



# REAL-TIME A IN THE CORPORATE WORLD

HOW APACHE PINOT®  
POWERS INDUSTRY LEADERS



# Who Much Pizza We Sold?

## Orders in the last 1 minute

Number of orders

948

↑ 16

Revenue in ₹

4,199,845.00

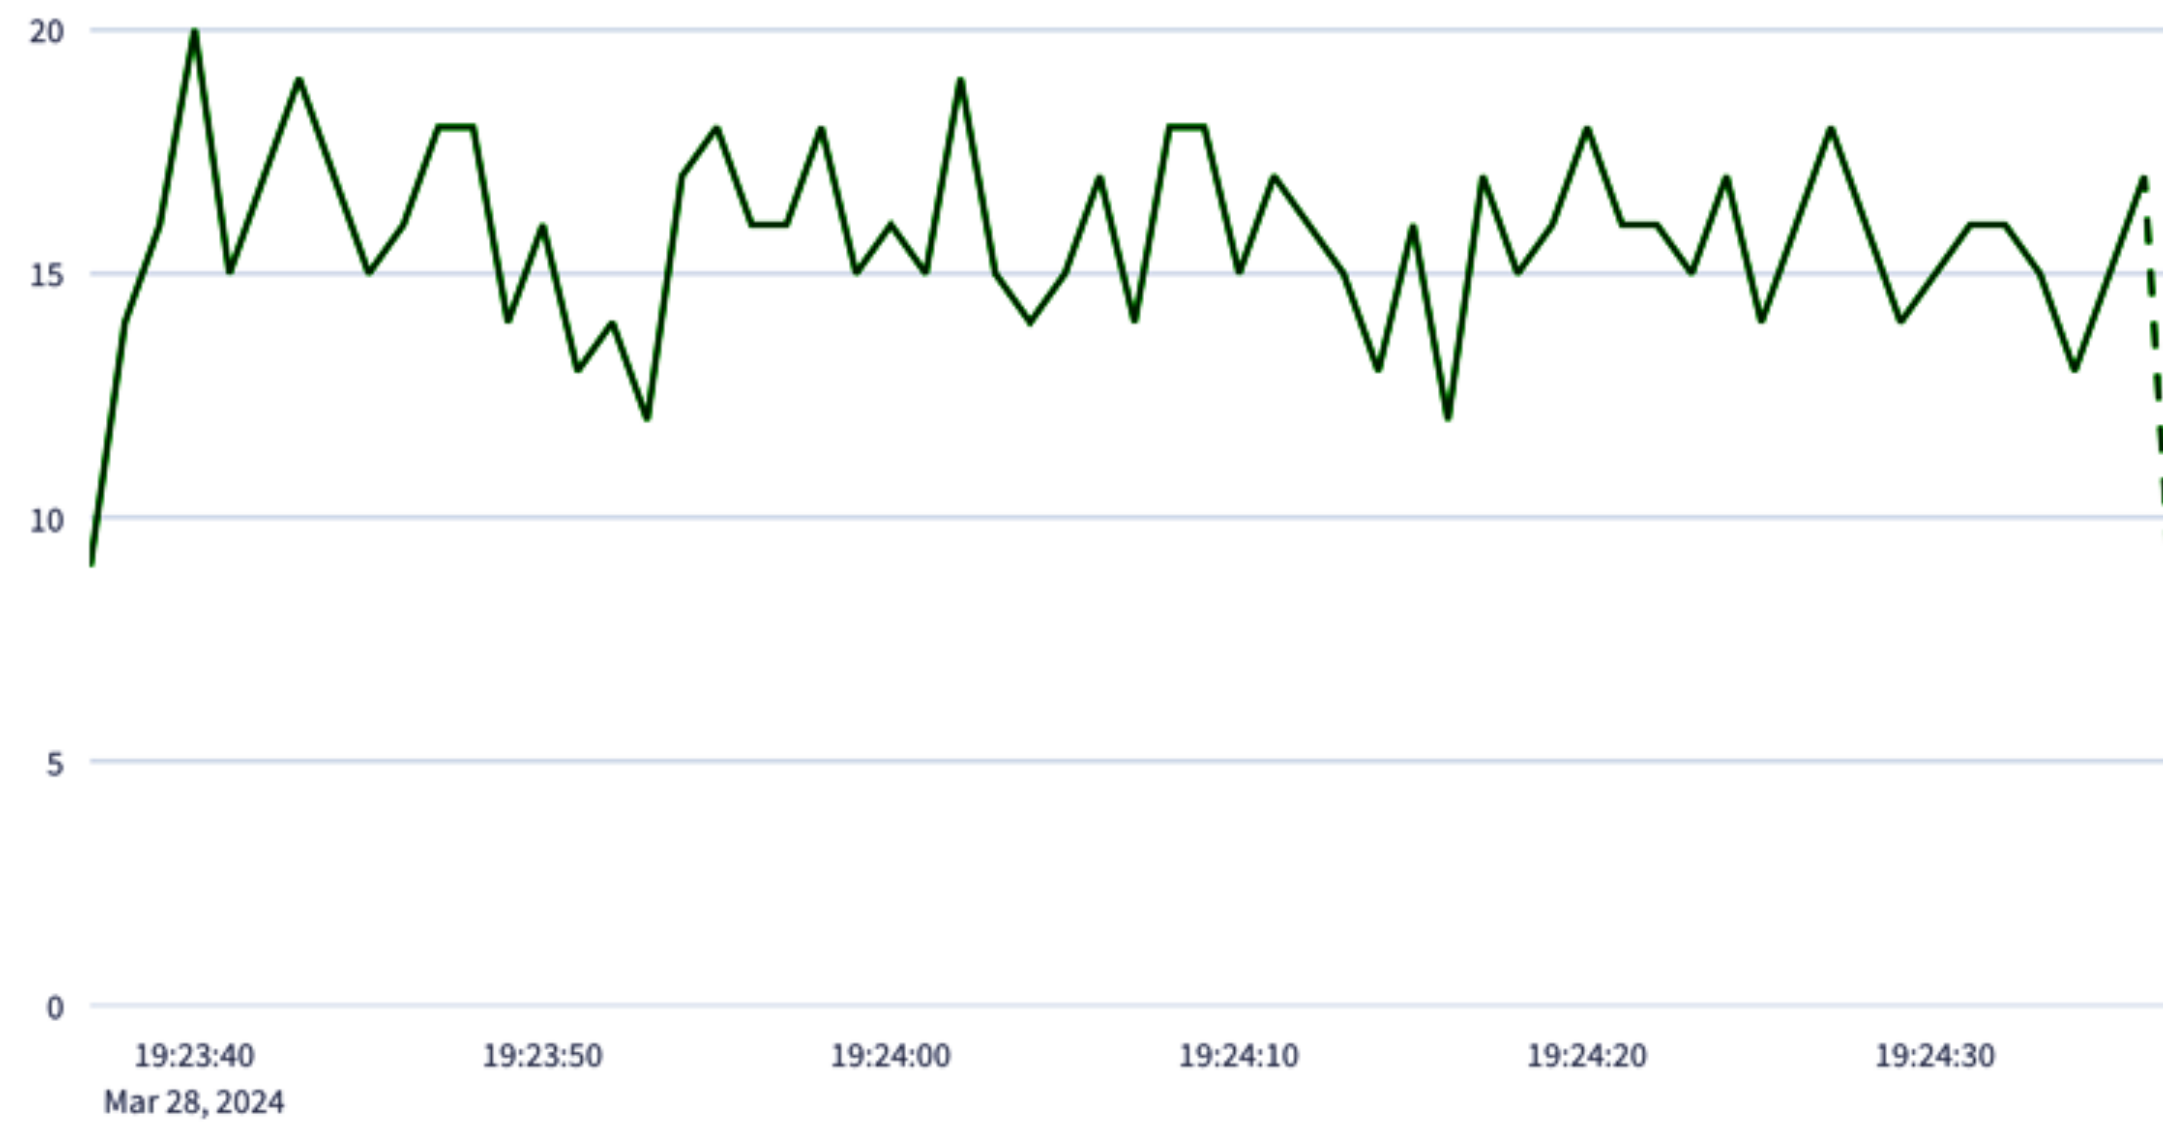
↑ 180,582.00

Average order value in ₹

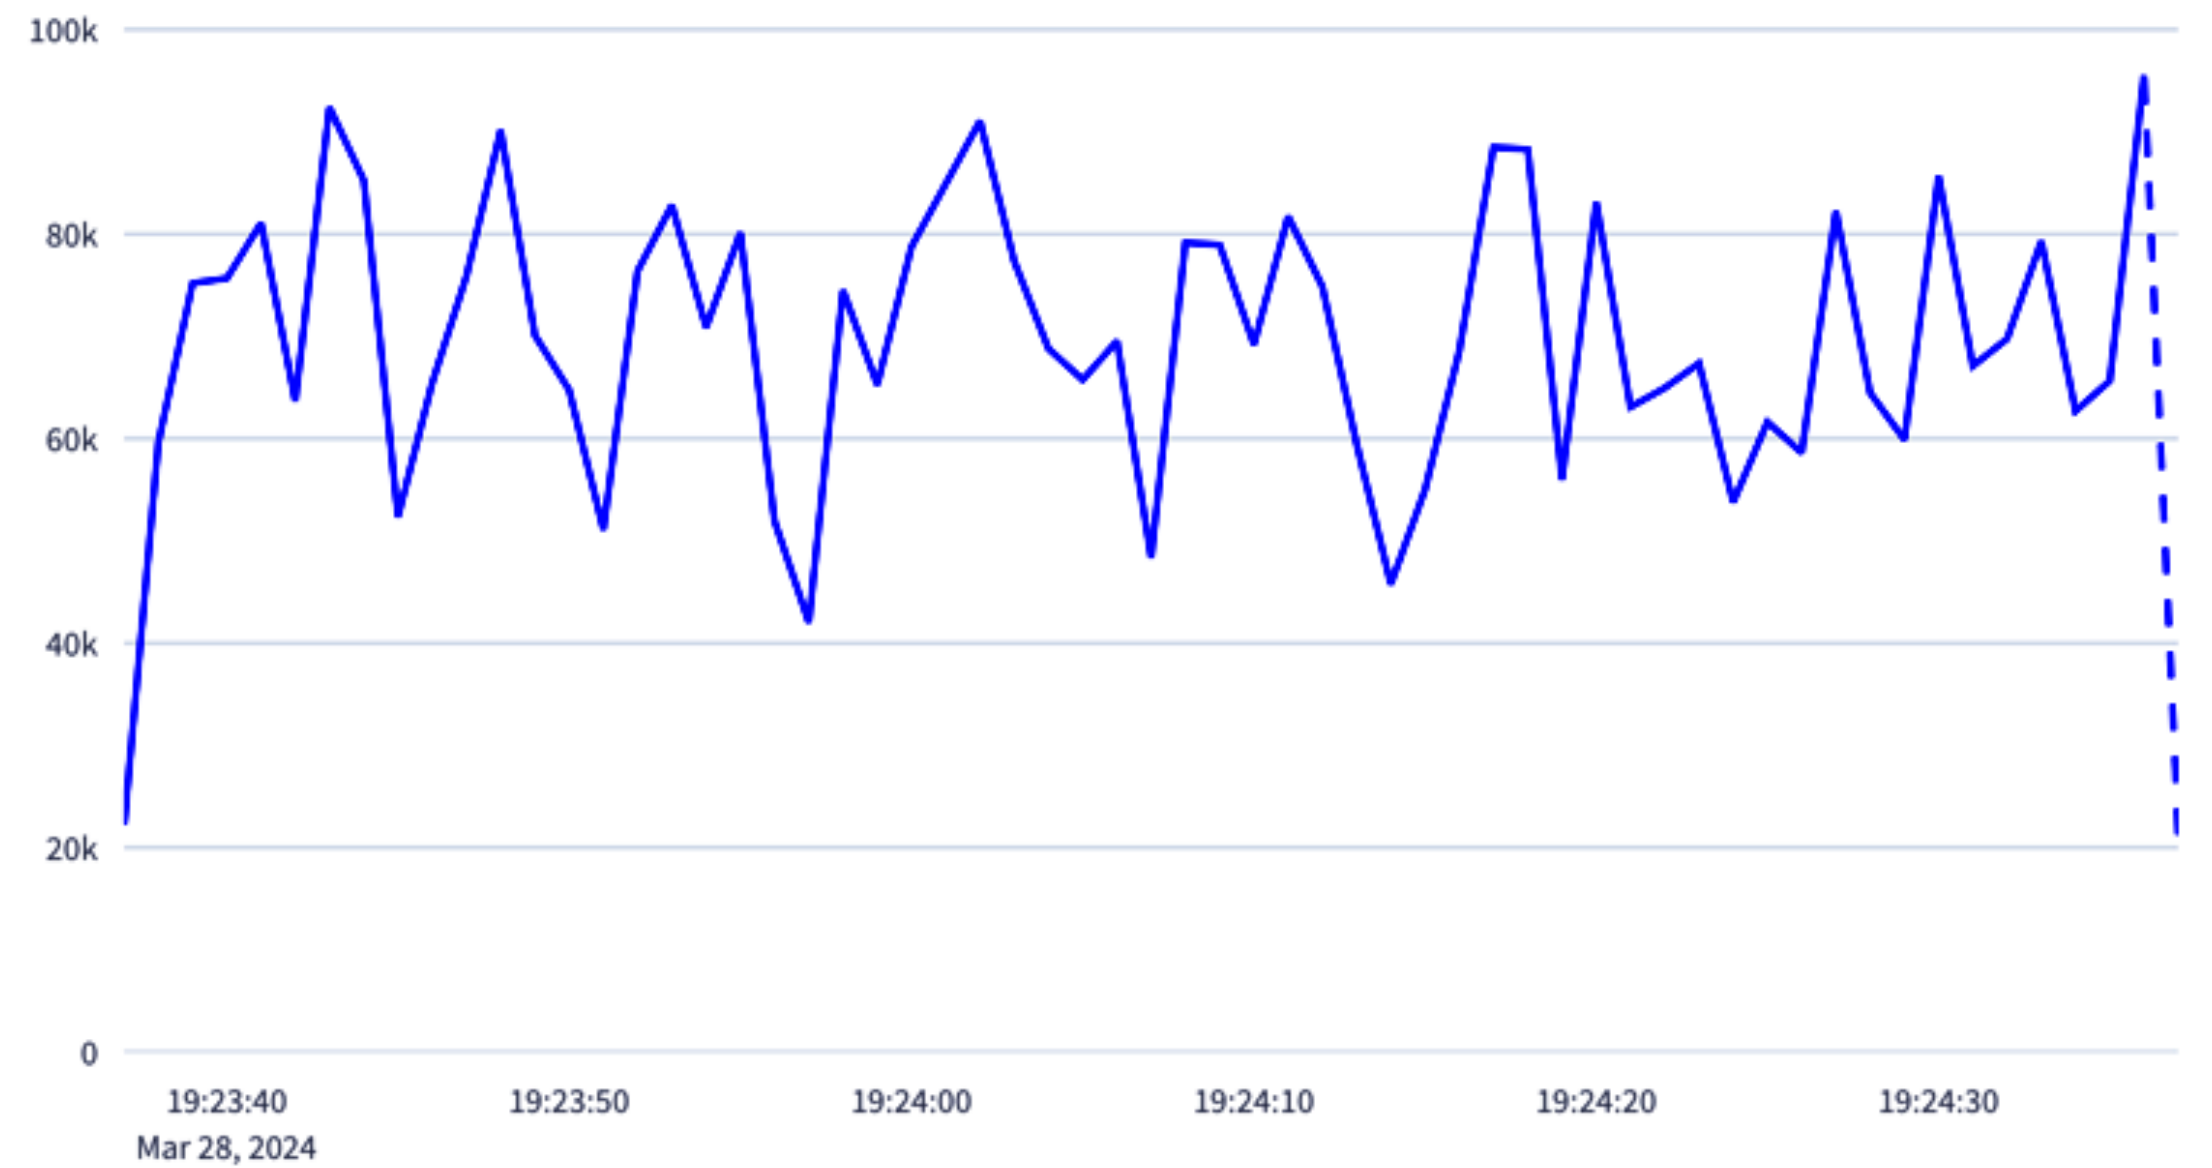
4,430.22

↑ 117.70

## Orders per second



## Revenue per second



## Latest orders

## Most popular categories

# Viktor

GAMOV

---

Head of Developer Advocacy | StarTree

Twitter X: @gamussa





«I Build Highly Scalable and Highly Available, (sometimes) Over-Engineered Hello World Applications»

@gamussa



# A Taxonomy of Analytics



OBSERVABILITY/  
MONITORING

USER-FACING  
ANALYTICS

INTERNAL

EXTERNAL

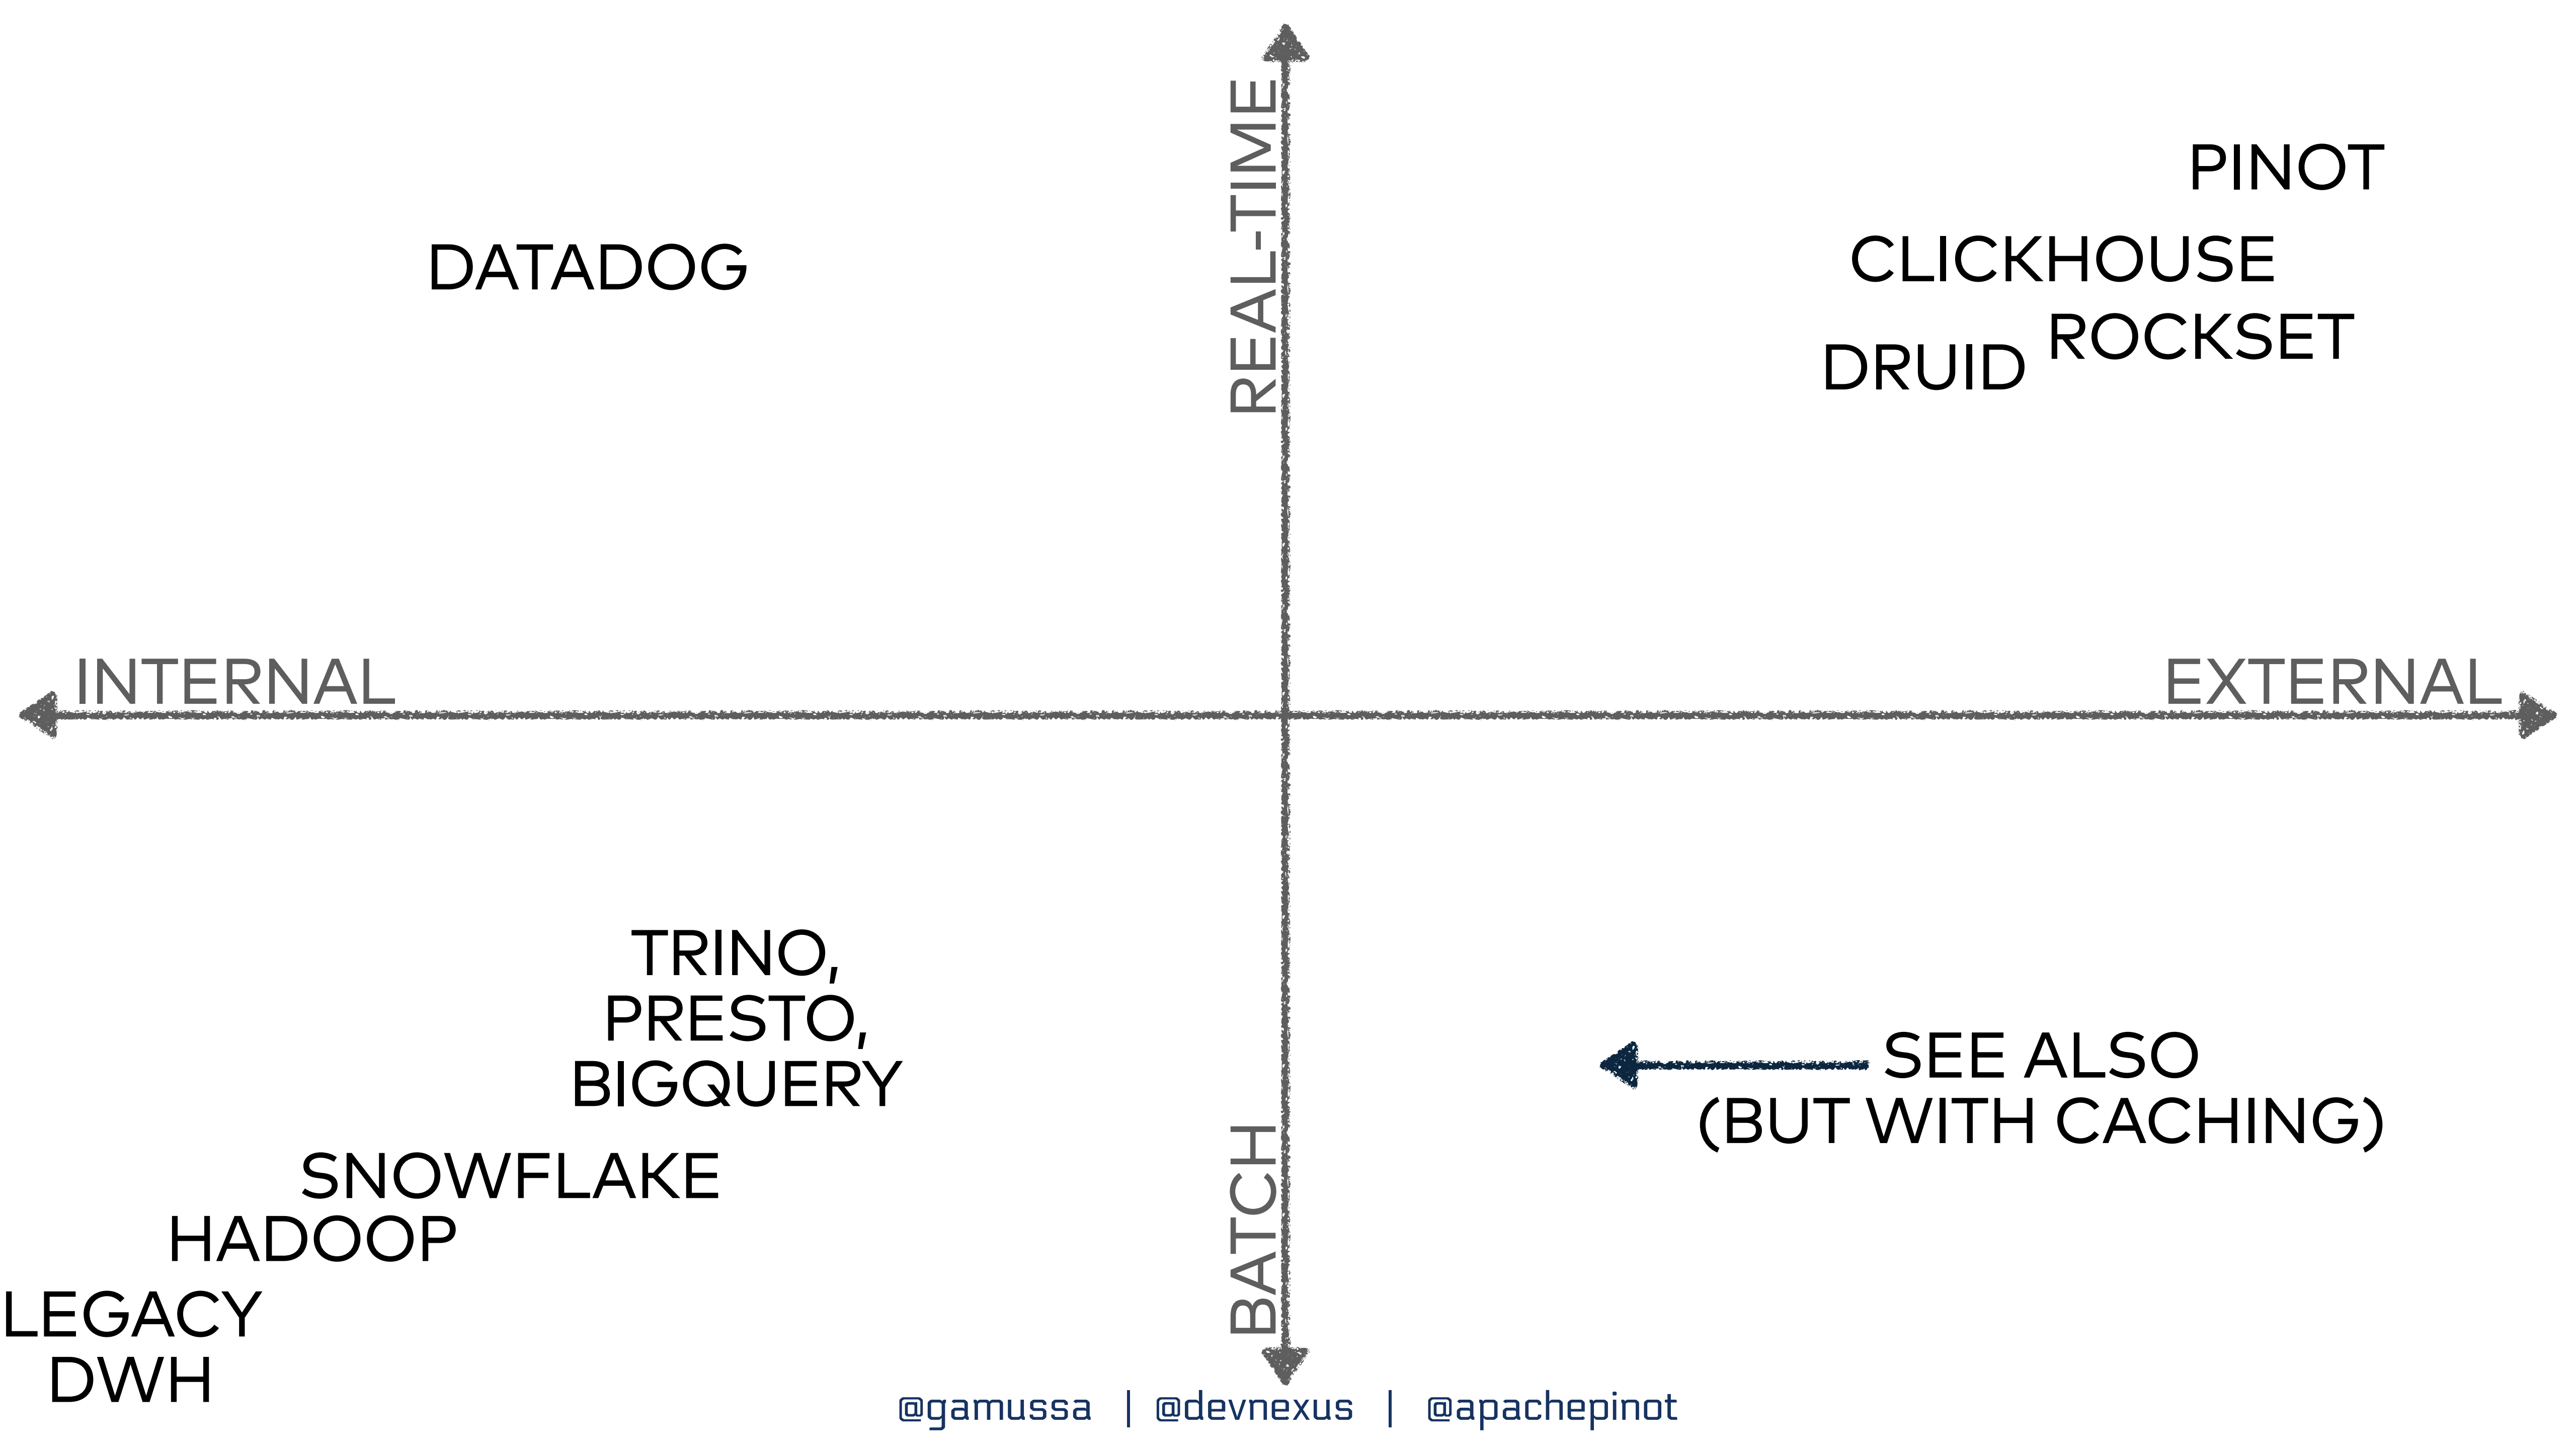
REAL-TIME

BATCH

DASHBOARDS  
REPORTING

REPORTING  
FEATURES





DATADOG

PINOT  
CLICKHOUSE  
DRUID ROCKSET

INTERNAL

EXTERNAL

REAL-TIME

BATCH

TRINO,  
PRESTO,  
BIGQUERY

SEE ALSO  
(BUT WITH CACHING)

SNOWFLAKE

HADOOP

LEGACY  
DWH



# Real-Time Analytics Summit 2024

May 7-9, 2024 | San Jose, CA

Hosted by **star+tree**




Use code  
**DELIGHT35**  
to register at  
**rtasummit.com**



# Who Does Real-Time Analytics?



# Who Viewed My Profile?



Total users	700 Million+
QPS	100,000s
Latency SLA	< 100 ms p99th
Freshness	Seconds

All profile viewers 77 interesting views 4 work at Pensando Systems 43 w >

10 people with the job title **Recruiter**

Get more views from recruiters by privately sharing you're open to opportunities

[Update settings](#)

3h

**Seunghyun Lee**  
Senior Software Engineer

[Connect](#)

11h

**Chinmay Soman**  
Founding Engineer

[Follow](#)

# LinkedIn Feed

**Chinmay Soman**  
Founding Engineer at Stealth

Who viewed your profile **431**  
Views of your post **564**

See all Premium features

Saved items

Recent

- LinkedIn Alumni Network (...)
- # careers
- # futurism
- # innovation
- # management

Groups

- LinkedIn Alumni Network (...)
- SNA Projects
- Stanford CS Alumni

Show more

Events +

Start a post

Write an article on LinkedIn

Sort by: Top

**kishore gopalakrishna** • 1st  
Builder  
21h •

Congratulations to the new Apache Pinot committers - [Kartik Khare](#), [Ting Chen](#), [Haibo Wang](#).

...see more

**Congrats to the new Apache Pinot Committers !**

<b>Haibo Wang</b>	<b>Kartik Khare</b>	<b>Ting Chen</b>
Launched Presto-Pinot connector at Uber Drove efforts to open source it Schema evolution	S3 support Protobuf input format Scalar transform functions JDBC driver	Bypass deep-store for Pinot Realtime Peer download scheme Enhancements to Realtime ingestion

In their own words, why they love contributing to Apache Pinot...

"Apache Pinot is battle-tested in large scale deployments and keeps attracting new diverse usecases. I love seeing the thriving & supportive community"	"Questions are quickly resolved on the community Slack channel, no matter how trivial they are. The reviews on PR also go deep into each line & are helping me upgrade my coding standard"	"The community is open to new ideas. There are established processes to contribute new features & bug fixes. I get to solve interesting and challenging problems in a distributed data system"
---	--	--

50 • 2 Comments

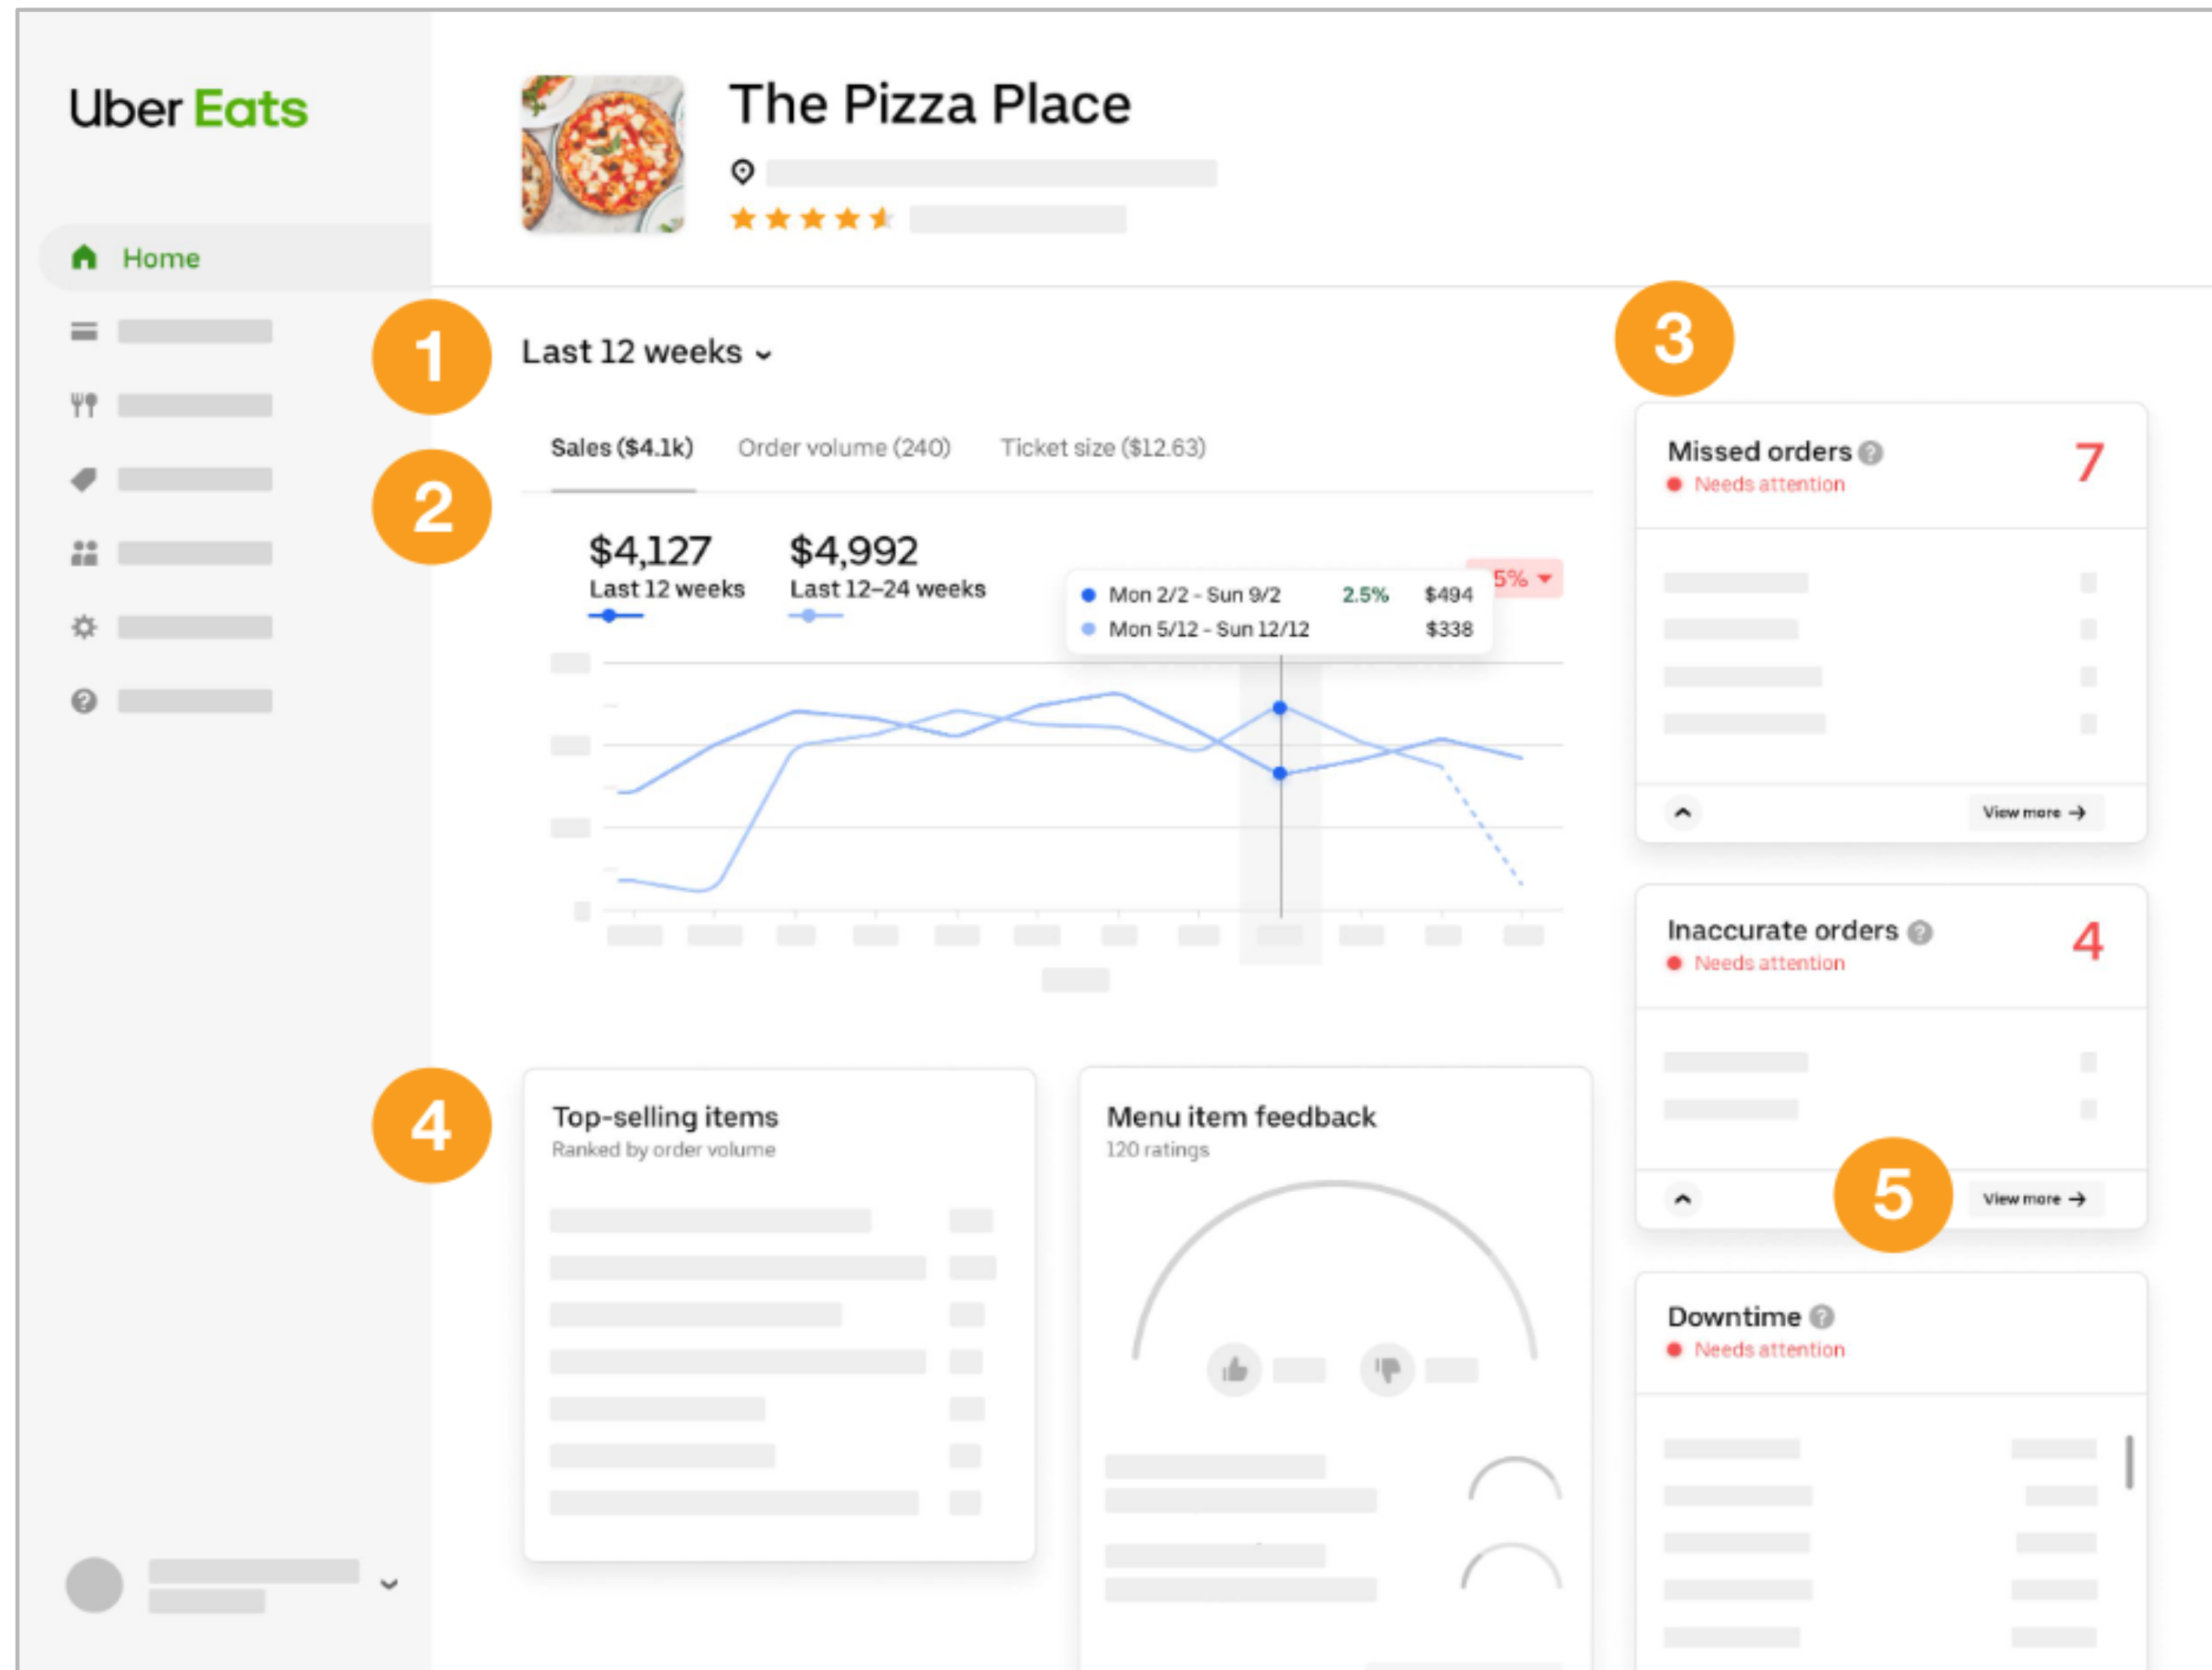
Like Comment Share Send

```
SELECT sum(count) from T
WHERE memberId = X
AND item in (list of 500-1500 items)
AND time >= (now - 14 days)
GROUP BY action, item, position, time
```

Total users	700 Million
QPS	100,000
Latency SLA	< 100 ms p99th
Freshness	Seconds



# Uber Eats Restaurant Manager



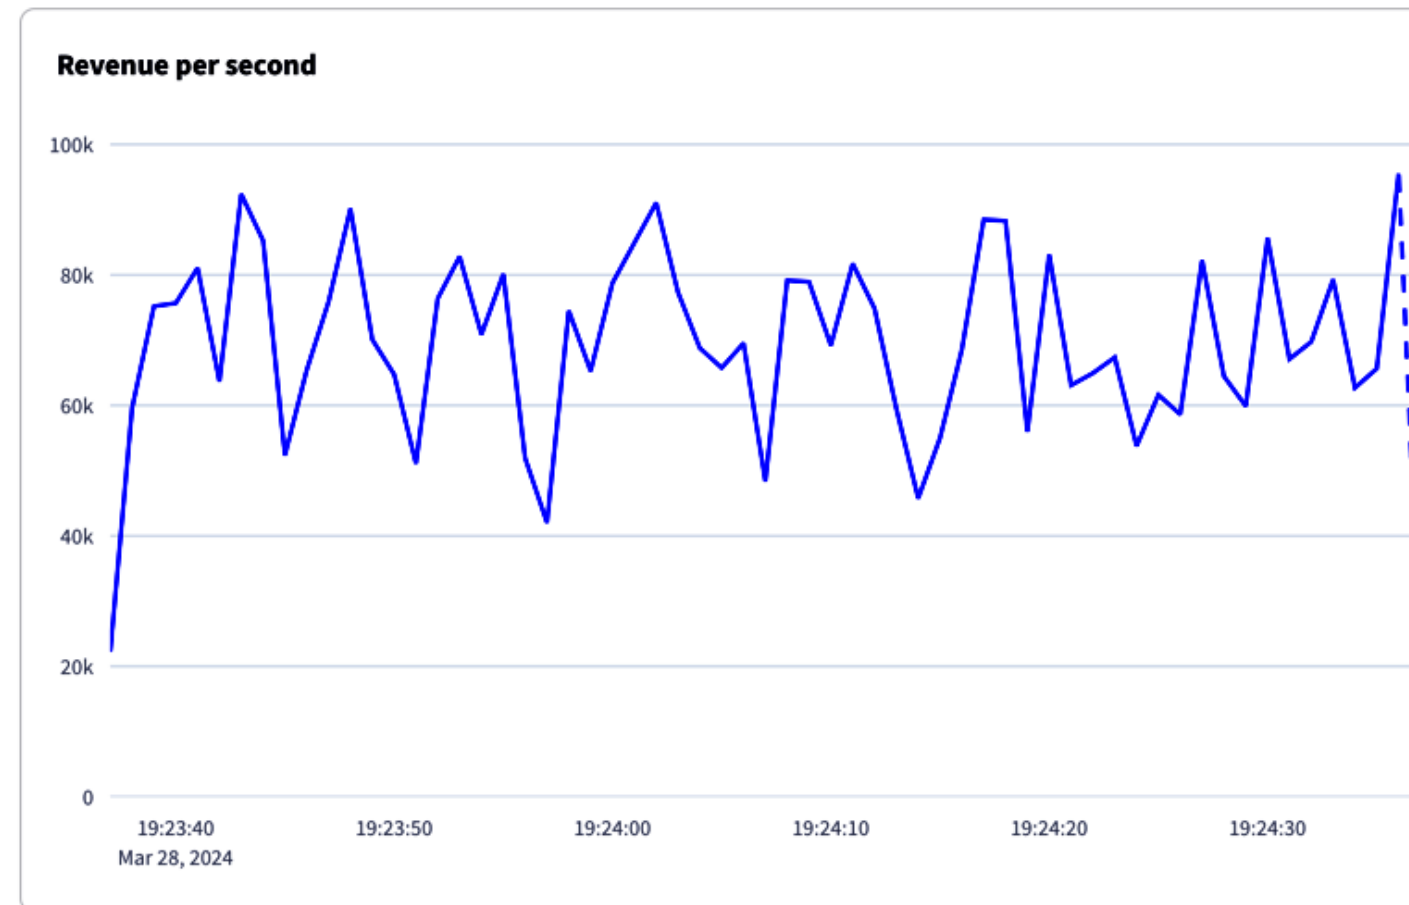

- Identify surge in realtime
- Detect missed/inaccurate orders in realtime

Total users	500000+
QPS	100s
Latency SLA	< 100 ms p99th
Freshness	Seconds - Minutes

# Who Much Pizza We Sold?

**Orders in the last 1 minute**

Number of orders	Revenue in ₹	Average order value in ₹
<b>948</b>	<b>4,199,845.00</b>	<b>4,430.22</b>
↑ 16	↑ 180,582.00	↑ 117.70



**Latest orders**

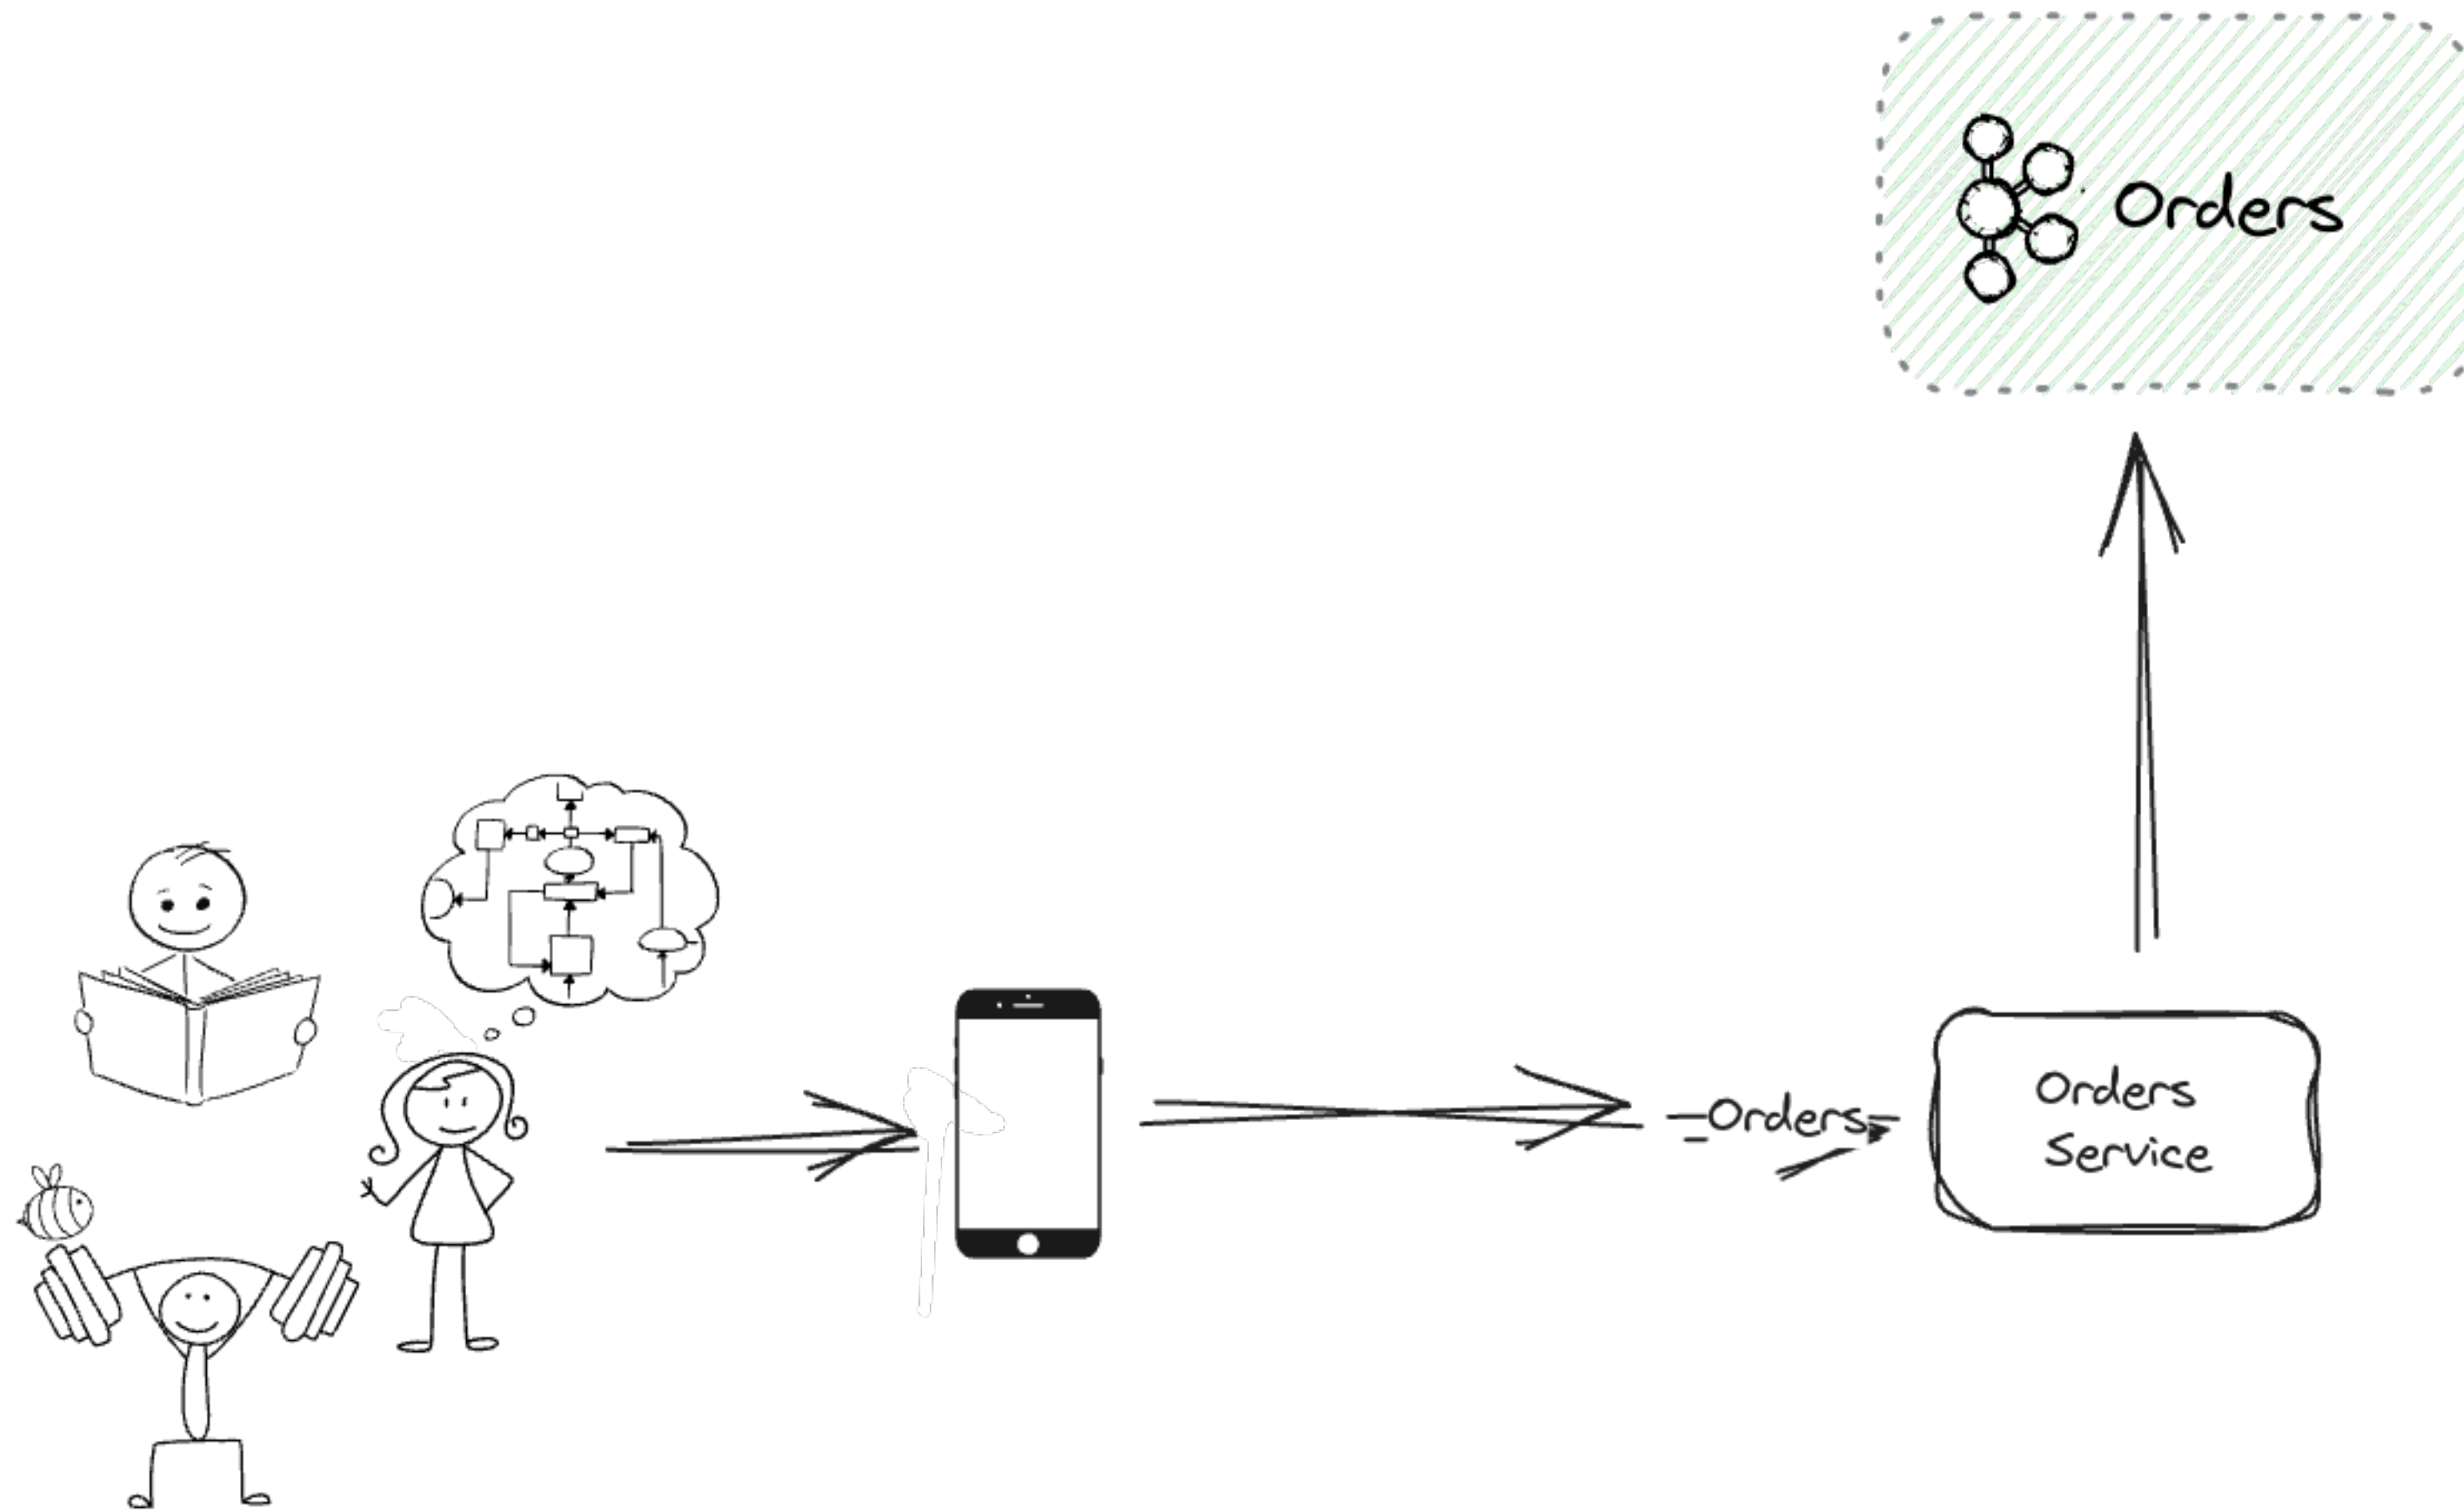
	Time	Status	Price	User ID	Quantity	Total quantity
0	19:24:37:410	PLACED_ORDER	₹5830.00	907	4	12
1	19:24:37:364	PLACED_ORDER	₹3909.00	562	5	11
2	19:24:37:305	PLACED_ORDER	₹1980.00	447	4	12
3	19:24:37:228	PLACED_ORDER	₹5870.00	824	8	30
4	19:24:37:173	PLACED_ORDER	₹180.00	581	1	3
5	19:24:37:112	PLACED_ORDER	₹930.00	708	2	4
6	19:24:37:024	PLACED_ORDER	₹8359.00	876	7	19
7	19:24:36:954	PLACED_ORDER	₹9591.00	329	9	36

**Most popular categories**

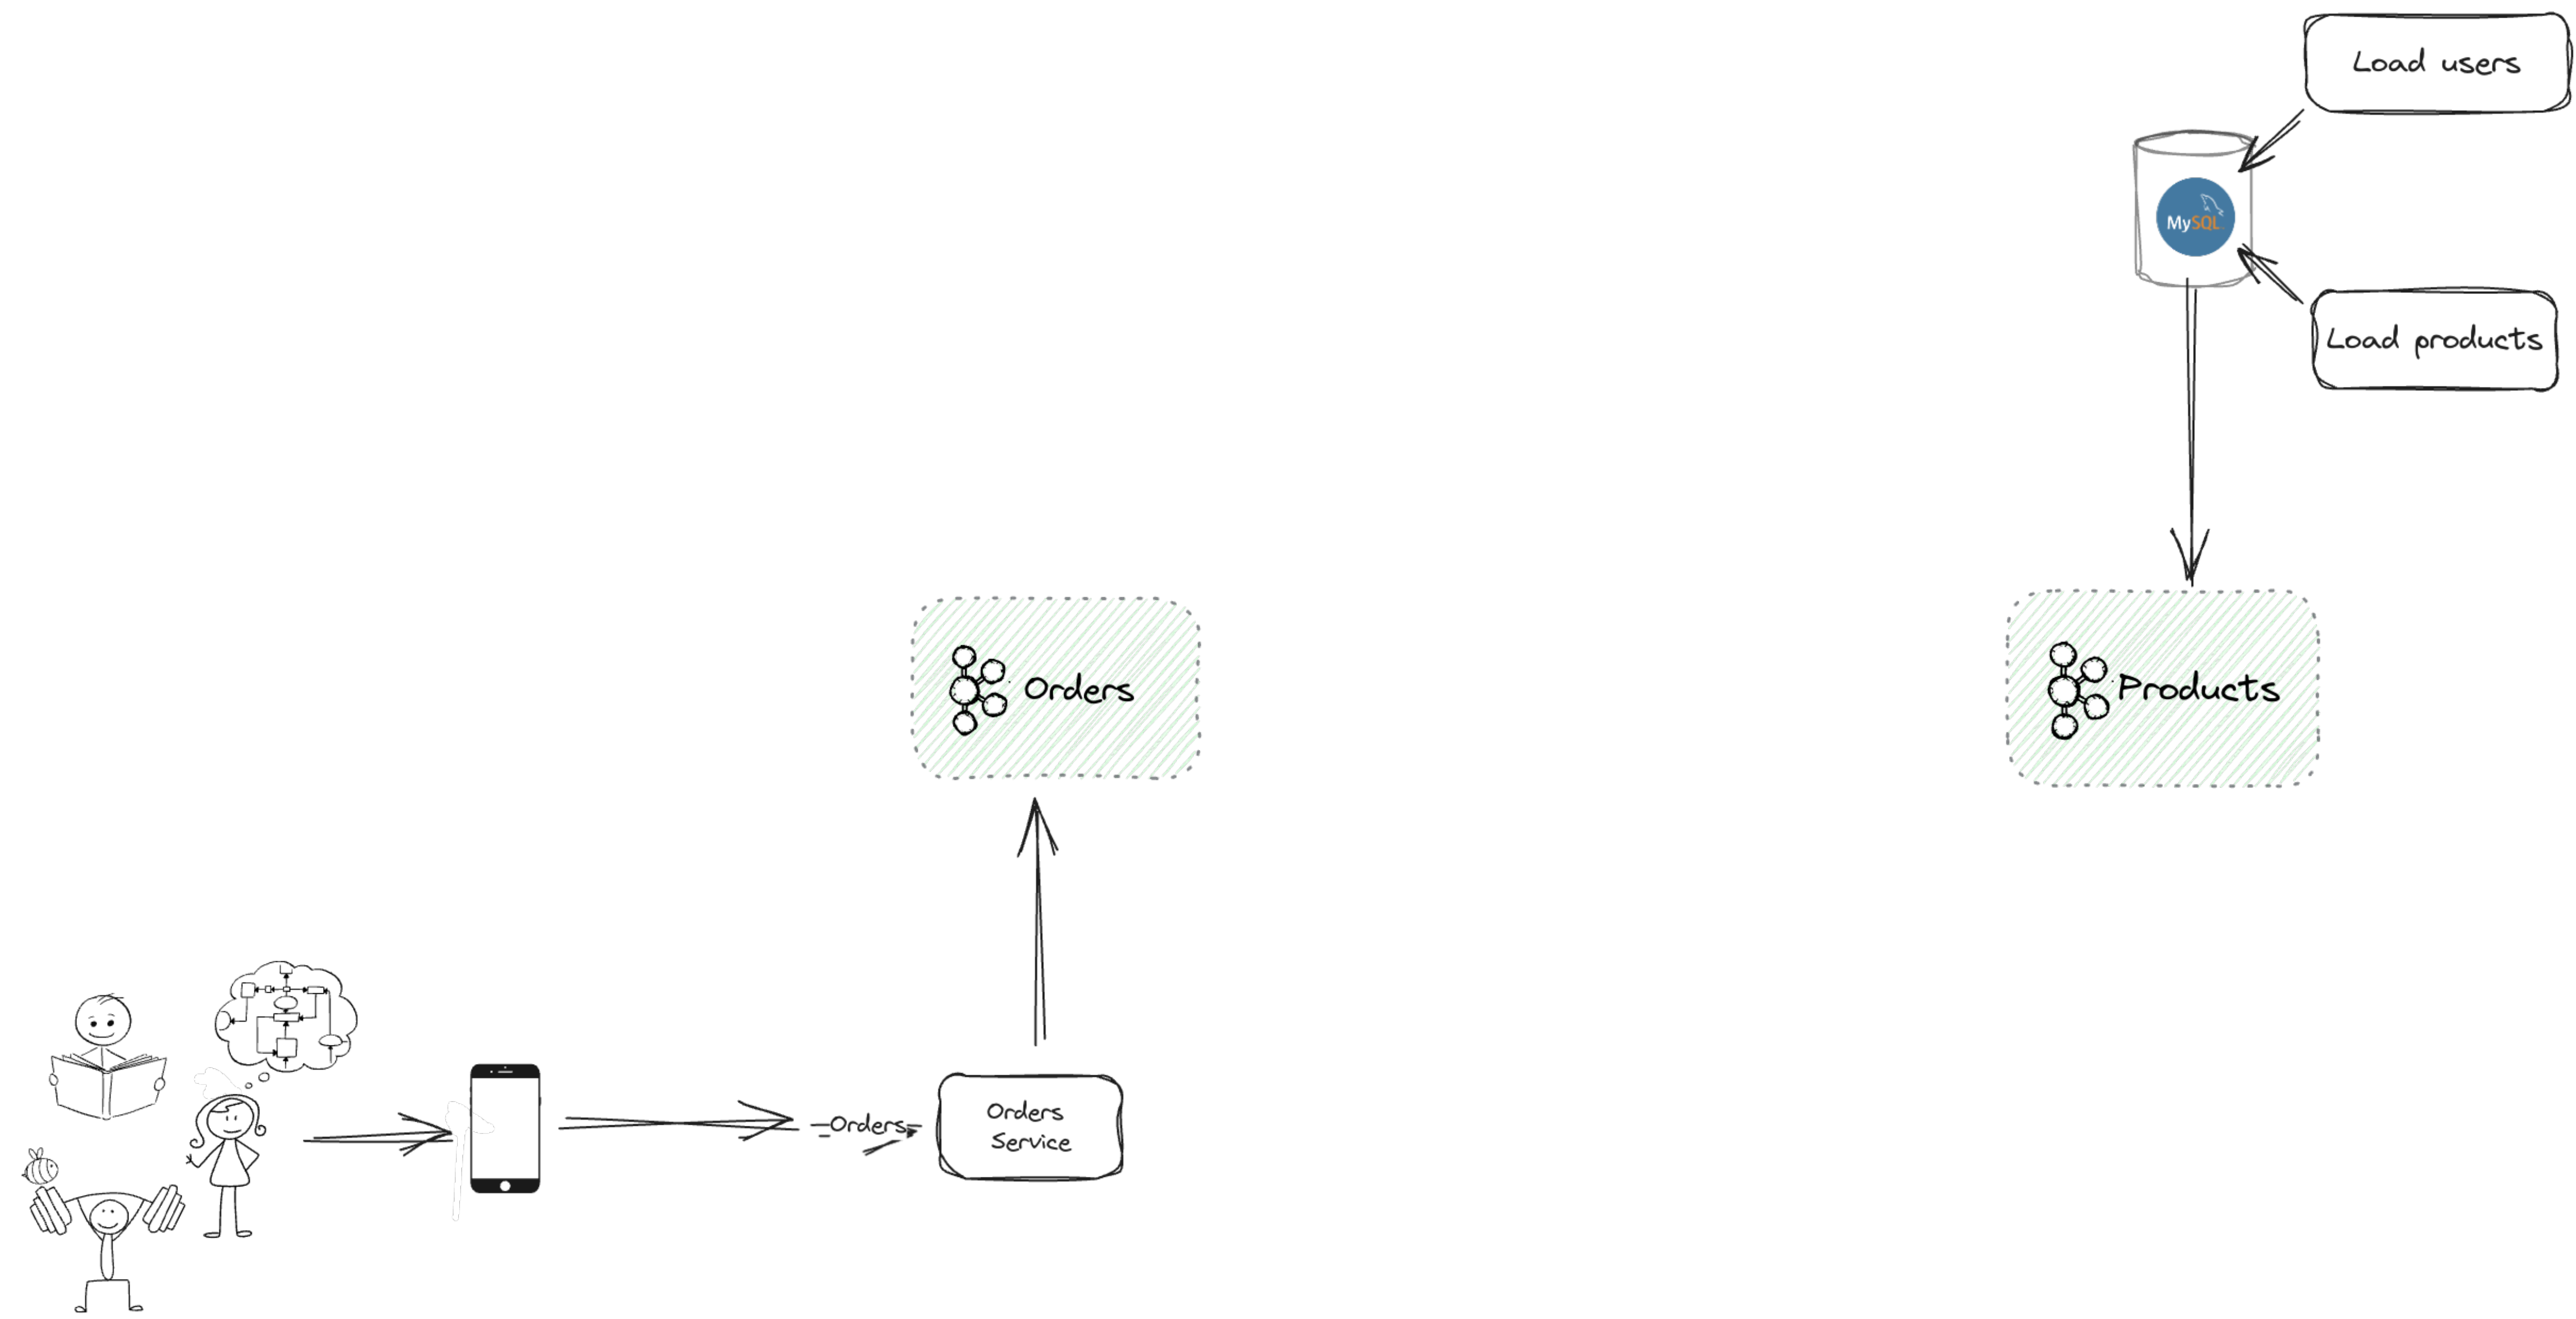
	Category	Number of orders	Total quantity ordered	Average quantity per order
0	side orders	666	\$3622	5
1	veg pizzas	571	\$2848	4
2	non veg pizzas	482	\$2201	4
3	pizza mania	449	\$2060	4
4	pasta	310	\$1144	3

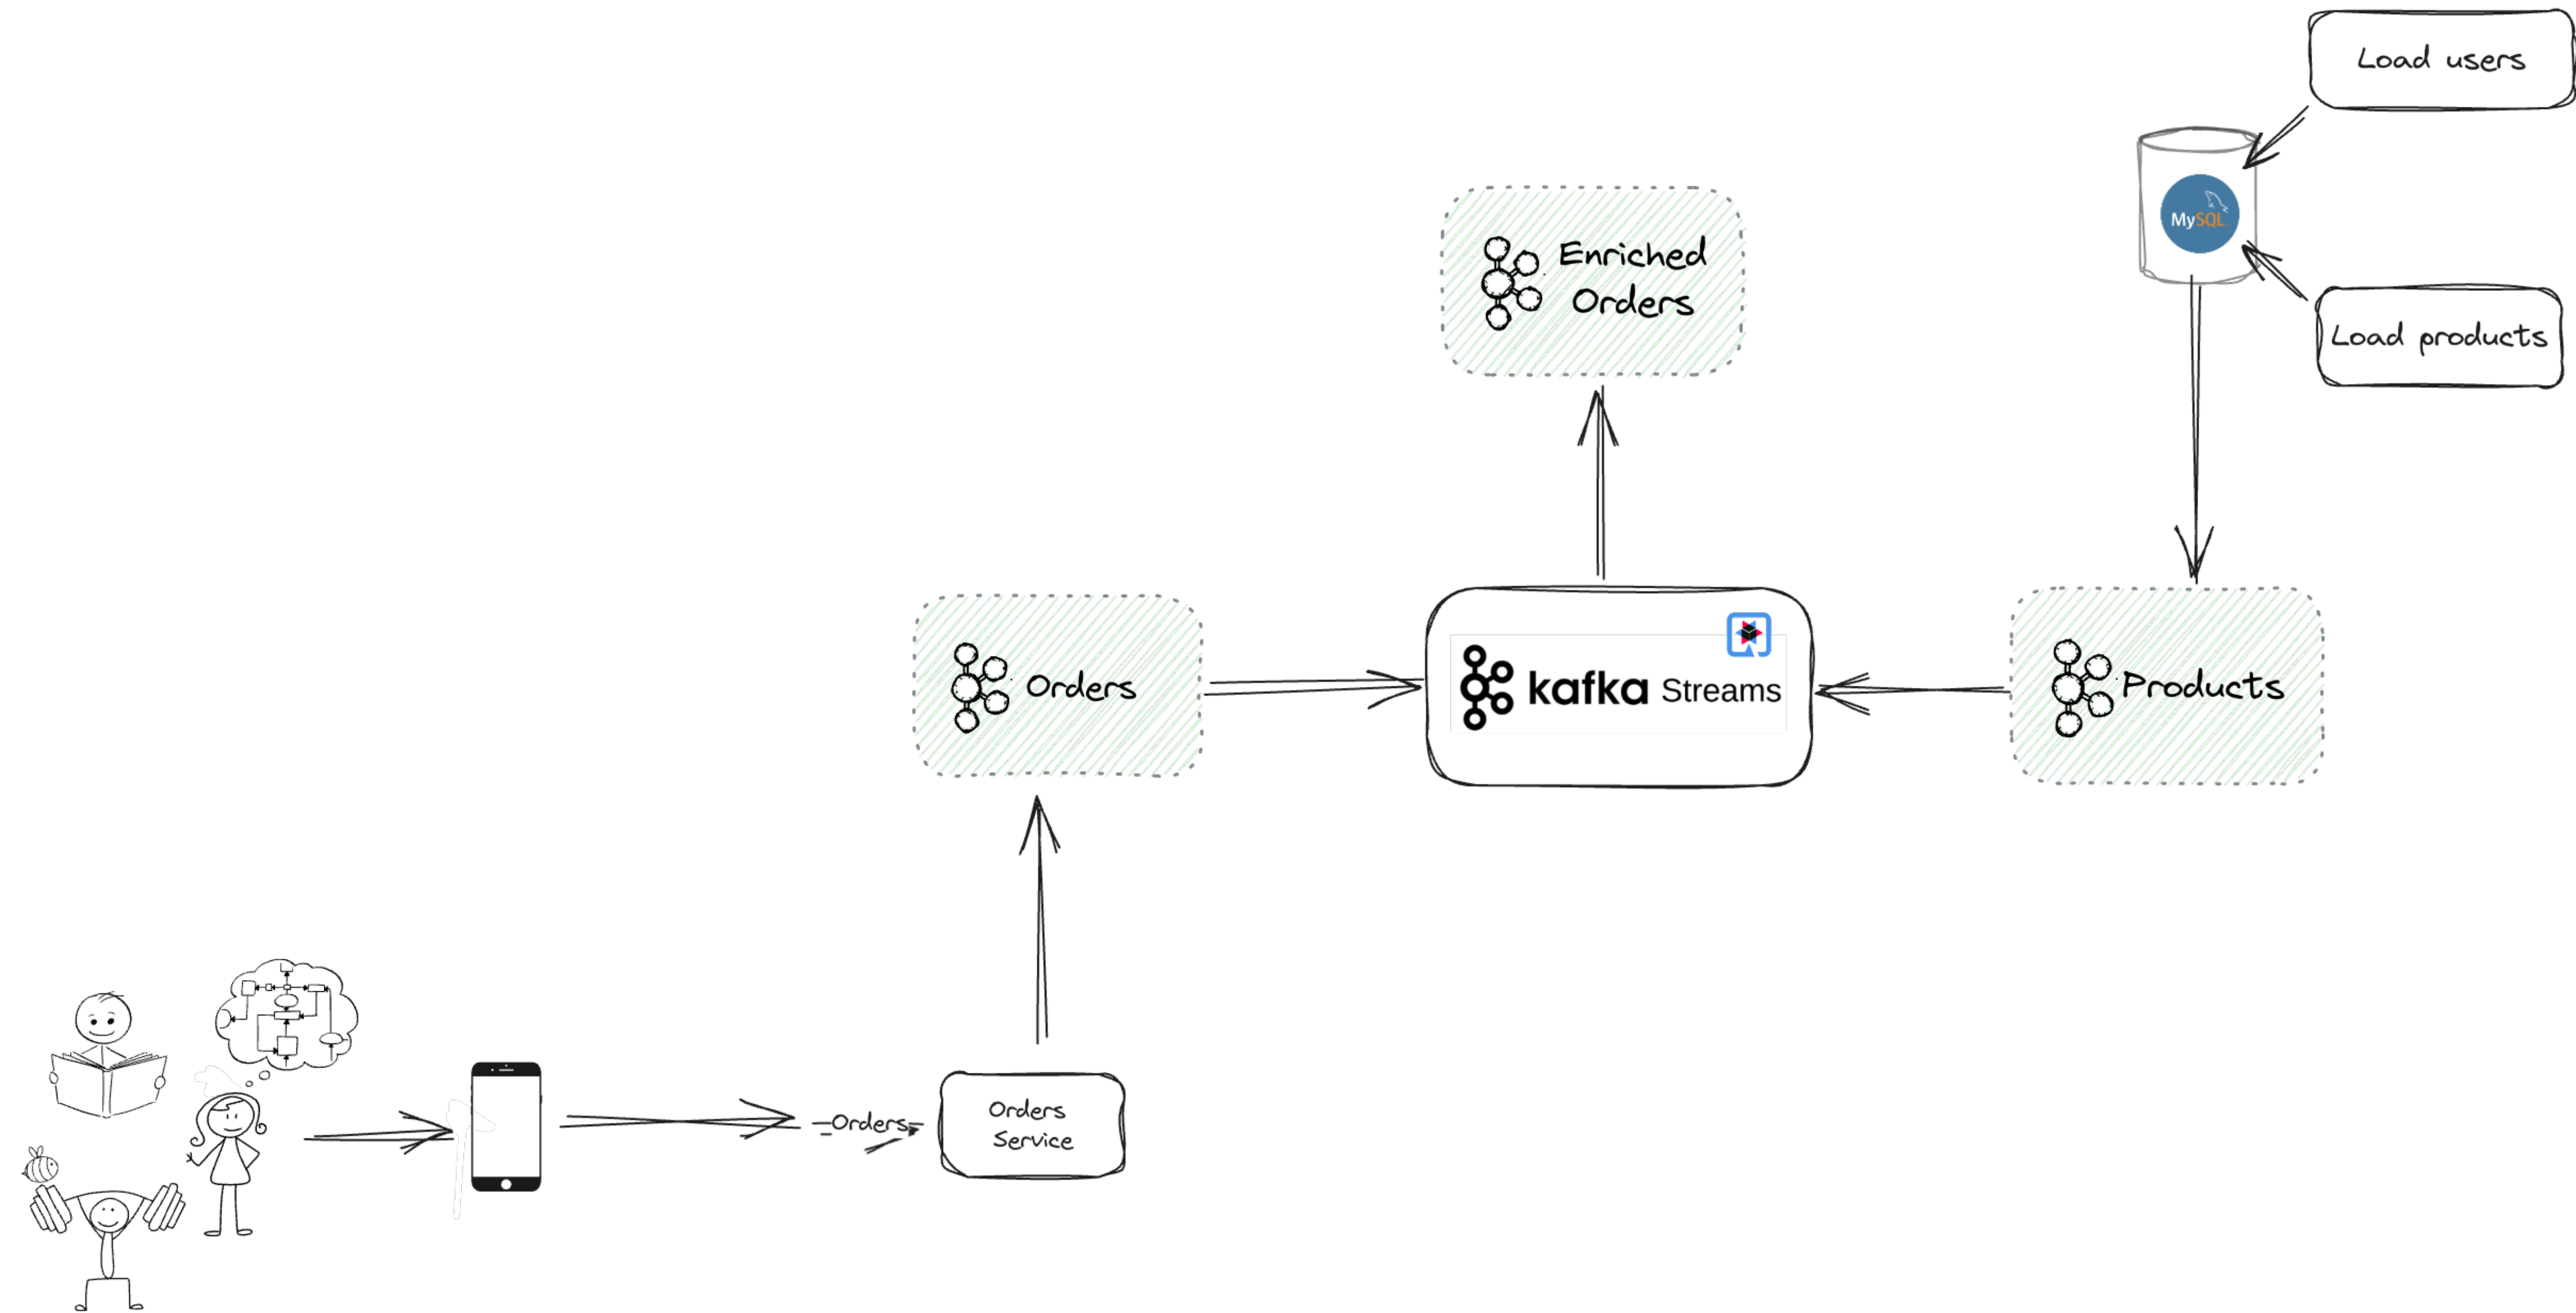


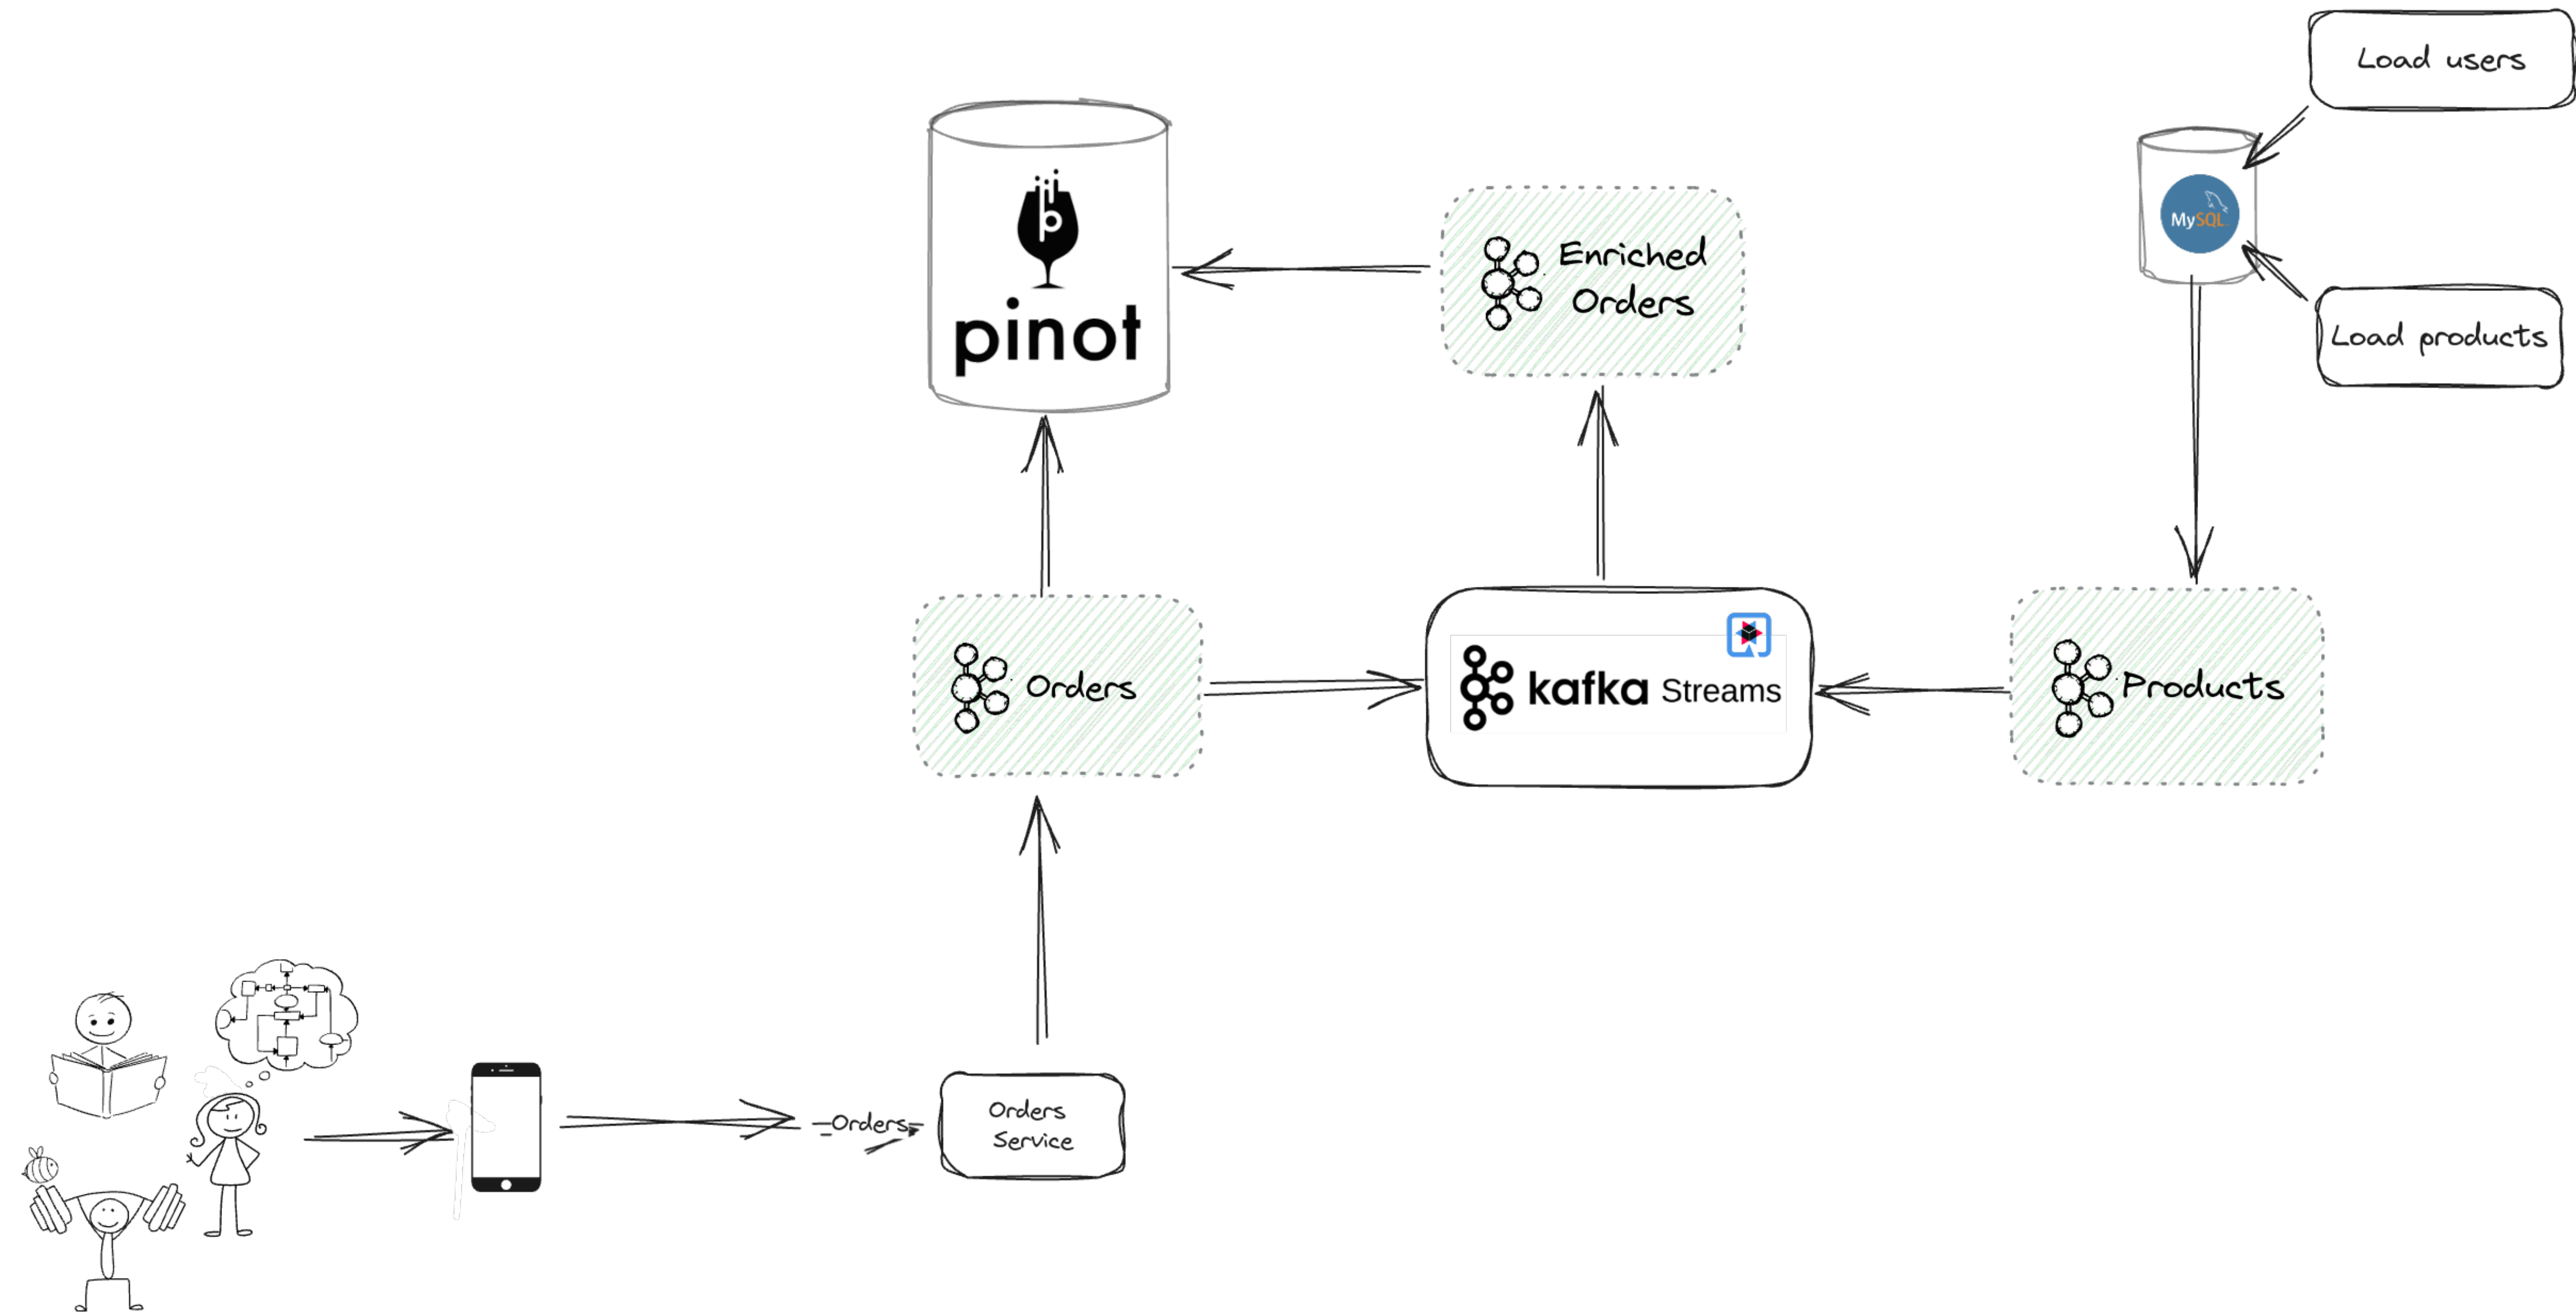
# Let's Break It All Down



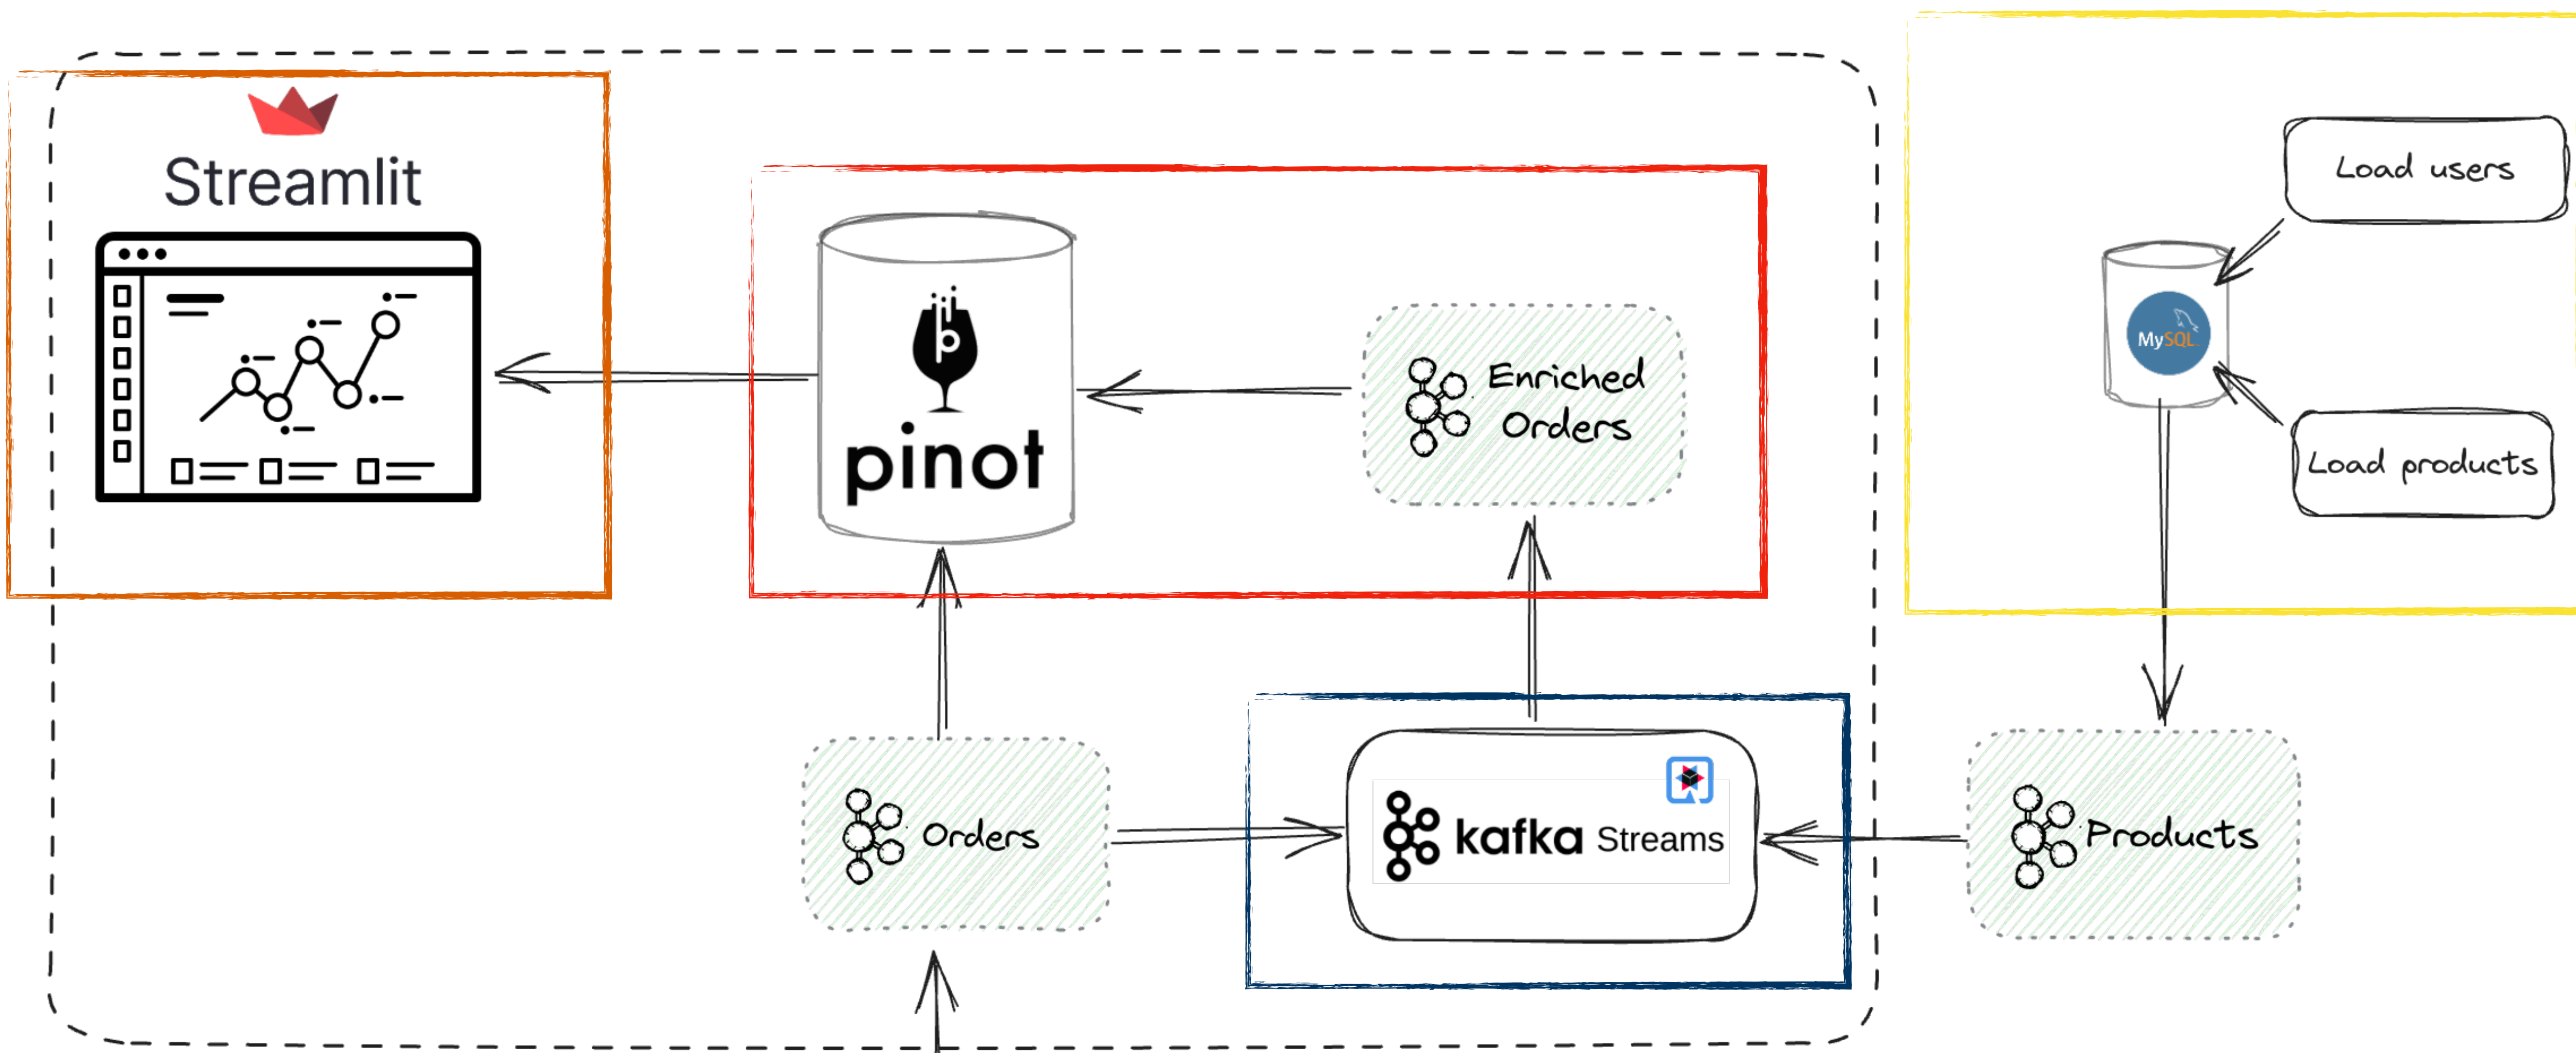












# Let's See It in Action



<https://stree.ai/pizza-demo>

# Pinot Architecture



# Batch Ingestion

```
[  
  ...,  
  {  
    "ip": "111.173.165.103",  
    "userid": 10,  
    "remote_user": "-",  
    "time": "3271",  
    "_time": 3271,  
    "request": "GET /site/user_status.html HTTP/1.1",  
    "status": "406",  
    "bytes": "1289",  
    "referrer": "-",  
    "agent": "Mozilla/5.0 (Windows NT 10.0; Win64;  
x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/  
59.0.3071.115 Safari/537.36"  
  },  
  ...  
]
```

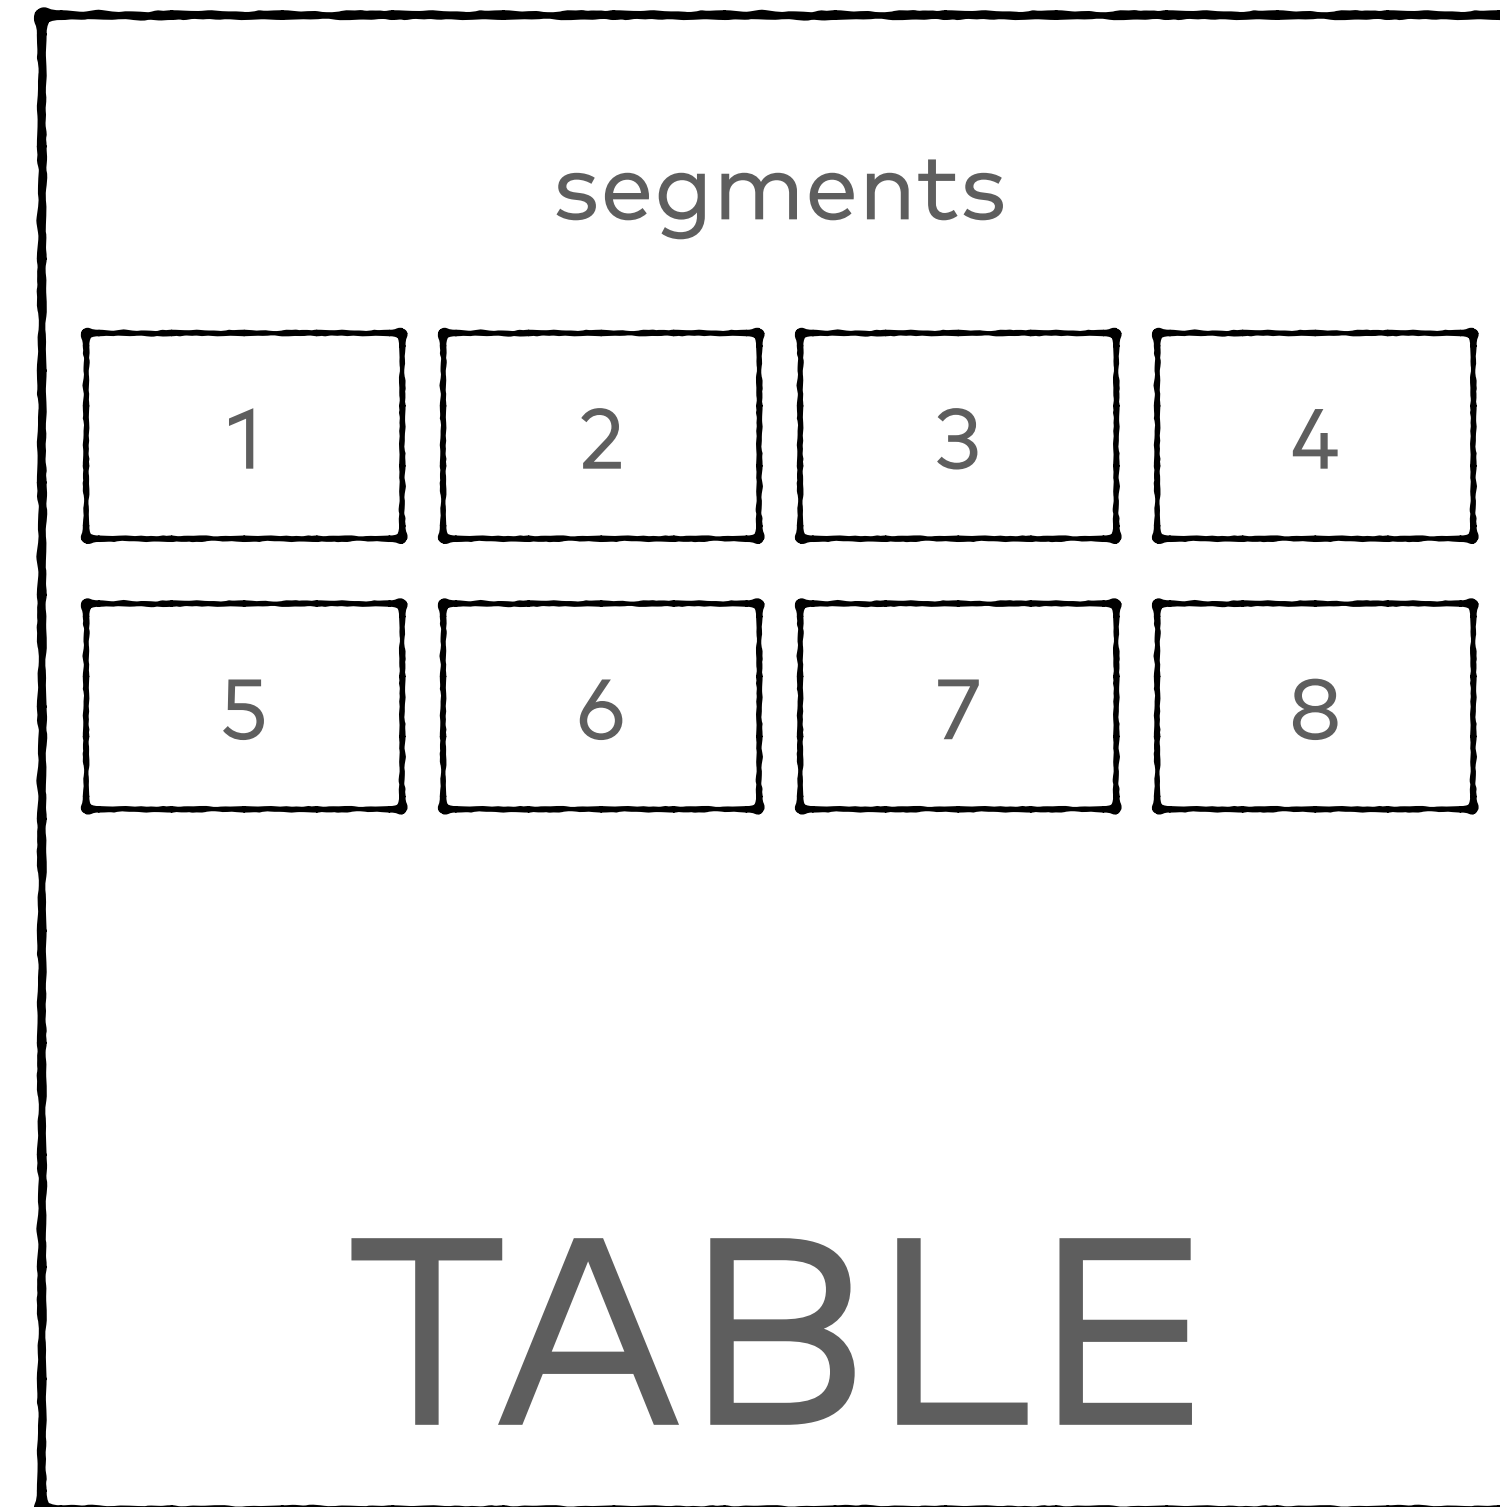
# Batch Ingestion

```
[  
  { ... },  
  { ... },  
  { ... },  
  { ... }  
]
```



# Batch Ingestion

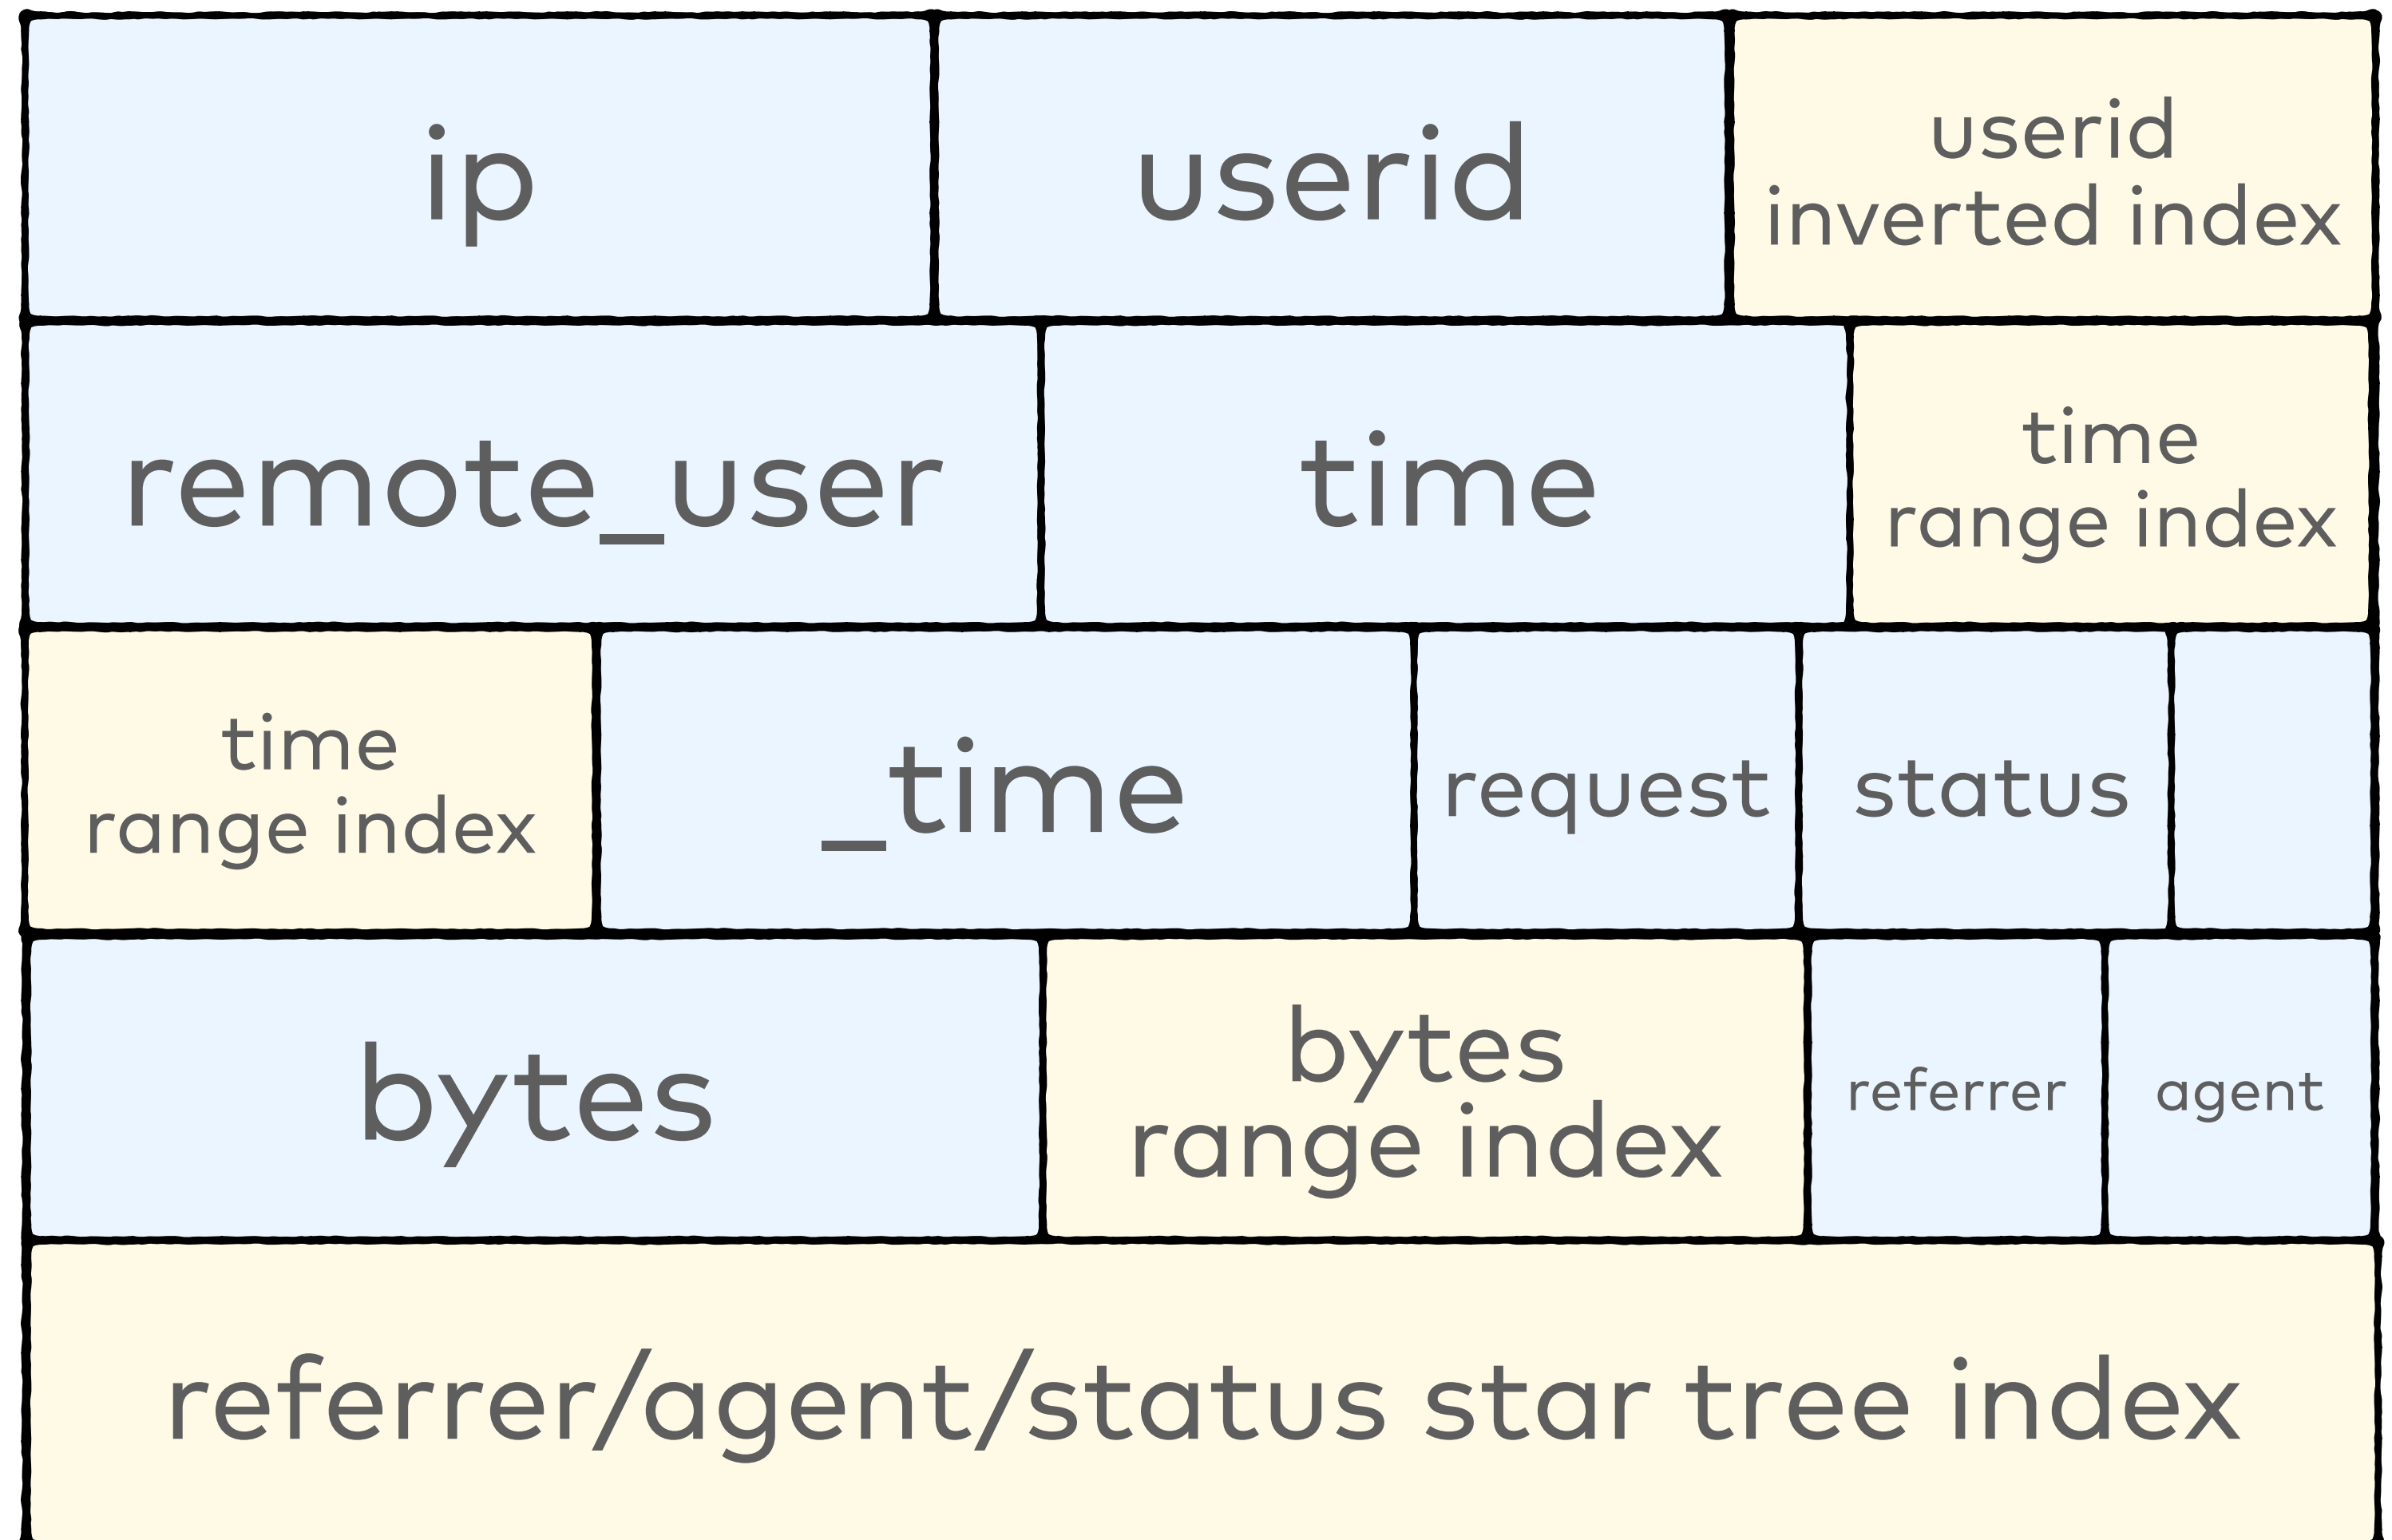
```
[  
  { ... },  
  { ... },  
  { ... },  
  { ... }  
]
```






# Segment Structure


```
...,  
{  
  "ip": "111.173.165.103",  
  "userid": 10,  
  "remote_user": "-",  
  "time": "3271",  
  "_time": 3271,  
  "request": "GET",  
  "status": "406",  
  "bytes": "1289",  
  "referrer": "-",  
  "agent": "Mozilla/5.0"  
},  
...,
```



# Segments




# Segments








# Segments



# Servers



# Servers








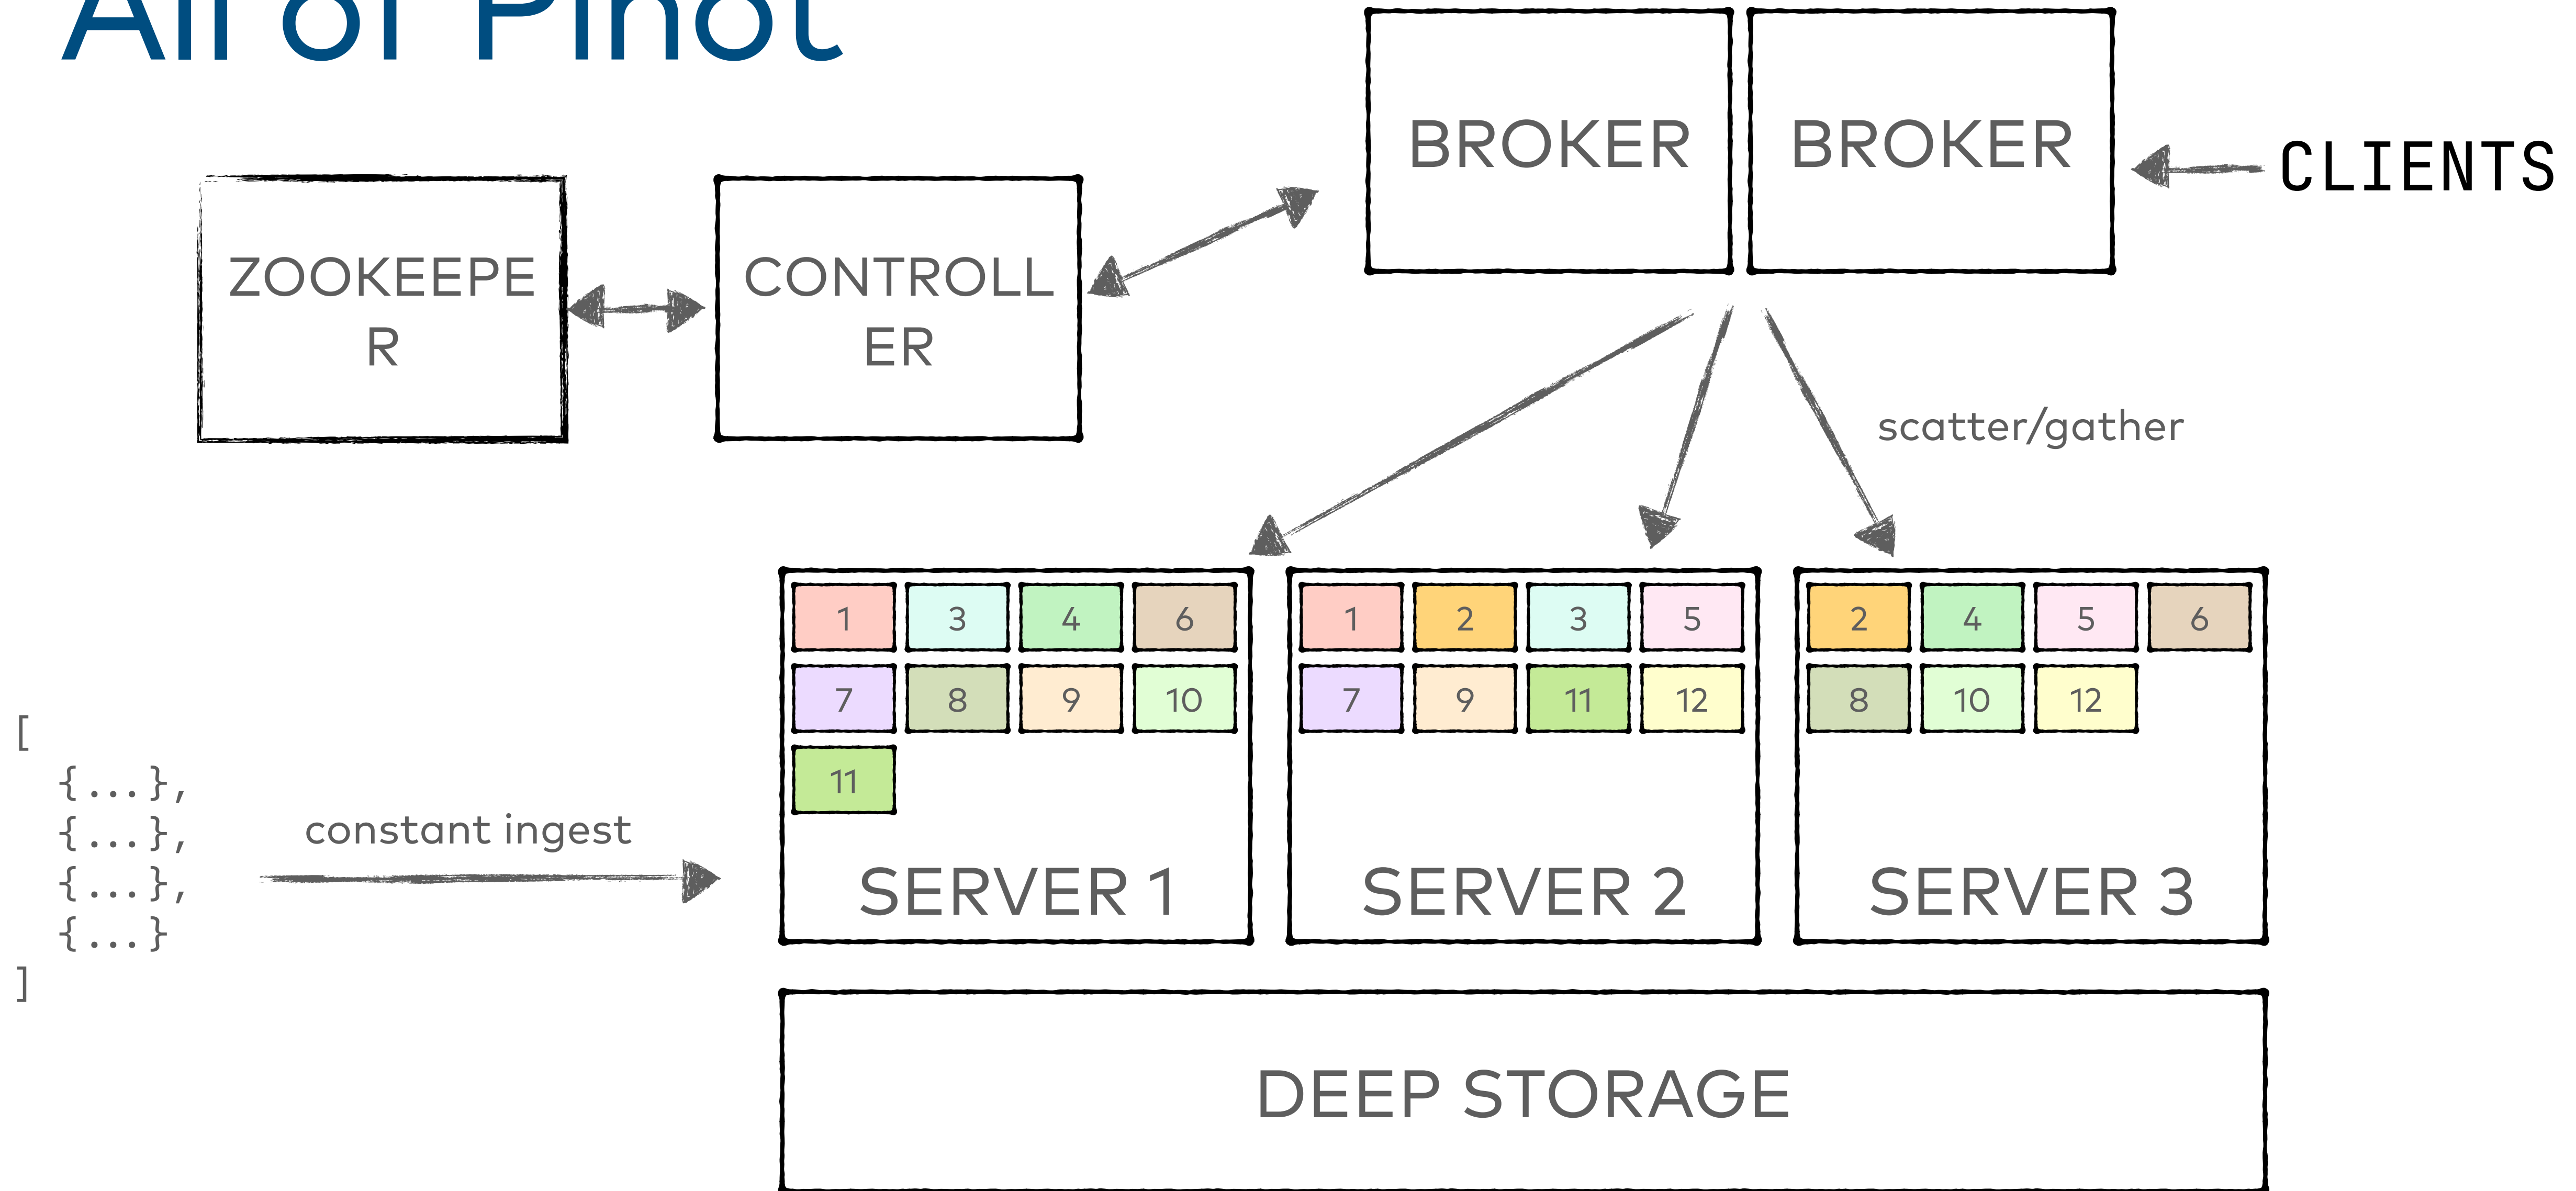
# Brokers

SELECT \*  
FROM  
SOMETHING

(query needs  
segments 2 and 8)

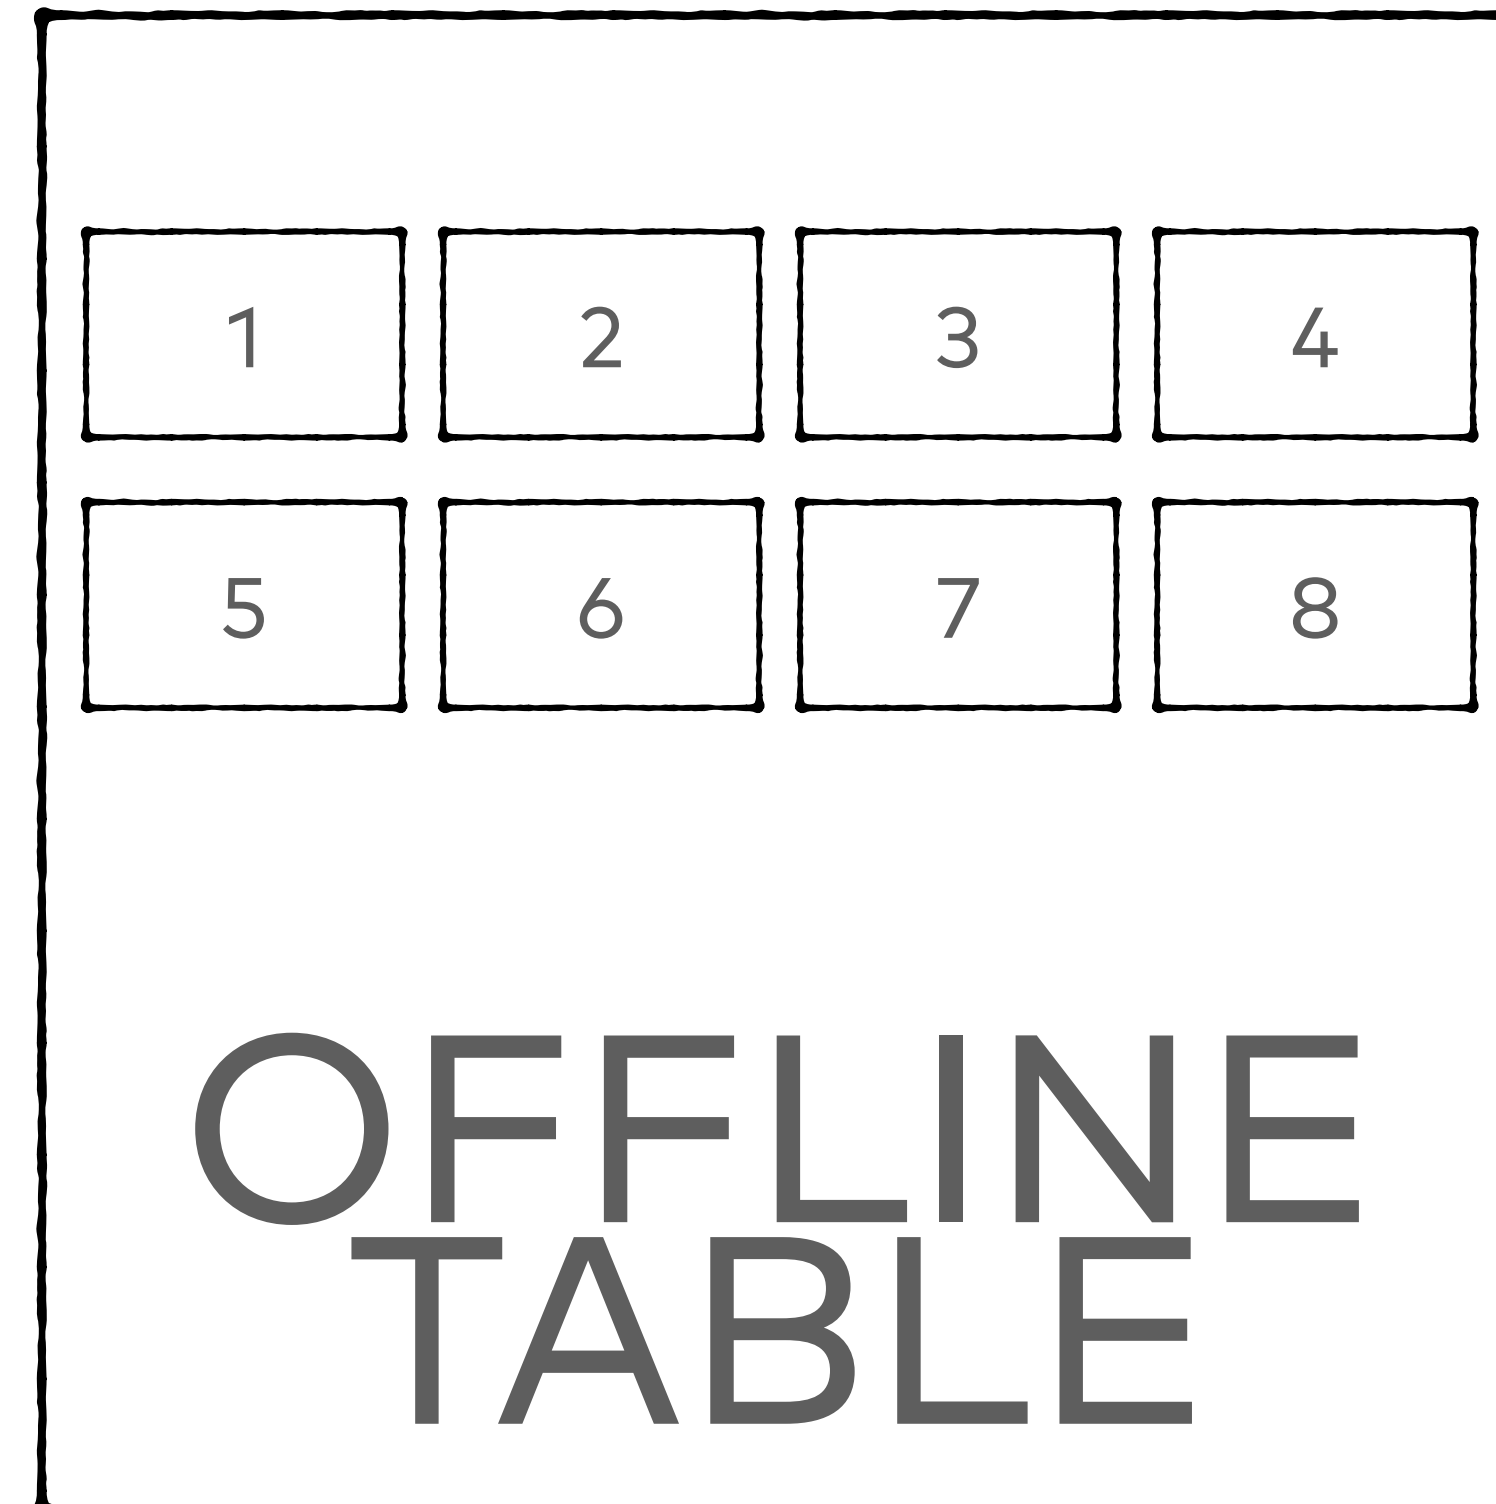


# All of Pinot





# Batch Ingestion

```
...,  
{  
  "ip": "111.173.165.103",  
  "userid": 10,  
  "remote_user": "-",  
  "time": "3271",  
  "_time": 3271,  
  "request": "GET",  
  "status": "406",  
  "bytes": "1289",  
  "referrer": "-",  
  "agent": "Mozilla/5.0"  
},  
...,
```


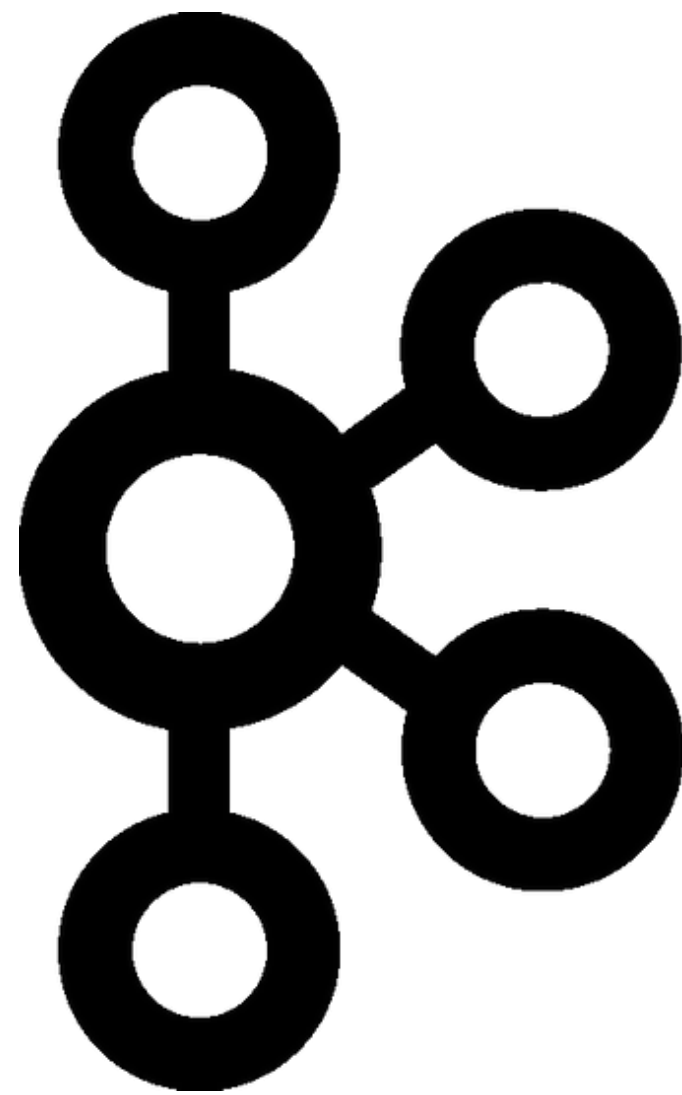






# Realtime Ingestion



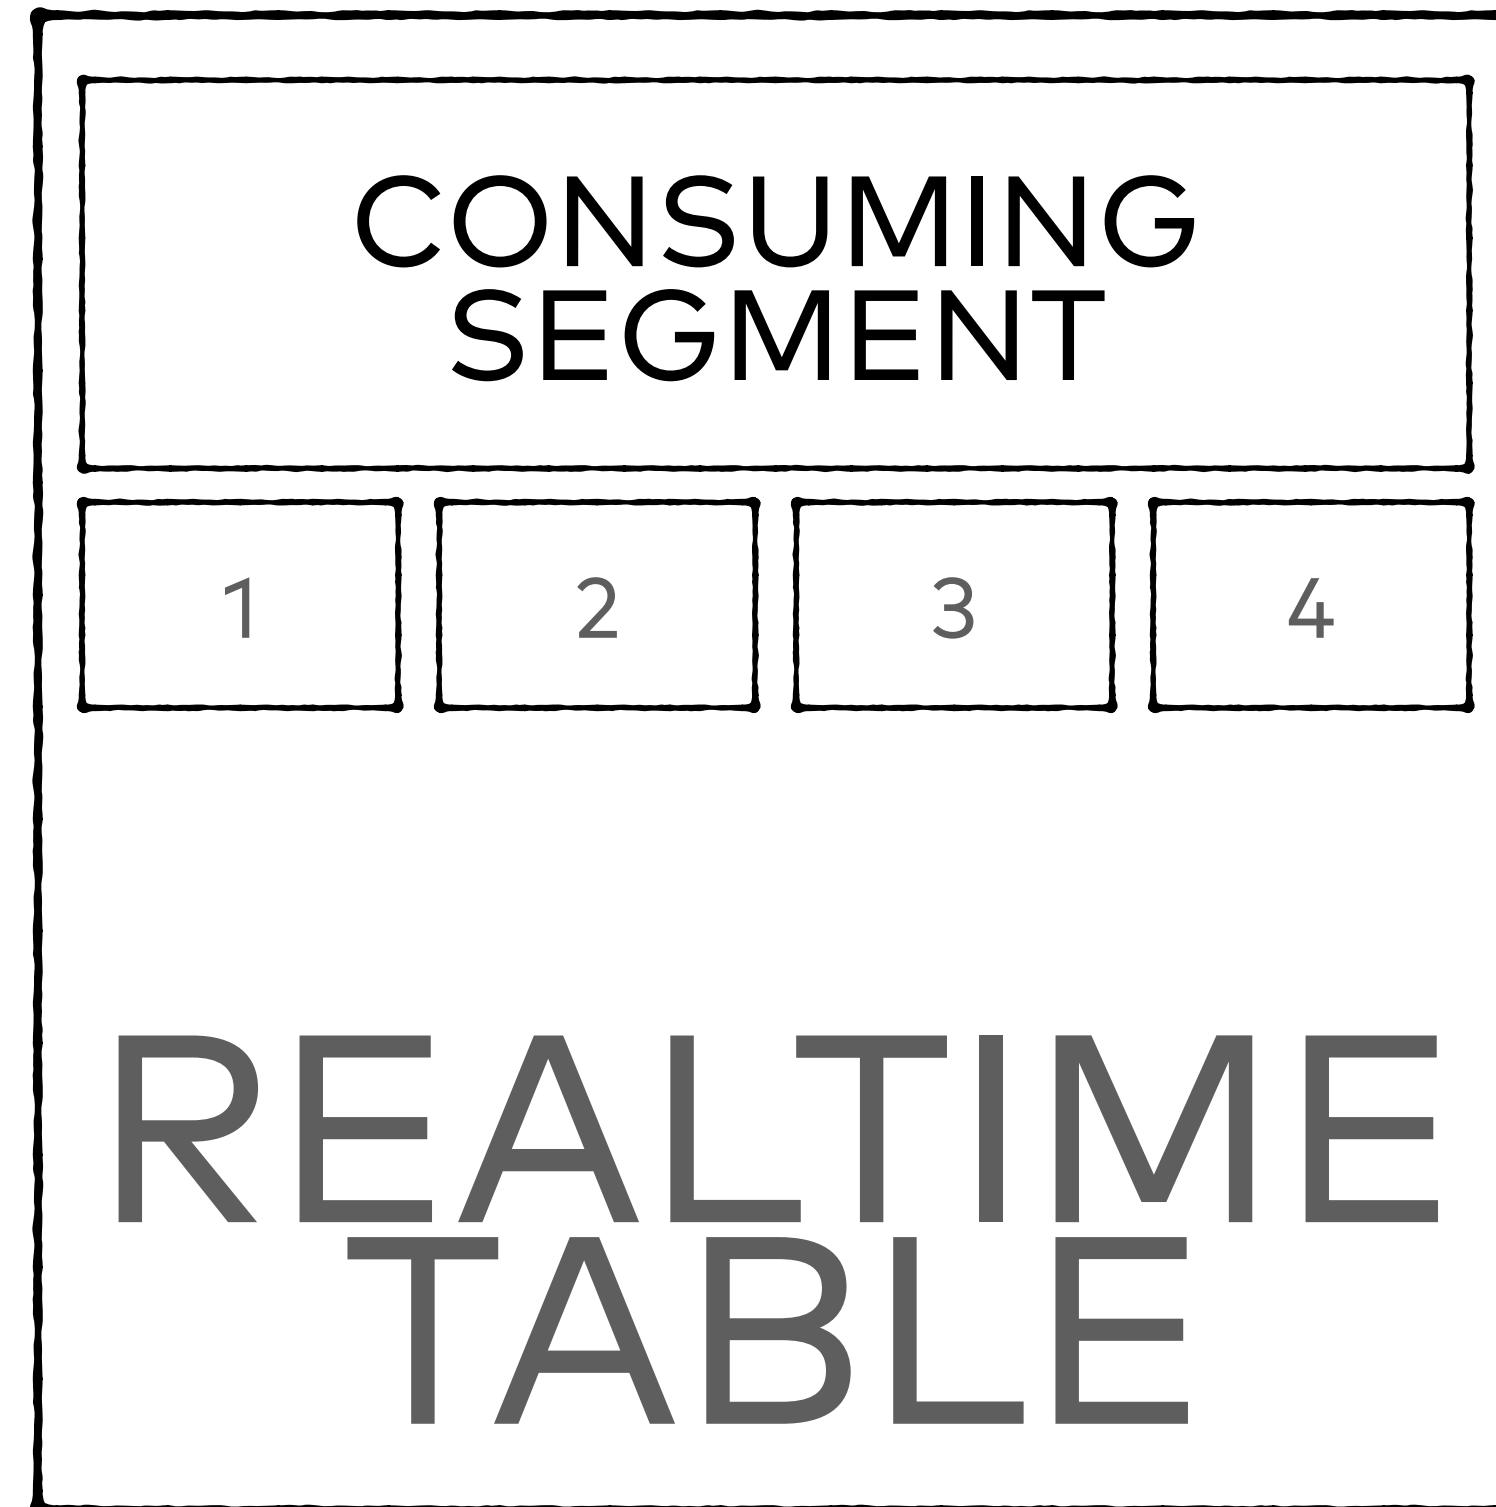

# Realtime Ingestion



# Realtime Ingestion





# Realtime Ingestion





# Realtime Ingestion



# Slack

<http://stree.ai/slack>

<https://stree.ai/pinot-slack>