

REAL-TIME ANALYTICS WITH APACHE FLINK AND DRUID

Jan Graßegger - @gesundkrank

Berlin Buzzwords 2016



DATA ENGINEER @

mbr targeting

70,000 EVENTS
PER
SECOND

50 DIMENSIONS
20 METRICS

OUR DATA



DRUID

- ▶ Online Analytical Processing (OLAP) System
- ▶ Column-oriented
- ▶ Distributed
- ▶ Built-in data sharding based on time windows
- ▶ JSON query language

DRUID

TOP PRIVATE DOMAIN

battle.net

battle.net

noxxic.com

noxxic.com

battle.net → 5

noxxic.com → 6

[5, 5, 6, 6]

Column

Strings to Integers

Encoded column data

DATA STRUCTURES

TOP PRIVATE DOMAIN

battle.net

battle.net

noxxic.com

noxxic.com

Column

battle.net → [1, 1, 0, 0]

noxxic.com → [0, 0, 1, 1]

Bitmap Indices


DATA STRUCTURES

FIREHOSES



APACHE FLINK





PROCESSING




TRANQUILITY



TRANQUILITY

- ▶ Helps ingesting real-time data into Druid
- ▶ Provides adapters for Samza, Spark, Storm and Flink
- ▶ Standalone HTTP and Kafka applications

TRANQUILITY



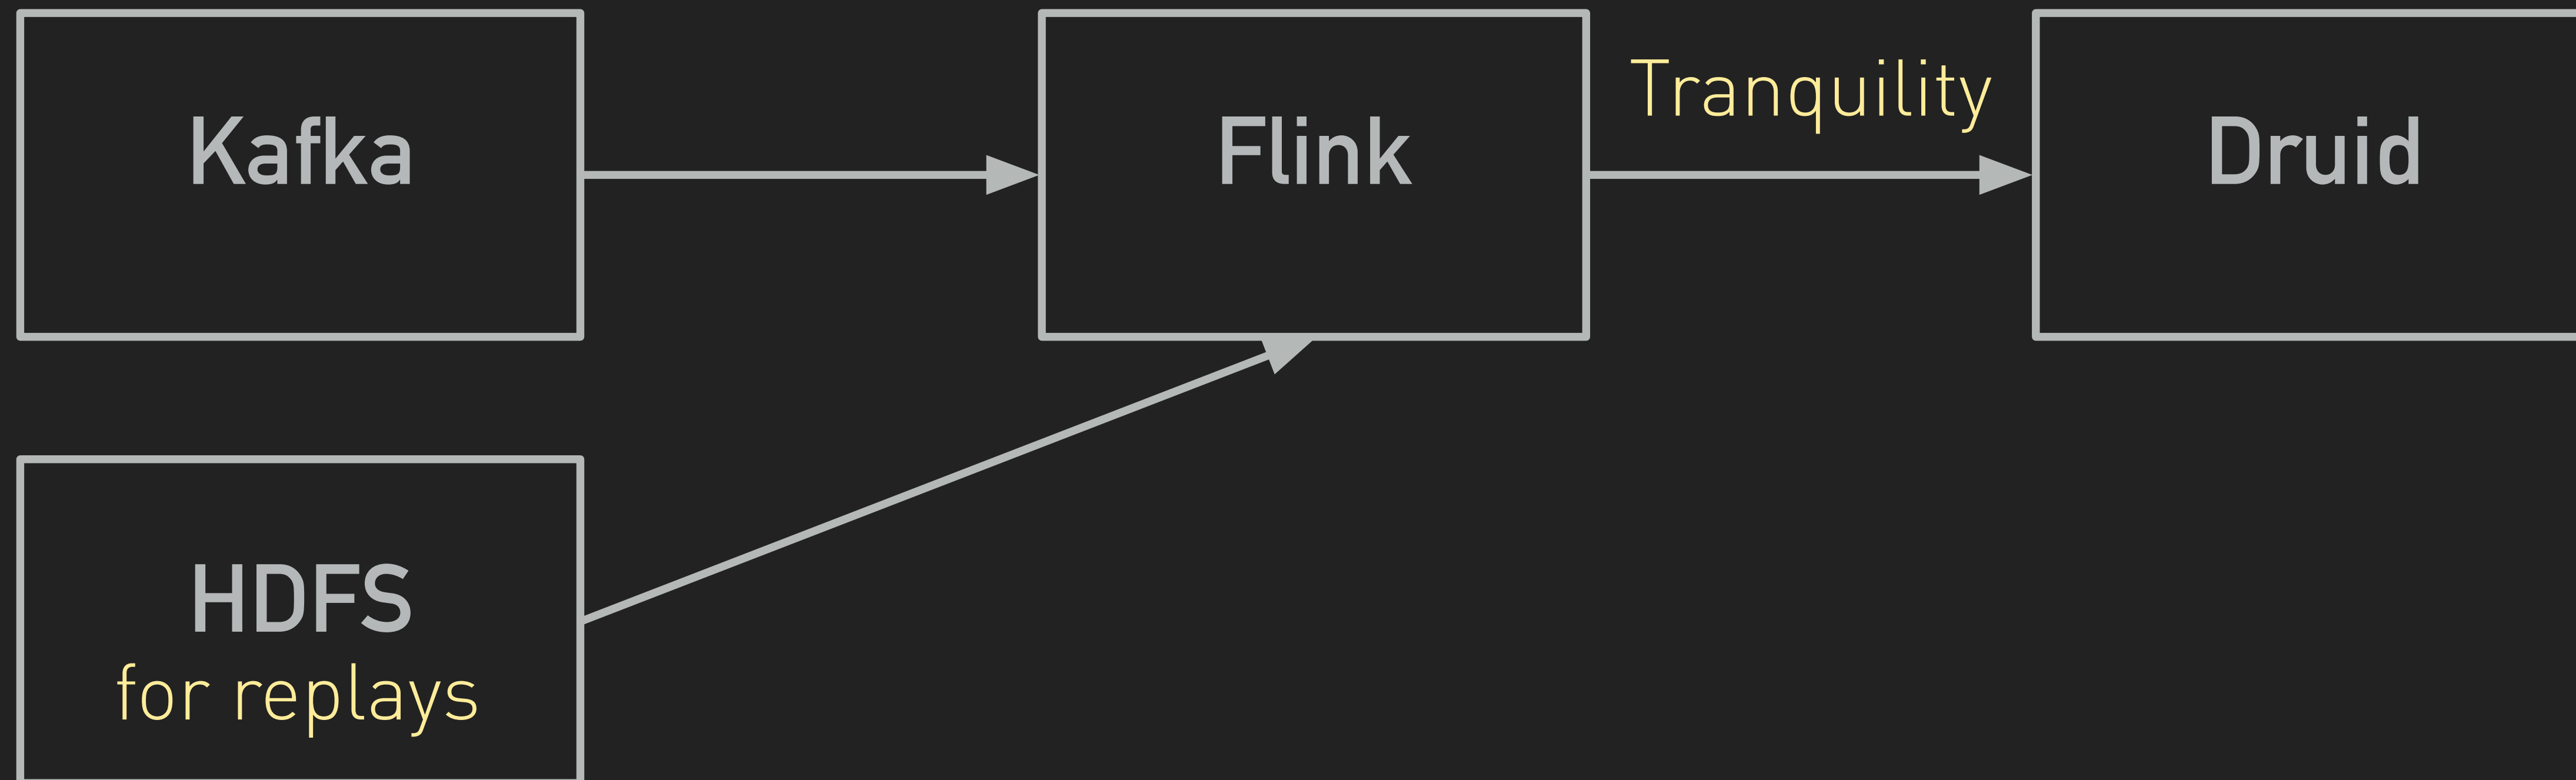
PROCESSING

LAMBDA





KAPPA



PROCESSING

- ▶ Kappa-like architecture that's able to do replays from HDFS & Kafka
 - ▶ Added Flink sink to Tranquility
 - ▶ "Hacked" replays into Tranquility
 - ▶ Real-Time Reporting
-

RESULTS



QUESTIONS?