

Optimizing Speed and Scale of User-Facing Real-time Analytics

Using Apache Kafka and Apache Pinot

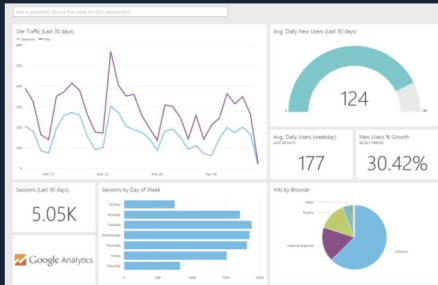
Neha Pawar Karin Wolok
@NehaPawar18 @KarinWolok

#DataAISummit

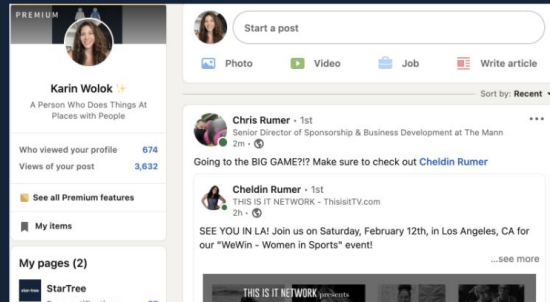
star⚡**tree**

Types of Analytics Use Cases

Dashboards / BI Tools



User-Facing Analytics



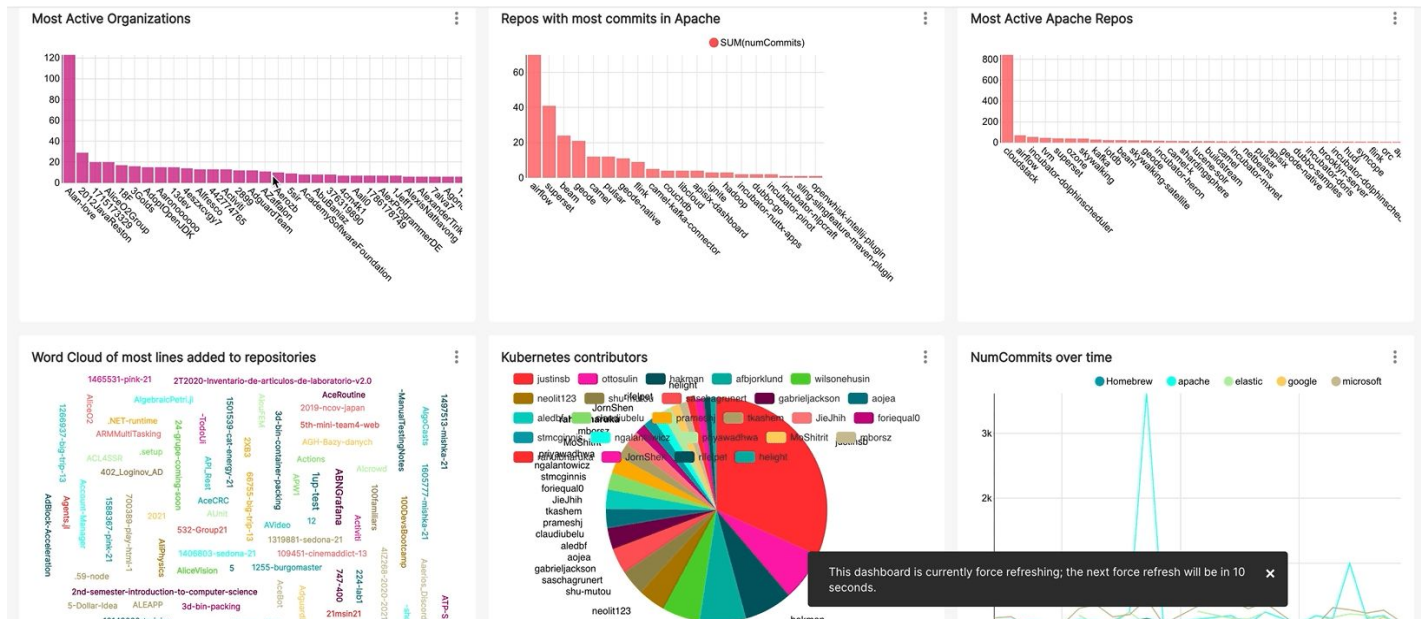
Machine Learning



Dashboards / BI Tools

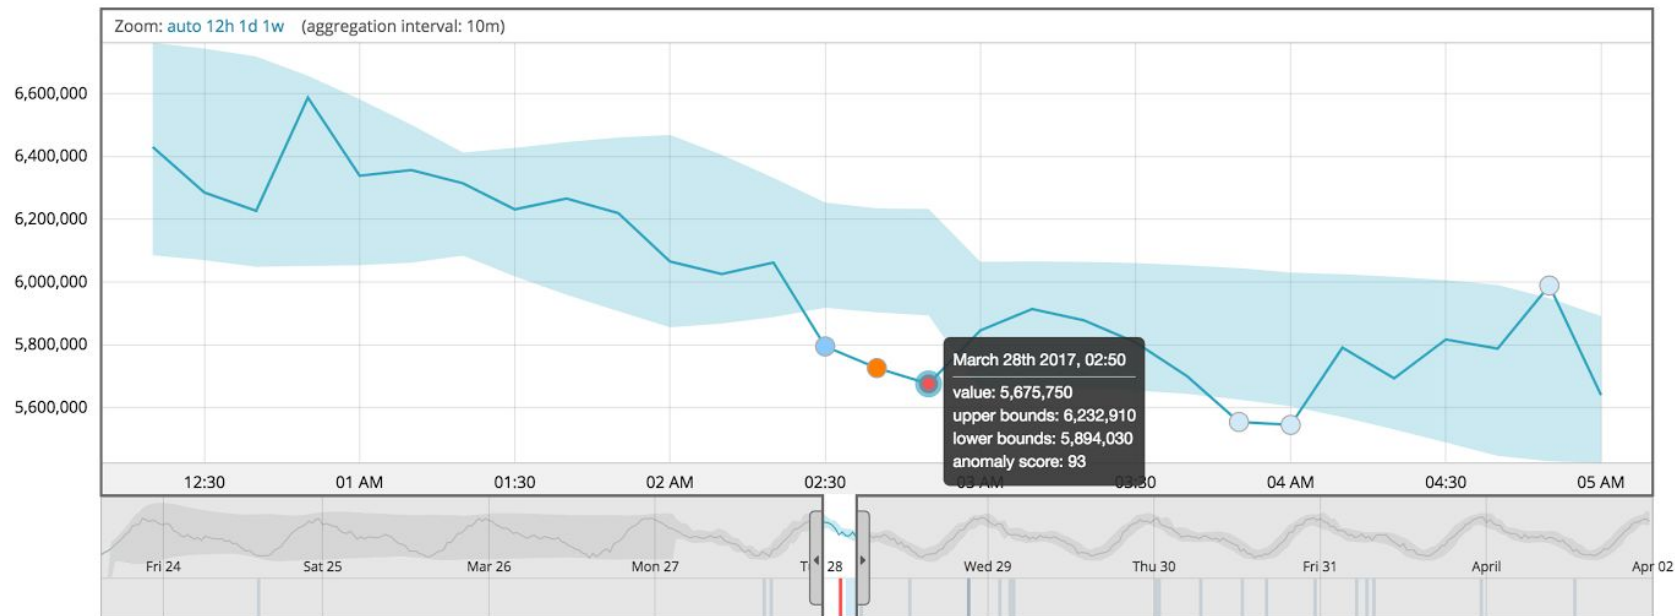
Github Analytics - Intro to Pinot Demo Draft ☆

✎ ...



This dashboard is currently force refreshing; the next force refresh will be in 10 seconds.

Machine Learning



User-Facing Analytics

The image shows a social media profile for Karin Wolok. The profile header includes a 'PREMIUM' badge, a profile picture, and the name 'Karin Wolok' with a verified status. Below the name is the bio: 'A Person Who Does Things At Places with People'. The analytics section shows 'Who viewed your profile' with 778 views and 'Impressions of your posts' with 539 impressions. There are buttons for 'See all Premium features' and 'My items'. The 'My pages (2)' section lists 'StarTree' with 173 page notifications. The post feed shows a post by Mattias Karlsson (1st) with the text 'Amazing video and thanks for joining!' and a post by Kaya Weers (2nd) with the text 'I had an amazing time speaking at Jfokus last week! For those who missed the conference, don't watch this video. You might get jealous 😊' and a video thumbnail.

PREMIUM

Karin Wolok ✨

A Person Who Does Things At Places with People

Who viewed your profile **778**

Impressions of your posts **539**

See all Premium features

My items

My pages (2)

StarTree
Page notifications **173**

Start a post

Photo Video Job Write article

Sort by: **Top**

Mattias Karlsson • 1st
Tech Leader, Community Geek & Agile Catalyst
1d • 🌐

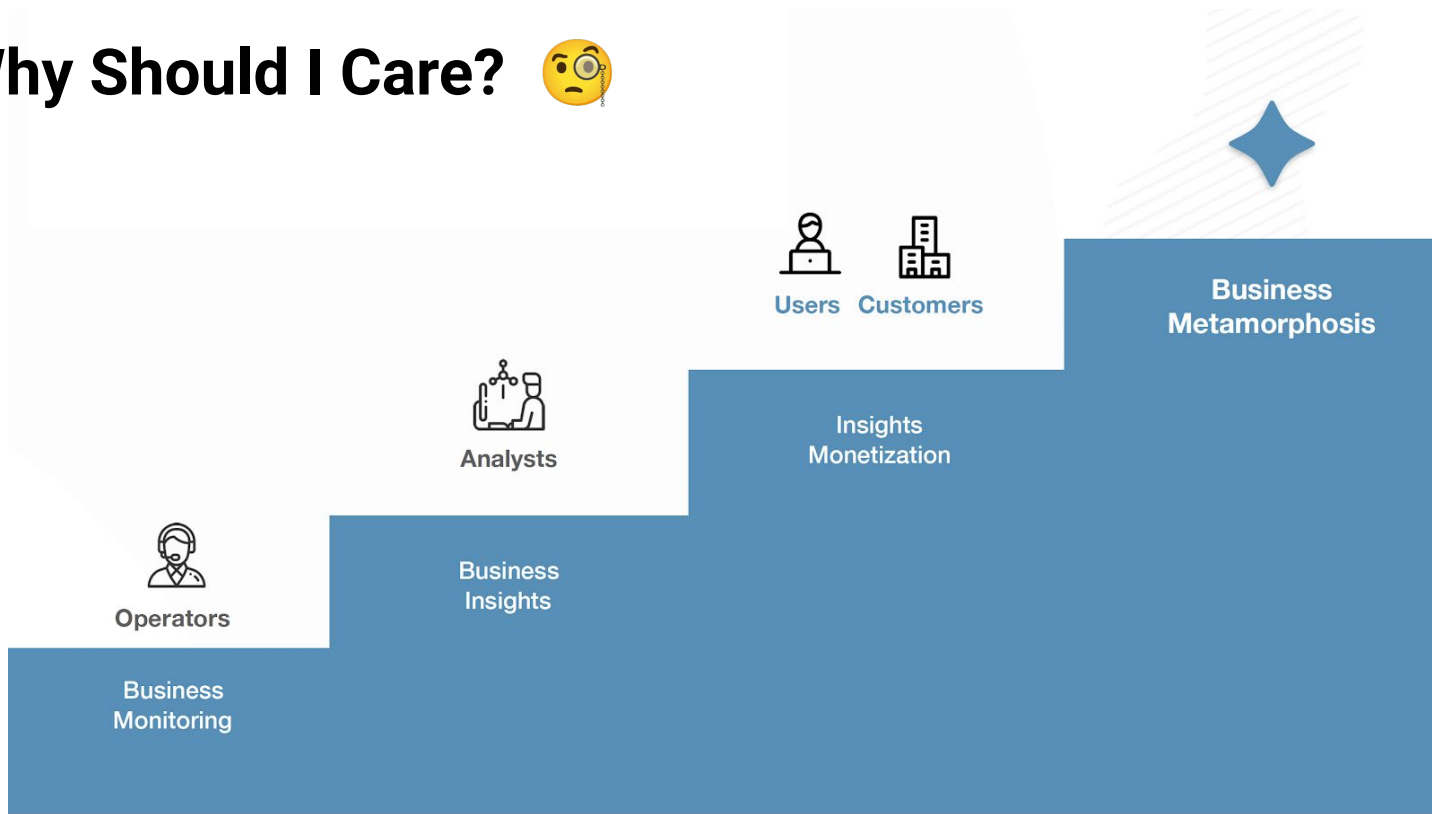
Amazing video and thanks for joining !

Kaya Weers • 2nd
Developer with a love for innovation
1d • 🌐 [+ Follow](#)

I had an amazing time speaking at **Jfokus** last week! For those who missed the conference, don't watch this video. You might get jealous 😊

...see more

Why Should I Care? 🤔



Examples of User-facing Real-Time Analytics

Talent Pool Report
266,301 professionals on LinkedIn

Showing data for: Skill: Product Development, Job title, Location, Industry: Consumer Goods

266,301 Professionals (+5% Changed jobs -17% Job posts 906 Engaged talent 21,036)

Where is this talent located?

Top locations	Professionals
Greater New York City Area	11,588
Greater Chicago Area	6,451
Greater Los Angeles Area	6,381
Indonesia (All regions)	4,375
Greater Minneapolis-St. Paul Area	4,052

Who is employing this talent?

Top companies	Professionals	1y growth	Job posts
Freshing	8,591	-1%	7
Zoomjax	7,697	-3%	3
Mintome	4,256	1%	69
Itkix	3,567	-5%	7
Techcore	3,220	0%	3

Hiring demand: High (This talent is hard to hire)

Key insights: 3 hidden gem locations, 2.8 years median tenure

Who viewed you: 128 profile viewers (Jun 16)

Not committers - Kartik Khare, Ting Chen, ...see more

Apache Pinot Committers!

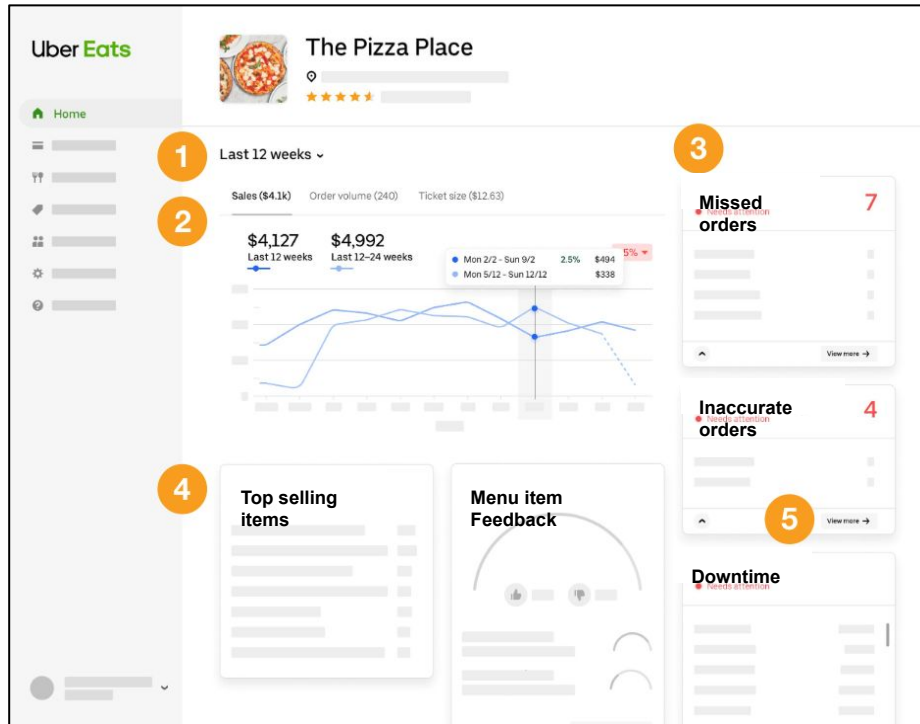
Kartik Khare, Ting Chen

Bypass deep-store for Pinot Realtime Peer download scheme Enhancements to Realtime ingestion C driver

Send

Examples of User-facing Real-Time Analytics

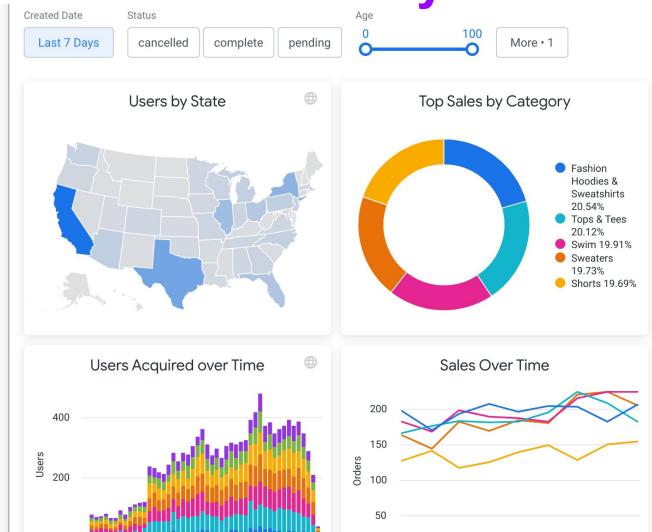
Uber Eats



Total users	500,000+
QPS	100s
Latency SLA	< 100 ms p99th
Freshness	Seconds - Minutes

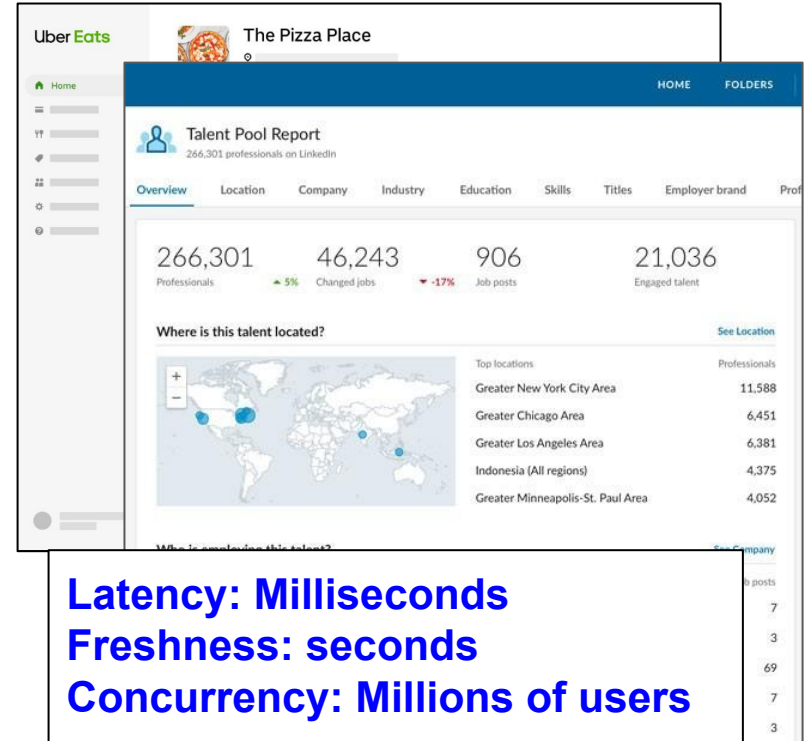
User-Facing Analytics

Internal Analytics



Latency: seconds
Freshness: seconds to minutes
Concurrency: 100s of users

External Analytics

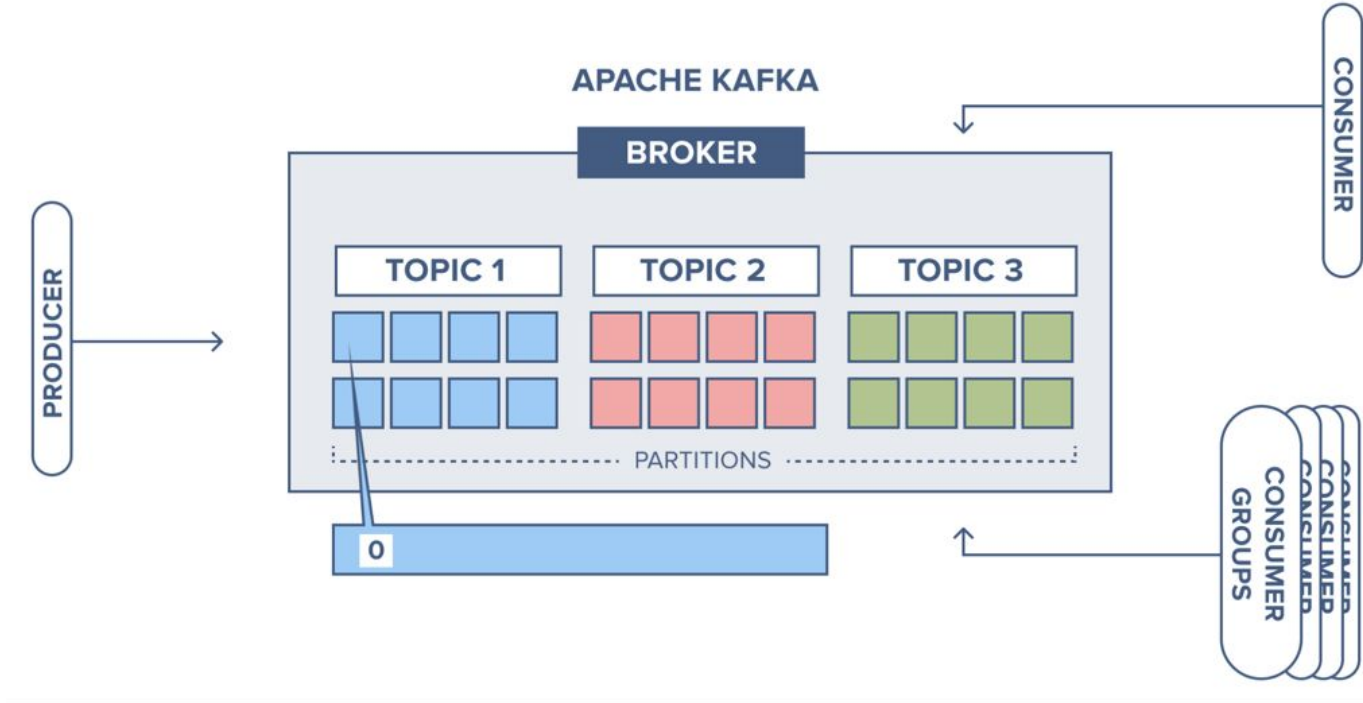


Latency: Milliseconds
Freshness: seconds
Concurrency: Millions of users


Building a User-facing Real-Time Analytics System



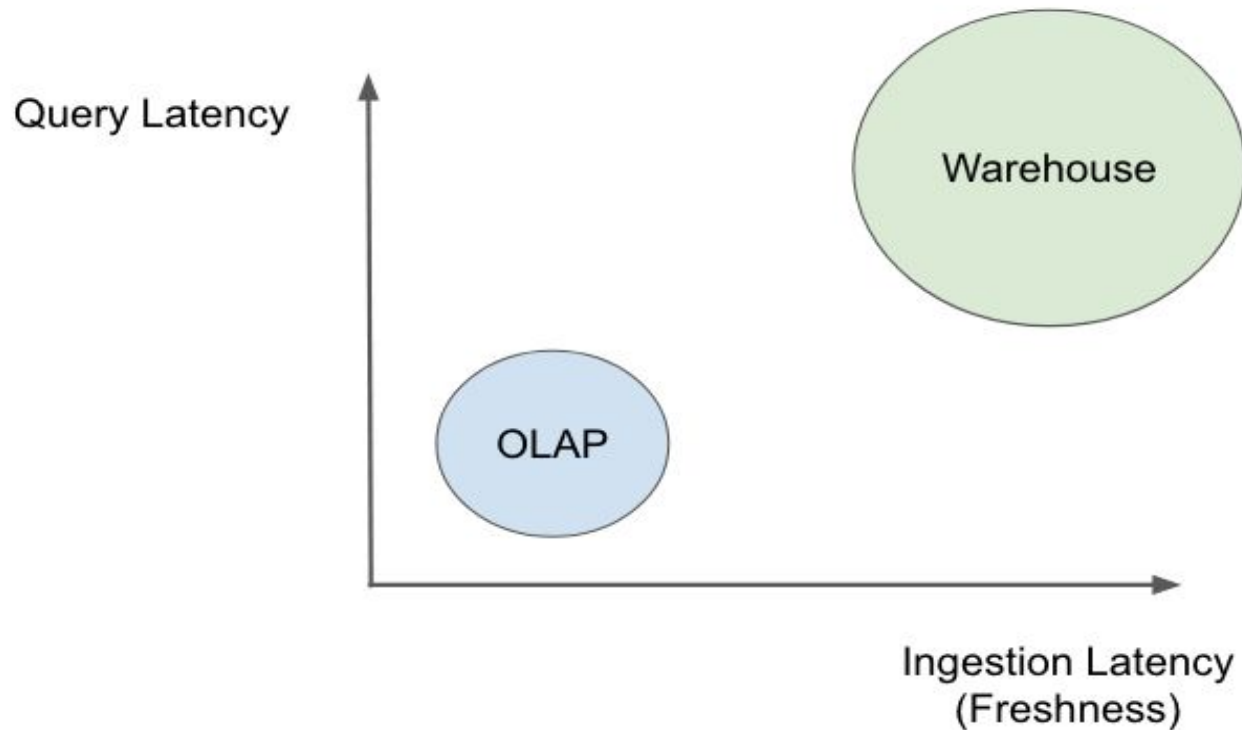
Apache Kafka



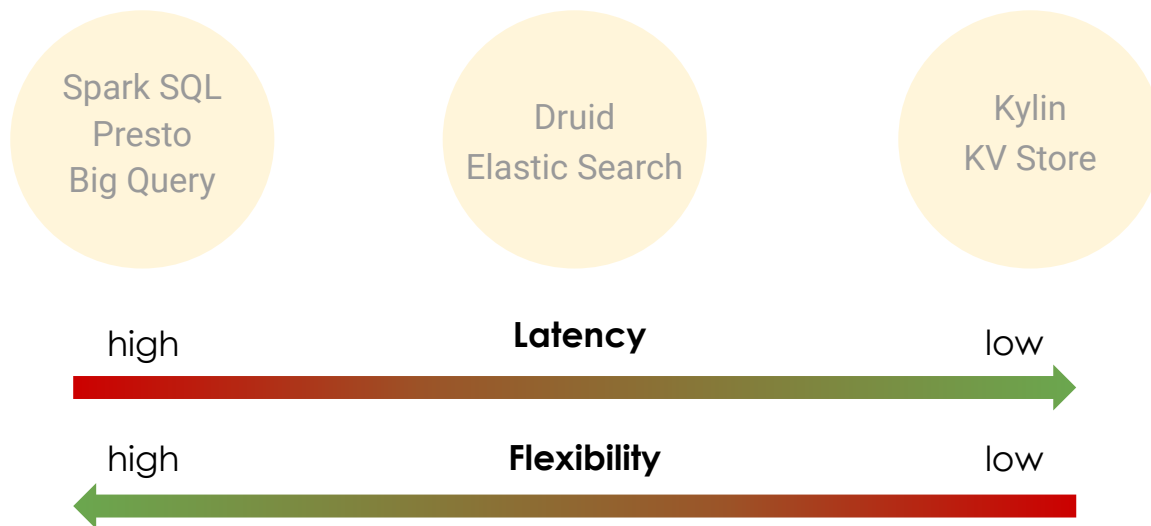
Building a User-facing Real-Time Analytics System



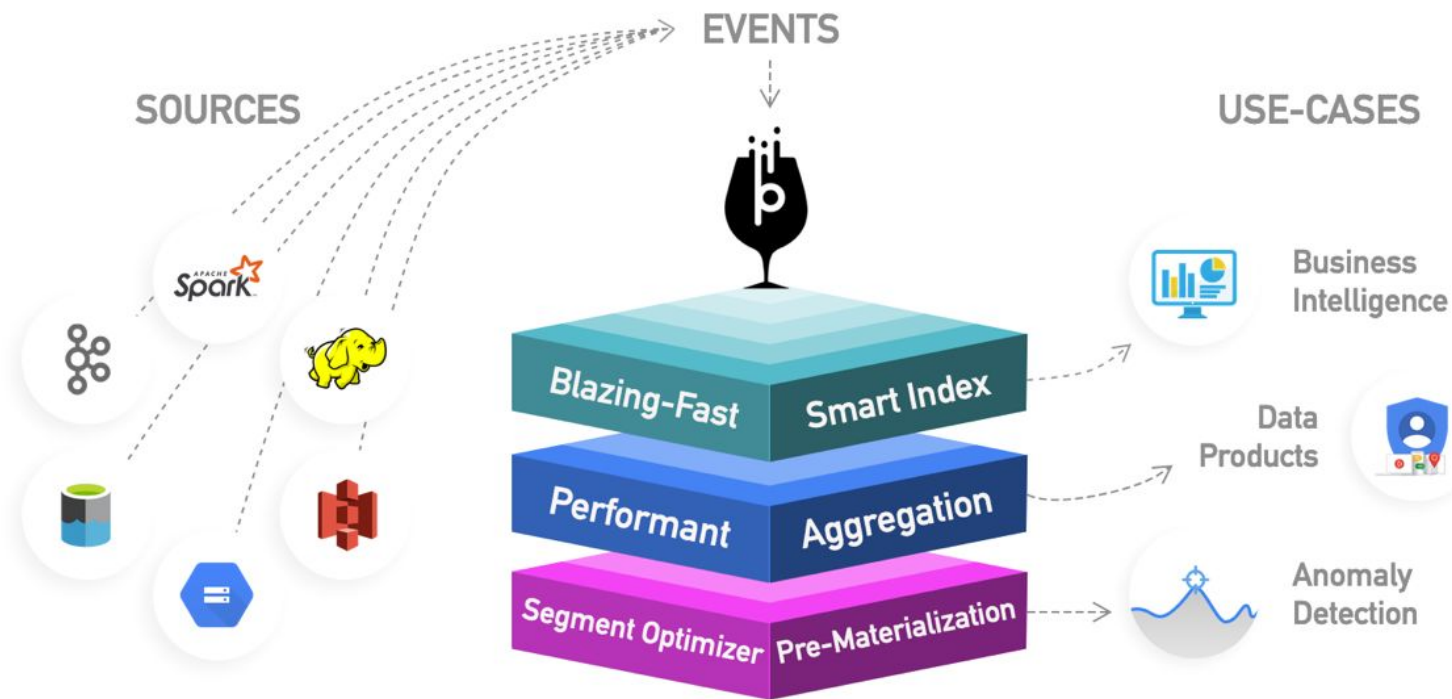
Analytics Landscape



Analytics Landscape



Apache Pinot



Apache Pinot Community

100+

Companies

stripe

wepay
a CHASE company

slack

weibo.com

confluent

factual



UBER



LinkedIn

Walmart
Save money. Live better.



Razorpay

nVIDIA

eero
an amazon company

Adbert

4k

Github Stars

star tree

Apache Pinot Community

1M+

Events/sec

stripe

wepay
a CHASE company

slack

weibo.com

confluent

factual

7 ELEVEN

TARGET

250k+

Peak QPS

CLOUD KITCHENS

UBER

Microsoft Teams

Guitar Center

LinkedIn

Walmart
Save money. Live better.

NVIDIA

eero
an amazon company

Adbert

Razorpay

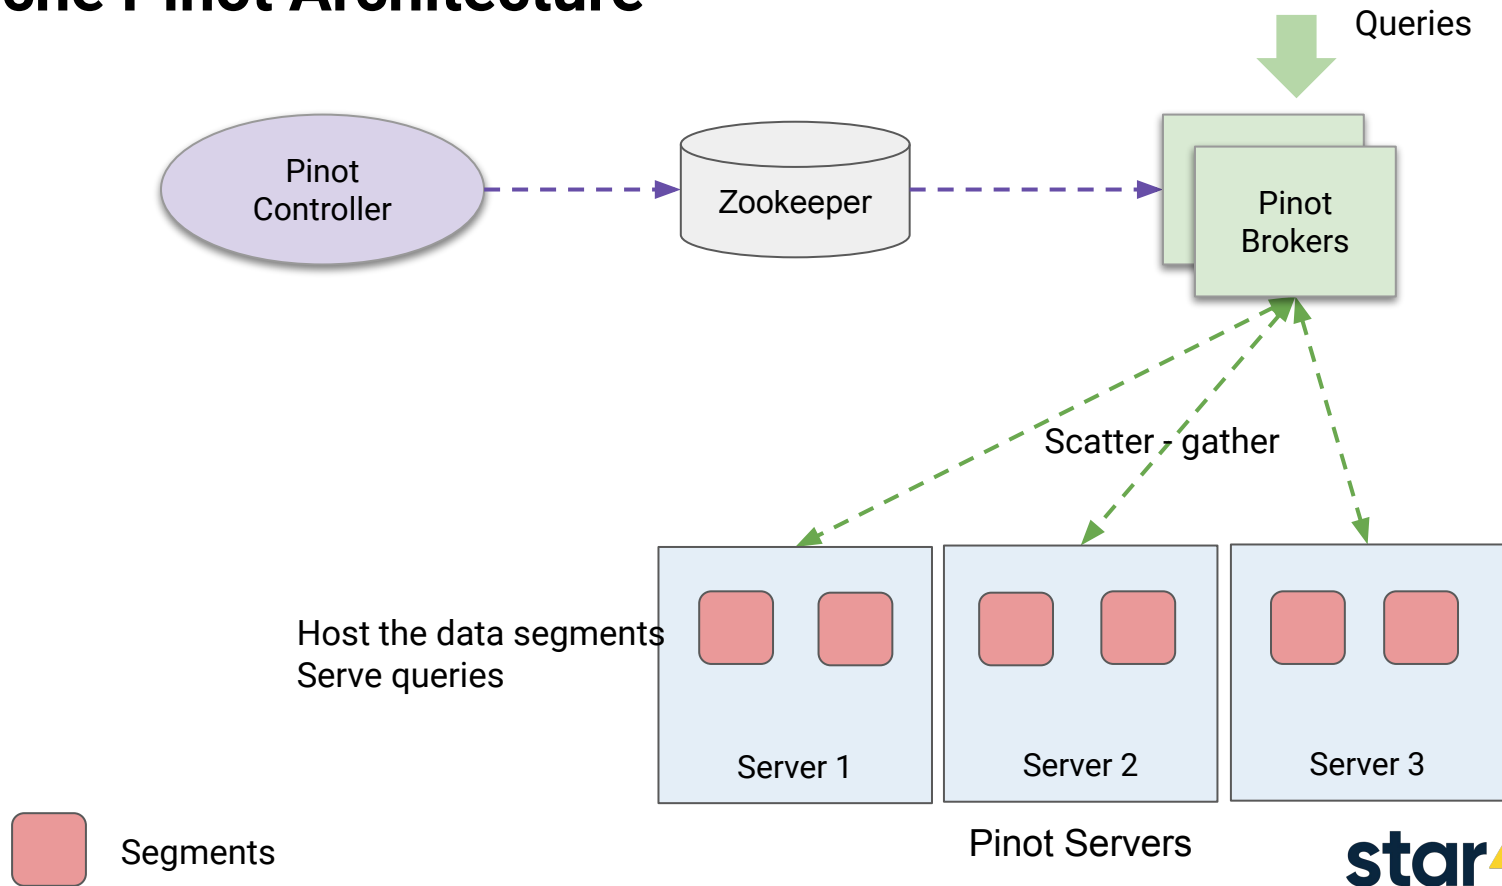
NVIDIA

ms

Query Latency

star tree

Apache Pinot Architecture





Speed of Apache Pinot

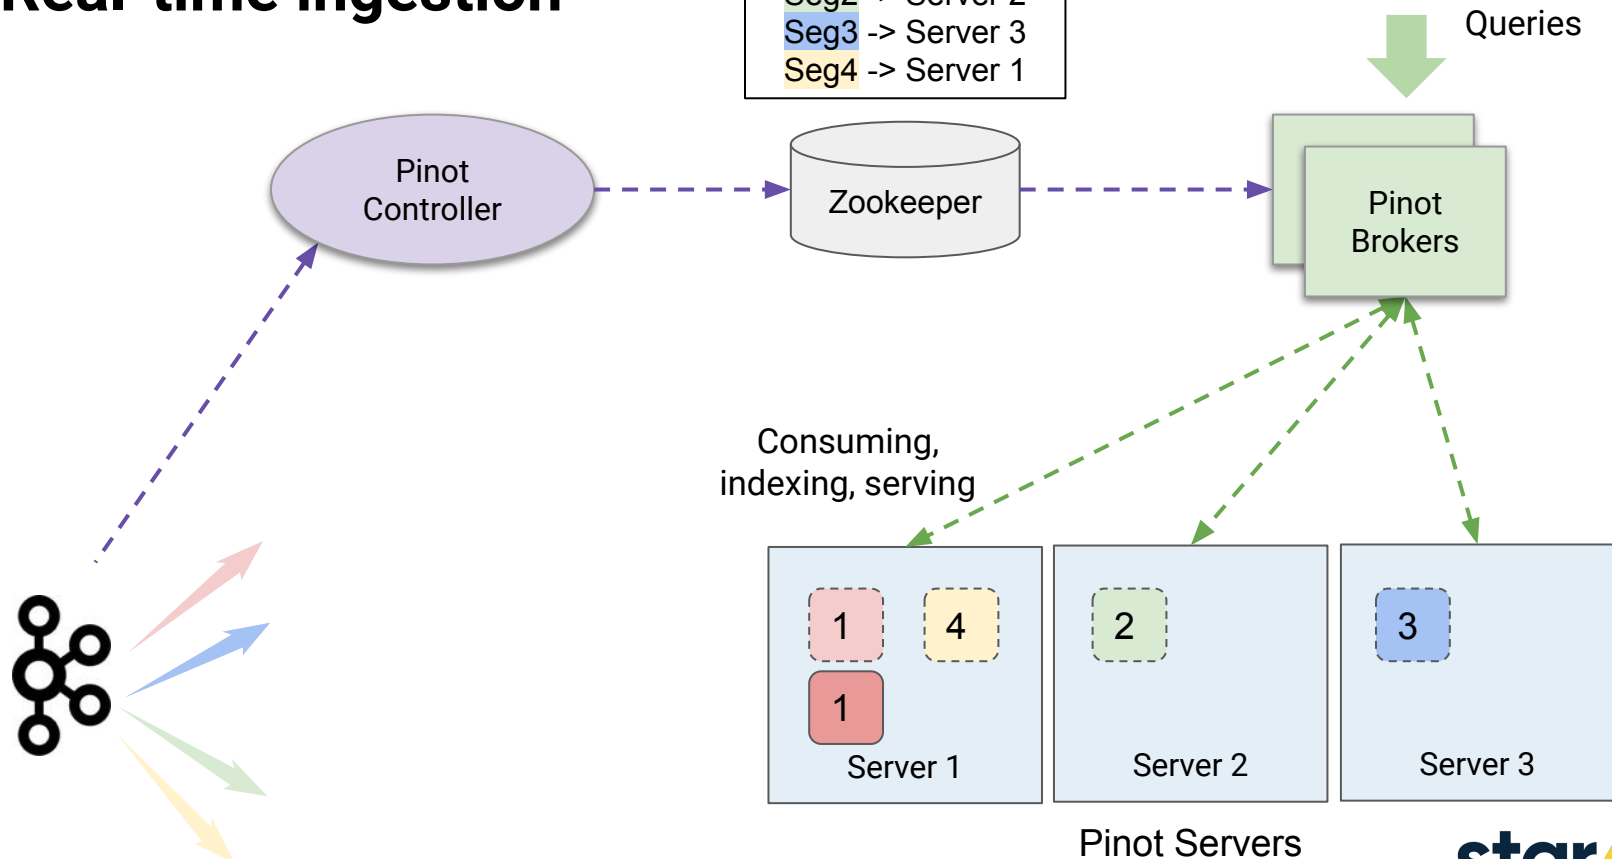
See more details in:

<https://www.startree.ai/blogs/what-makes-apache-pinot-fast-chapter-1>


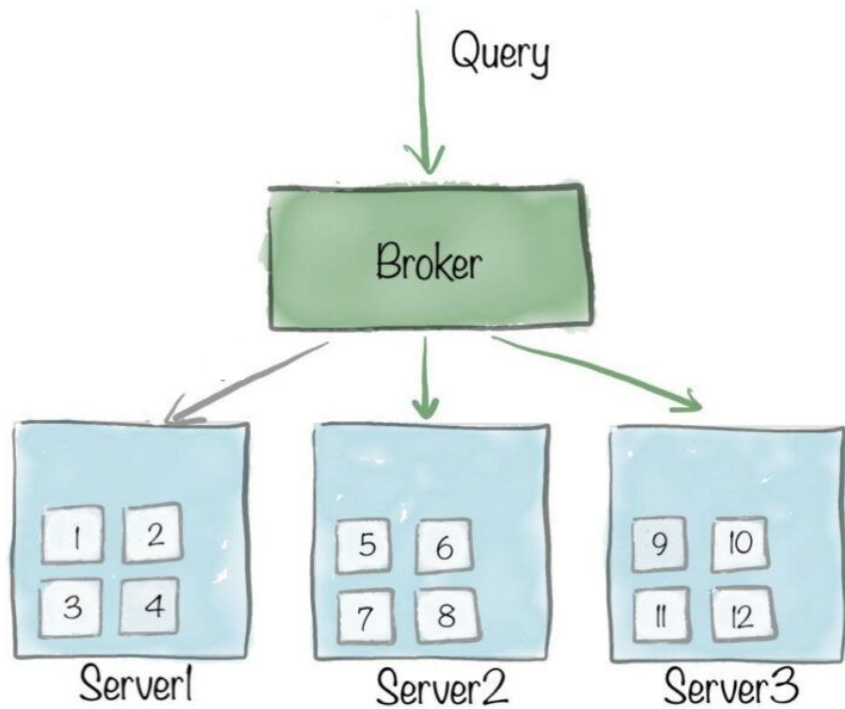
<https://www.startree.ai/blogs/what-makes-apache-pinot-fast-chapter-ii>

Real-time ingestion

Seg1 -> Server 1
Seg2 -> Server 2
Seg3 -> Server 3
Seg4 -> Server 1



Fast Queries



Segments

```
select count(*) from X  
where country = us
```

Raw Data

country	browser	...
us	chrome	...
ca	firefox	...
jp	ie	...
us	firefox	...
ca	ie	...
...



1



country
us
ca
jp
us
ca
...

browser
chrome
firefox
ie
firefox
ie
...

...
...
...
...
...
...
...

Inverted Index

```
select count(*) from X  
where country = us
```

Raw Data

country	browser	...
us	chrome	...
ca	firefox	...
jp	ie	...
us	firefox	...
ca	ie	...
...



1



country
us
ca
jp
us
ca
...

Inverted Index

country	docids
ca	1, 4...
jp	2...
us	0, 3...
...	...

Sorted Index

```
select count(*) from X  
where country = us
```

Raw Data

country	browser	...
us	chrome	...
ca	firefox	...
jp	ie	...
us	firefox	...
ca	ie	...
...



1



country
us
ca
jp
us
ca
...

Column Based

Sorted Index

country	start docid	end docid
ca	0	80
jp	81	100
us	101	300
...

Indexes: All supported index optimization in Pinot

Filtering Optimizations

Inverted Index

Sorted Index

Range Index

JSON Index

Text Index

Geo Index

StarTree Index


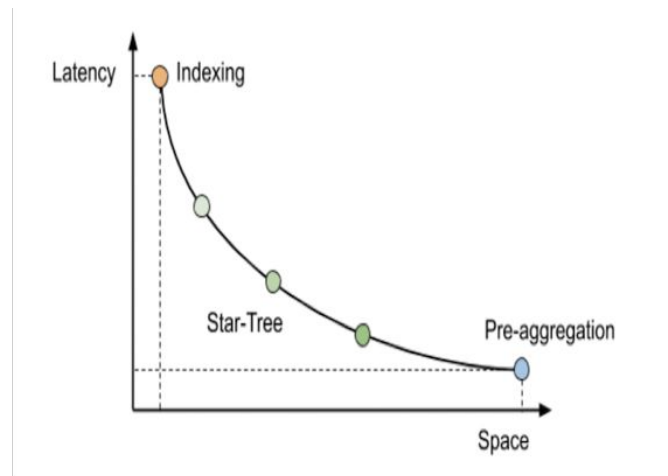
Aggregation Optimizations

Theta Sketches

HyperLogLog


Star-Tree Index

country	browser	device	os	clicks
us	chrome
ca	firefox
jp	ie
us	firefox
ca	ie
...



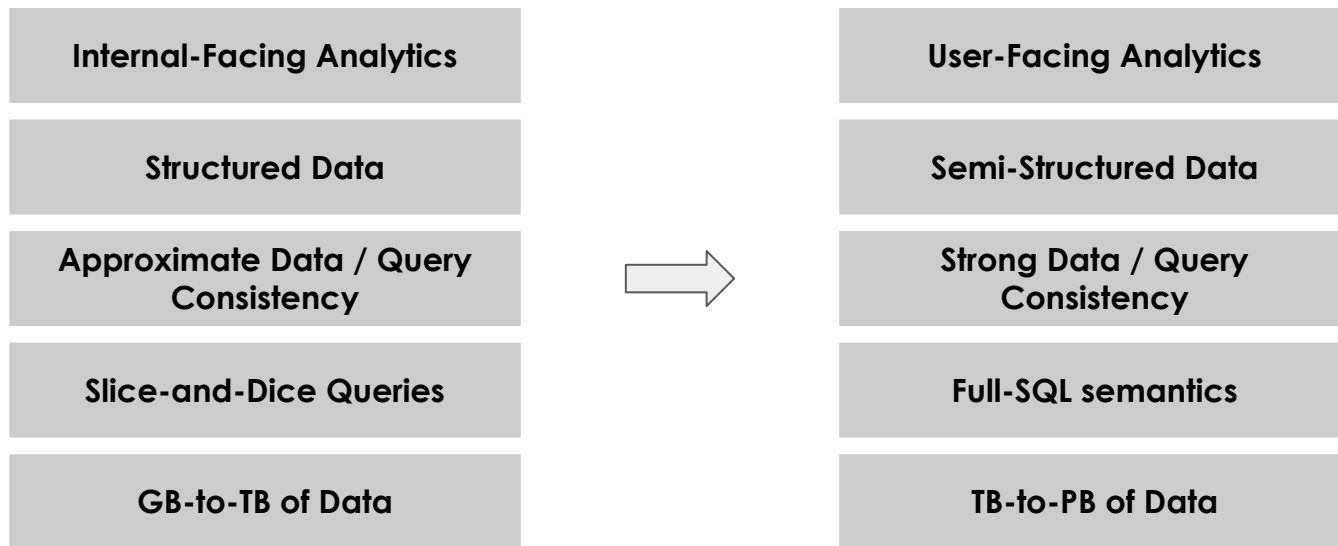
select count(*) from X
where country = us and
browser = chrome

Building a User-facing Real-Time Analytics System




Pushing the Boundaries of Real-time Analytics

- Real-Time analytics landscape is rapidly changing.
- Pinot evolving with these trends





Semi-Structured Data Support

Semi-Structured Data



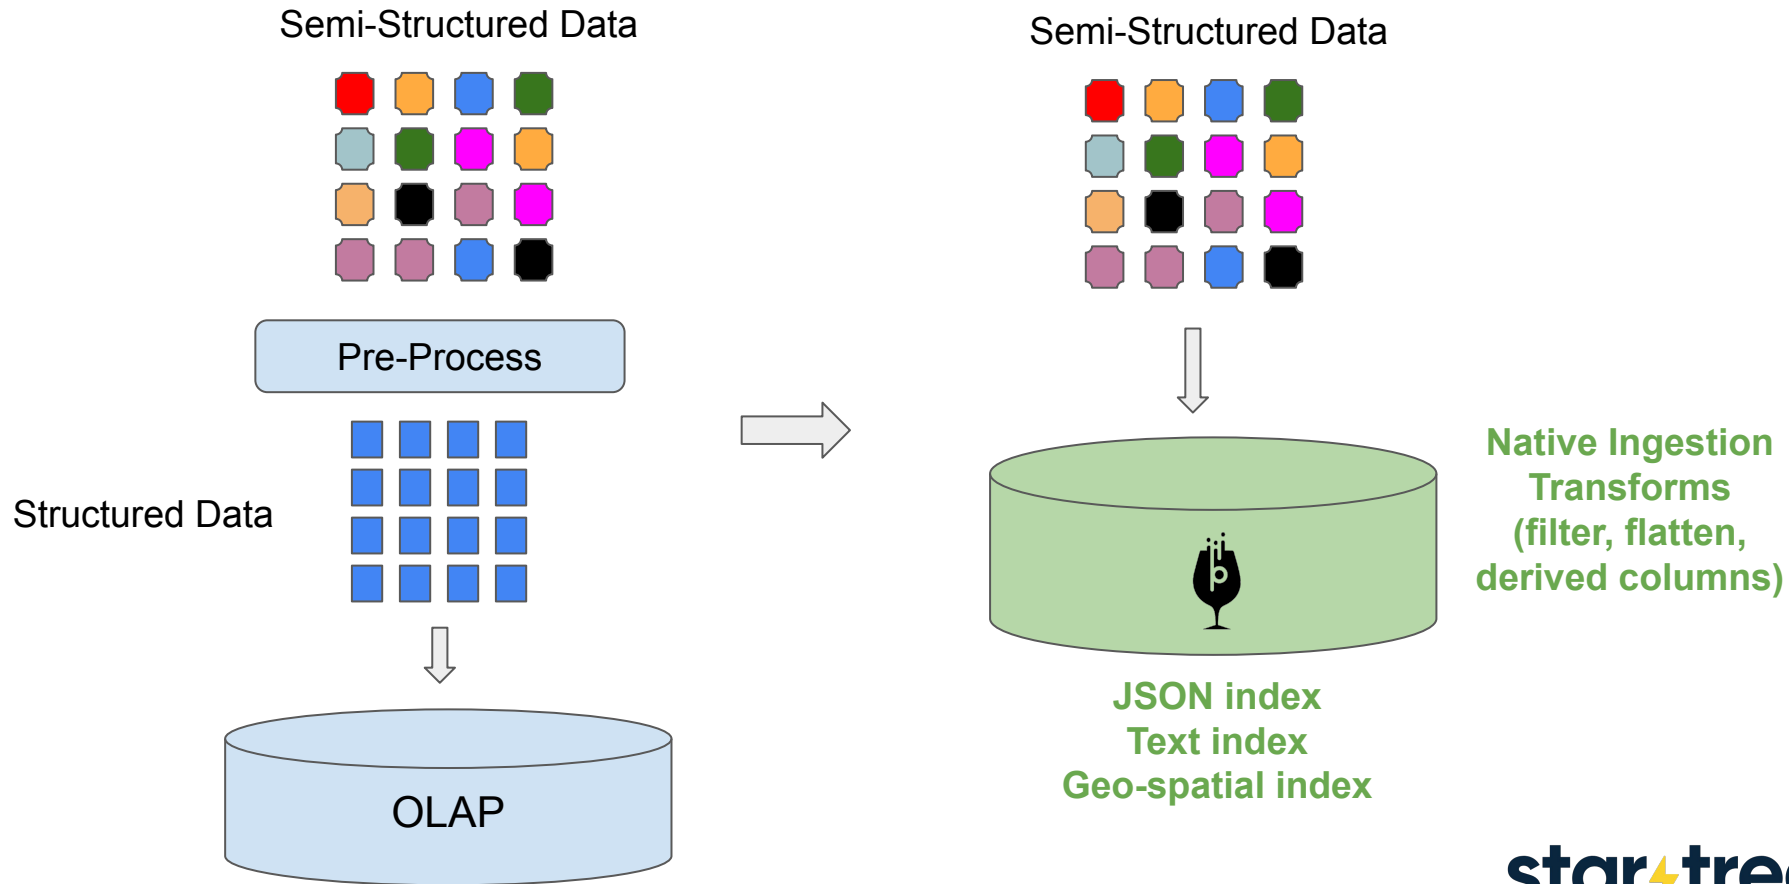
Pre-Process

Structured Data

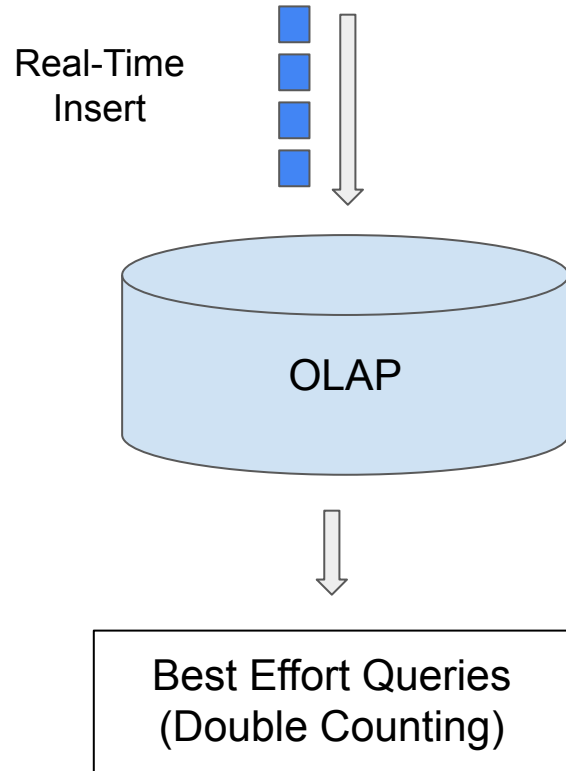


OLAP

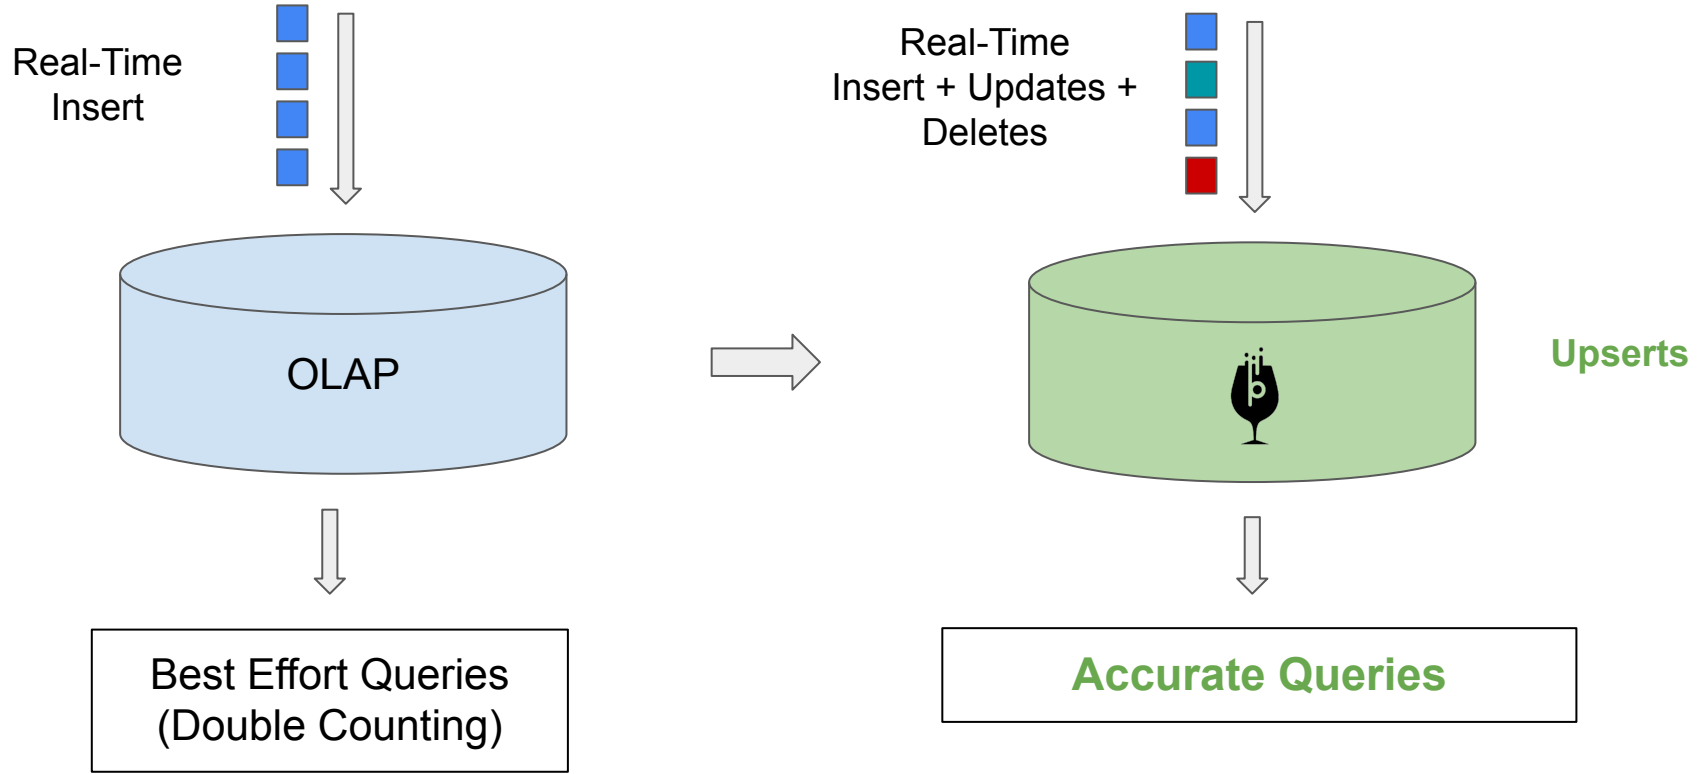
Semi-Structured Data Support




Strong Data/Query Consistency



Strong Data/Query Consistency - Upserts

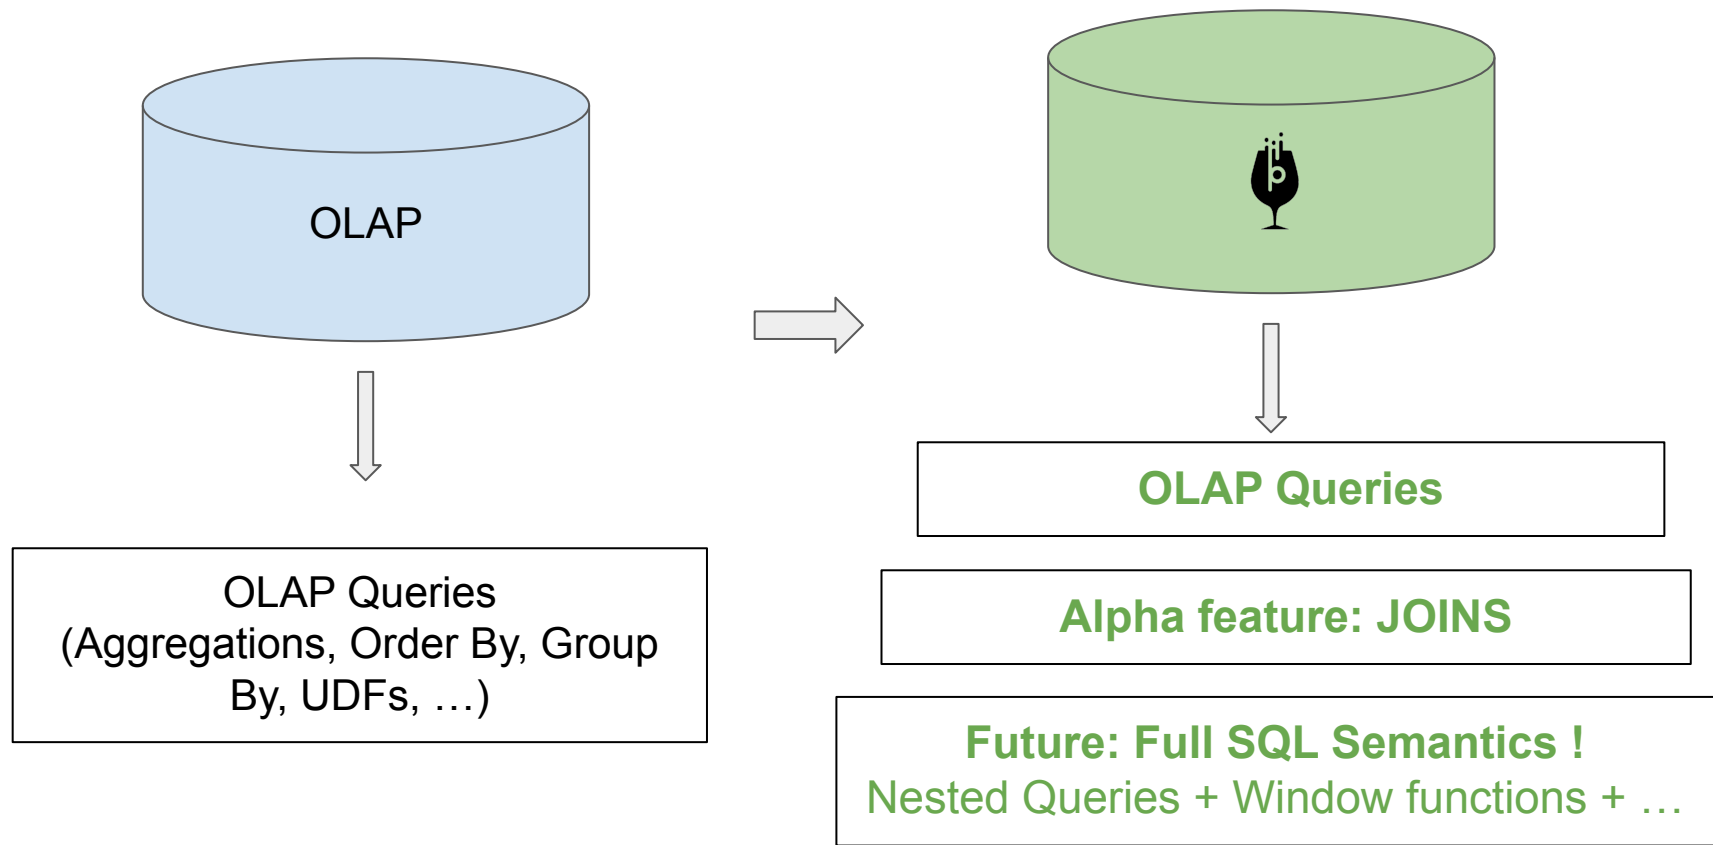


Full-SQL Semantics

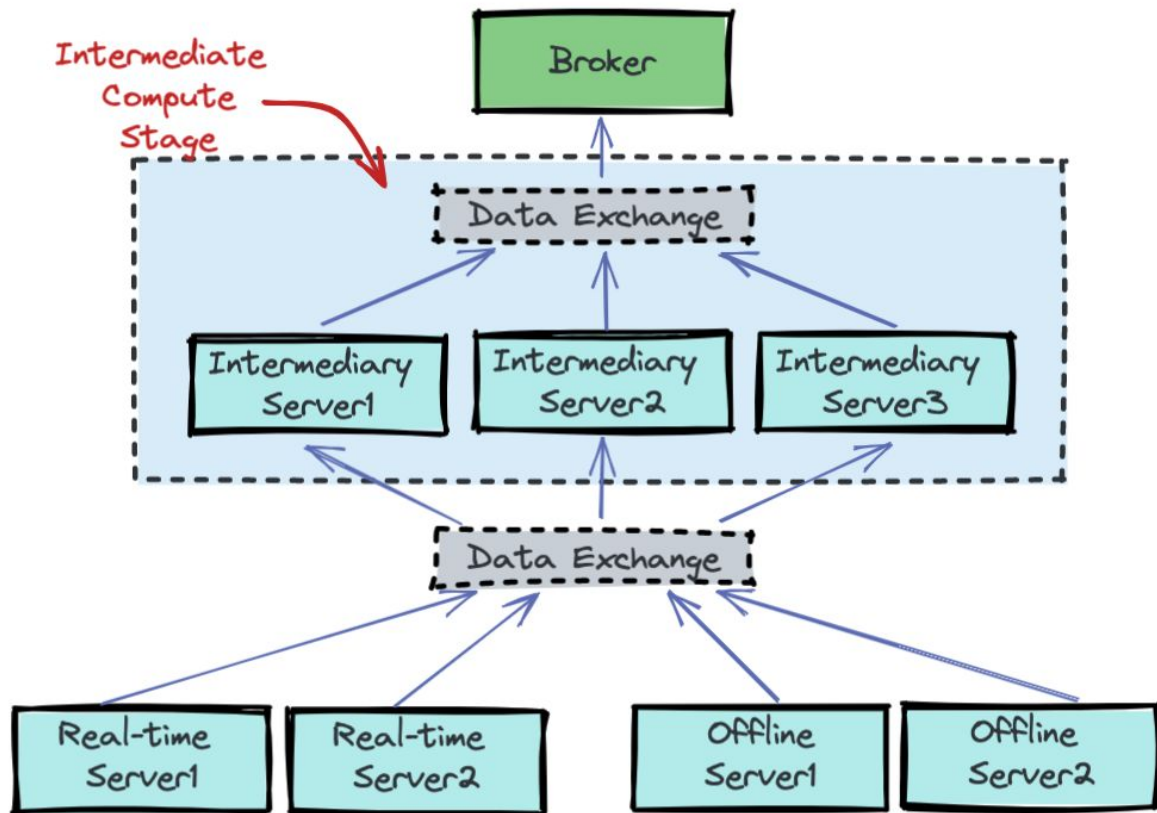


OLAP Queries
(Aggregations, Order By, Group
By, UDFs, ...)

Full-SQL Semantics

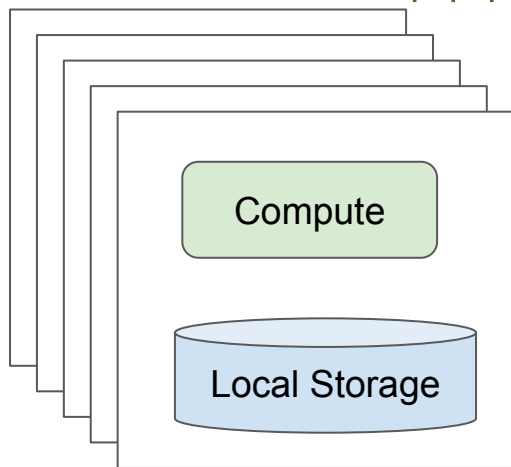


Multi-stage query execution



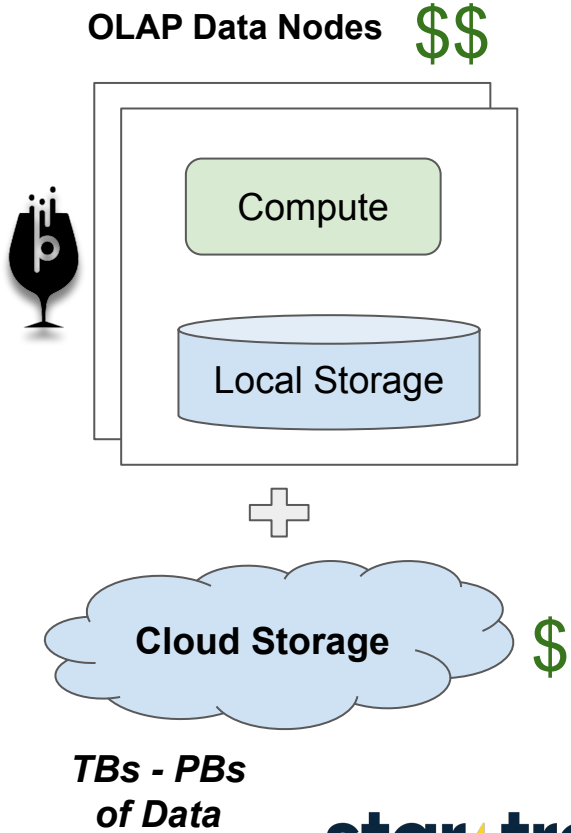

Data-Size Explosion

OLAP Data Nodes \$\$\$\$





***GBs - TBs
of Data***

Data-Size Explosion - Tiered Storage




Apache Pinot - Summary



Latency: Milliseconds
Freshness: seconds
Concurrency: Millions of users

OLAP Queries + Joins
(Single stage and multi-stage)



Thank you! 🙌



@KarinWolok
@NehaPawar18
#DataAISummit



Get Started:
dev.startree.ai

star⚡**tree**