Data Warehouse Service
8.2.0

Technical White Paper

Issue	01
Date	2023-03-30

Contents

1 GaussDB(DWS)

1.1 Product Positioning

GaussDB for Data Warehouse Service (GaussDB(DWS)) is an online data processing database based on the public cloud infrastructure and platform, and provides scalable, fully managed, and out-of-the-box analytic database services. It is a native cloud service based on the Huawei converged data warehouse GaussDB, and is fully compatible with ANSI SQL:99 and SQL:2003 standards, as well as the PostgreSQL and Oracle ecosystems. GaussDB(DWS) provides competitive solutions for PB-level big data analytics in various industries.

GaussDB(DWS) comes in three types: standard data warehouse, stream data warehouse, and hybrid data warehouse, helping you build the top data warehouse in terms of enterprise-level kernel, real-time analysis, collaborative computing, convergent analysis, and cloud native.

Oriented to data analysis scenarios, the standard data warehouse provides enterprise-level data warehouse services with high performance, high scalability, high reliability, high security, and easy O&M. It is capable of data analysis at a scale of 2048 nodes and 20 petabytes of data. Supports hot and cold data analysis, elastic scaling of storage and computing resources, and on-demand and pay-per-use pricing, providing users with elastic, flexible, and cost-effective experience. It is applicable to converged analysis services that integrate libraries, warehouses, cities, and lakes.

The stream data warehouse is built on top of the standard data warehouse. It integrates stream and time sequence engines to provide real-time data ingestion and high-concurrency real-time analysis capabilities. It can be used for IoT real-time analysis.

The hybrid data warehouse provides high-concurrency, high-performance, and low-latency transaction processing capabilities based on large-scale data query and analysis capabilities. It is suitable for hybrid transaction/analytical processing (HTAP). A database can be used for both production and analysis. The can be deployed in a cluster or a single-node system.

GaussDB(DWS) can be used in a wide range of fields, such as finance, Internet of Vehicles (IoV), government and enterprise, e-commerce, energy, and carrier. It has been listed in the Gartner Magic Quadrant for Data Management Solutions for

Analytics for two consecutive years, thanks to its large-scale scalability, enterprise-grade reliability, and higher cost-effectiveness over conventional data warehouses.

1.2 Application Scenario

Data Warehouse Migration

The data warehouse is an important data analytics system for enterprises. However, in this digital age, enterprises' own data warehouses fail to process the ever-growing service volumes due to their poor scalability and high costs. GaussDB(DWS) provides a good answer to these pain points. It is an enterprise-grade cloud data warehouse that boasts high performance, low cost, and smooth scale-out, helping enterprise data warehouses handle the growing data volumes at ease.

Figure 1-1 Data warehouse migration
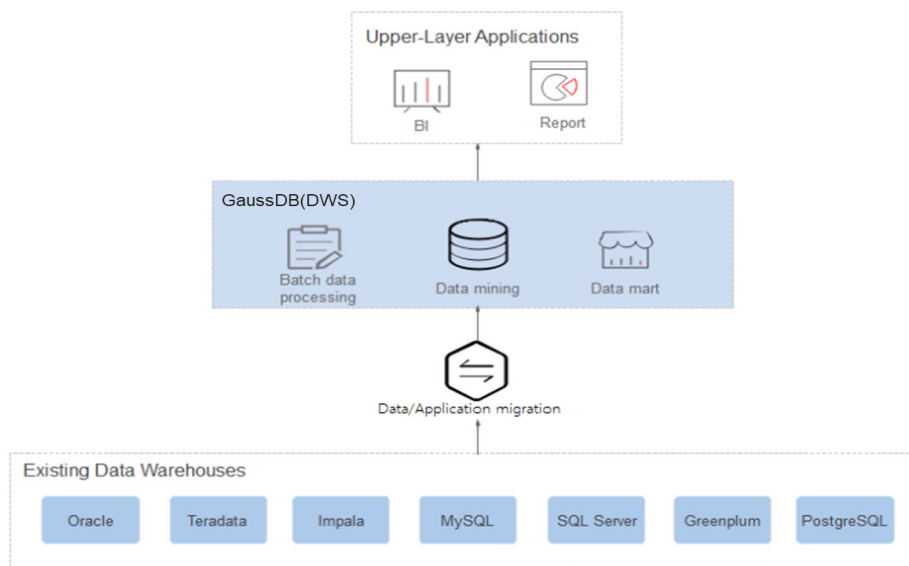
Advantages

● Seamless migration

GaussDB(DWS) provides various migration tools to ensure the smooth migration of popular data analytics systems, such as Teradata, Oracle, MySQL, SQL Server, PostgreSQL, Greenplum, and Impala.

● Compatibility with conventional data warehouses

GaussDB(DWS) supports the SQL 2003 standard and stored procedures. It is compatible with some Oracle syntax and data structures, and can be seamlessly connected to common BI tools, saving service migration efforts.

● Secure and reliable

GaussDB(DWS) supports data encryption and connects to Database Security Service (DBSS) to ensure data security on the cloud. In addition, GaussDB(DWS) supports automatic full and incremental backup of data, improving data reliability.

Convergent Analysis of Big Data

Data has become the most important asset. Enterprises must be able to integrate their data resources and build big data platforms to mine the full value of their data. How to quickly mine values from massive data becomes a key factor for customers to implement predictive analysis.

Figure 1-2 Convergent analysis of big data
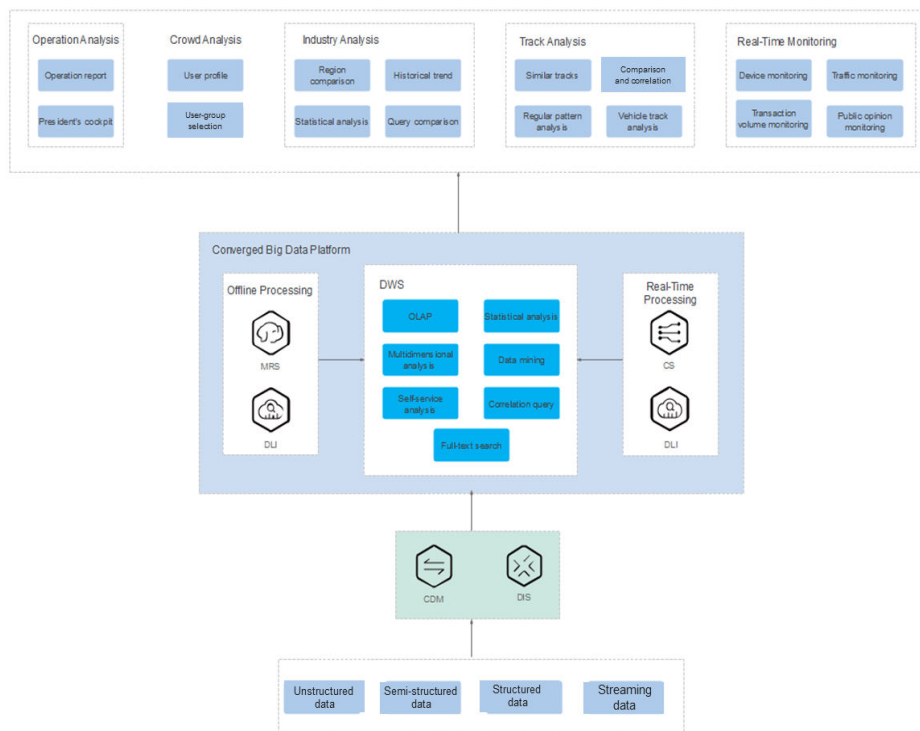
Advantages

● Unified analysis entrance

The GaussDB(DWS) SQL serves as the unified entrance of upper-layer applications, so that application developers can access all data using the SQL.

● Real-time interactive analysis

Analysis personnel can obtain immediately-actionable information from the big data platform in real time using broad analysis requests.

● Auto scaling

Adding nodes allows you to easily expand into PB-range capacity while enhancing query and analysis performance of the system.

Enhanced ETL + Real-Time BI Analysis

The data warehouse is the pillar of the BI system for collecting, storing, and analyzing massive volumes of data. It powers business decision analysis for the IoT, finance, education, mobile Internet, and Online to Offline (O2O) industries.

Advantages

● Data migration

Ability to import data in batches in real time from multiple data sources.

- High Performance

  Cost-effective PB-level data storage and response to correlation analysis of trillions of data records within seconds.

- Real-time processing

  Real-time consolidation of service data to produce actionable insights in operational decision-making.

Figure 1-3 Enhanced ETL + real-time BI analysis
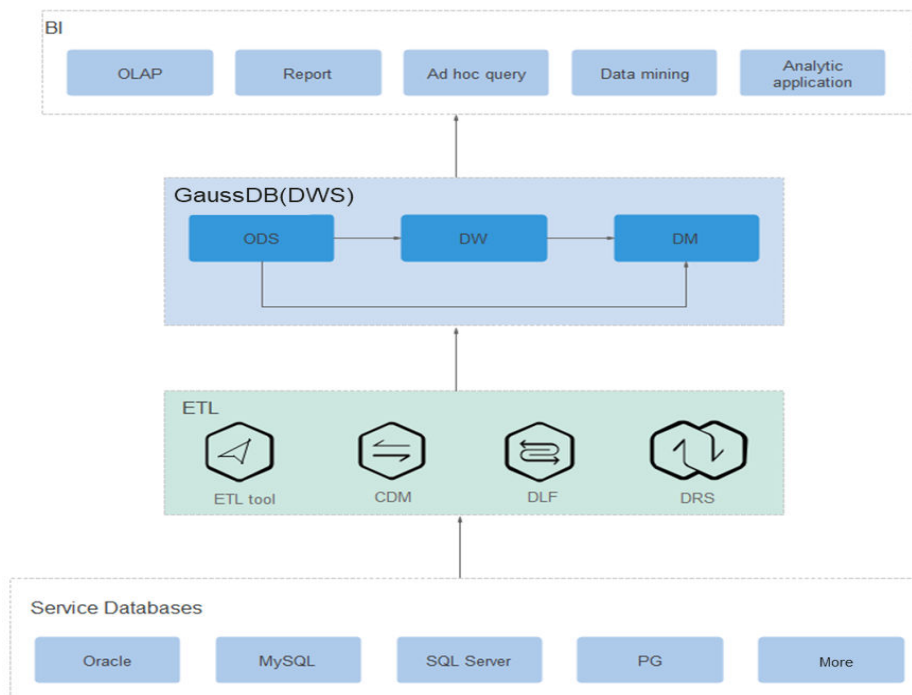
Real-Time Data Analysis

In the mobile Internet and IoT domains, huge volumes of data must be processed and analyzed in real time to extract the full value from data. GaussDB(DWS) quickly imports and queries data and supports real-time data analysis using times series, stream, and AI engines.

Figure 1-4 Real-time data analysis
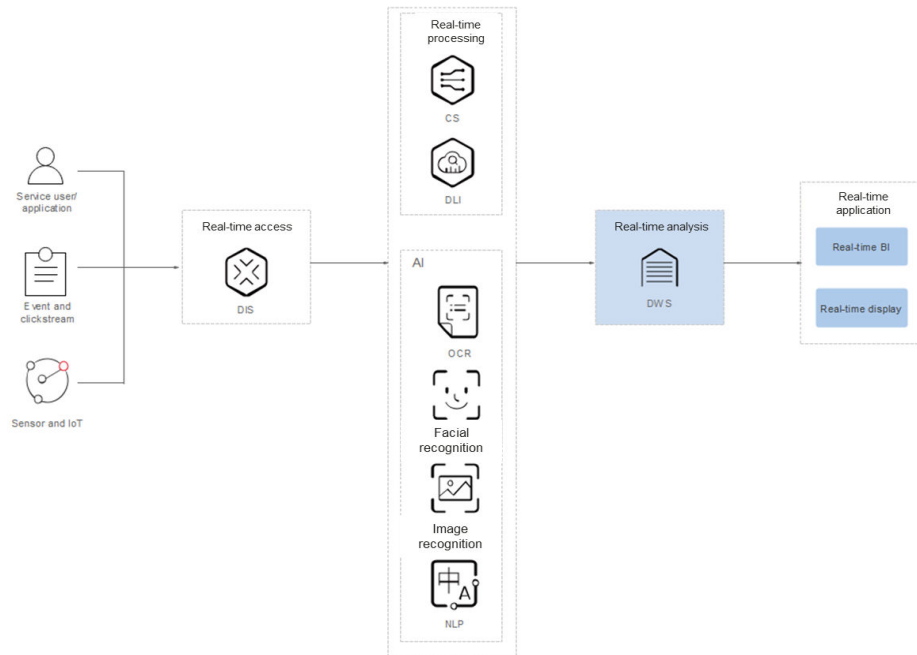
Advantages

- Real-Time import of streaming data

    Data from IoT and Internet applications can be written into GaussDB(DWS) in real time after being processed by the stream computing and AI services.

- Real-Time Monitoring and Prediction

    Device monitoring, control, optimization, supply, self-diagnosis, and self-healing based on data analysis and prediction.

- Convergent AI analysis

    You can conduct association analysis on results of AI-based image and text data analysis and other service data on GaussDB(DWS).

- **IoT**

Figure 1-5 IoT
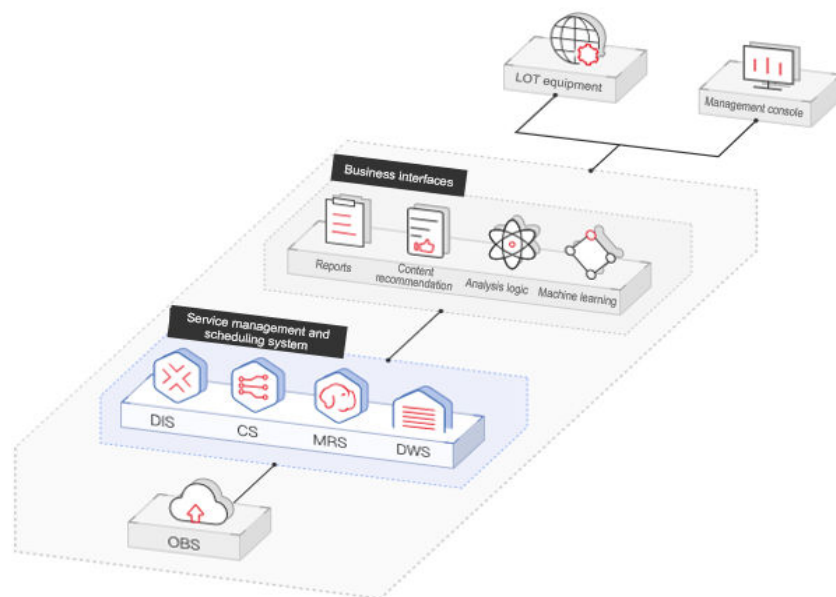
GaussDB(DWS) helps you analyze massive amounts of data from Internet of Things (IoT) in real time and perform optimization based on the results. It is widely used in industrial IoT, O2O service system, and IoV solutions.

Advantages of GaussDB(DWS) are as follows:

– Real-time archiving of stream data: importing stream data from IoT devices and the gateway to GaussDB(DWS) using HUAWEI CLOUD DIS

– Device monitoring and prediction: device monitoring, control, optimization, supply, self-diagnosis, and self-healing based on data analysis and prediction

– Information recommendation: information recommended to users based on the data collected by their networked devices.

1.3 GaussDB(DWS) Architecture and Advantages

GaussDB(DWS) clusters are distributed parallel database clusters based on the shared-nothing architecture.

Figure 1-6 GaussDB(DWS) architecture
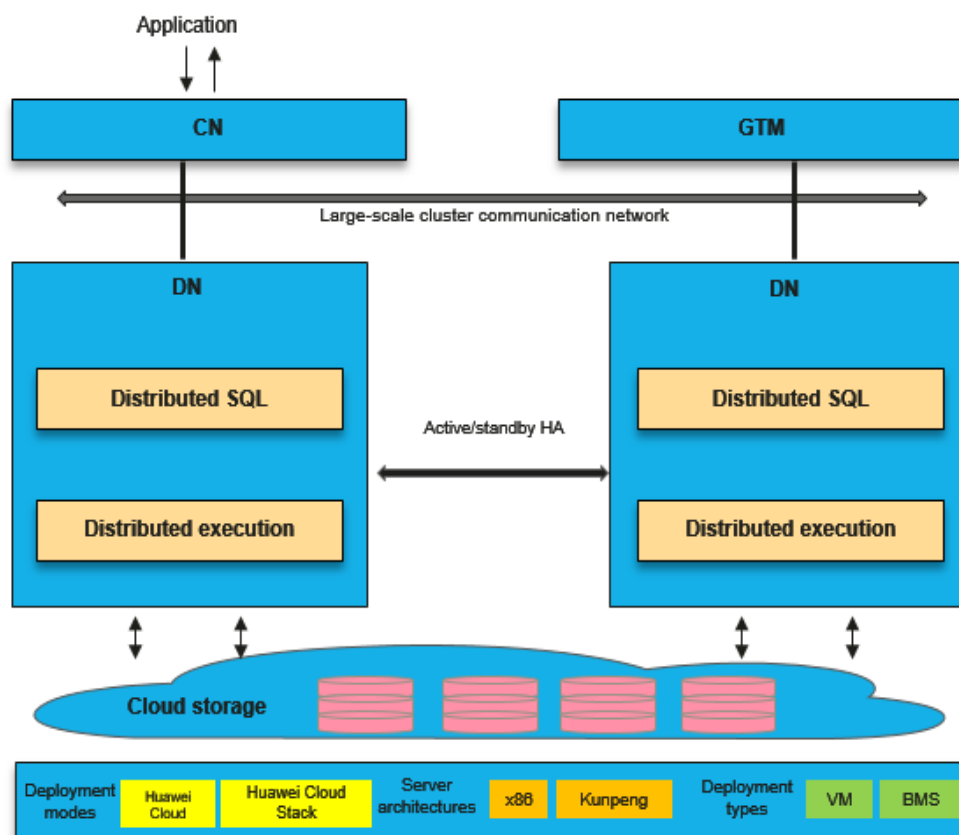
This architecture renders GaussDB(DWS) the following key advantages:

GaussDB(DWS) can be used in multiple cloud forms, including HUAWEI CLOUD, HUAWEI CLOUD Stack, HCSO, and Intelligent EdgeSite (IES), satisfying differentiated deployment and O&M demands.

GaussDB(DWS) uses the Huawei-developed GaussDB database kernel and is compatible with PostgreSQL. The GaussDB database is transformed from a single OLTP database to an enterprise-grade, MPP-based, and distributed OLAP database oriented to massive data analysis.

DWS comes in three types: standard data warehouse, stream data warehouse, and hybrid data warehouse, They provide processing capability of large volume of data across multiple industries and a unified management platform with the following advantages and features:

Ease of use

● Visualized one-stop management

GaussDB(DWS) allows you to easily complete the entire process from project concept to production deployment. With the GaussDB(DWS) management console, you do not need to install data warehouse software or deploy data warehouse servers. GaussDB(DWS) offers you a high-performance and high-availability enterprise-grade data warehouse cluster in a couple of minutes.

With just a few clicks, you can easily connect applications to the data warehouse, back up data, restore data, and monitor data warehouse resources and performance.

- Heterogeneous database migration tools

  GaussDB(DWS) provides various migration tools to migrate SQL scripts of Oracle and Teradata to GaussDB(DWS).

High performance

- Cloud-based distributed architecture

  GaussDB(DWS) adopts the MPP architecture so that service data is separately stored on numerous nodes. Data analytics tasks are quickly executed in parallel on the nodes where data is stored.

- Response to query of trillions of data records within seconds

  GaussDB(DWS) improves data query performance by executing multi-thread operators in parallel, running commands in registers in parallel with the vectorized computing engine, and using the Low Level Virtual Machine (LLVM) compiler to reduce redundant judgment conditions.

  GaussDB(DWS) provides you with a better data compression ratio (column-store), higher index performance (column-store), and better point update and query (row-store) performance.

- Fast data loading

  GaussDB(DWS) provides you with GDS, a high-speed parallel bulk data loading tool.

High scalability

- On-demand scale-out: With the shared-nothing open architecture, nodes can be added at any time to enhance the data storage, query, and analysis capabilities of the system. Up to 2048 nodes can be deployed.

- Enhanced linear performance after scale-out: The capacity and performance increase linearly with the cluster scale. The linear rate is 0.8.

- Service continuity: During scale-out, data can be added, deleted, modified, and queried, and DDL operations (**DROP**/**TRUNCATE**/**ALTER TABLE**) can be performed. Online table-level scale-out ensures service continuity.

Robust reliability

- ACID

  Support for the atomicity, consistency, isolation, and durability (ACID) feature, which ensures strong data consistency for distributed transactions.

- Comprehensive HA design

  All software processes of GaussDB(DWS) are in active/standby mode. Logical components, such as the CNs and DNs of each cluster, also work in active/standby mode. This ensures data reliability and consistency as well as service continuity when any single point of failure (SPOF) occurs.

- Security

  GaussDB(DWS) supports transparent data encryption and can connect to DBSS to better protect user privacy and data security with network isolation and security group rule setting options. In addition, GaussDB(DWS) supports automatic full and incremental backup of data for higher reliability.

Convergent analysis

- Convergence of multiple modes: You can perform direct calculation and convergent analysis of stream, time series, GIS, full-text, and AI data on GaussDB(DWS).

- Convergence of multiple sources: You can use standard SQL statements to query data on the Hadoop Distributed File System (HDFS) and object storage service (OBS) without data migration.

- Cluster acceleration: The shared cluster Express is provided based on OBS data access for more efficient convergent computing and analysis capabilities.

Strong security

- Transparent encryption: Database data files are encrypted to prevent malicious attackers from bypassing the database permission control mechanism at the OS layer or stealing disks to access user data.

- Data masking: Built-in masking functions for digits, characters, and time types are provided. In addition, masking rules can be customized to effectively protect sensitive data while efficiently accessing big data.

2 Platforms and Technical Specifications Supported by GaussDB(DWS)

2.1 Technical Specifications

Software Specifications

Table 2-1 lists the technical specifications of GaussDB(DWS).

Table 2-1 Technical specifications

Technical Specifications	Maximum Value
Data capacity	20 PB
Number of cluster nodes	2048
Size of a single table	1 PB
Size of data in each row	1 GB
Size of a single column in each record	1 GB
Number of records in each table	2^{55}
Number of columns in each table	1,600
Number of indexes in each table	Unlimited
Number of columns in the index of each table	32
Number of constraints in each table	Unlimited

Technical Specifications	Maximum Value
Number of concurrent connections	80 for analytical long transactions, 500 for online short queries, and 5000 for short transactions
Number of partitions in a partitioned table	32,768
Size of each partition in a partitioned table	1 PB
Number of records in each partition in a partitioned table	2^{55}

Product Performance

Table 2-2 lists the performance specifications of GaussDB(DWS).

Table 2-2 GaussDB performance

Performance		Server configurations:
	Data import: 80 MB/s for row-store tables and 150 MB/s for column-store tables per node	CPU: 2 x 12-core Intel Xeon E5-2690 CPUs Memory: 256 GB Disk: 20 x 600 GB SAS disks Network: 10GE
	Data export: 100 MB/s per node	
	Full table scan: 300 million records per second	
	Point query: precise query of trillions of records in seconds	
	GROUP BY: 20 million records per second per node	
	JOIN: 5 million records per second per node	
	ORDER BY: 3 million records per second per node	

3 GaussDB(DWS) Core Technologies

3.1 Shared-Nothing Architecture

GaussDB(DWS) is an MPP system employing the shared-nothing architecture. It consists of multiple independent logical nodes that do not share system resources with each other, such as CPUs, memory, and storage. In such an architecture, service data is stored on multiple physical nodes. Data analysis tasks are executed in parallel on the nodes where data is stored. The massively parallel data processing significantly improves response speed.

Figure 4-1 compares the shared-nothing architecture and other architecture.

Figure 3-1 Architecture comparison
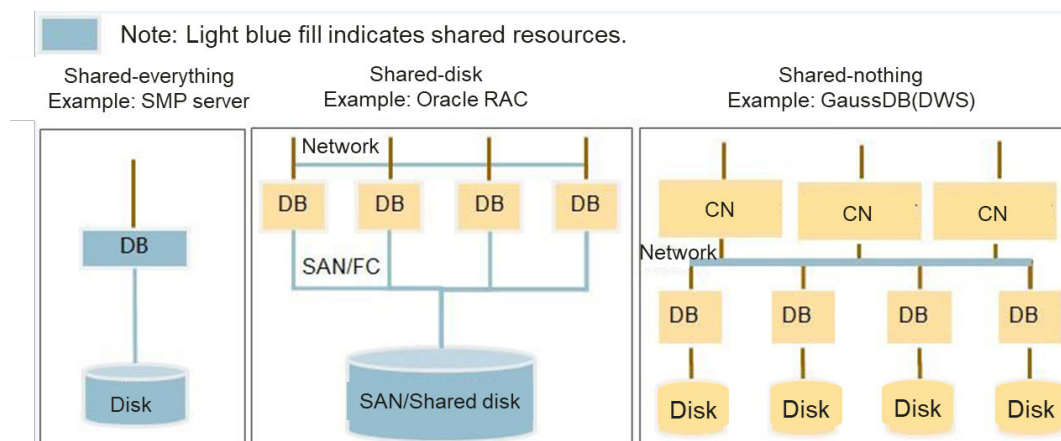
The shared-nothing architecture has the following advantages:

● High scalability
  – Provides on-demand scaling for BI and data analysis to process high-concurrency massive data.
  – Provides an automatic parallel processing mechanism.
● Implements automatic internal parallel processing without manual partitioning or optimization.

- – Allows you to load and access data as you would in a common database.
    - – Distributes data on all parallel nodes.
    - – Every node processes only partial data.
- Optimal I/O processing
    - – All nodes process data in parallel.
    - – Nodes share nothing with each other and have no I/O conflicts.
- Increased storage space and improved query and loading performance with added nodes

3.2 Data Distribution in a Distributed System

Background

DWS uses a two-layer data layout mechanism achieve high-performance query and import of PB-level data. At the first layer, users can specify a data distribution policy (hash distribution or replication distribution) when creating a table. When data is written to the system, the system determines the node where the data is stored based on the corresponding distribution policy. At the second layer, the node partitions its stored data according to partitioning rules.

3.2.1 Distributed Data Storage

GaussDB(DWS) horizontally partitions tables, distributing tuples in a table to multiple nodes. This allows you to filter out unnecessary data when querying to quickly locate data and significantly improve database performance.

Horizontal partitioning distributes data in a table to multiple nodes based on a specific data distribution policy.GaussDB(DWS) supports the data distribution policies described in **Table 3-1**. When executing the CREATE TABLE statement, you can configure the **DISTRIBUTE BY** parameter to enable data distribution on a specific table.

Table 3-1 Distribution policies

Policy	Description	Scenario	Advantage & Disadvantage
Replicatio n	Full data in a table is stored on each DN in the cluster.	Small tables and dimension tables	• The advantage of replication is that each DN has full data of the table. During the join operation, data does not need to be redistributed, reducing network overheads and reducing plan segments (each plan segment starts a corresponding thread). • The disadvantage of replication is that each DN retains the complete data of the table, resulting in data redundancy. Generally, replication is only used for small dimension tables.
Hash	Table data is distributed on all DNs in the cluster.	Fact tables containing a large amount of data	• The I/O resources of each node can be used during data read/write, greatly improving the read/write speed of a table. • Generally, a large table (containing over 1 million records) is defined as a hash table.

Policy	Description	Scenario	Advantage & Disadvantage
Polling (Round-robin) **Supported by 8.1.2 and later versions**	Each row in the table is sent to each DN in turn. Data can be evenly distributed on each DN.	Fact tables that contain a large amount of data and cannot find a proper distribution key in hash mode	• Round-robin can avoid data skew, improving the space utilization of the cluster. • Round-robin does not support local DN optimization like a hash table does, and the query performance of Round-robin is usually lower than that of a hash table. • If a proper distribution key can be found for a large table, use the hash distribution mode with better performance. Otherwise, define the table as a round-robin table.

3.2.2 Data Partitioning

Description

Most database products partition data. In the DWS distributed system, data partitioning is a process of horizontally partitioning data on a node based on a specified policy. A table is divided into multiple partitions based on a specified range, and data in different partitions does not overlap.

GaussDB(DWS) supports range partitioning and list partitioning. In range partitioning, records are divided and inserted into multiple partitions of a table. Each partition stores data of a specific range (ranges in different partitions do not overlap). If you configure the **PARTITION** parameter when running the **CREATE TABLE** statement, data in the table will be partitioned.

Benefits

Table 3-2 uses an xDR scenario to describe the benefits provided after data is partitioned based on time slices.

Table 3-2 Partitioning benefits

Scenarios	Benefits
The rows frequently accessed in a table are located in one or a few partitions.	Significantly reduces search space and improves access performance.
Most partition records need to be queried or updated.	Significantly improves performance because only one partition rather than the whole table needs to be scanned.
Records that need to be loaded or deleted in batches are located in one or a few partitions.	Improves processing performance because related partitions can be directly read or deleted. Reduces de-fragmentation workloads because records can be deleted in batches.

Data partitioning provides the following benefits:

- **Improves manageability:** Tables and indexes are divided into smaller and more manageable units, In this way, data management can be performed by partitions. Database administrators will perform maintenance in the designated area of the table.

- **Improves deletion performance:** You can delete an entire partition rather than delete data row by row.

  The syntax for deleting a partitioned table and a common table is the same: DROP TABLE.

- **Improves query performance:** You can restrict the volume of data to be checked or manipulated to make queries quicker.

  With partition pruning, also known as partition elimination, the CN filters out unexpected partitions and scans only the remaining partitions. Partition pruning greatly improves query performance.

  Intelligent partition connection: Partitioning can also improve the performance of multi-table joins by using a technique known as partition-wise joins. Partition-wise joins can be applied when two tables are joint and at least one of these tables is partitioned using a join key. Partition-wise joins break a large join into smaller joins of "identical" data sets. "Identical" here is defined as covering the same set of partitioning key values on both sides of the join, ensuring that only a join of these 'identical' data sets will produce a result and that other data sets do not have to be considered.

3.2.3 Parallel Data Import

Principles

Importing data in parallel on multiple nodes fully uses the computing and I/O capabilities of the nodes to maximize speed. The parallel data import function of GaussDB(DWS) implements high-speed and parallel import of external data in a specified format (CSV or TEXT).

Parallel data import is more efficient than the traditional data import method in which the INSERT statement is used to insert data. The procedure for importing data in parallel is as follows:

- The CN only plans and delivers data import tasks, and the DNs execute these tasks. This reduces CN resource usage, enabling the CN to process external requests.
- The computing capability and network bandwidth of all the DNs are fully utilized, improving data import performance.

The following uses the Hash distribution policy as an example to describe the GaussDB(DWS) data import process. **Figure 3-2** shows the parallel data import process.

Figure 3-2 Parallel data import
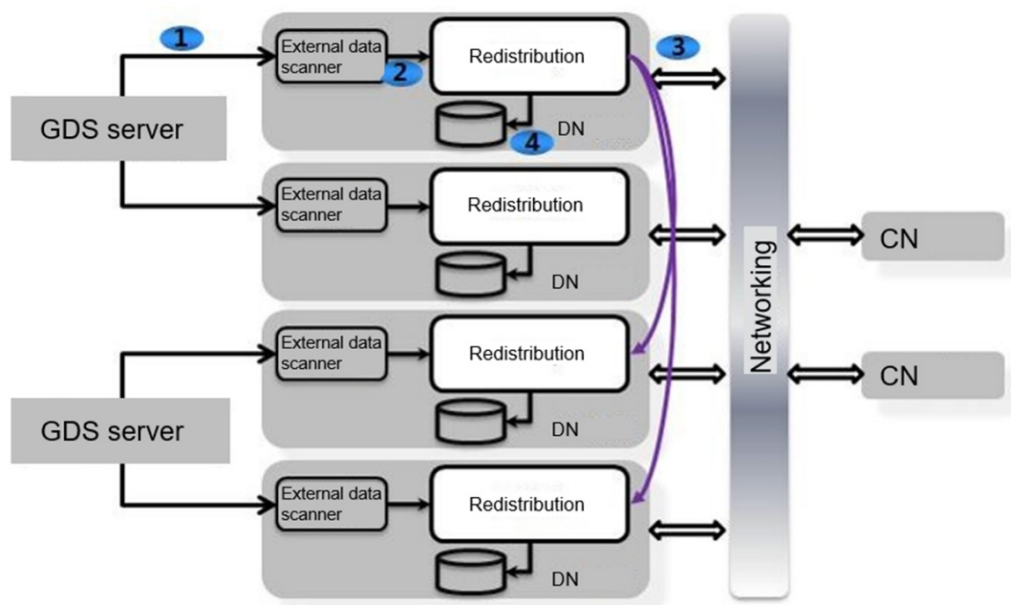
Table 3-3 Procedure description

Process	Description
Creating a table that complies with the Hash distribution policy	When running the **CREATE TABLE** statement, a service application presets the Hash distribution policy (specifies an attribute of a table as a distribution field).
Setting the partitioning policy	When executing the CREATE TABLE statement, a service application presets a partitioning rule (specifies an attribute of a table as a partitioning field). All Hash data in each DN is partitioned based on the preset partitioning rule.
①	During data import, GDS splits a specified data file into data blocks with a fixed size.

Process	Description
②	DNs download these data blocks from GDS in parallel.
③ ④	Each DN processes data blocks in parallel and parses out a data tuple from the data blocks. The physical location of each tuple is determined based on the Hash value obtained based on the distribution column. ● If data is distributed on remote nodes based on the Hash values, you need to redistribute it to target DNs. ● If data is distributed on local nodes based on the Hash values, store it on local DNs.
Writing data into partitions	After data is sent to the node where Hash is used, it is written into the partition data file based on the partitioning logic. While data is written into a partitioned table in GaussDB(DWS), you can exchange partitions to improve the writing performance.
General Data Service (GDS): Multiple GDSs can be deployed on a data server to improve the import performance.	

3.3 Fully Parallel Query

Description

Fully parallel distributed query processing is the core technology of GaussDB(DWS). It minimizes data flow between nodes during query, thereby making query more efficient.

To efficiently analyze data, GaussDB(DWS) employs a set of high-performance distributed executors, which input the execution plan generated by the SQL engine, process tuples based on the execution plan, and return the result to the client.

Technical Principles

Figure 3-3 demonstrates the fully parallel distributed query technology of GaussDB(DWS).

Figure 3-3 Fully parallel distributed query processing technology
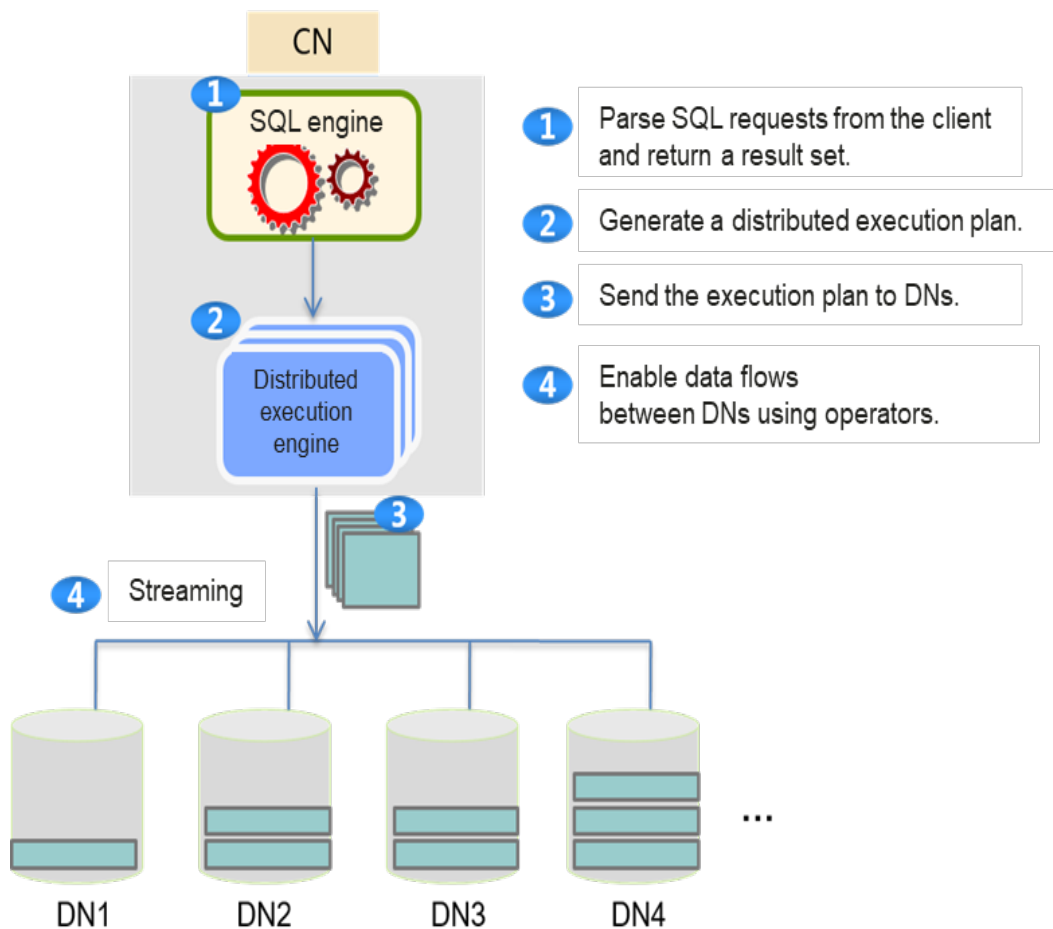
- The distributed executor running on the CN provides schedules and distributes executions.

- A new execution operator is introduced to support data flow between DNs. The new operator is called the data flow operator. Data flow can be classified into Gather, Broadcast, and Redistribution flows, based on the relationship between input and output. **Gather** combines multiple query fragments of data into one. **Broadcast** forwards the data of one query fragment to multiple query fragments. **Redistribution** reorganizes the data of multiple query fragments and then redistributes the reorganized data to multiple query fragments.

- Data transmission between DNs depends on the data stream topology constructed using the data distribution and the cost model during query and analysis. By the data stream topology, network connections are set up between DNs to drive data stream on the data flowing topology.

# 3.4 Vectorized Executor and Hybrid Row-Column Storage Engine

Background

In a wide table containing a huge amount of data, a query usually only involves certain columns. In this case, the query performance of the row-store engine is poor. For example, a single table containing the data of a meteorological agency has 200 to 800 columns. Among these columns, only 10 are frequently accessed. This is where vectorized executor technology steps in, by improving performance and saving storage space.

Vectorized Execution

Figure 3-4 shows a standard vectorized executor. Control flow travels in the downlink direction (shown as solid lines in the following figure) and data flow in the uplink direction (shown as dotted lines in the following figure). The upper-layer node invokes the lower-layer node to ask for data and the lower-layer node only returns one tuple to the upper-layer node at a time.

Instead of returning only one tuple at a time in the traditional executor, the vectorized executor returns a batch of tuples at a time, which significantly improves executor performance with the aid of the column storage feature.

Figure 3-4 Vectorized executor

## Hybrid Row-Column Storage Executor

GaussDB(DWS) supports both the row and column storage models. You can choose a row- or column-store table as needed.

Generally, if a table contains many columns (called a wide table) and its query involves only a few columns, column storage is recommended. If a table contains only a few columns and a query includes most of the fields, row storage is recommended.

The hybrid row-column storage engine achieves higher data compression ratio (column storage), index performance (column storage), and point update and point query (row storage) performance, as shown in **Figure 3-5**.

Figure 3-5 Hybrid row-column storage engine
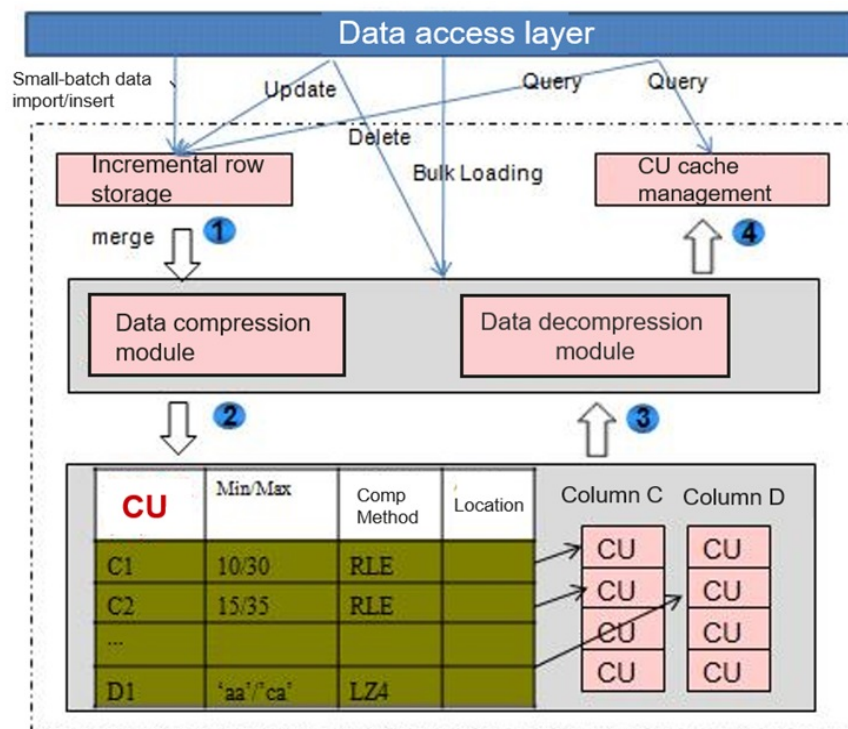
Data compression is supported for column storage. You can compress old, inactive data to free up space, reducing procurement and O&M costs.

In GaussDB(DWS), data can be compressed using the Delta Value Encoding, Dictionary, RLE, LZ4, and ZLIB algorithms. The system automatically selects a compression algorithm based on data characteristics. The average compression ratio is 7:1. Compressed data can be directly accessed and is transparent to services, greatly reducing the preparation time before accessing historical data.

The restrictions of the column storage engine are as follows:

- Only the CREATE TABLE, DROP TABLE, and TRUNCATE TABLE operations can be performed for DDL.

  Partition management using DDL statements (such as ADD PARTITION, DROP PARTITION, MERGE PARTITION, and EXCHANGE) can be used.

  The **CREATE TABLE LIKE** statement is supported.

  The **ALTER TABLE** statement is partially supported.

  Other DDL statements are not supported.

- Among DML syntax, UPDATE, COPY, BULKLOAD, and DELETE are supported.

- A trigger and primary foreign key are not supported.

- Psort indexes, B-tree indexes, and GIN indexes are supported. For details about the constraints, see "**CREATE INDEX**" in *Data Warehouse Service (DWS) Developer Guide*.

3.5 Resource Monitoring and Management

Background

To use resources appropriately, GaussDB(DWS) provides resource monitoring and management methods to allocate compute and storage resources that affect job running. This prevents performance deterioration or even system running problems caused by improper resource usage. Resource monitoring and management include resource monitoring, load management, and disk space management.

3.5.1 Load Management

Description

Load management balances system compute resources through service concurrency control to prevent resource contention between services, achieving harmonious coexistence of jobs and optimal resource utilization. In addition, the cgroup technology is introduced to manage CPU quotas.

Technical Principles

Load management is classified into static load management and dynamic load management, which are controlled by the **enable_dynamic_workload** parameter. In addition, GaussDB(DWS) provides priority control to control the priorities of tenant jobs.

- Static Load Management

  Each CN controls the memory and concurrency separately. The actual concurrency of a cluster is the sum of the concurrency of all CNs. The actual memory used by the cluster is the sum of the memory occupied by the jobs running on each CN. Because the memory of each CN is controlled separately, the actual memory used by a tenant on the DN may exceed the memory quota of the tenant.

- Dynamic Load Management

  The CCN is added to control the concurrency and memory of complex jobs. Each CN requests queuing information from the CCN. The CCN accumulates the estimated memory used by tenants to obtain the used memory of each tenant. When the available memory of a tenant is insufficient, the CCN triggers queuing. After the job running is complete, the CCN attempts to wake up the jobs in the queue. The sum of used memory of all tenants is the used memory of the cluster. In addition, the CCN periodically collects the memory usage information on DNs to update the available memory of the cluster. If the available memory is insufficient, jobs are queued. The job running must meet the memory limits of both the tenant and cluster.

- Priority Control

  GaussDB(DWS) implements load identification and intra-queue priority control based on query_band. It provides more flexible load identification methods and identifies load queues based on job types, application names,

and script names. Users can flexibly configure query_band identification queues based on service scenarios. In addition, priority control of job delivery in the queue is implemented. In the future, priority control of resources in the queue will be gradually implemented.

Benefits

Load management implements resource isolation and priority control between tenants. It ensures that jobs of tenants with high priorities run preferentially without affecting resources of tenants with low priorities, achieving effective and optimal resource utilization.

For more information, see **Resource Management**.

3.5.2 Space Control

Description

User space management includes permanent tablespace management, temporary tablespace management, and intermediate calculation result set flushing space management. When creating group users and service users, the database administrator can specify the permanent tablespace, temporary tablespace, and intermediate calculation result set flushing space limit. When a user's operation involves the increase or decrease of the managed space, the system checks whether the space exceeds the specified limit. If the space exceeds the limit, the user is not allowed to perform the operation and the available disk space is limited to the specified space limit.

In addition, the product supports schema-based permanent space control. When the space of objects in a schema are added or deleted, the operations that exceed the limit are controlled.

Technical Principles

When a service user executes a service, the size of the data space operated by the service user is recorded in the space quota of the service user. For a permanent tablespace, the quota is evenly allocated to each DN. DNs collect space statistics in real time when data files, temporary table files, and temporary files are added or expanded. When a service user performs a service operation, DNs check whether the space limit is reached and determine whether to perform the operation. If the space limit is exceeded, DNs roll back the operation. When a transaction or session ends, the space is released and the corresponding value is updated.

DNs collect statistics on disk changes involved in user operations in real time, save the statistics to the memory, wait for the CN to collect the statistics, and receive the information delivered by the CN to update the temporary space used by the entire system.

Benefits

By limiting the maximum storage space used by a user, the system prevents user services from squeezing or occupying too many storage resources to ensure that normal service data can be written, ensuring system stability and data availability.

For more information, see "Space Control"

3.6 Distributed Transactions

Background

In the distributed share nothing architecture, table data is distributed on different nodes. One or more statements on the client may modify data on multiple nodes at the same time. In this case, distributed transactions are generated. Pay attention to the following key points of distributed transactions:

1) Atomicity of transactions on each node: Distributed transactions are either all successful or all failed on all nodes.

2) Transaction consistency: The data returned on each node is the same. Data consistency cannot be ensured when a node is faulty.

The atomicity of distributed transactions must be ensured. The consistency of distributed transactions depends on the CAP theory. The common standards are CP systems that support strong consistency 2pc and 3pc protocols or AP systems that support eventual consistency TCC and message tables.

Technical Principles

GaussDB(DWS) supports strongly consistent distributed transactions and provides high-availability with consistency and partition tolerance (CP).

CSN
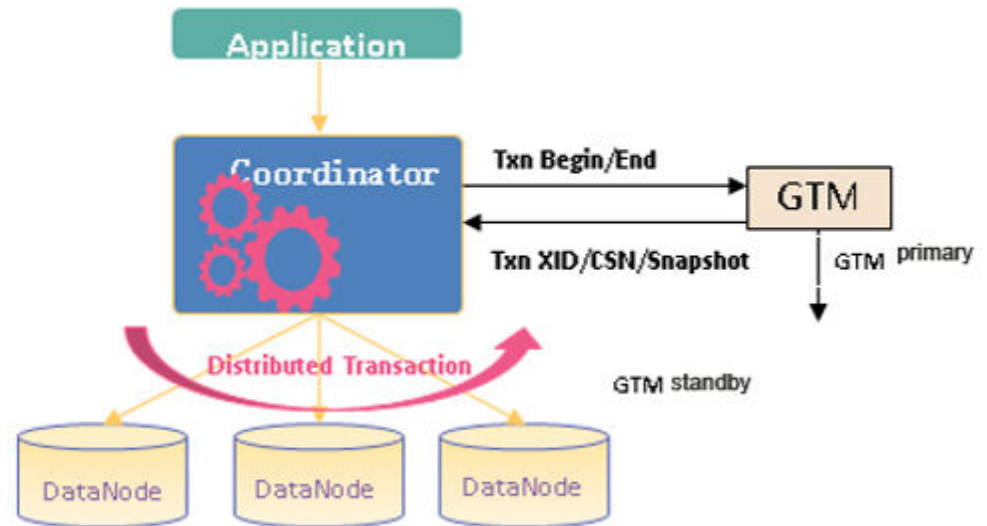
CSN update and mapping between CSNs and XIDs
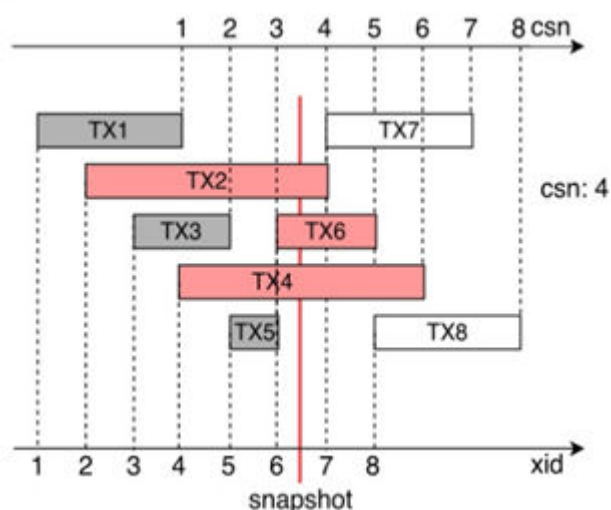
CSN – Commit Sequence Number: transaction submission number

The CSN is an 8-byte unsigned integer that increases monotonically and is maintained by the GTM.

2) When a transaction ends, the CSN value is updated from the GTM.

3) After the CSN mechanism is used, the CSN can be obtained from the GTM.

GTM

GTM is a component in the GaussDB(DWS) distributed framework. It has the following functions:

1. Manage and allocate transaction IDs (increases but does not decrease).

2. Manage and maintain CSN numbers (increase but not decrease)

When executing a modification operation, the CN obtains the transaction ID from the GTM.

At the beginning of the statement, the CN obtains a CSN from the GTM for query.

When a transaction starts or ends, the CN communicates with the GTM to register and destroy transaction information.

Troubleshooting

The gs_clean tool is used to automatically clear residual distributed transactions caused by node faults.

gs_clean queries the residual two-phase transactions on each node, checks whether the transactions are submitted or rolled back on other nodes based on the residual transaction IDs, and clears the residual two-phase transactions based on the final result.

Benefits

GaussDB(DWS) supports strongly consistent distributed transactions. You can use the GaussDB(DWS) database in the same way as using a standalone database.

After the CSN-based transaction mechanism is used, the concurrency performance is greatly improved.
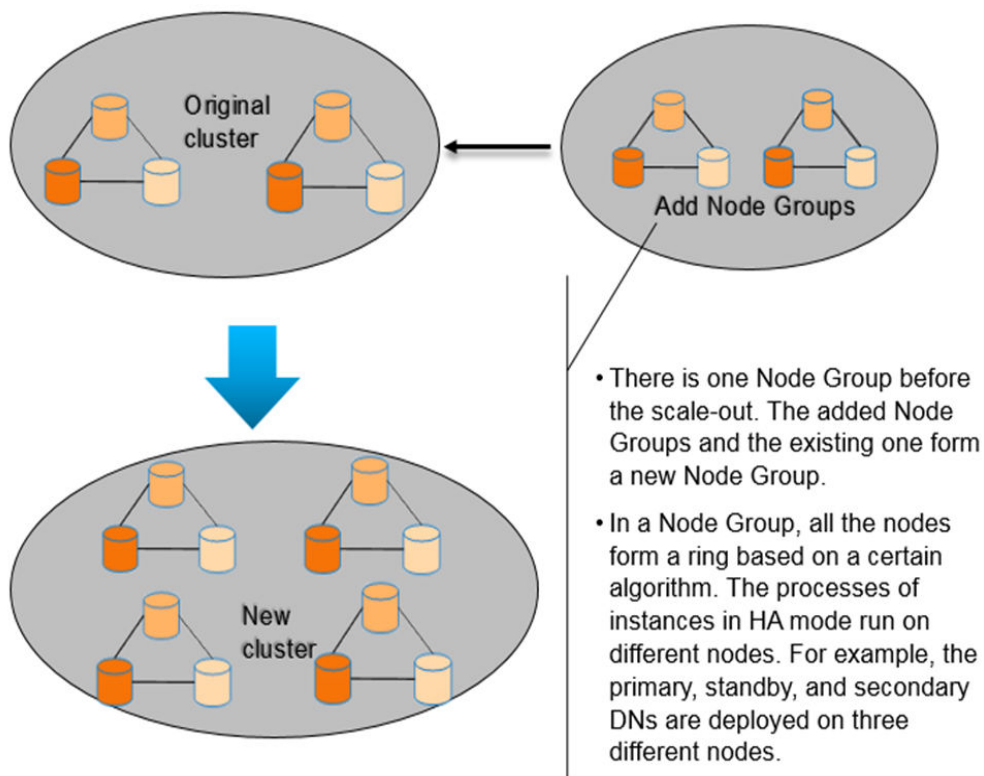
3.7 Online Scale-Out

Background

As your business expands, your current system may be unable to provide sufficient disk capacity and satisfactory performance. The distributed GaussDB(DWS) cluster allows you to fully utilize existing machines to scale the existing capacity out, meeting your business needs.

Technical Principles

The node group technology of GaussDB(DWS) supports the scale-out of multiple tables in parallel with a speed up to 400 GB per hour on each new node.

Figure **Figure 3-6** shows the GaussDB(DWS) scale-out process.

Figure 3-6 Scale-out

- There is one Node Group before the scale-out. The added Node Groups and the existing one form a new Node Group.

- In a Node Group, all the nodes form a ring based on a certain algorithm. The processes of instances in HA mode run on different nodes. For example, the primary, standby, and secondary DNs are deployed on three different nodes.

Benefits

Online scale-out of GaussDB(DWS) has the following advantages:

- Scale-out does not interrupt services.

  Data import and query in GaussDB(DWS) are not interrupted during scale-out.

- Consistent hashing and multi-table parallel scale-out technologies improve query performance.

  Hash technology is used to minimize the amount of data migrating during redistribution.

  Tables can be redistributed in parallel or in user-specified order.

  Users can query the scale-out progress in real time.

- As the number of nodes increases, the cluster performance improves linearly.

  As shown in **Figure 3-7**, in the full parallel distributed architecture, the data loading performance, service processing performance, and capacity of a GaussDB(DWS) cluster linearly improves with the number of nodes.

Figure 3-7 High-performance scalability
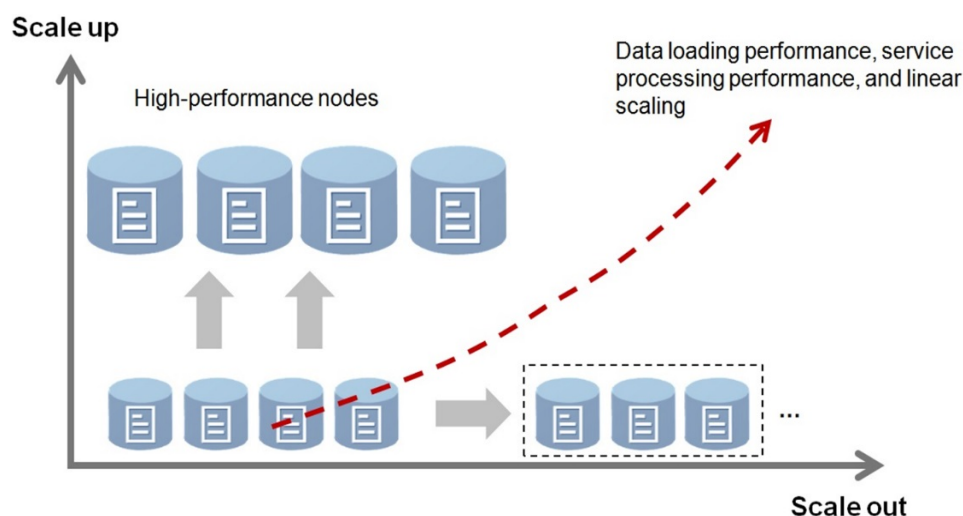
3.8 SQL on Anywhere

Background

Various types of engine components provide many interfaces for big data processing. However, for traditional database users, the SQL language is still the most familiar and convenient interface. Their operations will be efficient using SQL statements on a client to perform operations on big data components.

GaussDB(DWS) supports SQL on Anywhere, allowing users in a GaussDB(DWS) database to perform operations on the following systems to build a unified big data computing platform: Hadoop, Spark.

3.8.1 SQL on Hadoop

In DWS, you can directly read the structured data stored in HDFS, use the SQL query interfaces it provides, and perform complex analysis and query about the Hadoop native data using the vectorized executor.

Technical Architecture

DWS maps structured HDFS data to its foreign table. You can use the well-developed SQL compiler, vectorized executor, and SQL interfaces provided by the database to analyze and query native Hadoop data in HDFS.

The SQL compiler of DWS introduces a new component, Scheduler. If access to HDFS is required during SQL compiling, the Scheduler accesses the metadata of the HDFS NameNode to obtain the storage path of external files in HDFS. When the SQL engine of the CN is compiling query statements, it will invoke the Scheduler to obtain and specify the directory that contains the HDFS file to be accessed by each DN.

The allocation principles of the HDFS external data files (between DNs) are as follows:

- A DN gives priority to data stored on its own physical server node when reading data.

- Then, workload balancing between DNs is considered.

Figure 3-8 Logical architecture
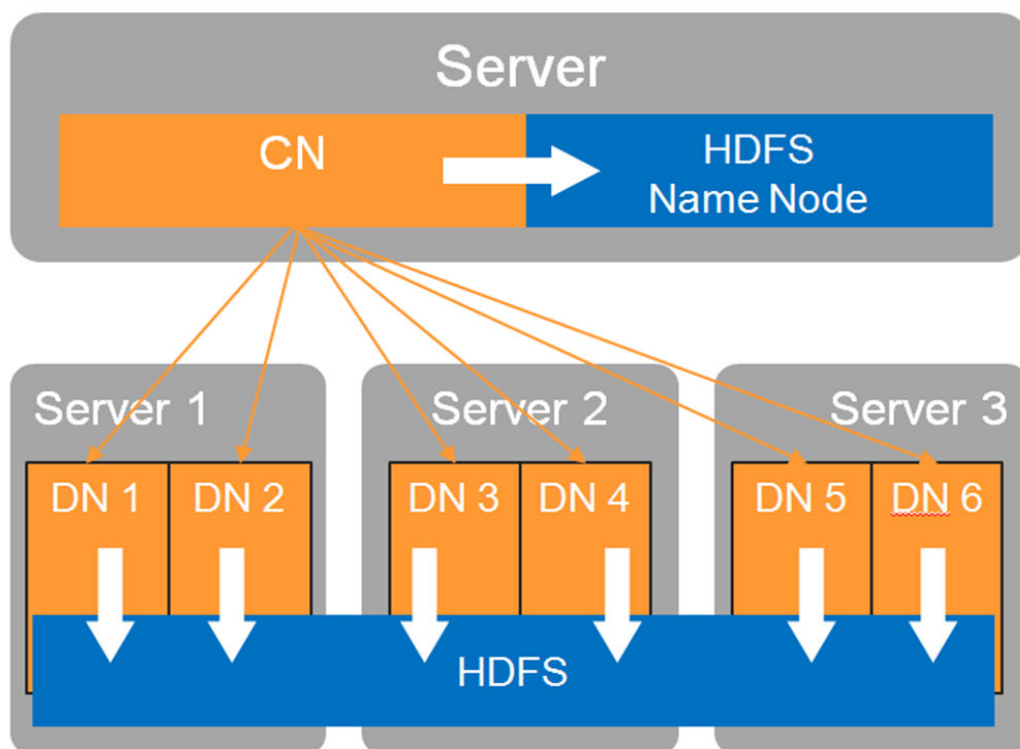
Technical Highlights

- **Standard SQL query syntax**

  The SQL query syntax inherited from GaussDB(DWS) supports structured data analysis and query in HDFS, supports system functions such as connection, aggregation, and character date, supports subqueries and joint access to HDFS structured data and GaussDB(DWS) local data, and supports window functions.

- **HDFS data cost estimation model**

  In addition to the cost-based SQL optimizer of DWS, the product uses the cost evaluation model for structured data access in the HDFS to formulate the best execution plan of HDFS data.

- **Intelligent scanning**

  DWS can directly push predicates to the native Hadoop data stored on HDFS, filter predicates in the compressed data, and perform late materialization for the Hadoop Optimized Row Columnar (ORC) storage, reducing the amount of data to be read in the HDFS.

- **Low Level Virtual Machine (LLVM) optimization**

  Predicates in native Hadoop data are scanned and optimized by LLVM. Then, intermediate representations (IRs) for their conditions are generated and further converted into machine codes. In this way, the performance of filtering and querying predicates is improved.

- **Informational constraints**

  If the columns in a table are unique, you can specify information constraints for certain columns when you create this table in the database, facilitating queries during execution.

- **Vectorized executor**

  DWS interconnects with its mature vectorized executor and improves analysis and query performance for structured ORC column data in the HDFS.

- **Partitioned tables**

  The product is adapted to the partitioned-table data defined by the Hive syntax in the HDFS. It automatically prunes partitioned tables using the DWS SQL optimizer, improving analysis and query performance.

- **Highly efficient distributed reads of HDFS data**

  The DWS SQL compiler uses the Scheduler to balance loads among DWS DNs that access HDFS data and uses the HDFS feature of short-circuit local read to improve data read performance.

  For more information, see **CREATE FOREIGN TABLE (SQL on Hadoop or OBS)**.

3.9 Cluster Management and HA

Description

GaussDB(DWS) provides the cluster manager (CM) module to manage and monitor the running status of each functional unit and physical resource in the distributed system, ensuring stable running of the entire system. CMs are classified into primary and standby CMs. In normal cases, only the primary CM provides the GaussDB(DWS) cluster management service. If the primary CM is faulty, the standby CM will be promoted to the primary to provide cluster management.

Technical Principles

The cluster management module consists of the CMServer, CMAgent, and Monitor components, and provides tools for querying cluster status, starting and stopping

a cluster, performing primary/standby switchover, and rebuilding instances. CMServer is deployed only on the primary and standby CMs. As the brain of the entire GaussDB(DWS) cluster, CMServer processes various status information reported by CMAgent and determines whether to change the status. Deployed on all nodes, CMAgent functions as an instance agent process, reports the status of CNs, DNs, GTMs, and other instances to CMServer, and receives and executes commands delivered by CMServer. Monitor is deployed on all nodes as a scheduled task that restarts CMAgent when it is stopped.

Hang Detection

The cluster management module uses the short connection mechanism to check whether the instance process is in an abnormal state such as network fault, disk I/O suspension, or process/thread suspension. If necessary, the cluster management module triggers the DN/GTM primary/standby switchover or CN removal process.

Take DNs as an example. By default, CMAgent creates instance connections every 180 seconds. If the connection fails, CMAgent retries every 84 seconds. If the connection fails for five consecutive times, the primary/standby DN switchover process is triggered. The complete hang detection period is about 600 seconds.


CN Retry

GaussDB(DWS) provides the CN Retry function to automatically retry SQL statements when an exception occurs, improving service continuity.

The CN triggers the retry mechanism when an error is reported during statement execution. For retryable errors, roll back the executed operation and execute the

statement again. If the operation still fails, the error information is reported to the client.

Users are unaware of the retry process of SQL statements.
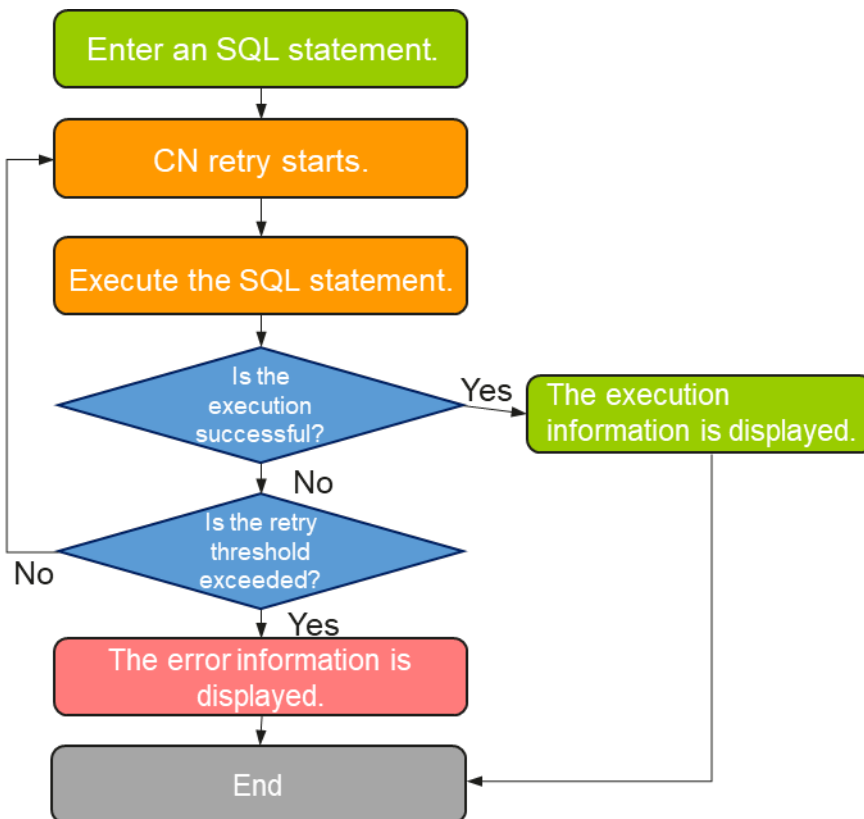

3.10 SQL Self-Diagnosis

Background

A large number of queries are involved in services. Query problems at the planning and execution phases may appear, for example, inaccurate estimation, data skews, and statistics not collected. SQL self-diagnosis provides users with a more efficient and easy-to-use method for locating performance problems. It helps users simplify the SQL optimization process of batch processing jobs. After entering SQL statements, users can easily obtain SQL problems and optimization suggestions in batches, instead of executing complex optimization such as extracting and rewriting SQL statements in the job set one by one, reproducing the SQL statements that have performance problems, checking Explain Performance to locate the problem, rewriting the SQL statements, and adjusting parameters.

Function

Before executing the job, configure the GUC parameters. For details, see sections "resource_track_level" and "resource_track_cost" in *Developer Guide*. After running a user job, you can view the related system views to obtain the possible performance problems of query jobs. These system views provide possible causes

of performance problems. Based on these causes, you can optimize job performance by referring to section "SQL Self-Diagnosis" in *Developer Guide*.

Technical Principles
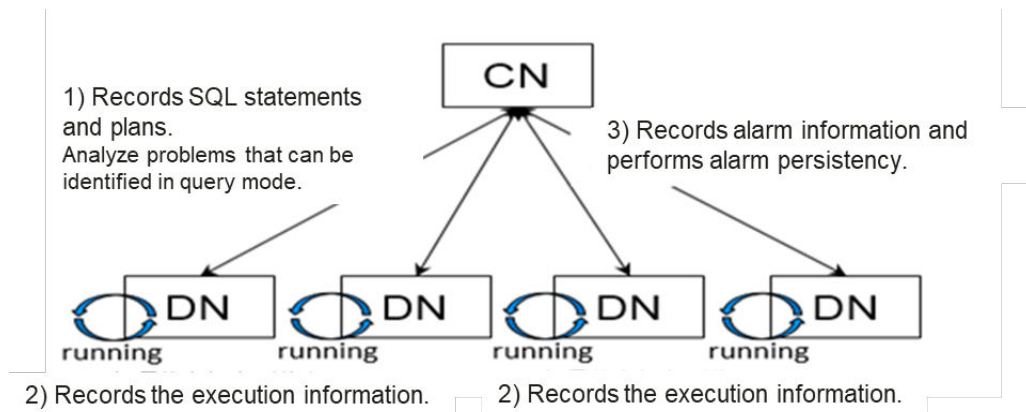
[1] The CN queries and compiles SQL statements to generate a plan tree. The diagnosis analyzer diagnoses the plan tree to identify problems in query mode (including alarms that statistics are not collected and that SQL statements cannot be delivered).

[2] DNs execute the SQL statement and writes the statistics to the shared memory of the DN. If the current SQL is running in operator mode (all supported alarm scenarios can be diagnosed), collect operator execution statistics during SQL running.

[3] DNs execute the SQL statement and returns the execution result to the CN. In the operator mode, the DN returns the collected runtime status to the CN. The diagnosis analyzer analyzes the status and records the analysis result.

[4] Statistics in the memory are cleared every 3 minutes. To record all historical information, enable the GUC parameter **enable_resource_record**. The statistics in the memory is persisted to a specific system catalog every 3 minutes.

Optimization

Scenario	Optimization
Statistics of a single column or multiple columns are not collected.	Run the analyze command to collect statistics and generate a better execution plan.
SQL statements are not pushed down.	Report the reason why the pushdown fails. Optimize SQL statements to push SQL statements down.
In a hash join, the larger table is used as the inner table.	A large amount of data is written to disks, which greatly affects the performance. Optimize SQL statements to avoid such situation.

Scenario	Optimization
nestloop is used in a large-table equivalent join.	Nestloop has a great impact on performance when the data volume is large. Optimize SQL statements to avoid such situation.
A large table is broadcasted.	A large amount of data is transmitted on the network, which greatly affects the performance. Optimize SQL statements to avoid such situation.
Data skew occurs.	Some nodes become the system bottleneck.
The index is improper.	The number of scanned rows is too large. Modify the index.
Estimation is inaccurate.	The difference between the estimated and actual number of rows is large. As a result, the selected plan is not the optimal one.

3.11 Transparent Data Encryption

Description

Transparent Data Encryption (TDE) encrypts GaussDB(DWS) data files. Generally, threat mitigation measures are taken to protect data security. For example, design a secure system, encrypt confidential assets, or build a firewall around database servers. However, in a scenario where the physical media (for example, disks) are stolen by attackers or internal personnel, the malicious party can just restore or attach the database and browse the data. One solution is to encrypt the sensitive data in the database and protect the keys that are used to encrypt the data. This prevents anyone without the keys from using the data, but this kind of protection must be planned. The GaussDB(DWS) provides a complete solution.

TDE performs real-time I/O encryption and decryption of the data. Users are unaware of the encryption. The encryption uses a database encryption key (DEK), which is not stored in the cluster. The DEK is a symmetric key secured by using the cluster encryption key (CEK) stored in the KMS server. Database servers store only DEK ciphertext. During database startup, each CN/DN uses identity authentication (for example, Kerberos) to connect to the KMS server, decrypts the DEK ciphertext to obtain the key plaintext, and caches it in the memory. Once the host is powered off or the cluster is shut down, keys are deleted. Therefore, do not lose key files in the cluster because they are irrecoverable.

Use Cases

In a traditional database cluster, user data is stored in plaintext in column-store or row-store files. Cluster maintenance personnel or malicious attackers can bypass the database permission control mechanism in the OS or steal disks to access user data. GaussDB(DWS) adopts and enhances the Hadoop KMS. The connection to third-party KMS helps GaussDB(DWS) achieve transparent data encryption for data security.

GaussDB(DWS) obtains keys from the third-party KMS through the Hadoop KMS.

Figure 3-9 Data storage encryption
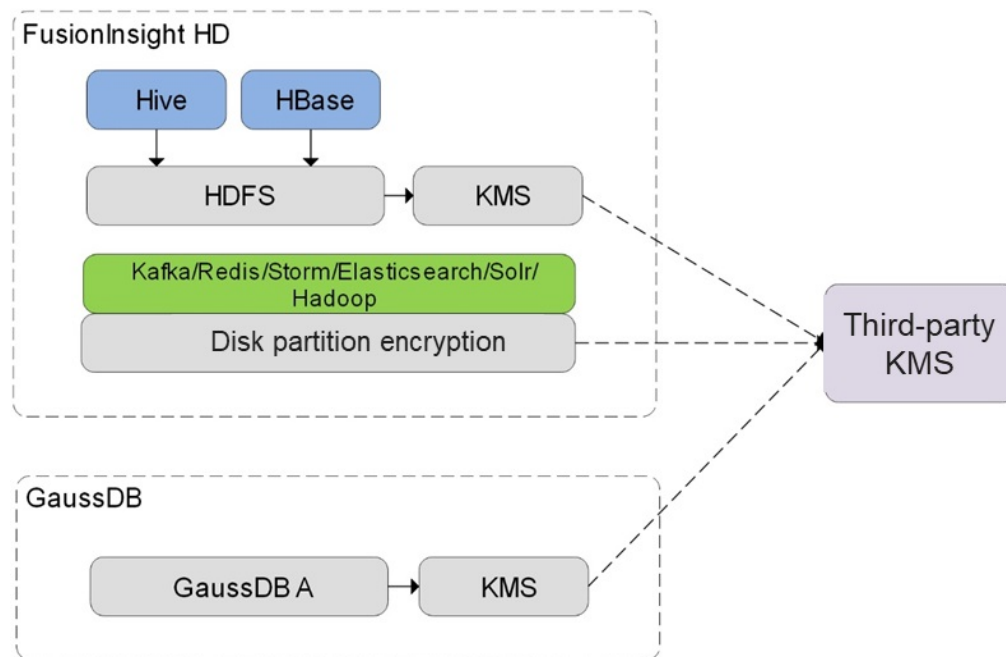
GaussDB(DWS) Transparent Encryption

In GaussDB(DWS) database-level transparent encryption, each GaussDB(DWS) cluster has a CEK, and each database is configured with a DEK. DEKs are encrypted using the CEK and their ciphertext is stored in GaussDB(DWS) clusters. Keys are applied for, encrypted, and decrypted through the KMS service. The encryption algorithm is configured using configuration items. Currently, the AES and SM4 algorithms are supported. The SM4 algorithm supports hardware acceleration in chips of Hi 1620 or later.

Currently, database-level transparent encryption is supported. You need to configure encryption when creating a cluster.

For details about how to configure transparent encryption, see **Encrypting GaussDB(DWS) Databases**.

3.12 Data Masking

Background

In the big data era, the huge value of data also brings difficulties in privacy protection. Data masking ensures efficient big data sharing and protects sensitive information.

Description

GaussDB(DWS) allows users to create data masking policies by column. Users can create policies for sensitive data in services. Sensitive data is identified by users based on their service scenarios. After a data masking policy is configured, only the administrator and table object owner can access raw data. In addition, the

data can be used for actual calculation and is masked only when the database service returns the final result.

Figure 3-10 Data masking effect

Technical Principles
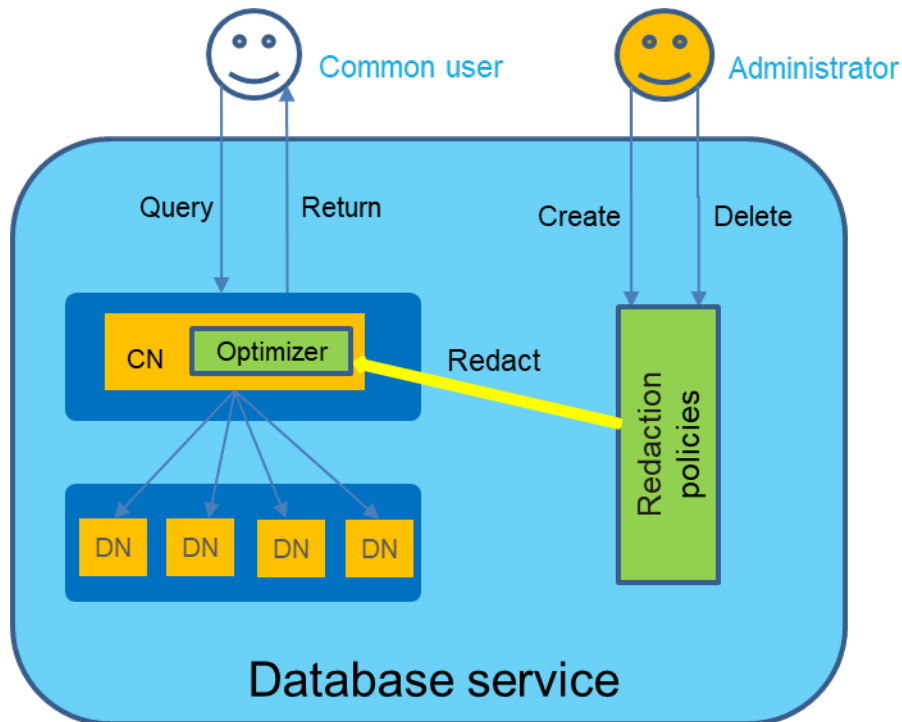
Figure 3-11 Technical principles

Table 3-4 Interface functions

External Interface	Function
add_policy	Creating a masking policy

External Interface	Function
alter_policy	Modifying a data masking policy
drop_policy	Deleting a data masking policy
enable_policy	Enabling a data masking policy
disable_policy	Disabling a data masking policy

After a data masking policy is created for a column in a table, all queries involving the column are affected by the policy. Only the administrator and table owner can query and return the original values of the column.

Benefits

Industries involving sensitive information have great requirements for data masking, such as finance, government, and healthcare. Data masking can be used to prevent sensitive information leakage in application development, testing, and training scenarios.

For details, see **Data Redaction**

3.13 Data Backup and Disaster Recovery

3.13.1 Backup and Restoration

The following figure shows the backup and restoration process of physical fine-grained backup and restoration:
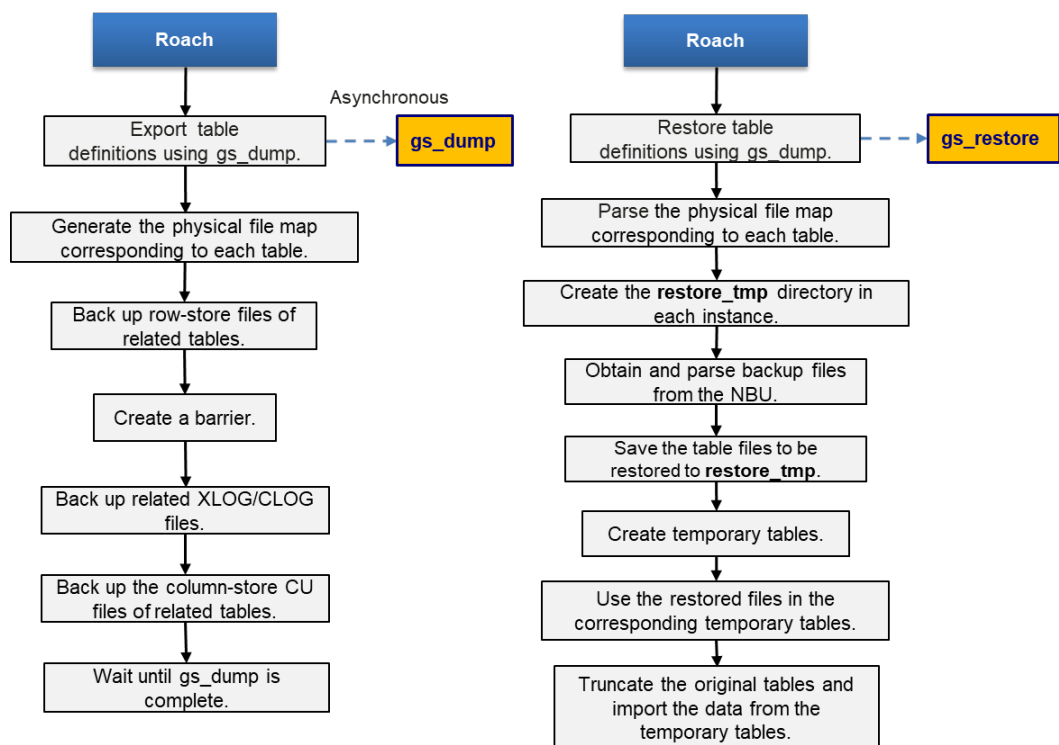
Snapshot

A snapshot is a complete backup that records point-in-time configuration data and service data of a GaussDB(DWS) cluster. A snapshot can be used to restore a cluster. Snapshots are stored on OBS.

📖 **NOTE**

- GaussDB(DWS) provides some free-of-charge storage space for storing snapshot data. However, when you use more storage space than the free-of-charge storage space, the excess space is billed on a pay-per-use basis.
- The free-of-charge space is the same as the size of the total storage space of a cluster (storage space of a single node x number of nodes).
- The snapshot management function depends on OBS.

A snapshot contains the data in databases running in a cluster and the cluster information, including the node quantity, node flavors, and administrator names. To restore a cluster from a snapshot, GaussDB(DWS) uses the cluster information to create a new cluster and then restores all databases from the snapshot. The new cluster created from the snapshot has the same configurations (including the number and flavor of nodes) as those of the original cluster. When restoring a cluster from a snapshot, the parameter values are consistent with those in the snapshot unless you specify them.

There are two types of snapshots: automated and manual.

Backup and Restoration Policies of Automated Snapshots

Automated snapshots adopt differential incremental backups. The automated snapshot created for the first time is a full backup (base version), and then the system creates full backups at a specified interval. Incremental backups are generated between two full backups. The incremental backup records change based on the previous backup. During snapshot restoration, GaussDB(DWS) uses all backups between the latest full backup and the current incremental backup to restore the cluster. Therefore, no data loss occurs. To ensure that every incremental snapshot can be used for data restoration, when its retention period exceeds the upper limit, GaussDB(DWS) does not delete the snapshot immediately. Instead, GaussDB(DWS) retains it for future cluster restoration using other incremental snapshots. GaussDB(DWS) deletes the previous full automated snapshots and related incremental snapshots only after a new full snapshot is created. If you disable the automated snapshot function for an existing cluster, all its automated snapshots will be deleted. However, manual snapshots will not be deleted.

Ecosystem Interconnection

Storage media for database backup include NetBackup, EISOO, A8000, OBS, and disk storage media. Local disk storage competes with database data for storage space on disks. Therefore, remote storage media, such as NetBackup, EISOO and A8000, are used to manage backup data. In addition, these storage services perform technical operations such as data deduplication, and provides information about backup space usage. OBS is a data storage service provided by Huawei. It is easy to use. During backup, the backup process directly sends data to OBS for storage. You do not need to set up an intermediate client on the local backup server. The mainstream backup software uses this architecture, as shown in the following figure. Each node has a backup client. When a backup task starts, the backup task is delivered to each client. Each client creates an **eefproc** process, which invokes the backup command of the database. After the backup data is generated, it is stored in the pipe. The **eefproc** process reads the data and sends it to the backup server.

For the backup and restoration ecosystem interconnection, GaussDB(DWS) uses the standard XBSA interface for backup, and backup vendors implement the XBSA protocol and use the non-intrusive Roach client provided by GaussDB(DWS) to connect to the GaussDB(DWS) backup storage service.

Figure 3-12 Backup and restoration ecosystem interconnection architecture
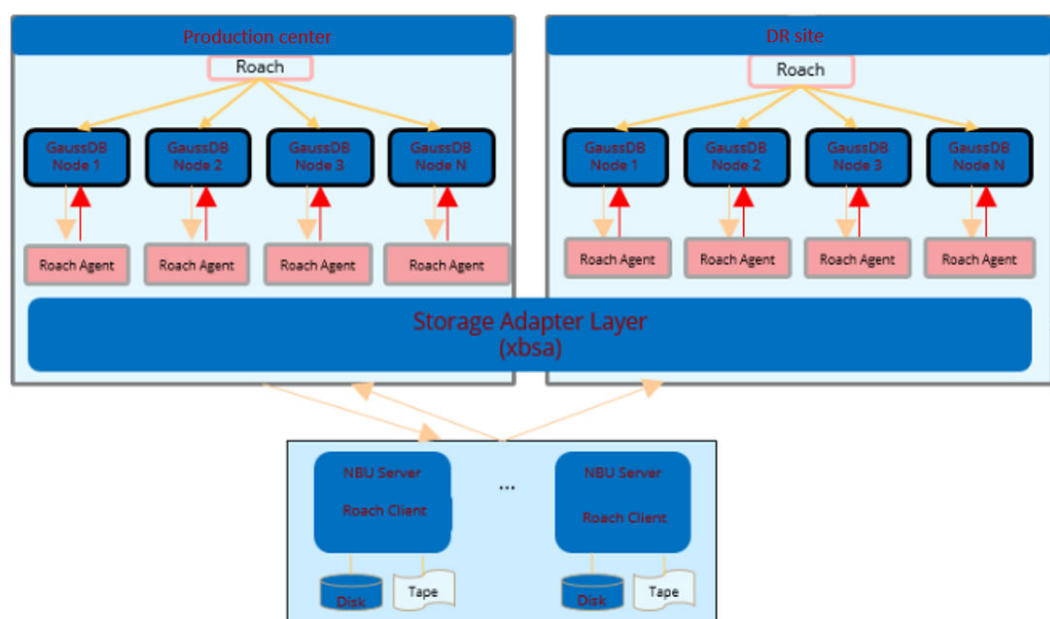
For details, see **Snapshots**.

## 3.13.2 Cluster DR

Background

GaussDB(DWS) clusters support disaster recovery. Data of the primary cluster is periodically synchronized to a standby cluster in another region or AZ. After considering several typical DR methods in the industry, such as dual-write on the application side and multiple remote copies, as well as the architecture

characteristics of GaussDB(DWS) and the replicability of the DR solution, we have chosen the dual-cluster DR.

Technical Principles

Dual-cluster DR architecture:

- Data synchronization: Node-to-node data synchronization is used between two clusters with the same number of DNs based on the MPP distribution characteristics. The network between the two clusters must be connected mutual trust must be configured.

- Periodic synchronization: Data is synchronized at a configurable interval.

- Data content: Cluster data including row-store data, column-store data, library data, and configuration information is synchronized.
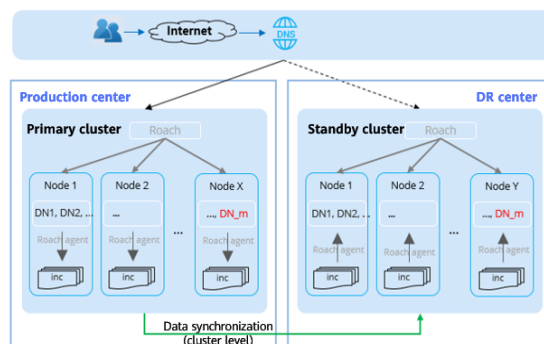
Architecture:
- Cross-DC (cross-AZ/cross-region/offline), dual-cluster (primary/standby), and asynchronous replication through Roach
- DR data: cluster-level

Application scenarios:
Intra-city and remote disaster recovery of enterprise-level data warehouses

RTO/RPO:
RTO/RPO: 30 minutes to hours, service-relevant

Constraints
- The primary and standby clusters must have the same topology and DNs.
- The network between the primary and standby clusters must be connected, and mutual trust can be configured.

Features:
- The two clusters are relatively independent. During non-recovery periods, the primary cluster can be read and written, and the standby cluster can be read.
- The DR architecture does not depend on underlying storage, services are transparent, and PB-scale data synchronization is supported.
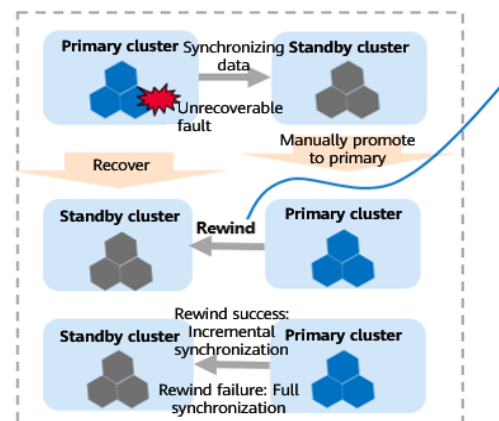- DR of row-store and column-store data is supported..

Cluster switchover:

Scenario 1: Planned switchover
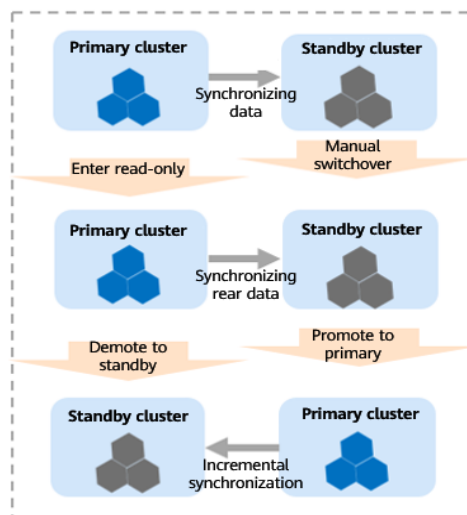
1. Used for DR drills.
2. RITO = Number of stacked backup sets x Recovery time of a single backup set + Cluster startup time;
3. RPO= 0

Scenario 2: Primary cluster failover

1. Used when unrecoverable faults occur in the primary cluster.
2. RITO = Rewind time + Number of unrestored backup sets x Recovery time of a single backup + Cluster startup time
3. RPO = Incremental data of the primary cluster that is not synchronized

There are two types of DR switchover. One is planned switchover, which is used for DR drills. The other is failover, which is performed upon cluster faults.

For a planned switchover, if the RPO is 0, the primary cluster synchronizes data to the standby cluster and becomes the standby cluster. During a failover, the standby cluster immediately becomes the primary cluster. In this case, the RPO is not 0.

Benefits

The dual-cluster HA DR is loosely coupled and is selected by vendors such as Oracle and DB2. The advantages are as follows: (a) The two clusters are independent of each other, and the standby cluster does not affect the primary cluster. (b) The dual-cluster switchover can easily support the software upgrade and application upgrade of the major versions. It helps with the data security protection of large-scale MPP clusters.

For more information, see .

4 GaussDB(DWS) Tools

4.1 Client Tools

4.1.1 gsql

Background

gsql is a database connection tool provided by GaussDB(DWS) and used by executing commands. You can use gsql to connect to a server and perform O&M on the server. In addition to the basic functions of operating on a database, gsql provides advanced features.

Function Description

- **Connecting to the database**: By default, you can connect to only the local server. To connect to a remote database, configure the server first.
- **Running SQL commands**: You can enter and run SQL commands interactively or run the SQL commands specified in a file.
- **Running meta-commands**: Meta-commands help administrators view database object information, query cache information, format SQL output, and connect to a new database.

4.1.2 Data Studio

Description

Data Studio is a GUI tool that connects to a database to help you execute and debug SQL statements and stored procedures. Data Studio supports basic features of GaussDB(DWS) and provides a GUI for database developers. This improves the efficiency of constructing application programs and simplifies database development.

Data Studio provides the following functions for database developers:

- Browsing database objects

- Creating and managing database objects, such as databases, users, tables, and indexes
- Editing and running PL/SQL statements
- Importing and exporting table data
- Debugging SQL statements and stored procedures

Technical Principles

**Figure 4-1** DataStudio component interaction
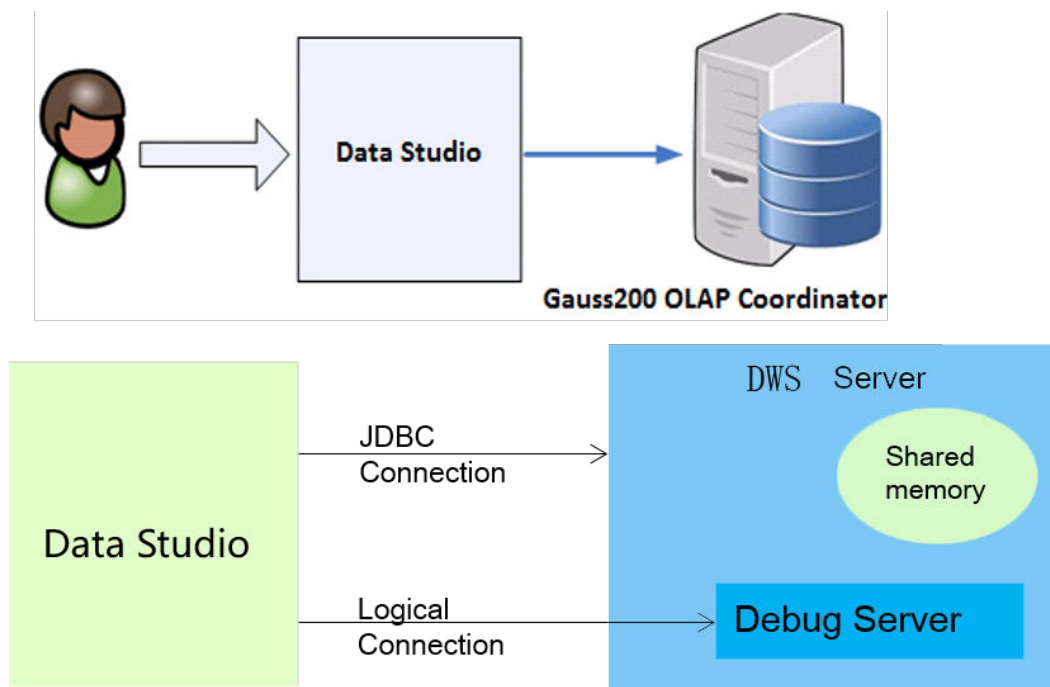
As shown in the preceding figure, Data Studio uses the C/S structure and communicates with the GaussDB(DWS) database through the JDBC driver.

During debugging, two connections are used:

- JDBC connection, which is used for query.
- Logical connection, which is used for other debugging operations, such as breakpoint and variable operations.

The database service and debugging service communicate with each other through the shared memory.

Figure 4-2 Data Studio GUI
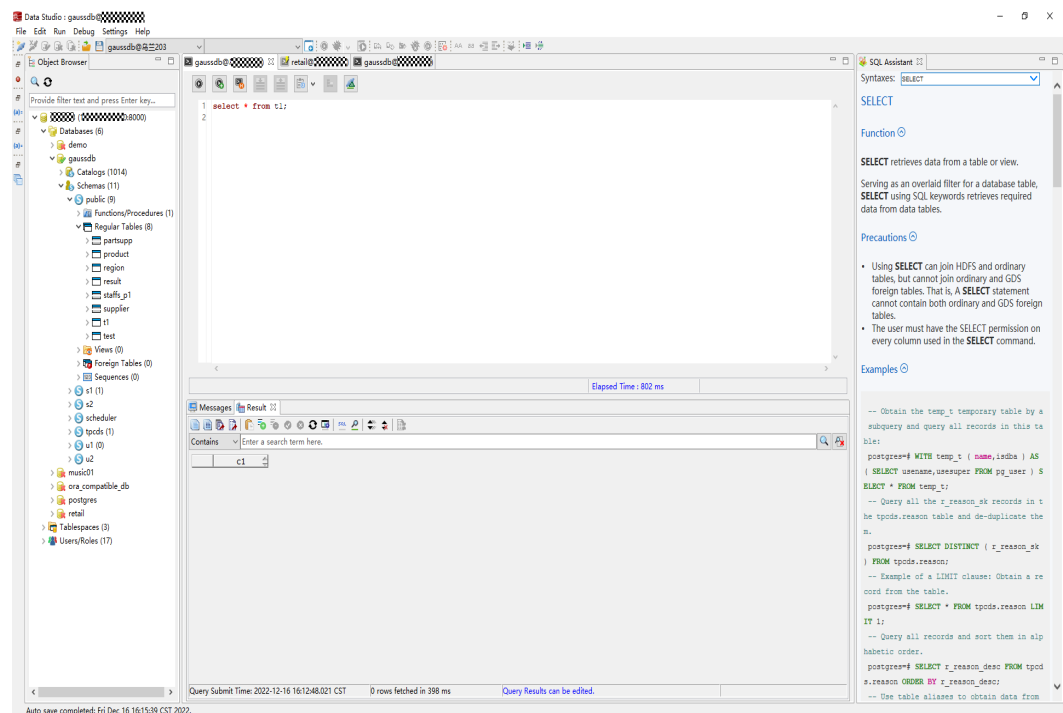
No.	Function	Description
1	Managing database objects	Manages database objects such as databases, schemas, tables, columns, constraints, indexes, views, tablespaces, and user roles.
2	SQL editor	Creates, edits, runs, and debugs PL/SQL stored procedures, formats query statements, automatically recommends SQL statements, and fills in templates.
3	Stored procedure debugging	Creates debugging links, uses breakpoints to control the execution of PL/SQL stored procedures, and displays debugging information such as call stacks and variables.
4	console	Views the execution plan and cost. Standard input and output are supported.
5	Query results	Displays, copies, exports, edits, and searches for query results.

For details about how to use Data Studio, see section **About Data Studio**.

4.1.3 Database Schema Convertor

Background

After customers switch to Huawei databases, database migration is required, which includes user data migration and application SQL script migration.

The migration of application SQL scripts is a complex, risky, and time-consuming process.

Description

Database Schema Convertor is a command-line tool running on the Linux or Windows OS. It is dedicated to providing customers with simple, fast, reliable application SQL script migration services. It parses SQL scripts into ones that compatible with GaussDB(DWS).

Database Schema Convertor does not require a connection to databases, and can perform migration offline. The tool also displays the status of a migration process and logs the errors that occur during the process, helping quickly locate faults.

Database Schema Convertor supports migration from Teradata and Oracle to GaussDB(DWS). It now supports the following objects:
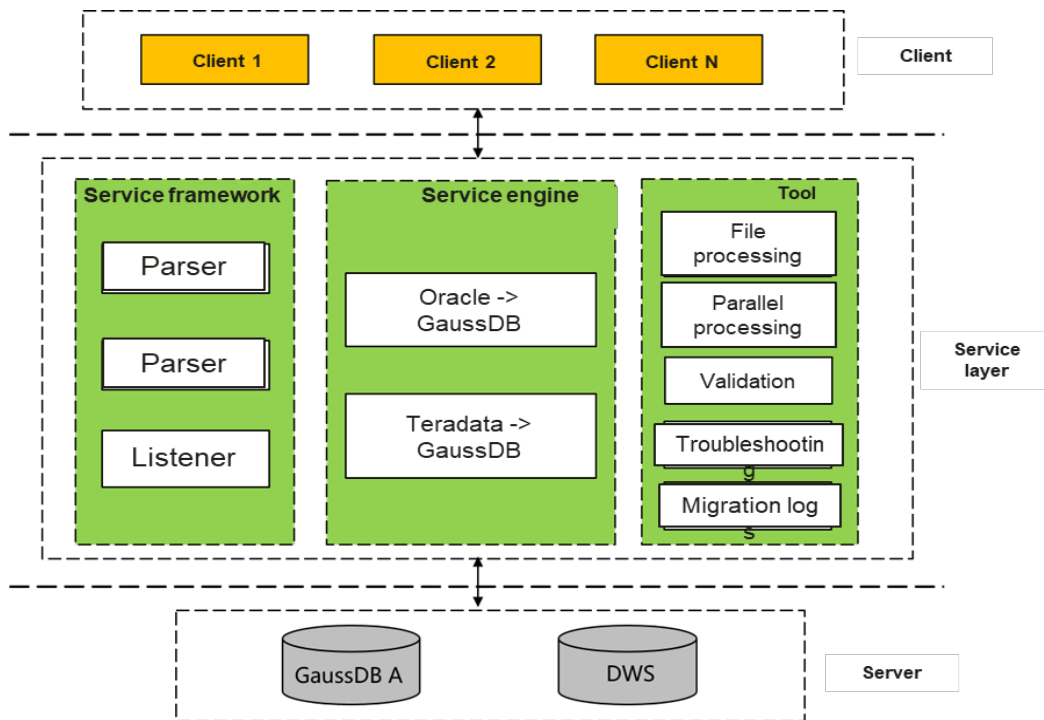
- General objects: SQL schemas and SQL queries
- Objects supported only by Oracle: PL/SQL objects
- Objects supported only by Teradata: Perl files containing **BTEQ** and **SQL_LANG** scripts

Target user groups

- Database administrators
- Database migration engineers

Technical Principles

Figure 4-3 Database Schema Convertor architecture

Database Schema Convertor supports the following types of source and target databases:

No.	Function	Description
1	Teradata Perl Migration	Migrating Teradata Perl files to GaussDB A, GaussDB 300, and GaussDB(DWS)
2	Teradata SQL Migration	Migrating Teradata SQL to GaussDB A, GaussDB 300, and GaussDB(DWS)
3	Oracle SQL Migration	Migrating Oracle SQL to GaussDB A, GaussDB 300, and GaussDB(DWS)
4	Oracle(beta) SQL Migration	Migrating Oracle SQL to GaussDB A, GaussDB 300, and GaussDB(DWS)

For details about how to use Database Schema Convertor, see the section **DSC: SQL Syntax Migration Tool**.

4.2 Compatibility with Third-Party Tools

To make the database more easy-to-use, DWS is developing its own tools and gradually interconnecting with mainstream database tools in the industry at home and abroad. There are various database tools and standards. Currently, DWS is

compatible with the following tools: BI tool, ETL tool, report tool, data mining, data replication, data integration, backup and restoration, security, O&M monitoring, cloud platform, and web middleware. For details about the tools, see documentation related to GaussDB ecosystem interconnection.

5 External APIs

Table 6-1 describes external APIs provided by the components of GaussDB(DWS).

Table 5-1 External APIs of the components

Component	Supported API Type
DWS	JDBC and ODBC
Management plane of DWS	REST/SNMP