




Succeeding with Apache Druid[®] and Clickstream Data

Peter Marshall




What is Apache Druid®? Who are Imply?



Apache Druid is a **high performance, real-time analytics** database ... where **fast** ad-hoc analytics, **instant** data visibility, or supporting high **concurrency** is important ... where an **interactive**, consistent user **experience** is desired.



We will devote our energy to making it as easy as possible for people to **use Druid** and **build awesome data applications** on top of it.






[Subscribe](#)
[Newsletter](#)
[Slack](#)
[Premium](#)

SOFTWARE ENGINEERING DAILY
The World Through the Lens of Software


October 1, 2021

Search...

[All Content](#) [Podcast](#) [Articles](#) [Cloud Engineering](#) [Business and Philosophy](#) [Greatest Hits](#) [Hackers](#) [Data](#) [Open Source](#) [Blockchain](#)



Druid: Event-Driven Data with Eric Tschetter


 By **SE Daily**
Podcast | Monday, August 16 2021



Podcast: [Play in new window](#) | [Download](#)

Subscribe: [Apple Podcasts](#) | [RSS](#)

Whether sending messages, shopping in an app, or watching videos, modern consumers expect information and responsiveness to be near-instant in their apps and devices. From a developer's perspective, this means clean code and a fast database.



POPULAR

-
-
-
-
-

Software Engineering Daily: an interview with Eric Tschetter

<https://softwareengineeringdaily.com/2021/08/16/druid-event-driven-data-with-eric-tschetter/>

Druid

A Real-time Analytical Data Store

Fangjin Yang
Metamarkets Group, Inc.
fangjin@metamarkets.com

Eric Tschetter
etcheddar@gmail.com

Nelson Ray
ncray86@gmail.com

Gian Merlino
Metamarkets Group, Inc.
gian@metamarkets.co

Xavier Léauté
Metamarkets Group, Inc.
xavier@metamarkets.com

Deep Ganguli
Metamarkets Group, Inc.
deep@metamarkets.com

ABSTRACT

ABSTRACT Druid is an open source¹ data store designed for real-time exploratory analytics on large data sets. The system combines a column-oriented storage layout, a distributed, shared-nothing architecture, and an advanced indexing structure to allow for the arbitrary exploration of billion-row tables with sub-second latencies. In this paper, we describe Druid's architecture, and detail how it supports fast aggregation, filtering, and low latency data ingestion.

Categories and Subject Descriptors

3.2.4 (Database Management): Systems—Distributed databases:

Keywords

Keywords
distributed; real-time; fault-tolerant; highly available; open source
mediation; column-oriented; OLAP

1. INTRODUCTION

[illegible]

³<http://druid.io/> <https://github.com/metamx/druid>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

SIGMOD'14, June 22–27, 2014, Snowbird, UT, USA.

Copyright is held by the owner(s) of the publication. Publication rights licensed to ACM.

ACM 978-1-4503-2776-5/14/06...\$15.00.

<http://doi.org/10.1145/2588555.2595632>

event streams into high-value aggregates for a variety of applications such as business intelligence and A-B testing.

As with many great systems, Hadoop has opened our eyes to a new space of problems. Specifically, Hadoop excels at storing and providing access to large amounts of data, however, it does not make any performance guarantees around how quickly that data can be accessed. Furthermore, although Hadoop is a highly available system, performance degrades under heavy concurrent load. Lastly, while Hadoop works well for storing data, it is not optimized for in-place access to that data immediately readable.

Early on in the development of the Metastore data product, we ran into each of these issues and came to the realization that Hadoop is a great back-office, batch processing, and data warehousing system. However, as a company that has product-level guarantees around query performance and data availability in a highly consistent environment (1000+ users), Hadoop wasn't going to meet our needs. We explored different solutions in the space, and after trying out a number of different products, we decided to build our own. As the Relational Database Management Systems and NoSQL architectures, we came to the conclusion that there was no single solution to our problem. We needed a system that could be fully leveraged for our online open source world that could be fully leveraged for our offline open source world. We ended up creating *Druid*, an open source, distributed column-oriented, real-time analytical data store. In many ways, *Druid* shares similarities with other OLAP systems [30, 35, 22]. *Druid* shares similarities with [28], main-memory databases [14], as well as interactive query engines [28], distributed data stores [7, 12, 23]. The distributed nature of *Druid* also borrows ideas from current generation search engines and query engines [25, 4, 3].

This paper describes the architecture of Druid, explores the various design decisions made in creating an always-on production system that powers a hosted service, and attempts to help other system owners face a similar problem about a potential technology of solution. Druid is deployed in production at several technology companies.¹ The structure of the paper is as follows: we first describe the problem in Section 2. Next, we detail system architecture from the point of view of how data flows through the system in Section 3. We then discuss how and why data gets converted into a binary format in Section 4. We briefly describe the query API in Section 5 and present performance results in Section 6. Lastly, we leave off with lessons from running Druid in production in Section 7, and future work in Section 8.

2. PROBLEM DEFINITION

Druid was originally designed to solve problems around ingesting and exploring large quantities of transactional events (log data, etc.). Some of the similarities data is commonly found in OLAP work

This form of timeseries is

NETFLIX



NETFLIX



YAHOO!





NETFLIX

CONDÉ NAST



dripstat



YAHOO!

hulu



criteo.





Fully scalable

Batch and real-time data

Ad-hoc statistical queries

Low latency delivery

log search

real-time ingest
flexible schema
text search



Fully scalable

Batch and real-time data

Ad-hoc statistical queries

Low latency delivery


log search

real-time ingest
flexible schema
text search



timeseries

low latency ingest
time-based storage
time functions



druid

Fully scalable

Batch and real-time data

Ad-hoc statistical queries

Low latency delivery


log search

real-time ingest
flexible schema
text search



timeseries

low latency ingest
time-based storage
time functions



druid

columnar

efficient storage
fast analytic queries
data distribution



Fully scalable

Batch and real-time data

Ad-hoc statistical queries

Low latency delivery


log search

real-time ingest
flexible schema
text search



timeseries

low latency ingest
time-based storage
time functions




columnar

efficient storage
fast analytic queries
data distribution



druid



High Performance
Real-time Analytics



Why Clickstream Analytics?

“It puts us closer to our users
and if you know what your users want,
you’re better able to serve them.”

Welcome to the BBC



LIVE Lifting curbs on 21 June risks U-turn - scientist

UK



House prices soar 10.9% in 'race for space'

Business




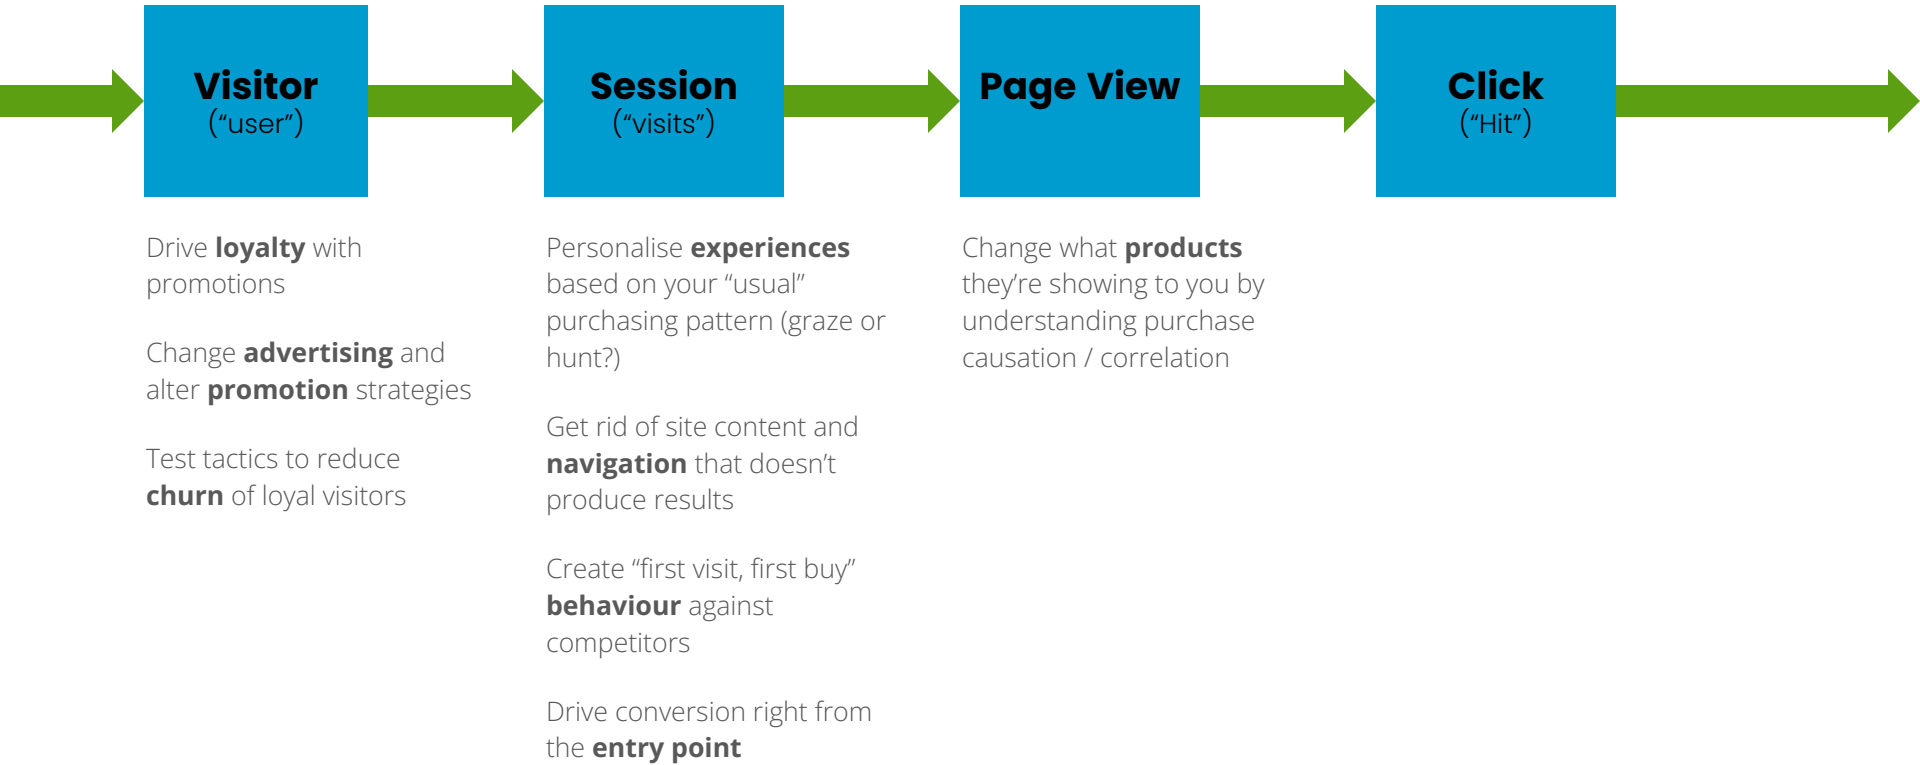
Greenwood pulls out of England squad

Football

More top stories

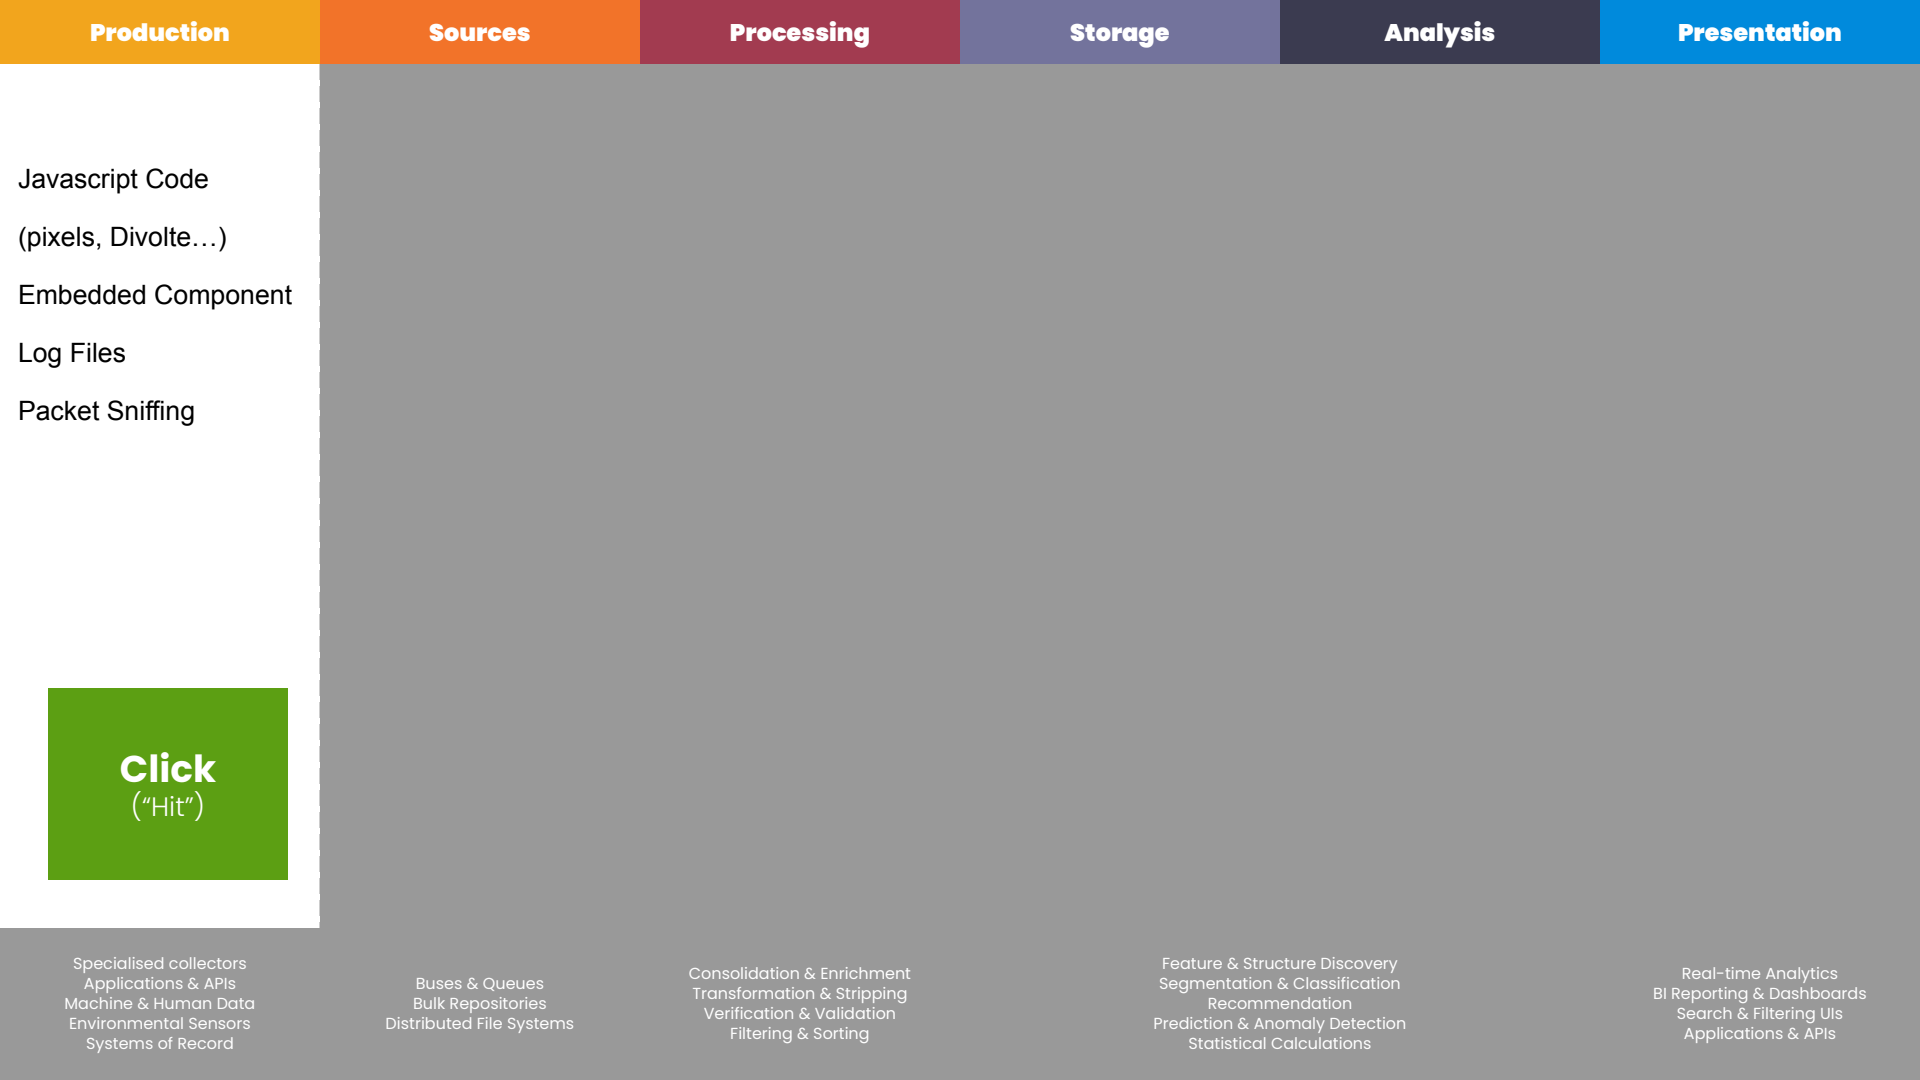








Clickstream Pipelines



Production

Sources

Processing

Storage

Analysis

Presentation

Javascript Code

(pixels, Divolte...)

Embedded Component

Log Files

Packet Sniffing



Click
("Hit")


Specialised collectors
Applications & APIs
Machine & Human Data
Environmental Sensors
Systems of Record

Buses & Queues
Bulk Repositories
Distributed File Systems

Consolidation & Enrichment
Transformation & Stripping
Verification & Validation
Filtering & Sorting

Feature & Structure Discovery
Segmentation & Classification
Recommendation
Prediction & Anomaly Detection
Statistical Calculations

Real-time Analytics
BI Reporting & Dashboards
Search & Filtering UIs
Applications & APIs

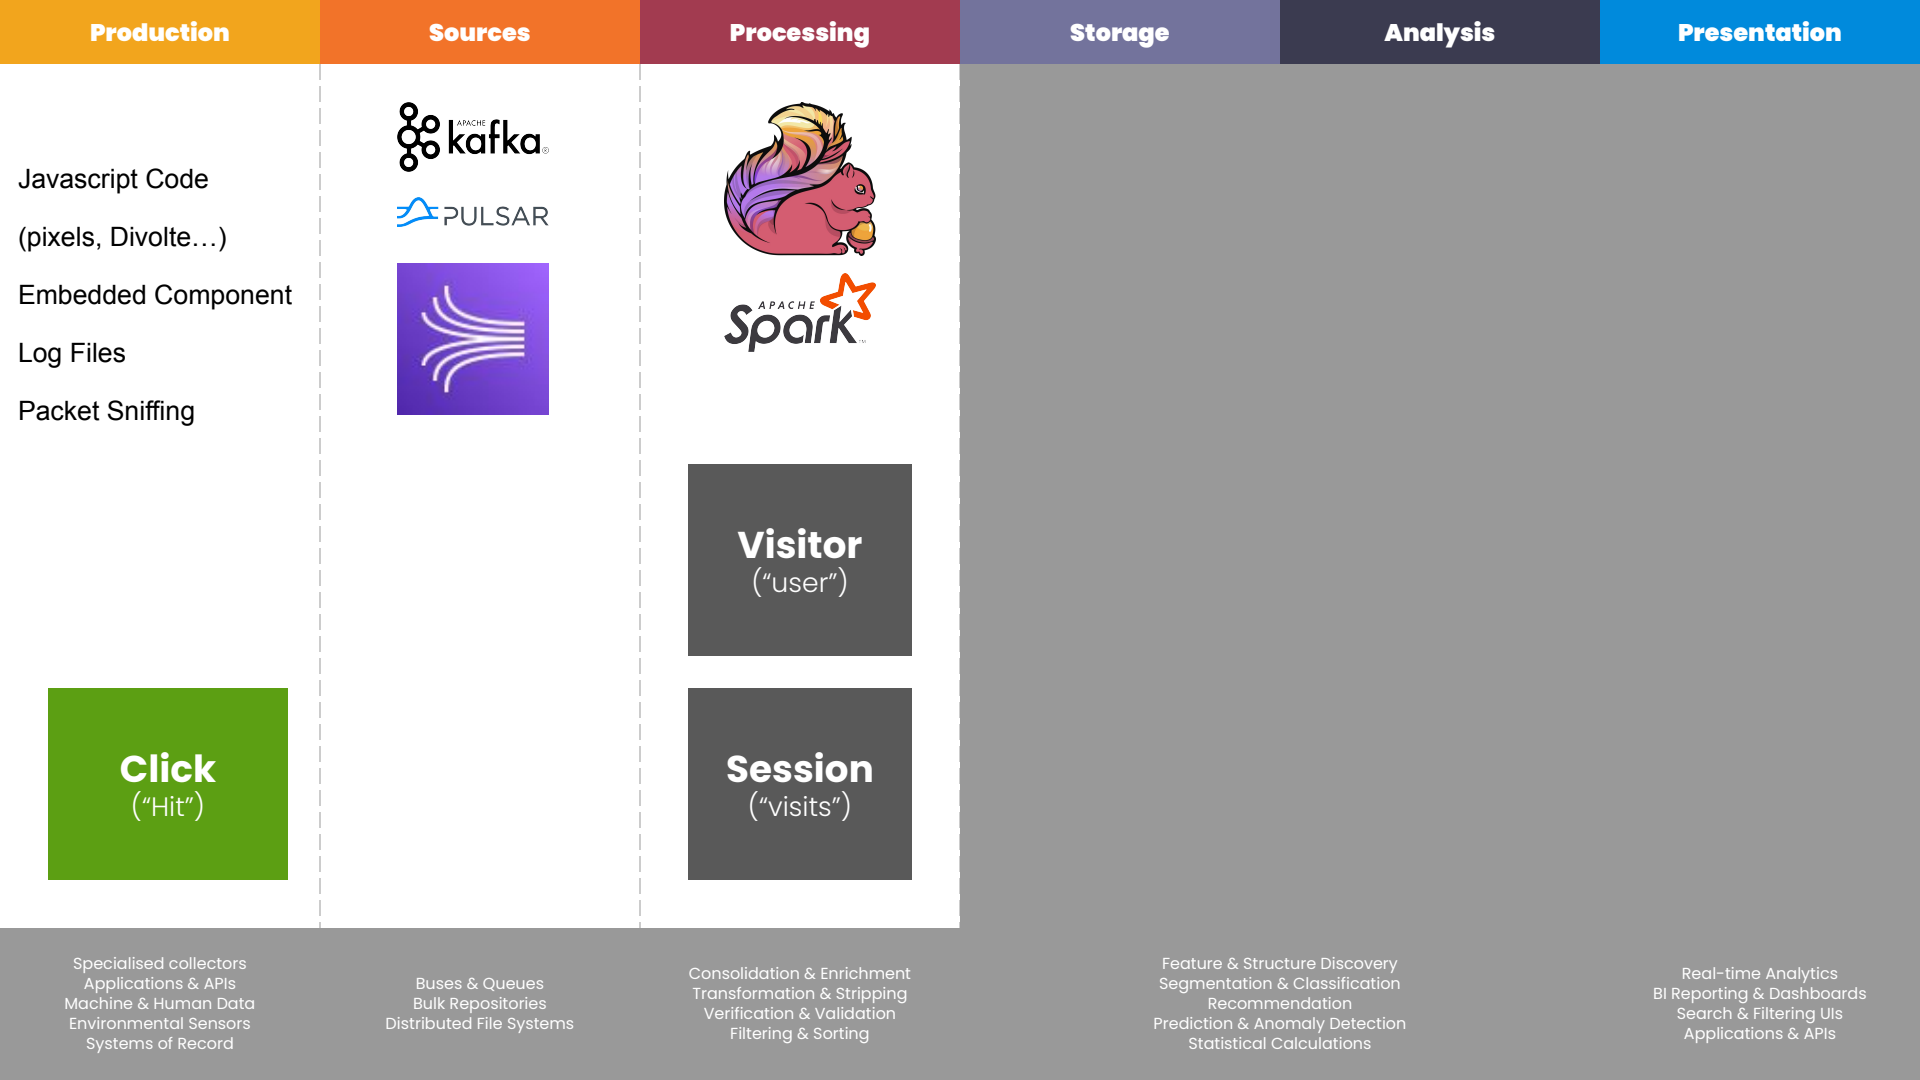


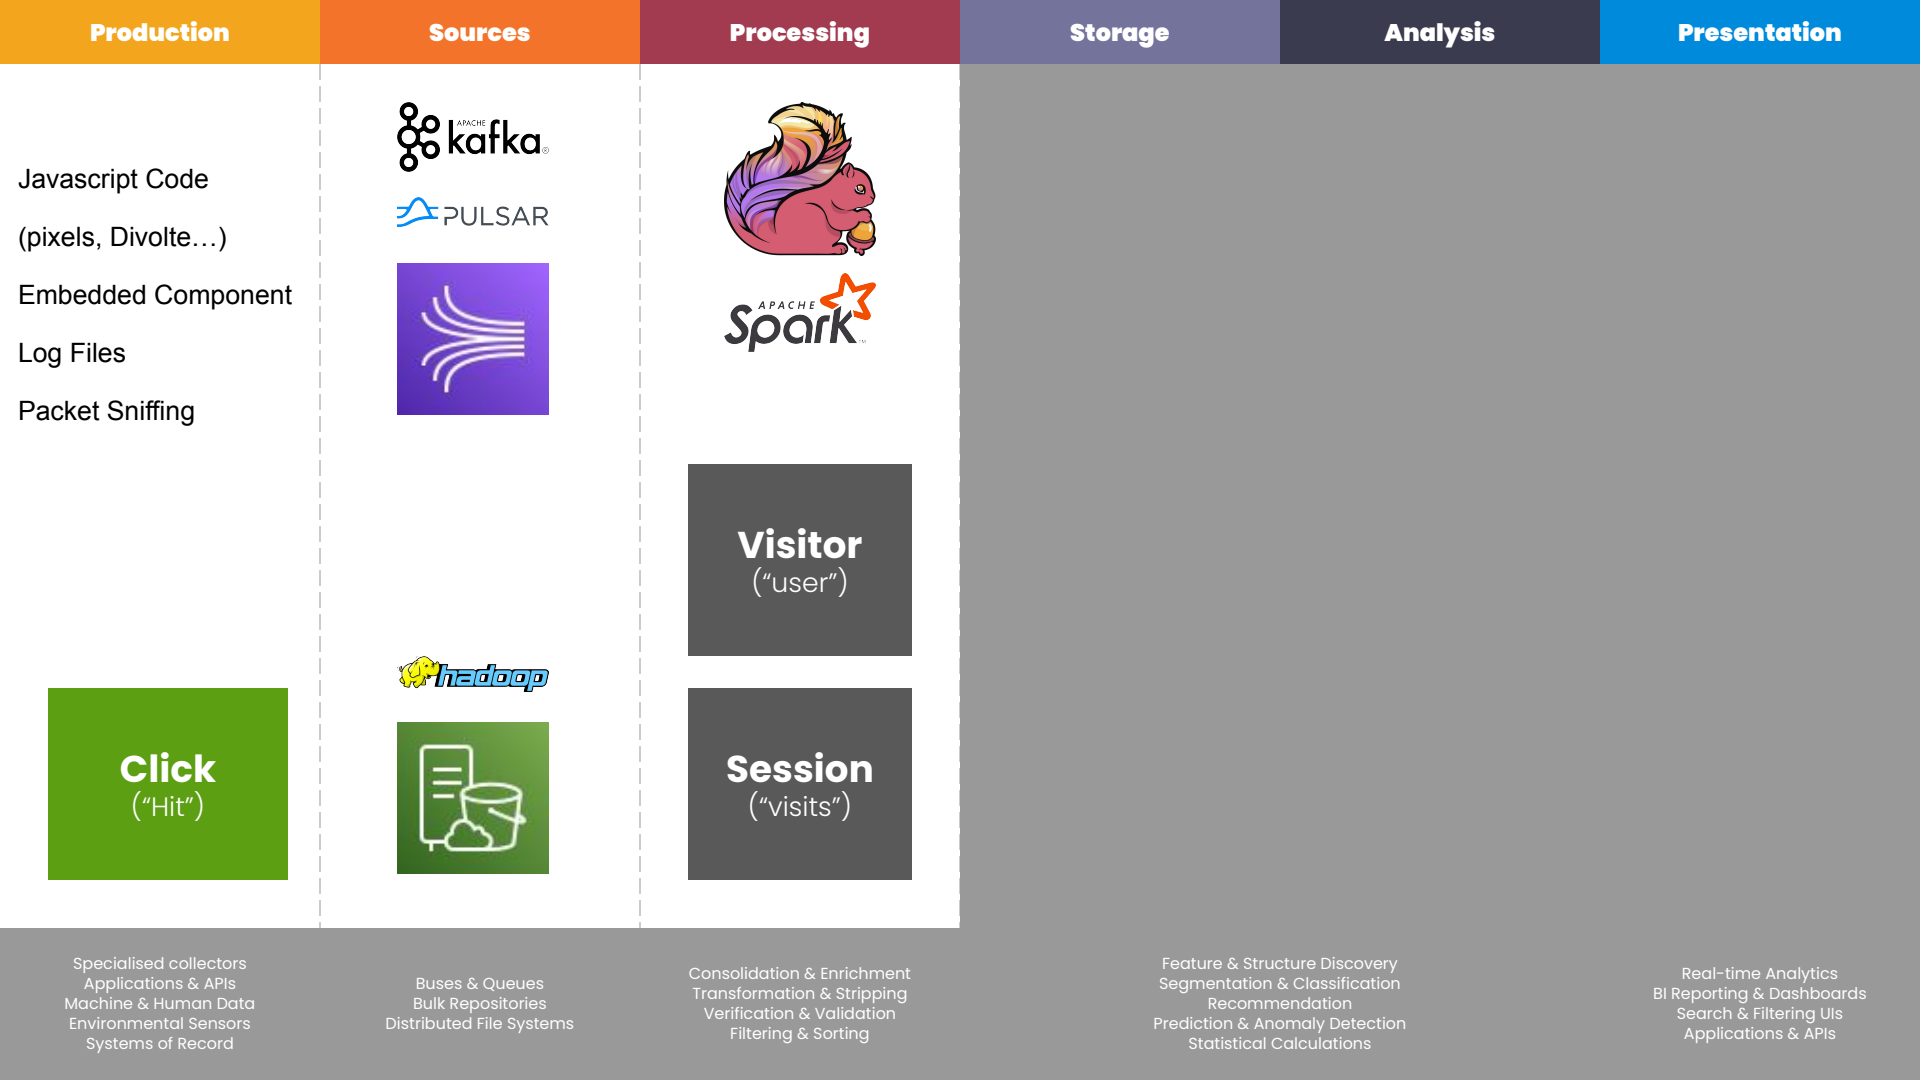
Clickstream is stateless

We don't know when a session ends

Clickstream is anonymous

We don't know who the visitor is



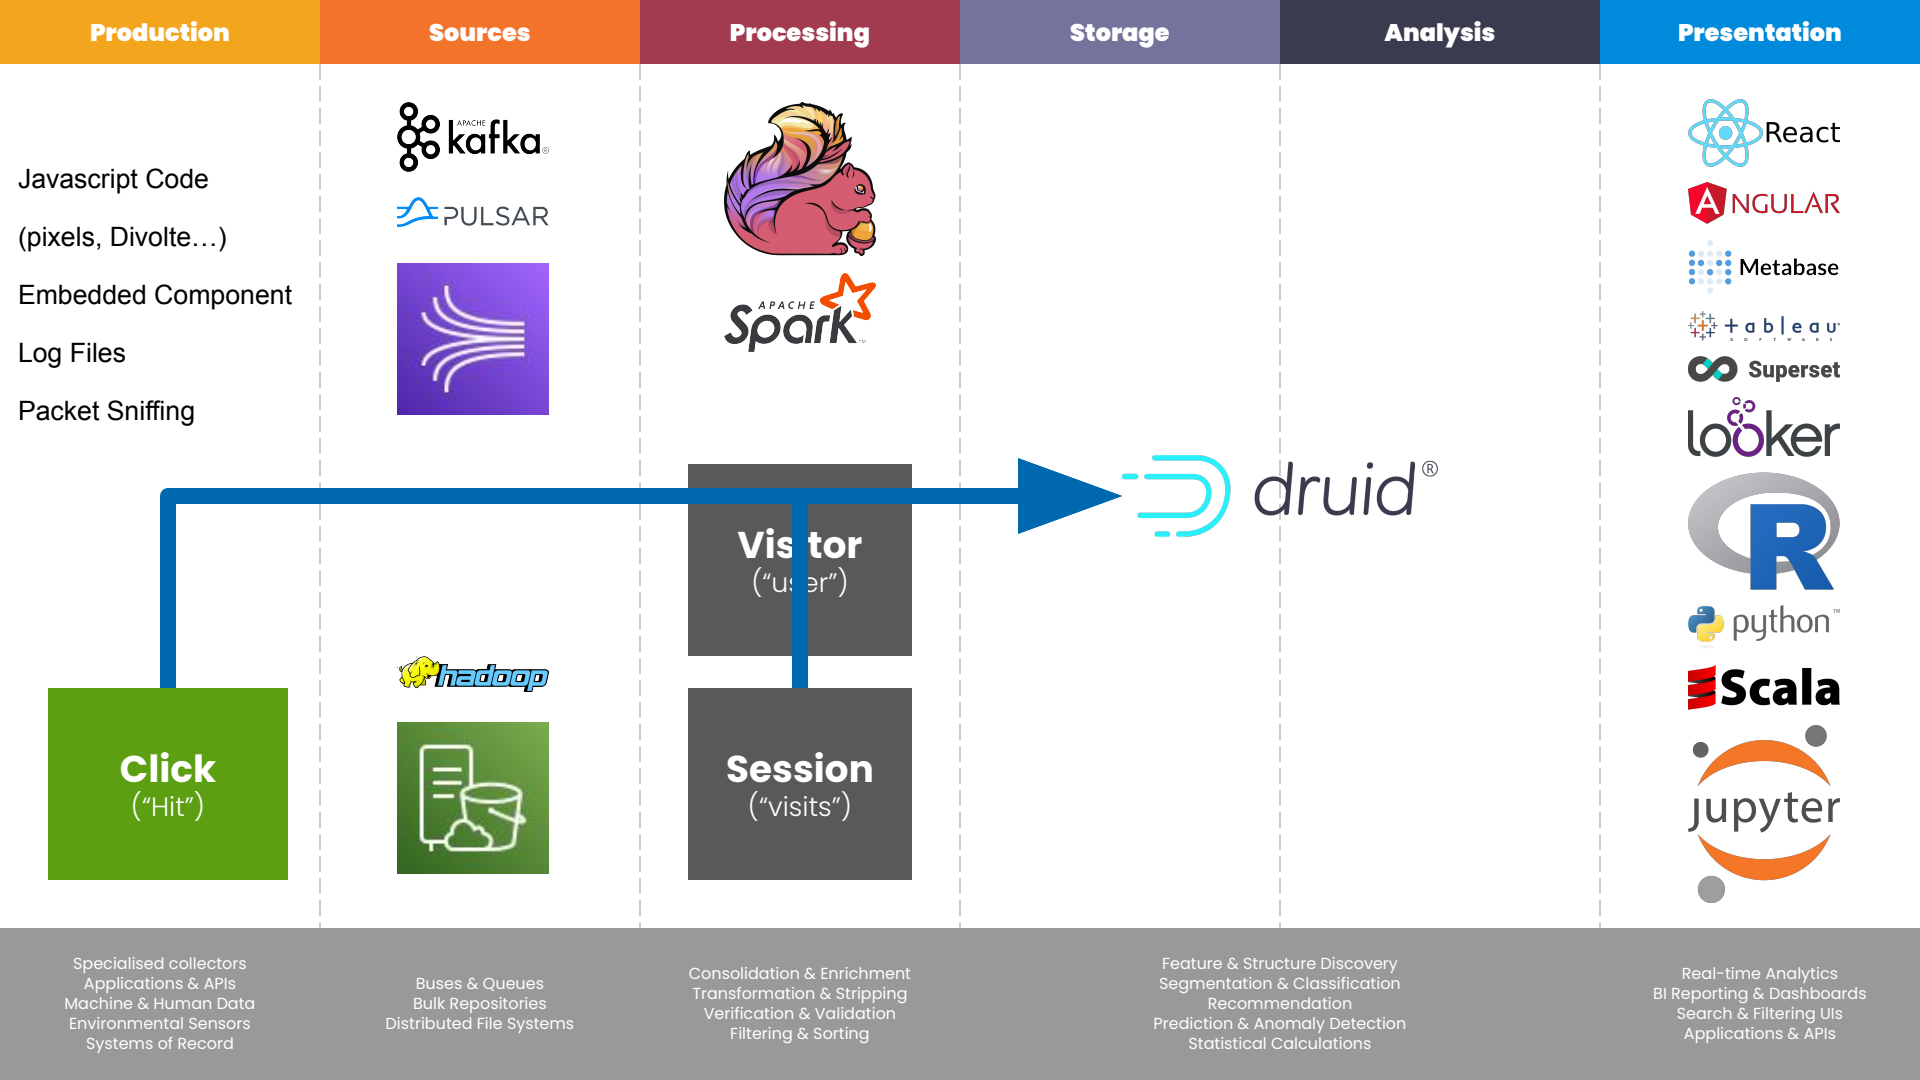


What problems do they face?

- Capturing data is hard
 - Data aggregation
 - Data scale

What problems do they face?

- Capturing data is hard
 - Data aggregation
 - Data scale
- The volume is scary
 - Filtering for the right stuff
 - Doing statistics ad-hoc
 - Solving the COUNT DISTINCT problem



What analytics are we talking about?

- Web Analytics
- Mobile App Analytics
- Advertising
- Streaming Video
- Process Mining



Critical Techniques

Druid functionality

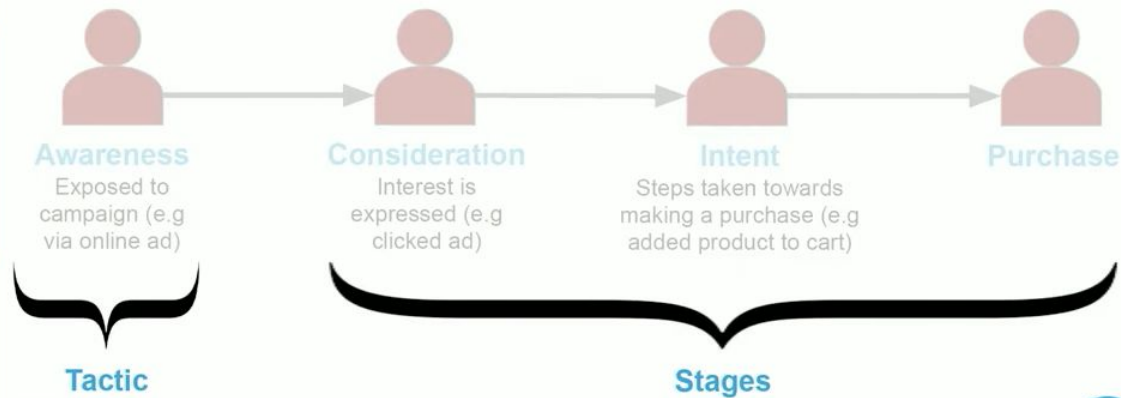
- Sub-partitioning: enhance pruning of data that's picked up (multitenant, market segmentation, end-user segmentation)
- Enrichment: ingestion time and upstream (e.g. visitor demographics)
- Rollup: match end user display requirements – the pixels on the screen!
- Approximation: Use HyperLogLog and Thetasketches for DISTINCT COUNT and for set analysis (funnels)
- Streaming: get data in FAST!
- Compaction: Dealing with late arriving and out of order data
- Changeable Schemas: Adapting to changes in upstream data
- Expressions: Ingestion-time or upstream calculation (e.g. RegEx)
- Hyperlean tables: Filtering ahead of time at ingestion or upstream



Critical Techniques: Funnel Analysis



Campaign phases - user's point-of-view

DATA+AI
SUMMIT 2021
FORMERLY SPARK+AI SUMMIT

@ItaiYaffe, @ettigur

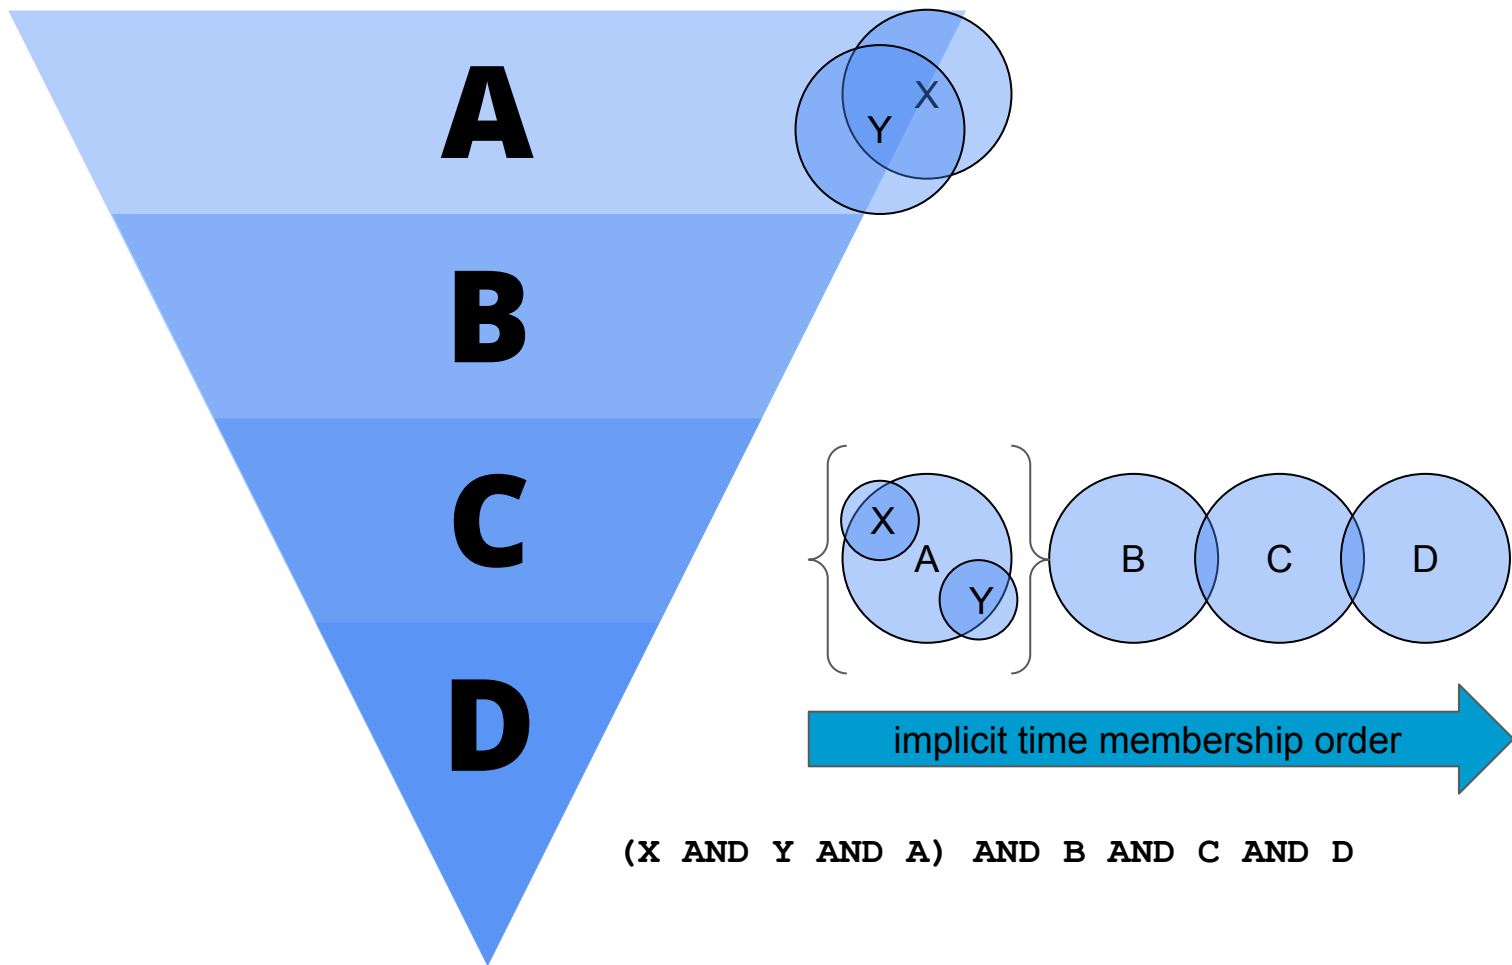
n

ORGANIZED BY  databricks

Funnel Analysis with Apache Spark and Druid

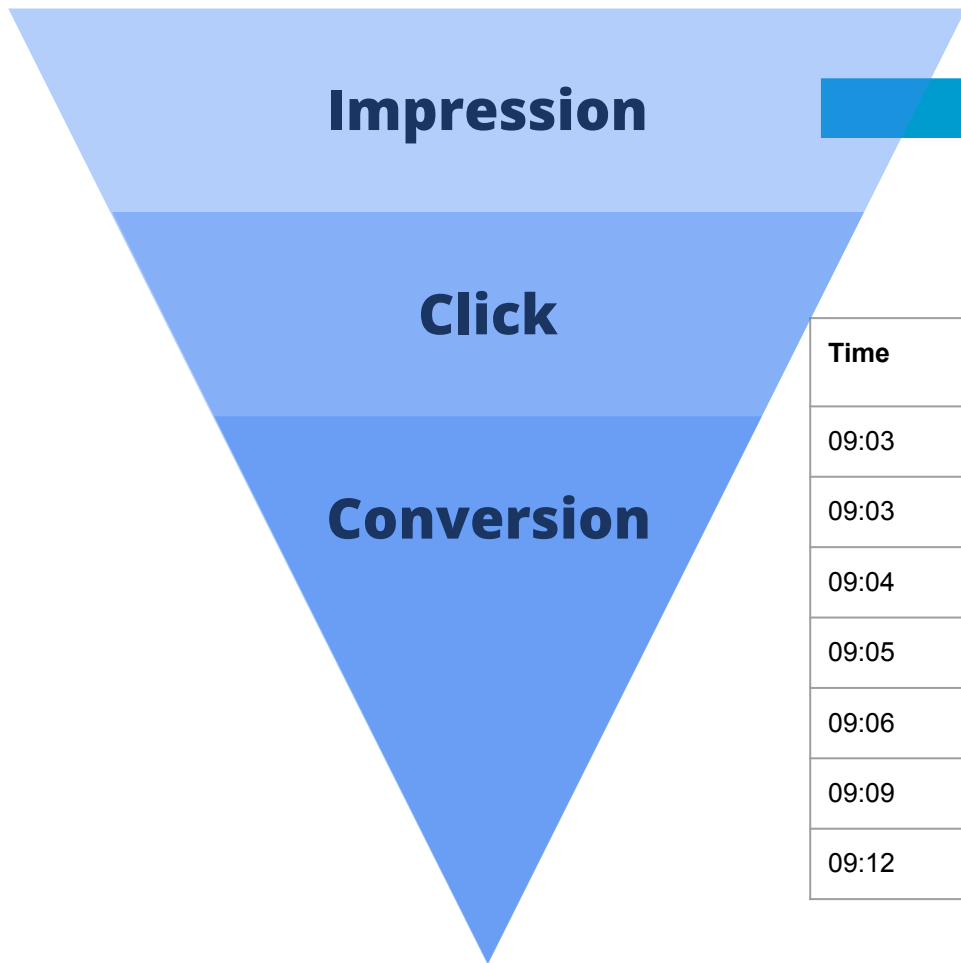
Funnel Patterns

- **Approximation with Set Analysis**
- Enriched Click data
- Enriched Session data

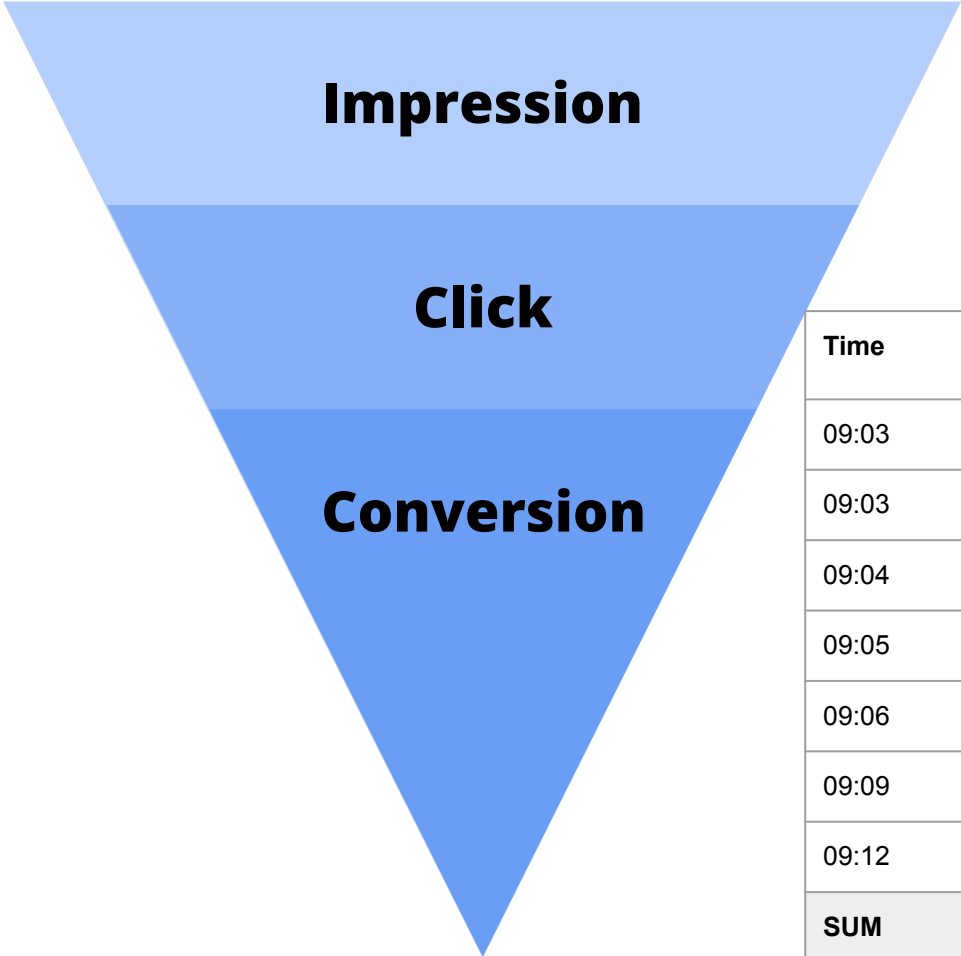


Funnel Patterns

- Approximation with Set Analysis
- **Enriched Click data**
- Enriched Session data




Time	Impression	Click	Conversion
09:03	1	0	0
09:03	0	0	1
09:04	0	0	0
09:05	1	0	0
09:06	0	1	0
09:09	0	1	1
09:12	0	0	1



Time	Impression	Click	Conversion
09:03	1	0	0
09:03	0	0	1
09:04	0	0	0
09:05	1	0	0
09:06	0	1	0
09:09	0	1	1
09:12	0	0	1
SUM	2	2	3

Funnel Patterns

- Approximation with Set Analysis
- Enriched Click data
- **Enriched Session data**



Session	Time	Last Stage
1142	09:04	A
2131	10:03	A
3112	11:43	B
7126	12:51	D

Stage	COUNT
A	5
B	250
C	500
D	1000



A

B

C

D

Session	Time	F1	F2	F3
1142	09:04	A	Z4	pl
2131	10:03	A	Z6	pl
3112	11:43	B	Z21	cart
7126	12:51	D	Z7	ad



Critical Techniques: Sessionisation

Sessions are long-lasting
Sessions are vague things
Some values *cannot* be known at 0s

Session Analytics

Difficulty

How much effort was required to move along?

Progress

What was the schedule from A to B?

Goals

Did they attain any achievements or rewards?

Cause

What kinds of things force a change in state?

Purpose

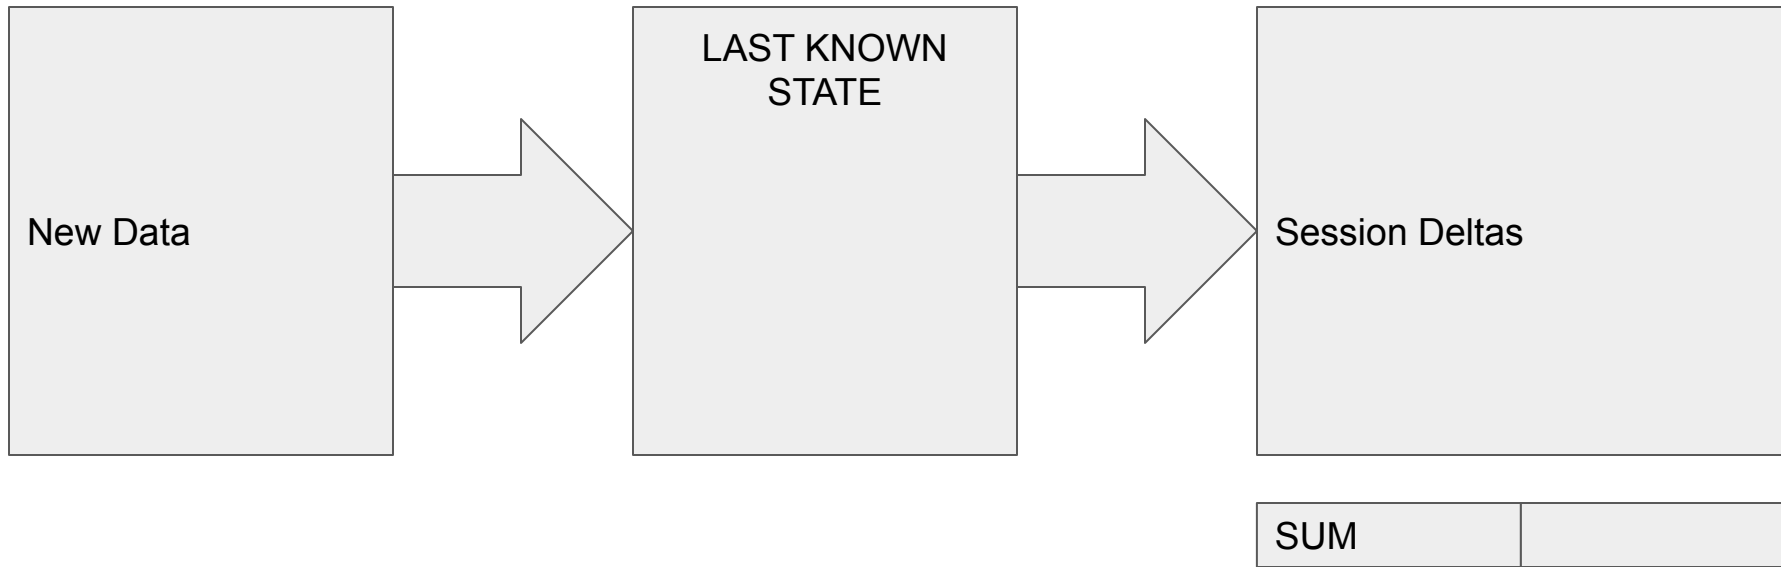
What external factors influence things towards their eventual state?


Plan

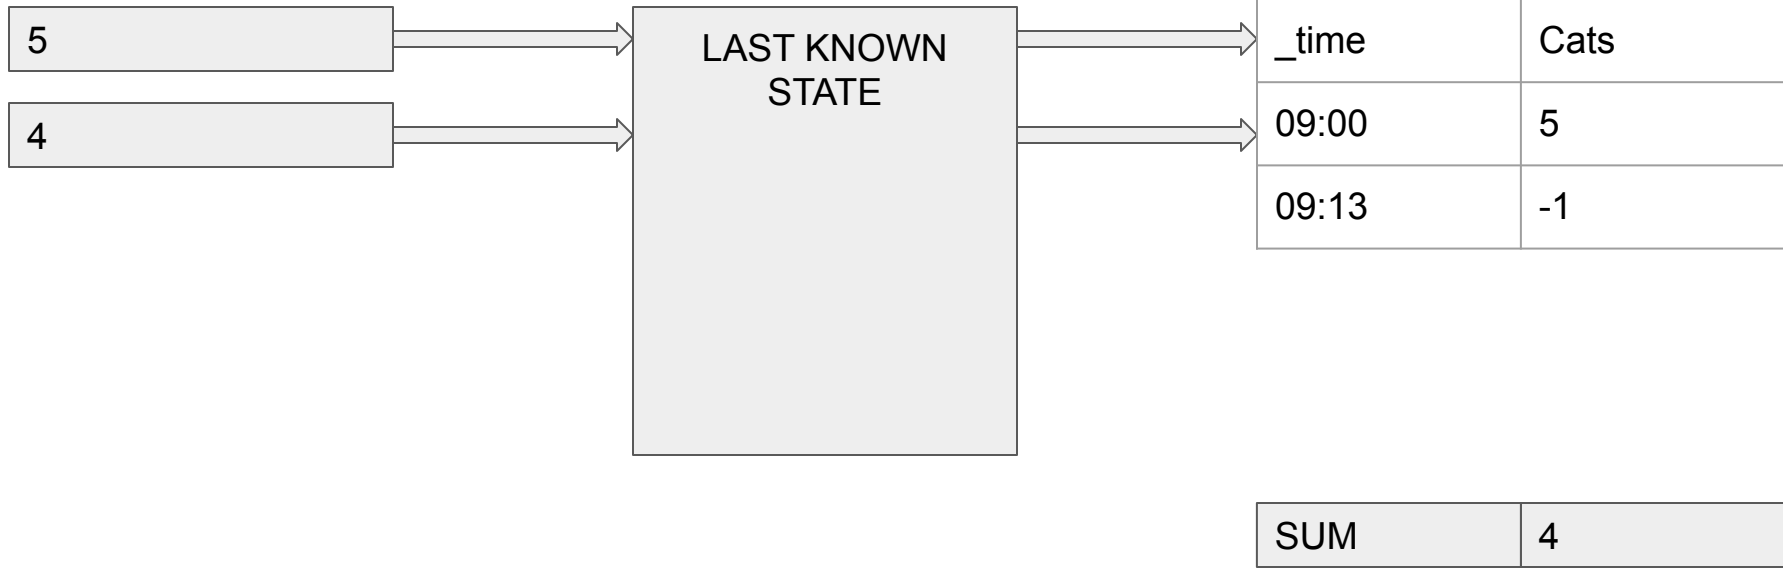
What route did the actor take to getting finished?

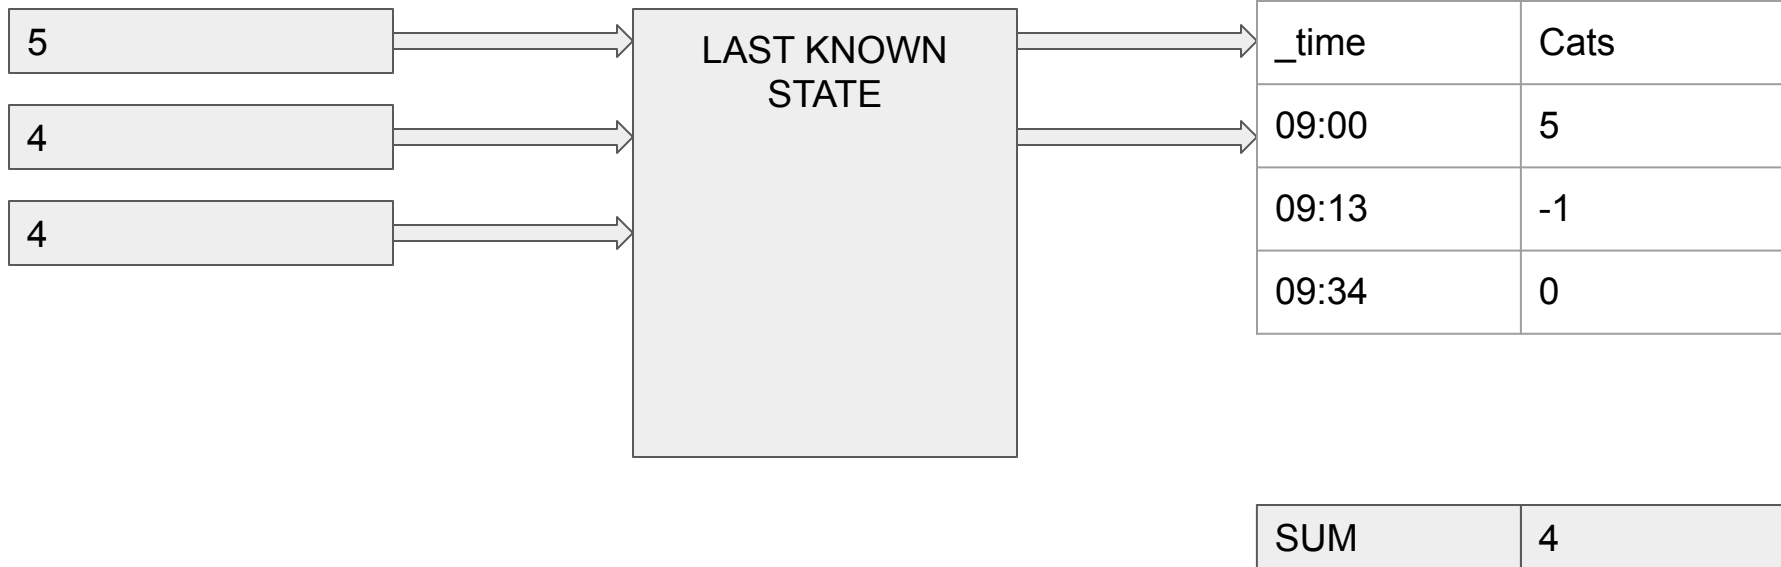
Acts

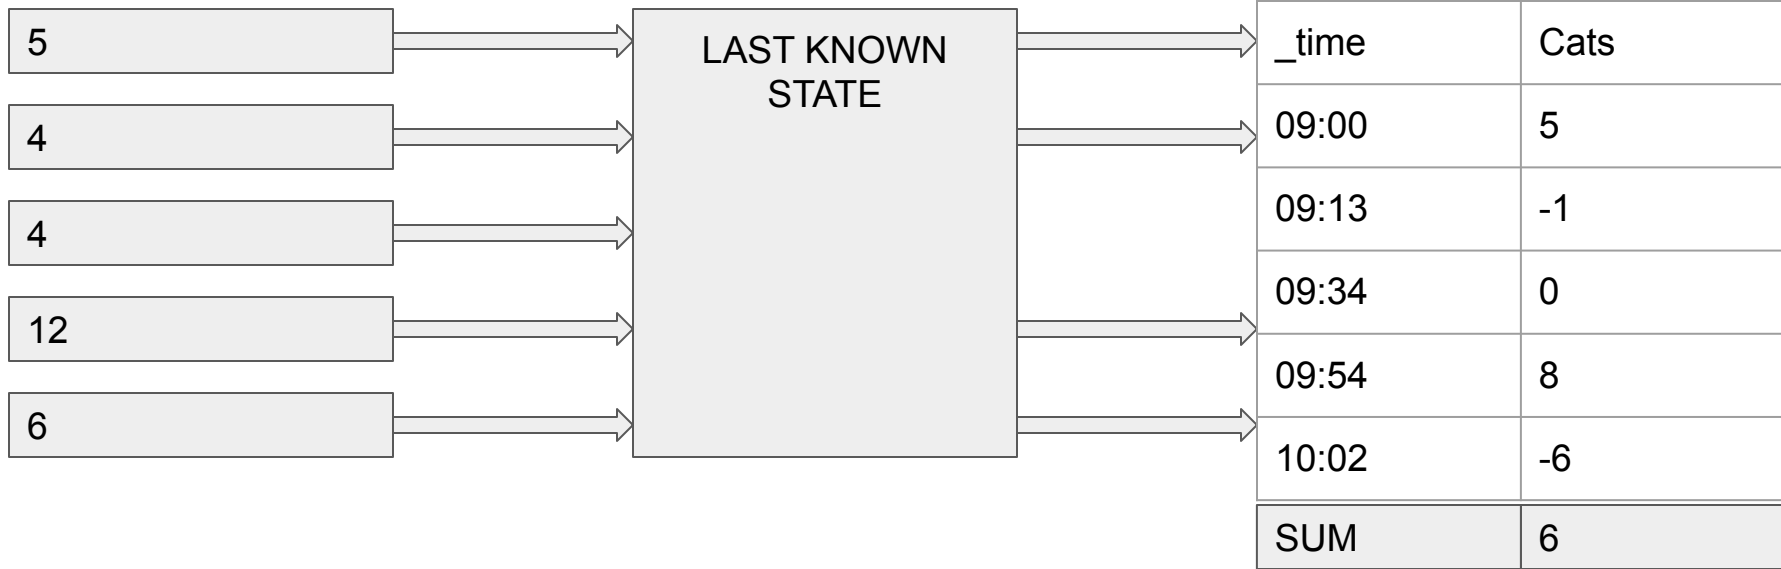
What actions did the actor take?











- Ingestion scalability
- On-demand aggregation
- Filtering efficiency
- Time-based comparison
- Approximation

Get it to the desk!



Thank You!