

Apache Doris: An Alternative Lakehouse Solution for Real-Time Analysis

Mingyu (Rayner) Chen

Apache Doris PMC Chair

VP of Technology at VeloDB



Contents

01 What is Apache Doris

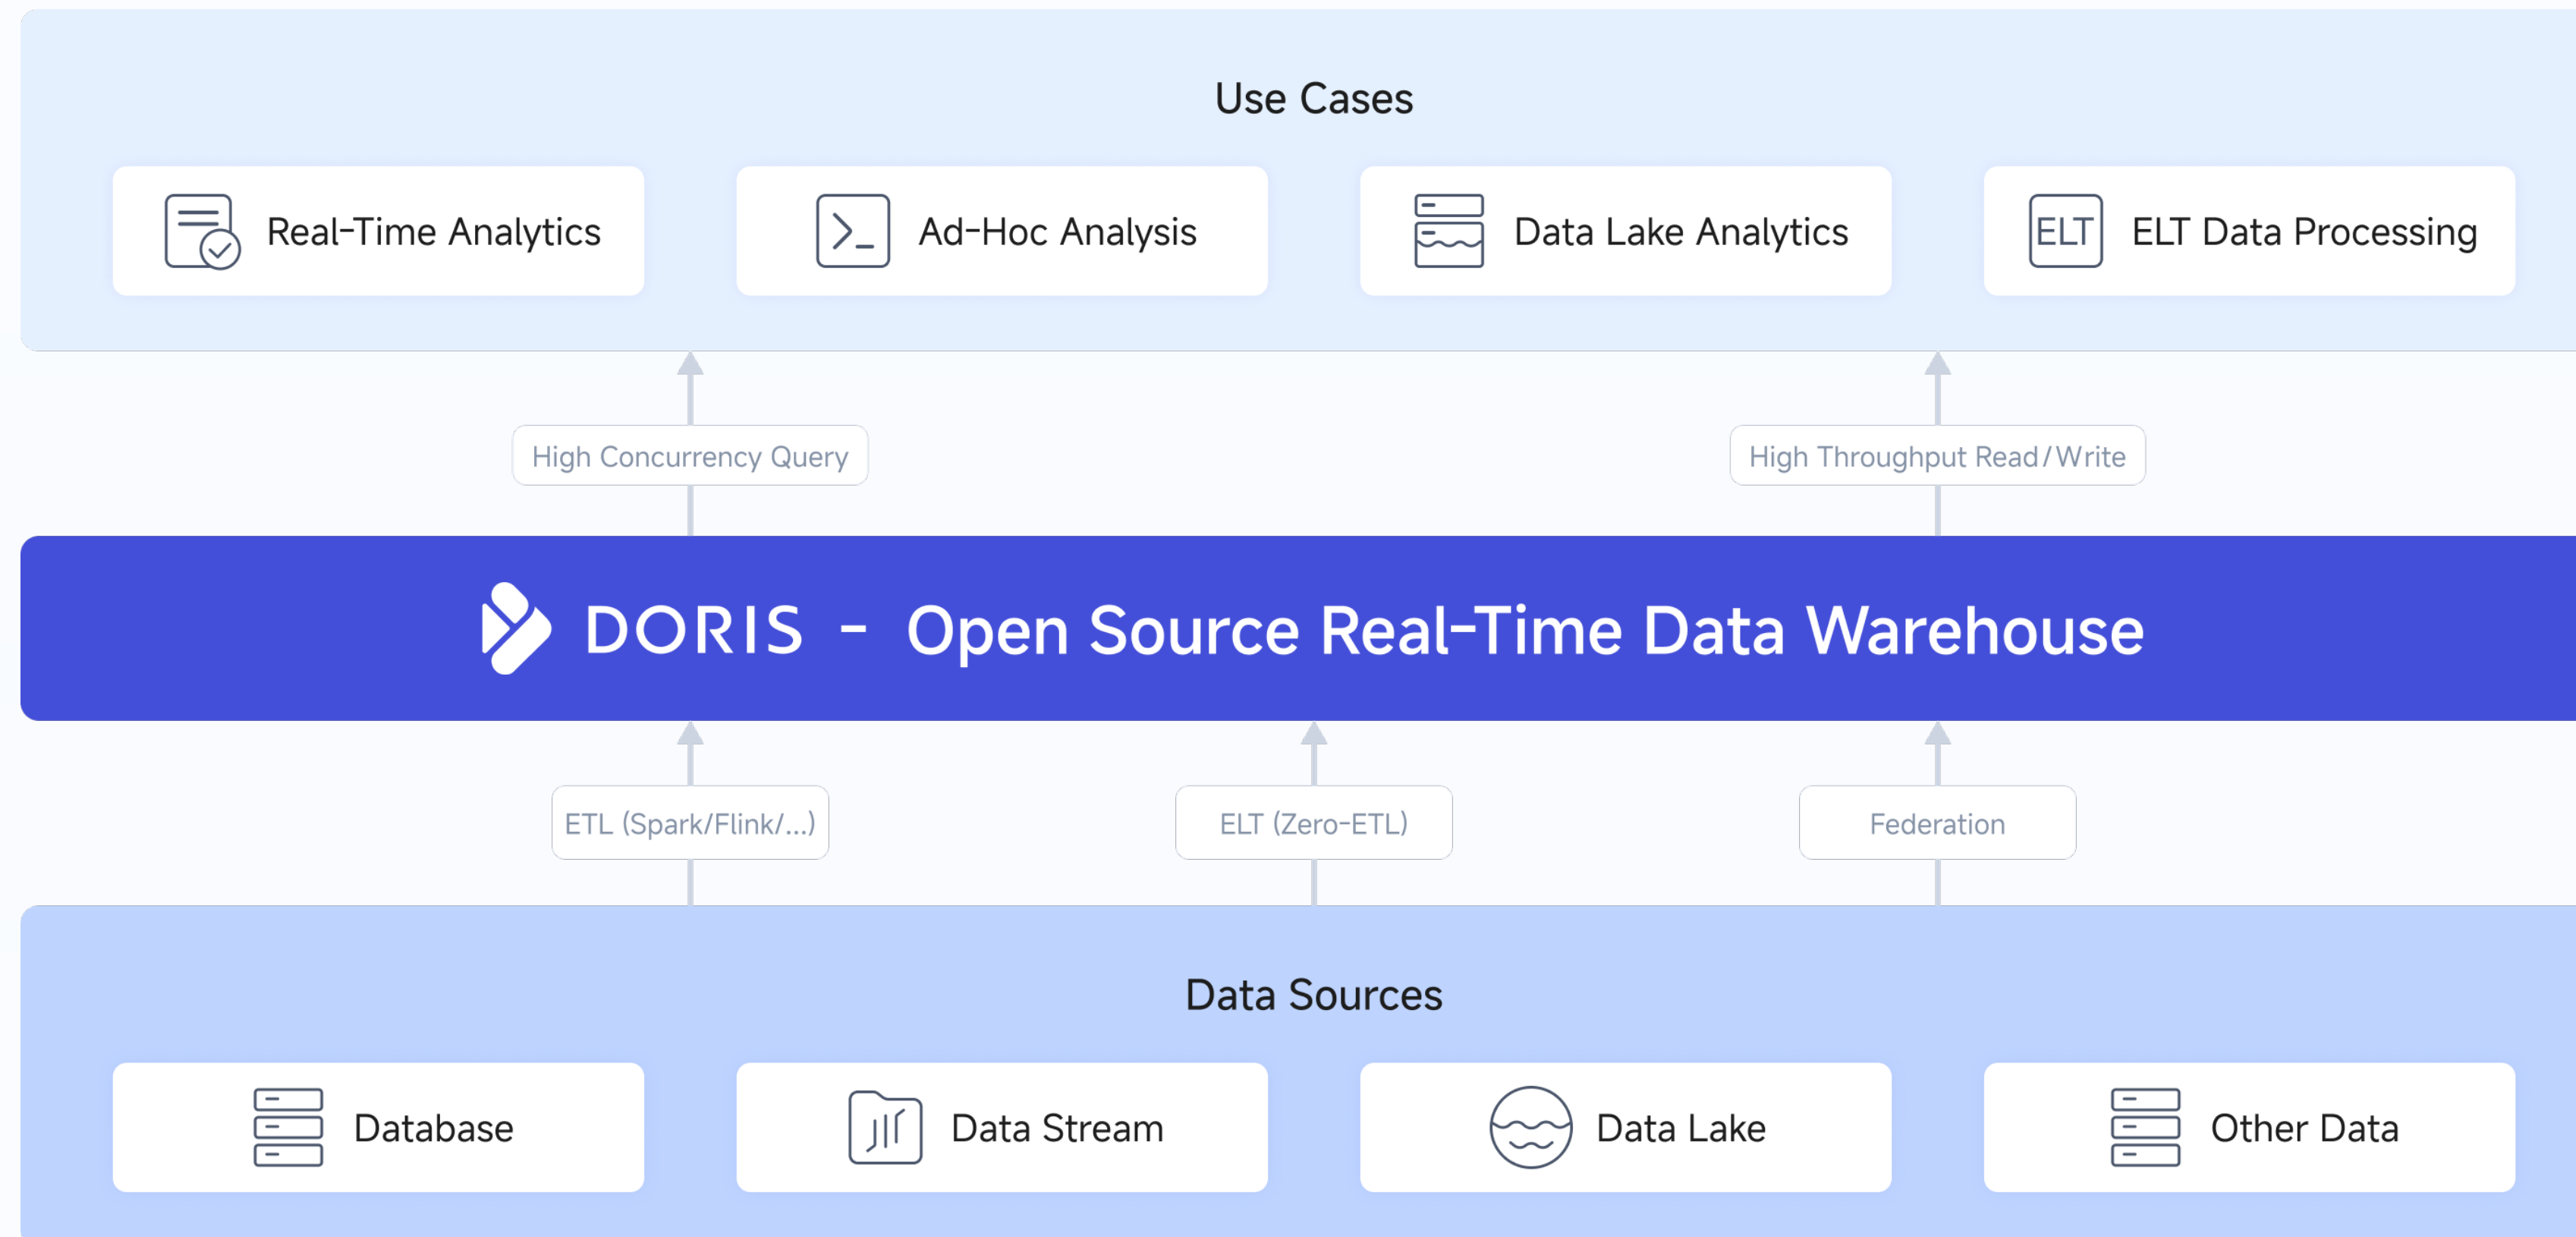
02 Building Lakehouse on Doris

03 Apache Doris Community

What is Apache Doris

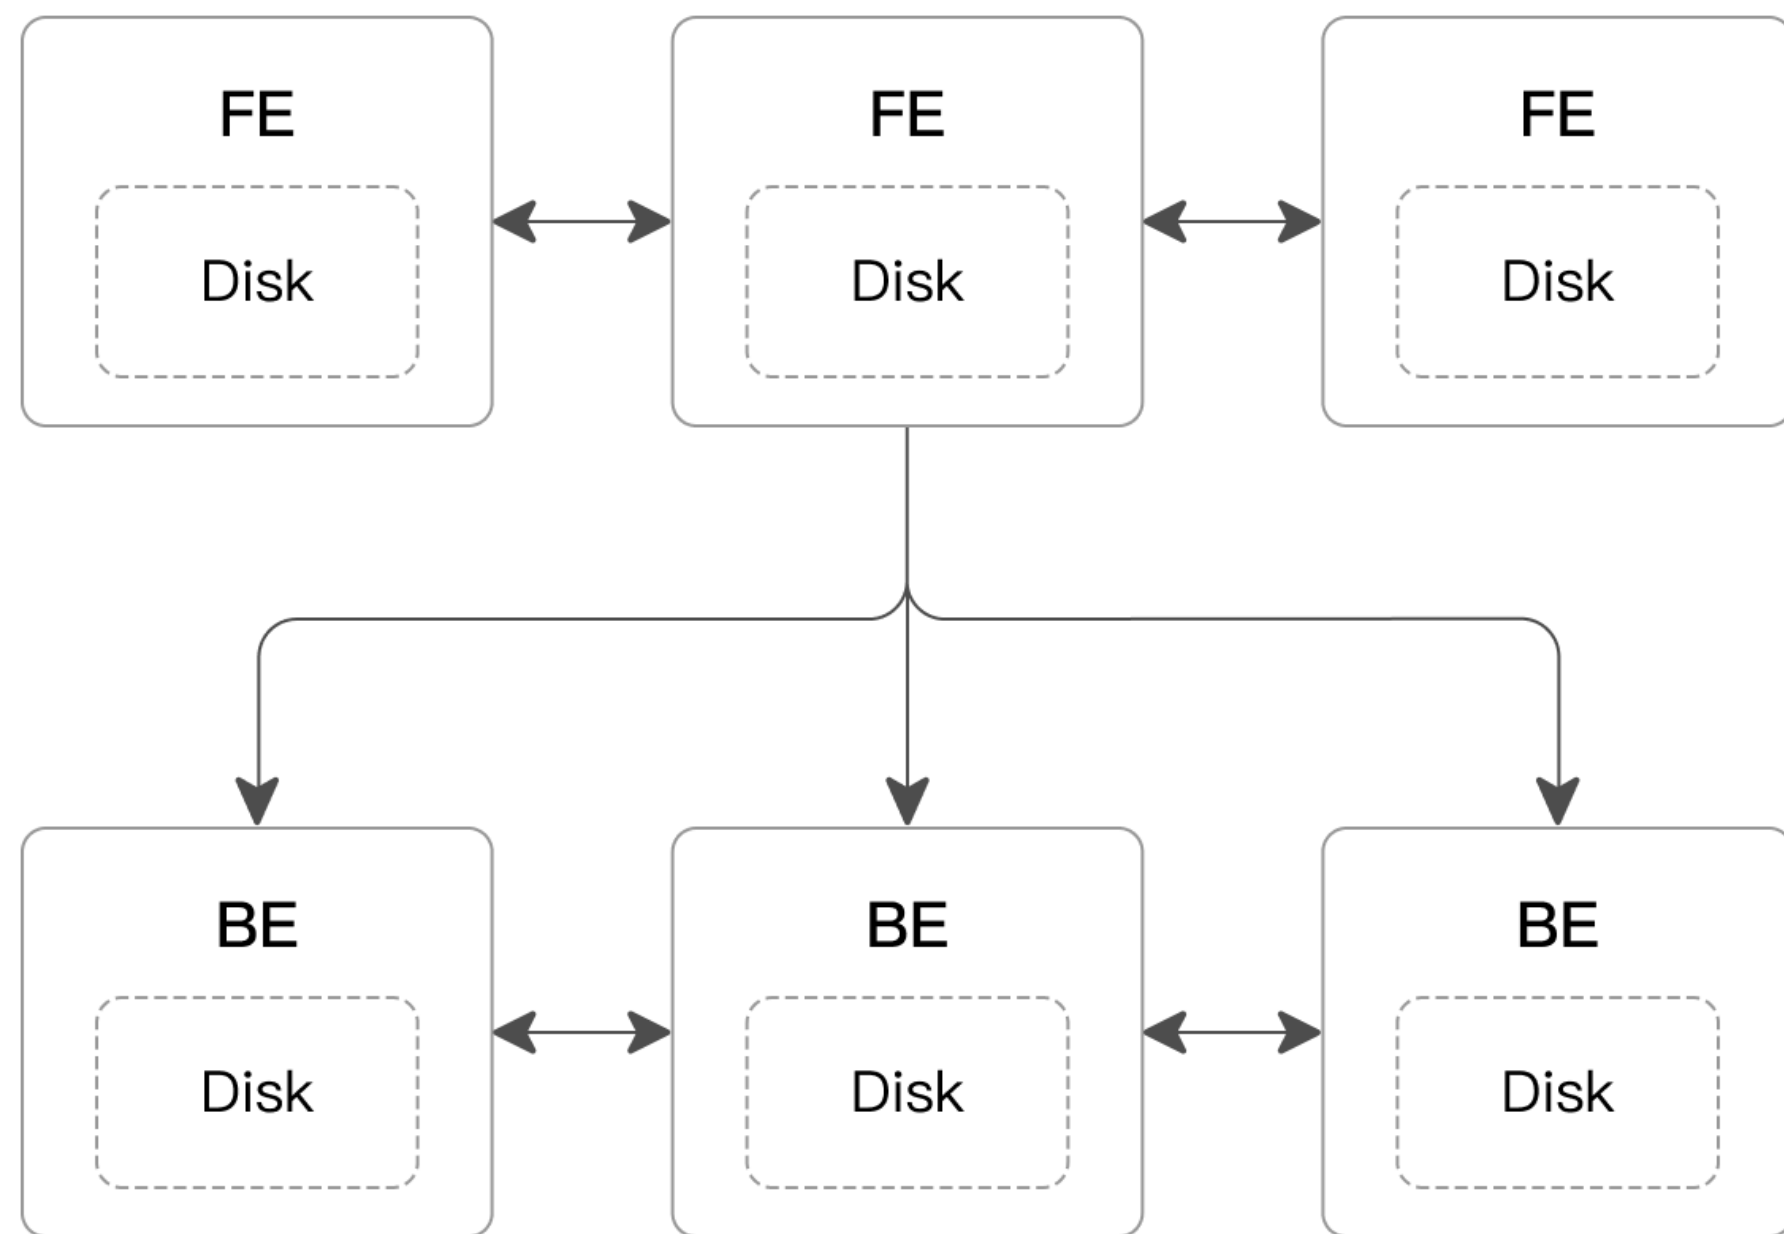
A Modern Data Warehouse

Offering Lightning-Fast Analysis on Large-Scale, Real-Time Data




Architecture

Compute-Storage Coupled




Simplicity

Compute-Storage Decoupled



Elasticity

Core Features of Apache Doris



Lightning Fast

- One of the world's fastest SQL query engines



Easy to Use

- Friendly for first-time user
- Low operational costs as a distributed system
- Flexible deployment options for various environments



Multi-Scenario

- Reporting & ad-hoc
- Semi-structured data analysis
- Lakehouse



Lightning Fast SQL Query Engine

ClickBench

- benchmark.clickhouse.com

System & Machine	Relative time (lower is better)
Umbra (c6a.metal, 500gb gp2):	x 1.61
ClickHouse (tuned, memory) (c6a.metal, 500gb gp2):	x 1.95
ClickHouse (tuned) (c6a.metal, 500gb gp2):	x 2.04
Apache Doris (c6a.metal, 500gb gp2):	x 2.15
ClickHouse (c6a.metal, 500gb gp2):	x 2.21
StarRocks (c6a.metal, 500gb gp2):	x 2.38
Umbra (c6a.4xlarge 500gb gp2):	x 2.40

TPC-H & TPC-DS



Behind the Lightning Fast SQL Query Engine

Cost-Based Optimizer

- Cost-based join reorder, runtime filter
- Short circuit plan for high-concurrency queries

Pipeline Execution

- Data-driven, no blocking of threads, fine-grained concurrency
- Self-adjusted parallelism level

Materialized Views

- Consistent single-table materialized views, support general aggregation functions
- Multi-table materialized views

Full Vectorization

- Reduce virtual function calls and cache miss
- Efficient use of SIMD instructions, supports X86 and ARM

Indexes

- BloomFilter, Min / Max / Sum
- Prefix Sorted Index
- Inverted Index

Smart Caching

- Caching of query results, data, metadata, and intermediate data
- Caching of internal and external tables

Massively Parallel Processing Architecture

- Parallelism within and between nodes to give full play to machines and cores
- Supports distributed join of large tables and operator materialization

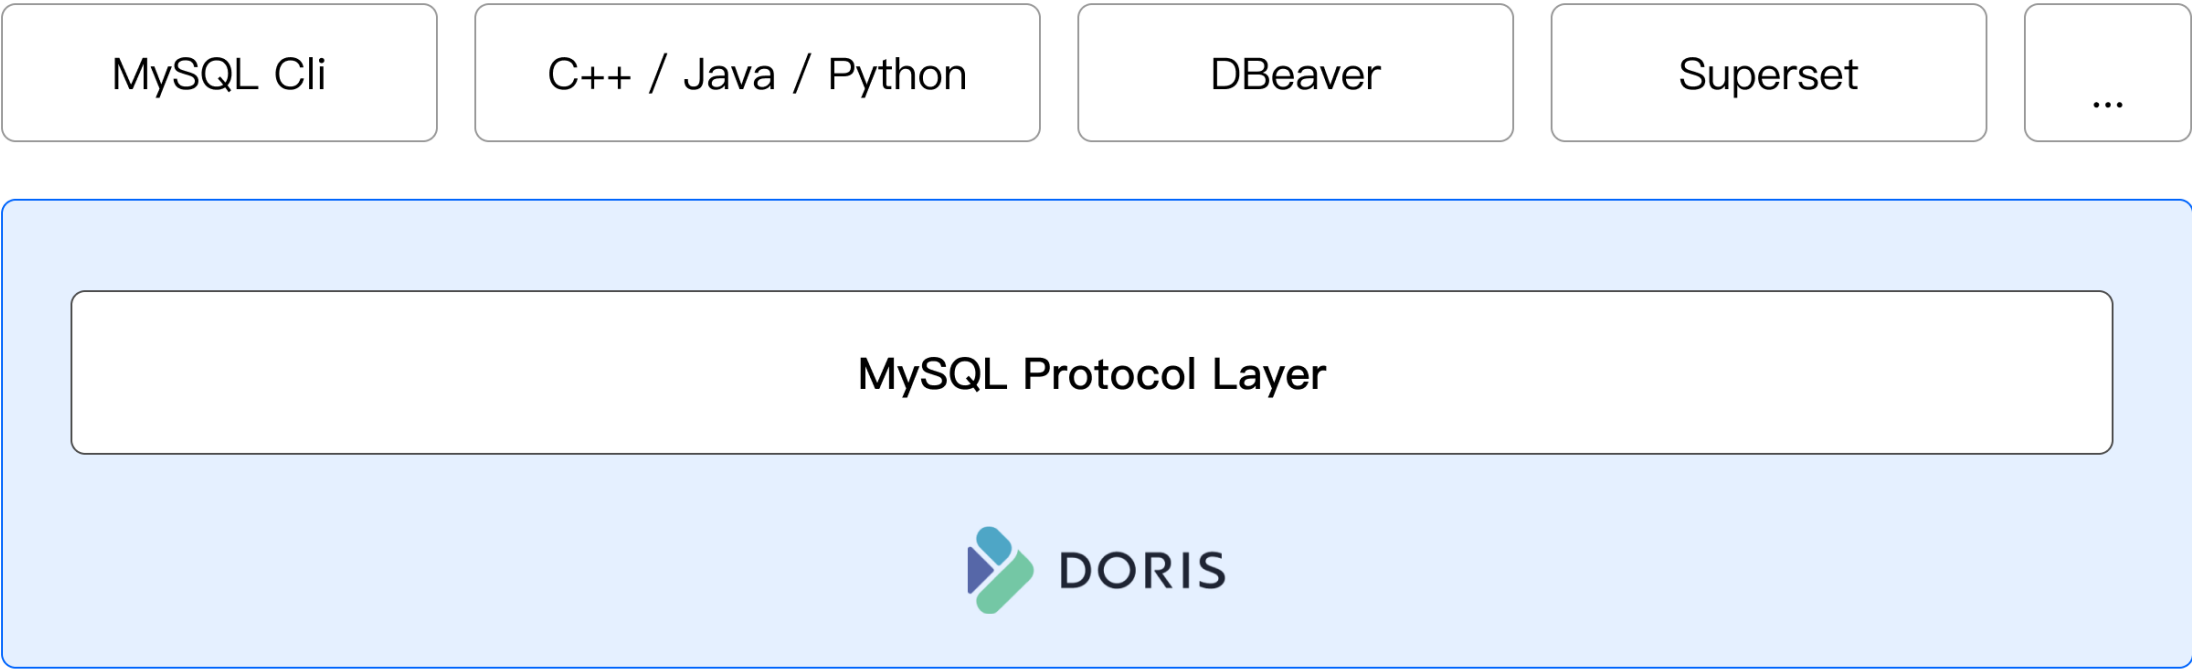
Columnar Storage & Hybrid Storage

- Columnar storage for efficient encoding, compression, and data sharding
- Row and columnar hybrid storage for flat tables to reduce IOPS amplification

Easy to Use


MySQL Protocol & ANSI SQL

```
CREATE TABLE doris  
(  
  col1 int,  
  col2 string  
) DISTRIBUTED BY RANDOM BUCKETS 10;  
  
SELECT * FROM doris WHERE col1 like "%kkey%";
```



Deployed Everywhere

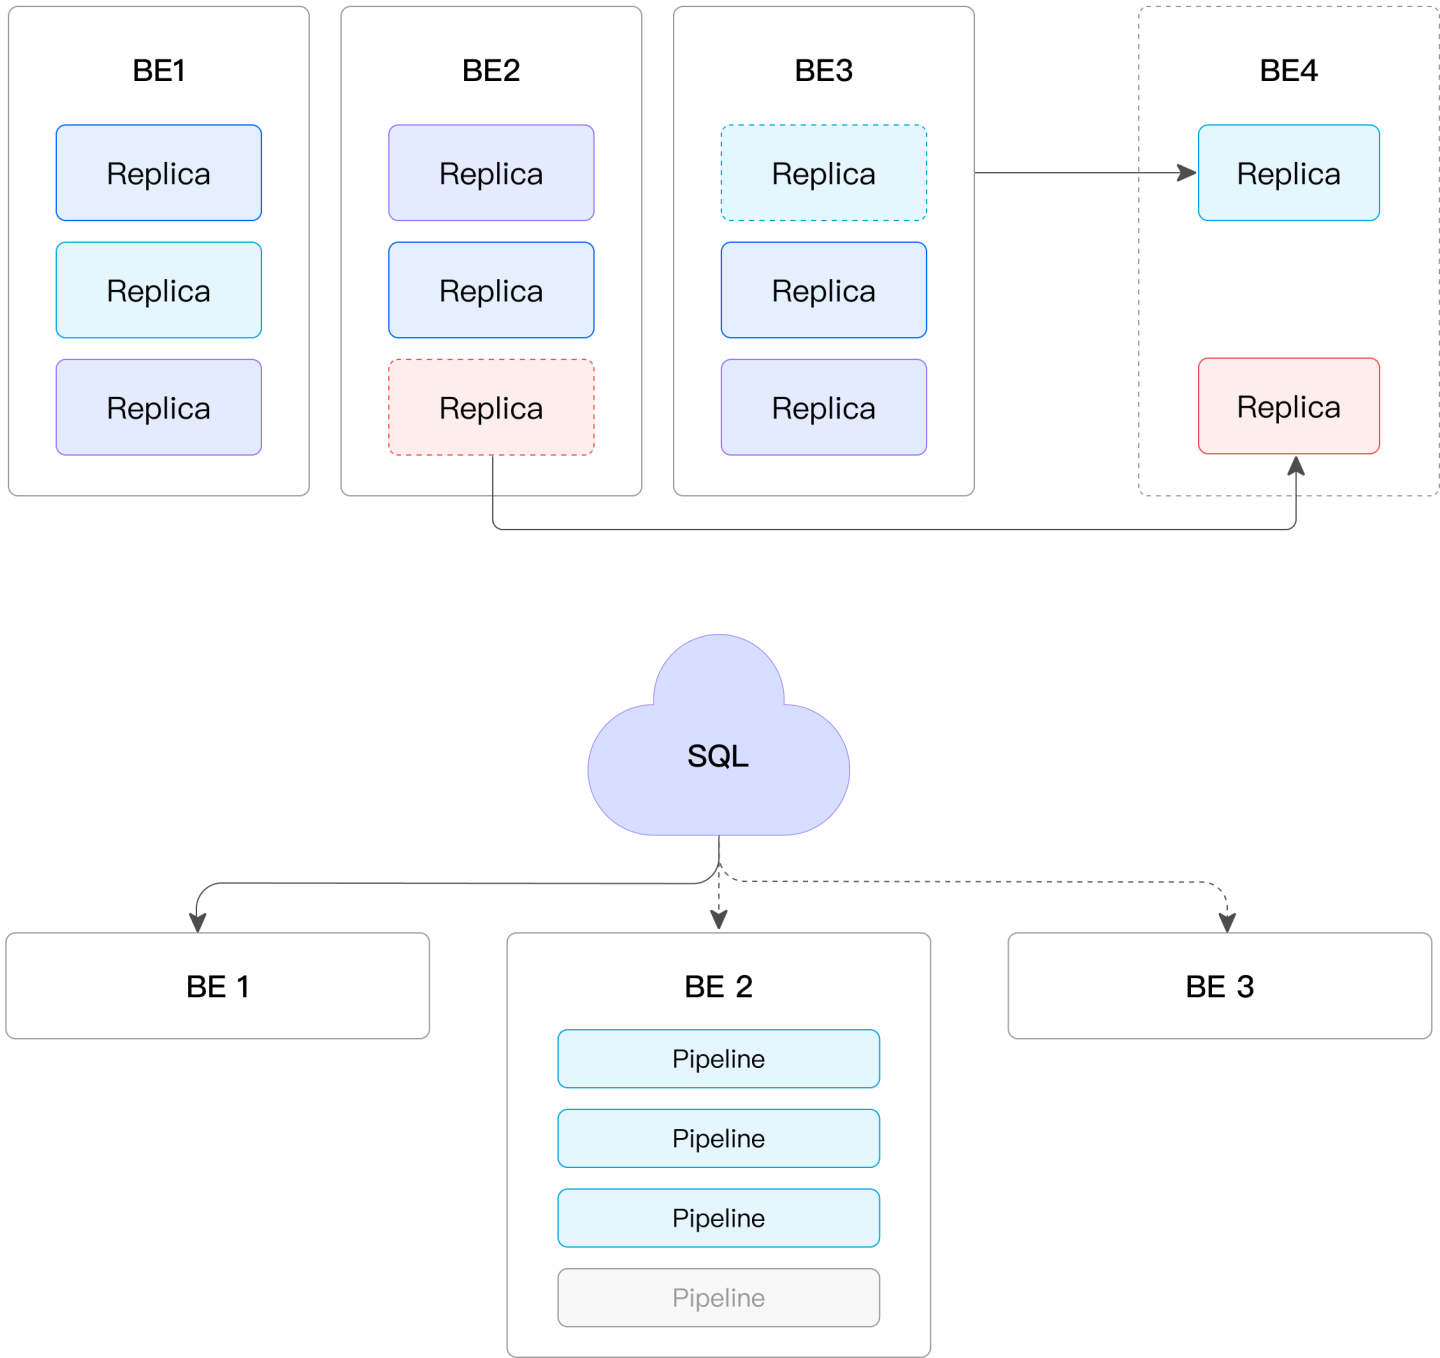
- Bare metal
- EC2
- K8s
- BYOC / SaaS



Easy to Use

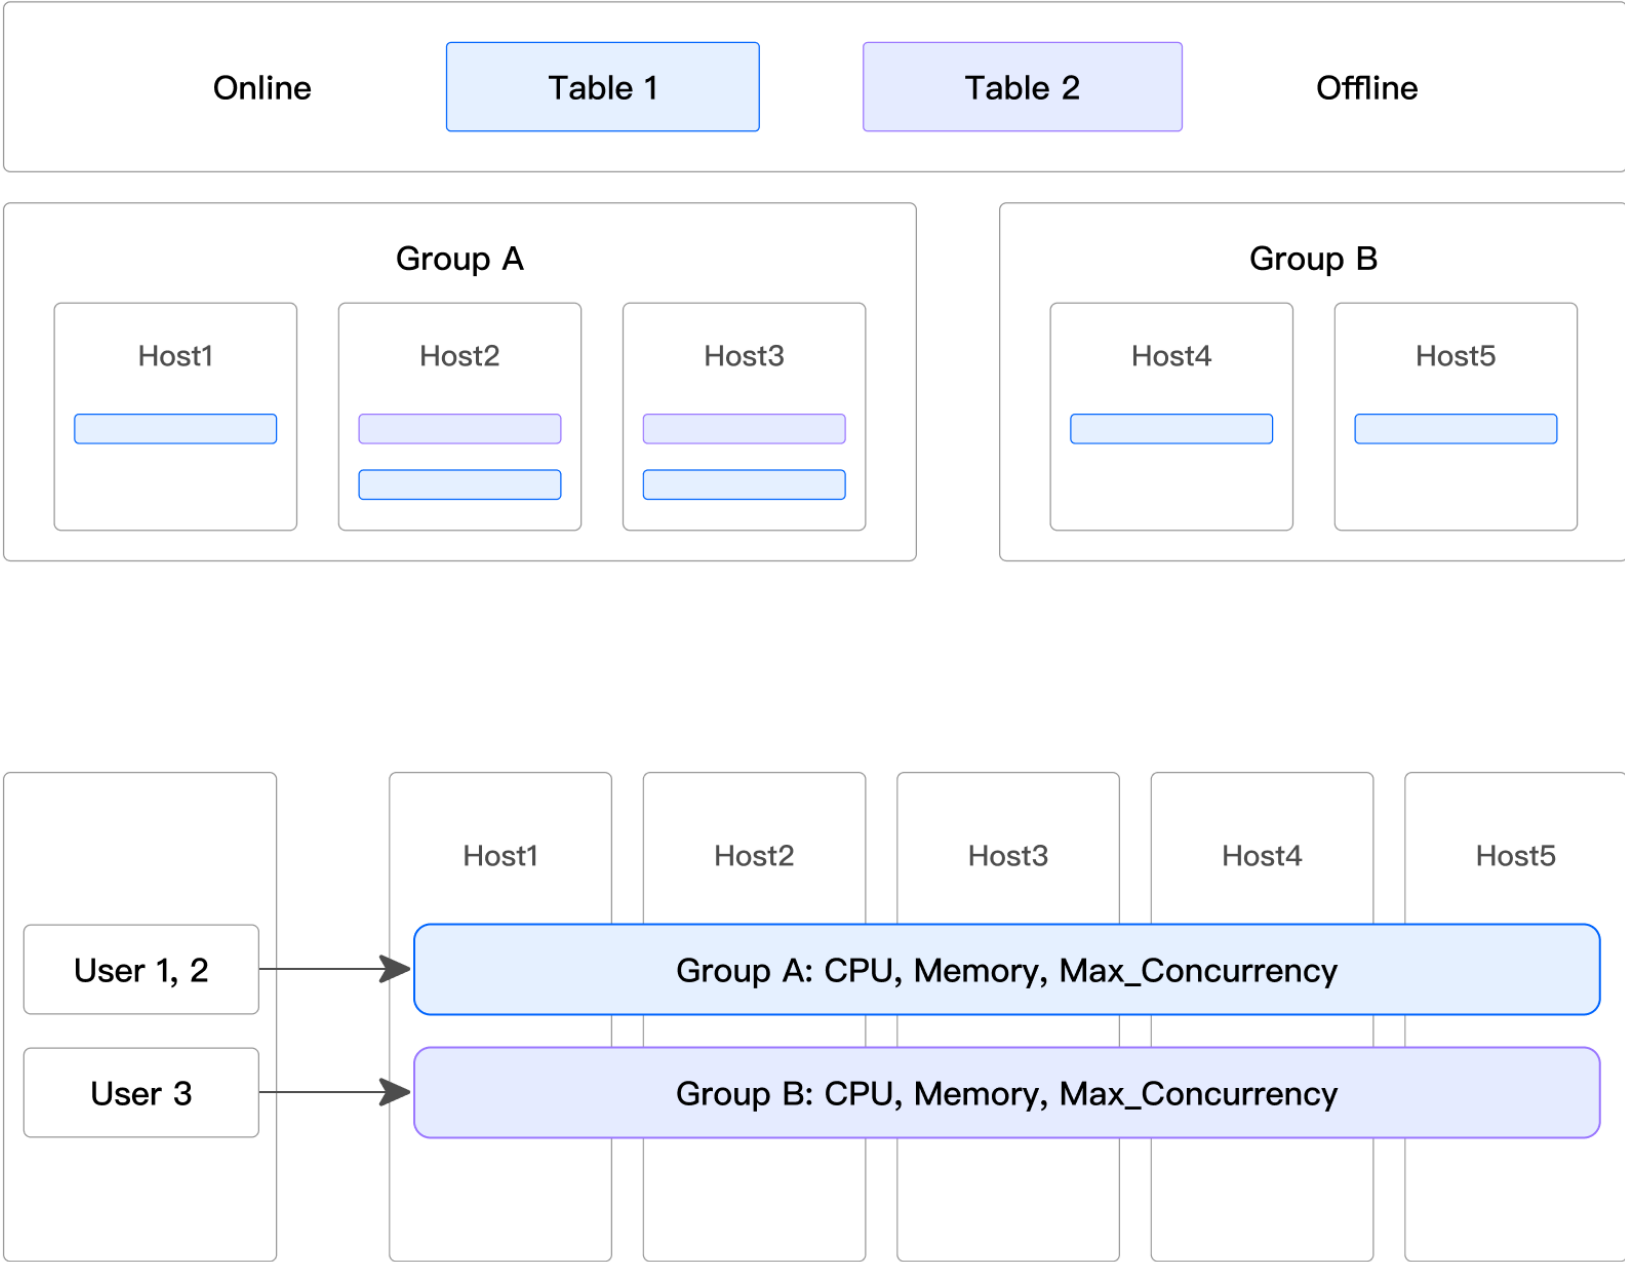
Easy Operation and Maintenance

- Auto Balance
- Adaptive Concurrency
- Auto Replica Repair
- Fault Tolerant



Multi-Tenancy and Resource Isolation

- Resource Group
- Virtual Cluster
- Workload Group
- Query Fuse



Multi-Scenario

Reporting

- Pre-aggregation data model (Rollup)
- Query Cache

```
SELECT Department, SUM (Salary)  
FROM EMPLOYEE  
GROUP BY Department
```

Department	SUM (Salary)
RD	38000
QA	19000

Name	Department	Salary
John	RD	20000
Bob	RD	18000
Alice	QA	19000

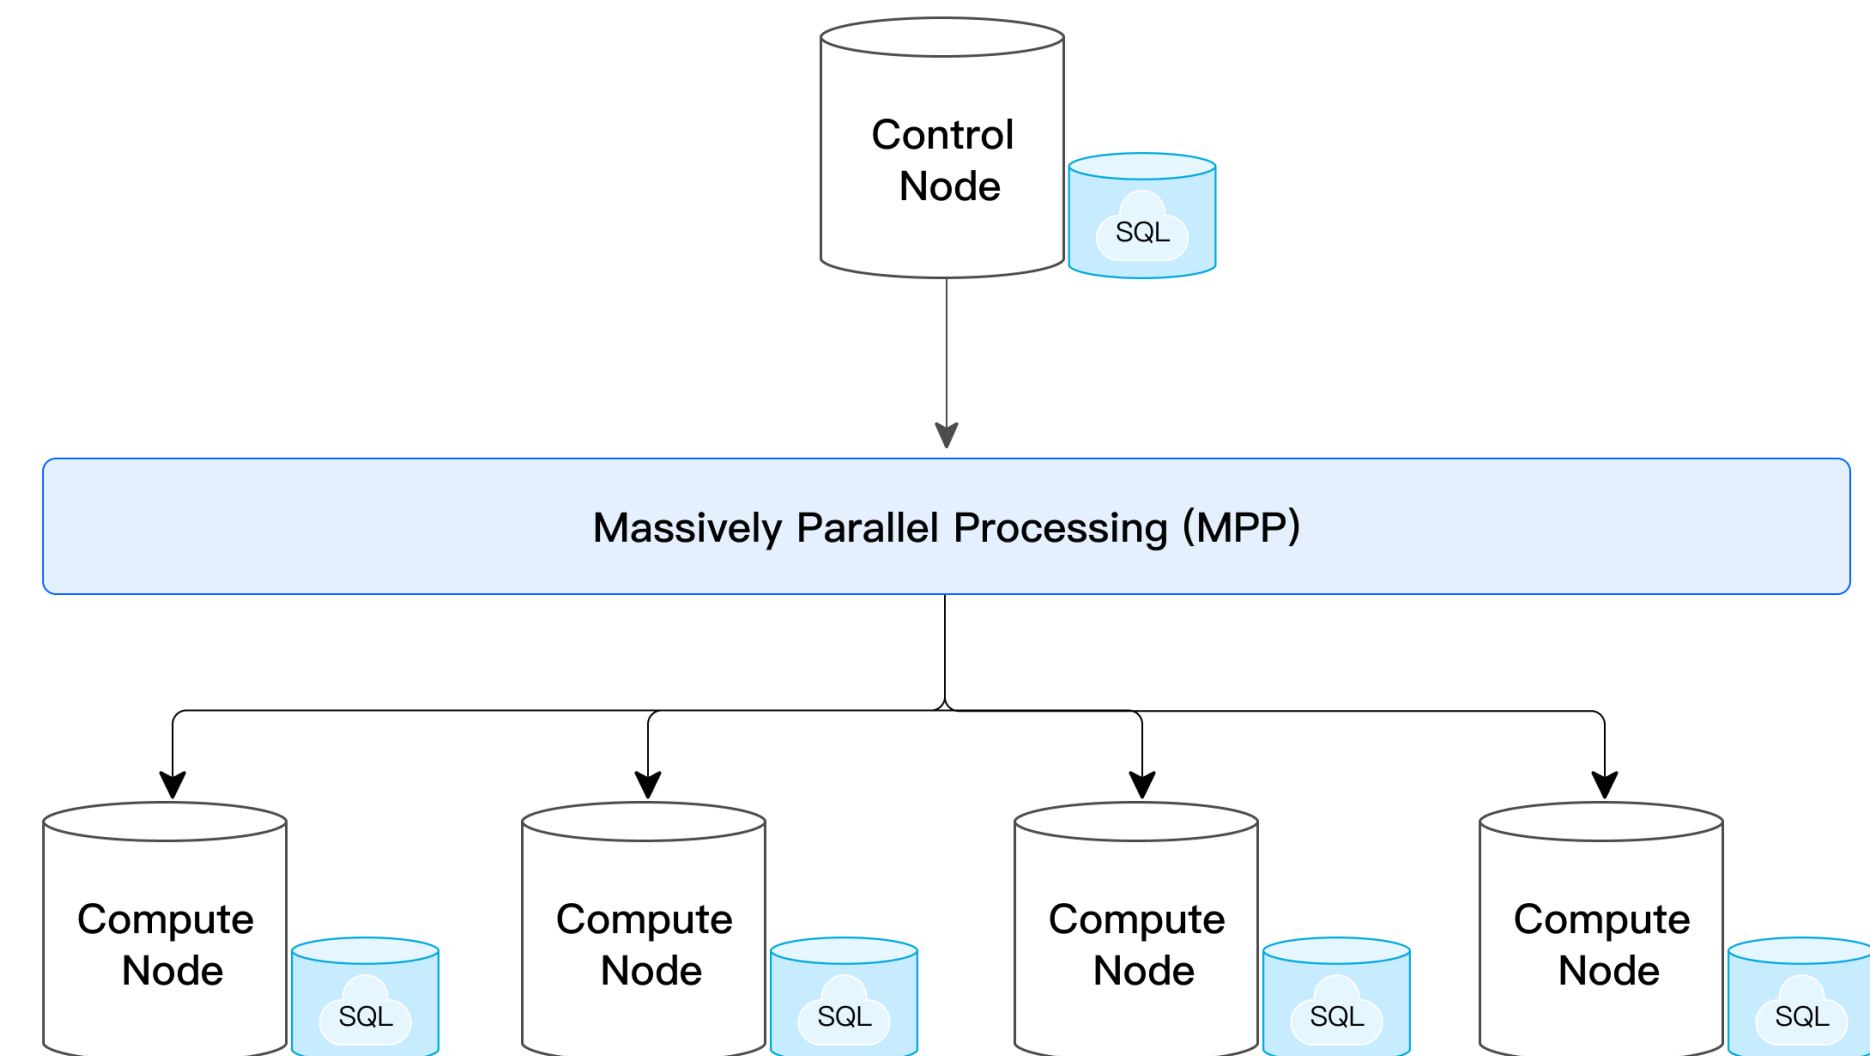
Result Cache

Partition Cache

Page Cache

Ad-Hoc Query

- Massively parallel processing
- Adaptive pipeline execution engine
- Spill to disk

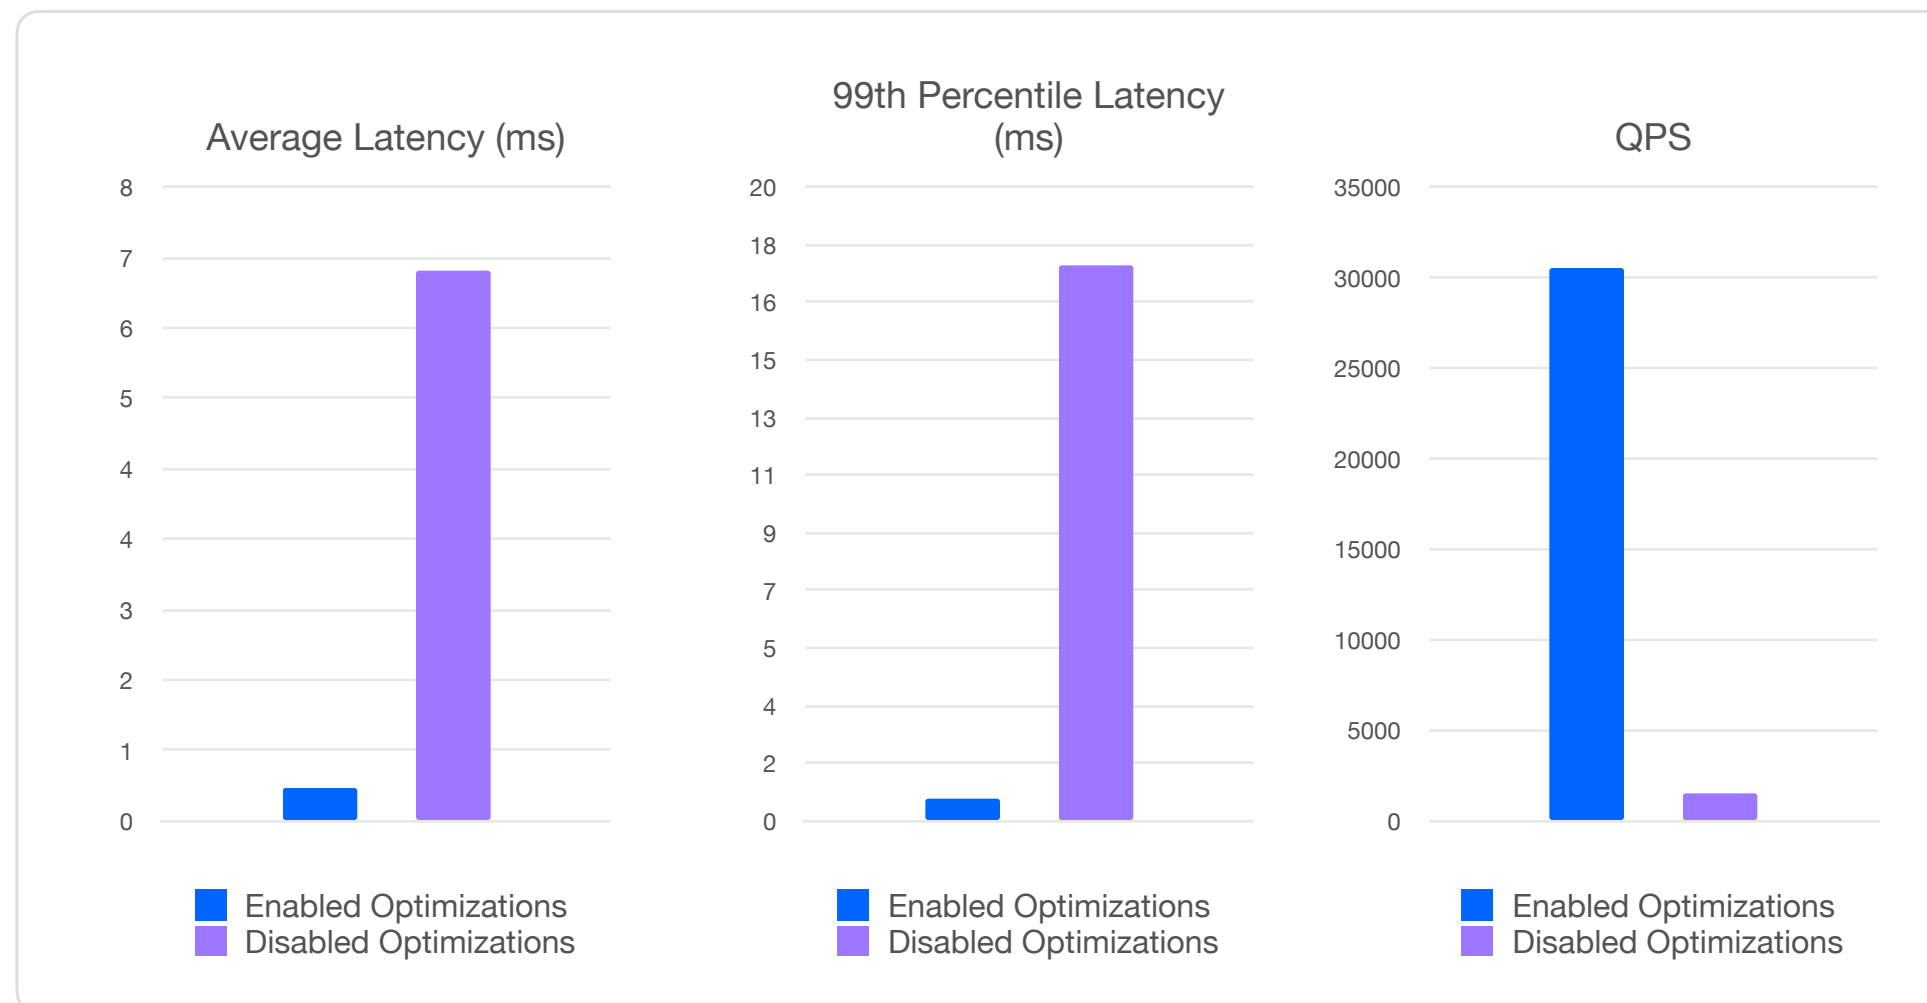


Multi-Scenario

High Concurrency Point Query

Small amount of data retrieved from a massive dataset

```
SELECT * FROM billing WHERE user_id=123
```

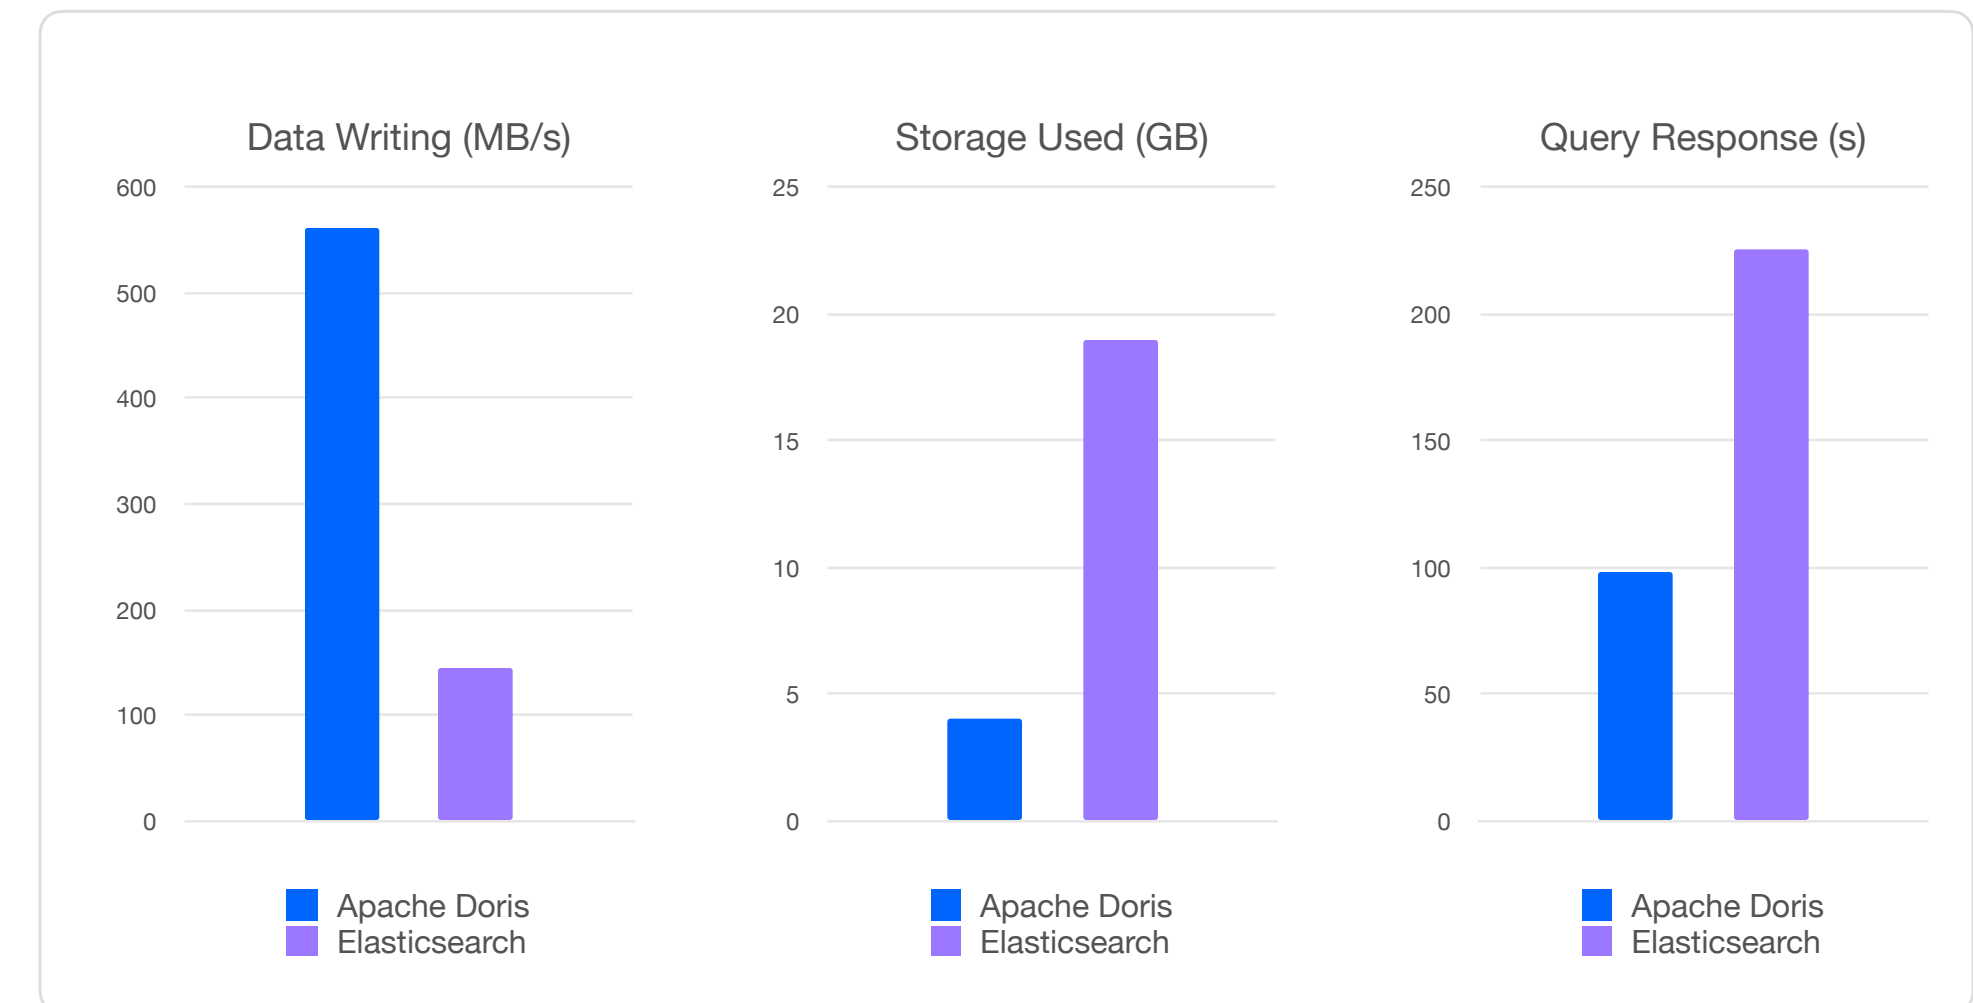


- Row storage
- Prepared statement
- Short circuit query plan

Semi-Structured Data Analysis

Log Management

Compared to Elasticsearch



- Inverted index
- Full-text search
- JSON / VARIANT data type


Contents

01 What is Apache Doris

02 Building Lakehouse on Doris

03 Apache Doris Community

Lakehouse Challenges



Performance

- How to speed up the query on lake data?



Diversity

- Semi-structured data support
- Insertion, deletion and update



Openness

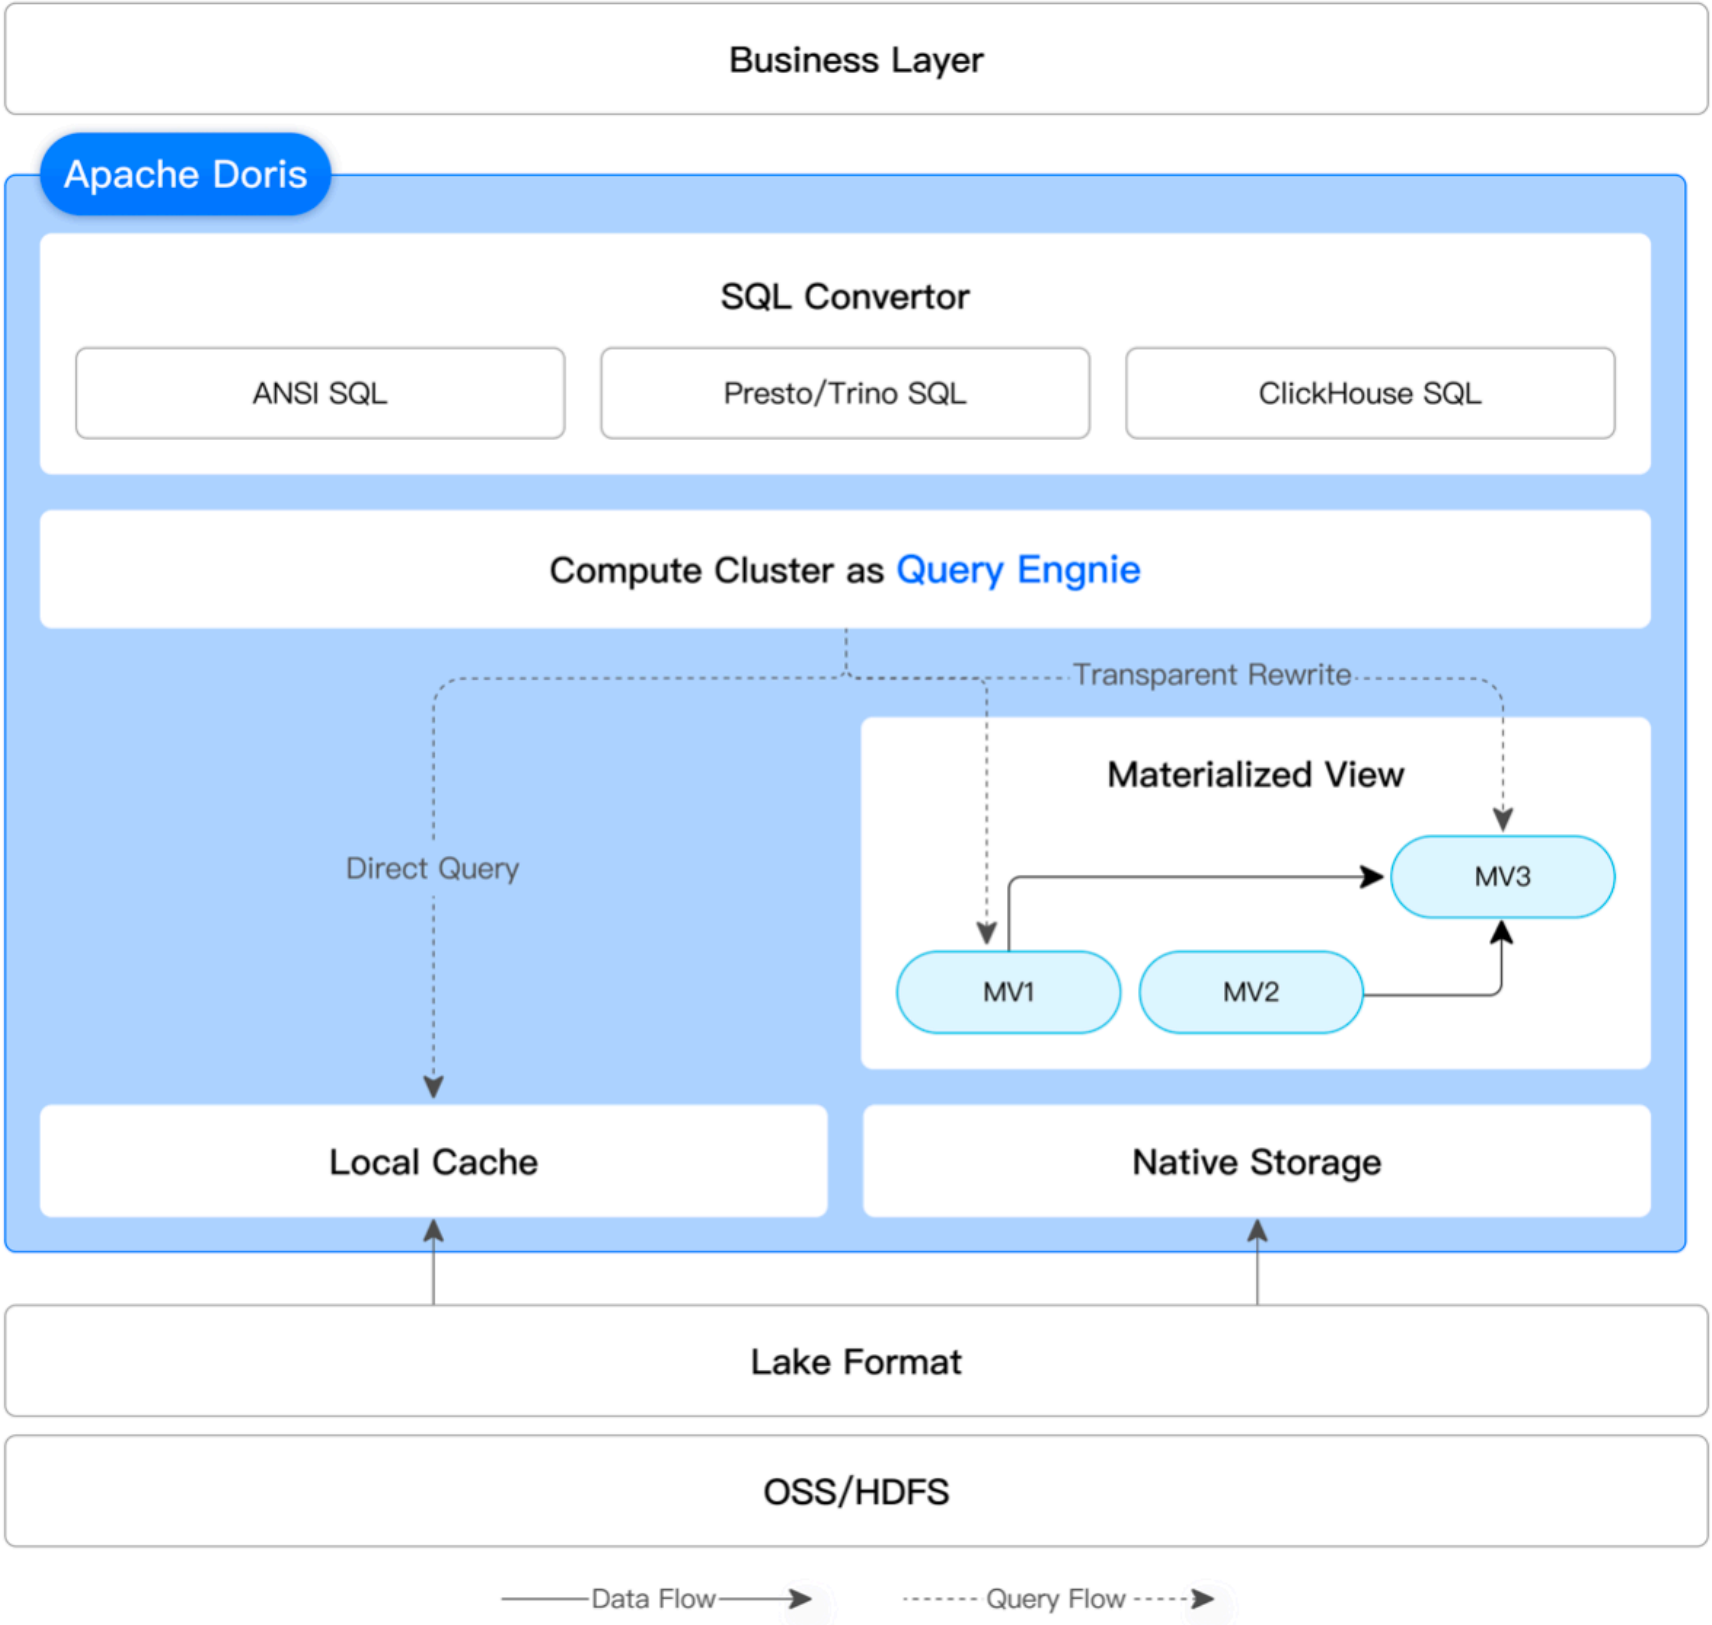
- No vendor lock-in
- Support various engine

Apache Doris Lakehouse Solution

Scenario 1: Query Engine

Query Engine

- Hive, Iceberg, Hudi
- Materialized view
- File Caching
- Query rewriting



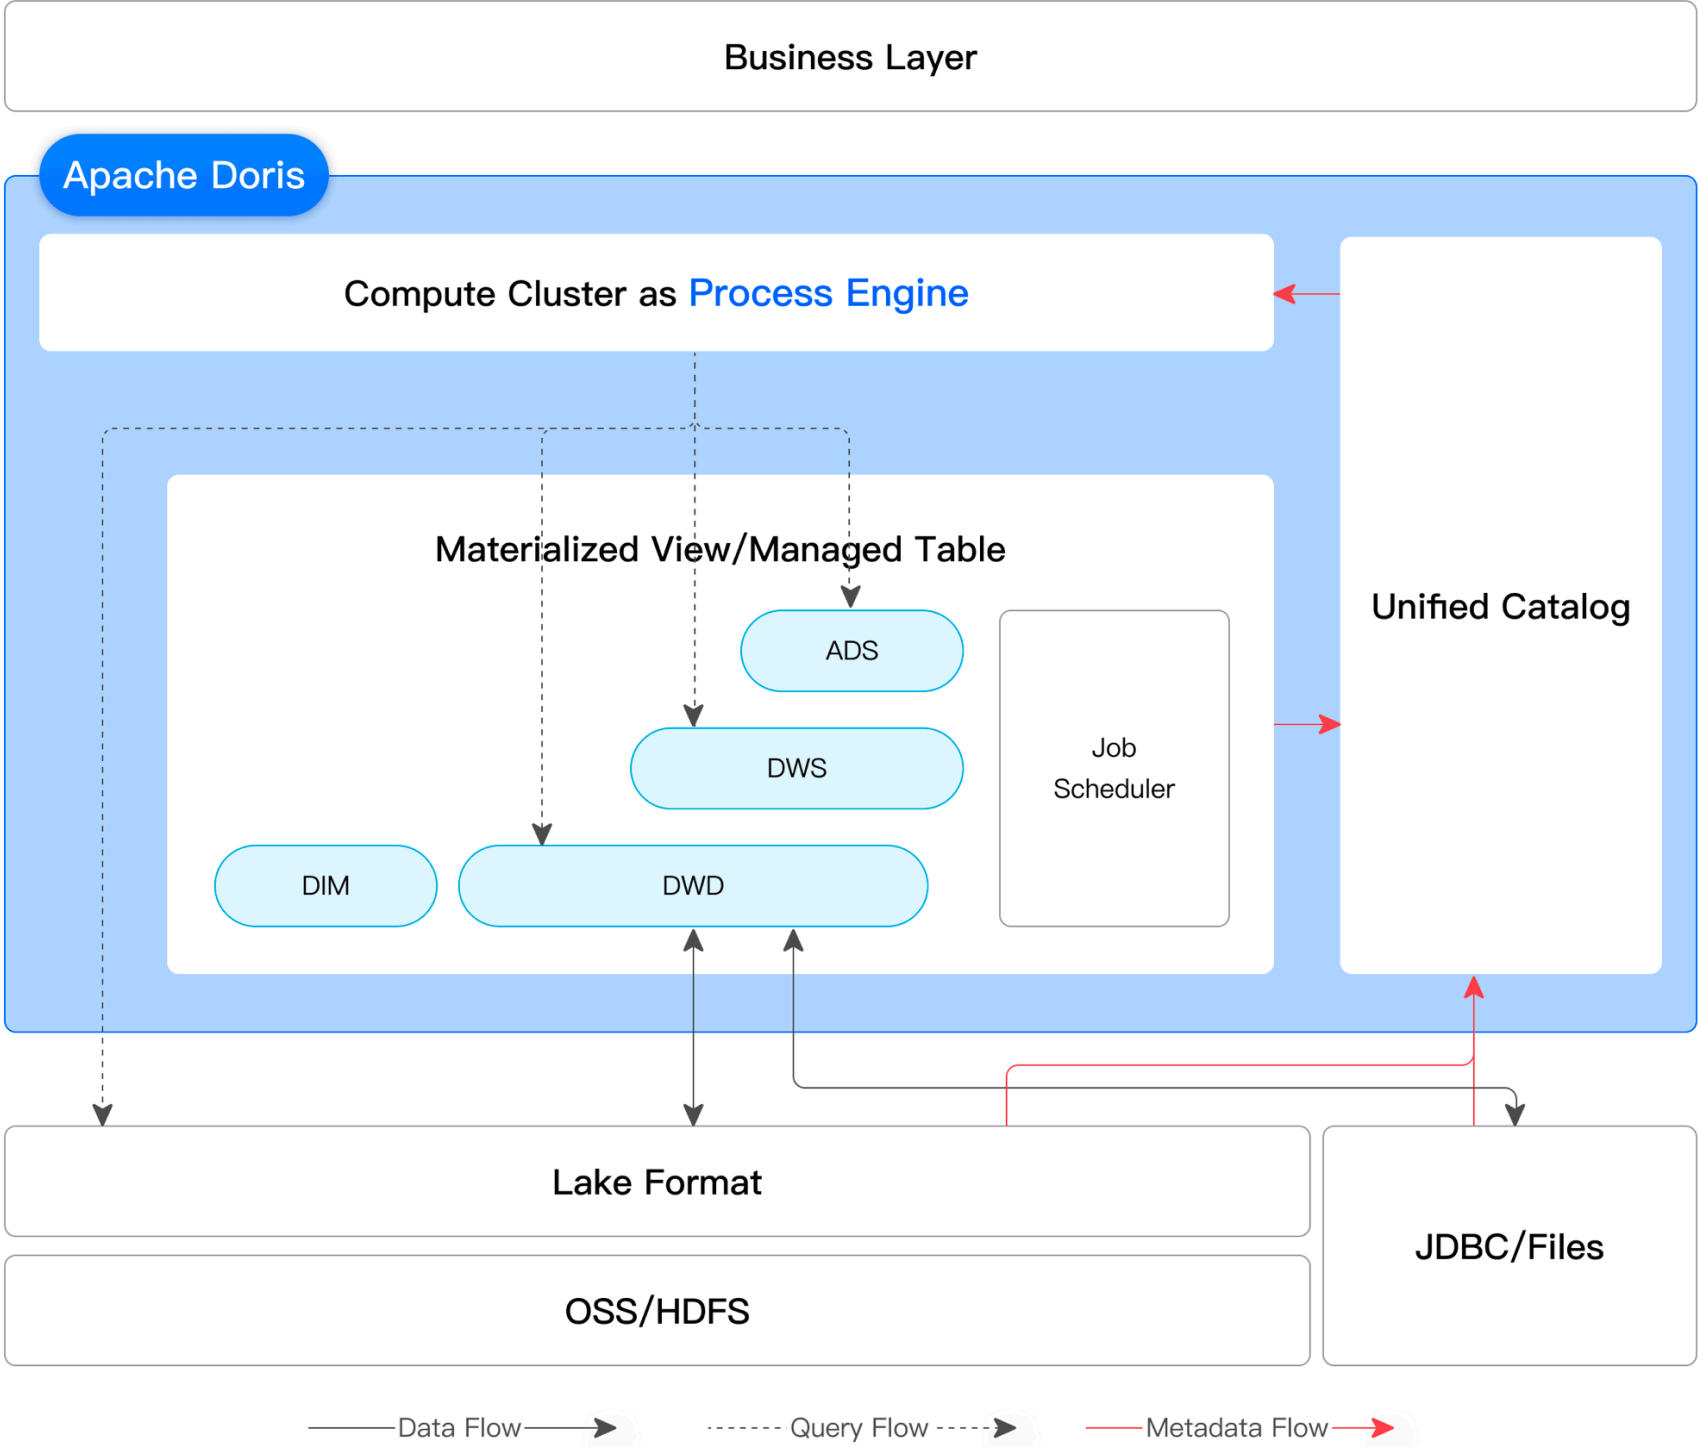
Apache Doris Lakehouse Solution

Scenario 2: Process Engine

Data Process Engine

- Write data to Hive/Iceberg
- Job scheduler
- Spill to disk

```
CREATE JOB my_job  
ON SCHEDULE EVERY 1 DAY STARTS '2024-11-18 00:00:00' DO  
INSERT INTO hive.db1.table1 SELECT * FROM doris.db.table2  
WHERE create_time >= days_add(now(),-1);
```

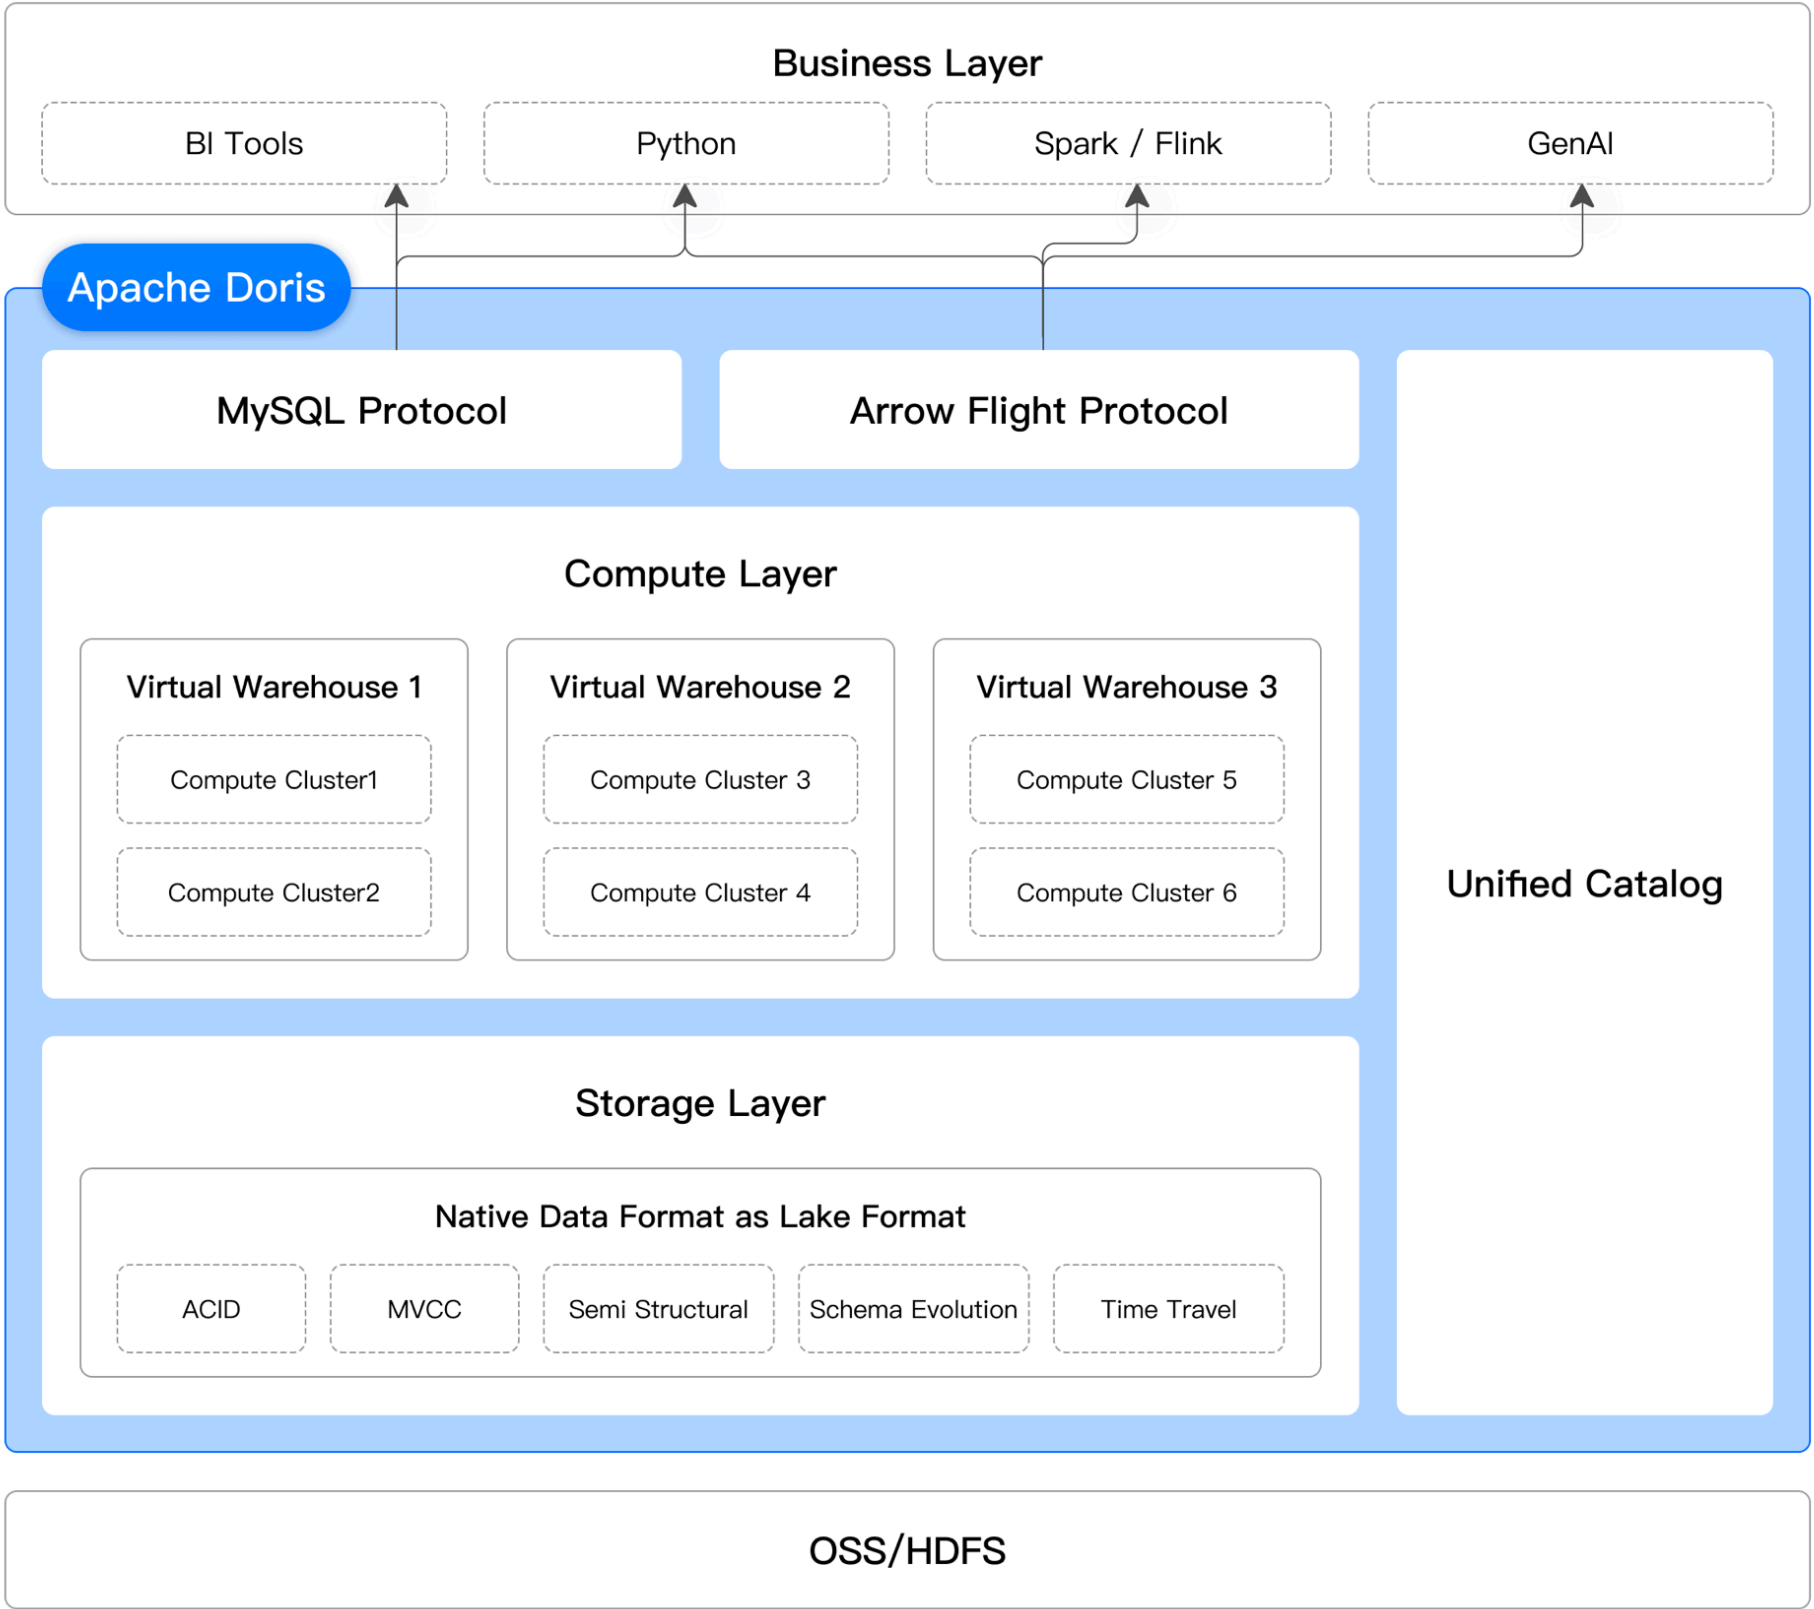


Apache Doris Lakehouse Solution

Scenario 3: Lakehouse Engine

Open Lake Format



- MVCC
- Data insert/delete/update
- Open Storage API
- Unified Catalog



User Case: Building Lakehouse Engine on Doris

Kwai: a leading short-video app provider

Lakehouse Query Engine & Auto Materialized Data Management



Contents

01 What is Apache Doris

02 Building Lakehouse on Doris

03 Apache Doris Community

One of the world's most active open source communities in big data

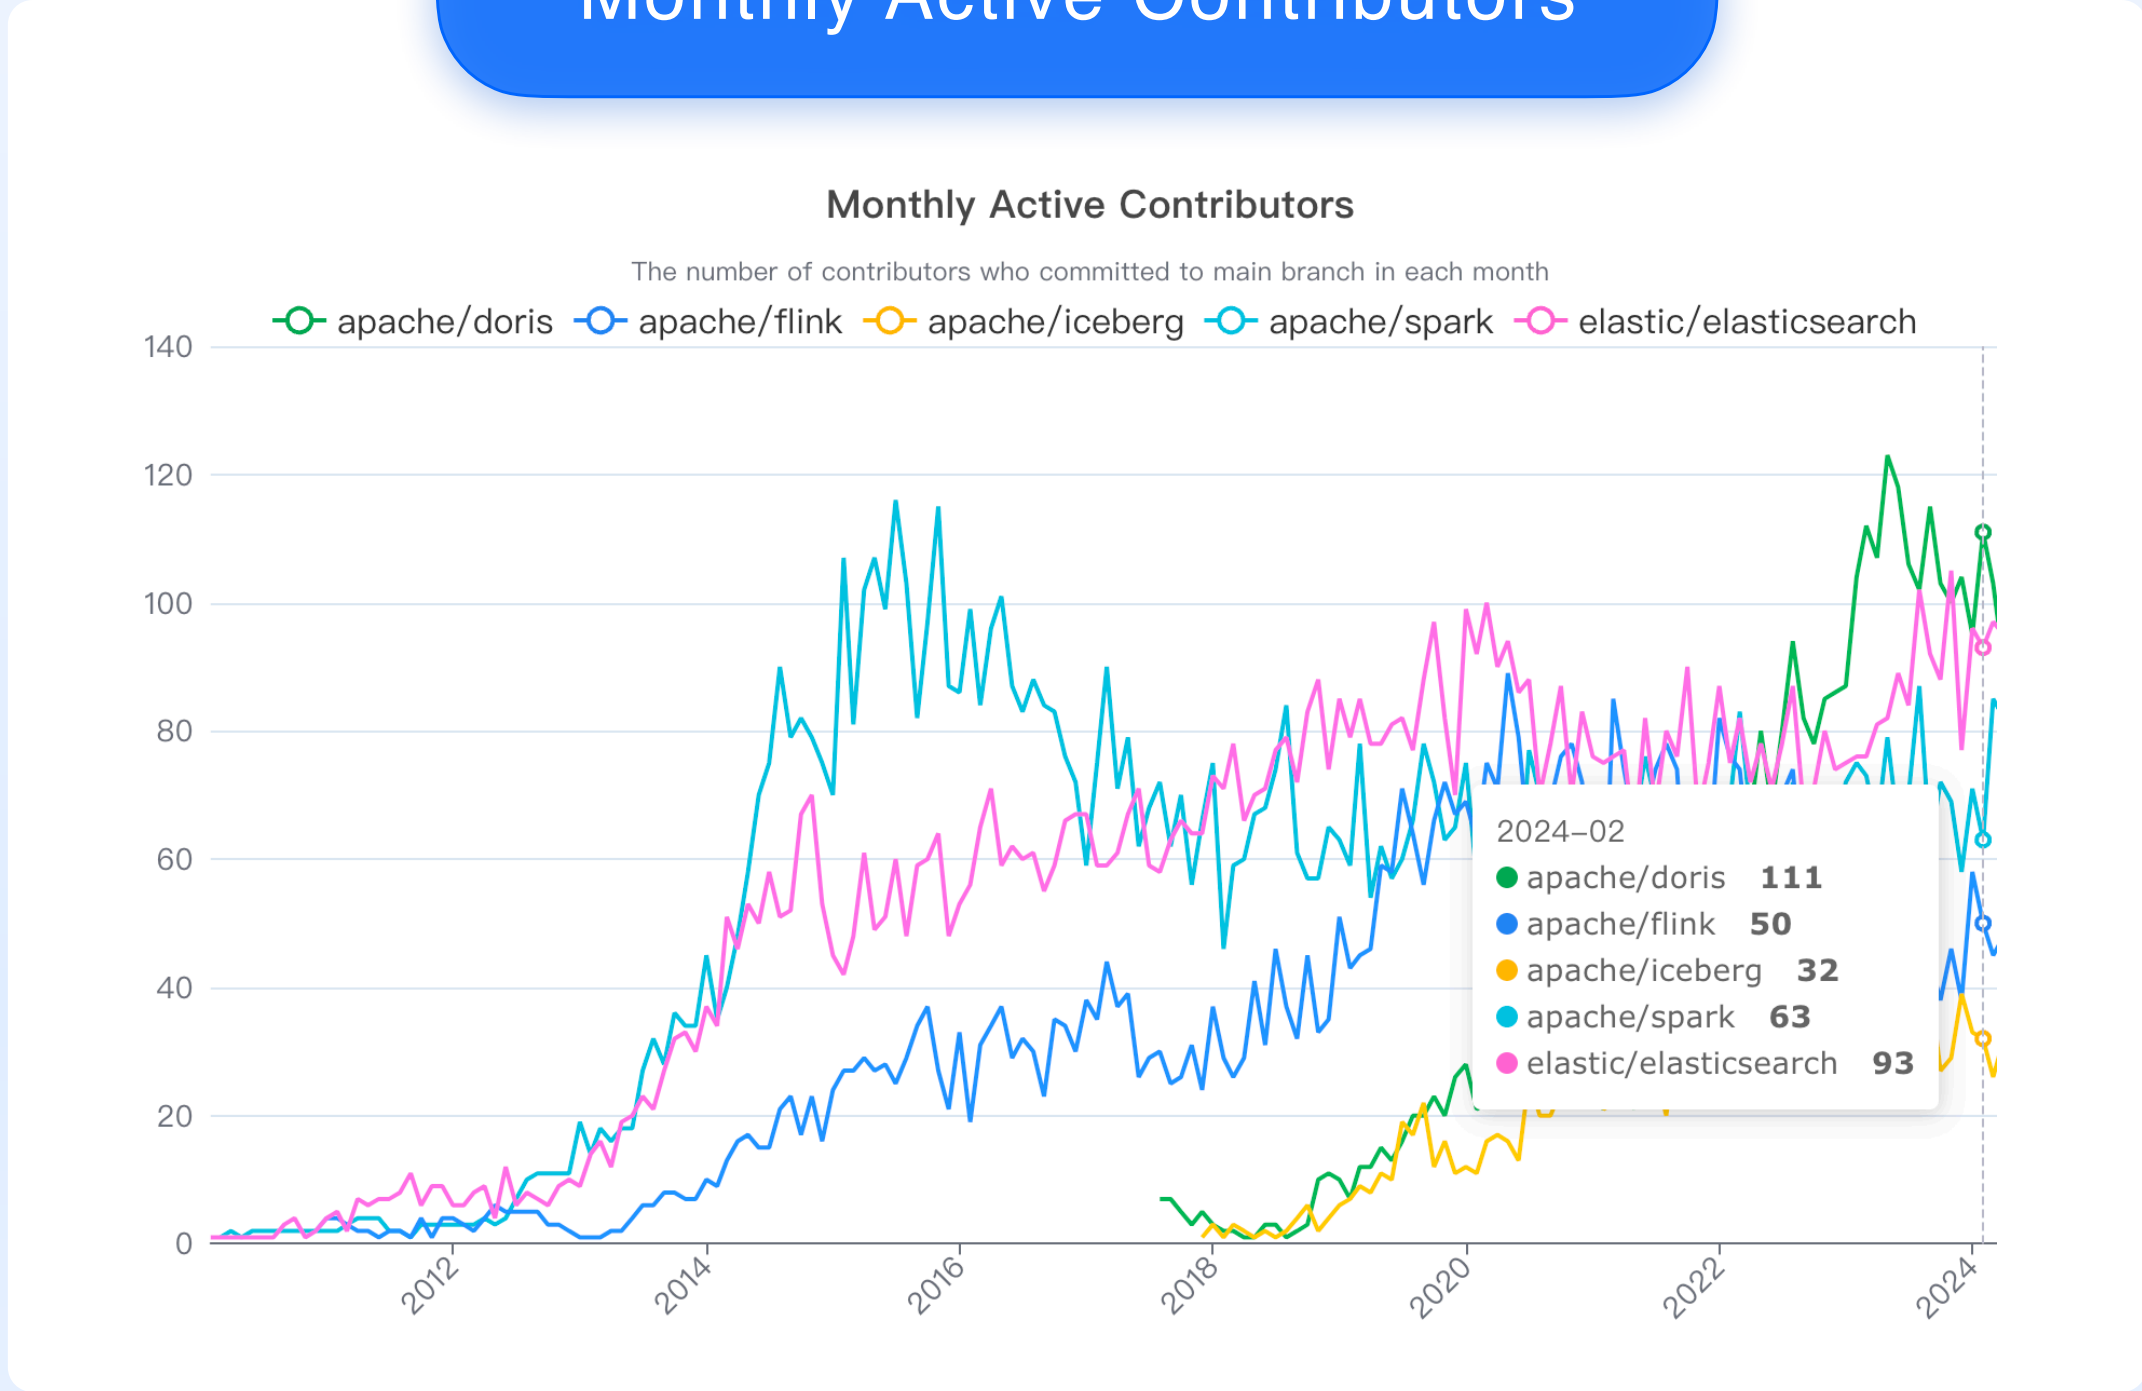
The data is current as of March 2024

Total Contributors



650+ Total Contributors

Monthly Active Contributors



100+ 100+ monthly active contributors

Trusted by over 5000 enterprises worldwide for online analytics

Apache Doris is used worldwide in industries like Retail, Finance, Internet, Gaming, Telecommunications, etc.




It's never too late to join the Apache Doris Community

- Subscribe to our mailing list and join our discussion: dev@doris.apache.org
- Get technical support on Slack apachedoriscommunity.slack.com
- Give us a star on GitHub: [apache/doris](https://github.com/apache/doris)
- Follow us on LinkedIn, Twitter and YouTube [@ApacheDoris](#) [@VeloDB](#)



WhatsApp



Telegram

Thanks !

