LakeFormation

Product Introduction

Issue 01

Date 2024-01-31

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 What Is LakeFormation?

DataArts Lake Formation (LakeFormation) is an enterprise-level one-stop data lake construction service. It adopts a storage-compute decoupling architecture and provides GUIs and APIs for unified lake metadata management. It is compatible with Hive metadata models and Ranger permission models and can interconnect with MapReduce Service (MRS),Data Lake Insight (DLI), ModelArts, DataArts Studio, and GaussDB(DWS). LakeFormation helps you to easily and efficiently build data lakes and operation services, accelerating the release of service data value.

LakeFormation is a serverless service that uses underlying resources to implement cross-AZ deployment, high reliability, auto scaling, unified metadata management, association between metadata and file directories, and interconnection with multiple compute engines.

LakeFormation Architecture

Figure 1-1 shows the LakeFormation architecture.

Figure 1-1 LakeFormation service architecture

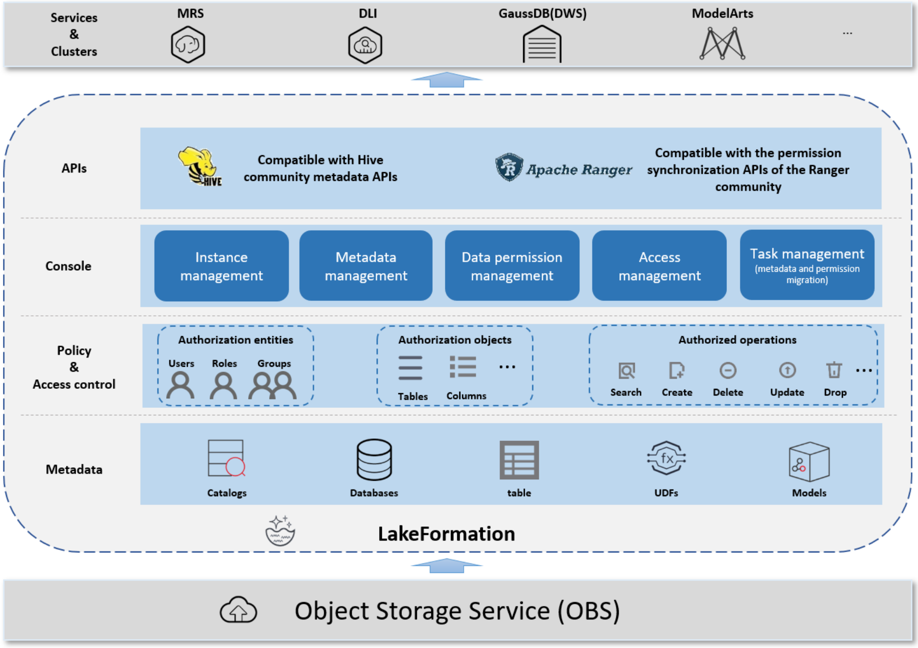

LakeFormation provides metadata management, data permission management, console, and API functions.

- The metadata management system of LakeFormation is built based on the Hive metadata model and supports metadata objects such as catalogs, databases, tables, and functions.
- LakeFormation allows you to configure permission policies and the corresponding access permission control.
  - Authorization entities include IAM users, user groups, and LakeFormation roles.
  - You can grant permissions to metadata objects such as catalogs, databases, tables, columns, functions, and OBS parallel file system paths.
  - Authorized operations include operations related to metadata objects and read and write operations on OBS paths.
- On the LakeFormation console, you can manage instances, metadata, data permissions, data access, and tasks.
- The API layer provides metadata APIs compatible with Hive and permission synchronization APIs compatible with Ranger to facilitate integration and interconnection with services such as MRS and GaussDB(DWS).

Advantages

- Open ecosystem

  LakeFormation complies with open source standards and supports seamless service evolution.
  - Smooth interconnection with Hive/Spark/Flink/Trino using metadata APIs
  - One-time authorization with Ranger
  - Smooth metadata migration with effective migration tools

- Data-AI convergence

  LakeFormation eliminates the barriers in big data and implement data-AI convergence.

  - Unified management of databases, tables, function models, and unstructured datasets
  - Secure data sharing across services and clusters based on unified fine-grained data permission management

- Large specifications (HA)

  LakeFormation supports high reliability of ultra-large-scale big data services.

  - Effective management for ultra-large-scale metadata
  - Unified fine-grained permission management at scale
  - Multi-AZ disaster recovery

- Ease of use

  LakeFormation provides metadata-based management capabilities.

  - User-friendly serverless architecture
  - Management capabilities such as data lake management and metadata statistics

Functions

Table 1-1 lists the common functions provided by LakeFormation.

Learning about the basic concepts of LakeFormation LakeFormation in advance helps you choose the optimal functions in actual situations.

Table 1-1 LakeFormation LakeFormation Functions

Function	Description
Instance Type	LakeFormation provides different types of instances to meet customers' different requirements for performance and costs in different scenarios. For details, see **Comparison Between Shared and Exclusive Instances**.
Instance management	LakeFormation provides basic functions such as instance creation, overview, and deletion, helping you easily manage instances and accelerate the planning and deployment of services carried by the data lake.
Metadata management	LakeFormation allows you to create, modify, view, and delete data lake metadata catalogs, databases, and tables. It also supports metadata life cycle configuration. Easy management helps you initialize and operate the data lake with ease, manage all LakeFormation metadata in a unified manner, and quickly plan and deploy data lake services.

Function	Description
Data permission management	LakeFormation allows you to authorize, cancel, and view data resources such as catalogs, databases, and tables. It helps you implement convenient and unified data permission management for the data lake.
Task management	LakeFormation supports full or incremental migration of metadata and permissions from external services to the current LakeFormation instance for unified management.
Access management	LakeFormation provides unified access management capabilities. You can create an access client to establish a network connection channel for a specified client environment. In addition, you can view information such as the access IP address and access domain name in the client details for other cloud services to access LakeFormation instances.

Access Methods

You can access the LakeFormation service from the web-based console or using HTTPS APIs. In addition, LakeFormation provides an SDK client to facilitate interconnection and integration with compute engines.

- Using APIs

  If you want to integrate LakeFormation instances on the public cloud platform into a third-party system for secondary development, use the APIs to access LakeFormation instances. For details about the operations, see ***API Reference***.

- Web-based console

  You can access LakeFormation on the console by selecting **Analytics** > **LakeFormation** from the service list if you have registered with the public cloud.

- Using SDKs
  - LakeFormation provides an SDK client compatible with Hive metadata models. If you need to connect compute engines such as Hive and Spark to LakeFormation for unified metadata management, you can use SDKs to access LakeFormation instances.
  - LakeFormation provides REST APIs, allowing you to call APIs using HTTPS.

2 Application Scenarios

2.1 Data Lake Construction and Continuous Operations

Scenario

To ensure fast data lake building and easy daily management of mass metadata and permissions, customers need convenient and efficient methods.

Shortcomings of Traditional Methods

- You can only execute SQL statements in compute engines (such as Hive and Spark) to define, modify, and query metadata. This requires certain skills. In addition, GUI is not supported, resulting in poor usability.
- If an authorization is required, both the compute engine and OBS need to be authorized, which is inconvenient.

Advantages of LakeFormation

- One-stop visualized lakehouse construction: A GUI is provided for unified metadata definition and authorization, ensuring fast data lake construction by simplifying operations.
- Associated authorization: During metadata authorization, the file directory mapped to the metadata can be automatically authorized, which is convenient and efficient.
- Fine-grained access control: LakeFormation implements fine-grained access control on metadata in databases, tables, and columns, ensuring transaction data security.

Recommended Services

MRS

GaussDB(DWS)

DataArts Studio

DLI

NOTE

For details, contact the corresponding service personnel.

2.2 Metadata Sharing

Scenario

Multiple services and clusters use unified metadata to maximize data sharing, avoid unnecessary duplicate data, and maximize the value of service data.

Advantages

- Being compatible with the Hive metadata model, the SDK client supports easy and fast interconnection between compute engines and LakeFormation.
- The API for querying permissions is compatible with the Ranger permission model.

Recommended Services

MRS

GaussDB(DWS)

DLI

NOTE

For details, contact the corresponding service personnel.

3 Comparison Between Shared and Exclusive Instances

LakeFormation provides different types of instances to meet customers' requirements on performance and costs in different scenarios.

NOTE

The exclusive feature is available only to whitelisted users.

- Billing

  For details about the billing items and description of shared and exclusive instances, see **Table 4-1** in **Billing**.

- Performance

Table 3-1 Performance

Item	Shared Instance	Exclusive Instance
Deployment	Physical resources are shared, but instances are logically isolated.	Physical resources are exclusively used. The performance of an instance is not affected by other instances. You can select different specifications based on your requirements.
QPS	2,000 requests can be sent to a shared instance per second.	The value fluctuates based on the QPS setting selected during the instance's creation.

- Functions

Table 3-2 Functions

Item	Description	Shared Instance	Exclusive Instance
Catalog management	Allows you to create, modify, delete, and view catalogs of data lake metadata.	√	√
Database management	Allows you to create, modify, delete, and view databases of data lake metadata.	√	√
Table management	Allows you to create, modify, delete, and view tables of data lake metadata.	√	√
Function management	Allows you to create, modify, delete, and view functions of data lake metadata.	√	√
Metadata lifecycle management	Allows you to configure data deletion policies, saving space and costs and improving system flexibility.	√	√
Metadata permission management	Allows you to authorize, remove, and view metadata.	√	√
Data migration	Supports full or incremental migration of metadata from external services to the current LakeFormation instance for unified management.	√	√
Permission migration management	Supports full or incremental migration of metadata permissions from external services to the current LakeFormation instance for unified management.	√	√
Client access management	Supports unified access management. You can create an access client to establish a network connection channel for a specified client environment so that other cloud services can access LakeFormation instances.	√	√

4 Billing

Billed Items

Huawei Cloud LakeFormation service LakeFormation are billed based on the selected specifications and service duration.

For details, see **Table 4-1**.

You can also use the **Price Calculator** provided by LakeFormation to select the required instance specifications and service duration to quickly calculate the reference price for purchasing LakeFormation instances.

Table 4-1 Billing items

Type	Item	Description
Exclusive instance	Metadata object	You are charged based on the number of metadata objects, which is the sum of the numbers of catalogs, databases, tables, partitions, indexes, and functions. The minimum unit is 10,000 metadata objects for one hour; for example, if you used 9,500 metadata objects, you are billed for 10,000 metadata objects for one hour.
	QPS	You are billed based on the QPS selected during purchase. Five QPS specifications are available and the value range is 10,000 to 50,000.
Shared instance	Metadata object	You are charged based on the number of metadata objects, which is the sum of the numbers of catalogs, databases, tables, partitions, indexes, and functions. There is no charge for the first 1 million metadata objects. After that, the usage fee is based on a minimum unit of 10,000 metadata objects per hour. If the number of users is less than 10,000, it is rounded up to 10,000.
	API call	You are charged based on the number of metadata-related API calls. The first 1 million API calls in each month are free of charge, and the subsequent API calls are charged per time.

Billing Modes

LakeFormation instances can be billed on a pay-per-use basis. For details, see **Table 4-2**.

Table 4-2 Billing modes

Billing Mode	Pay-per-use
Payment	Postpayment You are billed by the usage duration of LakeFormation instances
Billing Usage Period	Billed by the second. A bill is generated on the hour.
Billing Mode Change	Not supported.
Specification Change	Instance specifications can be changed.
Application Scenario	This mode is suitable when you want more flexibility and control on LakeFormation usage.

Expiration and Overdue Payment

If your account is in arrears, you can view the arrears details in the billing center. To prevent related resources from being stopped or released, top up your account in time. If your account is in arrears, top up your account within the specified period. For details, see **Top-Up and Repayment**.

If you do not renew or top up your account in time, your resources enter a grace period. If you still do not complete the payment or renewal after the grace period has ended, you will enter a retention period, during which the resources will be suspended. If you still do not complete the payment or renewal after the retention period has ended, your data stored in the cloud service will be deleted and the resource will be released.

5 Security

5.1 Asset Identification and Management

Asset Identification

- Asset information: metadata and data permission policy information.
- Account information: Users are unaware of the account information in LakeFormation.
- API mapping table: For details, see **API Reference**.
- Tenant resources: LakeFormation needs to read user group and user information, create and delete OBS file directories, and access OBS tag permission APIs.

Recommended Security Configuration

N/A

Infrastructure Security

- LakeFormation instances run in a cross-AZ cluster. The failure of a single AZ does not affect the running of LakeFormation instances.
- LakeFormation instances use cross-AZ highly reliable storage media to store data persistently. The failure of a single AZ does not cause data loss of LakeFormation instances.

5.2 Identity Authentication and Access Control

Identity Authentication

- IAM users of the current tenant access LakeFormation on the console.

  LakeFormation authenticates IAM tokens in HTTPS requests delivered by the console to identify tenants and IAM users. If the authentication fails, the request is rejected.

- On the console, IAM users of other tenants switch to the agency role of the current tenant to access LakeFormation.

  LakeFormation authenticates the IAM token in the HTTPS request delivered by the console to identify the delegating tenant, agency, delegated tenant, and delegated IAM user. If the authentication fails, the request is rejected.

- Instances or clusters of other cloud services (such as MRS) access LakeFormation as an agency of the current tenant.

  LakeFormation authenticates the IAM token in the HTTPS request delivered by the console to identify the delegating tenant (local tenant), agency, delegated tenant (ECS account), and delegated IAM user (built-in user of ECS). If the authentication fails, the request is rejected.

Asset Access Control

- Metadata

  When you request metadata access from the console or other cloud services, you first need to verify your identity. Then, IAM authentication checks if you have the permission to operate on the metadata in the request. Finally, fine-grained authentication further verifies your permission to operate on the specific metadata in the request. If the authentication fails, the request is rejected.

- Data permission policy

  When you request metadata access from the console or other cloud services, you first need to verify your identity. Then, IAM authentication checks if you have the operation permissions specified in the request. If the authentication fails, the request is rejected.

5.3 Data Protection

Transmission encryption (HTTPS)

To ensure data transmission security, LakeFormation APIs use HTTPS. Therefore, the console or other cloud services need to use HTTPS to access LakeFormation.

Data backup

LakeFormation supports data backup of LakeFormation instances.

5.4 Auditing

Cloud Trace Service (CTS) is a Huawei Cloud log audit service, which allows you to collect, store, and query cloud resource operation records (traces). You can use these traces to perform security analysis, track resource changes, audit compliance, and locate faults.

CTS can be used to manage permissions on LakeFormation instances and metadata.

5.5 Update Management

LakeFormation instances automatically update SSL certificates on a scheduled basis to provide continuous and stable HTTPS services.

6 Permission Management

If you need to assign different permissions to employees in your enterprise to access your LakeFormation resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permission management, and access control, helping you secure access to your Huawei Cloud resources.

With IAM, you can create IAM users for your employees, and assign permissions to these users to control their access to specific resource types. For example, if you want them to use LakeFormation but must not delete the databases or perform any high-risk operations, you can create IAM users and grant them only the permissions to query LakeFormation instances but not to delete them.

If your HUAWEI CLOUD account does not need individual IAM users for permission management, you may skip this section.

IAM is a free service. You only pay for the resources in your account. For more information about IAM, see **What Is IAM?**.

LakeFormation Service Permission

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

LakeFormation permissions are assigned to users in the global project, and users do not need to switch regions when accessing OBS.

You can grant permissions by using roles and policies.

- Roles: A coarse-grained authorization mechanism provided by IAM to define permissions based on job responsibilities. Only a limited number of service-level roles are available for authorization. If one role has a dependency role required for accessing SA, assign both roles to the users. Roles are not ideal for fine-grained authorization and least privilege access.
- Policies: A fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This type of authorization is more flexible and is ideal for secure access control. For example, you can grant users only permission to manage cloud servers of a certain type.

Table 6-1 LakeFormation fine-grained permissions

Operation Type	Item	Description
Read-only	lakeformation:instance:describe	Permission to query LakeFormation instances.
	lakeformation:catalog:describe	Permission to query the data catalogs of LakeFormation metadata.
	lakeformation:database:describe	Permission to query the databases of LakeFormation metadata.
	lakeformation:table:describe	Permission to query the data tables of LakeFormation metadata.
	lakeformation:function:describe	Permission to query the functions of LakeFormation metadata.
	lakeformation:policy:describe	Permission to query LakeFormation permission policies.
	lakeformation:policy:export	Permission to query LakeFormation permission policies in batches.
	lakeformation:agency:describe	Permission to query LakeFormation agencies.
	lakeformation:credential:describe	Permission to obtain the authentication information for accessing LakeFormation.
	lakeformation:group:describe	Permission to obtain the relationship between a LakeFormation user group and its associated roles.
	lakeformation:user:describe	Permission to obtain the relationship between a LakeFormation user and its associated roles.
	lakeformation:role:describe	Permission to query LakeFormation roles.
	lakeformation:configuration:describe	Permission to query user configurations.
	lakeformation:access:describe	Permission to query the client access permission.

Operation Type	Item	Description
	lakeformation:job:describe	Permission to query LakeFormation tasks.
Write	lakeformation:instance:create	Permission to create LakeFormation instances.
	lakeformation:role:create	Permission to create LakeFormation roles.
	lakeformation:policy:create	Permission to create LakeFormation permission policies.
	lakeformation:function:create	Permission to create the functions of LakeFormation metadata.
	lakeformation:catalog:create	Permission to create the data catalogs of LakeFormation metadata.
	lakeformation:database:create	Permission to create the databases of LakeFormation metadata.
	lakeformation:table:create	Permission to create the tables of LakeFormation metadata.
	lakeformation:access:create	Permission to create the client access permission.
	lakeformation:agency:create	Permission to create LakeFormation agencies.
	lakeformation:job:create	Permission to create LakeFormation tasks.
	lakeformation:instance:alter	Permission to modify LakeFormation instances.
	lakeformation:catalog:alter	Permission to modify the data catalogs of LakeFormation metadata.
	lakeformation:database:alter	Permission to modify the databases of LakeFormation metadata.
	lakeformation:table:alter	Permission to modify the tables of LakeFormation metadata.
	lakeformation:function:alter	Permission to modify the functions of LakeFormation metadata.

Operation Type	Item	Description
	lakeformation:role:alter	Permission to modify the relationship between a LakeFormation role and its associated user groups.
	lakeformation:group:alter	Permission to modify the relationship between a LakeFormation user group and its associated roles.
	lakeformation:user:alter	Permission to modify the relationship between a LakeFormation user and its associated roles.
	lakeformation:job:alter	Permission to modify LakeFormation tasks.
	lakeformation:instance:drop	Permission to delete LakeFormation instances.
	lakeformation:role:drop	Permission to delete LakeFormation roles.
	lakeformation:policy:drop	Permission to delete LakeFormation permission policies.
	lakeformation:function:drop	Permission to delete the functions of LakeFormation metadata.
	lakeformation:catalog:drop	Permission to delete the data catalogs of LakeFormation metadata.
	lakeformation:database:drop	Permission to delete the databases of LakeFormation metadata.
	lakeformation:table:drop	Permission to delete the tables of LakeFormation metadata.
	lakeformation:access:delete	Permission to delete the client access permission.
	lakeformation:agency:drop	Permission to delete LakeFormation agencies.
	lakeformation:job:drop	Permission to delete LakeFormation tasks.
	lakeformation:transaction:operate	Permission to operate LakeFormation transactions.

Operation Type	Item	Description
	lakeformation:instance:access	Permission to query a LakeFormation instance or apply for the access to it.
	lakeformation:job:exec	Permission to execute LakeFormation tasks.

Table 6-2 LakeFormation system permissions

Role/Policy Name	Description	Type	Dependency
LakeFormation FullAccess	Administrator permissions for LakeFormation. Users granted these permissions can use all LakeFormation functions.	System policy	N/A
LakeFormation ReadOnlyAccess	Read-only permissions for LakeFormation. Users granted these permissions can query LakeFormation data.	System policy	N/A
LakeFormation CommonAccess	Basic permissions for LakeFormation, including viewing, authorizing, and canceling the LakeFormation service agreement and basic permissions for dependent services such as OBS and TMS.	System policy	N/A

7 Constraints and Limitations

- After a IAM user group is deleted, you need to manually delete the related permission policies in LakeFormation data permissions.
- The path selected during database creation cannot be the parent path or the same path of the catalog where the database is located, or the parent path, subpath, or the same path of other databases (except the default database) in the same catalog.
- The storage location of the created database must be under that of the catalog to which the database belongs.
- Authorization and fine-grained permission control are not supported for catalog objects and their sub-metadata objects created by users.
- LakeFormation data permissions can be granted to a maximum of 20 entities or 10 metadata objects at a time.
- The number of partitions cannot exceed 1,000,000,000.
- LakeFormation does not support unified management of metadata and permissions across regions.
- LakeFormation does not support unified management of metadata and permissions across instances.
- In a data table, the combination of partition value corresponding to each partition must be unique.
- A partition name consists of partition keys and partition values and its total length cannot exceed 1,000 characters.
- In the parameter description of metadata, one Chinese character contains three bytes.
- LakeFormation needs the parallel file system of OBS. OBS nodes should be separately deployed based on the storage-compute decoupling architecture. The storage location of LakeFormation metadata corresponds to the OBS path and is interconnected with big data clusters such as MRS, which also adopts the storage-compute decoupling architecture. The OBS parallel file system must support the AccessLabel feature.
- In LakeFormation, roles sharing identical names across different instances are associated with the same OBS AccessLabel during the authorization process. It is recommended to avoid creating roles with duplicate names in separate instances within the same region to prevent conflicts.

8 Product Lifecycle

A lifecycle indicates the LakeFormation instance statuses recorded from the time when the node is created through the time when the node is deleted or released. For details about LakeFormation statuses, see **Table 8-1**.

Table 8-1 LakeFormation statuses

Status	Description
Preparing resources	After a LakeFormation instance is created, resources are being prepared for the instance.
Resource preparation failed	After a LakeFormation instance is created, resources fail to be prepared for the instance.
Running	The LakeFormation instance is running properly. Only instances in the **Running** state can provide services.
Releasing resources	Resources are being released after the LakeFormation instance is deleted.
Deleting	The instance is being deleted.
Deleted	The LakeFormation instance has been deleted.
Restoring	The deleted instance is being restored from the recycle bin.
Frozen	If your account is in arrears or violates regulations, LakeFormation instances will be frozen. The DB instance is in the read-only state and cannot be modified or deleted.

9 Relationship with Other Cloud Services

The following table describes the relationships between LakeFormation and other services.

Table 9-1 Relationships with other services

Service Name	Relationships
Identity and Access Management (IAM)	IAM authenticates IAM users or agencies and controls some access.
Cloud Trace Service (CTS)	CTS records LakeFormation operations for query, auditing, or backtracking.
Object Storage Service (OBS)	The actual service data mapped by LakeFormation metadata is stored in the directories and files of the OBS parallel file system.
MapReduce Service (MRS)	LakeFormation interconnects with Ranger, Hive, and Spark in MRS clusters to implement unified management of lake and warehouse metadata.
GaussDB(DWS)	LakeFormation interconnects with GaussDB(DWS) to implement unified management of lake and warehouse metadata.

10 Basic Concepts

10.1 Metadata

Data Catalogs

A top-level resource in the metadata resources of a LakeFormation instance and multiple catalogs can be created in a LakeFormation instance. Metadata information such as name, description, and location are included in catalogs. Catalogs can be created, modified, and deleted.

Location indicates the file directory of the OBS parallel file system mapped to the catalog.

Databases

Databases are stored in the data catalogs of a LakeFormation instance and multiple databases can be created under a catalog. Metadata information such as name, description, and location are included in databases. You can create, modify, and delete databases, as well as grant and check databases permissions.

Location indicates the file directory of the OBS parallel file system mapped to the databases.

Tables

You can create multiple tables in a database. Metadata such as basic information, format and serialization information, fields, and attributes are included in tables. You can create, modify, and delete tables, as well as grant and check permissions.

Functions

Functions are used to perform specific processing on data in SQL queries, including built-in functions and user-defined functions (UDFs).

User-defined functions are classified into the following types:

- Common UDFs: used to perform operations on a single data row and export a single data row.

- User-defined aggregating functions (UDAFs): used to input multiple data rows and export a single data row.
- User-defined table-generating functions (UDTFs): used to perform operations on a single data row and export multiple data rows.

Partitions

Partitioning is to split a data table by row to reduce the total amount of data read and write operations in specific SQL operations, and therefore shortening the response time.

10.2 Data Permissions

Permissions Policies

On the **Instances** page of the LakeFormation console, you can grant fine-grained data access permissions to user groups for all data resources such as catalogs, databases, and tables in an instance.

After the preceding authorization operations, one or more permission policies are generated.

A permission policy contains the authorization entity, authorization object, permissions, and authorization permissions. You can cancel a permission policy.

Authorization Entities

You can specify any user, user group, or role to be the authorization entity.

You can select **GROUP**, **ROLE**, and **USER** in the **Entity Type**.

- **USER**: Huawei Cloud IAM user
- **GROUP**: Huawei Cloud IAM user group
- **ROLE**: LakeFormation role

Authorization Objects

Metadata objects managed in LakeFormation, including data resources such as catalogs, databases, and tables. For instance, you can authorize permissions on the columns of a database a data table. The values of **Resource Type** include **CATALOG**, **DATABASE**, **TABLE**, **COLUMN**, and **FUNC**.

- **CATALOG**: A data catalog stores multiple databases.
- **DATABASE**: A database contains multiple data tables or functions.
- **TABLE**: A data table contains multiple columns.
- **COLUMN**: Columns in a LakeFormation table.
- **FUNC**: Functions managed by LakeFormation.

Permissions

You can grant different access and operation permissions on a data resource to an authorization entity, such as **ALTER**, **DROP**, and **ALL**. The permissions that can be granted to each authorization object are as follows:

- **CATALOG**: **ALL**, **ALTER**, **CREATE_DATABASE**, and **DROP**
- **DATABASE**: **ALL**, **ALTER**, **DROP**, **DESCRIBE**, **LIST_TABLE**, **LIST_FUNC**, **CREATE_TABLE**, **CREATE_FUNC**
- **TABLE**: **ALL**, **ALTER**, **DROP**, **DESCRIBE**, **UPDATE**, **INSERT**, **SELECT**, and **DELETE**
- **COLUMN**: **SELECT**
- **FUNC**: **ALL**, **ALTER**, **DROP**, **DESCRIBE**, and **EXEC**

Authorization Permission

You can select **Grant Authorization Permission** to enable a user to grant the permissions that he or she has to others.

10.3 Regions and AZs

A region or an availability zone (AZ) identifies the location of a data center. You can create resources in a specific region or an AZ.

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. After a resource is created, its region cannot be changed.
- An AZ is a physical location with independent power supplies and network in a region. A region can contain multiple AZs, which are physically isolated but interconnected through internal networks. This ensures the independence of AZs and provides low-cost and low-latency network connections.

Figure 10-1 Regions and AZs

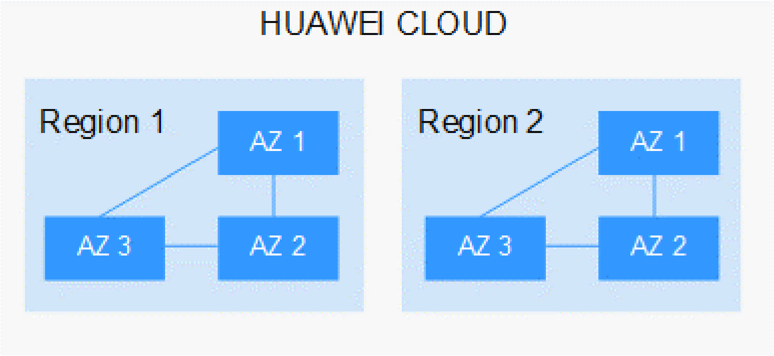

Huawei Cloud provides services in many regions worldwide. You can select a region and AZ as required. For more information, see **Global Products and Services**.

11 Change History

Released On	Change History
2024-02-01	This issue is the first official release.