Software project

Analysis and monitoring of product upload dynamics on the marketplace platform



About

This project develops a system to collect, process, and analyze large volumes of seller productupload events on the Ozon marketplace





About

This project develops a system to collect, process, and analyze large volumes of seller product-upload events on the Ozon marketplace

Focus areas

- Development of architecture
- Statistical metrics computation
- Visualization
- Automated alerting

Technologies





Problem statement

Ozon's marketplace generates a high-velocity stream of product load events from thousands of sellers.



Key challenges

- Handling millions of small, frequent events per day
- Ensuring data consistency and replication
- Delivering real-time analytics and alerts



Tasks

Step 1. Draft requirements

Domain description, functional & non-functional specs, constraints

Step 2. Compare alternatives

Step 4. Test the system

Step 3. Implement solution

- Implement database sharding locally (via zookeeper)
- Deploy ClickHouse, Grafana, Prometheus (via docker-compose)
- Create microservices that will simulate data stream and compute metrics
- Connect & register microservices (via consul)
- Write simulation script
- Create leader election for microservices
- Develop ETL pipelines & pre-aggregations
- Build anomaly detection modules

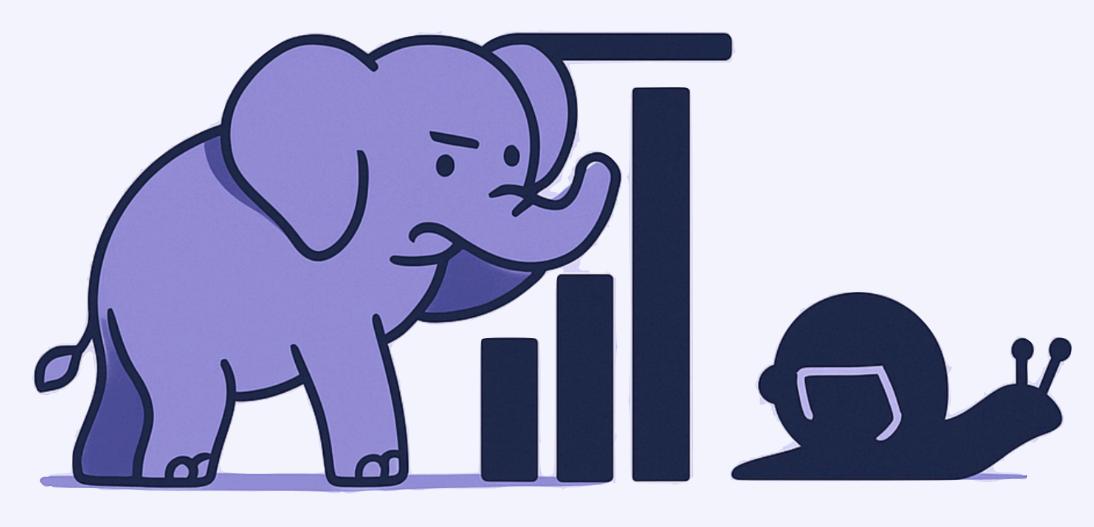


Limitations of PostgreSQL

Vertical scaling limits

WAL/MVCC overhead under heavy write load

Slow OLAP queries on growing datasets







Limitations of PostgreSQL

Scalability



only vertical or external sharding (Citus/Postgres-XL)

Storage



row-oriented, suboptimal for analytics.

Transactions



single-node ACID; no built-in 2PC across cluster

Maintenance



long VACUUM, index rebuilds on TB-scale tables

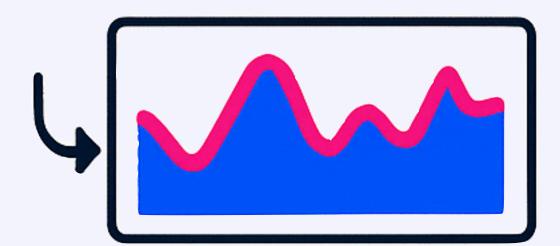
Write throughput



WAL contention and locking delays



Functional Requirements



The system has to:

- Survive the ingestion of maximum 1000 rps
- Keep one year of history online
- Provide control-charts for upload latency, failure rate, and seller anomalies
- Trigger Prometheus alerts when specific metrics exceed the limit



Alternative approaches



A highly scalable database optimized for fast writes, but less suitable for real-time analytics due to eventual consistency.



Apache Druid

A fault-tolerant, distributed analytics engine built for fast, interactive queries over large-scale time-series data.



Implementation



Architecture Overview

Synthetic data generation



id Unique identifier of the item

event

attempt_id

is_created

origin

up_ts

country

data

Type of event: "SAVE", "UPDATE", "ERROR"

Identifier for this upload attempt

Boolean flag

Source of the operation

If is_created=1, a ns ts when creation occurred; or 0

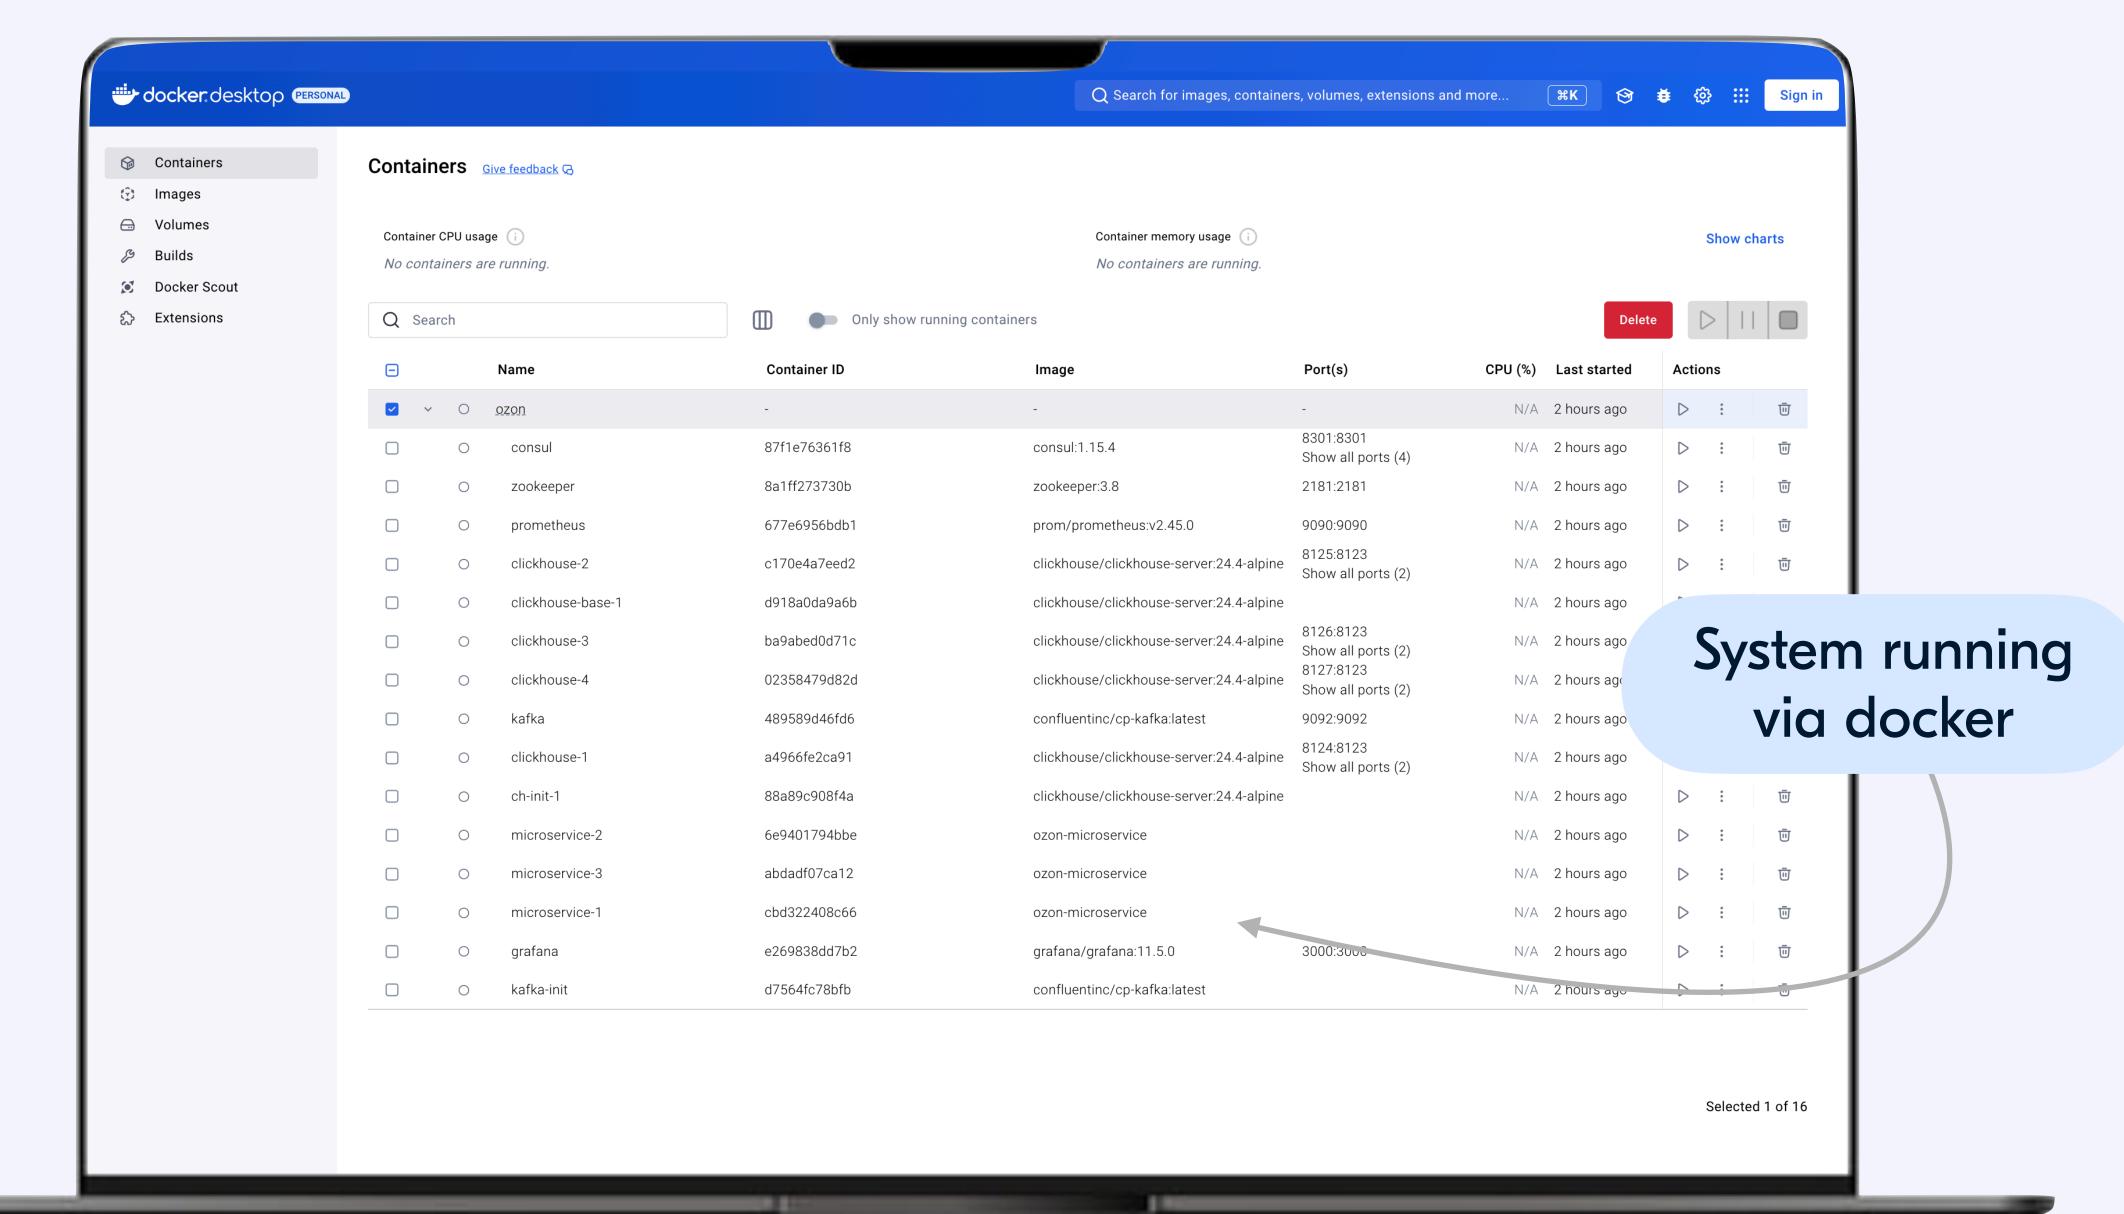
Production country code

JSON: info, company ID and list of media files



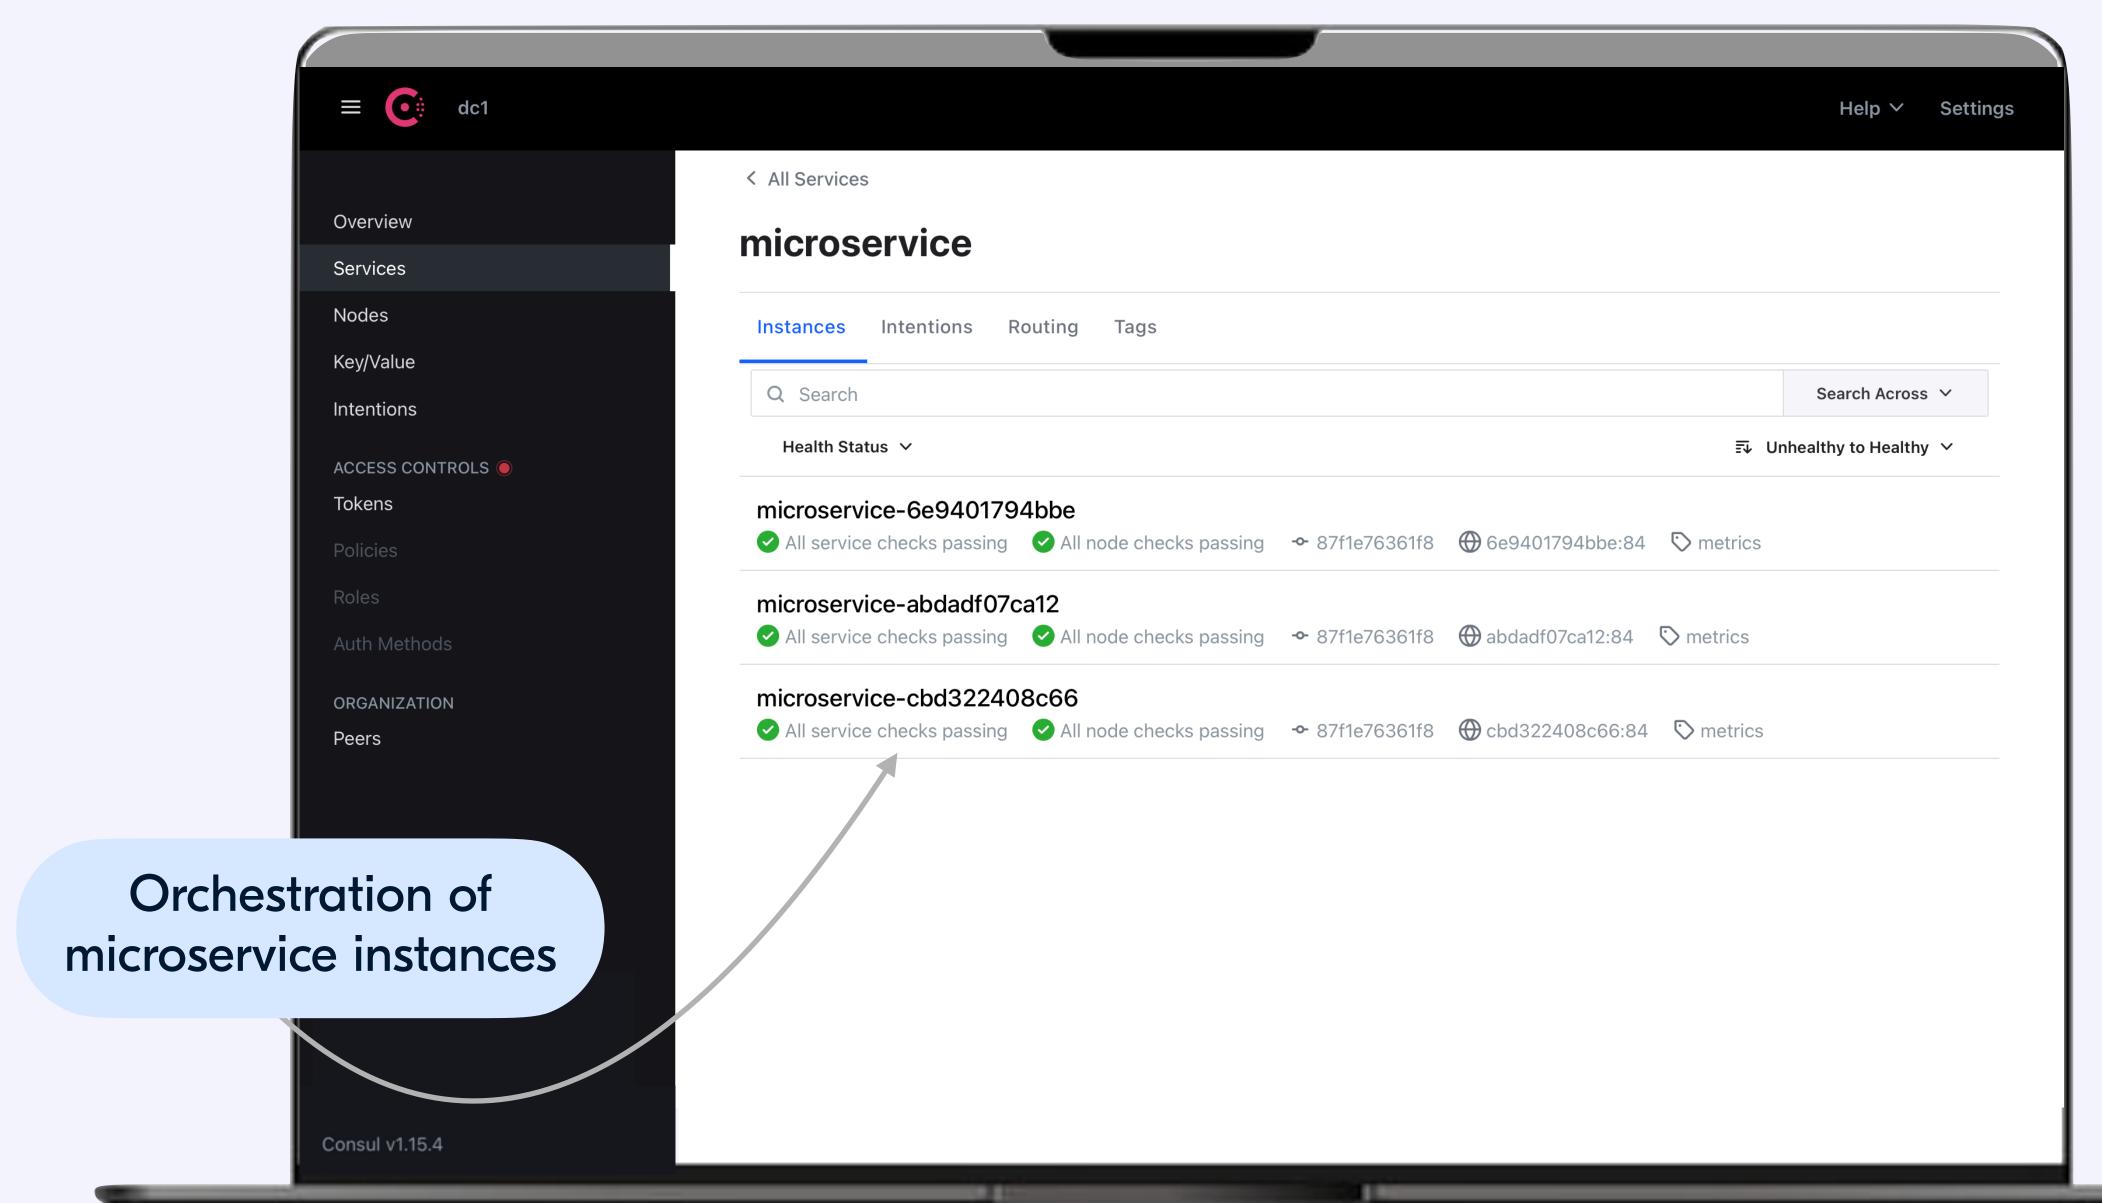
Architecture Overview

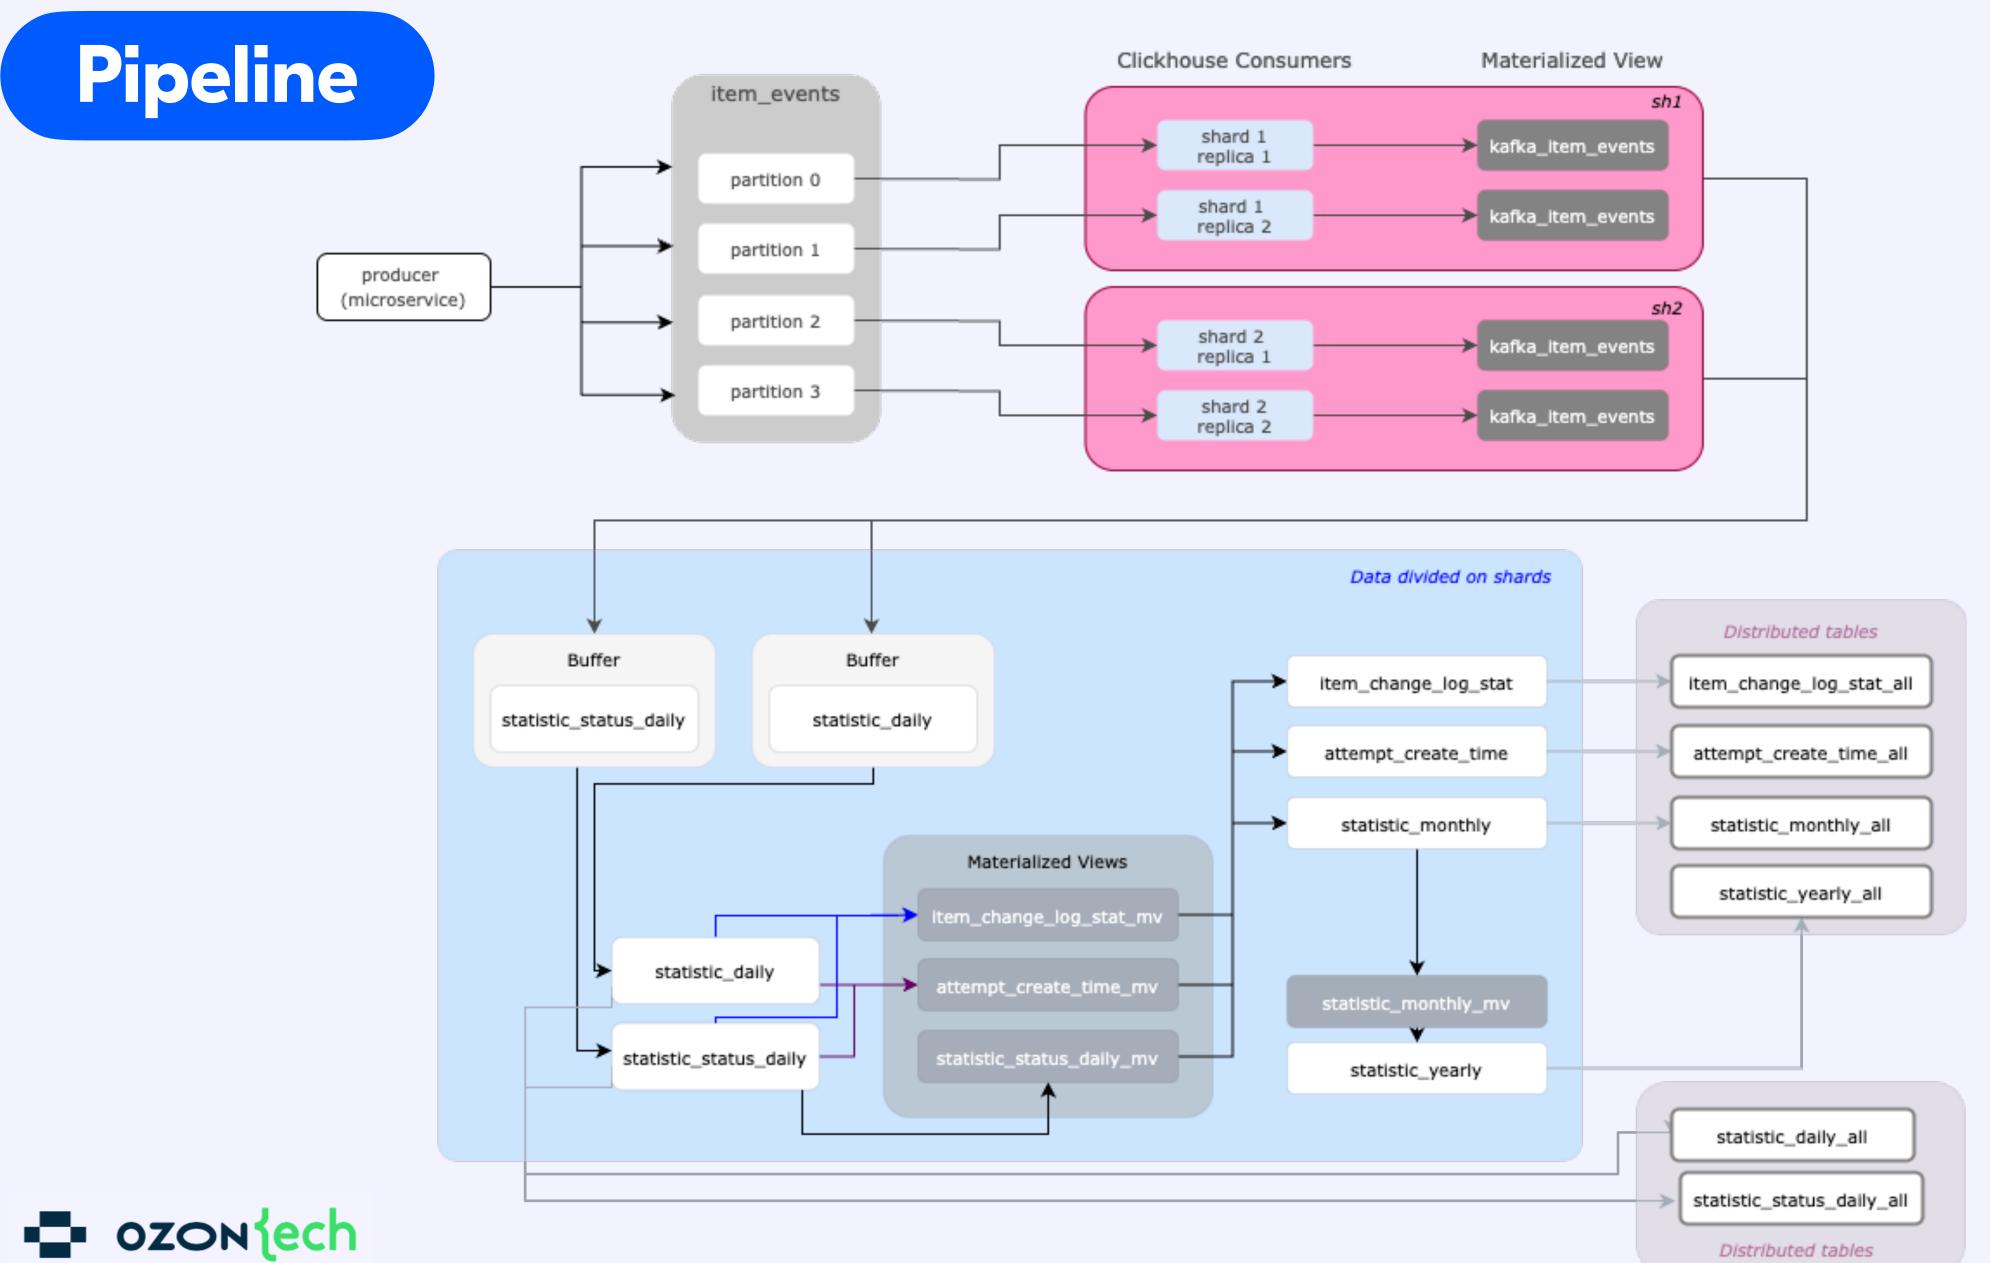




Architecture Overview









Pipeline

```
\leftarrow \rightarrow
                                                                                                                                                                                    \wp ozon
                                                                                                                                                                                                                                                                                                                    ! alert_rule: ▷ ∨ □ ···
                                                                       ddl.sql
                                                                                                                     simulate_data.py
                                                                                                                                                                                  compute_metrics.py ×
                                                                                                                                                                                                                                                         ! prometheus.yml
      EXPLORER
                                                                         microservice >  compute_metrics.py
\vee OZON
                                                                                                                                                                                                                                                                                                                                                        whether the second seco
                                                                                           def on_lose_leadership():

∨ clickhouse

                                                                           81
                                                                                                       logger.warning(

√ node01

                                                                            83
          macros.xmi
                                                                            84
        > node02
                                                                            85
       > node03
                                                                            86
                                                                                           def leader_election_loop():
                                                                                                      session_id = create_session()
                                                                           87
       > node04
                                                                                                      is_leader = False
                                                                            88
     > demo
                                                                                                      while True:
                                                                            89
     > grafana
                                                                            90
                                                                                                                 try:

→ microservice

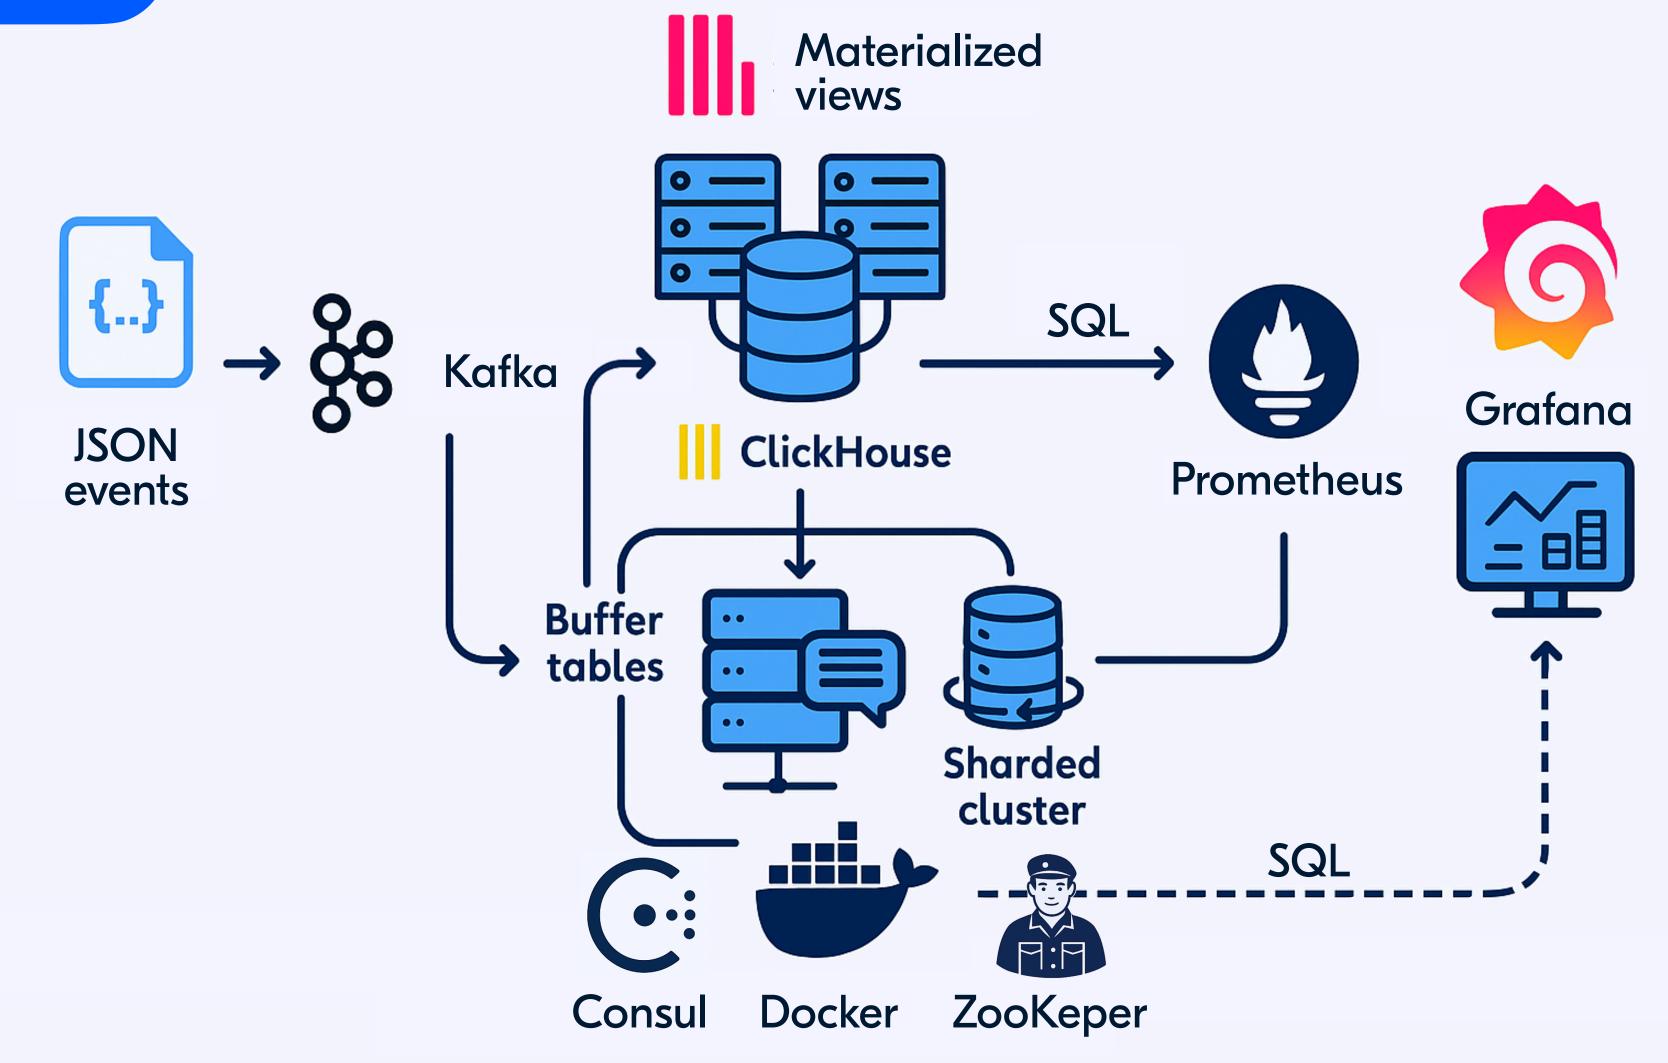
                                                                                                                            got = acquire_lock(session_id)
                                                                            91
                                                                           92
                                                                                                                            if got and not is_leader:
       compute_metrics.py
                                                                                                                                       is_leader = True
                                                                            93
      config.py
                                                                                                                                      on_become_leader()
                                                                            94
       connect.py
                                                                                                                            elif not got and is_leader:
                                                                            95
      ddl.sql
                                                                                                                                       is_leader = False
                                                                            96
      Dockerfile
                                                                                                                                       on_lose_leadership()
                                                                            97
                                                                           98
       register_to_consul.py
                                                                                                                            if not renew_session(session_id):
                                                                            99
         = requirements.txt
                                                                                                                                       logger.warning(f"

Session expired: {session_id}, creating new one")
                                                                         100
       simulate_data.py
                                                                                                                                      session_id = create_session()
                                                                         101

    ∨ prometheus

                                                                         102
                                                                                                                                      if is_leader:
                                                                         103
                                                                                                                                                 is_leader = False
        > data
                                                                         104
                                                                                                                                                 on_lose_leadership()
        > etc
                                                                         105
    • .gitignore
                                                                                                                 except Exception as e:
                                                                         106
   docker-compose.yml
                                                                                                                            logger.exception("△ Leader election loop error")
                                                                         107
   (i) README.md
                                                                         108
                                                                                                                 delta = BASE * JITTER_FACTOR
                                                                         109
                                                                                                                 sleep_time = BASE + random.uniform(-delta, delta)
                                                                         110
> VS CODE PETS
```

Pipeline





Results & Testing



Vizualization K6





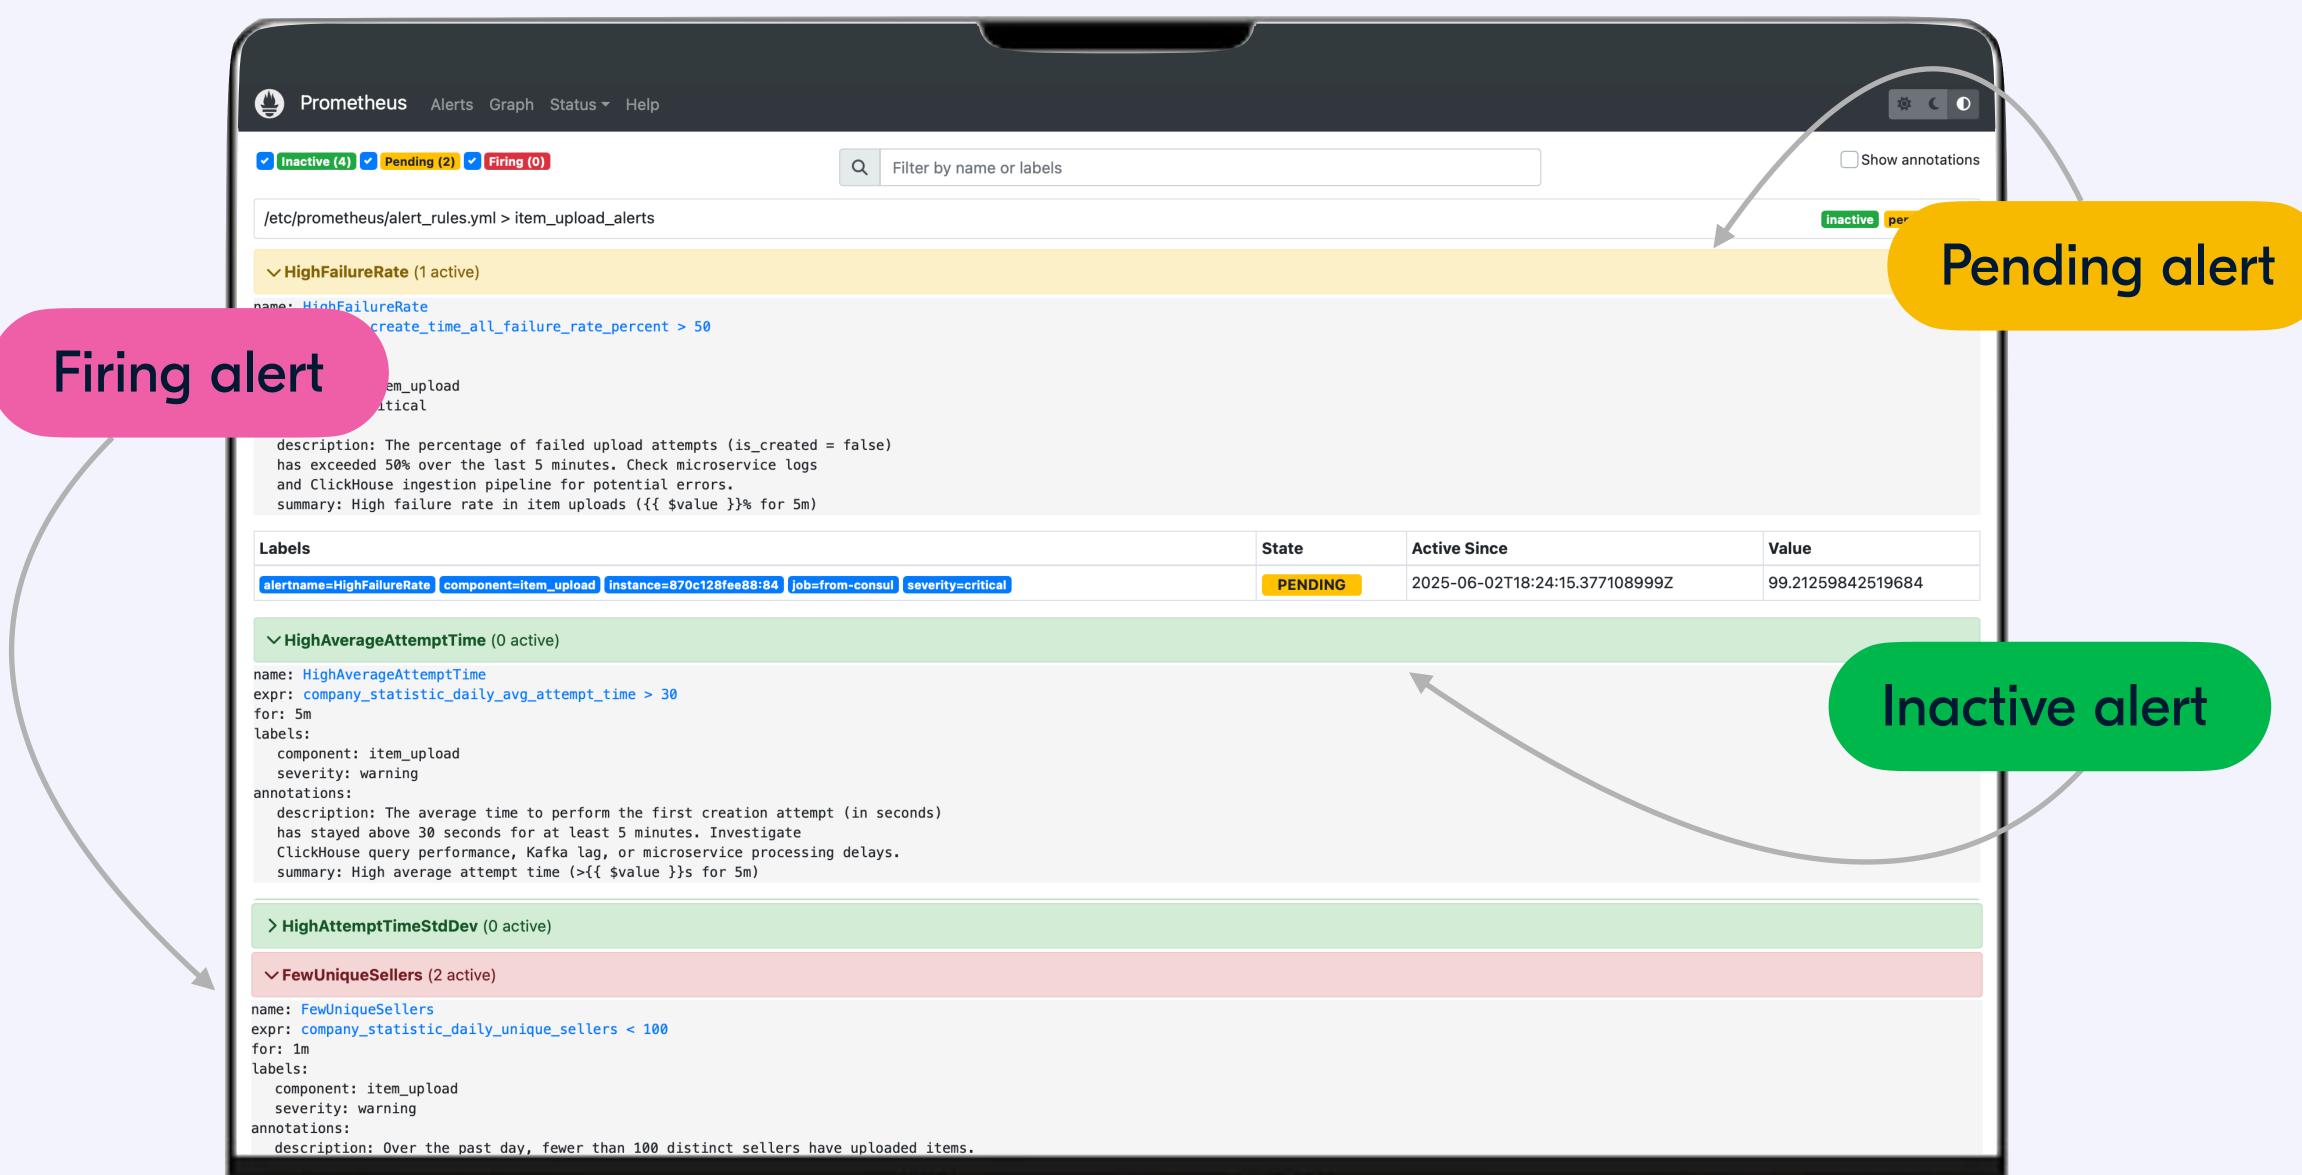
Vizualization K 6





Alerts





Reference:

Ozon Production Cluster

- 2 replicas (non-sharded)
- 8 vCPU, 48 GB RAM, 700 GB SSD
- 1,000 RPS peak

Test setup:

Local Sharded Cluster

- 2 shards × 2 replicas
- Load balanced across shards
- No increase in latency

- Linear scalability by adding shards;
- Ready to handle higher loads



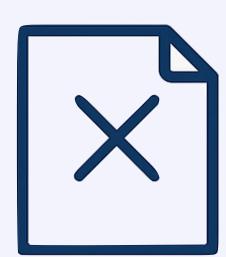
Issues

Data skew



Effect of random distribution in Python

Synthetic traffic



Lacks specific patterns

Limited testing



Pipeline with scripted chaos tests required

Cluster orchestration



Deployment with Kubernetes may be an improvement



References

- [1] ACID-transactions in PostgreSQL. Accessed: 2025-05-10. 2025. url: https://blog.ydb. tech/sharded-is-not-distributed
- [2] Denis Sheahan Adrian Cockcroft. Benchmarking Cassandra Scalability on AWS. url: https://netflixtechblog.com/benchmarking-cassandra-scalability
- [3] T. Akidau et al. Streaming Systems: The What, Where, When, and How of Large-Scale Data Processing. O'Reilly Media, 2018.
- [4] Apache Cassandra Documentation. Accessed: 2025-02-03. 2025. url: https://cassandra.apache.org/doc/latest/.
- [5] Apache Druid Documentation. Accessed: 2025-02-03. 2025. url: http://druid.apache.org/docs/latest/.
- [6] J. Carpenter and E. Hewitt. Cassandra: The Definitive Guide: Distributed Data at Scale. 2nd. O'Reilly Media, 2016.
- [7] Citus. Accessed: 2025-20-05. 2025. url: https://demirhuseyinn-94.medium.com/scaling-horizontally-on-postgresql-cituss-impact-on-database-architecture
- [8] ClickHouse Documentation. Accessed: 2025-02-03. 2025. url: https://clickhouse.com/docs/en/.
- [9] Horizontal Scaling PostgreSQL. Accessed: 2025-28-05. 2025. url: https://www.cockroachlabs.com/blog/limitations-of-postgres/.
- [10] Multiversion concurrency control. Accessed: 2025-18-05. 2025. url: https://en.wikipedia.org/wiki/PostgreSQL#Multiversion_concurrency_control_(MVCC).
- [11] Postgre XL, github. Accessed: 2025-18-05. 2025. url: https://github.com/pharosnet/postgres-xl.
- [12] Round-Robin scheduling algorithm. Accessed: 2025-07-10. 2025. url: https://en.wikipedia. org/wiki/Round-robin_scheduling.
- [13] Write-ahead logic. Accessed: 2025-18-05. 2025. url: https://en.wikipedia.org/wiki/PostgreSQL#Storage_and_replication.

