

Deep Dive into the Apache Doris Compute-Storage Decoupled Mode

Rayner Chen

Apache Doris PMC Chair

Tech VP at VeloDB

March 27, 2025


Contents

01 Capabilities of the Decoupled Mode


02 Applicable Scenarios of the Decoupled Mode

Apache Doris Compute-Storage Decoupled Mode


Coupled Mode



Decoupled Mode – Current




Decoupled Mode – Future



Query Performance of the Decoupled Mode

Multi-Layered Caching



Doris Page Cache: decompressed data




Linux Page Cache: compressed data



Local Disk Cache: compressed data




Persistent storage



- **Full Cache Hit:** All relevant data resides in Doris Page Cache or Linux Page Cache after warmup.
- **Partial Cache Hit:** Clear all caches, execute the TPC-DS queries sequentially, and record the performance metrics during the initial run.
- **No Cache Hit:** Clear all caches before executing each TPC-DS query.


Data Freshness of the Decoupled Mode



32 concurrent writes from Flink with a checkpoint interval of 5 seconds

- A: after introducing Meta Service
- B: with meta data written into object storage

Count of Small Files Generated



Write to a 32-bucket table

- A: after introducing Meta Service
- B: with meta data written into object storage

Capabilities of the Decoupled Mode



Low Cost

Elasticity

**Workload
Isolation**


Low Cost: 90% Cost Reduction

Single-Replica Storage on Affordable and Reliable Infra



SSD gp3 \$0.08/GB-Month * 3

HDD st1 \$0.045/GB-month * 3



Amazon S3

\$0.023 /GB-Month

cache with gp3: \$0.08/GB-Month

Resource Consumption from Data Compaction Reduced by 66%

Compaction




Compaction



Compaction




Compaction



Amazon S3


Freeing up Compute Resources by Offloading Cold Data



Amazon S3


Elasticity: Stateless Compute Nodes

Independent Scaling of Compute & Storage




1. Pooling compute and storage resources to improve utilization
2. Shared storage for simpler operation and maintenance

Dynamic Resource Scheduling



Dynamically scaling compute resources for peak demand & cloud cost optimization

Workload Isolation by Compute Groups



Current

- Isolated Cache
- Isolated BE
- **Shared FE**


Future

- Isolated Cache
- Isolated BE
- **Isolated FE**

02

Applicable Scenarios


Typical Use Case: Cost Reduction for Online Business



Cost calculation is based on AWS pricing.


Typical Use Case: Cost Reduction for **Historical Data**

200 TB of Historical Data



Cost calculation is based on AWS pricing.

Typical Use Case: Elastic Scaling to **Handle Traffic Spikes During Big Sales**




Customer-Facing Analytics
(Queries from consumers)

Customer Segmentation
(Queries from businesses)

Logistics Services


Typical Use Case: **Dynamic Compute Resource Allocation**

Phase One: Warm up G2




Phase Two: Switch to G2

2. Set Default Compute Group to G2




Typical Use Case:

Multi-Service Data Sharing with Workload Isolation



Compute-Storage Coupled Resource Group


- Each business line requires a data replica.
- Each business line requires the same number of machines, as storage capacity is tied to computing nodes.




Compute-Storage Decoupled Resource Group

- A single data replica in object storage can serve multiple business lines.
- Computing resources can be allocated flexibly based on actual needs.

Typical Use Case: Isolation of Online and Offline Workloads



Typical Use Case: Read-Write Isolation



Typical Use Case: Next-Gen Data Foundation



Key Capabilities

- High-concurrency Queries
- Real-Time Data Warehouse
- Log Storage & Analysis
- Compute-Storage Decoupled
- Workload Isolation
- Data Lakehouse

Welcome to the Apache Doris Community

Get technical support

- Slack: apachedoriscommunity.slack.com



Give us a star on GitHub

- Repo: <https://github.com/apache/doris>



Subscribe

- Mailing list: dev@doris.apache.org



Follow us on:

- LinkedIn: <https://www.linkedin.com/company/doris-apache/>
- Youtube: https://www.youtube.com/@VeloDB_IO
- X: https://x.com/doris_apache



Thanks for Watching