

# Distributed Database Middleware

## Service Overview

**Issue** 01  
**Date** 2022-09-30



**Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei 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

---

<b>1 Overview.....</b>	<b>1</b>
<b>2 Basic Concepts.....</b>	<b>3</b>
<b>3 Core Functions.....</b>	<b>5</b>
<b>4 Product Specifications.....</b>	<b>7</b>
<b>5 Usage Constraints.....</b>	<b>8</b>
5.1 Network Access.....	8
5.2 Data Nodes.....	8
5.3 Unavailable Features and Limitations.....	9
5.4 High-risk Operations.....	15
<b>6 Regions and AZs.....</b>	<b>16</b>
<b>7 Application Scenarios.....</b>	<b>18</b>
<b>8 Permissions Management.....</b>	<b>19</b>
<b>A Change History.....</b>	<b>27</b>

# 1 Overview

---

## Definition

Distributed Database Middleware (DDM) is a MySQL-compatible, distributed middleware service designed for relational databases. It can resolve distributed scaling issues to break through capacity and performance bottlenecks of databases, helping handle highly concurrent access to massive volumes of data.

DDM uses a decoupled compute and storage architecture developed by Huawei Cloud and provides functions such as database and table sharding, read/write splitting, elastic scaling, and sustainable O&M. Management of instance nodes has no impacts on your workloads. You can perform O&M on your databases and read and write data from and to them on the DDM console, just like as operating a single-node MySQL database.

## Advantages

- **Automatic Database and Table Sharding**

MySQL databases are usually deployed on single nodes. Once a fault occurs, all data may be lost, and your workloads are 100% affected.

DDM supports automatic database and table sharding to distribute data across multiple data nodes, so impacts on your services are greatly reduced once a fault occurs. It also supports explosive growth of services.
- **Read/Write Splitting**

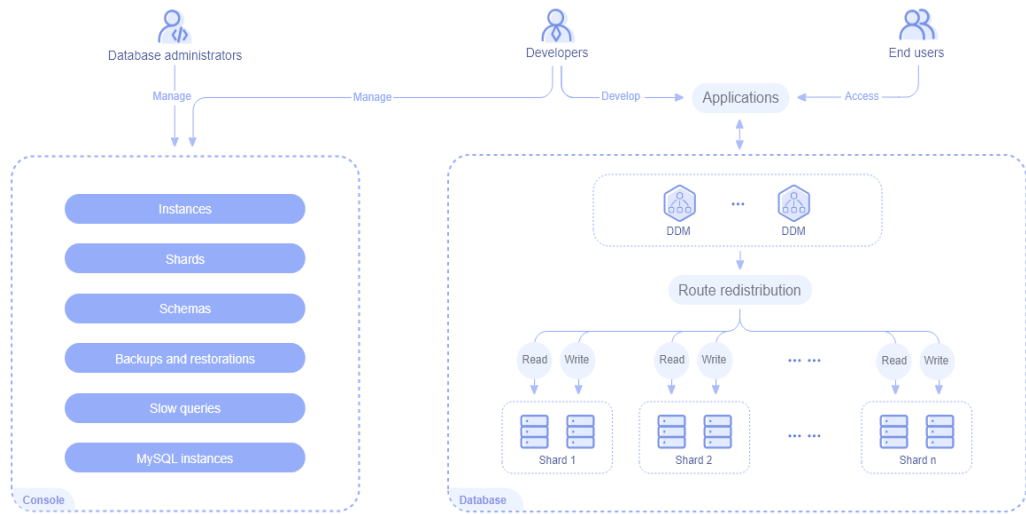
DDM can leverage data nodes. If there is still great query pressure after horizontal sharding, you can enable read/write splitting to speed up database processing and access, without the need to reconstruct your service system.
- **Elastic Scaling**

MySQL databases can support only medium- and small-scale service systems because their CPU, memory, and network processing are limited by server configurations and their storage depends on the size of SSD or EVS disks.

DDM supports both compute and storage scaling. You can add nodes to a DDM instance or scale up its node class. Alternatively, increase shards or data nodes to distribute data from one large table to multiple tables or scale out storage resources as services grow, without worrying about O&M.

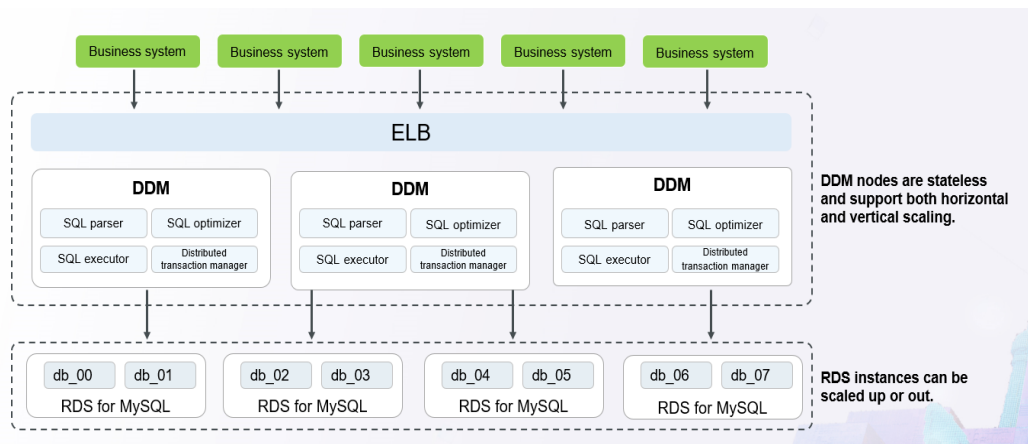
## Service Architecture

Figure 1-1 DDM service architecture



## How DDM Works

Figure 1-2 DDM working diagram



# 2 Basic Concepts

---

## Data Node

A data node is the minimum management unit of DDM. Each data node represents an independently running database, and it may be an RDS for MySQL or GaussDB(for MySQL) instance that is associated with your DDM instance. You can create multiple schemas in a DDM instance to manage data nodes and access each data node independently.

### NOTE

DDM instances do not store service-related data, which is stored in shards of data nodes.

## VPC

A Virtual Private Cloud (VPC) is a private and isolated virtual network. You can configure IP address ranges, subnets, and security groups, assign EIPs, and allocate bandwidth for DDM instances.

## Subnet

A subnet is a range of IP addresses, a logical subdivision of an IP network. Subnets are created for a VPC where you will place your DDM instances. Every subnet is defined by a unique CIDR block which cannot be modified once the subnet is created.

## Security Group

A security group is a collection of rules for ECSs that have the same security protection requirements and are mutually trusted. After a security group is created, you can add different access rules to the security group, and these rules will apply to all ECSs added to this security group.

Your account automatically comes with a security group by default. The default security group allows all outbound traffic and denies all inbound traffic. Your ECSs in this security group can communicate with each other without the need to add rules.

## Parameter Template

A parameter template acts as a container for configuration values that can be applied to one or more DDM instances. If you want to use your own parameter template, you only need to create a custom parameter template and select it when creating a DDM instance. You can also apply the parameter template to an existing DDM instance.

## EIP

The Elastic IP (EIP) service provides independent public IP addresses and bandwidth for Internet access. EIPs can be bound to and unbound from DDM instances.

## Region and Endpoint

Before using an API to call resources, you need to specify its region and endpoint. For more details, see [Regions and Endpoints](#).

# 3 Core Functions

DDM supports a variety of functions, including horizontal sharding, shard configuration, highly compatible SQL syntax, read/write splitting, and global sequence.

**Table 3-1** DDM main functions

Function	Description
Horizontal sharding	Select a sharding key when creating a logical table. DDM will generate a sharding rule and horizontally shard data.
Flexible shard configuration	DDM supports both compute and storage scaling. You can add nodes to a DDM instance or scale up its node class. Alternatively, increase shards or data nodes to distribute data from one large table to multiple tables or scale out storage resources. Compute scaling is undetectable to your applications. Storage scaling minimizes service interruption to seconds.
Distributed transactions	DDM processes three types of transactions, including single-shard, FREE, and Extended Architecture (XA). <ul style="list-style-type: none"><li>• Single-shard: Transactions cannot be committed across shards.</li><li>• FREE: Transactions are committed across shards. A transaction is not rolled back when it fails to be executed by any shard, causing data inconsistency.</li><li>• XA: Transactions are committed in two phases. If a transaction fails to be executed by any shard, all work done will be rolled back to ensure data consistency.</li></ul>
Data import and export	External data can be imported into DDM instances to help you migrate databases to the cloud, and DDM instance data can be exported based on service requirements.



Function	Description
Highly compatible SQL syntax	DDM is highly compatible with the MySQL licenses and syntax.
Read and write splitting	Read and write requests can be split without modifying the application code, and this is totally transparent to applications. You only need to create read replicas for a MySQL instance associated with your DDM instance and configure a read policy, and a large number of concurrent requests can read data from those read replicas.
Global sequence	DDM allows you to use globally unique, distributed, and ascending SNs as primary or unique keys or to meet your requirements in specific scenarios.
Online monitoring	DDM monitors reads, writes, and slow query logs of your DDM instance on the console, helping you detect resource and performance bottlenecks as fast as possible.
DDM console	The DDM console enables you to manage and maintain DDM instances, schemas, and accounts.

# 4 Product Specifications

There are two types of DDM instances: general-enhanced and Kunpeng general computing-plus.

- General-enhanced DDM instances use Intel® Xeon® Scalable processors. Working in high-performance networks, these DDM instances offer high and stable computing performance, meeting enterprise-class application requirements.
- Kunpeng general computing-plus DDM instances use Kunpeng 920 processors and 25GE high-speed intelligent NICs to offer powerful computing and high-performance networks. These instances provide governments and Internet enterprises with cost-effective, secure, and reliable cloud services.

**Table 4-1** Supported specifications

Specification	Architecture	vCPUs	Memory (GB)
General-enhanced	x86	8	16
		16	32
		32	64
Kunpeng general computing-plus	Arm	8	16
		16	32
		32	64

# 5 Usage Constraints

## 5.1 Network Access

Restrictions on network access are as follows:

- The data nodes and ECSs running your applications must be in the same VPC as your DDM instance.
- To access DDM from your computer, you need to bind an EIP to your DDM instance and then use the EIP to access the DDM instance.

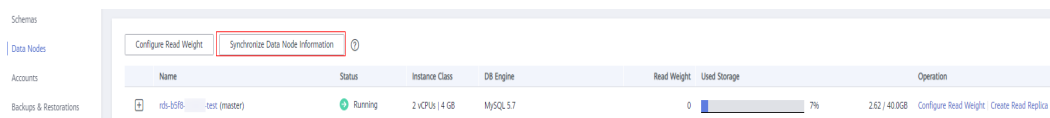
## 5.2 Data Nodes

Constraints on data nodes are as follows:

- Data nodes can be only RDS for MySQL and GaussDB(for MySQL) instances of versions 5.7 and 8.0.
- DDM cannot connect to MySQL instances using SSL connections.
- Case sensitivity support cannot be enabled for MySQL instances.

### NOTE

- If you are using MySQL 5.7, select **Case insensitive** for **Table Name** when you create a MySQL instance, or set **lower\_case\_table\_names** to **1** on the **Parameters** page after you complete the creation.
- If you are using MySQL 8.0, select **Case insensitive** for **Table Name** when you create a MySQL instance.
- Modifying configurations of a data node may result in an exception in using your DDM instance. After the modification, click **Synchronize Data Node Information** on the **Data Nodes** page to synchronize changes from the data node to DDM.



- Character set GBK is not allowed for data nodes.

## 5.3 Unavailable Features and Limitations

### Unsupported Features

- Stored procedures
- Triggers
- Views
- Events
- User-defined functions
- Foreign key reference and association
- Full-text indexes and SPACE functions
- Temporary tables
- Compound statements such as BEGIN...END, LOOP...END LOOP, REPEAT...UNTIL...END REPEAT, and WHILE...DO...END WHILE
- Process control statements such as IF and WHILE
- RESET and FLUSH statements
- BINLOG statement
- HANDLER statement
- INSTALL and UNINSTALL PLUGIN statements
- Character sets other than ASCII, Latin1, binary, utf8, and utf8mb4
- SYS schema
- Optimizer Trace
- X-Protocol
- CHECKSUM TABLE syntax
- Table maintenance statements, including ANALYZE, CHECK, CHECKSUM, OPTIMIZE, and REPAIR TABLE
- Statements for assigning a value to or querying variable **session**, for example, `set @rowid=0;select @rowid:=@rowid+1,id from user`
- SQL statements that use `--` or `/.../` to comment out a single line or multiple lines of code
- DDM provides incomplete support for system variable queries. The returned values are variable values of RDS instances, instead of DDM kernel variable values. For example, the returned values of `SELECT @@autocommit` do not indicate the current transaction status.
- Executing SET Syntax to modify global variables
- PARTITION syntax. Partitioned tables are not recommended.
- LOAD XML statement

### Unsupported Operators

- Assignment operator (`:=`) is not supported. After this operator is expected, it does not take effect even if no errors are reported.

- Operator (->) is not supported. This operator can be executed successfully in a single table. An error is reported when this operator is executed in other types of tables.
- Operator (->>) is not supported. This operator can be executed successfully in a single table. An error is reported when this operator is executed in other types of tables.
- Expression IS UNKNOWN

## Unsupported Functions

Do not use a function if you are not sure whether it can be pushed down to RDS. The compute layer of DDM does not support the following functions:

- XML functions
- Function **ANY\_VALUE()**
- Function **ROW\_COUNT()**
- Function **COMPRESS()**
- Function **SHA()**
- Function **SHA1()**
- Function **MD5()**
- Function **AES\_ENCRYPT()**
- Function **AES\_DECRYPT()**
- Aggregate function **JSON\_OBJECTAGG()**
- Aggregate function **JSON\_ARRAYAGG()**
- Aggregate function **STD()**
- Aggregate function **STDDEV()**
- Aggregate function **STDDEV\_POP()**
- Aggregate function **STDDEV\_SAMP()**
- Aggregate function **VAR\_POP()**
- Aggregate function **VAR\_SAMP()**
- Aggregate function **VARIANCE()**

## SQL Syntaxes

### SELECT

- DISTINCTROW
- Configuring options [HIGH\_PRIORITY], [STRAIGHT\_JOIN], [SQL\_SMALL\_RESULT], [SQL\_BIG\_RESULT], [SQL\_BUFFER\_RESULT], and [SQL\_NO\_CACHE] [SQL\_CALC\_FOUND\_ROWS] in SELECT statements on DDM instances
- SELECT ... GROUP BY ... WITH ROLLUP
- SELECT ... ORDER BY ... WITH ROLLUP
- WITH
- Window functions

- SELECT FOR UPDATE supports only simple queries and does not support statements such as JOIN, GROUP BY, ORDER BY, and LIMIT. Option [NOWAIT | SKIP LOCKED] for modifying FOR UPDATE is invalid for DDM.
- DDM does not support multiple columns with the same name for each SELECT statement in UNION. Duplicate column names are used in the following SELECT statement:  

```
SELECT id, id, name FROM t1 UNION SELECT pk, pk, name FROM t2;
```

### **SORT and LIMIT**

- LIMIT/OFFSET, value range: 0–2147483647

### **Aggregation**

- Function **asc/desc** cannot be used in the GROUP BY statement to sort out results.

#### **NOTE**

DDM automatically ignores keyword **asc/desc** after GROUP BY. In MySQL versions earlier than 8.0.13, function **asc/desc** can be used in the GROUP BY statement to sort out results. In MySQL 8.0.13 or later, a syntax error is reported if you use function **asc/desc** this way. ORDER BY is recommended for sorting.

### **Subqueries**

- Subqueries that join grandparent queries are not supported.
- Using subqueries in the HAVING clause and the JOIN ON condition is not supported.
- Each derived table must have an alias.
- A derived table cannot be a correlated subquery.

### **LOAD DATA**

- LOW\_PRIORITY is not supported.
- CONCURRENT is not supported.
- PARTITION (partition\_name [, partition\_name] ...) is not supported.
- LINES STARTING BY 'string' is not supported.
- User-defined variables are not supported.
- ESCAPED BY supports only '\\\'.
  - If you have not specified a value for your auto-increment key when you insert a data record, DDM will not fill a value for the key. The auto-increment keys of data nodes of a DDM instance all take effect, so the auto-increment key values may be duplicate.
- If the primary key or unique index is not routed to the same physical table, REPLACE does not take effect.
- If the primary key or unique index is not routed to the same physical table, IGNORE does not take effect.

### **INSERT and REPLACE**

- INSERT DELAYED is not supported.
- Only INSERT statements that contain sharding fields are supported.

- PARTITION syntax is not supported. Partitioned tables are not recommended.
- Setting **datetime** to **1582** or any value smaller in INSERT statements is not supported.
- Nesting a subquery in ON DUPLICATE KEY UPDATE of an INSERT statement is not supported. The following is an example:

```
INSERT INTO t1(a, b)
SELECT * FROM(SELECT c, d FROM t2 UNION SELECT e, f FROM t3) AS dt
ON DUPLICATE KEY UPDATE b = b + c;
```

Subquery c is used in the ON DUPLICATE KEY UPDATE clause.
- The sharding key values in INSERT and REPLACE statements cannot be **DEFAULT**.

### UPDATE and DELETE

- Updating a sharding key value to **DEFAULT** is not supported.
- Repeatedly updating the same field in one SQL statement is not supported.
- Updating a sharding key using UPDATE JOIN is not supported. The following is an example:

```
UPDATE tbl_1 a,tbl_2 b set a.name=b.name where a.id=b.id;
```

**name** indicates the sharding key of table **tbl\_1**.
- Updating a sharding key by executing INSERT ON DUPLICATE KEY UPDATE is not supported.
- Updating self-joins is not supported.

```
UPDATE tbl_1 a,tbl_1 b set a.tinyblob_col=concat(b.tinyblob_col,'aaabbb');
```
- UPDATE JOIN supports only joins with WHERE conditions. The following is an example:

```
UPDATE tbl_3,tbl_4 SET tbl_3.varchar_col='dsgfdg';
```
- Referencing other object columns in assignment statements or expressions is not supported when UPDATE JOIN syntax is used. The following is an example:

```
UPDATE tbl_1 a,tbl_2 b SET a.name=concat(b.name,'aaaa'),b.name=concat(a.name,'bbbb') ON
a.id=b.id;
```
- You can update a sharding field by two steps: delete the original sharding field and then insert a new field. During this process, the results of querying the sharding fields involved in the target table may be inconsistent.

### DDL

- SQL statements for modifying database names and sharding field names and types
- SQL statements for creating and deleting schemas
- Index FULL\_TEXT
- AS SELECT clause of the CREATE TABLE statement
- CREATE TABLE ... LIKE statement
- Dropping multiple tables with one SQL statement
- Executing multiple SQL statements at the same time
- Creating foreign keys for broadcast and sharded tables
- Creating tables whose names are prefixed by **\_ddm**
- Creating temporary sharded or broadcast tables

- Specifying globally unique keys in the CREATE TABLE statement

## Indexes

- Global secondary indexes
- Global unique indexes. Unique keys and primary keys may not be globally unique.

## Table Recycle Bins

- Hints
- Deleting tables by schema
- Deleting tables by logical table
- After a table is recovered, its globally unique sequence increases automatically but may not follow the last sequence value.
- Shard configuration
- Retaining copies with no time limit
- Recovering data to a table with any name
- Unlimited copies

## Transactions

- Savepoints
- XA syntax. DDM has implemented distributed transactions through XA, so the user layer does not need to process the syntax.
- Customizing the isolation level of a transaction. Currently, DDM supports only the READ COMMITTED isolation level. In consideration of compatibility, DDM does not report errors for any SQL statement (such as SET GLOBAL TRANSACTION ISOLATION LEVEL REPEATABLE READ) to set the database isolation level, but will ignore the modifications to the transaction isolation level.
- Setting a transaction to read-only (START TRANSACTION READ ONLY). DDM can enable read/write of a transaction, instead of enabling read-only, to ensure compatibility.

## Permissions

- Column-level permissions
- Subprogram-level permissions

## Database Management Statements

- SHOW TRIGGERS
- Most of SHOW statements such as SHOW PROFILES, SHOW ERRORS, and SHOW WARNINGS
- The following SHOW statements are randomly sent to a database shard. If database shards are on different RDS for MySQL instances, the returned variables or table information may be different.
  - SHOW TABLE STATUS;



- SHOW VARIABLES Syntax;
- SHOW WARNINGS Syntax does not support the combination of LIMIT and COUNT.
- SHOW ERRORS Syntax does not support the combination of LIMIT and COUNT.

## INFORMATION\_SCHEMA

- Only simple queries of SCHEMATA, TABLES, COLUMNS, STATISTICS, and PARTITIONS are supported. No subqueries, JOINS, aggregate functions, ORDER BY, and LIMIT are allowed.

## Broadcast Tables

If a broadcast table is used, do not use any function that has different results returned each time it is executed. Otherwise, data inconsistency will occur between different shards. If such functions are indeed required, calculate their results, write the results to your SQL statements, and then execute the SQL statements on the broadcast table. Functions of this type include but are not limited to the following:

- CONNECTION\_ID()
- CURDATE()
- CURRENT\_DATE()
- CURRENT\_TIME()
- CURRENT\_TIMESTAMP()
- CURTIME()
- LAST\_INSERT\_ID()
- LOCALTIME()
- LOCALTIMESTAMP()
- NOW()
- UNIX\_TIMESTAMP()
- UTC\_DATE()
- UTC\_TIME()
- UTC\_TIMESTAMP()
- CURRENT\_ROLE()
- CURRENT\_USER()
- FOUND\_ROWS()
- GET\_LOCK()
- IS\_FREE\_LOCK()
- IS\_USED\_LOCK()
- JSON\_TABLE()
- LOAD\_FILE()
- MASTER\_POS\_WAIT()
- RAND()

- RELEASE\_ALL\_LOCKS()
- RELEASE\_LOCK()
- ROW\_COUNT()
- SESSION\_USER()
- SLEEP()
- SYSDATE()
- SYSTEM\_USER()
- USER()
- UUID()
- UUID\_SHORT()

## 5.4 High-risk Operations

Pay attention to the following when you use DDM:

- Do not connect to any data node for data operations to avoid deleting by mistake system catalogs or metadata.
- Do not clear system tables **TBL\_DRDS\_TABLE** and **MYCAT\_SEQUENCE** to prevent metadata loss.

# 6 Regions and AZs

---

## Concepts

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

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides services of the same type only or for specific tenants.
- An AZ covers one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, compute, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

HUAWEI CLOUD provides services in many regions around the world. You can select a region and AZ as needed.

## Selecting a Region

When selecting a region, consider the following factors:

- Location  
Select the region closest to you or your target users to reduce likelihood of latency issues. Chinese mainland regions provide basically the same infrastructure, BGP network quality, as well as operations and configurations on resources. If you or your target users are in the Chinese mainland, it is not necessary to consider network latency differences when selecting a region.
- Resource price  
Resource prices may vary in different regions.

## Selecting an AZ

When determining whether to deploy resources in the same AZ, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs in the same region.
- For low network latency, deploy resources in the same AZ.

# 7 Application Scenarios

---

As an OLTP-oriented service, DDM allows you to access distributed relational databases across a variety of industries.

It is especially suitable for applications requiring high-concurrency access to large volumes of data. Typical application scenarios are as follows:

- **Internet**  
E-commerce, finance, O2O, retail, and social networking applications usually face challenges such as large user base, frequent marketing events, and slow response of core transactional systems. DDM can scale compute and storage resources to improve database processing of high-concurrency transactions and ensure fast access to data.
- **IoT**  
In industrial monitoring, remote control, smart city extension, smart home, and Internet of Vehicles (IoV) scenarios, a large number of sensors and monitoring devices frequently collect data and generate huge amounts of data, which may exceed the storage capability of single-node databases. DDM provides horizontal expansion to help you store massive data at low costs.
- **Traditional sectors**  
Government agencies, large-sized enterprises, banks, and the like usually use commercial solutions to support high-concurrency access to large volumes of data. These solutions are expensive because they need to rely on mid-range computers and high-end storage devices. DDM, deployed in clusters with common ECSs, provides cost-efficient database solutions with the same or even higher performance than traditional commercial database solutions.

# 8 Permissions Management

---

If your account does not need individual IAM users for permissions management, you may skip over this chapter.

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

You can create IAM users for your employees, and assign permissions to these users to control their access to specific types of resources. For example, you can create IAM users for software developers and assign specific permissions to allow them to use DDM resources but disallow them to delete the resources or perform any high-risk operations.

IAM is free of charge. You pay only for the resources in your account.

## DDM Permissions

By default, new IAM users do not have any permissions assigned. To assign permissions to these new users, you need to add them to one or more groups, and attach permissions policies or roles to these groups.

DDM is a project-level service deployed in specific physical regions. When assigning DDM permissions to a user group, you need to specify region-specific projects where the permissions will take effect. If you select **All projects**, the permissions will be granted for all region-specific projects. To access DDM, you need to switch to the region where you are authorized.

You can grant users permissions using roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism that provides only a limited number of service-level roles. When using roles to grant permissions, you also need to assign other dependent roles. Roles are not ideal for fine-grained authorization and secure access control.
- **Policies:** A fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization and more secure access control. For example, you can grant IAM users only the permissions for managing a certain type of DDM resources.

**Table 8-1** System-defined policies

Policy Name	Description	Type	Dependency
DDM FullAccess	Full permissions for Distributed Database Middleware	System-defined policy	None
DDM CommonOperations	Common permissions for Distributed Database Middleware, excluding the permissions to create, delete, and add nodes, configure shards, and roll back shard configuration tasks	System-defined policy	None
DDM ReadOnlyAccesses	Read-only permissions for Distributed Database Middleware	System-defined policy	None

The following are permission configurations of supported system-defined policies:

- DDM FullAccess

```
{
  "Version": "1.1",
  "Statement": [{
    "Action": ["ddm:*:*",
      "rds:instance:list",
      "rds:instance:modify",
      "rds:instance:modifyParameter",
      "vpc:*:*",
      "ecs:*:get*",
      "ecs:*:list*",
      "ecs:cloudServerNics:update",
      "ecs:serverInterfaces:use"],
    "Effect": "Allow"
  }]
}
```

- DDM CommonOperations

```
{
  "Version": "1.1",
  "Statement": [{
    "Action": [
      "vpc:*:list*",
      "vpc:*:get*",
      "vpc:ports:update",
      "ecs:*:get*",
      "ecs:*:list*",
      "rds:instance:list",
      "rds:instance:modify",
      "rds:instance:modifyParameter"
    ],
    "Effect": "Allow"
  }],
  {
    "Condition": {
      "StringEqualsIgnoreCase": {
        "g:ServiceName": [
          "ddm"
        ]
      }
    }
  }
},
```

```

        "NotAction": [
            "ddm:instance:create",
            "ddm:instance:delete",
            "ddm:database:migrate*",
            "ddm:instance:resize",
            "ddm:instance:extendNode"
        ],
        "Effect": "Allow"
    }
}

```

- DDM ReadOnlyAccess

```

{
    "Version": "1.1",
    "Statement": [{
        "Action": [
            "rds:instance:list",
            "vpc:*:list*",
            "vpc:*:get*",
            "ecs:*:get*",
            "ecs:*:list*",
            "ddm:*:list",
            "ddm:*:get",
            "ddm:instance:listParameter",
            "ddm:instance:listRwInfo",
            "ddm:instance:listSlowSqlInfo",
            "ddm:rds:connectivity"
        ],
        "Effect": "Allow"
    }]
}

```

**Table 8-2** lists the common operations supported by each DDM system-defined policy. Choose appropriate system-defined policies based on your requirements.

**Table 8-2** Common operations supported by each system-defined policy

Operation	DDM FullAccess	DDM CommonOperations	DDM ReadOnlyAccess
Querying DDM instances	Supported	Supported	Supported
Querying details of a DDM instance	Supported	Supported	Supported
Modifying instance information, including the name and security group	Supported	Supported	Not supported
Restarting a DDM instance	Supported	Supported	Not supported
Creating a DDM instance	Supported	Not supported	Not supported
Deleting a DDM Instance	Supported	Not supported	Not supported
Changing node class	Supported	Not supported	Not supported
Scaling out a DDM instance	Supported	Not supported	Not supported



Operation	DDM FullAccess	DDM CommonOperations	DDM ReadOnlyAccess
Creating a schema	Supported	Supported	Not supported
Querying schemas	Supported	Supported	Supported
Querying details of a schema	Supported	Supported	Supported
Performing a rollback if configuring shards fails Deleting source data if configuring shards fails Retrying if configuring shards fails	Supported	Not supported	Not supported
Deleting a schema	Supported	Supported	Not supported
Querying accounts	Supported	Supported	Supported
Creating an account	Supported	Supported	Not supported
Modifying an account	Supported	Supported	Not supported
Resetting a password	Supported	Supported	Not supported
Deleting an account	Supported	Supported	Not supported
Synchronizing data node information	Supported	Supported	Not supported
Querying data nodes	Supported	Supported	Supported
Querying details of a data node	Supported	Supported	Supported
Modifying the read policy of a data node	Supported	Supported	Not supported
Viewing products	Supported	Supported	Supported
Creating a parameter template	Supported	Supported	Not supported
Deleting a parameter template	Supported	Supported	Not supported
Applying a parameter template	Supported	Supported	Not supported
Modifying a parameter template	Supported	Supported	Not supported
Replicating a parameter template	Supported	Supported	Not supported

Operation	DDM FullAccess	DDM CommonOperations	DDM ReadOnlyAccess
Comparing two parameter templates	Supported	Supported	Supported
Querying parameter templates	Supported	Supported	Supported

**Table 8-3** Common operations and supported actions

Operation Category	Operation	Action
DDM routine operations	As required	<p>ddm:instance:create</p> <p>Before you a DDM instance, obtain the following dependent permissions:</p> <ul style="list-style-type: none"> <li>• ecs:*.get*</li> <li>• ecs:*.list*</li> <li>• vpc:vpcs:list</li> <li>• vpc:securityGroups:get</li> <li>• vpc:subnets:get</li> <li>• ecs:cloudServerNics:update</li> <li>• ecs:serverInterfaces:use</li> <li>• <b>vpc:ports:*</b> when you want to buy a global or regional DDM instance</li> </ul>
	Querying DDM instances	ddm:instance:list
	Querying details of a DDM instance	<p>ddm:instance:get</p> <p>To view details of a DDM instance, you need to configure the following permissions:</p> <ul style="list-style-type: none"> <li>• vpc:*.get*</li> <li>• vpc:*.list*</li> </ul>
	Modifying instance information , including the name and security group	<p>ddm:instance:modify</p> <p>To modify a security group, you need to configure the following permissions:</p> <ul style="list-style-type: none"> <li>• vpc:*.get*</li> <li>• vpc:*.list*</li> <li>• vpc:ports:update</li> </ul>

Operation Category	Operation	Action
	Restarting a DDM instance	ddm:instance:reboot
	Deleting a DDM instance	ddm:instance:delete vpc:ports:delete
	Changing node class	ddm:instance:resize
	Scaling out a DDM instance	ddm:instance:extendNode
	Monitoring the read/write ratio	ddm:instance:listRwInfo
	Monitoring slow SQL queries	ddm:instance:listSlowSqlInfo
DDM routine operations	Renewing a yearly/monthly DDM instance	Configure policies <b>BSS Finance</b> and <b>BSS Operator</b> as follows: <ol style="list-style-type: none"> <li>1. Log in to the IAM console.</li> <li>2. In the navigation pane, click <b>User Groups</b>.</li> <li>3. Choose <b>More &gt; Assign Permissions</b>.</li> <li>4. Click <b>Attach Policy</b> in the same row as the project for which you want to edit the permissions.</li> <li>5. In the <b>Available Policies</b> area, select <b>BSS Finance</b> and <b>BSS Operator</b>.</li> </ol>
DDM routine operations	Changing to yearly/monthly billing	Configure policies <b>BSS Finance</b> and <b>BSS Operator</b> . The procedure is the same as that for renewing an instance.
DDM schema operations	Creating a schema	ddm:database:create
	Querying schemas	ddm:database:list
	Querying details of a schema	ddm:database:get

Operation Category	Operation	Action
	Performing a rollback if configuring shards fails Deleting source data if configuring shards fails Retrying if configuring shards fails	ddm:database:migrateRollback
	Deleting a schema	ddm:database:delete
DDM account operations	Querying accounts	ddm:user:list
	Creating an account	ddm:user:create
	Modifying an account	ddm:user:modify
	Resetting a password	ddm:user:modify
	Deleting an account	ddm:user:delete
DDM and data node management (using an RDS for MySQL instance as an example)	Synchronizing data node information	ddm:rds:synchro To synchronize data node information, you need to configure the following permissions: <ul style="list-style-type: none"> <li>• rds:instance:list</li> <li>• rds:instance:modify</li> <li>• rds:instance:modifyParameter</li> </ul>
	Querying data nodes	ddm:rds:list
	Querying details of a data node	ddm:rds:get
	Modifying the read policy of a data node	ddm:rds:modifyReadPolicy

Operation Category	Operation	Action
DDM product operations	Viewing products	ddm:product:list
DDM parameter template operations	Creating a parameter template	ddm:param:create
	Deleting a parameter template	ddm:param:delete
	Applying a parameter template	ddm:param:apply
	Modifying a parameter template	ddm:param:update
	Replicating a parameter template	ddm:param:create
	Comparing two parameter templates	ddm:param:list
	Querying parameter templates	ddm:param:list

# A Change History

---

Released On	Description
2022-09-30	This is the first official release.