Kafka + Druid = Immediate Intelligence

Apache Kafka® and Apache Druid® as the platform for modern analytics

Kafka + Druid = Immediate Intelligence

Apache Druid
- Use Cases
- Architecture
- What Makes it Fast

Kafka Integration
- Streaming at any Scale
- Data Enhancement in the Pipeline

Learn more...

Kafka + Druid = Immediate Intelligence

Apache Druid

- Use Cases
- Architecture
- What Makes it Fast

Kafka Integration

- Streaming at any Scale
- Data Enhancement in the Pipeline

Learn more...

druid®

High performance.	*Low-latency, distributed query execution and high throughput ingestion*
Real-time.	*Event data (clickstream, network flows, user behavior, programmatic advertising, server metrics, IoT…)*
Analytics.	*Counting, ranking, statistics...*
Database.	*Highly-available, time-sharded, partitioned, columnar, indexed, compressed, versioned materialized view*

The adoption and maturity of our technology

Modern Analytics Applications

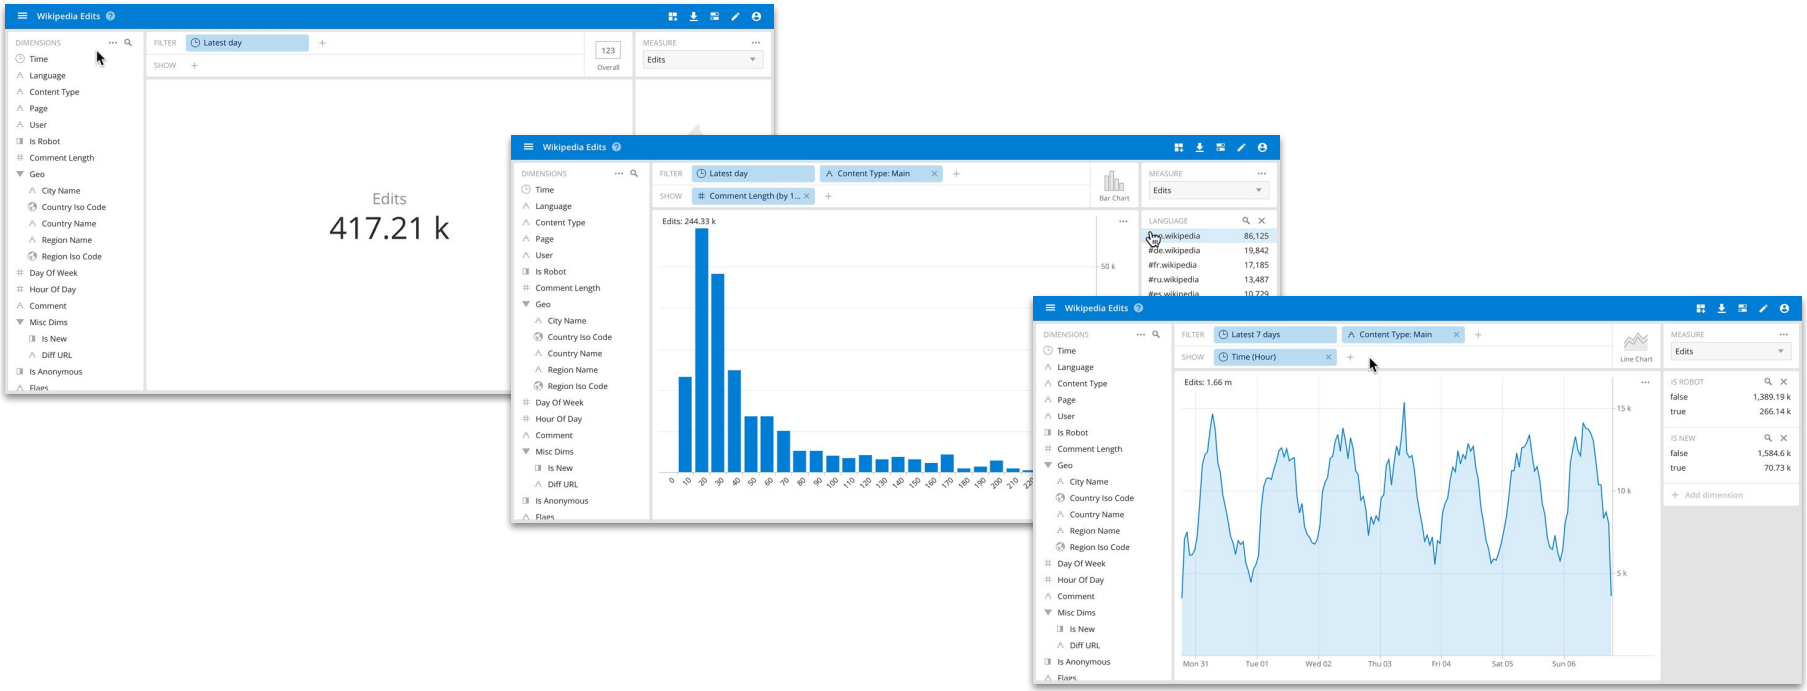

⏱️ 0.1–3s query

🏋️‍♀️ high concurrency

🚰 fresh data

🚴‍♀️ highly interactive

Where Immediate Intelligence matters...

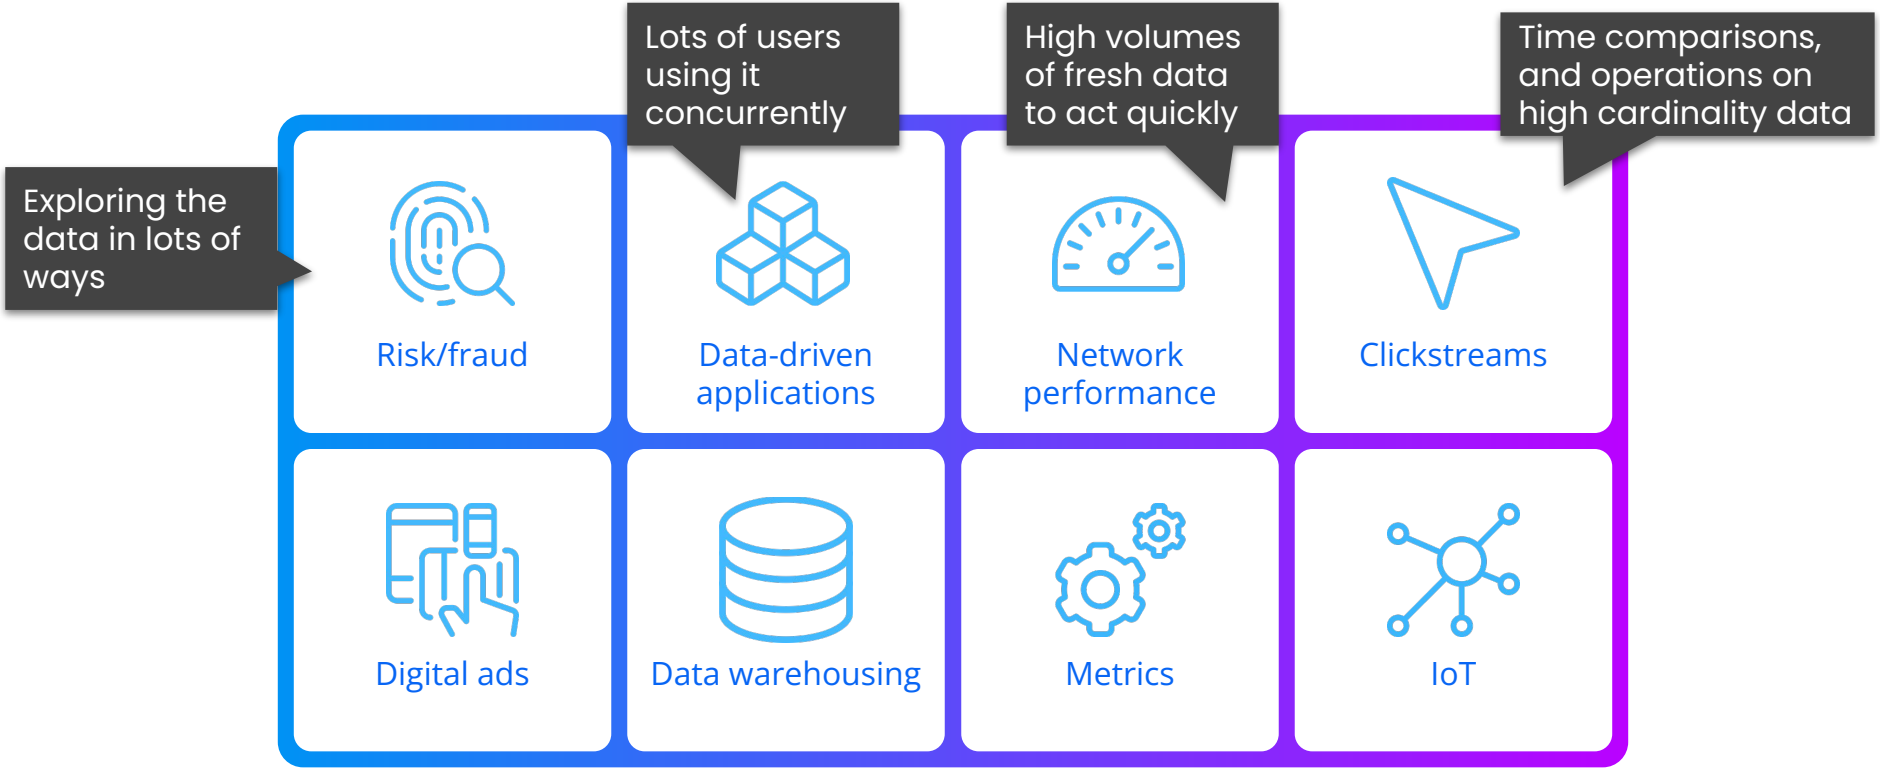

Architecture – Druid uses a microservice architecture

Scalable & Resilient
- Ingestion
- Data Management
- Query

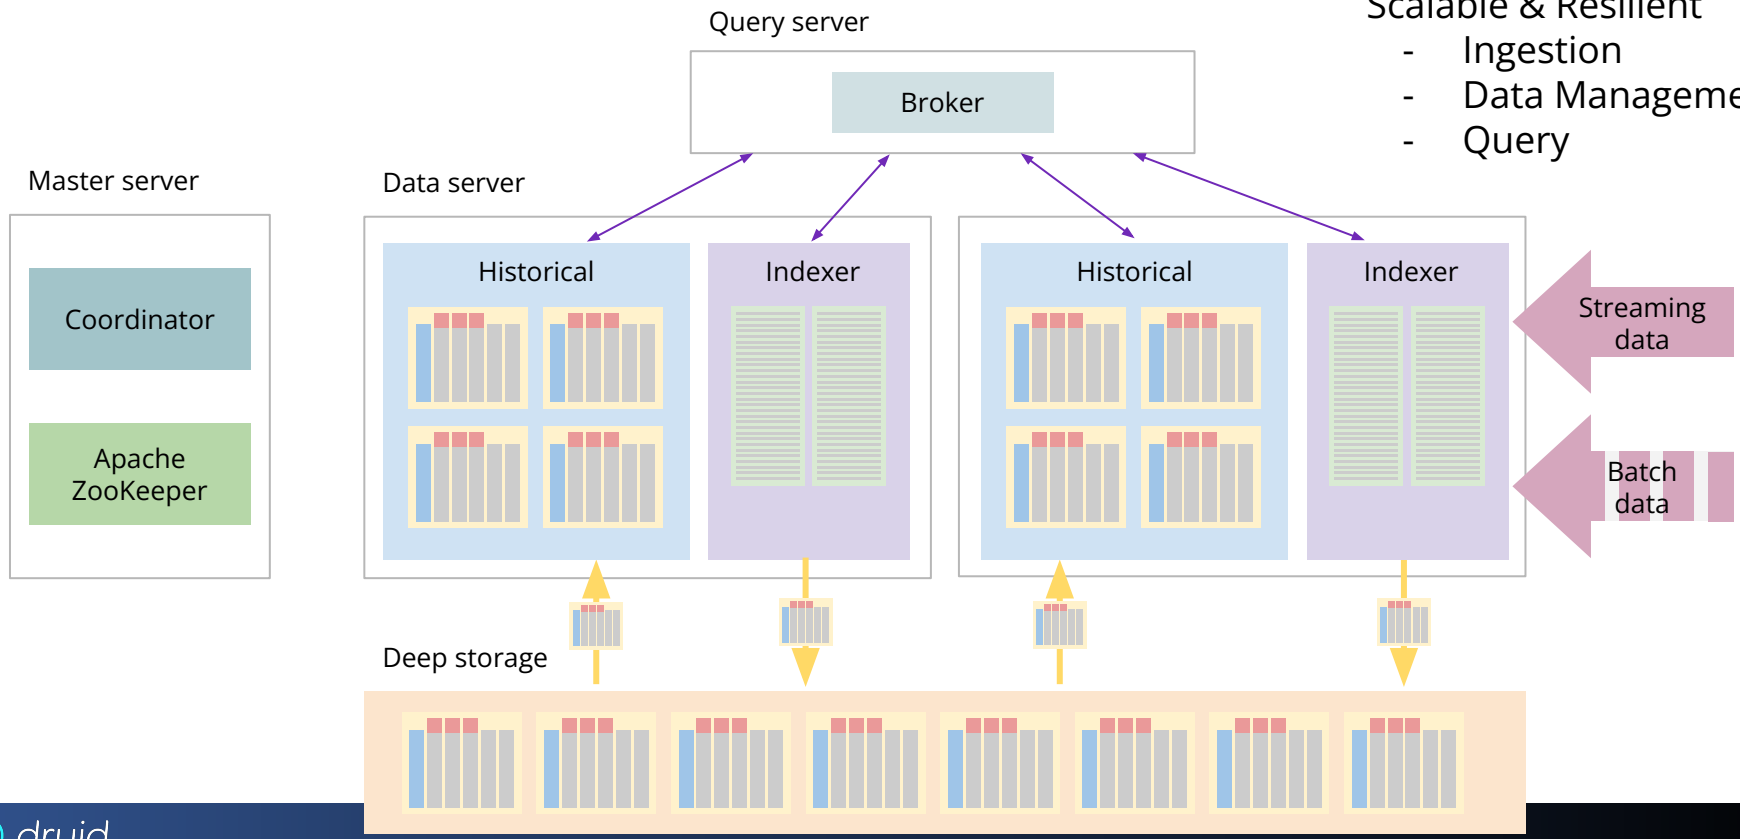

What Makes it Fast

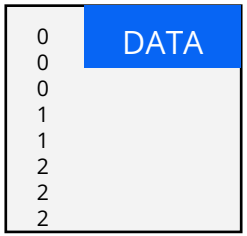

Compression

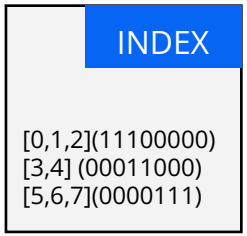

Secondary indexes

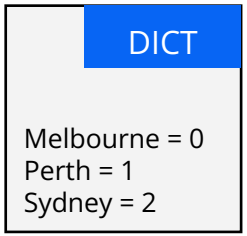

Operate on compressed data

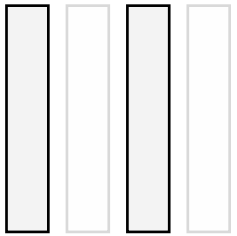

Late materialization

Engine and data format are tightly integrated

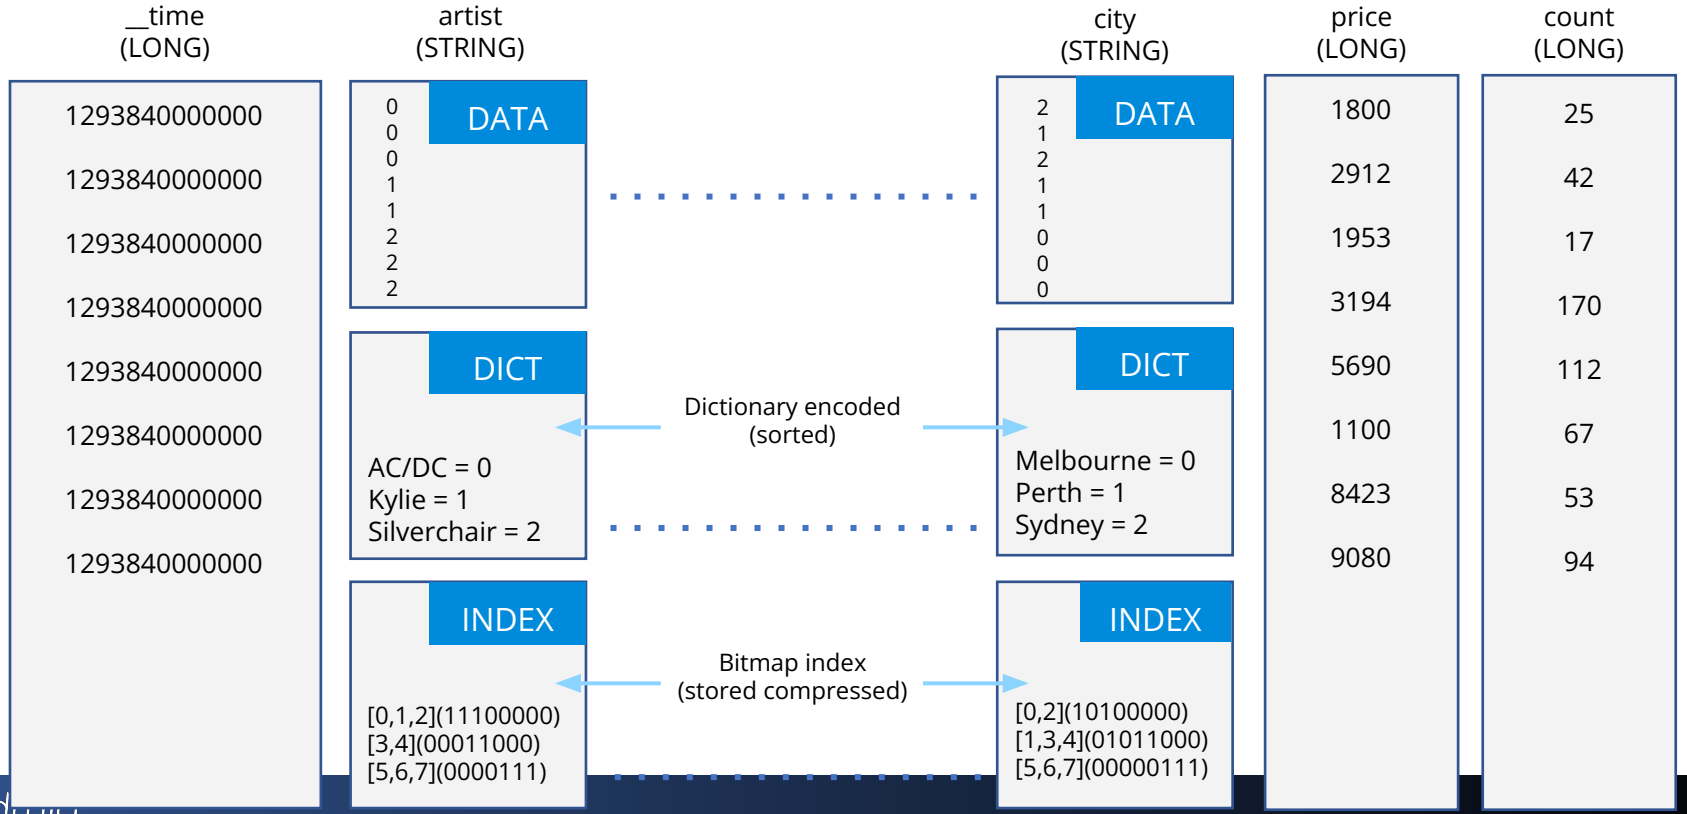

Engine and data format are tightly integrated

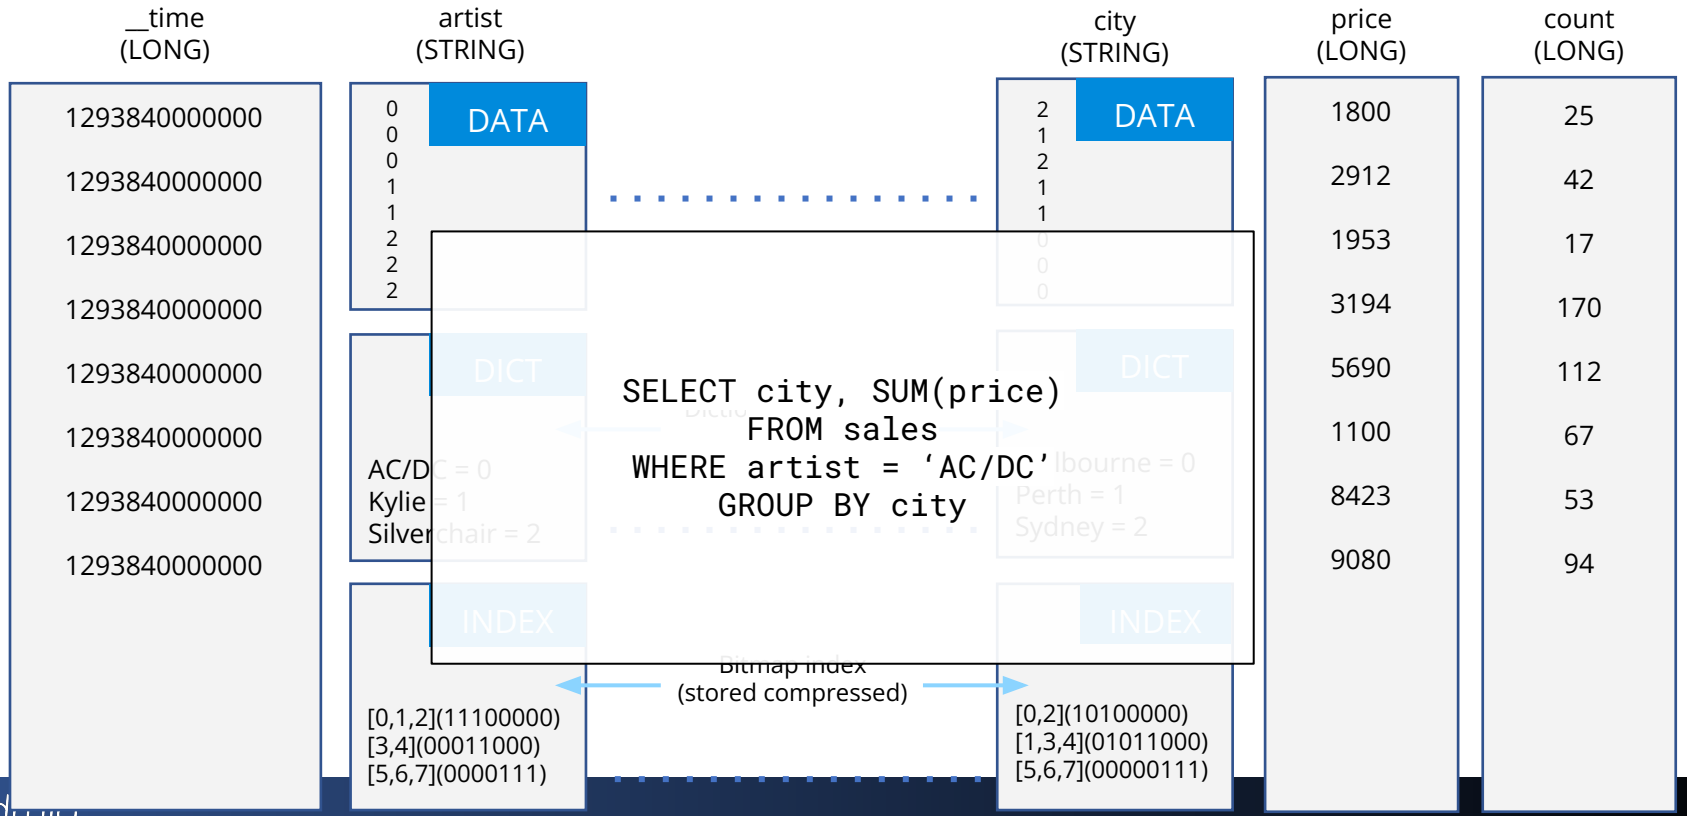

Engine and data format are tightly integrated

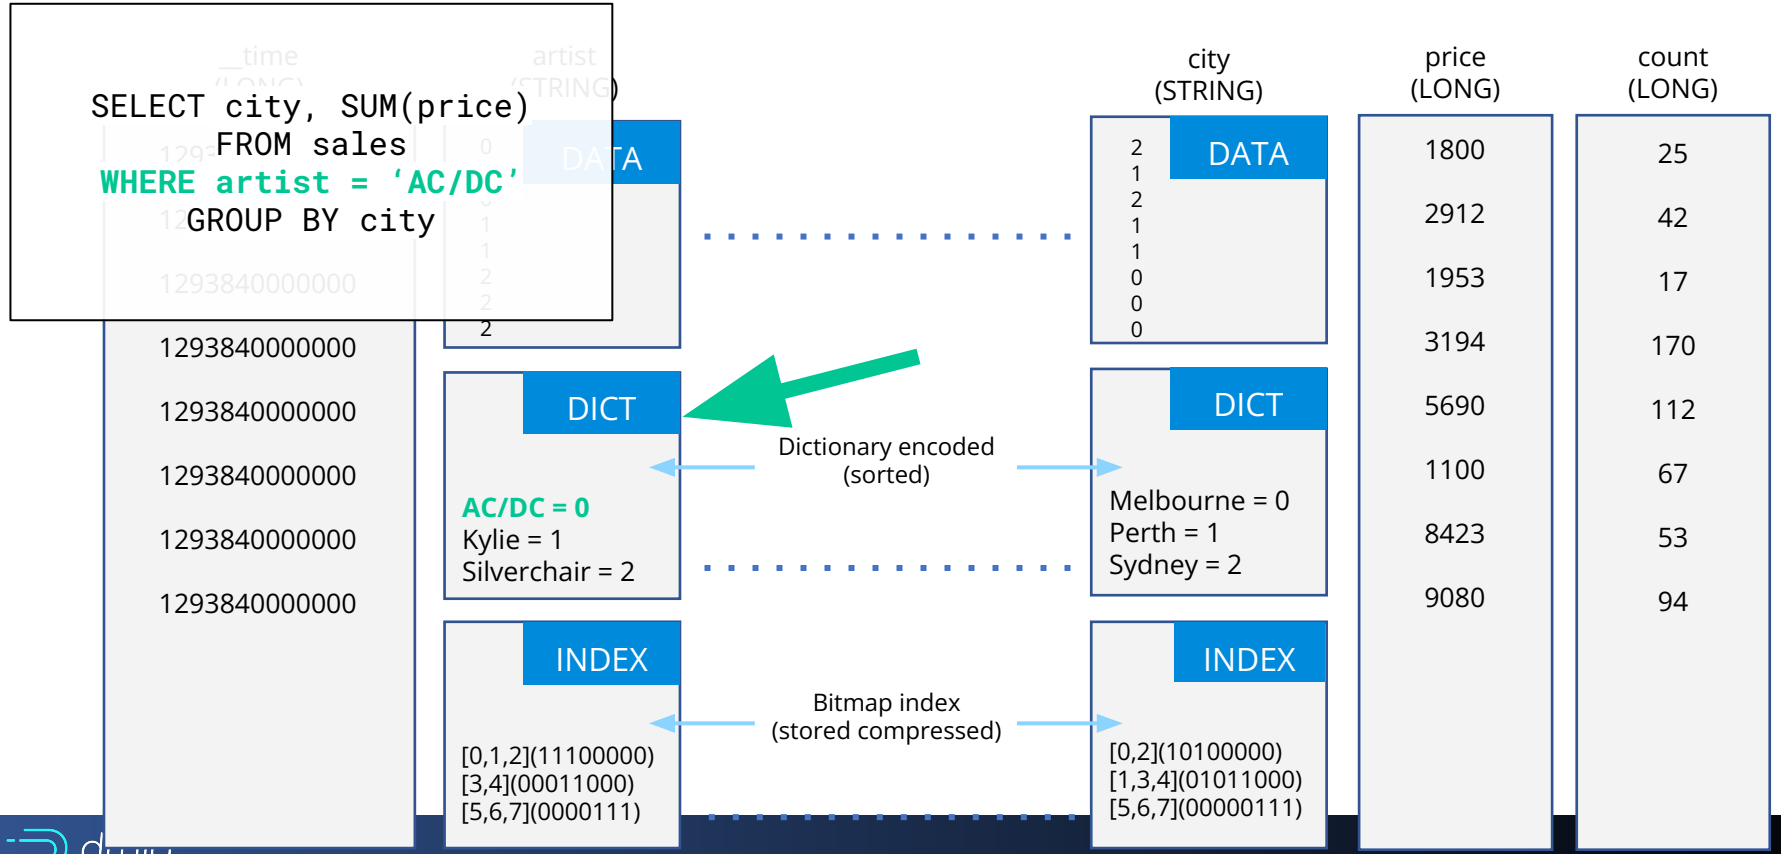

Engine and data format are tightly integrated

```
SELECT city, SUM(price)
    FROM sales
WHERE artist = 'AC/DC'
    GROUP BY city
```

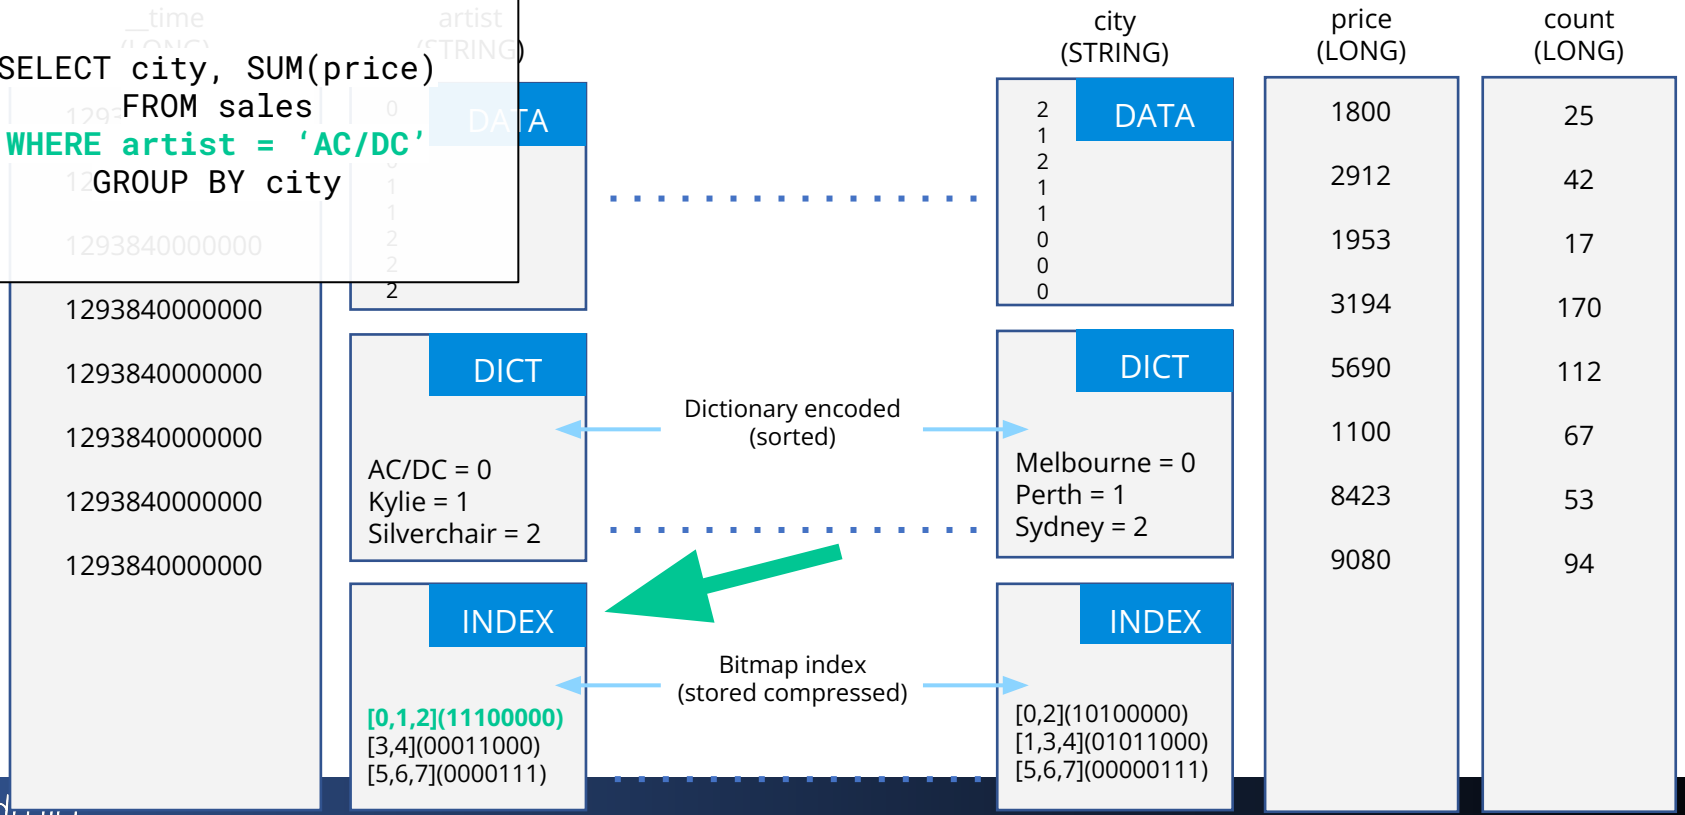

Engine and data format are tightly integrated

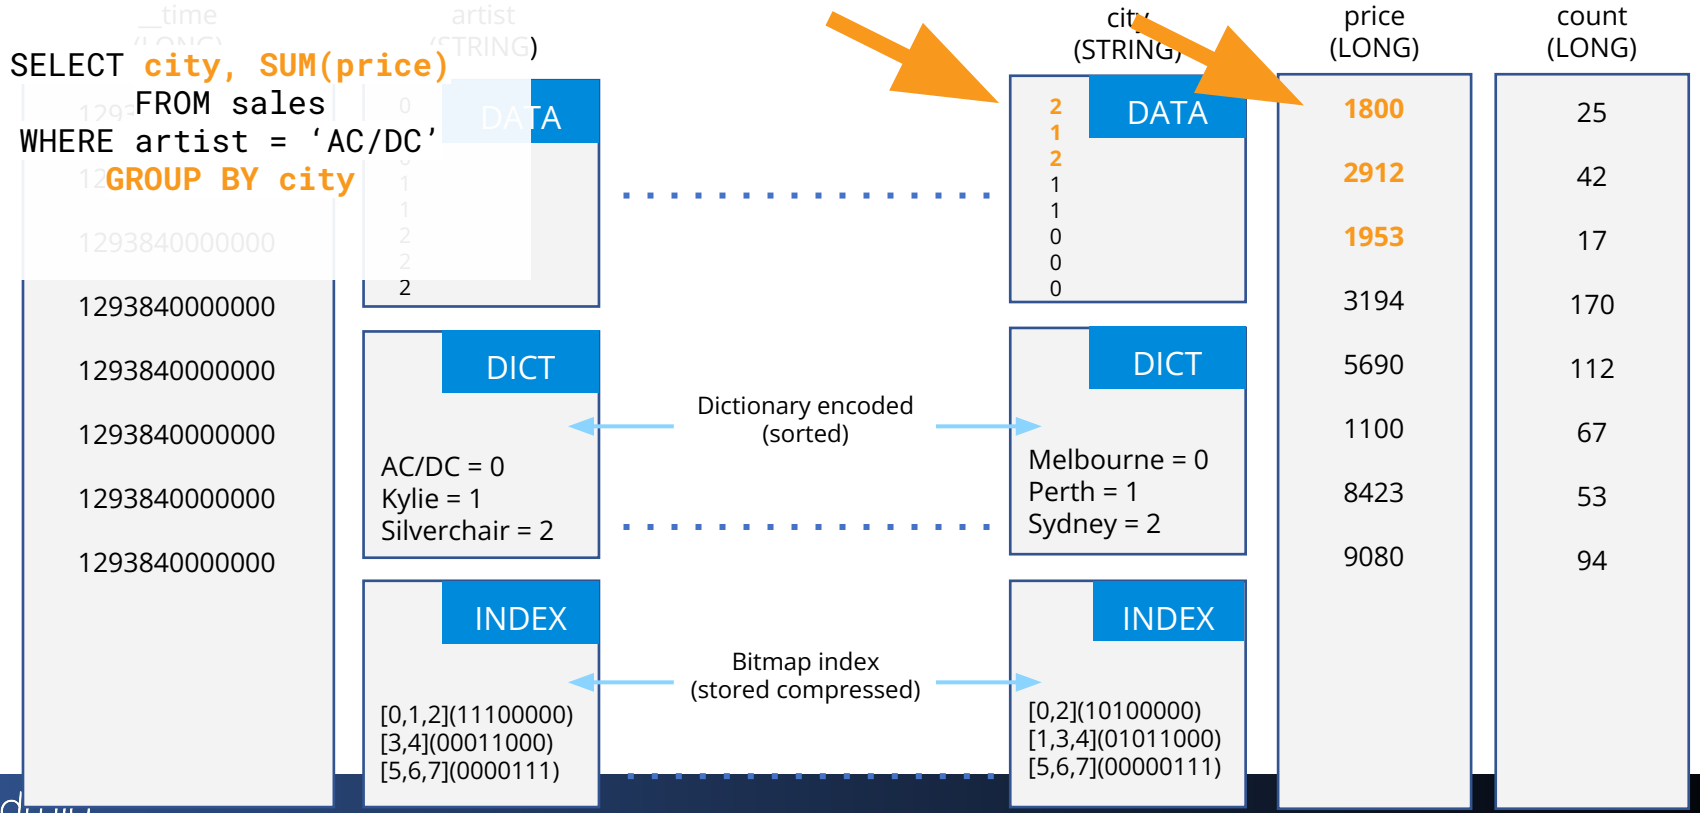

Engine and data format are tightly integrated

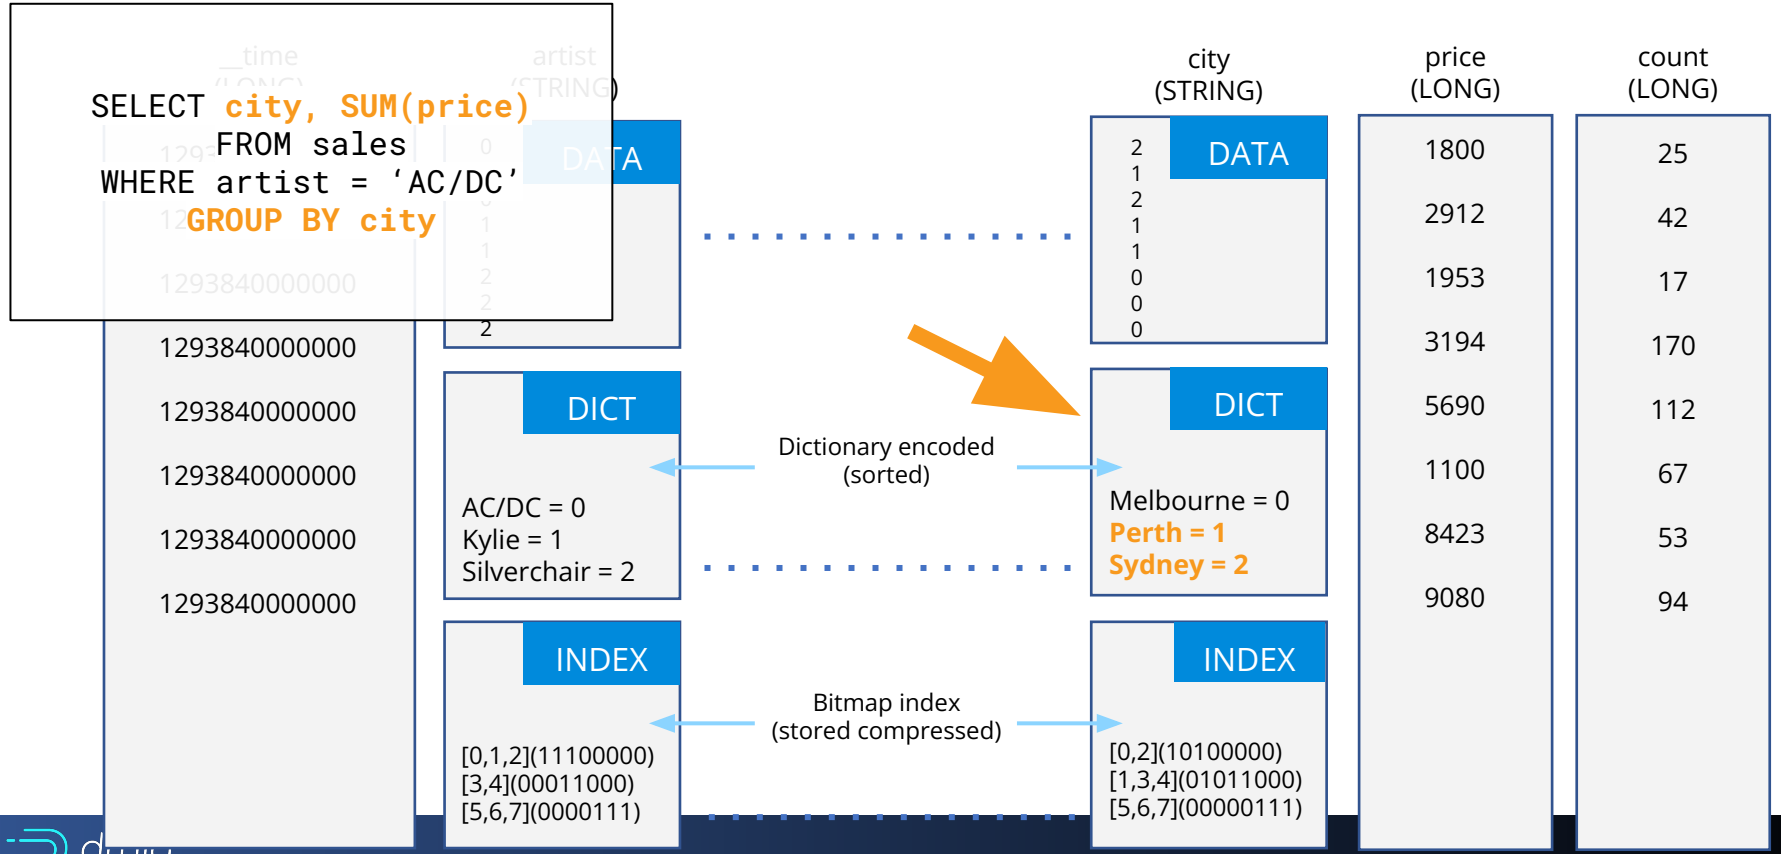

Engine and data format are tightly integrated

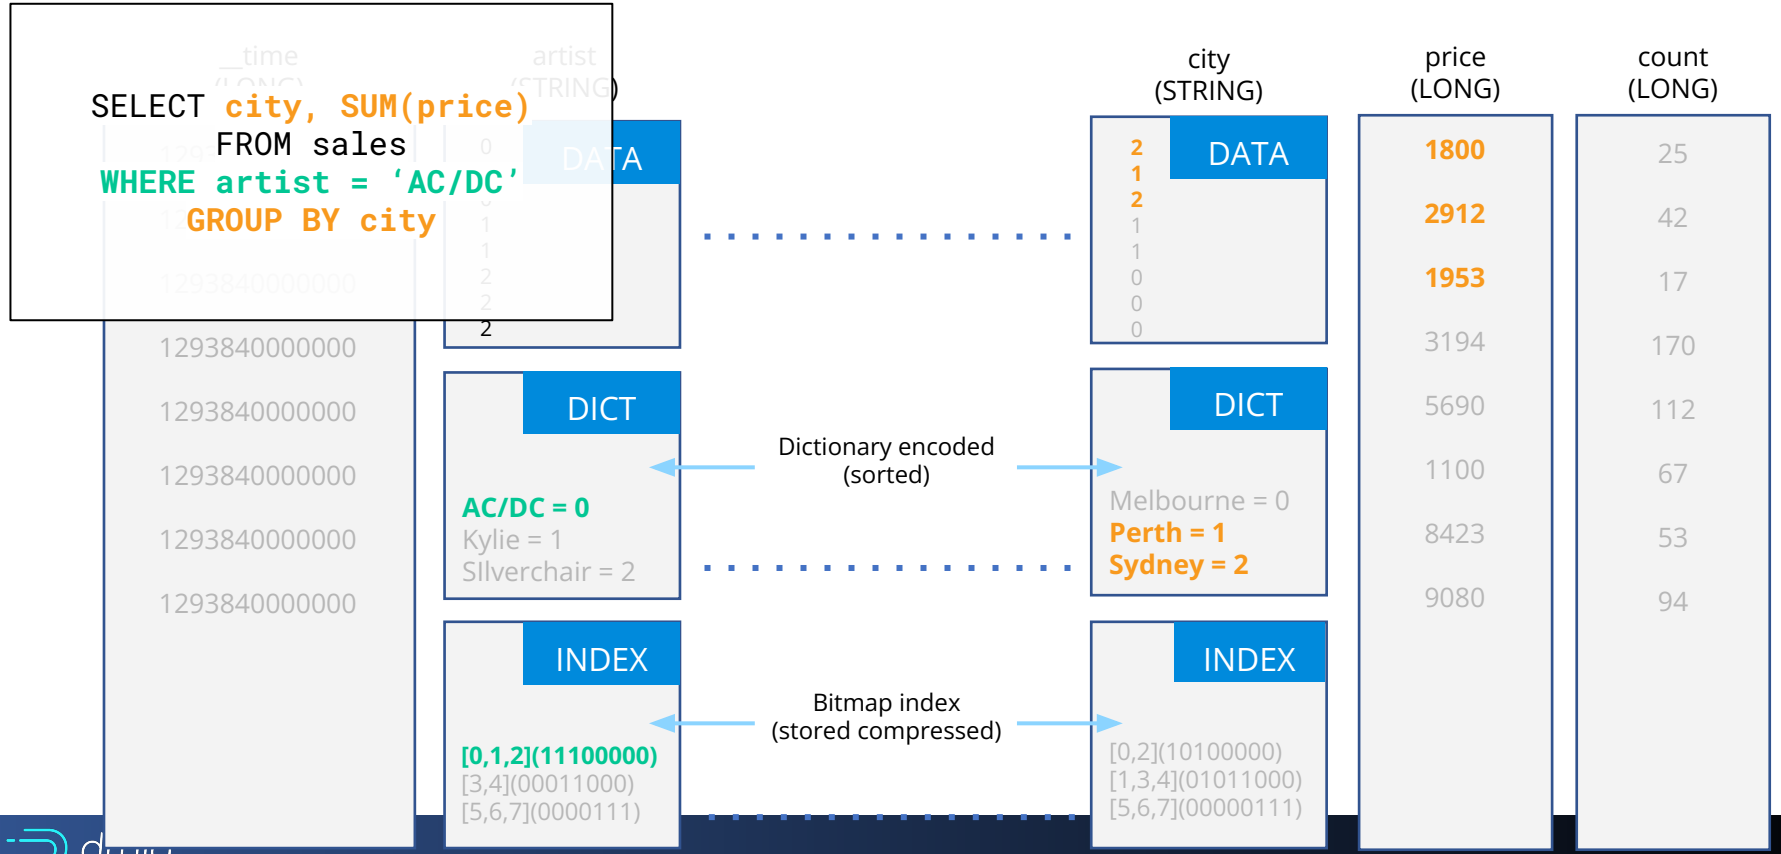

Druid uses a microservice architecture

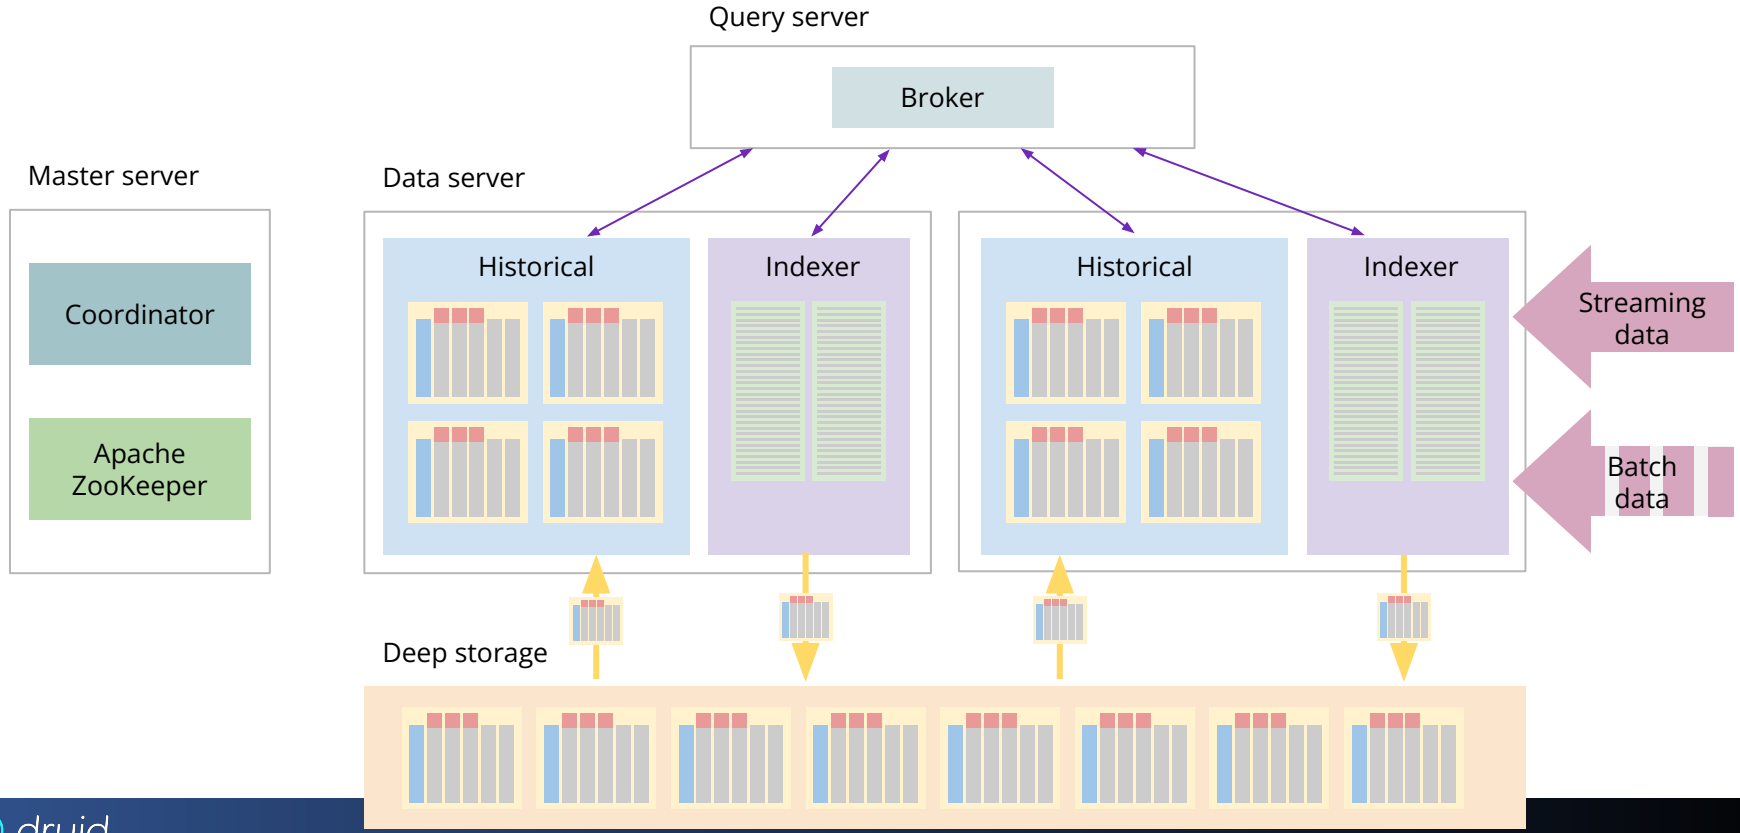

Kafka + Druid = Immediate Intelligence

Apache Druid
- Use Cases
- Architecture
- What Makes it Fast

Kafka Integration
- Streaming at any Scale
- Data Enhancement in the Pipeline

Learn more...

Scalable Data Ingestion

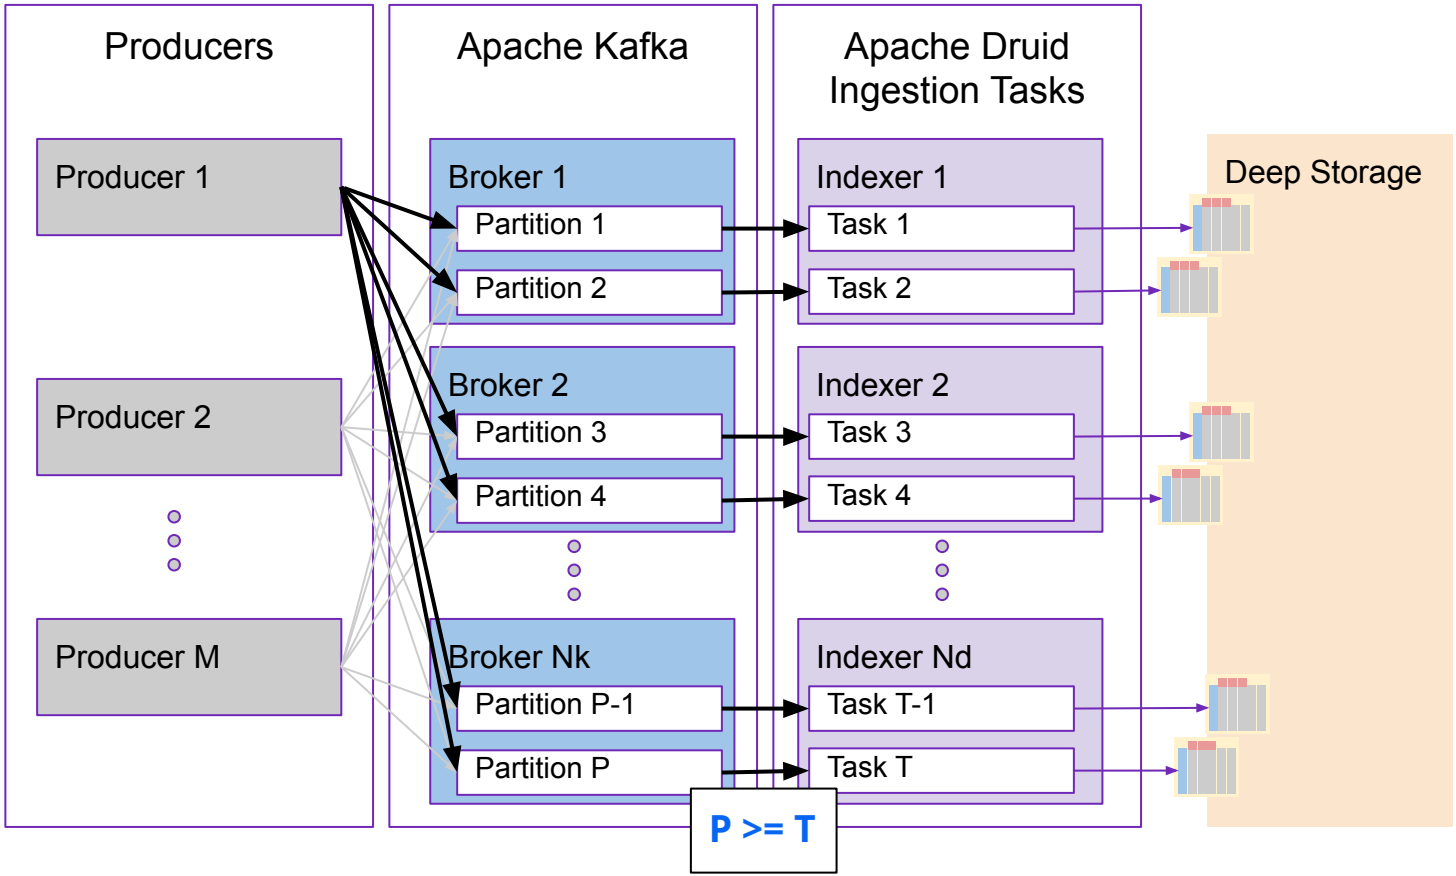

Scalable Real-time Queries

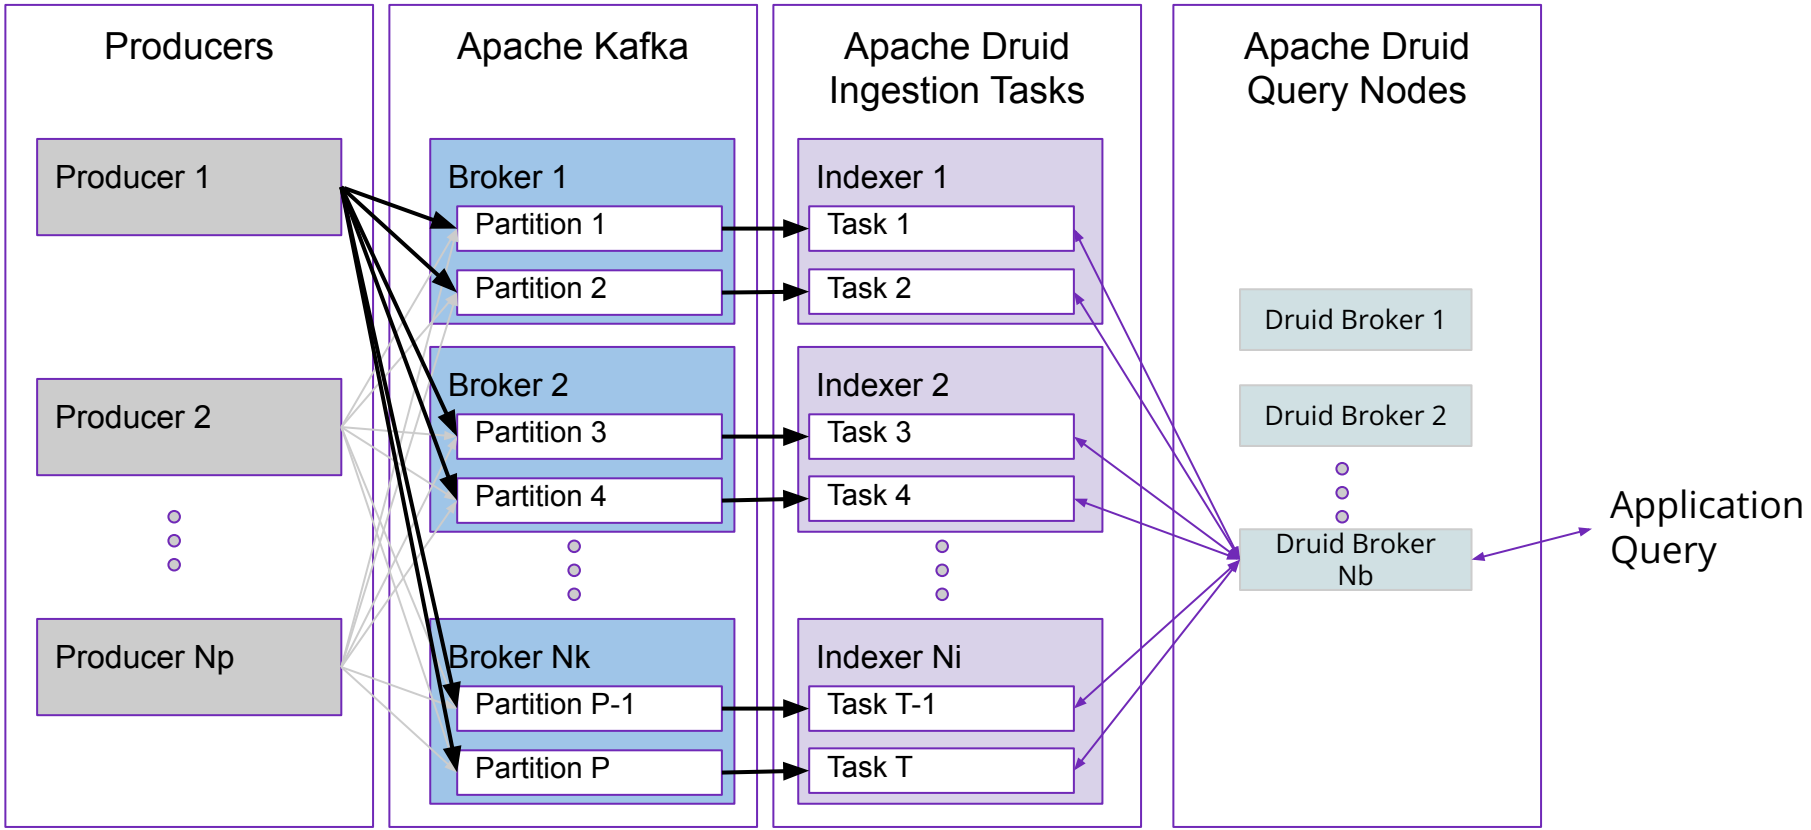

Scalable Real-time + Historical Queries

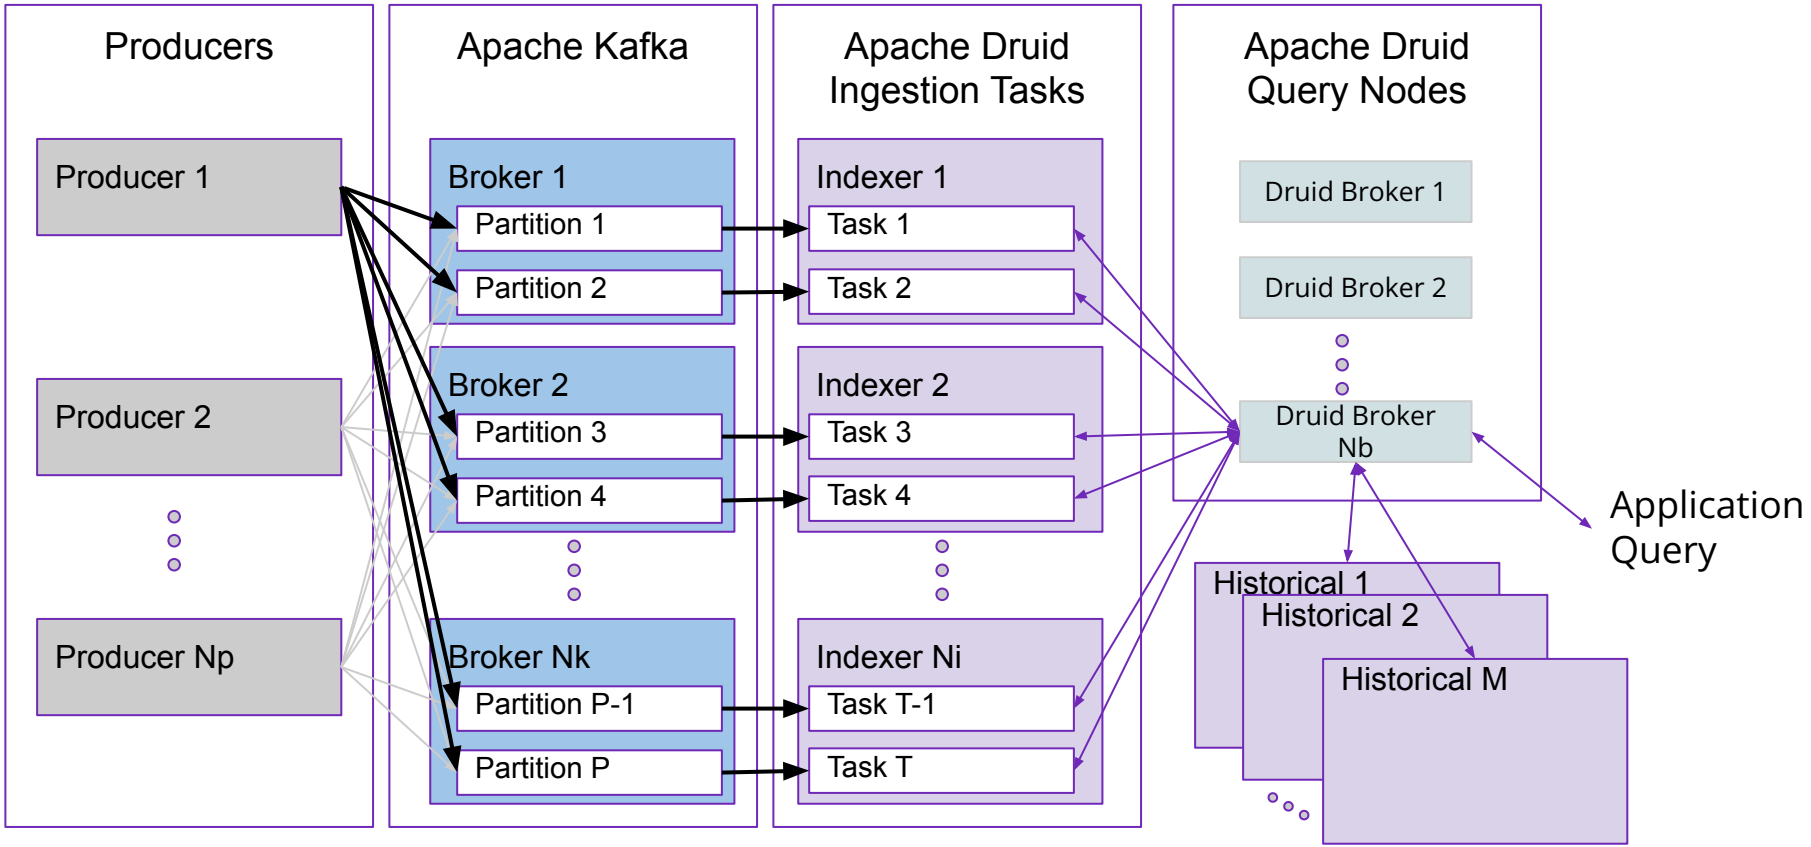

Best Practice: Scalable Data Ingestion + Autocompaction

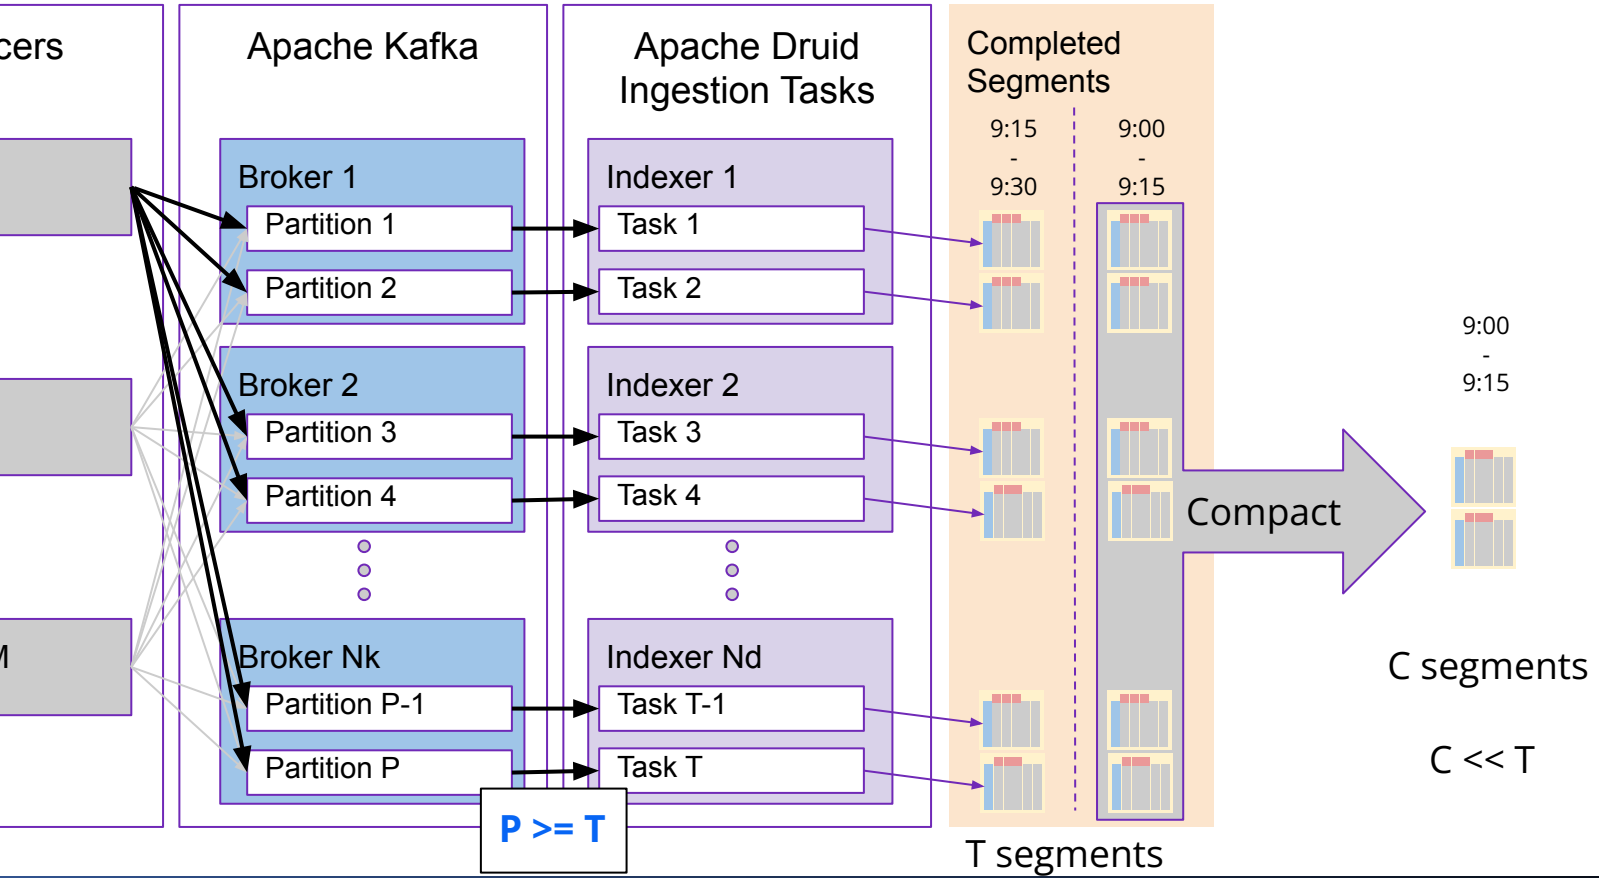

Data Enhancement in the Pipeline - Kafka SQL

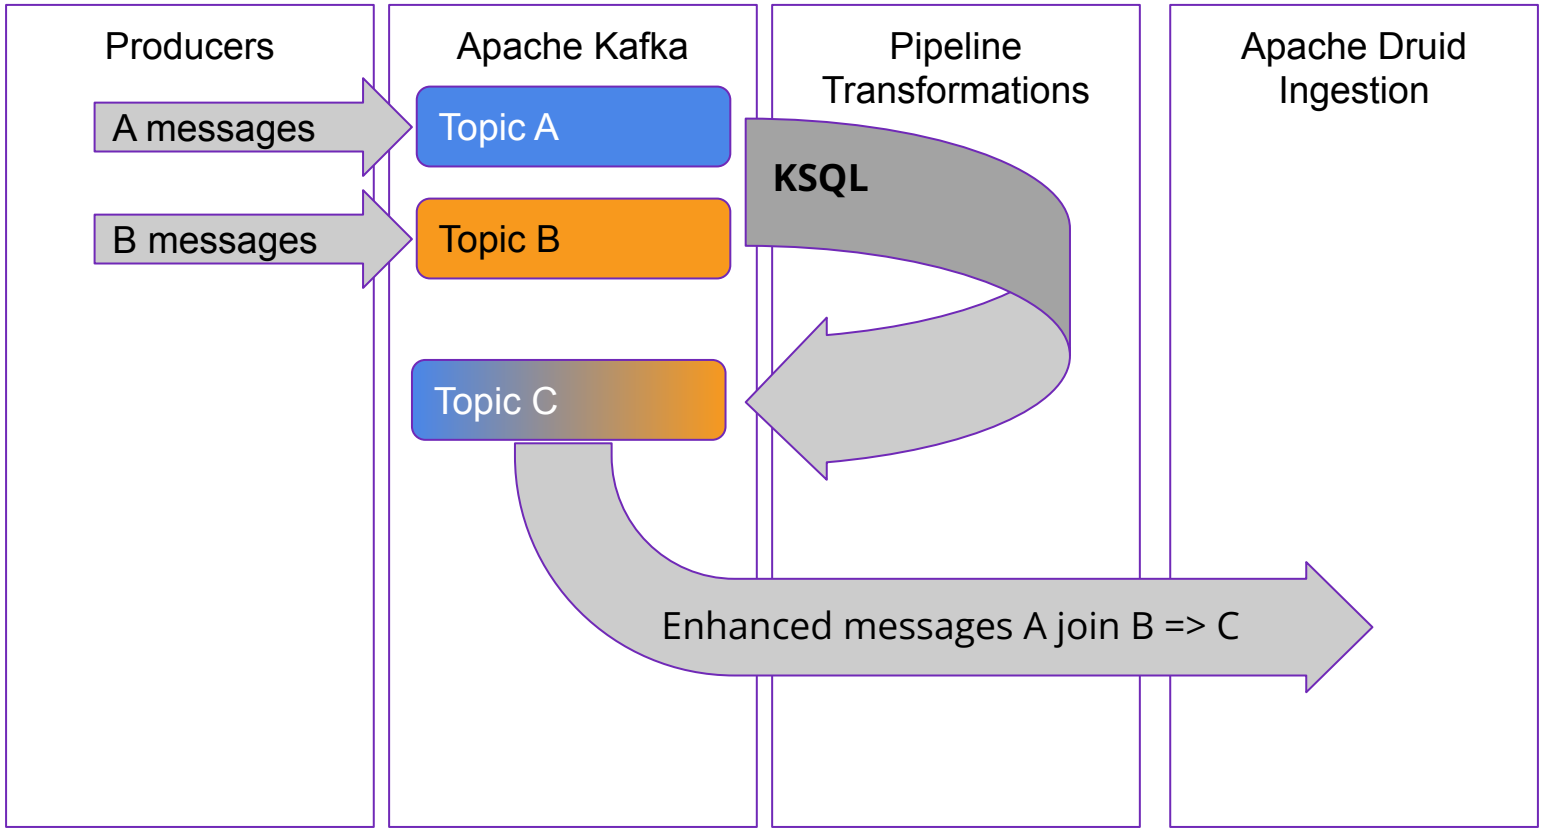

Data Enhancement in the Pipeline - Other Tech

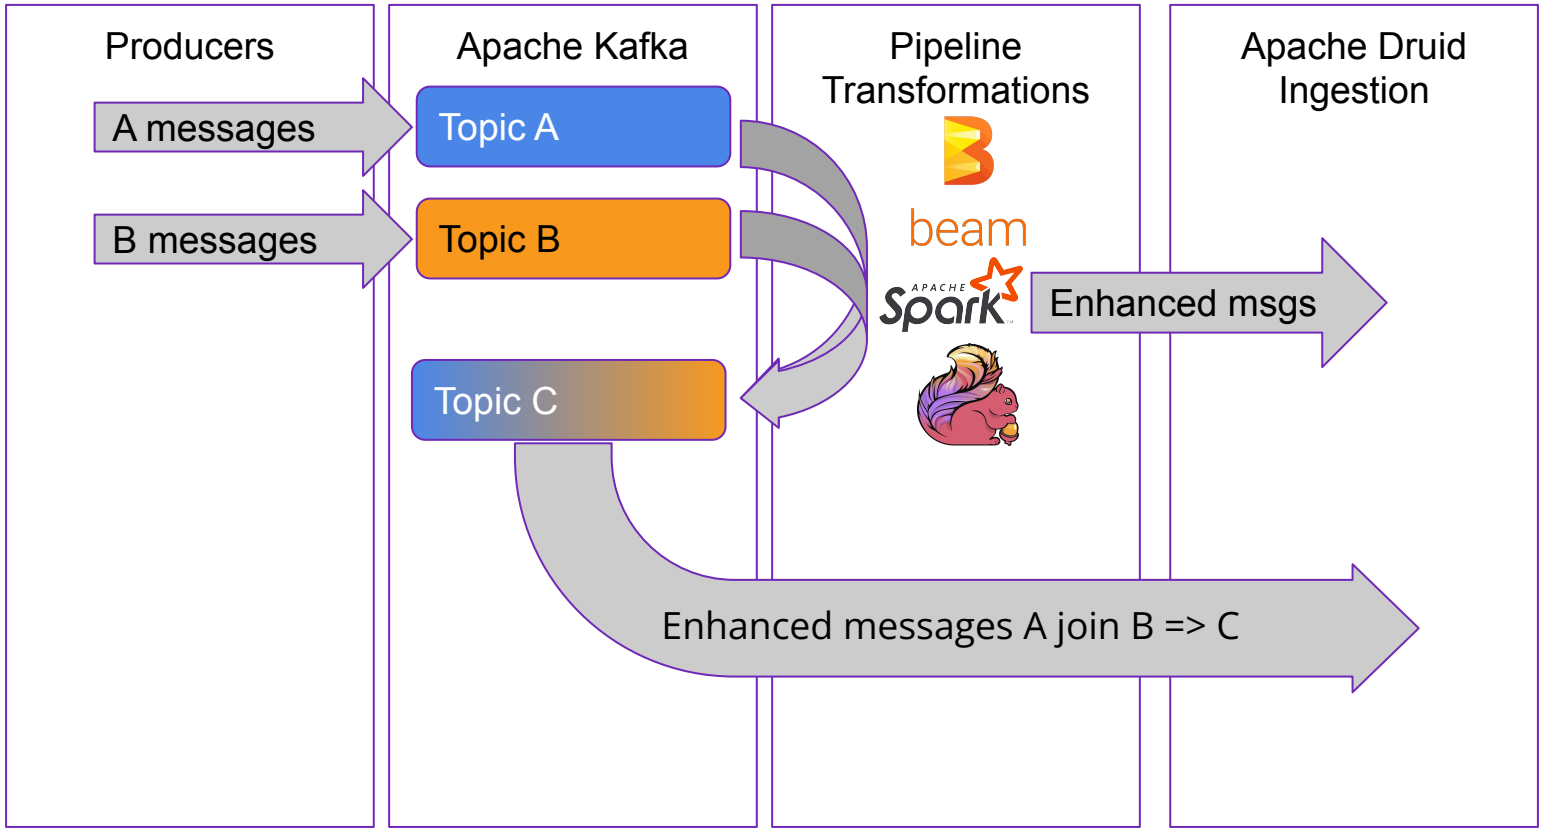

Kafka + Druid = Immediate Intelligence

Apache Druid
- Use Cases
- Architecture
- What Makes it Fast

Kafka Integration
- Streaming at any Scale
- Data Enhancement in the Pipeline

Learn more...

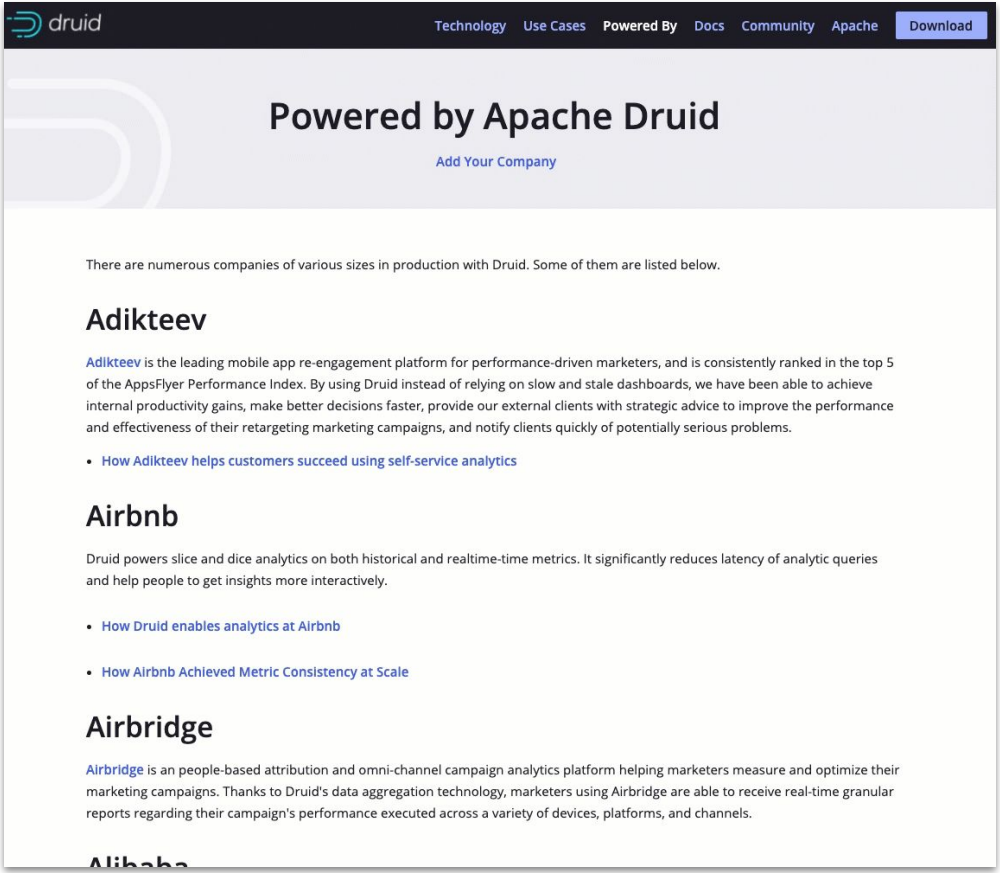
druid
Technology
Use Cases
Powered By
Docs
Community
Apache
Download
Powered by Apache Druid
Add Your Company
There are numerous companies of various sizes in production with Druid. Some of them are listed below.
Adikteev
Adikteev is the leading mobile app re-engagement platform for performance-driven marketers, and is consistently ranked in the top 5 of the AppsFlyer Performance Index. By using Druid instead of relying on slow and stale dashboards, we have been able to achieve internal productivity gains, make better decisions faster, provide our external clients with strategic advice to improve the performance and effectiveness of their retargeting marketing campaigns, and notify clients quickly of potentially serious problems.
How Adikteev helps customers succeed using self-service analytics
Airbnb
Druid powers slice and dice analytics on both historical and realtime-time metrics. It significantly reduces latency of analytic queries and help people to get insights more interactively.
How Druid enables analytics at Airbnb
How Airbnb Achieved Metric Consistency at Scale
Airbridge
Airbridge is an people-based attribution and omni-channel campaign analytics platform helping marketers measure and optimize their marketing campaigns. Thanks to Druid's data aggregation technology, marketers using Airbridge are able to receive real-time granular reports regarding their campaign's performance executed across a variety of devices, platforms, and channels.

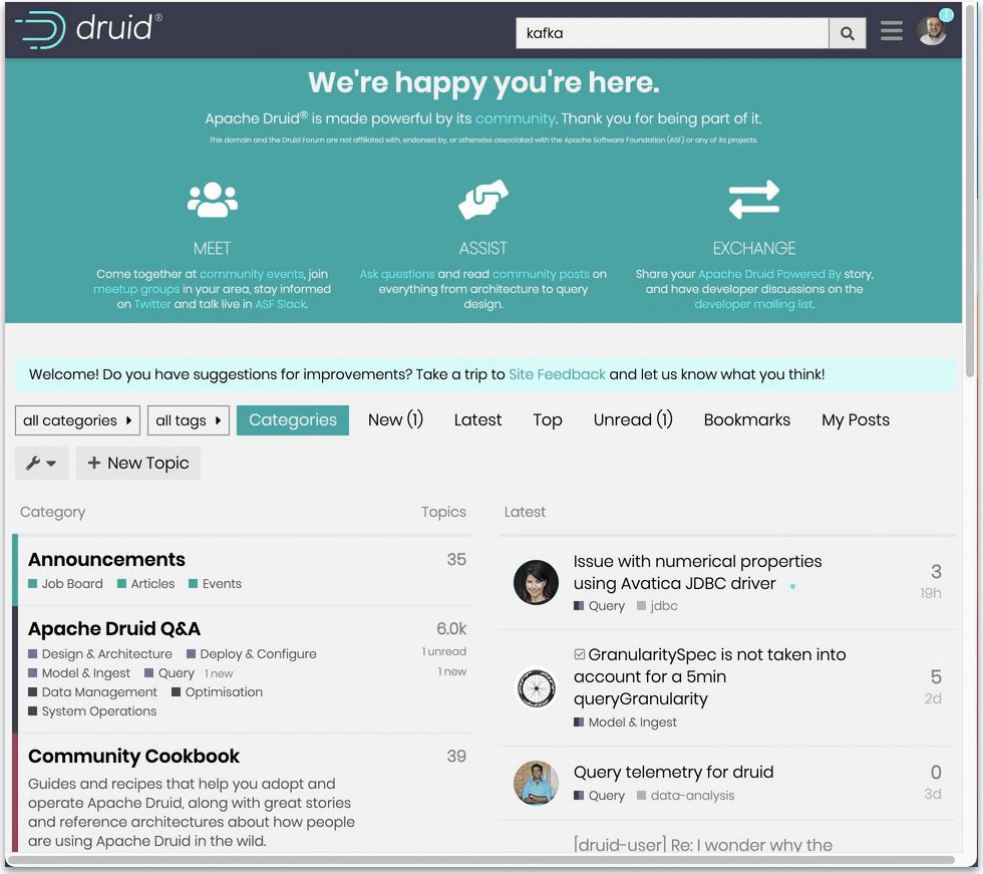

Druid Forum “Tag Browser”

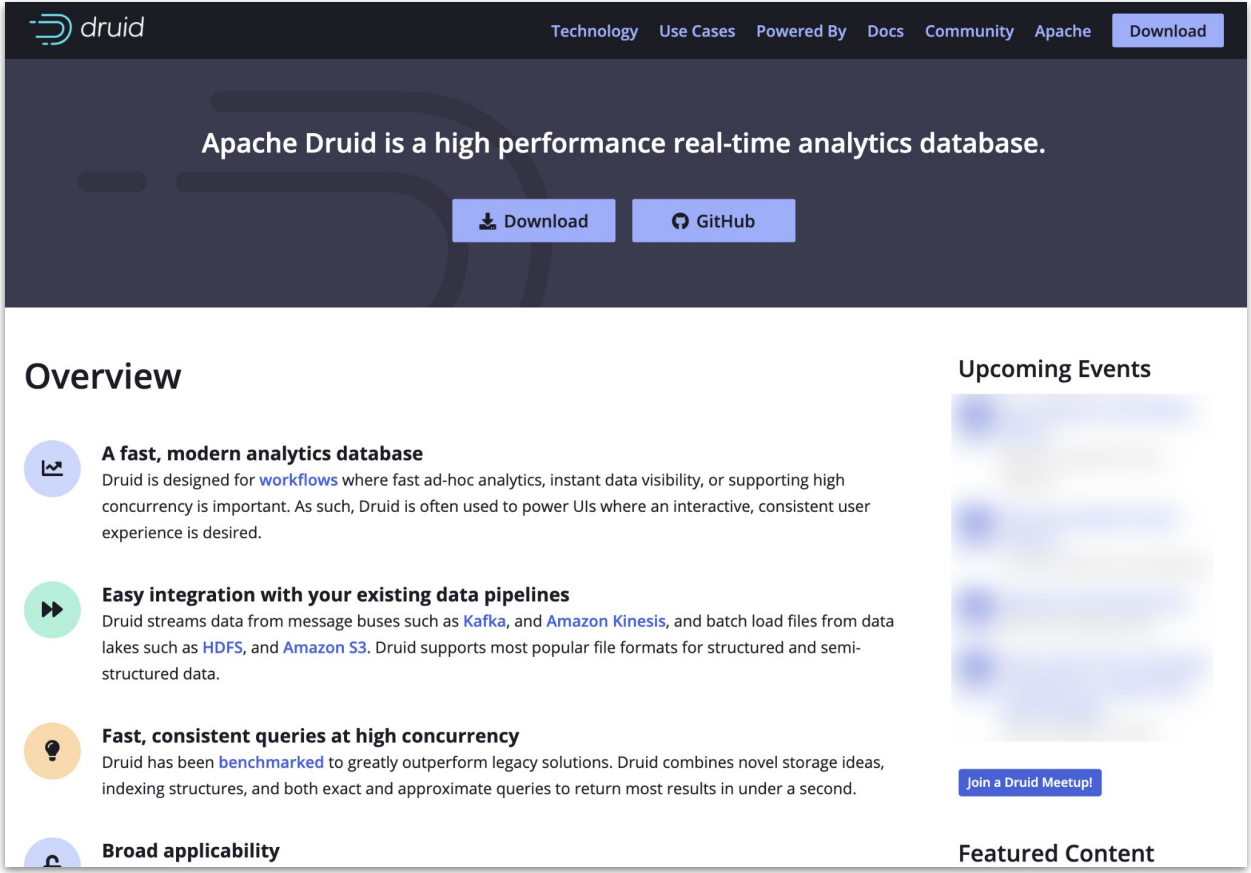
druid
Technology
Use Cases
Powered By
Docs
Community
Apache
Download
Apache Druid is a high performance real-time analytics database.
Download
GitHub
Overview
A fast, modern analytics database
Druid is designed for workflows where fast ad-hoc analytics, instant data visibility, or supporting high concurrency is important. As such, Druid is often used to power UIs where an interactive, consistent user experience is desired.
Easy integration with your existing data pipelines
Druid streams data from message buses such as Kafka, and Amazon Kinesis, and batch load files from data lakes such as HDFS, and Amazon S3. Druid supports most popular file formats for structured and semi-structured data.
Fast, consistent queries at high concurrency
Druid has been benchmarked to greatly outperform legacy solutions. Druid combines novel storage ideas, indexing structures, and both exact and approximate queries to return most results in under a second.
Broad applicability
Upcoming Events
Join a Druid Meetup!
Featured Content

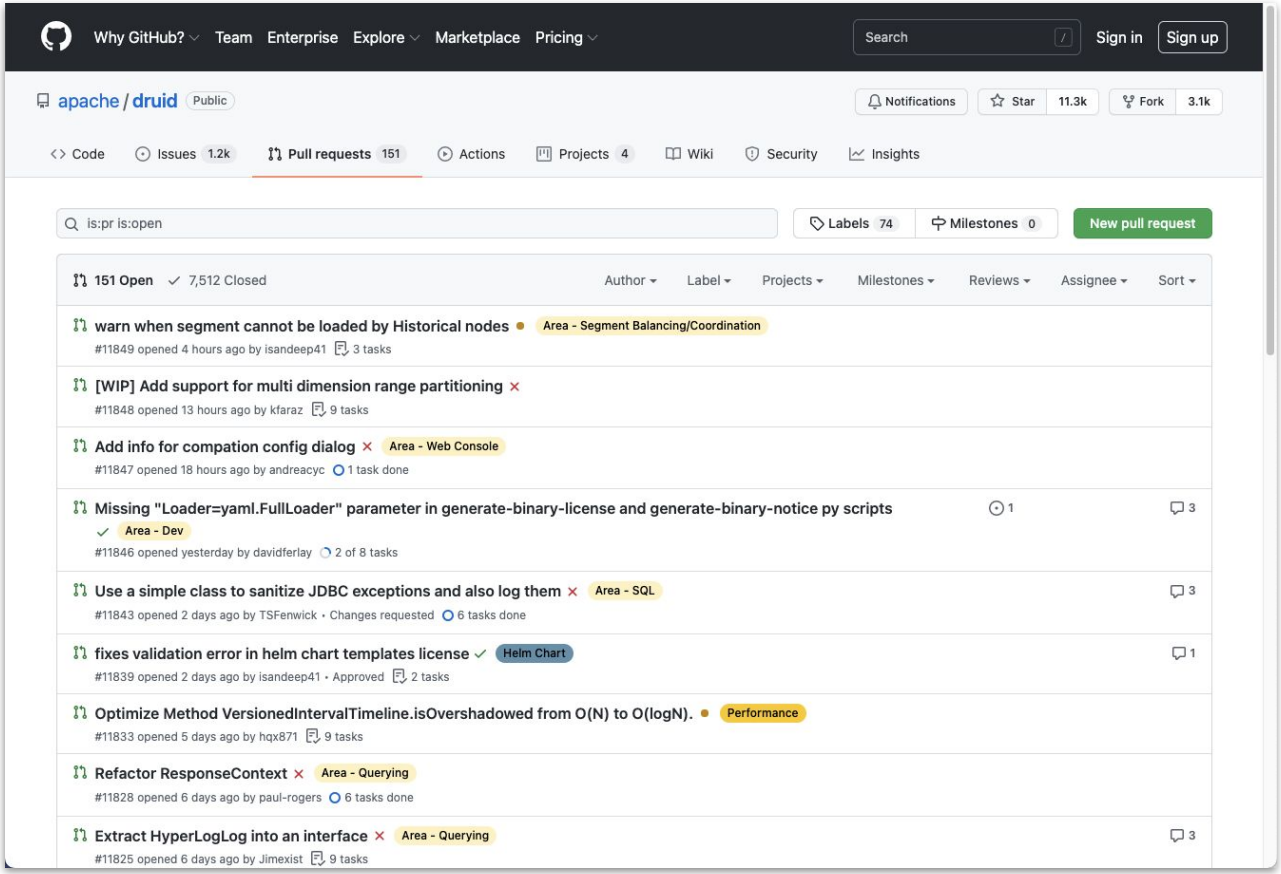
Why GitHub? Team Enterprise Explore Marketplace Pricing
Search
Sign in
Sign up
apache / druid Public
Notifications
Star 11.3k
Fork 3.1k
Code
Issues 1.2k
Pull requests 151
Actions
Projects 4
Wiki
Security
Insights
is:pr is:open
Labels 74
Milestones 0
New pull request
151 Open
7,512 Closed
Author
Label
Projects
Milestones
Reviews
Assignee
Sort
warn when segment cannot be loaded by Historical nodes
Area - Segment Balancing/Coordination
#11849 opened 4 hours ago by isandeep41 3 tasks
[WIP] Add support for multi dimension range partitioning
#11848 opened 13 hours ago by kfaraz 9 tasks
Add info for compation config dialog
Area - Web Console
#11847 opened 18 hours ago by andreacyc 1 task done
Missing "Loader=yaml.FullLoader" parameter in generate-binary-license and generate-binary-notice py scripts
Area - Dev
#11846 opened yesterday by davidferlay 2 of 8 tasks
Use a simple class to sanitize JDBC exceptions and also log them
Area - SQL
#11843 opened 2 days ago by TSFenwick • Changes requested 6 tasks done
fixes validation error in helm chart templates license
Helm Chart
#11839 opened 2 days ago by isandeep41 • Approved 2 tasks
Optimize Method VersionedIntervalTimeline.isOvershadowed from O(N) to O(logN).
Performance
#11833 opened 5 days ago by hqx871 9 tasks
Refactor ResponseContext
Area - Querying
#11828 opened 6 days ago by paul-rogers 6 tasks done
Extract HyperLogLog into an interface
Area - Querying
#11825 opened 6 days ago by Jimexist 9 tasks

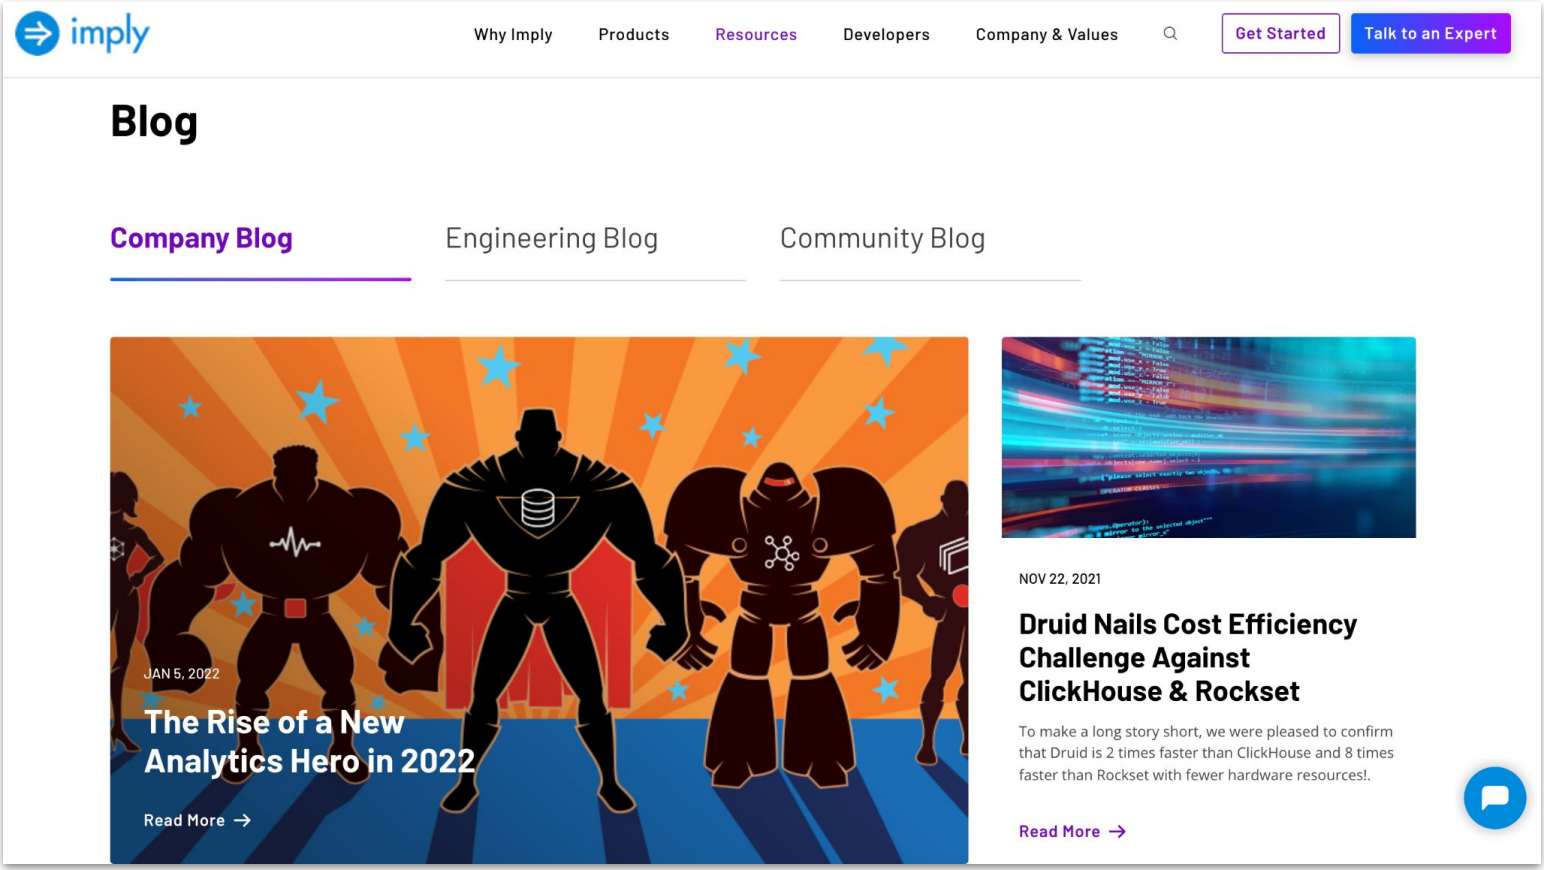

Imply blog “Community Spotlights” and “Discovery Series”

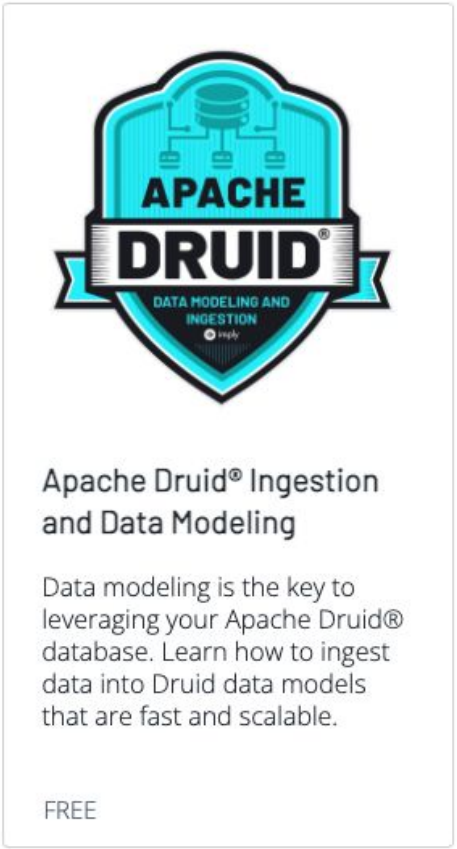

Beginner and Intermediate level lessons with accreditation exams

imply

Analytics in ***Motion***

Thank you

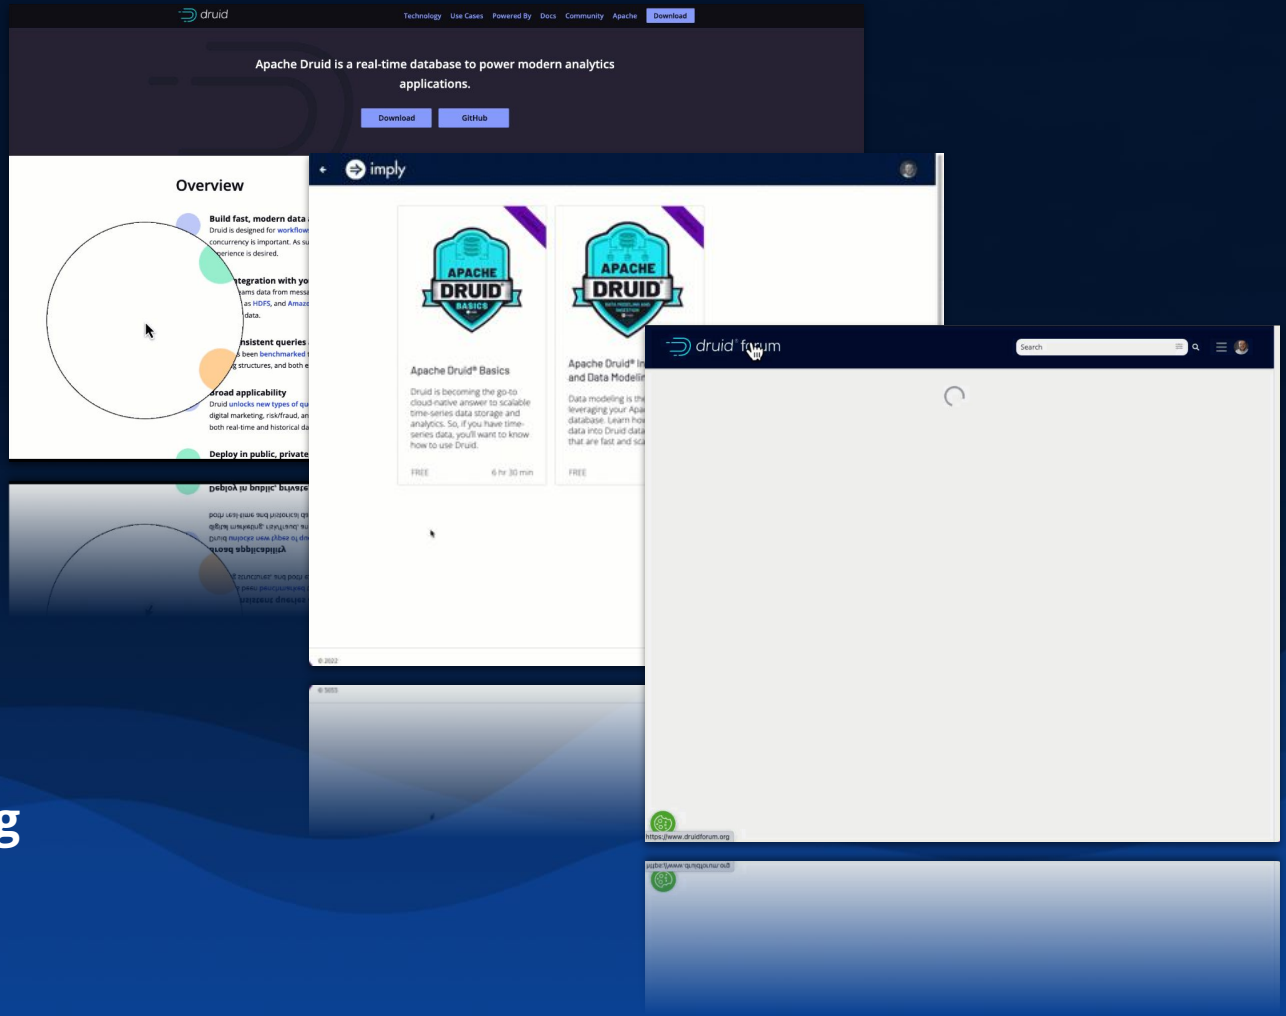

Apache Druid training and labs

learn.imply.io

Q&A, video library, and events

www.druidforum.org

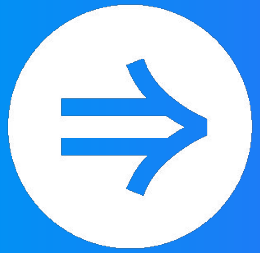

imply & druid®

Imply Community Team
community@imply.io

Imply Training Program
https://learn.imply.io

Druid User Forum by Imply
https://www.druidforum.org

Druid Professionals Group
https://www.linkedin.com/groups/8791983/

Druid Meetups
https://www.meetup.com/pro/apache-druid/

Druid News & Info
@druidio #apachedruid @implydata

Apache Druid Slack
#druid

Druid @ Imply
https://imply.io/what-is-druid/

Thank you!

Backup Slides

The hot / immediate pipeline!

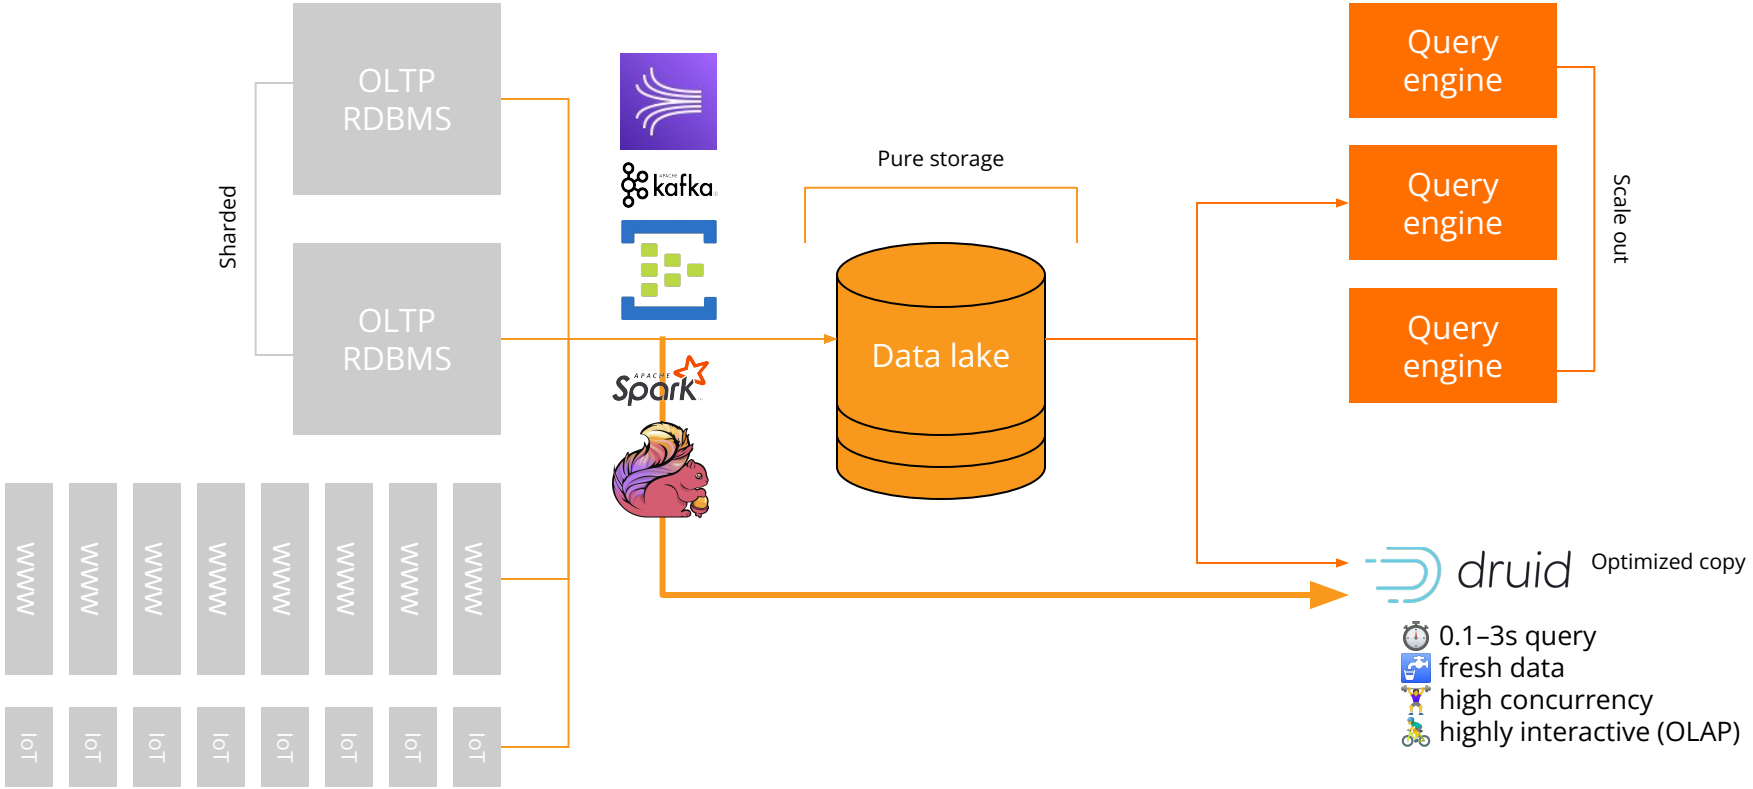

DECISIONS

DEVELOPMENT

DATA SCIENCE

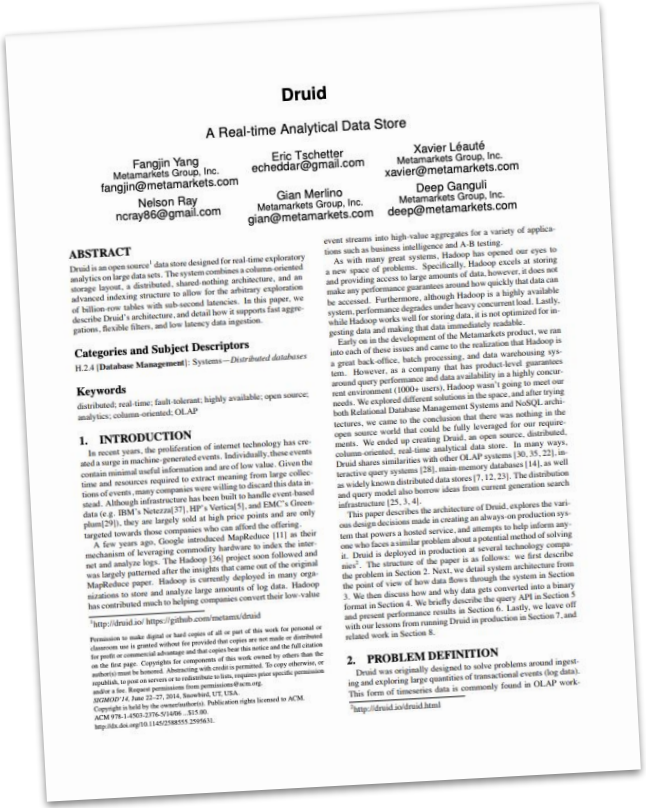

Druid

A Real-time Analytical Data Store

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

Eric Tschetter
echeddar@gmail.com

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

Nelson Ray
ncray86@gmail.com

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

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

ABSTRACT

Druid is an open source[1] 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 aggregations, flexible filters, and low latency data ingestion.

Categories and Subject Descriptors

H.2.4 [**Database Management**]: Systems—*Distributed databases*

Keywords

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

1. INTRODUCTION

In recent years, the proliferation of internet technology has created a surge in machine-generated events. Individually, these events contain minimal useful information and are of low value. Given the time and resources required to extract meaning from large collections of events, many companies were willing to discard this data instead. Although infrastructure has been built to handle event-based data (e.g. IBM's Netezza[37], HP's Vertica[5], and EMC's Greenplum[29]), they are largely sold at high price points and are only targeted towards those companies who can afford the offering.

A few years ago, Google introduced MapReduce [11] as their mechanism of leveraging commodity hardware to index the internet and analyze logs. The Hadoop [36] project soon followed and was largely patterned after the insights that came out of the original MapReduce paper. Hadoop is currently deployed in many organizations to store and analyze large amounts of log data. Hadoop has contributed much to helping companies convert their low-value 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 ingesting data and making that data immediately readable.

Early on in the development of the Metamarkets 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 concurrent environment (1000+ users), Hadoop wasn't going to meet our needs. We explored different solutions in the space, and after trying both Relational Database Management Systems and NoSQL architectures, we came to the conclusion that there was nothing in the open source world that could be fully leveraged for our requirements. 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], interactive query systems [28], main-memory databases [14], as well as widely known distributed data stores [7, 12, 23]. The distribution and query model also borrow ideas from current generation search infrastructure [25, 3, 4].

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 inform anyone who faces a similar problem about a potential method of solving it. Druid is deployed in production at several technology companies[2]. 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 our lessons from running Druid in production in Section 7, and related 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). This form of timeseries data is commonly found in OLAP work-

[1] http://druid.io/ https://github.com/metamx/druid

[2] http://druid.io/druid.html

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

ACM 978-1-4503-2376-5/14/06...$15.00.
http://dx.doi.org/10.1145/2588555.2595631.

Druid: A real-time analytical data store

NETFLIX


NETFLIX
ebay
druid®
YAHOO!
PayPal
CISCO


airbnb
NETFLIX
CONDÉ NAST
dripstat
Alibaba.com
WIKIMEDIA FOUNDATION
ebay
druid®
YAHOO!
hulu
SK telecom
Time Warner Cable
criteo
PayPal
CISCO
triplelift