Tencent Cloud TCHouse-D

Product Introduction

Product Documentation

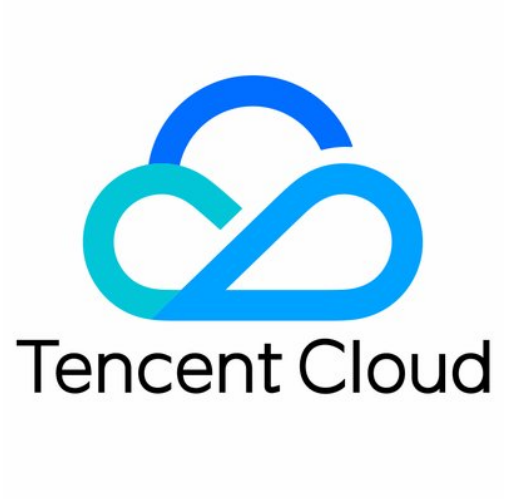

Contents

Product Introduction
Overview

Last updated：2024-06-27 10:36:25

Tencent Cloud TCHouse-D is built based on the Apache Doris kernel used by leading Online Analytical Processing (OLAP) databases and is compatible with the MySQL protocol. It integrates with the cloud big data ecosystem and provides comprehensive cluster management capabilities and a perfect inspection and alarm system, offering customers a simple-to-use, easy-to-maintain, fully managed cloud service to help quickly perform real-time OLAP data analysis.

Product Features

MySQL Protocol Compatibility

A connection API compatible with the MySQL protocol is provided. This enables users to directly use MySQL-related libraries or tools, eliminating the need to deploy new client libraries or tools separately.

High Throughput for Large Queries

The MPP architecture allows the concurrent execution of queries on multiple nodes in a distributed manner, making full use of the overall compute resources of clusters and improving the throughput for large queries.

High Concurrency Point Query

Using technologies such as partition pruning, pre-aggregation, predicate pushdown, vectorized execution, and asynchronous RPC, Tencent Cloud TCHouse-D supports high-concurrency point query scenarios. The concurrency of a 100-node cluster can reach 100,000 queries per second (QPS).

Data Update and Deletion Support

Tencent Cloud TCHouse-D supports deleting and updating data by primary key. It conveniently synchronizes real-time updated data from transactional databases like MySQL.

High Availability and Reliability

Data and metadata are stored in three-replica storage by default. This ensures data reliability even if some nodes go down. Tencent Cloud TCHouse-D automatically checks and repairs corrupted data while routing requests to healthy nodes, ensuring data availability around the clock.

Horizontal Scaling and Data Balancing

Both frontend (FE) and backend (BE) nodes can be scaled horizontally. Users can flexibly scale nodes based on computation and storage needs. When BE nodes are scaled, Tencent Cloud TCHouse-D automatically balances data shards based on the load on the nodes, eliminating the need for manual intervention.

Pre-aggregation Engine

Tencent Cloud TCHouse-D supports storing pre-aggregated data results using a rollup table, improving query efficiency in some aggregation scenarios.

Rich Data Import Features and Transaction Guarantees

Tencent Cloud TCHouse-D supports multiple import methods, including real-time streaming import and large-volume data import. It also supports direct subscription to and consumption of data in Kafka. It provides import transaction support along with the import label mechanism, ensuring uniqueness and integrity of imported data and atomic consistency.

Efficient Column-oriented Storage Engine and Primary and Secondary Indexes

Tencent Cloud TCHouse-D adopts proprietary column-oriented storage format. This format is coupled with various encoding methods, such as dictionary encoding and Run-Length Encoding (RLE), to provide a high data compression ratio that helps save storage space. Moreover, various query acceleration technologies like smart min/max index, sparse index, Bloom filter, and bitmap inverted index are used to boost query efficiency.

Online Table Structure Modification

You can modify the table structure after data is imported, including adding columns, deleting columns, modifying column types, and changing the column order. These operations will not affect the current database queries and writes.

Ecosystem Support and Compatibility with Peripheral Components

Data from Cloud Object Storage, HDFS, and Kafka can be easily imported into Tencent Cloud TCHouse-D, while Flink or Spark can directly write ETL-processed data into Tencent Cloud TCHouse-D. Users can also directly query data in Tencent Cloud TCHouse-D using Spark. Tencent Cloud TCHouse-D can read data from external sources such as MySQL, PostgreSQL, SQLServer, and Oracle through Java DataBase Connectivity (JDBC). It can also read data from Elasticsearch, providing a powerful distributed SQL query layer.

Concepts

Last updated：2024-06-27 10:36:35

Region

A region is a geographical area where a physical Tencent Cloud TCHouse-D server is located. Select a region with caution as you cannot change it after service purchase and networks are completely isolated between regions. For a lower access latency and higher read and write speeds, it is advisable to select a region closest to your users.

Availability Zone (AZ)

An availability zone (AZ) is a physical IDC with isolated resources within a region.

Cluster

A cluster of Tencent Cloud TCHouse-D typically consists of 1-3 frontends (FE) and several backends (BE).

FE

FE is an abbreviation for frontend, which manages metadata, client connections, query planning, and query scheduling.

BE

BE is an abbreviation for backend, which is responsible for data storage, computation execution, compaction, and replica management.

Broker

The Broker process is a transit service to bridge Tencent Cloud TCHouse-D with external data such as HDFS and Cloud Object Storage, providing import and export features.

Tablet

In Tencent Cloud TCHouse-D, a tablet is the logical shard of a table and the basic unit for replica management. According to the partitioning and bucketing mechanism, each table is divided into multiple tablets, which are stored on different BE nodes.

Cluster Architecture

Last updated：2024-06-27 10:36:57

Tencent Cloud TCHouse-D is a fully managed cloud product developed based on open-source Apache Doris. It improves ease of use, security, and stability on top of the open-source version while adopting an architectural design consistent with the open-source Apache Doris system.

Tencent Cloud TCHouse-D mainly includes three components:

Frontend (FE) refers to the frontend nodes, mainly responsible for receiving and returning client requests, metadata and cluster management, and generation of query plans.

Backend (BE) refers to the backend nodes, mainly responsible for data storage and management and execution of query plans.

Broker is an optional process within a cluster, primarily aiding in reading and writing files and directories on remote storage, such as HDFS and Cloud Object Storage.

The overall architecture diagram of Tencent Cloud TCHouse-D is as follows:

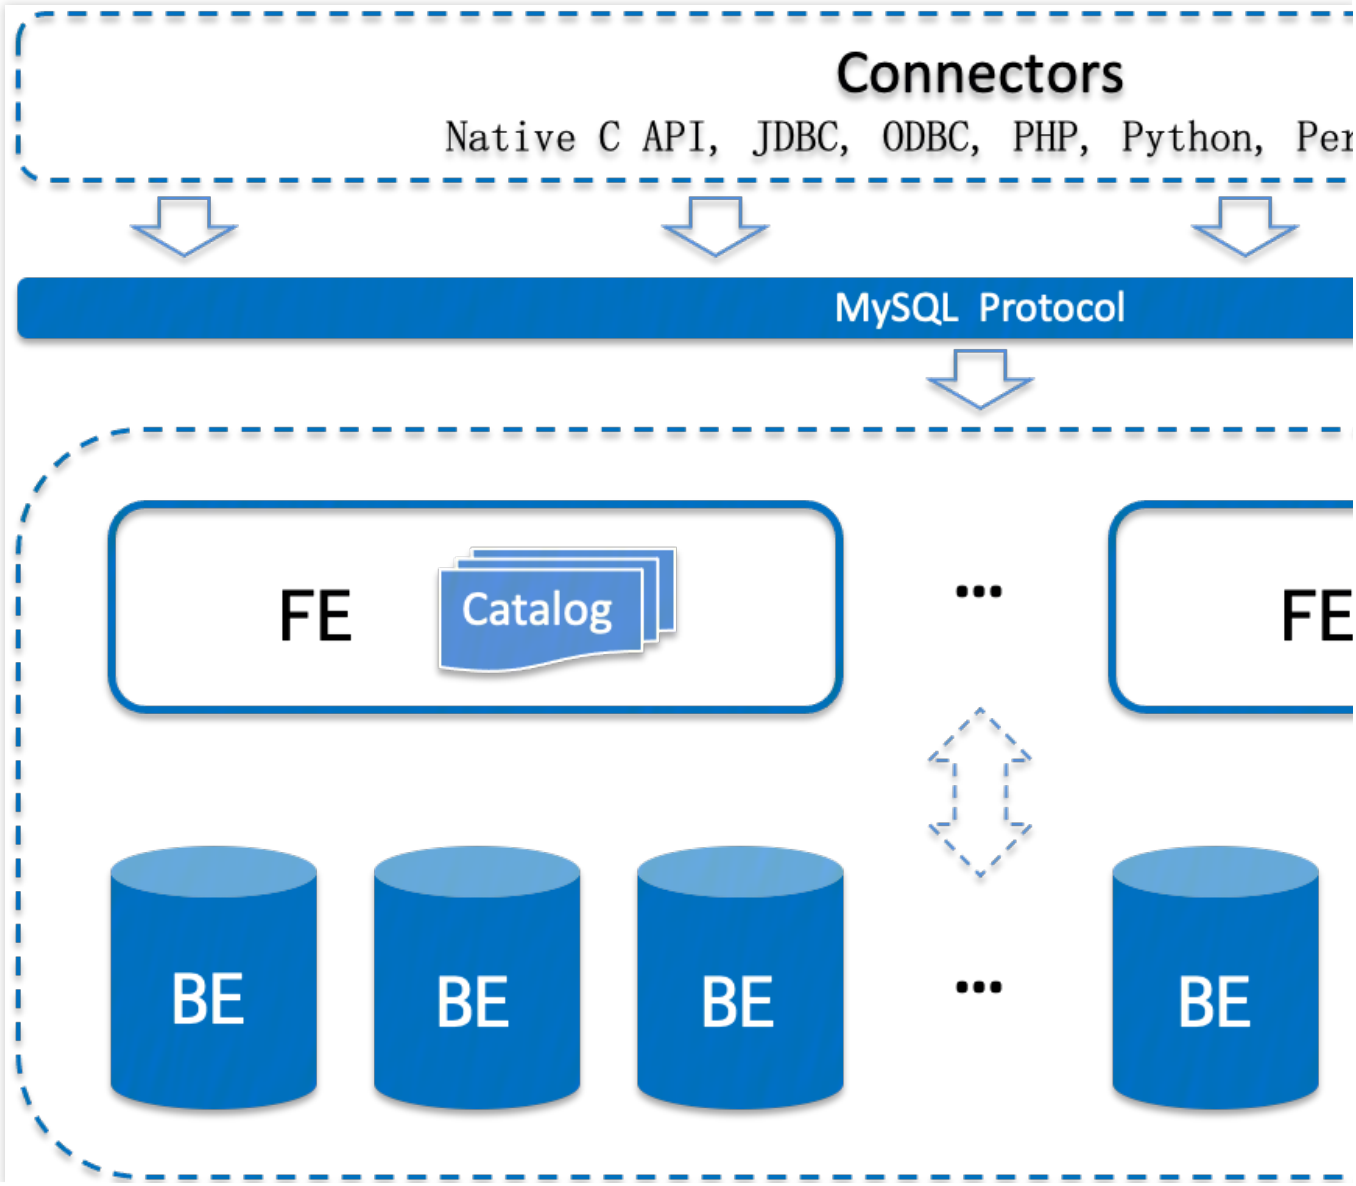

Strengths

Last updated：2024-06-27 10:38:03

Cloud Hosting

Tencent Cloud TCHouse-D seamlessly integrates with cloud-based virtual hosts, cloud disks, Cloud Object Storage, and cloud-based MySQL instances, achieving the elasticity, scalability, security, high availability, and reliability of cloud-native data warehouses. For example, data can be stored in cloud disks and Cloud Object Storage systems, and metadata can be stored in cloud-based MySQL instances.

OPS and Monitoring

Tencent Cloud TCHouse-D comes with a comprehensive set of cluster management features, reducing the need for laborious manual operations. Nearly all OPS tasks, including cluster scaling, up/downgrading, and parameter configuration, can be accomplished through console operations. Moreover, Tencent Cloud TCHouse-D offers visual monitoring features to peek into cluster operations, aiding businesses to be aware of cluster status timely.

Security and Reliability

Tencent Cloud TCHouse-D allows users to separately deploy clusters and supports isolation by Virtual Private Cloud (VPC). This enhances data access security. It also supports a data replica mechanism to ensure seamless service disaster recovery and failover.

MySQL Protocol Compatibility

A connection API compatible with the MySQL protocol is provided. This enables users to directly use MySQL-related libraries or tools, eliminating the need to deploy new client libraries or tools separately. The API provides easy compatibility with upper-layer applications. This also lower the learning curve and make it easy for users to get started.

High Throughput for Large Queries

The Massively Parallel Processing (MPP) architecture allows concurrent execution of queries on multiple nodes in a distributed manner, making full use of the overall compute resources of clusters and improving the throughput for large queries.

High Concurrency for Small Queries

Using technologies such as partition pruning, pre-aggregation, predicate pushdown, vectorized execution, and asynchronous RPC, Tencent Cloud TCHouse-D supports high-concurrency point query scenarios.

Data Update

Tencent Cloud TCHouse-D supports deleting and updating data by primary key. It conveniently synchronizes real-time updated data from transactional databases like MySQL.

High Availability and Reliability

Data and metadata are stored in three-replica storage by default (this feature requires at least 3 backend (BE) nodes). This ensures data reliability even if some nodes go down. In addition, Tencent Cloud TCHouse-D automatically checks and repairs corrupted data while routing requests to healthy nodes, ensuring data availability around the clock.

Horizontal Scaling and Data Balancing

Both frontend (FE) and backend (BE) nodes can be scaled horizontally. Users can flexibly scale nodes based on computation and storage needs. When BE nodes are scaled, Tencent Cloud TCHouse-D automatically balances data shards based on the load on the nodes, eliminating the need for manual intervention.

Pre-aggregation Engine

Tencent Cloud TCHouse-D supports storing pre-aggregated data results using a rollup table, improving query efficiency in some aggregation scenarios.

Efficient Column-oriented Storage Engine

The column-oriented storage format is used to enhance the query efficiency in the Online Analytical Processing (OLAP) field. This format is coupled with various encoding methods, such as dictionary encoding and Run-Length Encoding (RLE), to provide a high data compression ratio that helps save storage space. Moreover, various query acceleration technologies like smart min/max index, sparse index, Bloom filter, and bitmap inverted index are used to boost query efficiency.

Online Table Structure Modification

Tencent Cloud TCHouse-D allows users to modify the table structure after data is imported, including adding columns, deleting columns, changing column types, and changing the column order. These operations will not affect existing query and write operations on the database.

Scenarios

Last updated：2024-06-27 10:38:26

Tencent Cloud TCHouse-D, as an analytical database, is suitable for almost all analytical data scenarios, including the following four main scenarios:

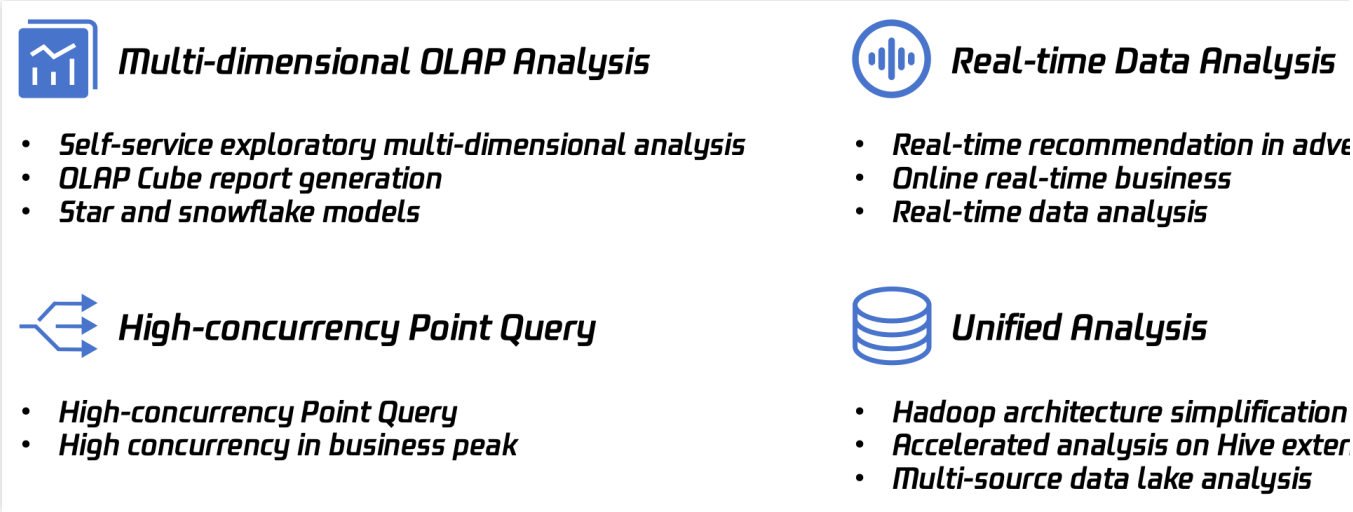

These four scenarios are most anticipated by users in actual business when it comes to analytical databases. However, facing large data volumes, these scenarios demand a highly performant system. Tencent Cloud TCHouse-D meets these stringent requirements on system features and performance through the following technologies.

Multi-dimensional OLAP Analysis—High-speed and Flexible Probing on Multi-dimensional Tables

In relational databases, operations like drilling, rolling up, slicing, dicing, and pivoting on data cubes (CUBE) for multi-dimensional analysis definition are achieved through dimensional modeling.
The most common models in dimensional modeling are the star model and the snowflake model.
In dimensional modeling, columns in a table can be divided into dimension columns and metric columns. Tencent Cloud TCHouse-D supports defining dimension columns and metric columns during table creation, as shown in the figure below.

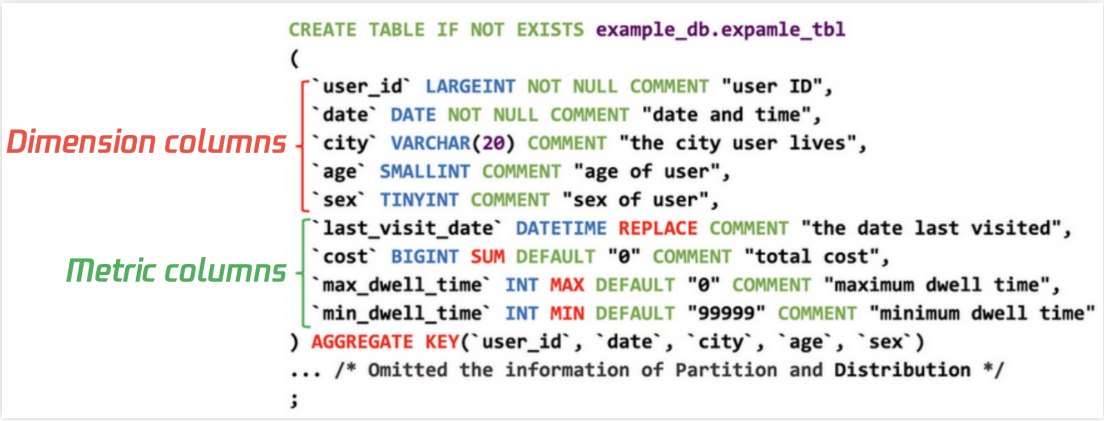

Functions can be defined on metric columns. During data import, data can be categorized according to dimension columns, and then aggregated according to the functions specified by the metric columns. This pre-aggregation capability greatly reduces the amount of data scanned by queries, thereby accelerating aggregate queries.
Also, Tencent Cloud TCHouse-D supports materialized views, rollup indexes, and CUBE syntax. The rollup corresponds to the operation of rolling up, and CUBE can be created through Grouping Set syntax.

Real-time Data Analysis—Real-time Addition, Deletion, Modification, and Query on PB-level Data

The following figure shows a table creation statement of Tencent Cloud TCHouse-D. From this statement, it can be seen that data can be partitioned (Partition) and can also be bucketed (DISTRIBUTED BY). Through partitioning and bucketing, data in a table (Table) can be split into multiple tablets.

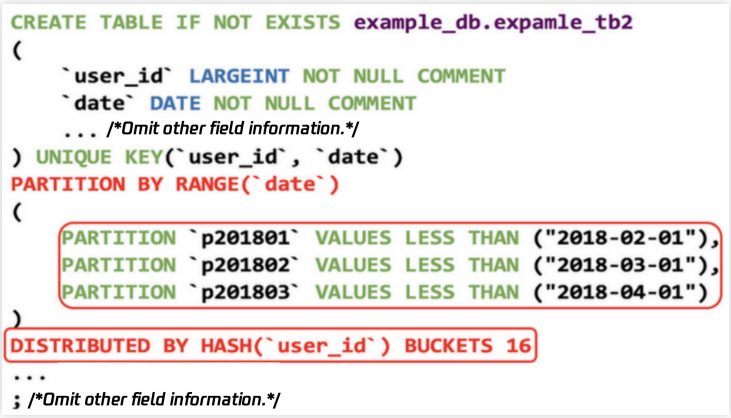

In Tencent Cloud TCHouse-D, each tablet can be set with multiple replicas. These tablets and their replicas can be stored in different backends (BEs), ensuring high data availability and reliability.
Physically, each tablet is split into multiple segment files of a certain size (256 MB). Segment is a column-stored LSM-Tree, also known as Log Structured Merge Tree, which is a layered, orderly, disk-oriented data structure. The

theoretical basis of this structure is that the performance of bulk sequential writing on disk is much higher than random writing.

Also, real-time data input is a very crucial link. To achieve real-time data input, Tencent Cloud TCHouse-D supports multiple input methods:

Real-time data writing through Stream Load;

Real-time access to MySQL binlog through the built-in Canal client;

Precision data import through Doris Flink Connector's integration with Flink CDC;

Real-time data update through subscription to Kafka Topic using the built-in Kafka client.

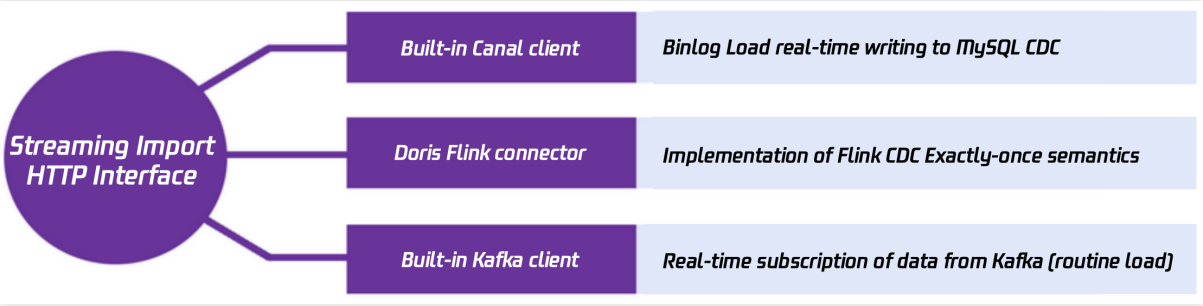

High-concurrency Point Query—Low Latency for High-concurrency Queries

Tencent Cloud TCHouse-D is a modern Massively Parallel Processing (MPP) query engine, which completely implements the Exchange node. With the Exchange node, queries can be decomposed for parallel data processing on nodes.

At the same time, the easy horizontal scaling of frontends (FEs) and backends (BEs) theoretically enables Tencent Cloud TCHouse-D to handle the increase in concurrency, thereby meeting high concurrency requirements.

Additionally, Tencent Cloud TCHouse-D provides a wealth of index structures to help expedite data reading and filtering. Indexes can generally be divided into smart indexes and secondary indexes.

Smart indexes are automatically generated during data writing without user intervention, including prefix sparse indexes and MinMax indexes.

Secondary indexes are auxiliary indexes that users can selectively add to certain columns.

Tencent Cloud TCHouse-D also supports the dynamic partition trimming and predicate push technologies. These technologies can effectively reduce the data volume scanned from the disk, thereby accelerating query execution.

Unified Analysis—Hadoop Ecosystem Compatibility and High-performance External Table Querying

Tencent Cloud TCHouse-D is not dependent on Hadoop components, but it fully supports the Hadoop ecosystem. Apart from writing through Flink and Spark, it can import HDFS data or query Hive data directly by establishing a Hive external table.

The following figure is a panoramic view of Tencent Cloud TCHouse-D's support for the Hadoop ecosystem.