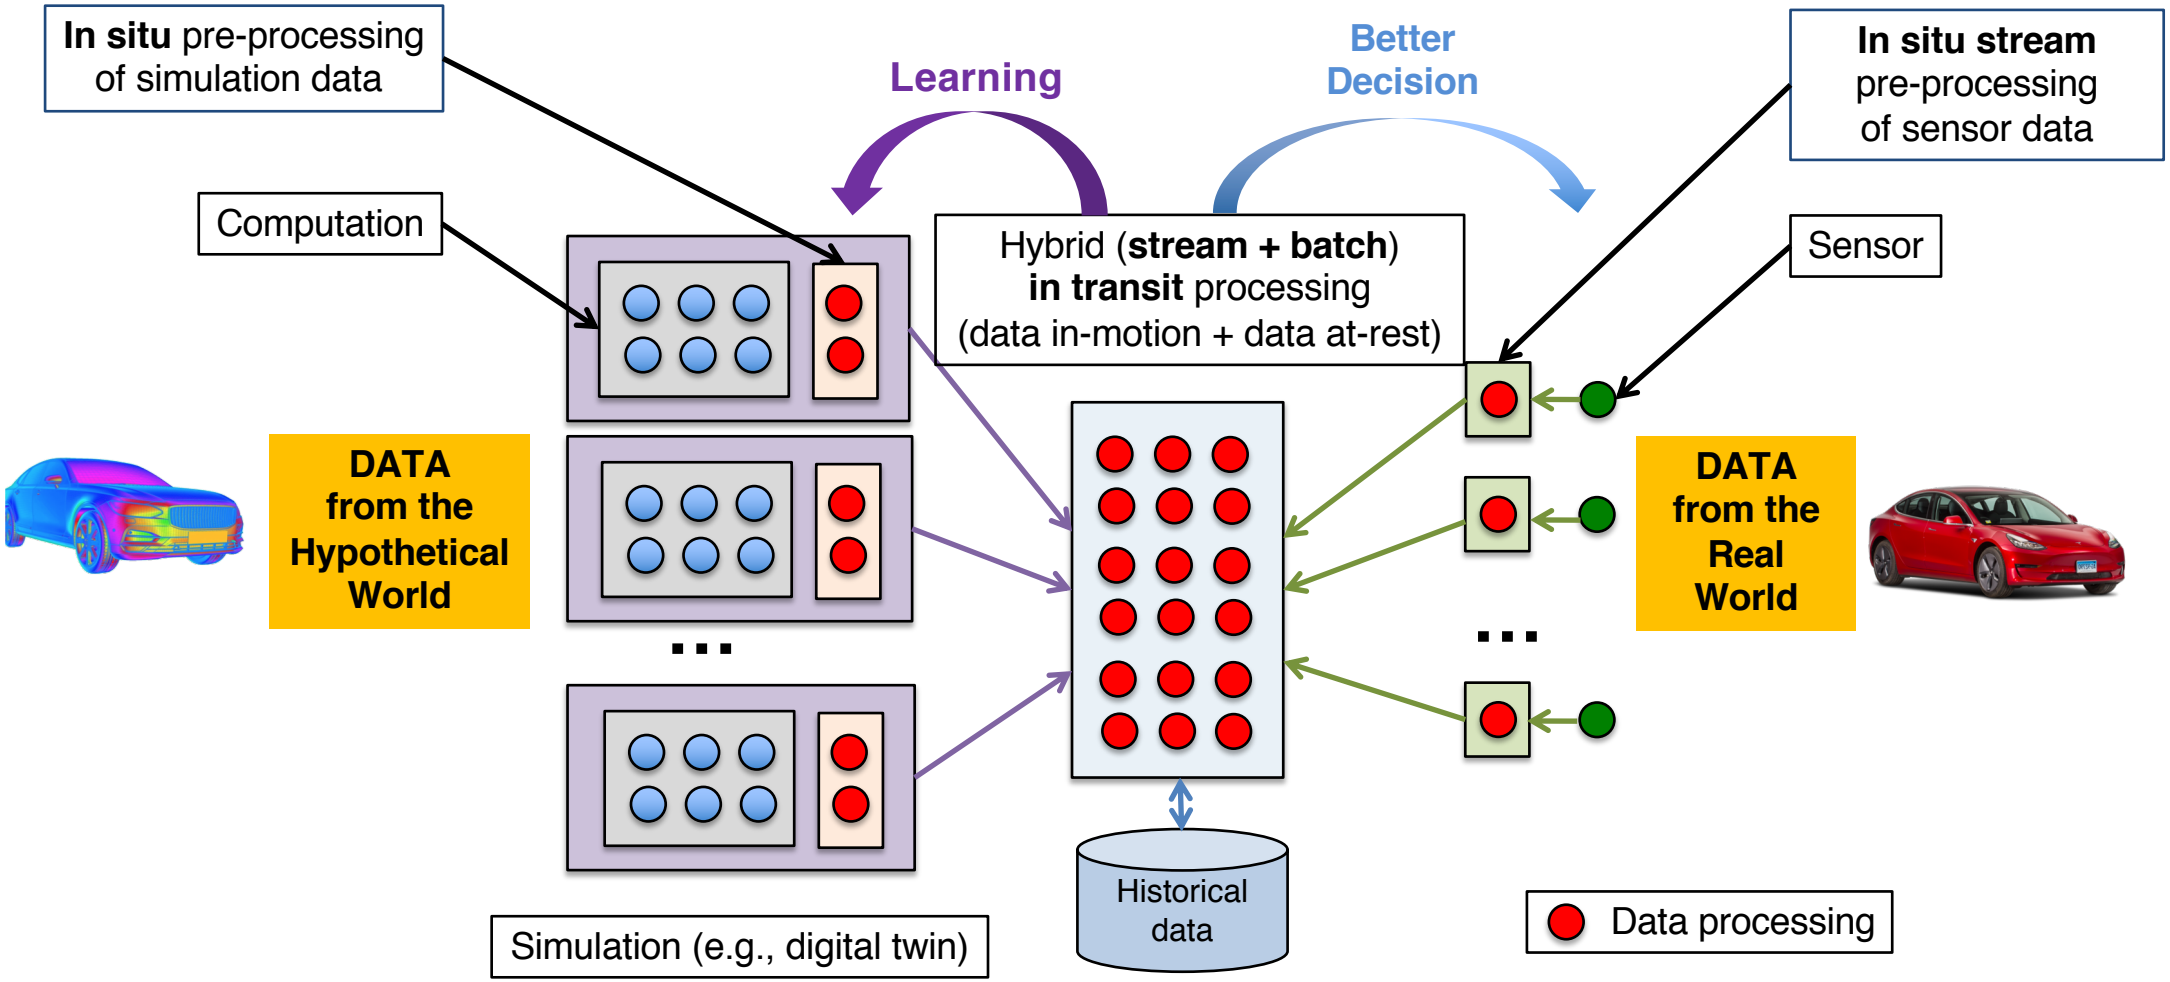

Hybrid analytics architecture

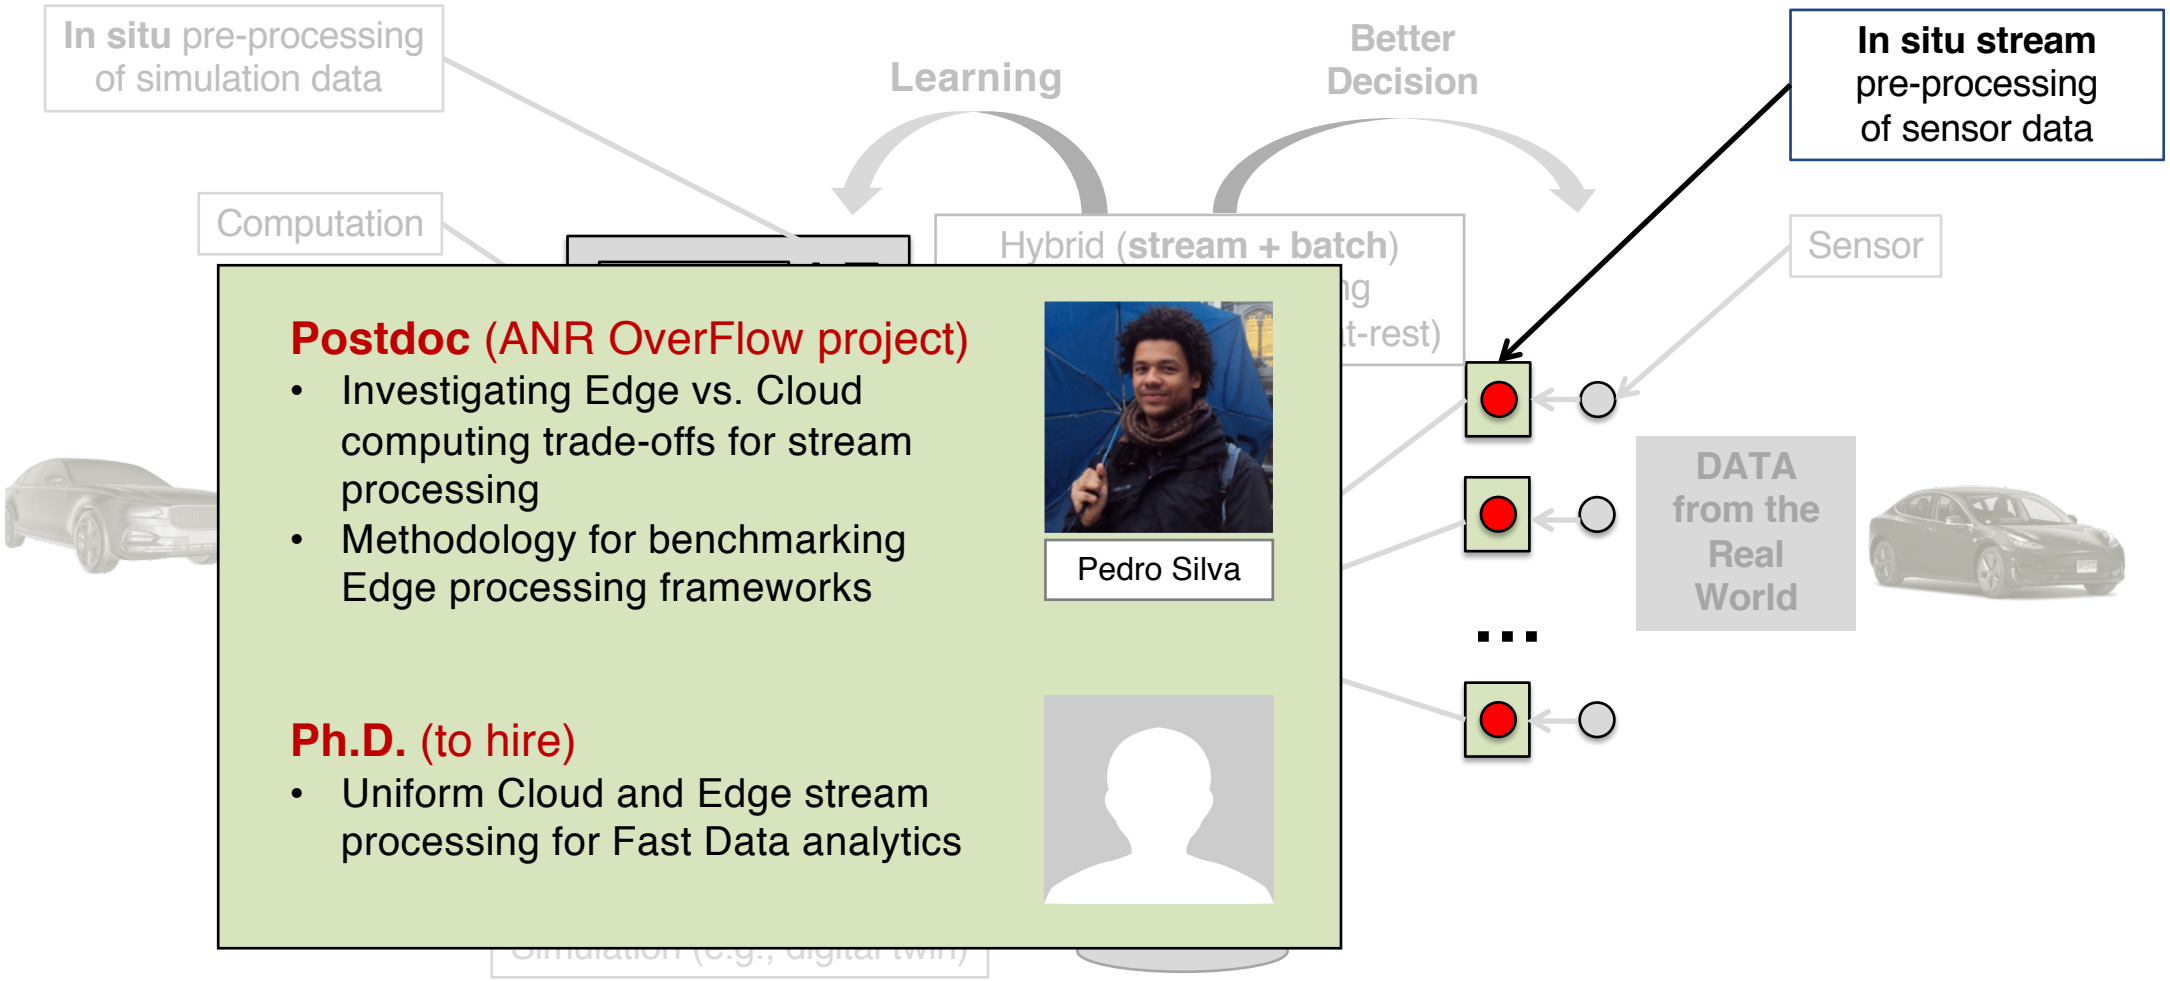

Postdoc (ANR OverFlow project)

- Investigating Edge vs. Cloud computing trade-offs for stream processing
- Methodology for benchmarking Edge processing frameworks

Pedro Silva

Ph.D. (to hire)

- Uniform Cloud and Edge stream processing for Fast Data analytics

Hybrid analytics architecture

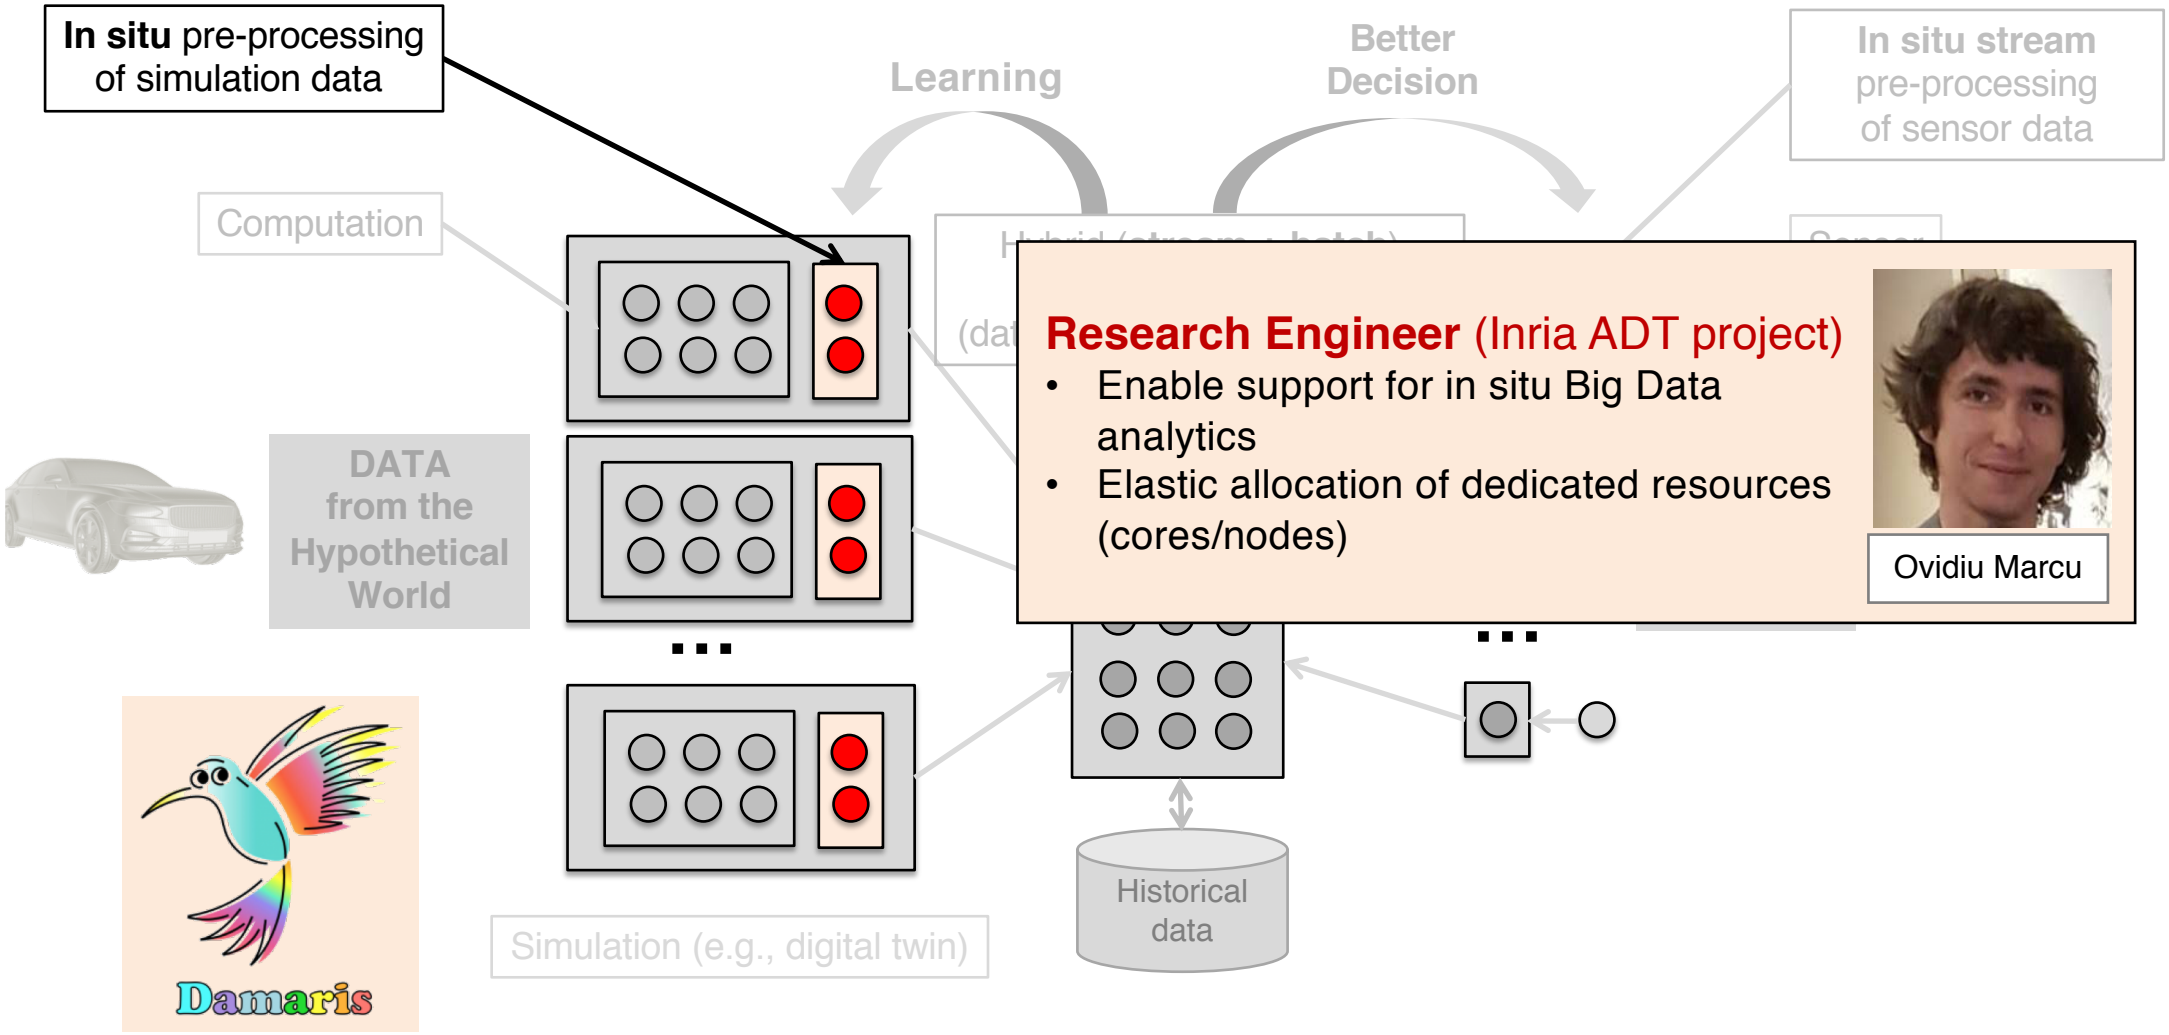

Research Engineer (Inria ADT project)

- Enable support for in situ Big Data analytics
- Elastic allocation of dedicated resources (cores/nodes)

Ovidiu Marcu

Hybrid analytics architecture

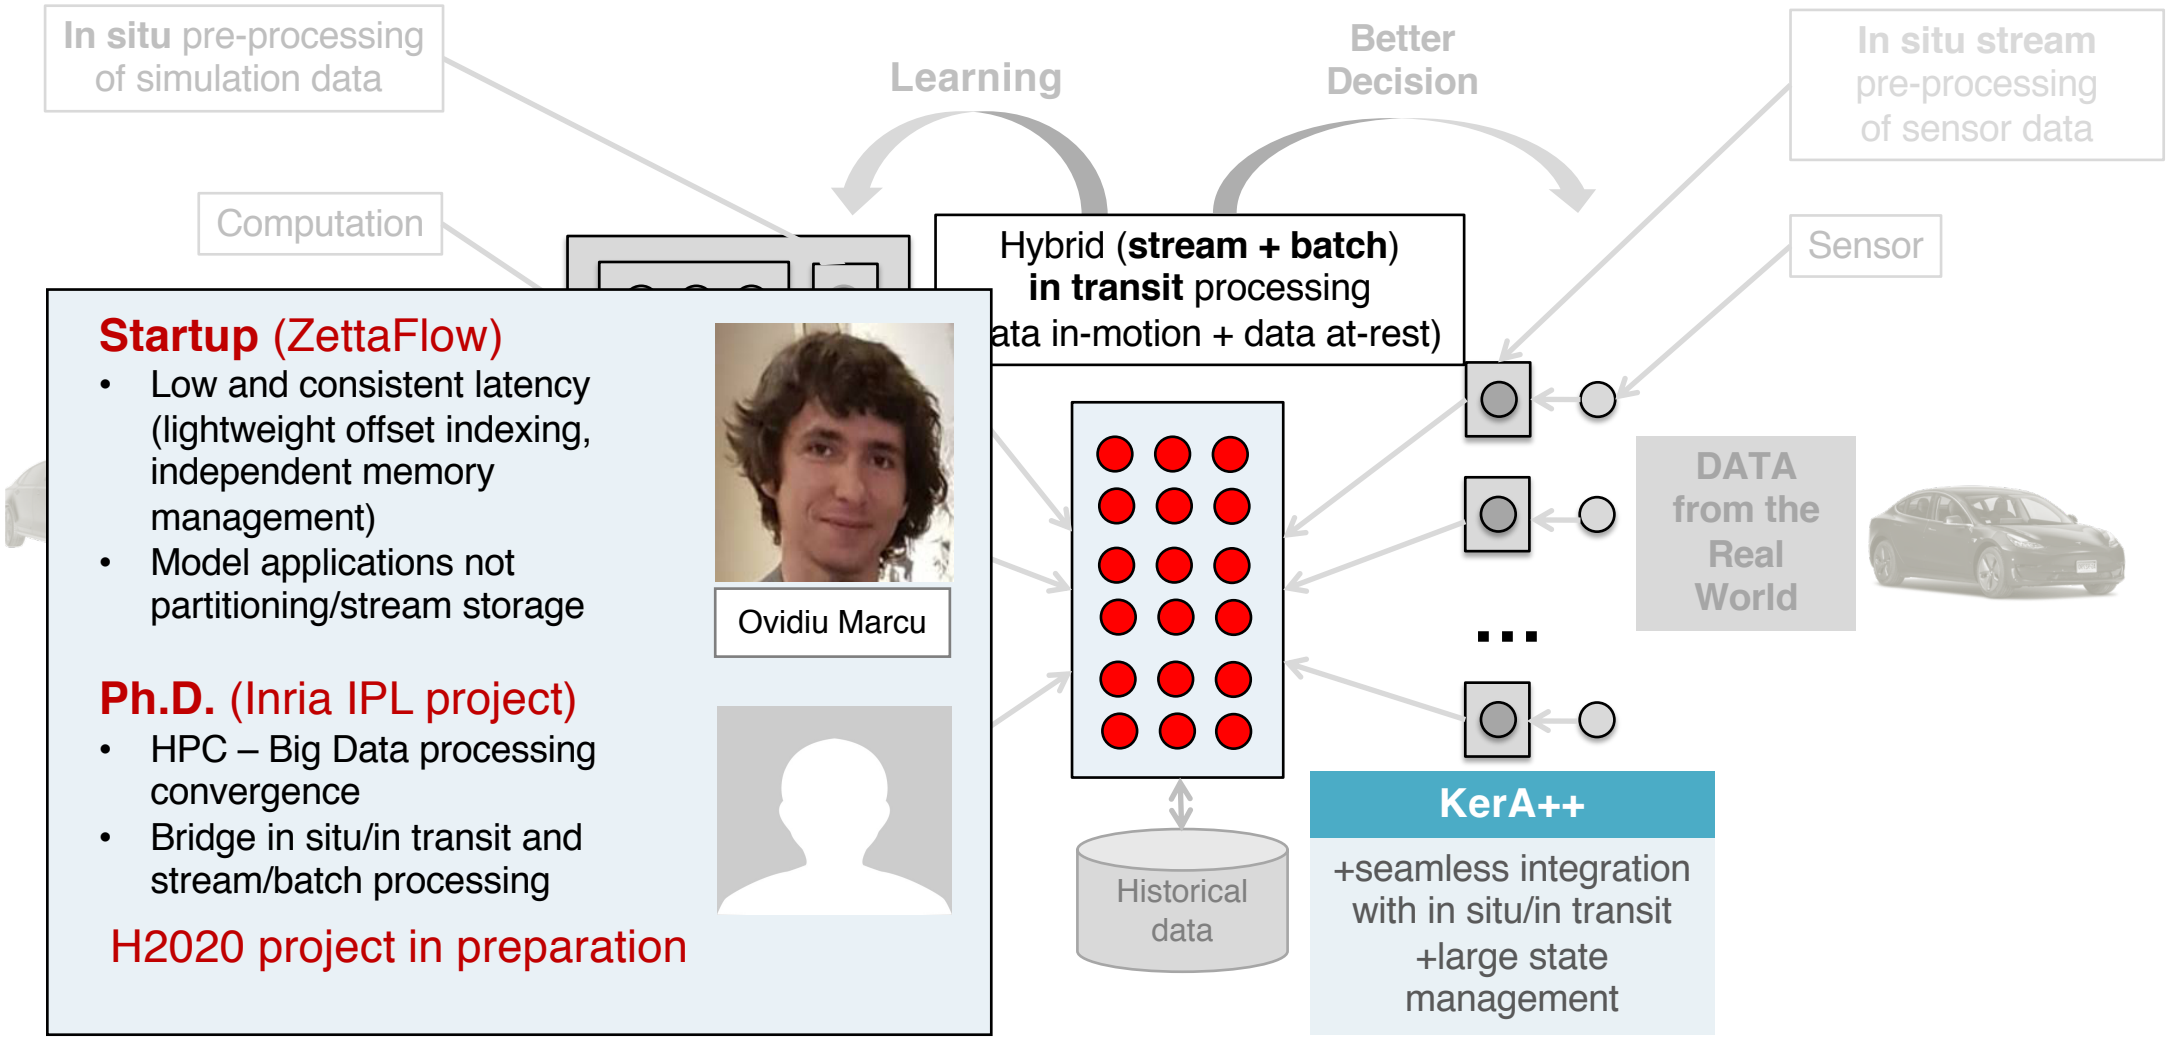

Startup (ZettaFlow)

- Low and consistent latency (lightweight offset indexing, independent memory management)
- Model applications not partitioning/stream storage

Ovidiu Marcu

Ph.D. (Inria IPL project)

- HPC – Big Data processing convergence
- Bridge in situ/in transit and stream/batch processing

H2020 project in preparation

Thank You!