Welcome to the second Workshop on Big data Open Source Systems (BOSS)

September 10th, 2016

Co-located with VLDB 2016

Tilmann Rabl & Sebastian Schelter

Hands on Big Data

- 6 parallel tutorials
- 6 systems
  - Open source
  - Publicly available
- Presenters
  - System experts
- Hands on
  - This is not a demo!
- You can pick two!

But why?

- Mike Carey
  - Doing It On Big Data: a Tutorial/Workshop
  - Driving force
- Other people involved
  - Volker Markl
  - Kerstin Forster
- Second instance
  - Last time: 8 systems
  - Tell us what you think
  - Email: rabl@tu-berlin.de

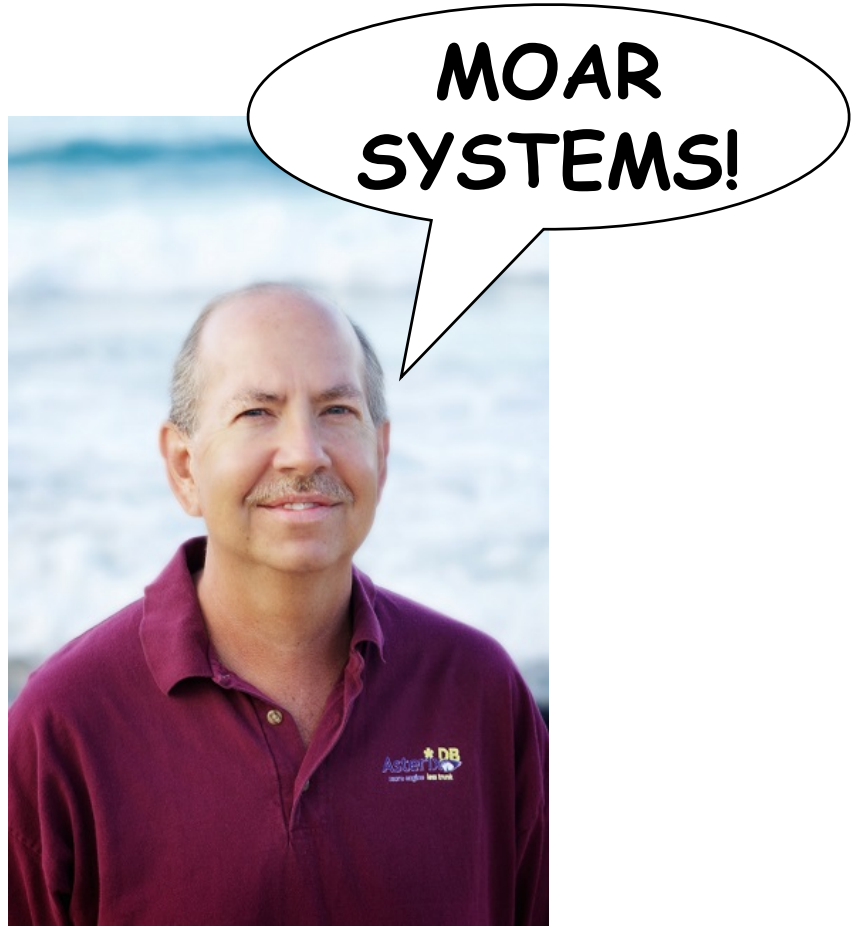

Public Voting

- 9 Submissions, 6 tutorials selected
- Google forms vote
- 236 votes, 137 individuals
- Max 46, min 11 votes

- **Did you vote?**

Presented Systems

- Apache Flink

- Apache SystemML

- HopsFS & ePipe

- LinkedIn‘s Open Source Analytics Platform

- rasdaman

- RHEEM

Massively Parallel Program

- Bulk Synchronous Parallel
- People Flow
- Danger of skew!

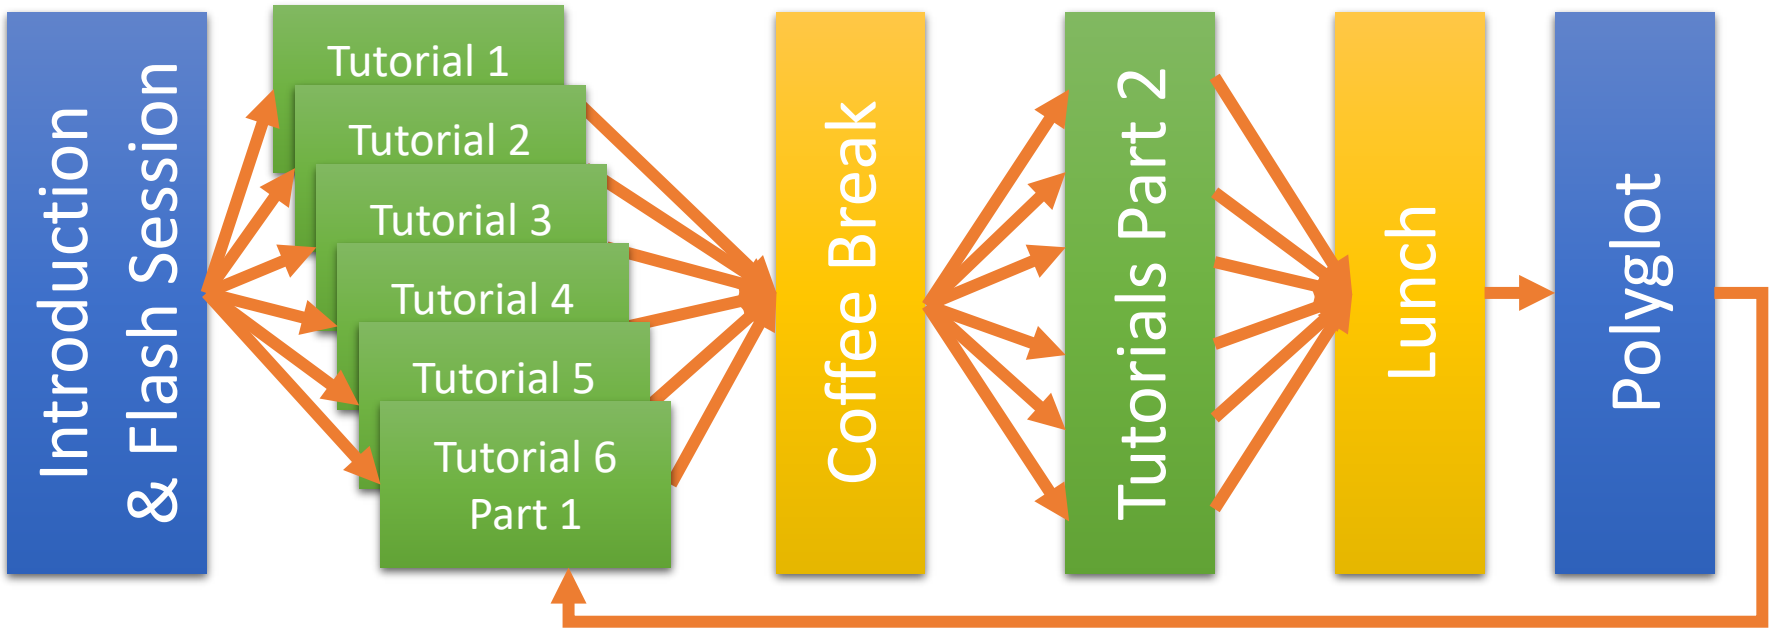

Heterogeneous Runtime Environment

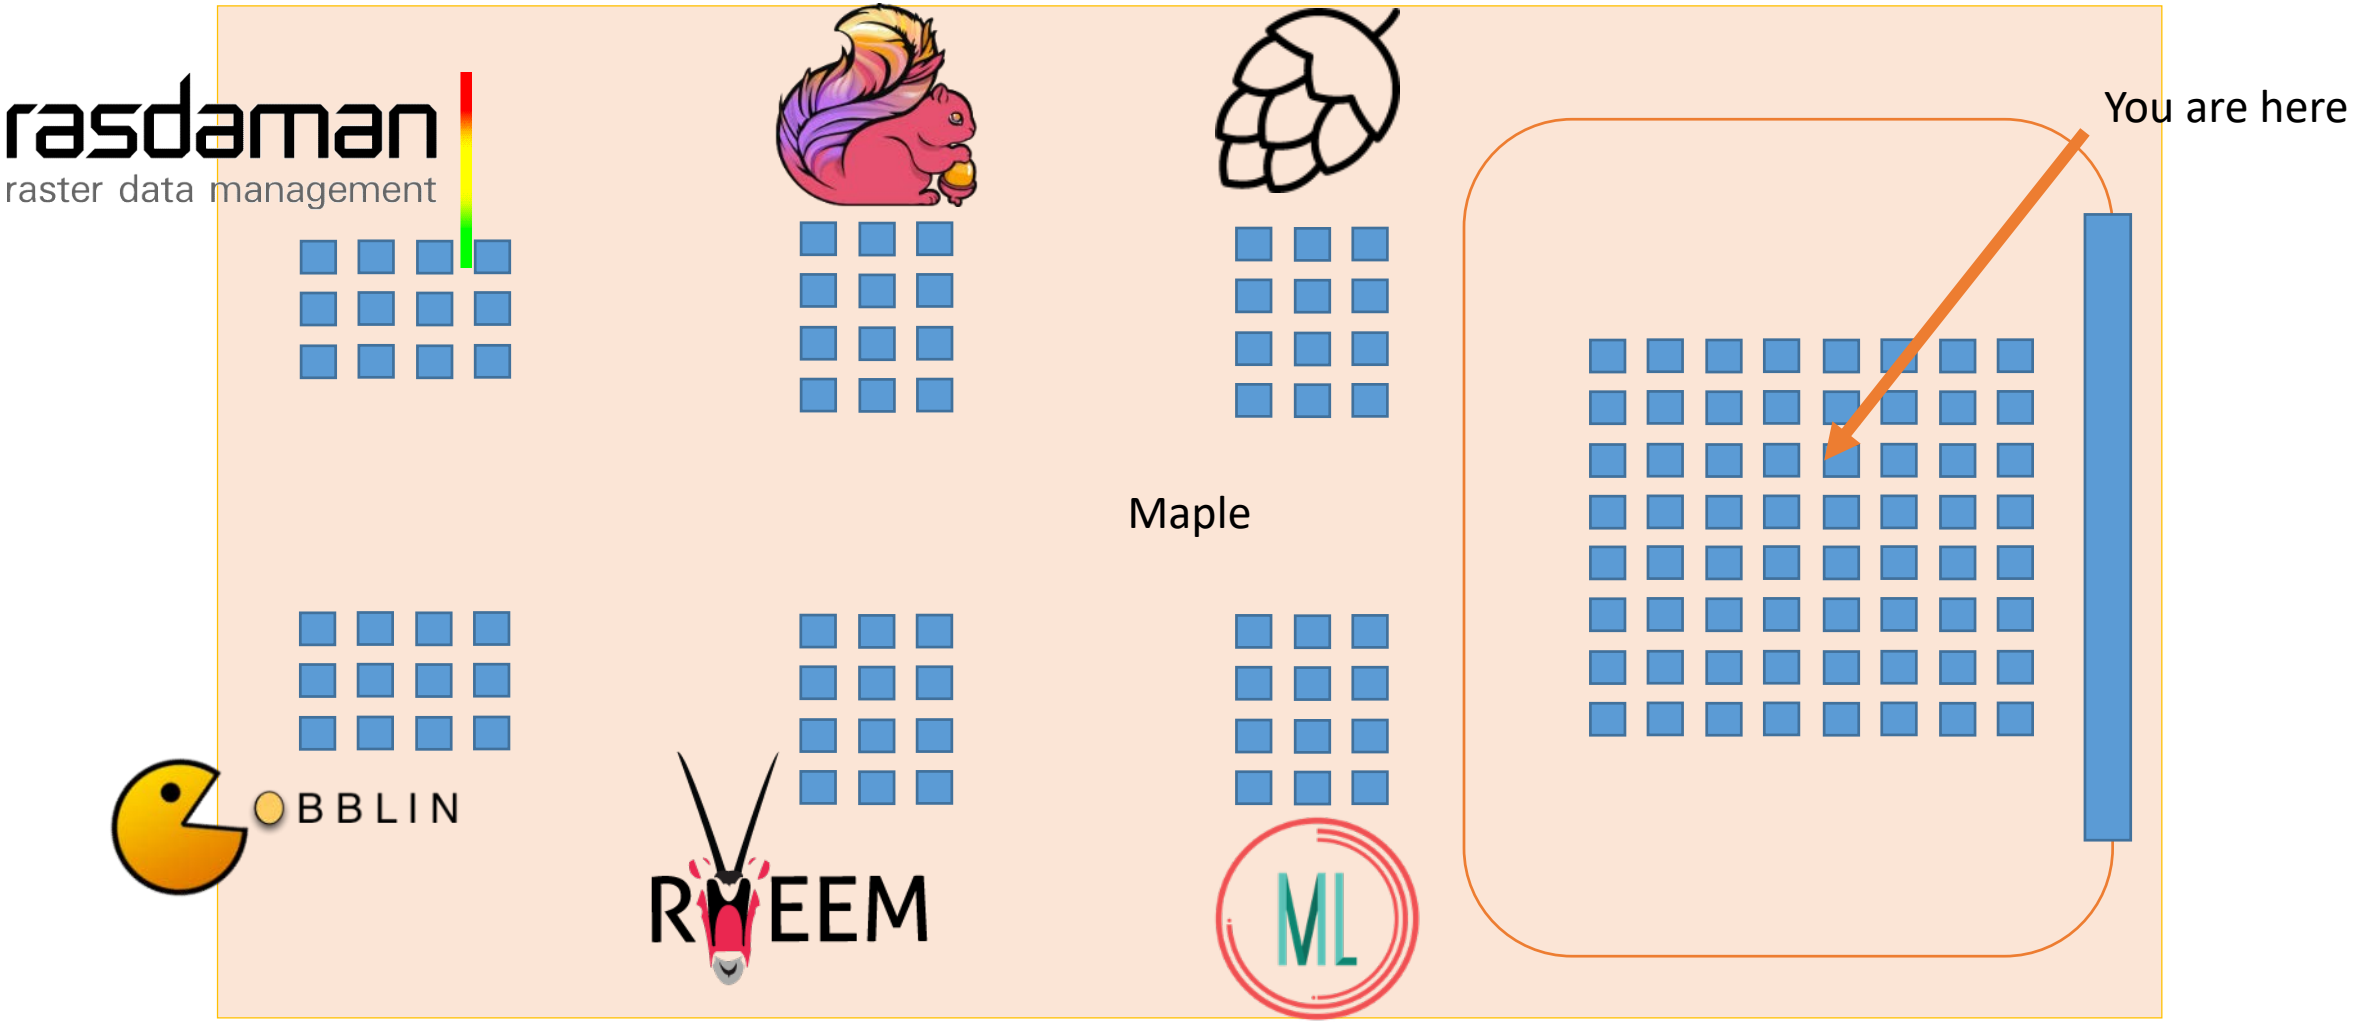

Polyglot Session

Big Data processing using Polybase. Karthik Ramachandra (Microsoft Gray Systems Lab)

Multistore Systems: Retrospection on CloudMdsQL. Jose Pereira (Univ. do Minho & INESC)

Exploiting the data center in contemporary commodity boxes: The scaling-in approach. Jignesh Patel (Univ. of Wisconsin-Madison)

LeanBigData: Blending OLTP and OLAP to Deliver Real-Time Analytical Queries. Ricardo Jimenez-Peris (LeanXcale)

Flash Intro

Apache Flink

Introduction to Stream Processing with Apache Flink®

Kostas Kloudas
Vasia Kalavri
Jonas Traub

dataArtisans

Overview

- What is Stream Processing?
- What is Apache Flink?
- Windowed computations over streams
- Handling time
- Handling node failures
- Handling planned downtime
- Handling code upgrades

A data processing engine

Apache Flink is an open source platform for distributed stream and batch processing

What does Flink provide?

- High Throughput and Low Latency
- Event-time (out-of-order) processing
- Exactly-once semantics
- Flexible windowing
- Fault-Tolerance

The Apache Flink Ecosystem

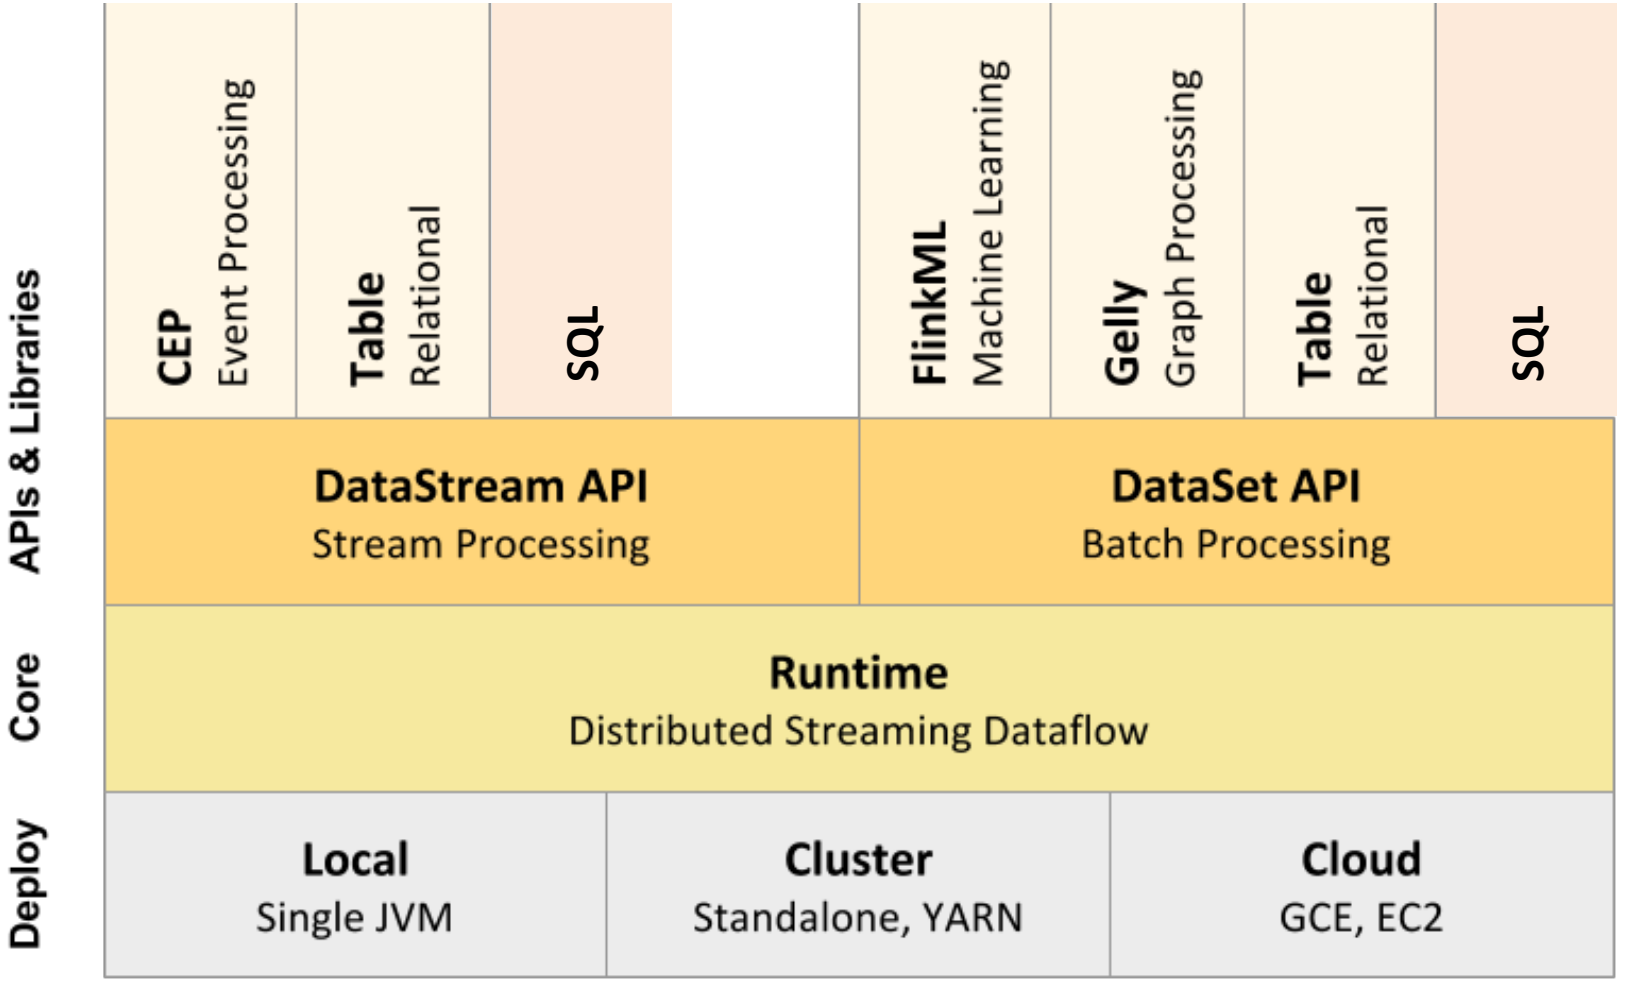

Its Users

ERICSSON

…https://flink.apache.org/poweredby.html

Time for demo…

Robust Stream Processing with Apache Flink®: A Simple Walkthrough
http://data-artisans.com/robust-stream-processing-flink-walkthrough/#more-1181

Apache SystemML

Apache SystemML: Declarative Large-Scale Machine Learning

Matthias Boehm

IBM Research – Almaden

Acknowledgements:
A. V. Evfimievski, F. Makari Manshadi, N. Pansare, B. Reinwald, F. R. Reiss, P. Sen, S. Tatikonda,

M. W. Dusenberry, D. Eriksson, N. Jindal, C. R. Kadner, J. Kim, N. Kokhlikyan, D. Kumar, M. Li, L. Resende, A. Singh, A. C. Surve, G. Weidner, and W. P. Yu

Case Study: An Automobile Manufacturer

- **Goal:** Design a model to predict car reacquisition

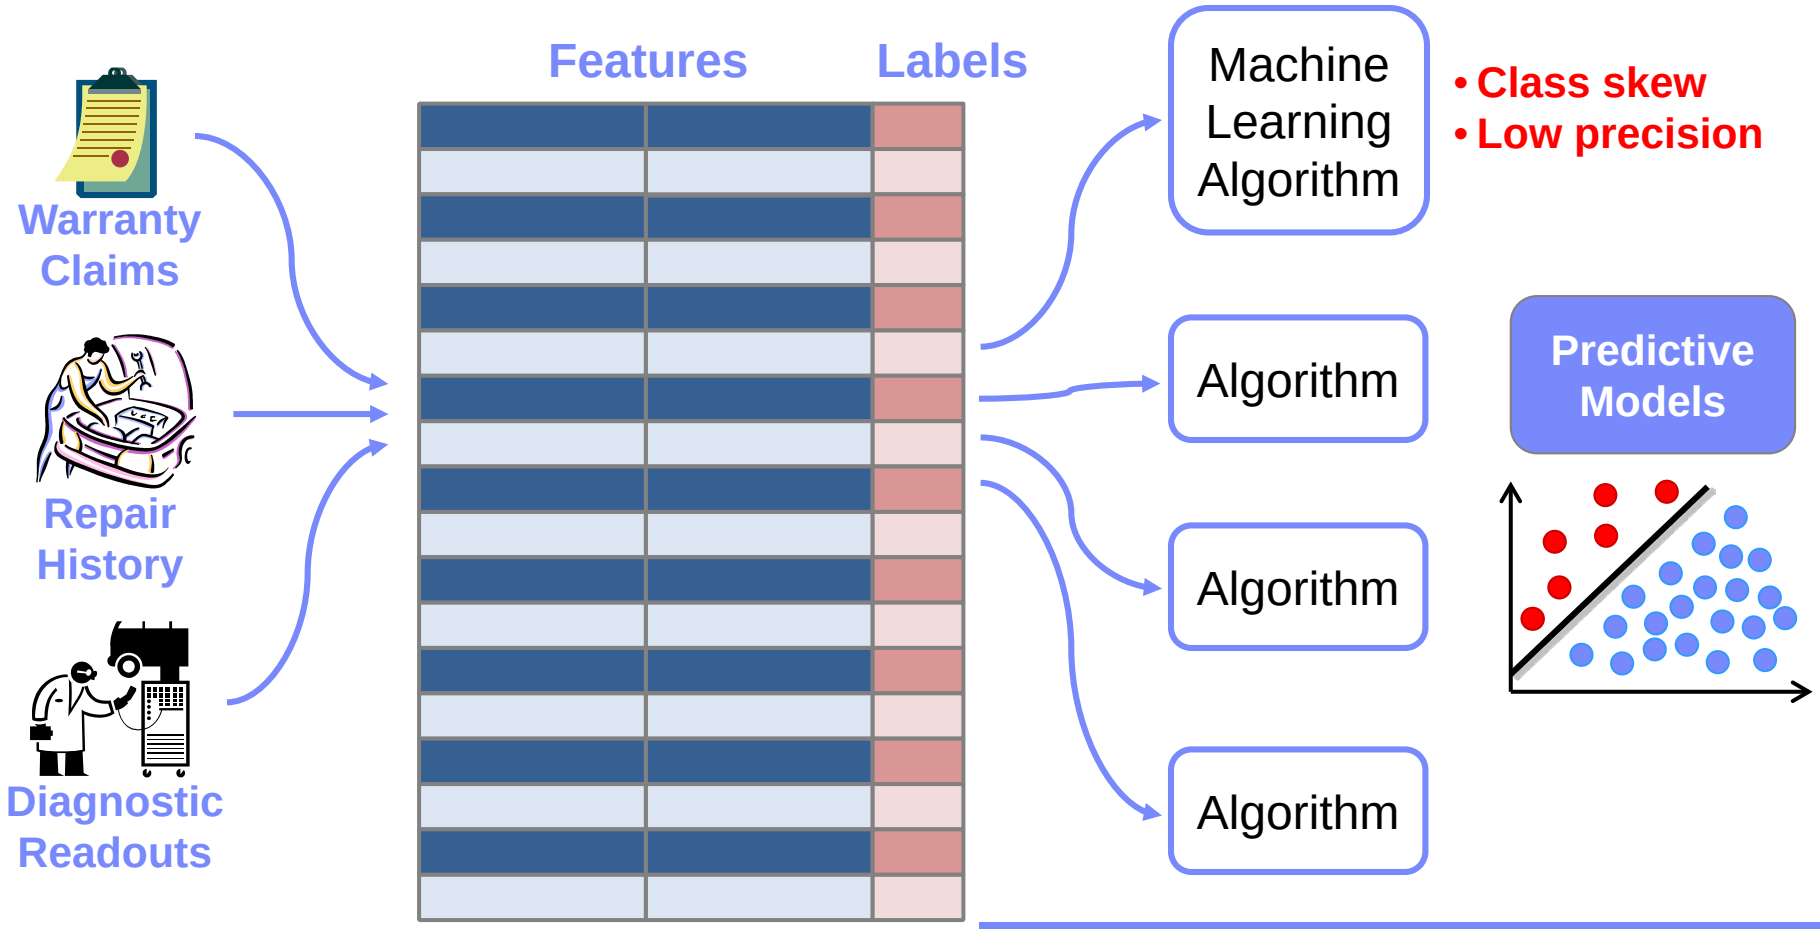

Common Patterns Across Customers

- **Algorithm customization**
- **Changes in feature set**
- **Changes in data size**
- **Quick iteration**

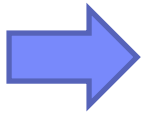

Custom Analytics

Declarative Machine Learning

Abstraction: The Good, the Bad and the Ugly

[adapted from Peter Alvaro: "I See What You Mean“, **Strange Loop, 2015**]

Platform Independence

Data Independence **Adaptivity**

Simple & Analysis-Centric **Efficiency & Performance**

```
q = t(X) %*% (w * (X %*% v))
```

The Ugly: Expectations ≠ Reality

(Missing) Rewrites

(Missing) Size Information

Complex Control Flow

Operator Selection

(Implicit) Copy-on-Write

Distributed Operations

Data Skew

Latency

Local / Remote Memory Budgets

Load Imbalance

Distributed Storage

➔ **Understanding of optimizer and runtime techniques underpinning declarative, large-scale ML**

Tutorial Outline

- **Case Study and Motivation (Flash)** 5min
- **SystemML Overview, APIs, and Tools** 30min
- **Common Framework** 15min
- **SystemML's Optimizer (w/ Hands-On-Labs)** 45min

HopsFS & ePipe

HopsFS & ePipe

Mahmoud Ismail
<maism@kth.se>
Gautier Berthou
<gautier@sics.se>

From HDFS to HopsFS

- Scale to a million operations/sec
- Scale to billions of files/directories
- Search with sub second latency (ePipe)

Tutorial

- Introducing Github style for Hadoop projects (HopsWorks)
- Installation of Hops on AWS using Karamel
- Managing Datasets
  - create, attach metadata, and search
- Running sample programs on HopsWorks

Goblin & Pinot

Open Source Analytics Pipeline at LinkedIn

Issac Buenrostro
Jean François Im

BOSS Workshop, 2016

LinkedIn

Large Scale Analytics

1. Analyze many TB data daily.
2. Multiple, heterogeneous sources, with varying data quality.
3. Fast querying for offline and real-time needs.
4. Integrate with other data processing jobs (MR, Hive, Spark, etc.).
5. Fault tolerance, scalability, manageability, …

Solution: Gobblin + Pinot

- Universal data ingestion framework.
- Extract, transform, quality check, and write data from/to a large variety of data storage technologies: HDFS, S3, Kafka, JDBC, Rest, …

- Distributed near-realtime OLAP data store.
- Index and combine data from offline data sources (e.g. Hadoop) and real time data sources (e.g. Kafka).
- SQL query interface.

In This Workshop

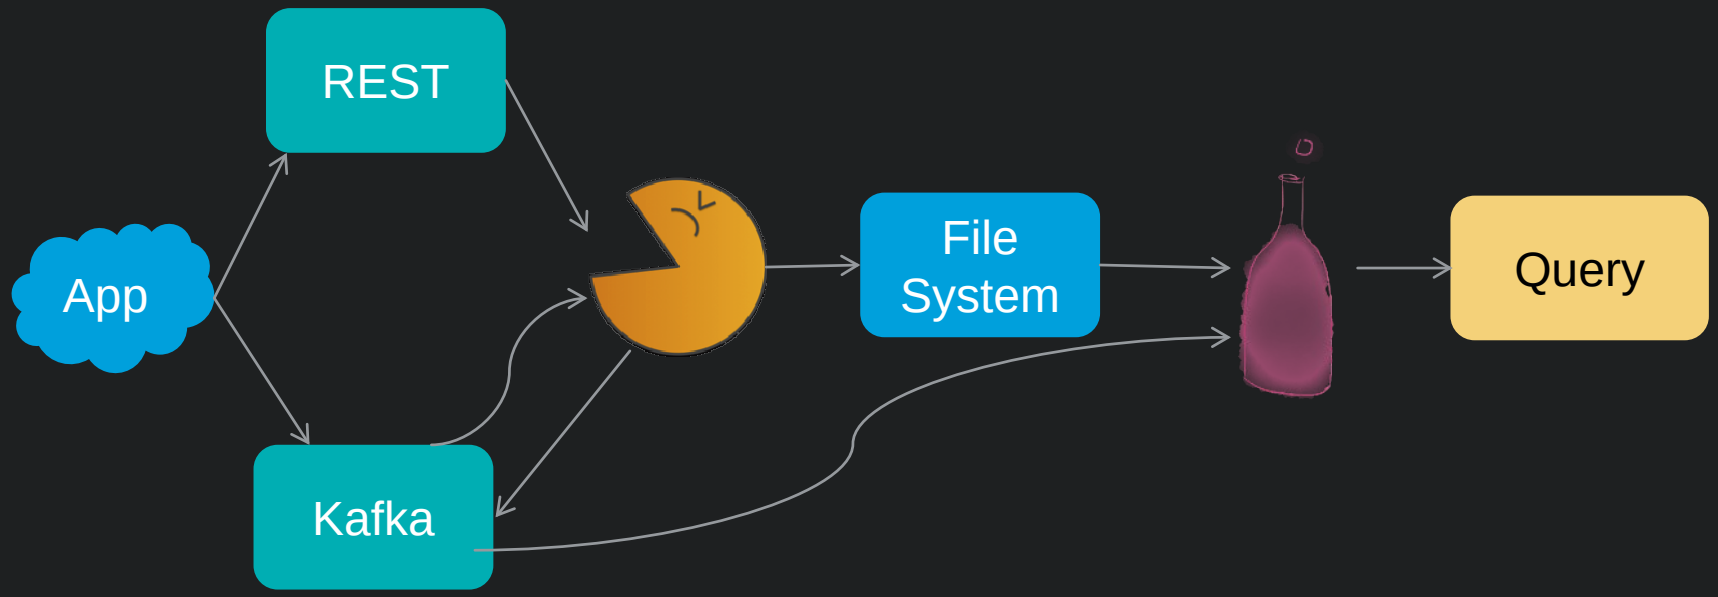

Find out more:

https://github.com/linkedin/gobblin
http://gobblin.readthedocs.io/
gobblin-users@googlegroups.com

https://github.com/linkedin/pinot
pinot-users@googlegroups.com

https://engineering.linkedin.com/

rasdaman

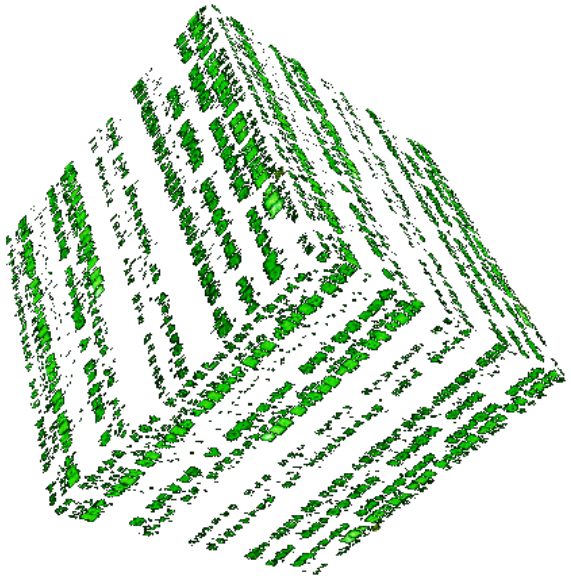

rasdaman @ BOSS'16

New Delhi, India, 09-sep-2016

Dimitar Mišev <misev@rasdaman.com>
Jacobs University | rasdaman GmbH

[gamingfeeds.com]

Research supported by EU through EarthServer

Array Analytics Research @ Jacobs U

- Large-Scale Scientific Information Systems research group
  - Flexible, scalable n-D array services
  - www.jacobs-university.de/lsis
- Most visible results:
  - Pioneer Array DBMS, rasdaman
  - Standardization: OGC Big Geo Data, ISO SQL

rasdaman: Agile Array Analytics

- „raster data manager“: n-D arrays in SQL
  - [VLDB 1994, VLDB 1997, SIGMOD 1998, VLDB 2003, …]
- Array Algebra [NGITS 1998]
  - SQL/MDA [SSDBM 2014, DOLAP 2015]
- Scalable, parallel „tile streaming“ architecture
- 130+ TB installations in operational use

CIO 100 MOST PROMISING BIG DATA Review SOLUTION PROVIDERS 2016

The Earth Monitoring Competition

Copernicus masters

WINNER

T-SYSTEMS BIG DATA CHALLENGE

2014

T · · Systems ·

Tutorial outline

- Installation & deployment
  - RPM/DEB, VM download, build from source
- Data modelling and concepts
  - What kind of data is supported?
- Query language
  - Typical array analytics queries, hands on
- Storage management
  - Single array datacubes can reach hundreds of TB
  - Learn how rasdaman scales to such volumes
- Domain application: Geo services

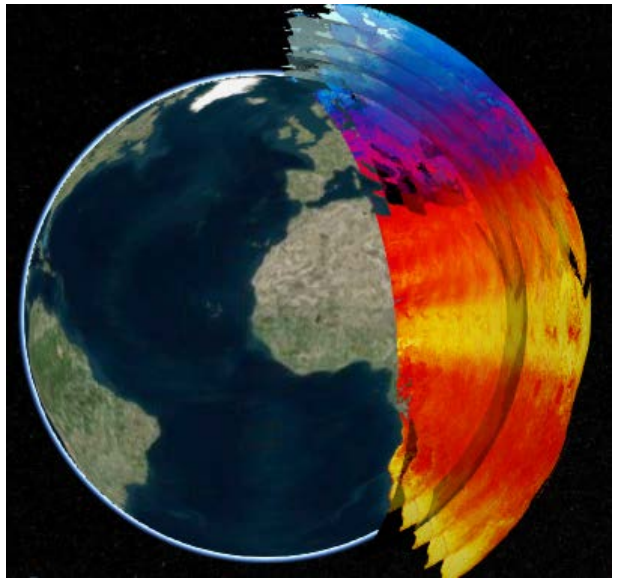

RHEEM

Zoi Kaoudi
Sebastian Kruse
Jorge Quiané

Turning a Zoo into a Circus

RHEEM

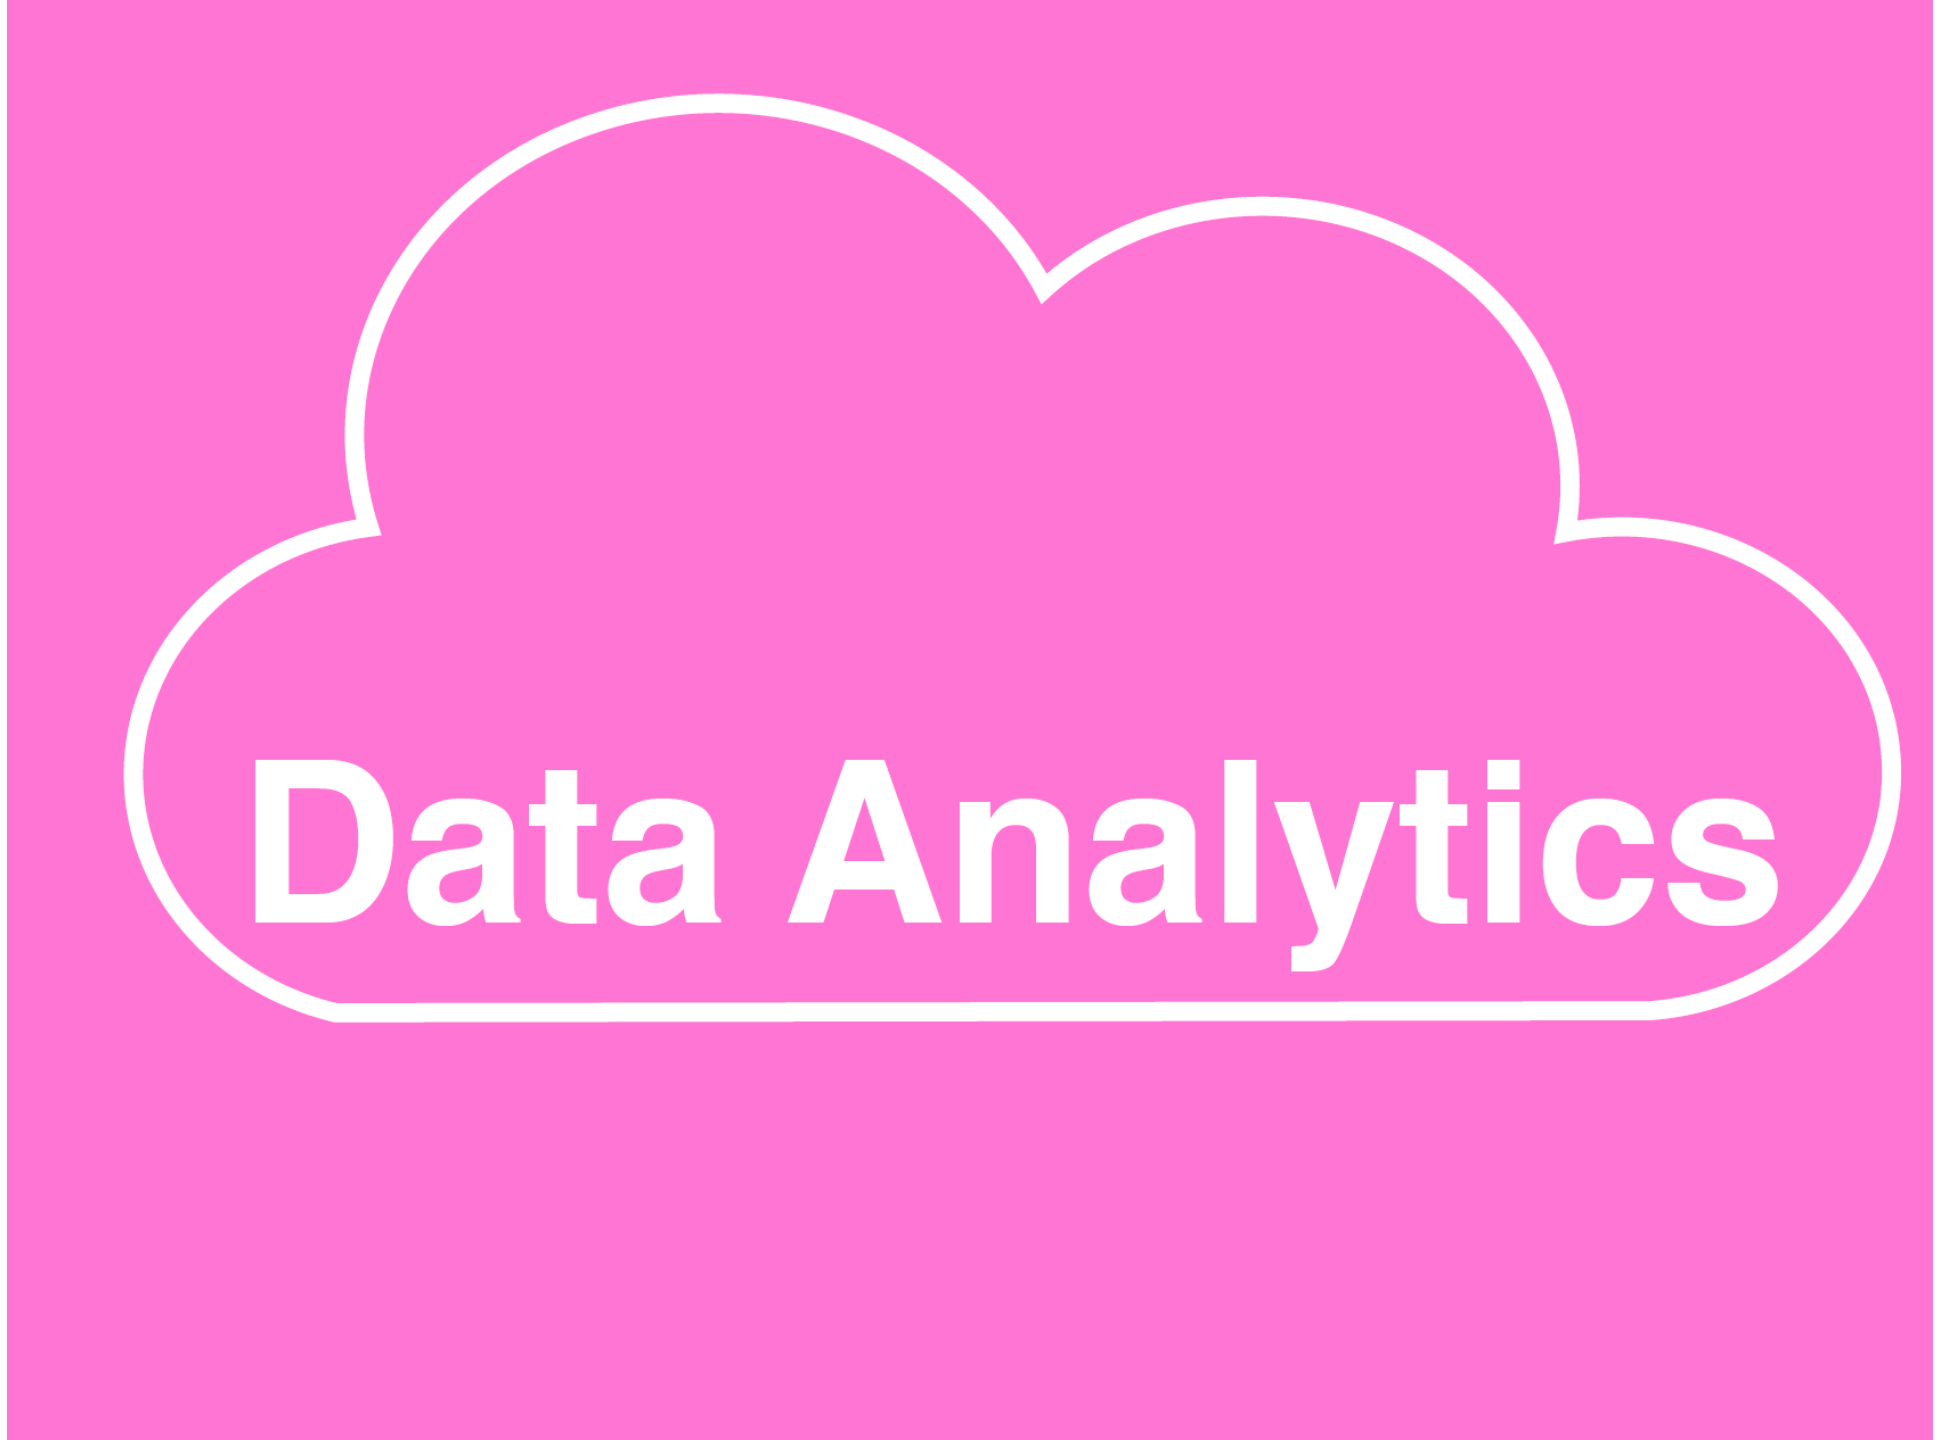
Data Analytics

System-X

Data Analytics


cloud
25
50
15
10

Data Scientist

- sexy job -

Big Data Landscape 2016

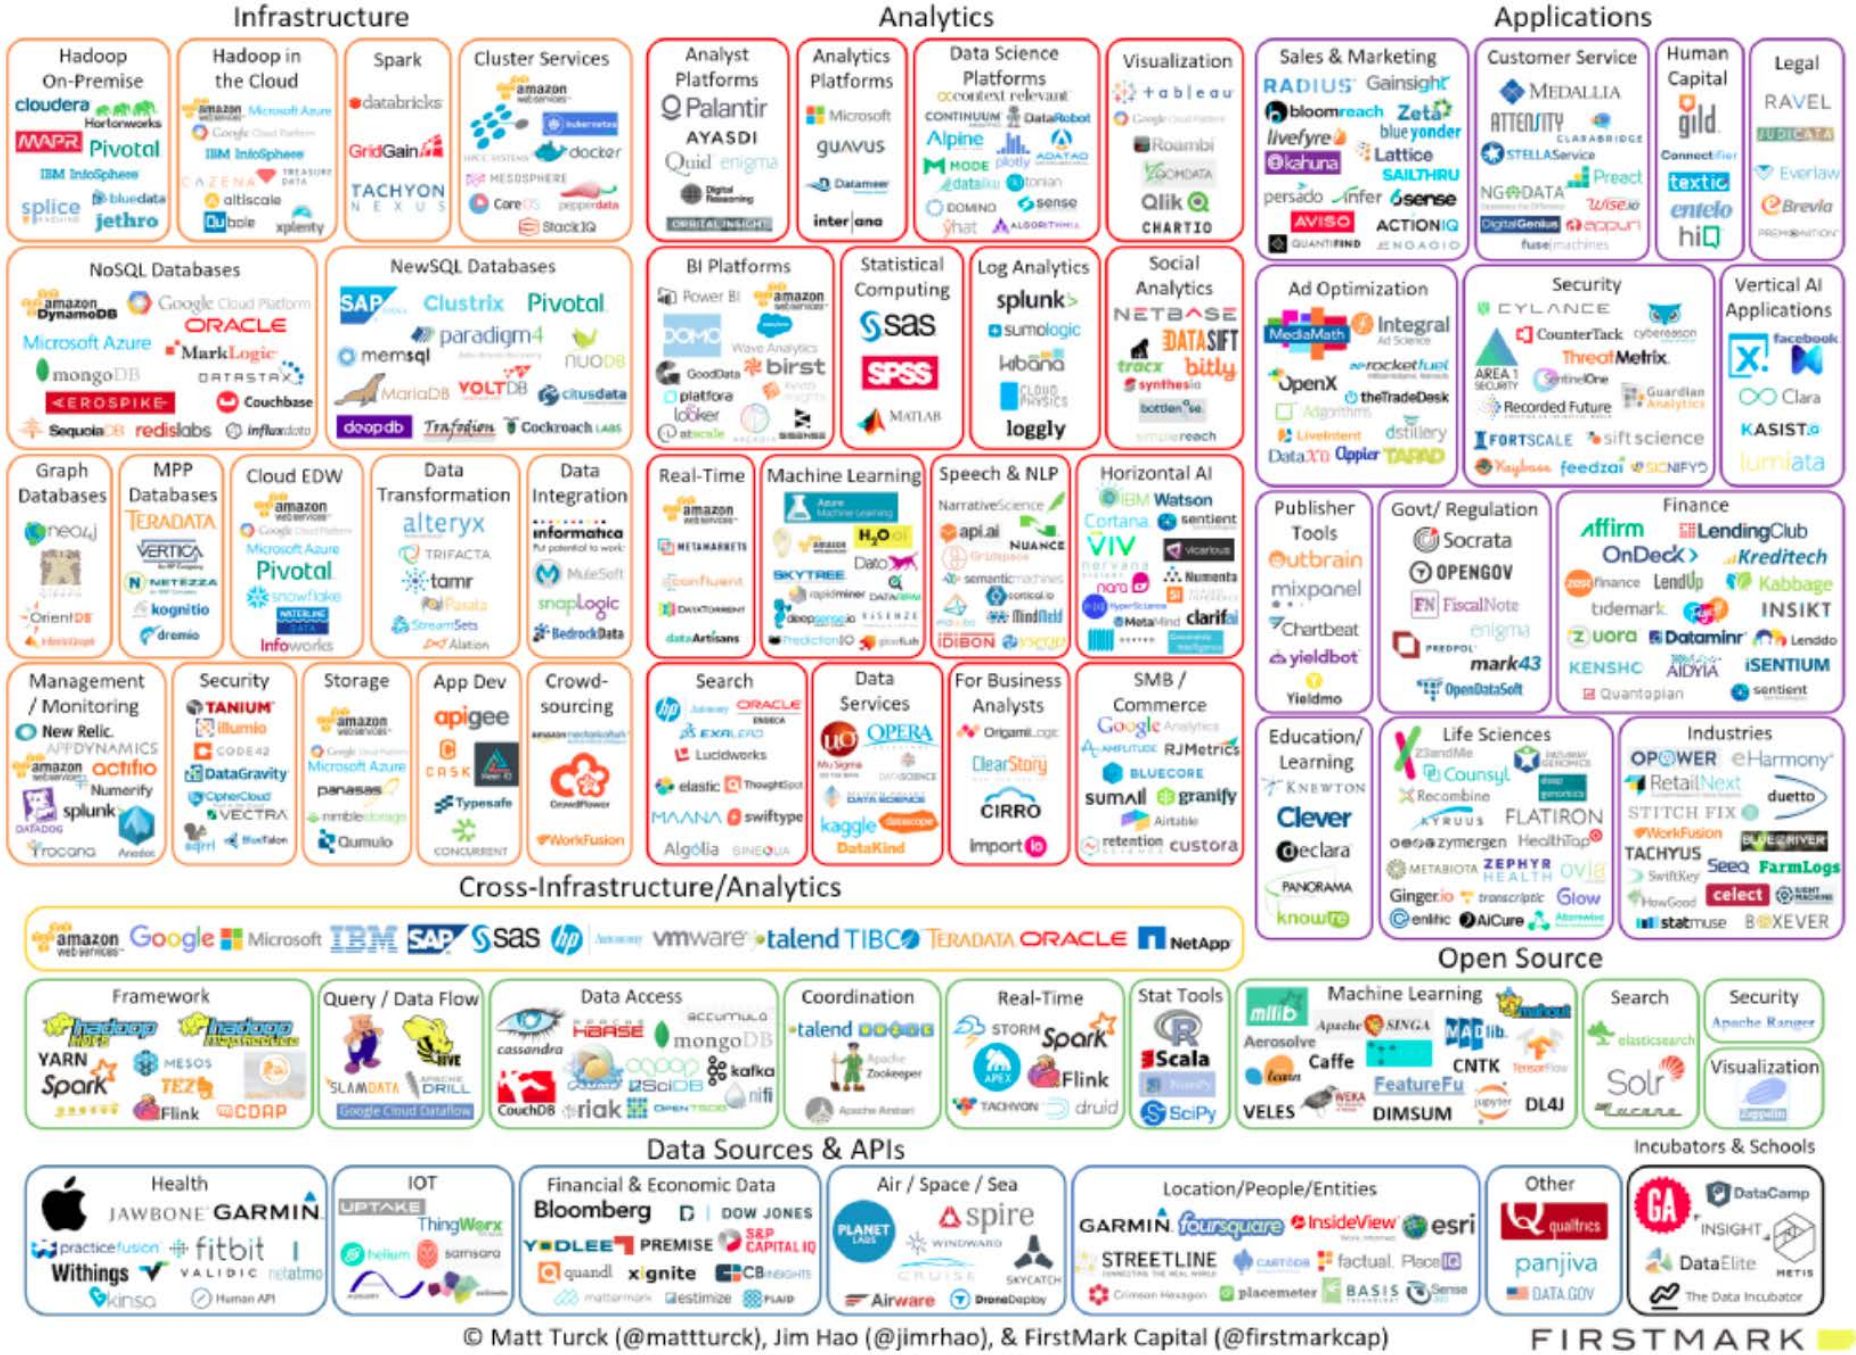

FIRSTMARK

Anything but sexy!

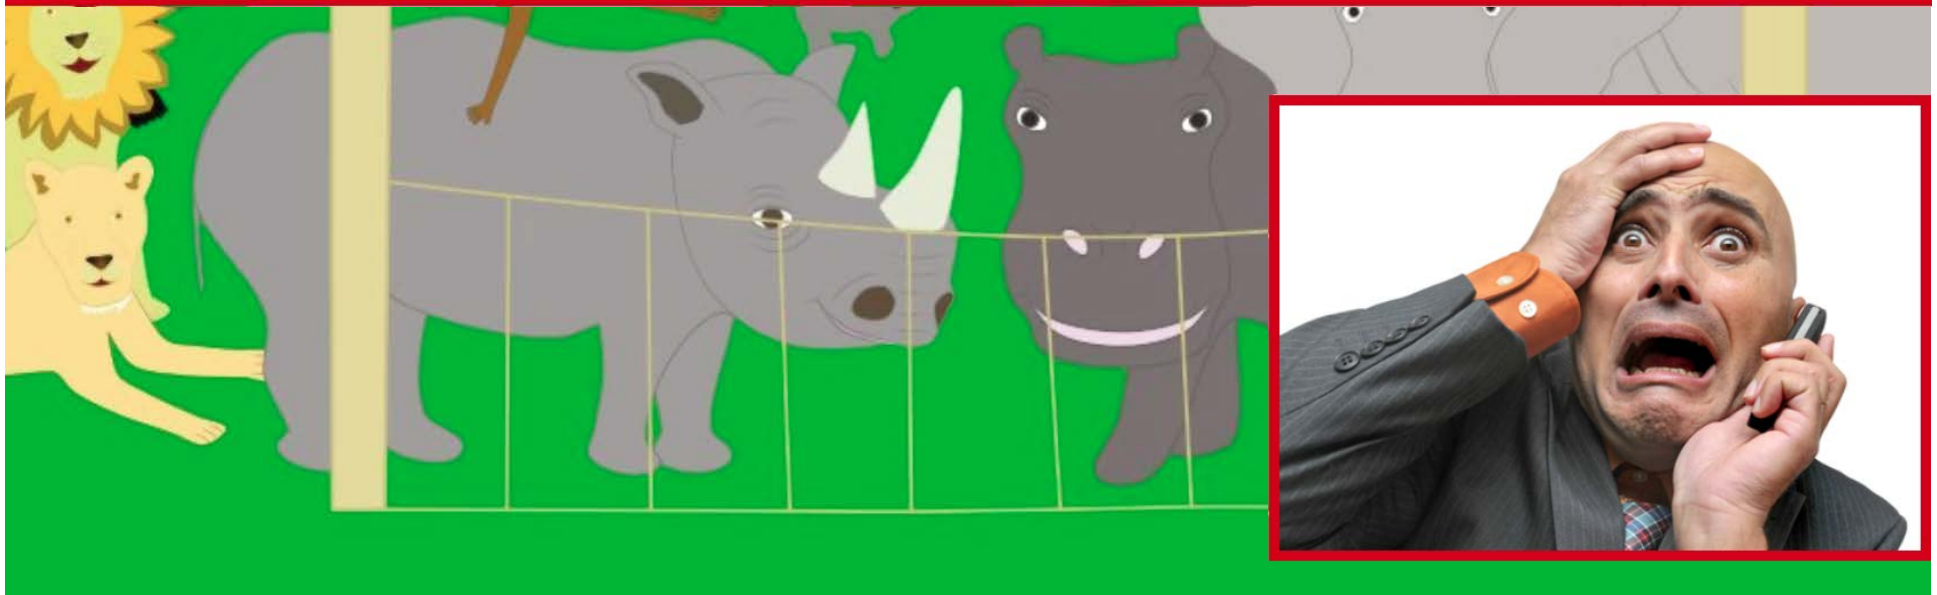


CIRCUS
CIRCUS
HOT
FRESH
POPCORN
PLEASE DO NOT FEED

RHEEM

A Cross-Platform Data Processing Framework

Let‘s go!

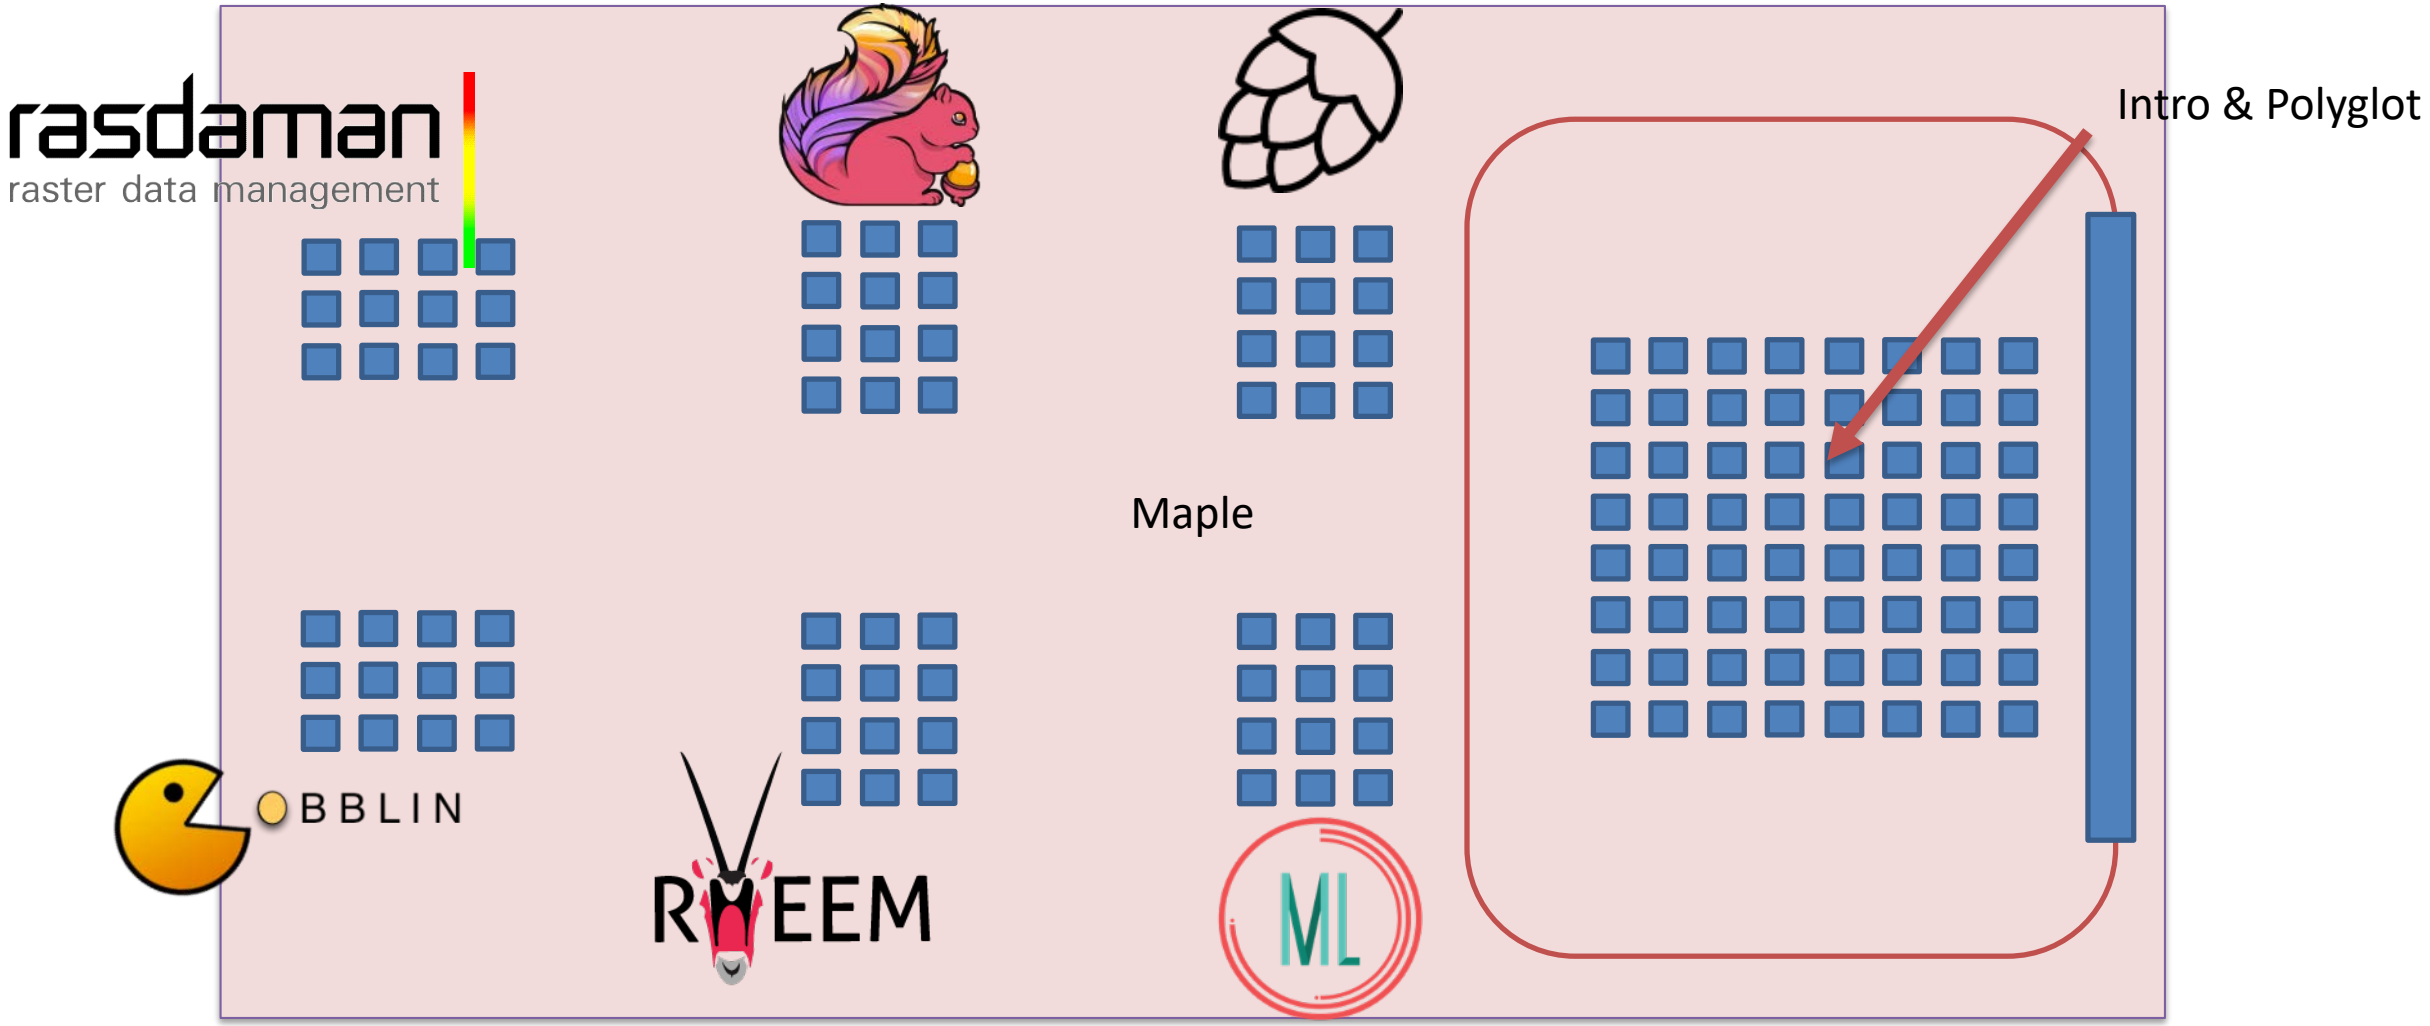
rasdaman
raster data management
Intro & Polyglot
Maple
OBBLIN
RHEEM
ML