



# Powering Real-Time Workloads — Anytime, Anywhere.

Timeplus Introduction

June 2025

CONFIDENTIAL

"What's happening?"  
is overtaking  
"What happened?"



#### Capital Market

Monitor real-time market data for P&L, signals, risk and compliance.



#### Fraud Detection

Identify and mitigate fraudulent transactions as they happen.



#### User Engagement

Analyze live customer interactions and sentiments for faster support or services.



#### Cybersecurity

Prevent costly breaches by detecting and mitigating cyber threats in real time.



#### Observability

Monitor infrastructure telemetry for real-time pipeline, alerting, and action.



#### AI Inference

Generate instant responses and actions using real-time context for AI agents.



# The **Fastest** Unified Streaming Data Processing for *Real-Time Analytics, Telemetry and AI/ML*

High Throughput  
& Cardinality

Incremental  
Transformation

Mutable  
Processing

Stream & Batch  
Processing

Routing  
& Alerting

2021: Seed Round

2023: pre-A

Global Clients: 20+

**Low Latency**

Incremental processing for telemetry, transactions/CDC, any events, in milliseconds.

**High Efficiency**

SQL-native, resource-efficient, auto-scaling, by a single binary engine. 10x less TCO


**Local First**

Keep data & compute close to the source, across the edge, cloud, or BYOC

Headquarter in Santa Clara, California

TIMEPLUS IS TRUSTED BY

## Finance



## Online Entertainment and Technology




## Manufacturing



# Real-Time, Real Opportunity, Real Challenge

*Legacy historical analytics aren't designed for operational estate in real-time*



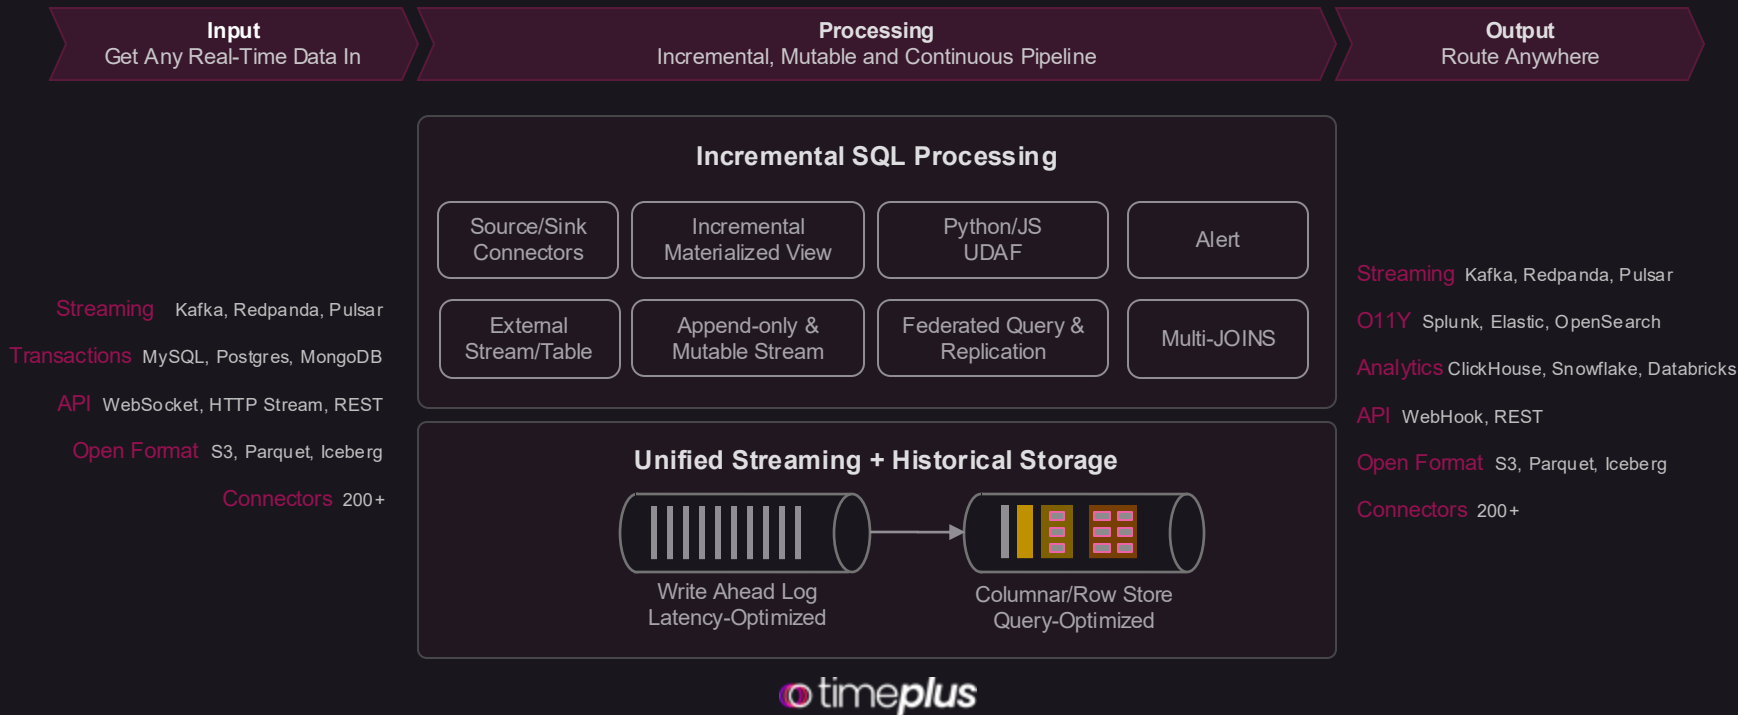
- Data is constantly generated, expanding and evolving at an unprecedented pace
- To stay ahead, companies need to ensure data quality from the source, gain rapid insights, and act fast
- Take control of this data chaos and avoid unnecessary \$\$\$

- Historical processing can't keep up with fast-changing, high-throughput data
- Centralized data creates privacy, security, and compliance challenges
- Snowflake or Databricks is too expensive for real-time workloads

**Your Real-time Workloads Demand A Purpose-built Platform  
Designed for Low-Latency, Efficiency and Local-First**

# Architecture Deep-dive

The only unified SQL platform for streaming, columnar and row stores in a single-binary



# One Engine for Any Real-Time Pipeline

## Real-time Analytics

Real-Time Computation and Data Pipelines for Analytics and Data Warehousing

OLTP/OLAP, Gaming, Fraud Detection, Trading

## Security/O11Y

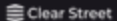



Real-Time CEP and Telemetry Pipelines for Logs, Metric, Trace and More

Telemetry Data Transformation, Enrichment, Routing and Alerting

## AI/ML

Real-Time Data Context for AI, ML and AI Agents

Real-Time Online Features, Continuous and Contextual Window

 Clear Street  HTSC  سلة  Byxtra

 Sony Interactive Entertainment  ROCNET  FTI  Li Auto

 HTSC  Sony Interactive Entertainment  CRESTA  WYNN PALACE



**One Real-Time Incremental SQL Engine for "Collect. Transform. Route. Alert — All in Milliseconds."**

- Supports millions of events per second with ultra-low latency
- Powers dynamic, multi-stream JOINS for complex real-time scenarios

# Timeplus Enterprise Features

Security  
Authn/Authz

API Gateway  
(HTTP/WebSocket Gateway)

Console  
(Data Exploration, Lineage,  
Dashboards, Alerts)

Timeplus Connectors  
(200+ Sources/Sinks/Alerts)

System Monitoring Dashboards and Plugins

High Availability  
Sharding/Replication/Recovery (Multi-RAFT)

Data Federation and Virtualization  
(across clusters)

Timeplus Core  
Fast Streaming Engine + Multi-modal storage (indexes)

Mutable Streams  
(indexing, column families)

Bare Metal, Kubernetes, BYOC



## Collection

Built-in external streams and external tables to natively collect real-time data from, or send data to: Apache Kafka, Apache Pulsar, Confluent Cloud, Redpanda, ClickHouse, or another Timeplus instance. Timeplus also supports a wide range of data sources through sink/source connectors. Plus, our partnership with Redpanda Connect brings you another 200+ systems.



## Transformation

With a powerful streaming SQL console, users can leverage their preferred query language to create Streams, Views, and incremental Materialized Views. This enables them to transform, roll up, join, correlate, enrich, aggregate, and downsample real-time data, generating meaningful outputs for real-time alerting, analytics, or any downstream systems.



## Routing

Timeplus allows data to be routed to different sinks based on SQL-based criteria and provides a data lineage view of all derived streams in its console. A single data result can generate multiple outputs for various scenarios and systems, such as analytics, alerting, compliance, etc., without any vendor lock-in.



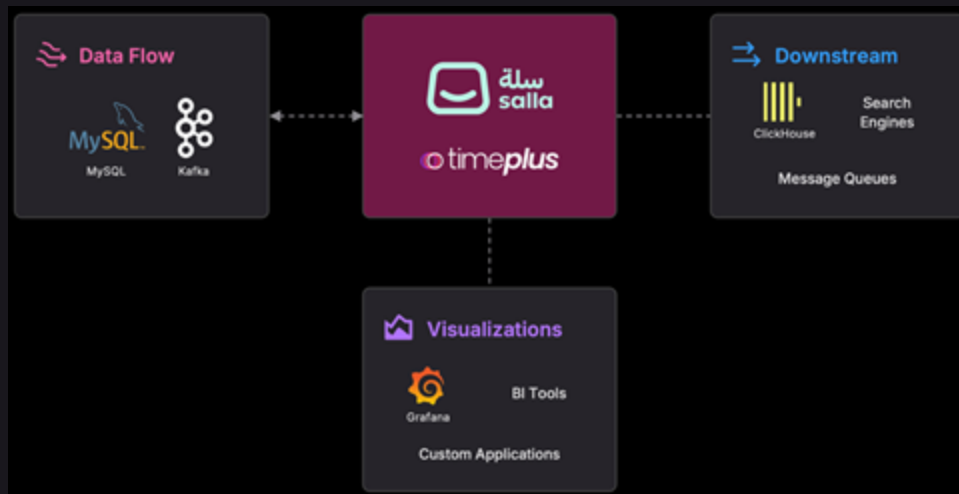
## Analytics and Alerting

Powered by SSE (Server-Sent Events), Timeplus supports push-based, low-latency dashboards to visualize real-time insights through data pipelines or ad-hoc queries. Additionally, users can easily build observability dashboards using Grafana plugins. SQL-based rules can be used to trigger or resolve alerts in systems such as PagerDuty, Slack, and other downstream platforms.



# Timeplus + Real-time Incremental Data Pipeline

Large Scale Incremental, Denormalized Streaming ETL in E-Commerce



## Challenges

- MySQL CDC transformations to update data in real-time
- Reduce Latency serving to different downstream systems
- Data Deduplication & Replay to reduce inconsistency and resource usage on Analytical systems
- Data Enrichment
- Incremental Update of Denormalized Data Views for reports, billing, AI analytics
- Streaming data to other downstream (Search, Queues, Alerting)

## Results After Timeplus

↓ 67%

Reduction in CDC pipelines overhead



800+ Materialized Views

Migrated from ClickHouse, freeing up CPU/memory



Minimized latency

Processing lag reduced from minutes to seconds

# Timeplus + Real-time Trading Monitoring

10x performance improvement over existing investment, market making and risk control monitoring

## Investment Analysis

Long Position Yield

- Input: Transaction data, position data, market data
- Time window: 1-minute rolling window
- Compute: Combination of 3000 stocks \* 20 accounts

## Market Making

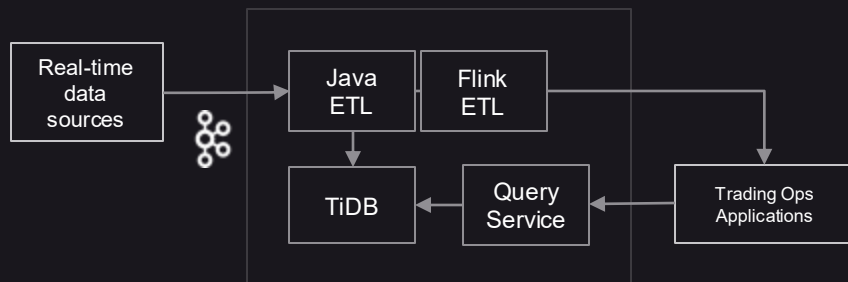
Continuous Auction Participation Rate

- Input: Order data, market making parameters
- Time window: 1-second rolling window
- Compute: 3000 individual stocks

## Risk Control

Intraday Price Manipulation Monitoring

- Input: Transaction data, market data
- Time window: 3 minute event-triggered sliding window
- Calculation: one strategy for each combination of 3000 stocks



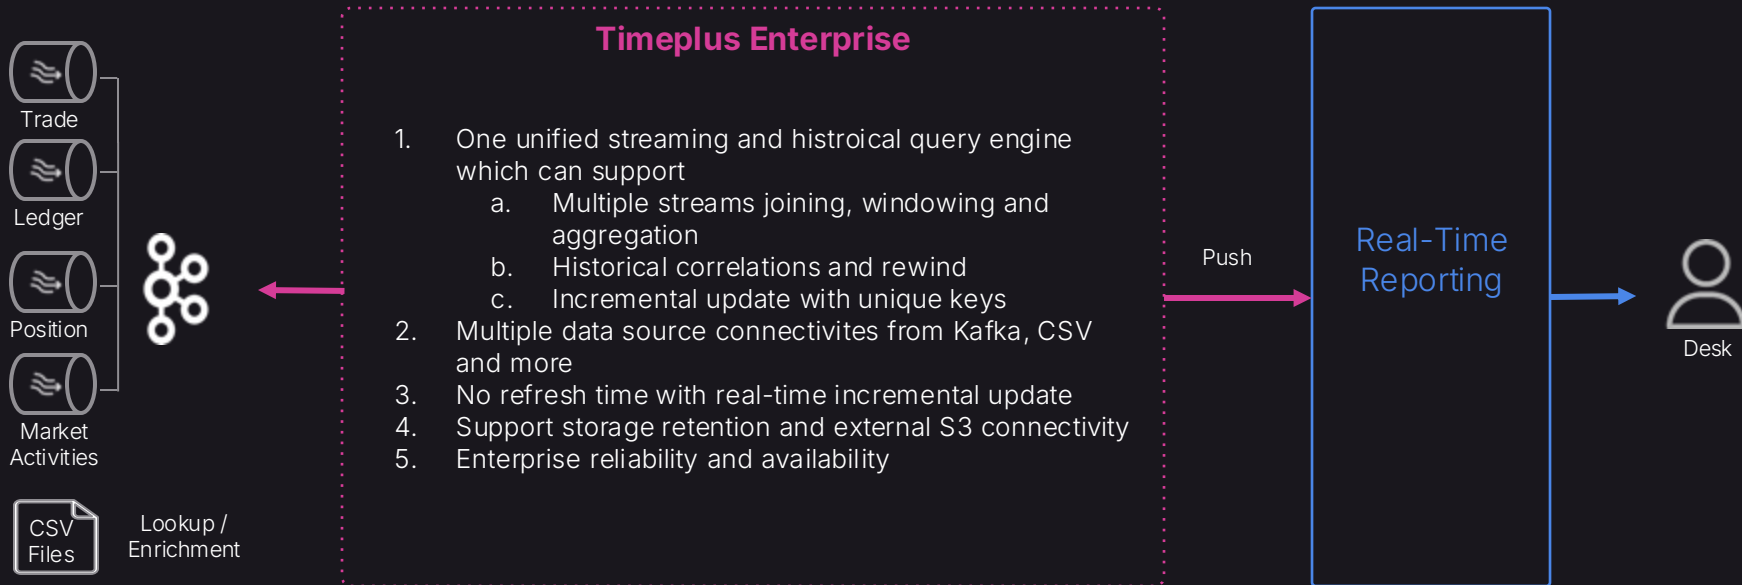
Before



After

# Timeplus + Real-Time Trading Desk Insights

A real-time reporting system which supports trading desk operations. It aggregates live data on trades, settlements, stock borrowing, and loans to help traders make time-sensitive decisions such as recalling loans or lending assets.



# Timeplus + Real-time Wagering Monitoring

14x performance improvement than existing wagering monitoring solution

## Odd Analysis

- **Input** Horse bet ticket data(WIN, DBL, QUARTET), 50,500+ tickets/s
- **Process** Real-time Wagering Metrics Monitoring
- **Challenges** Existing infrastructure is too slow (30+ s) to get the betting odds calculated) and too complicated



1. Amount wagered per bet
2. Number of bets placed in a betting pool
3. Total amount wagered in a pool
4. Total commission collected from a pool
5. Total number of bets placed in a race
6. Total amount wagered in a race
7. Total commission collected from a race



Race Day Data (Win)

Start Stop Delete

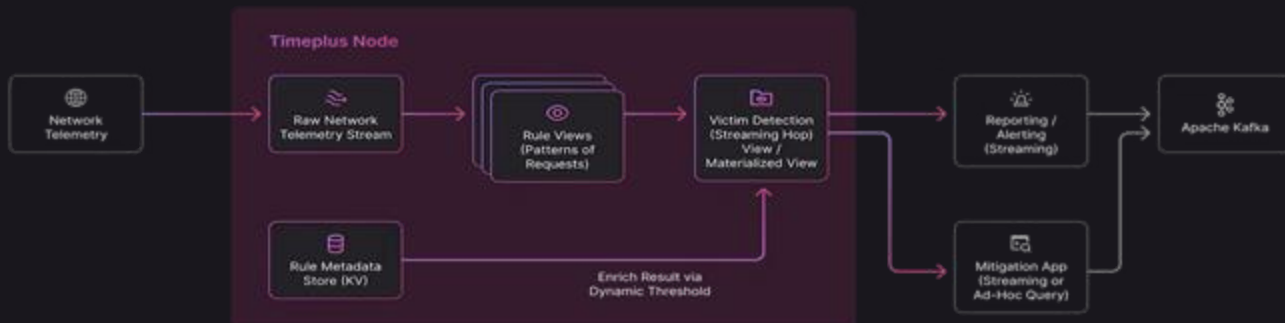
Horse	Race 1	Race 2	Race 3	Race 4	Race 5	Race 6	Race 7	Race 8	Race 9	Race 10	Race 11
01	14.73	-	60.66	59.12	9.72	45.62	17.16	-	62.35	123.75	159.93
02	34.94	26.92	2.96	58.26	5.67	10.58	15.75	14.22	85.71	10.77	1.99
12	9.45	-	58.72	24.57	11.71	6.33	117.89	-	2.08	31.35	18.87
03	13.6	12.93	16.92	59.28	58.64	20.57	1.89	63.67	6.35	61.87	7.39
13	8.69	-	69.08	44.24	97.19	5.59	164.03	-	79.39	-	26.88
14	53.13	-	47.9	30.09	96.73	24.73	34.3	-	153.52	5.69	99.99
03	13.96	41.96	9.51	12.18	78.95	123.34	164.53	65.31	94.68	17.66	23.54
04	46.86	65.53	61.53	10.01	3.33	11.53	24.06	1.78	16.66	26.83	42.67
05	5.48	1.54	26.04	78.86	47.53	65.37	13.56	70.01	7.29	32.66	25.25
06	19.11	6.41	56.42	112.92	8.95	41.58	19.04	15.89	41.96	5.03	31.67
07	8.46	24.01	6.49	53.75	13.19	3.42	33.93	14.62	27.69	5.55	113.97
08	91.14	-	71.54	27.51	61.15	60.63	80.55	45.07	36.26	9.12	14.78
09	9.03	-	21.88	1.96	19.35	130.1	141.15	8.49	163.29	46.3	52.88
10	43.98	-	7.02	11.87	117.63	143.29	53.37	19.11	75.63	25	57.73



# Timeplus + Real-time Cybersecurity

Detect attacks at a global scale in the hierarchical and distributed network infrastructure

To have a continuous DDoS Detector and Mitigator which runs on a set of User Defined Domain Rules to detect and block attacks of different types DDoSing machines in the customer network.



```


1 WITH aggregated AS (
2   SELECT
3     window_start AS attack_start_time,
4     ipv4_dst_ip AS victim_ip,
5     sum(byte_count) AS sum_bytes,
6     'ddos' AS rule_id
7   FROM hop(ipfix_flows, 1s, 30s)
8   SHUFFLE BY ipv4_dst_ip
9   GROUP BY window_start, ipv4_dst_ip
10  SETTINGS max_threads=8
11 )
12 SELECT attack_start_time, victim_ip, sum_bytes, rules.id, rules.version, rules.bps_threshold
13 FROM aggregated LATEST JOIN rules ON aggregated.rule_id = rules.id
14 WHERE sum_bytes >= rules.bps_threshold;
  
```

- Very low latency to process extreme high throughput calculations (millions events per second)
- Candidate detections in under 1 min with very high cardinality of destination IPs
- Latest join with dynamic rules which can be updated on the fly
- Deeper look with fast Top K queries for alerts and reporting
- Very low footprint: 100MB memory with 2m cardinalities aggregations

# Timeplus + Real-time Cybersecurity

Detect attacks with real-time CEP

MITRE T1218.010: Squiblydoo attack : To have a continuous threat detection with UDF based CEP rules.




```
SELECT raw:`process`.`pid` as pid, mt_1218_attack(_tp_time,
raw) FROM threat GROUP BY pid
```

## UDF: mt\_1218\_attack.js

```
var PROTON_FUNC = {
  has_customized_emit: true,
  initialize: function() {
    this.infected = false;
    .....
    this.result = [];
  },
  process: function (event_ts, events) {
    for (let i = 0; i < event_ts.length; i++) {
      let event = JSON.parse(events[i]);
      if (this.infected == false && event['event']['category'] == 'library' &&
event['dll']['name'] == 'scrobj.dll') {
        this.infected = true;
        this.infect_raw_event = JSON.stringify(event);
        this.infect_ts = event_ts[i];
      } else if (this.infected == true && event['event']['category'] == 'network') {
        this.attack_ts = event_ts[i];
        this.attack_raw_event = JSON.stringify(event);
        this.result.push({
          'infect_ts': this.infect_ts,
          'infect_evt': this.infect_raw_event,
          'attack_ts': this.attack_ts,
          'attack_evt': this.attack_raw_event
        });
      }
    }
  }
}
```

# Timeplus + Real-Time Device Monitoring



## Challenges

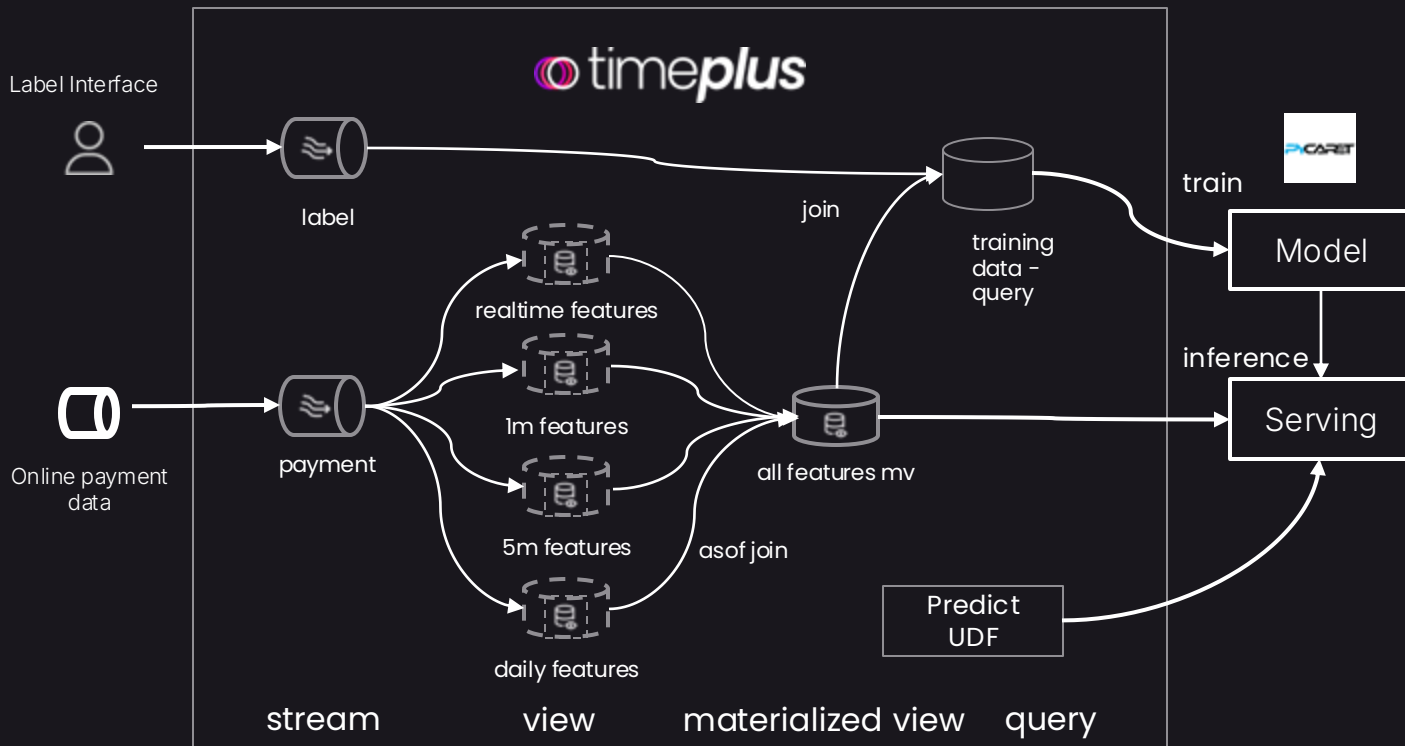
- Streaming Tech debt from previous suppliers/contractors needs to be understood and maintained
- Complex architecture and data migration to own cloud account requires safe migration path
- Spiraling costs in SaaS solution (Multiple idle environments for QA/Staging/UAT/Prod)

## Why Timeplus?

- Enrichment capability directly from CRM/CMDB to save duplication of storage in Kafka
- SQL easy to train up new team as well as maintain
- SQL for stream processing and analysis can be used to validate migration successful
- Data migration path (Kafka topics) to self-hosted doable in reversible steps with Timeplus

# Timeplus + Real-time Context for AI

Empowering a unified online + offline context/feature stores for AI/ML models



## Feature freshness

Low latency streaming processing  
Support those time range related, stateful features processing with TVF

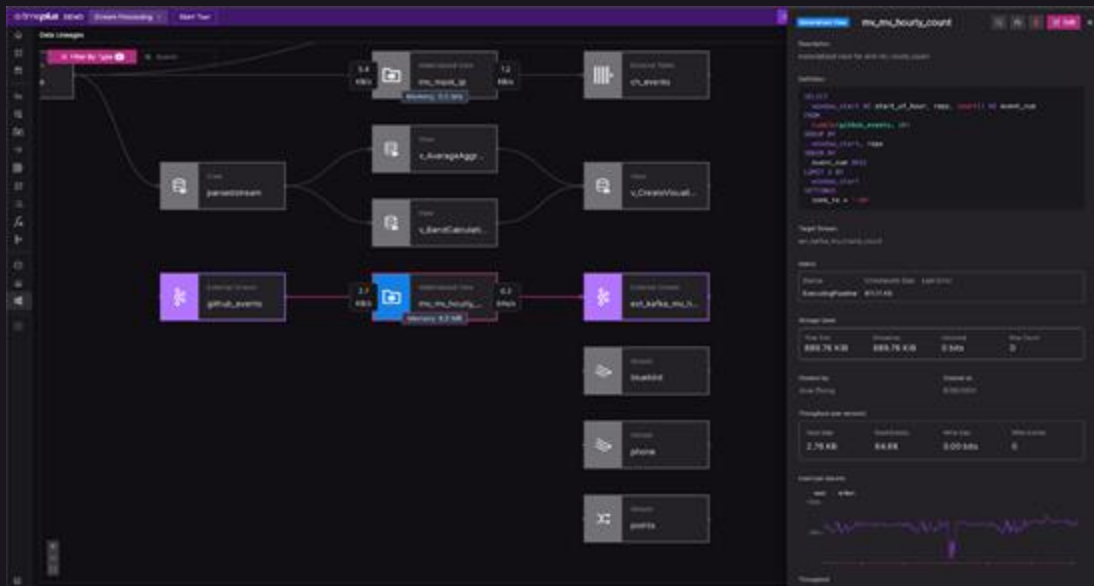
## Feature consistency

Unified streaming/historical processing  
Time travel/ASOF Join for PITC

## Simplicity

Solving the online and offline requirements for a unifying data layer, simplified overall ML application deployment  
Provide most of functionalities for the feature generation in SQL





Experience the Future of Real-time Intelligence.  
Anytime, Anywhere

<https://demos.timeplus.com>

<https://www.timeplus.com/download>