ALIBABA CLOUD

Alibaba Cloud

Tablestore

FAQ

Document Version: 20220325

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.General FAQ

1.1. What are the differences between Tablestore and traditional database services such as MySQL and SQL Server?

Tablestore is a multi-model NoSQL database service that is developed by Alibaba Cloud. It can store a large amount of structured data and supports fast query and analysis. Tablestore and traditional database services (RDBMS such as MySQL and SQL Server) differ in data models and technologies that are implemented.

Tablestore and traditional database services have the following differences:

- You need to deploy and maintain the software and hardware of traditional database services. You are charged only for the actual resources you use such as storage resources and read/write capacity units (CUs) when you use Tablestore.
- Traditional database services provide a variety of features such as such as view, indexing, and transaction. Tablestore provides features such as global secondary index, search index, and local transaction, which is superior in scalability to support large amounts of data (several hundred TB) and higher concurrency level of access (100,000 QPS per table).
- Traditional database services use SQL statements. Tablestore provides standard HTTP RESTful APIs and SDKs in multiple languages.
- Traditional database services require strict schemas. Tablestore does not use the tabular schema of rows and columns that are found in most traditional database systems. Instead, Tablestore uses a storage model that is optimized for the specific requirements of the type of data being stored. Columns in each row can be different. You can add or remove attribute columns.

1.2. What do I need to specify when I create a table?

Tablestore can store semi-structured data. When you create a table, you need only to specify one to four primary key columns. Attribute columns are optional.

Each table can contain an unlimited number of attribute columns. Rows in each table can contain attribute columns that differ in column quantity and supported data types. When you use an application to write data, Tablestore requires that the application have the names of primary key columns and attribute columns and the values in all these columns specified.

What is partition key?

The first primary key column is called the partition key. When the size of data in a table reaches a specified value, Tablestore distributes data to different partitions based on the range of values in the partition key to implement load balancing

By default, data in a table is distributed to one partition when the table is created. When multiple partitions are created for the table, data distributed to a partition shares the same specified range of the values in the partition key. The range of values in a partition key is sorted by the system based on different data types of the values such as INTEGER and STRING.

The number of partitions affects performance when you access data and the utilization of reserved read/write throughput. When a table contains multiple partitions, the reserved read/write throughput is pre-distributed to each partition based on a proportion.

How can I select a partition key?

When you create a table, select a partition key based on the following rules to ensure access performance if the table stores large amounts of data.

- Do not use primary key columns whose range of values is narrow, such as customer gender Male or Female.
- Do not use primary key columns whose values may be frequently accessed. For example, the primary key column that contains timestamps is used as the partition key.
- We recommend that you use primary key columns that contain data that is less frequently accessed as the partition key. For example, you can select the UserID column.

What do I do if I cannot predict which data will be frequently accessed when I select a partition key?

We recommend that you calculate the hash value of data in a column before you write data to a partition key. For example, when you write a row of data, you can use simple algorithms to calculate the hash value of characters from a user ID, concatenate the hash value and the actual value, and store the concatenation result of the user ID in the table. You can use this lightweight operation to distribute access. Note that a value in a partition key is concatenated by a hash value and actual value. You cannot perform the GetRange operation based on the partition key.

Why Tablestore limits the number of tables that an account can have?

Each Tablestore account can create a maximum of 10 instances. Each instance can have 64 tables created. In other words, a maximum of 640 tables can be created in an account.

The following situations are considered to increase the number of tables:

- Large amounts of data and high access performance

  Unlike traditional SQL-based database services such as MySQL that addresses access to large amounts of data by using sharding, Tablestore resolves these bottlenecks by implementing a distributed solution.

  You can store structured or semi-structured data in a large sparse table without having to worry about performance issues when the amount of data to access is extremely large.

- Rapid business increase

  In addition to the increase based on existing data and increase of visits, you may use Tablestore to provide services for your customers such as the third-party partners and providers. To provide services for the providers, you can deploy a table for each provider after you activate Tablestore. The number of tables may quickly reach the upper limit. The constant increase of the upper limit can cause maintenance costs uncontrollable and make global data analysis more difficult.

  We recommend that you use Tablestore in a nontraditional way and store large amounts of structured and semi-structured data in a table.

- Considerations for Tablestore

The distributed mode of Tablestore makes the number of tables a resource attribute of Tablestore. When the scale of the Tablestore cluster is specified, Tablestore imposes no limits on the maximum number of tables. The scalability of Tablestore allows you to increase the upper limit. Tablestore specifies the upper limit of tables in an account in terms of controllability of Tablestore resources.

To increase the upper limit of tables in an account, submit a ticket.

1.3. What are the naming conventions of tables, columns, and instances?

This topic describes the naming conventions of Tablestore tables, columns, and instances.

Names of tables and columns

Conventions:

- The name can contain only letters, digits, and underscores (_).
- The name is case-sensitive.
- The name cannot start with a digit.
- The name must be 1 to 255 characters in length.

Names of instances

The name of an instance must be unique within a region. The name can be the same across different regions.

Conventions:

- The name can contain only letters, digits, and hyphens (-).
- The name is case-insensitive.
- The name must start with a letter and cannot end with a hyphen (-).
- The name must be 3 to 16 bytes in length.

1.4. What are primary key, partition, and partition key?

Primary key

Each row is uniquely identified by the primary key. You must specify columns that compose a primary key when you create a table. These columns are called primary key columns. Each primary key column must contain values. You must ensure that the combination of values of primary key columns can identify a row. The data types of values in primary key columns cannot be modified after the data types are specified.

Partition and partition key

Tablestore automatically splits tables into different partitions to balance data loads. The first column of a primary key is called the partition key.

Data that shares the same partition keys is distributed to the same partition. You can use the partition keys to modify data in the same partition, which ensures consistency.

The following figure shows a part of a mailing list from an email system. Information about the primary key and partition key:

- UserID, ReceiveTime, and FromAddr are primary key columns that uniquely identify an email. UserID is the partition key.
- ToAddr, MailSize, Subject, and Read are attribute columns that are used to store information related to emails.
- Data whose user ID is U0001 and U0002 is distributed to a partition. Data whose user ID is U0003 and U0004 is distributed to another partition.

UserID	ReceiveTime	FromAddr	ToAddr	MailSize	Subject	Read
U0001	1998-1-1	eric@demo.com	bob@demo.com	10000	Hello	Y
U0001	2011-10-20	alice@demo.com	bob@demo.com; vivian@demo.com	15000	Fw: Greetings	Y
U0001	2011-10-21	alice@demo.com	team1@demo.com	8900	Review	N
U0001	2011-10-24	lucy@demo.com	vteam@demo.com	500	Meeting Request	N
U0001	2011-11-9	alice@demo.com	bob@demo.com; team@demo.com; robin@demo.com	1250	Re: Review	N
U0001	2011-11-11	alice@demo.com	team@demo.com	3300	Re: Review	Y
U0002	1999-12-1	bill@demo.com	tom@demo.com; windy@demo.co; bill@demo.com	4300	Re: Hello	Y
U0002	2010-3-18	windy@demo.com	tom@demo.com	500	Meeting Request	Y
U0003	2010-11-11	robin@demo.com	vteam@demo.com	21000	Have a nice day	Y
U0003	2010-12-10	bob@demo.com	vteam@demo.co; alice@demo.com	10000	Re: help	N
U0004	1999-6-29	vivian@demo.com	vivian@demo.com	50	Test	N

数据分片键	主键	列	数据分片

1.5. What are the differences between Tablestore and spreadsheets in tables, rows, columns, and values?

Tablestore tables store structured data. You can query, insert, modify, and delete data in the tables. One user can create multiple tables. Data in the tables is organized in forms of columns, rows, and values.

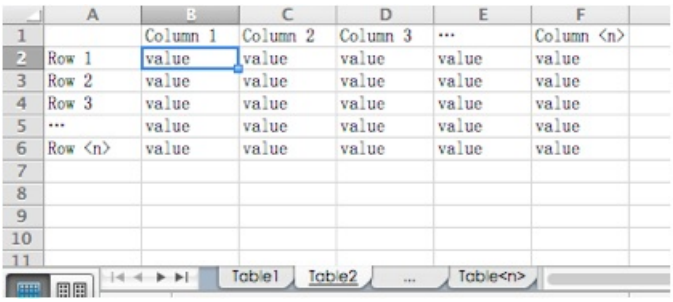

The preceding figure compares concepts such as table between Tablestore and spreadsheets.

- **table**: similar to sheets in spreadsheets. Each sheet corresponds to a table.
- **row**: similar to rows in spreadsheets. Each row consists of column names and values in the columns.
- **column**: similar to the column in spreadsheets. All data in a column shares the same dimension attributes.
- **value**: similar to the cell in spreadsheets. Each value indicates the value of a column in a specific row. Unlike spreadsheets, Tablestore can contain columns that contain null values. The following data types are supported for values: STRING, INTEGER, BOOLEAN, DOUBLE, and BINARY. Data types of values in primary key columns can only be STRING, INTEGER, or BINARY.

1.6. How does Tablestore verify a user?

Tablestore uses symmetric signatures to verify whether a request is sent by the resource owner and whether the response is sent by Tablestore.

After you create an Alibaba Cloud account, you can create an AccessKey pair in the User Management Console. For more information, see Obtain an AccessKey pair. An AccessKey ID identifies a user. An AccessKey secret is used to sign and verify requests and responses. Keep the AccessKey secret secure.

When you send a request, you must include the request plaintext, AccessKey ID, and the verification code generated based on the signature calculated by using AccessKey secret and request plaintext.

After Tablestore receives the request, Tablestore uses the AccessKey ID to obtain the corresponding AccessKey secret and signs the plaintext by using the same method. If the calculated verification code is the same as what is provided by the client, the request is valid.

The client uses the same method to verify whether the calculated verification code is the same as what is provided by the server. If the verification codes are the same, the response is valid.

1.7. How can I obtain an AccessKey pair?

You can create an AccessKey pair for an Alibaba Cloud account or a RAM user. When you call API operations, you must use the AccessKey pair to complete identity verification.

Context

An AccessKey pair consists of an AccessKey ID and an AccessKey secret.

- The AccessKey ID is used to identify a user.
- The AccessKey secret is used to verify the identity of the user. You must keep your AccessKey secret strictly confidential.

Warning If the AccessKey pair of your Alibaba Cloud account is leaked, your resources are exposed to potential risks. We recommend that you use the AccessKey pair of a RAM user to call API operations. This prevents the AccessKey pair of your Alibaba Cloud account from being leaked.

Procedure

1. Log on to the Alibaba Cloud Management Console by using your Alibaba Cloud account.
2. In the upper-right corner of the page, move the pointer over the profile picture and click **AccessKey Management**.
3. In the **Note** dialog box, click Use Current AccessKey Pair or Use AccessKey Pair of RAM User.

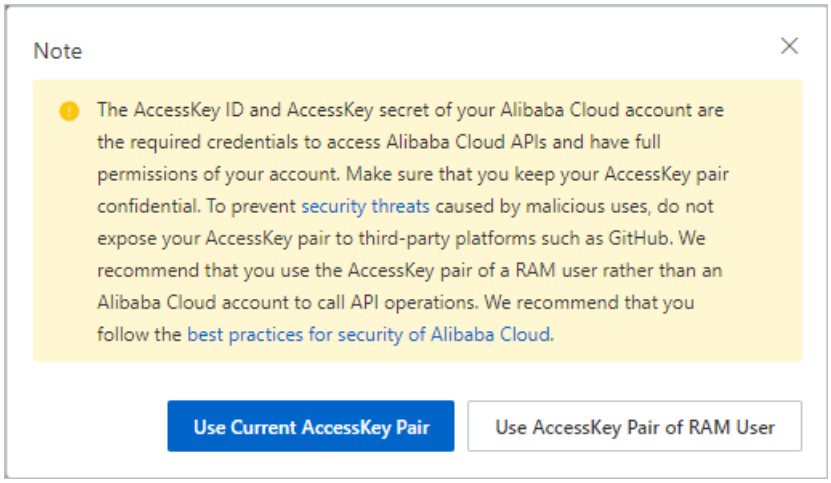

- Obtain the AccessKey pair of the Alibaba Cloud account
  a. In the Note dialog box, click **Use Current AccessKey Pair**.
  b. On the **AccessKey Pair** page, click **Create AccessKey Pair**.
  c. In the **View Secret** dialog box, view the AccessKey ID and AccessKey secret. Click **Download CSV File** to download the AccessKey pair or click **Copy** to copy the AccessKey pair.

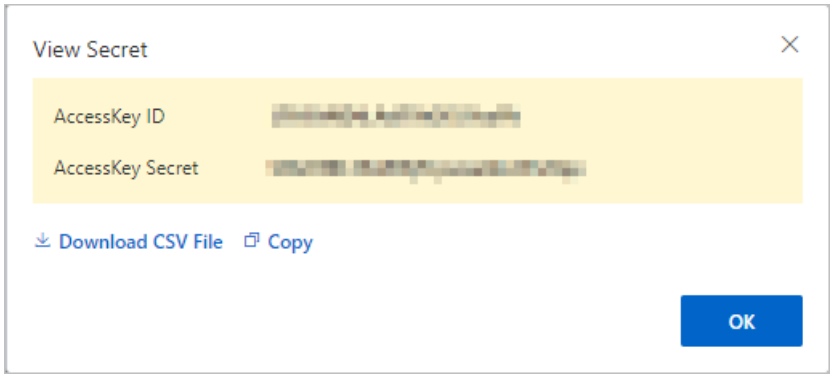

- Obtain the AccessKey pair of a RAM user
  a. Click **Use AccessKey Pair of RAM User**. Then, you are redirected to the RAM console.

b. On the **Users** page of the RAM console, find the RAM user whose AccessKey pair you want to obtain.

> **Note** If you do not have a RAM user, create one first. For more information, see Create a RAM user.

c. Click the name of the RAM user in the User Logon Name/Display Name column.

d. In the **User AccessKeys** section of the **Authentication** tab, click **Create AccessKey**.

e. In the **View Secret** dialog box, view the AccessKey ID and AccessKey secret. Click **Download CSV File** to download the AccessKey pair or click **Copy** to copy the AccessKey pair.

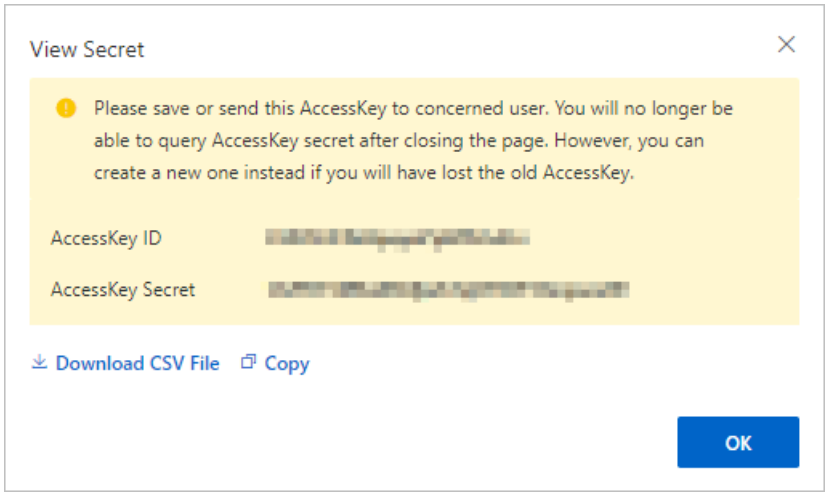

> **Note**
>
> - An AccessKey secret is displayed only after you click Create AccessKey. You cannot query the AccessKey secret in subsequent operations. Therefore, you must back up your AccessKey secret.
> - If your AccessKey pair is leaked or lost, you must create another AccessKey pair. You can create a maximum of two AccessKey pairs for each RAM user.

1.8. How does Tablestore use AccessKey pairs to control access?

Tablestore uses an AccessKey pair that consists of an AccessKey ID and AccessKey secret to authenticate the identity of a user. Users who pass the authentication can access all Tablestore resources owned by your account.

An account may have different AccessKey pairs. However, the Tablestore resources that you can access are the same when you use different AccessKey pairs.

To access Tablestore as a RAM user, the RAM user must be authorized by an Alibaba Cloud account to access specified resources.

1.9. Does the amount of data affect the query speed?

No. The performance of single-row query and range query is not affected by the amount of data.

Tablestore is a NoSQL database service. The size of data stored in Tablestore tables can be linearly scaled up to the cluster scale without affecting the speed of single-row query and range query. The query speed is not affected even if the count of rows reaches several hundred million or tens of billions.

The time to query a high-performance instance (SSD-based) is within several milliseconds. If the size of a single row is small, the time for the query is within 10 ms.

For more information about query-related operations, see GetRow, GetRange, and BatchGetRow.

1.10. How do I view the total count of rows in a table?

You can obtain the total count of rows in a table by calling the GetRange operation, using search index, or using Data Lake Analytics (DLA).

- Call the GetRange operation

  Call the GetRange operation to calculate the count of rows in a table. Low-concurrency levels and poor performance may result in latencies when you call this operation.

- Use search index

  > **Note** To use search index, you must Create search indexes.

  - Call MatchAllQuery of search index to return the total count of rows in a table within milliseconds. For more information, see Match all query.
  - Call aggregation of search index to return the total count of rows in a table within milliseconds. For more information, see Aggregation.

1.11. What do I do if "OTSErrorMsg: Disallow read index table in building base state" is displayed?

Problem description

Similar information is displayed when I read data from an index table for which existing data is being read from the base table:

```
OTSErrorMsg: Disallow read index table in building base state
```

Cause

Existing data is read from the base table and synchronized to the index table. Data cannot be read from the index table before existing data has been synchronized. You can read data from the index table only after existing data is synchronized to the index table. The time used to synchronize existing data is related to the size of data in the base table.

1.12. Does Tablestore support the IN and BETWEEN...AND operators?

If you want to perform an IN query, use one of the following methods that best fit your business requirements.

- Method 1: Use the terms query feature of search index to query multiple terms in any column. Terms query provides the same functionality as the IN query of MySQL, but delivers better performance.
- Method 2:Call the BatchGetRow operation to query only terms of a complete primary key.

If you want to perform a BETWEEN...AND query, use one of the following methods that best fit your business requirements.

- Method 1: Use the range query feature of search index to query a specific range of any column.
- Method 2: Call the GetRange operation to query only a specific range based on complete primary keys or primary key prefixes.

1.13. How do I view the size of data stored in a table?

This topic describes how to view the size of data stored in a table.

Procedure

- Log on to the Tablestore console. Go to the **Details** tab of a table. View the size of data stored in the table in the Description section.

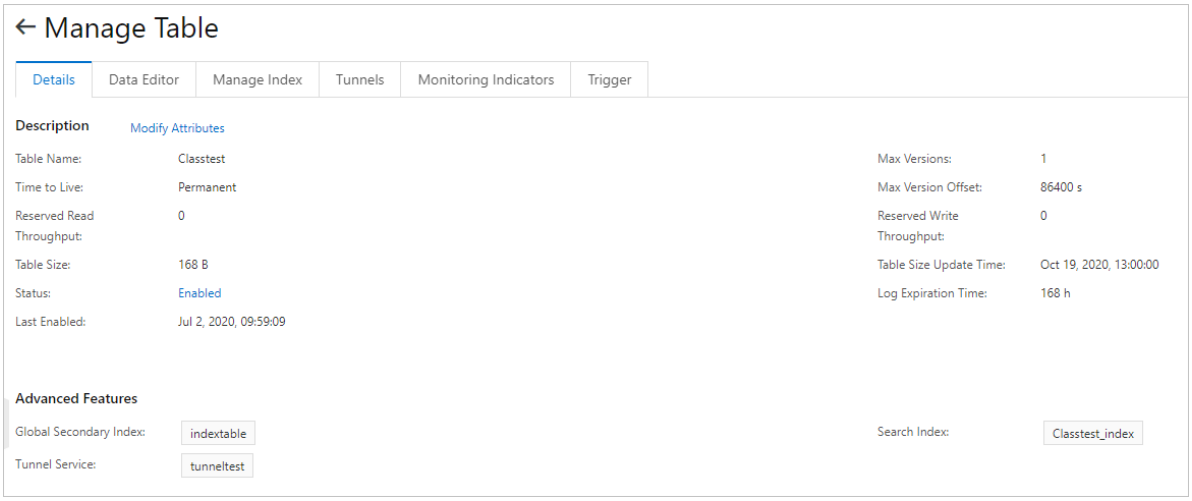

- Log on to the Tablestore console. Go to the **Monitoring Indicators** tab of a table. In the Catergories section, click **Table Size** to view the size of data stored in the table.

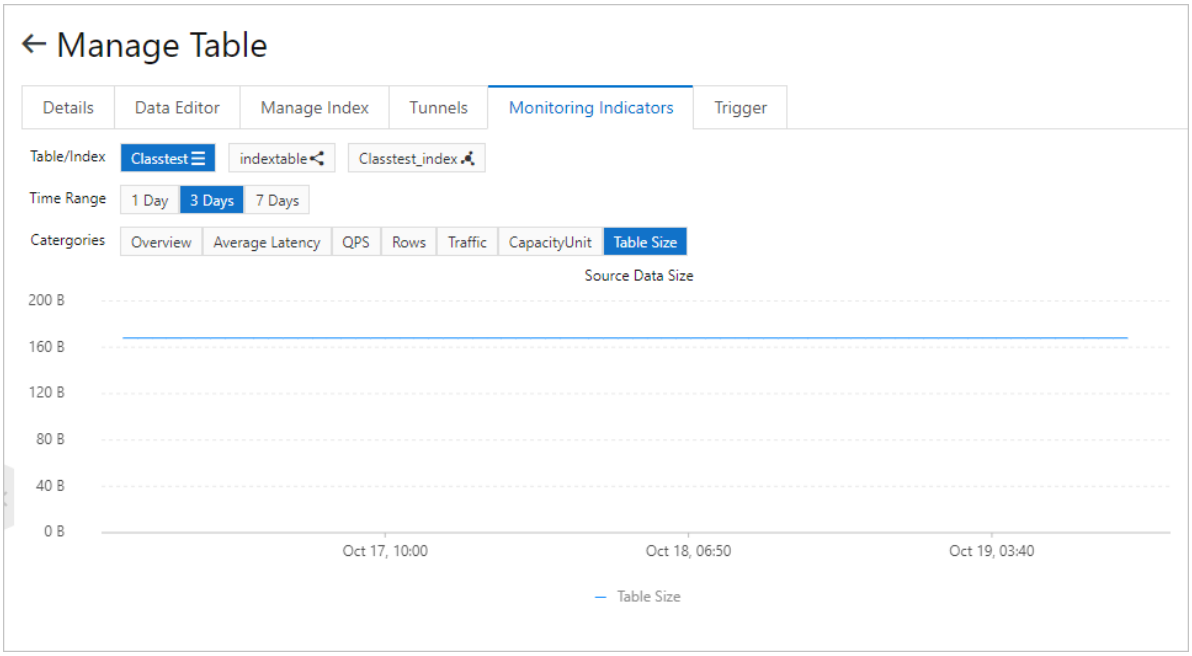

>

1.14. How do I delete all data from a table?

You can delete all data from a table by using one of the following methods: 1. Delete the table and create a table that has the same name. 2. Obtain information about the primary key of the table. Then, delete the data in the table.

- Delete a table and create a table of the same schema

  To delete a table, you can use Tablestore SDKs to call the DeleteTable operation. You can also delete a table in the Tablestore console. After you delete the table, you must create a table that has the same schema.

  > Note
  >
  > - If the table has indexes, you must delete the indexes before you delete the table.
  > - After the indexes of the table are deleted, you must create indexes for the created table if you want to use indexes.

- Obtain the primary keys of all rows in the table to delete data from the table

  Obtain the primary keys of all rows in a table by using Tablestore SDK to call the GetRange operation. And then call the BatchWriteRow or DeleteRow operation to delete all data from the table.

1.15. How do I increase the value of the limit parameter to 1000 when I call the Search operation of the search index feature to query data?

This topic describes how to increase the value of the limit parameter to 1000 when you call the Search operation of the search index feature to query data.

To increase the number of returned rows in a query, call the Search operation in search indexes. If only data in search indexes is queried, the value of the limit parameter automatically increases to 1000. If data in tables is queried, the value of the limit parameter is 100.

Procedure

To increase the value of the limit parameter to 1000, perform the following operations:

1. When you create a search index, you must set the value of the store parameter to true for the specified column.
   - If you create a search index in the Tablestore (OTS) console, the default value of the store parameter is true. You do not need to specify this parameter.
   - If you use SDKs to create a search index, you must set the value of the store parameter in FieldSchema to true for the specified column.
2. When you call the Search operation to query data, you must specify the ColumnsToGet parameter in SearchRequest.

   The ColumnsToGet parameter returns only columns for which indexes are created in search indexes. If the data type of the returned columns is not ARRAY, GEOPOINT, or NEST, the value of the limit parameter automatically increases to 1000.

   > **Note** If the ColumnsToGet parameter returns columns whose data type is ARRAY, GEOPOINT, or NEST, the table is queried when you call the Search operation. In this case, the value of the limit parameter remains at 100.

Examples

In this example, OTS SDK for Java is used to describe how to specify the ColumnsToGet parameter. The configuration method for this SDK is similar to that for OTS SDKs for other languages. You need to modify only the ColumnsToGet parameter in SearchRequest.

```
SearchQuery searchQuery = new SearchQuery();
searchQuery.setQuery(new MatchQuery());
searchQuery.setLimit(1000);
SearchRequest searchRequest = new SearchRequest(tableName, indexName, searchQuery);
ColumnsToGet columnsToGet = new ColumnsToGet();
columnsToGet.setReturnAll(false);
columnsToGet.setColumns(Arrays.asList("field_1", "field_2", "field_3"));  // Set the names
of the returned columns. The data type of the returned columns cannot be ARRAY, GEOPOINT, o
r NEST. Otherwise, the table is queried.
searchRequest.setColumnsToGet(columnsToGet);
SearchResponse response = client.search(searchRequest);
// OTS SDK for Java V5.6.1 and later allow you to specify the returnAllColumnsFromIndex par
ameter in ColumnsToGet to query all attribute columns in search indexes.
ColumnsToGet columnsToGet = new ColumnsToGet();
columnsToGet.setReturnAllFromIndex(true);
searchRequest.setColumnsToGet(columnsToGet);
SearchResponse response = client.search(searchRequest);
```

1.16. Can I query joined Tablestore tables?

No. Tablestore does not support queries based on joined tables. You can combine Tablestore with Data Lake Analytics (DLA) and use standard SQL statements that are compatible with most SQL syntax of MySQL 5.7 to analyze data in Tablestore. For more information, see DLA-based analysis of data in Tablestore by using SQL statements.

1.17. How do I configure split points when I use Tablestore Reader to synchronize full data?

This topic describes the causes of and solutions to slow synchronization when Tablestore Reader is used to synchronize full data.

Problem description

When Tablestore Reader is used to synchronize full data, data is synchronized at a low speed. The following script provides an example on how to configure full data synchronization:

```
"reader": {
  "plugin": "ots",
  "parameter": {
    "datasource": "",
    "table": "",
    "column": [],
    "range": {
      "begin": [
        {
          "type": "INF_MIN"
        }
      ],
      "end": [
        {
          "type": "INF_MAX"
        }
      ]
    }
  }
}
```

Cause

A large amount of data must be synchronized. However, no split points are configured in the script, and only one thread is created to obtain the data, which affects the speed at which the data is synchronized.

Solution

If you want to synchronize a large amount of data by using Tablestore Reader, configure split points in the script. To configure split points in the script, perform the following steps:

1. Obtain the information about the required split points by using one of the following methods:
   - Use Tablestore SDK to call the ComputeSplitPointsBySize operation. For more information, see Split data by a specified size.

     Sample response:

     ```
     LowerBound:pkname1:INF_MIN, pkname2:INF_MIN
     UpperBound:pkname1:cbcf23c8cdf831261f5b3c052db3479e, pkname2:INF_MIN
     LowerBound:pkname1:cbcf23c8cdf831261f5b3c052db3479e, pkname2:INF_MIN
     UpperBound:pkname1:INF_MAX, pkname2:INF_MAX
     ```

   - Download the Tablestore CLI tool. Then, run the following command: **points -s splitSize -t tablename**. For more information, see Tablestore CLI.

     > **Note** The unit of the splitSize value is 100 MB. If the amount of data that you want to synchronize is small, you do not need to configure split points. If the amount of data that you want to synchronize is large, we recommend that you specify a value for the splitSize parameter based on the maximum number of concurrent threads supported in your environment.

     Sample response:

     ```
     [
         {
             "LowerBound": {
                 "PrimaryKeys": [
                     {
                         "ColumnName": "pkname1",
                         "Value": null,
                         "PrimaryKeyOption": 2
                     },
                     {
                         "ColumnName": "pkname2",
                         "Value": null,
                         "PrimaryKeyOption": 2
                     }
                 ]
             },
             "UpperBound": {
                 "PrimaryKeys": [
                     {
                         "ColumnName": "pkname1",
                         "Value": "cbcf23c8cdf831261f5b3c052db3479e\u0000",
                         "PrimaryKeyOption": 0
                     },
                     {
                         "ColumnName": "pkname2",
                         "Value": null,
                         "PrimaryKeyOption": 2
                     }
                 ]
             },
             "Location": "80310717938EDF503FB1E26F70710391"
         },
     ```

```
    {
        "LowerBound": {
            "PrimaryKeys": [
                {
                    "ColumnName": "pkname1",
                    "Value": "cbcf23c8cdf831261f5b3c052db3479e\u0000",
                    "PrimaryKeyOption": 0
                },
                {
                    "ColumnName": "pkname2",
                    "Value": null,
                    "PrimaryKeyOption": 2
                }
            ]
        },
        "UpperBound": {
            "PrimaryKeys": [
                {
                    "ColumnName": "pkname1",
                    "Value": null,
                    "PrimaryKeyOption": 3
                },
                {
                    "ColumnName": "pkname2",
                    "Value": null,
                    "PrimaryKeyOption": 3
                }
            ]
        },
        "Location": "80310717938EDF503FB1E26F70710391"
    }
]
```

Find the values of the first primary key columns. For example, the pkname1 value of the first LowerBound is null, the pkname1 value of the first UpperBound is "cbcf23c8cdf831261f5b3c052db3479e\u0000", the pkname1 value of the second LowerBound is "cbcf23c8cdf831261f5b3c052db3479e\u0000", and the pkname1 value of the second UpperBound is null. To synchronize full data, configure the following settings in the script:

```
"split" : [
   {
       "type":"STRING",
       "value":"cbcf23c8cdf831261f5b3c052db3479e\u0000"
   }
]
```

When you run the script, Tablestore splits full data into two parts and concurrently obtains data based on the (INF_MIN,cbcf23c8cdf831261f5b3c052db3479e\u0000) and [cbcf23c8cdf831261f5b3c052db3479e\u0000,INF_MAX) ranges. This way, data synchronization is accelerated.

2. Configure split points in the script used to synchronize data. The following script provides an example on how to configure split points:

```
"range": {
      "begin": [
        {
          "type": "INF_MIN"
        }
      ],
      "end": [
        {
          "type": "INF_MAX"
        }
      ],
      "split": [
        {
          "type": "STRING",
          "value": "splitPoint1"
        },
        {
          "type": "STRING",
          "value": "splitPoint2"
        },
        {
          "type": "STRING",
          "value": "splitPoint3"
        }
      ]
}
```

If the synchronization remains slow after you configure split points, submit a ticket to contact the technical support.

1.18. Does Tablestore allow me to migrate data across regions?

Yes, Tablestore allows you to migrate data across regions. You can use DataWorks or Tunnel Service to migrate data.

When you migrate data across regions, make sure that the machine where scheduling tasks are executed is connected to the networks where the source and destination instances are deployed. You can use the following methods to test the network connectivity:

- Use DataWorks to migrate data: Create data sources for OTSReader and OTSWriter.
- Use Tunnel Service to migrate data: Log on to the machine that is connected to a tunnel and use cURL to connect to the endpoints of the source and destination instances. If Unsupported Operation is returned from the preceding test results, the networks are connected.

1.19. How do I delete multiple rows of data?

After you obtain the primary key information of the data you want to delete, you can call the BatchWriteRow operation to delete the data.

Tablestore allows you to use primary key information to delete multiple rows of data. To delete multiple rows of data, perform the following operations:

1. Select a method to query the primary key information of the data you want to delete.
   - To query the data within the specified primary key range and obtain the primary key information of the data, call the GetRange operation. Then, you can delete the data. For more information about specific operations, see GetRange.
   - To delete the data that meets the specified conditions, use search indexes to query the data. Then, obtain the primary key information of the data. For more information about specific operations, see Create search indexes and Search.
2. To delete multiple rows of data by using primary key information, call the BatchWriteRow operation. For more information about specific operations, see BatchWriteRow.

1.20. What do I do if errors occur when OTSStream Reader is running?

This topic describes errors that may occur when OTSStream Reader is running and how to resolve the errors based on the error messages.

Error message: "Must set date or time range millis or time range string, please check your config"

- Problem description

  The "Must set date or time range millis or time range string, please check your config" error message appears when OTSStream Reader is running. The following figure shows the error message.

```
com.alibaba.datax.common.exception.DataXException: [ code:OTSStreamReaderError, messageOTS Stream Reader Error] - com.alibaba.datax.plugin.reader
.otsstreamreader.internal.OTSStreamReaderException: Must set date or time range millis or time range string, please check your config. - com.alib
aba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderException: Must set date or time range millis or time range string, please check
your config.
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.config.OTSStreamReaderConfig.load(OTSStreamReaderConfig.java:201)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReader$Job.init(OTSStreamReader.java:30)
        at com.alibaba.datax.core.job.JobContainer.initJobReader(JobContainer.java:1085)
        at com.alibaba.datax.core.job.JobContainer.init(JobContainer.java:497)
        at com.alibaba.datax.core.job.JobContainer.start(JobContainer.java:220)
        at com.alibaba.datax.core.Engine.start(Engine.java:100)
```

- Possible cause

  The parameters that specify the timestamp range of incremental data are not configured in the configuration script of OTSStream Reader.

- Solution

  Specify the timestamp range of incremental data by using one of the following methods:

  - Set the value of the parameter.date parameter to the date of incremental data that you want to export.
  - Set the value of the parameter.startTimestampMillis parameter to the start timestamp of the timestamp range of incremental data, and the value of the parameter.endTimestampMillis parameter to the end timestamp of the timestamp range.

  For more information about how to configure the timestamp range of incremental data, see Synchronize incremental data in script mode.

Error Message: "The stream of data table is not enabled"

- Problem description

  The "The stream of data table is not enabled" error message appears when OTSStream Reader is running. The following figure shows the error message.

```
[ code:OTSStreamReaderError, messageOTS Stream Reader Error] - com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderException:
 The stream of data table is not enabled. - com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderException: The stream of data
 table is not enabled.
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.core.OTSStreamReaderChecker.checkStreamEnabledAndTimeRangeOK(OTSStreamReaderC
hecker.java:49)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderMasterProxy.init(OTSStreamReaderMasterProxy.java:38)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReader$Job.init(OTSStreamReader.java:31)
        at com.alibaba.datax.core.job.JobContainer.initJobReader(JobContainer.java:1085)
        at com.alibaba.datax.core.job.JobContainer.init(JobContainer.java:497)
        at com.alibaba.datax.core.job.JobContainer.start(JobContainer.java:220)
        at com.alibaba.datax.core.Engine.start(Engine.java:100)
        at com.alibaba.datax.core.Engine.entry(Engine.java:321)
        at com.alibaba.datax.core.Engine.main(Engine.java:354)
```

- Possible cause

  Stream is disabled for the data table that is specified by the paramter.table parameter in OTSStream Reader.

- Solution

  Enable Stream for the data table that is specified in OTSStream Reader by using one of the following methods:

  - Log on to the Tablestore console, find the data table for which you want to enable the Stream feature, go to the **Manage Table** page, click the **Tunnels** tab, and then enable the Stream feature.

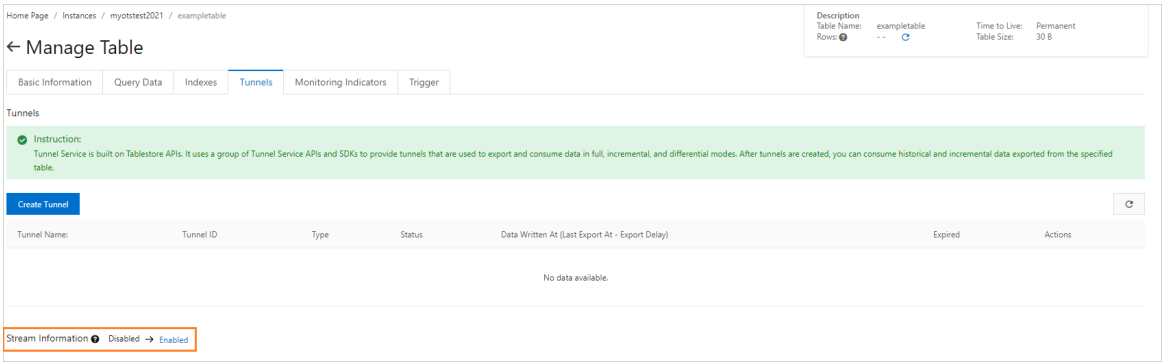

  - Modify data table attributes by using a suitable SDK to call the UpdateTable operation. For more information, see UpdateTable.

Error Message: "As expiration time is xx, so the start timestamp must greater than xx"

- Problem description

  The "As expiration time is xx, so the start timestamp must greater than xx" error message appears when OTSStream Reader is running. The following figure shows the error message.

```
com.alibaba.datax.common.exception.DataXException: [ code:OTSStreamReaderError, messageOTS Stream Reader Error] - com.alibaba.datax.plugin.reader
.otsstreamreader.internal.OTSStreamReaderException: As expiration time is 25200000, so the start timestamp must greater than 2021-09-25T20:09Z(16
32600562281) - com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderException: As expiration time is 25200000, so the start ti
mestamp must greater than 2021-09-25T20:09Z(1632600562281)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.core.OTSStreamReaderChecker.checkStreamEnabledAndTimeRangeOK(OTSStreamReaderC
hecker.java:60)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderMasterProxy.init(OTSStreamReaderMasterProxy.java:38)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReader$Job.init(OTSStreamReader.java:31)
        at com.alibaba.datax.core.job.JobContainer.initJobReader(JobContainer.java:1085)
        at com.alibaba.datax.core.job.JobContainer.init(JobContainer.java:497)
        at com.alibaba.datax.core.job.JobContainer.start(JobContainer.java:220)
        at com.alibaba.datax.core.Engine.start(Engine.java:100)
        at com.alibaba.datax.core.Engine.entry(Engine.java:321)
```

- Possible cause

  The startTime parameter that specifies the start timestamp of the timestamp range of incremental logs is not correctly configured based on the log expiration time.

- Solution

  Set the startTime parameter in the synchronization script to a value that is greater than the value that is obtained by using the following formula: `Synchronization task start time - Log expiration time + 10 minutes`.

  To obtain the log expiration time, log on to the Tablestore console, find the data table whose log expiration time you want to obtain, go to the **Manage Table** page, click the **Basic Information** tab, and then view the value of Log Expiration Time.

Error message: "To avoid timing error between different machines, the end timestamp must smaller than xx"

- Problem description

  The "To avoid timing error between different machines, the end timestamp must smaller than xx" error message appears when OTSStream Reader is running. The following figure shows the error message.

```
com.alibaba.datax.common.exception.DataXException: [ code:OTSStreamReaderError, messageOTS Stream Reader Error] - com.alibaba.datax.plugin.reader
.otsstreamreader.internal.OTSStreamReaderException: To avoid timing error between different machines, the end timestamp must smaller than 2021-09
-26T03:59Z(1632628796433) - com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderException: To avoid timing error between diff
erent machines, the end timestamp must smaller than 2021-09-26T03:59Z(1632628796433)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.core.OTSStreamReaderChecker.checkStreamEnabledAndTimeRangeOK(OTSStreamReaderC
hecker.java:65)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderMasterProxy.init(OTSStreamReaderMasterProxy.java:38)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReader$Job.init(OTSStreamReader.java:31)
        at com.alibaba.datax.core.job.JobContainer.initJobReader(JobContainer.java:1085)
        at com.alibaba.datax.core.job.JobContainer.init(JobContainer.java:497)
        at com.alibaba.datax.core.job.JobContainer.start(JobContainer.java:220)
        at com.alibaba.datax.core.Engine.start(Engine.java:100)
        at com.alibaba.datax.core.Engine.entry(Engine.java:321)
        at com.alibaba.datax.core.Engine.main(Engine.java:354)
        at com.alibaba.datax.common.exception.DataXException.asDataXException(DataXException.java:41)
```

- Possible cause

  An invalid value is specified for the endTime parameter.

- Solution

  Set the endTime parameter in the synchronization script to a value that is less than the value that is obtained by using the following formula: `Synchronization task start time - 5 minutes`.

Error message: "The item of column must be map object"

- Problem description

  The "The item of column must be map object" error message appears when OTSStream Reader is running. The following sample shows the error message.

```
[ code:OTSStreamReaderError, messageOTS Stream Reader Error] - com.alibaba.datax.plugin.r
eader.otsstreamreader.internal.OTSStreamReaderException: Parse column fail, please check
your config. - com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReaderEx
ception: Parse column fail, please check your config.
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.config.OTSStreamReade
rConfig.parseConfigForSingleVersionAndUpdateOnlyMode(OTSStreamReaderConfig.java:178)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.config.OTSStreamReade
rConfig.load(OTSStreamReaderConfig.java:267)
        at com.alibaba.datax.plugin.reader.otsstreamreader.internal.OTSStreamReader$Job.i
nit(OTSStreamReader.java:30)
        at com.alibaba.datax.core.job.JobContainer.initJobReader(JobContainer.java:1083)
        at com.alibaba.datax.core.job.JobContainer.init(JobContainer.java:497)
        at com.alibaba.datax.core.job.JobContainer.start(JobContainer.java:220)
        at com.alibaba.datax.core.Engine.start(Engine.java:100)
        at com.alibaba.datax.core.Engine.entry(Engine.java:318)
        at com.alibaba.datax.core.Engine.main(Engine.java:351)
Caused by: java.lang.IllegalArgumentException: The item of column must be map object, ple
ase check your input.
```

- Possible cause

  An invalid format is used to specify the value for the parameter.column parameter.

- Solution

  Make sure that the parameter.column parameter is configured based on a valid format. The following code provides examples of the format:

  - Incorrect example

    ```
    "column":[
         "col1",
         "col2"
    ]
    ```

  - Correct example

    ```
    "column":[
         {
              "name" : "col1"
         },
         {
               "name" : "col2"
         }
    ]
    ```

1.21. How do I resolve common errors of Tablestore Reader?

This topic describes the errors that may occur when Tablestore Reader is running. Troubleshoot the errors based on the error messages.

Error message: "Input size of values not equal size of primary key. input size: xx, primary key size: xx"

- Problem description

  The "Input size of values not equal size of primary key. input size: xx, primary key size: xx" error message appears when Tablestore Reader is running. The following figure shows the error message.

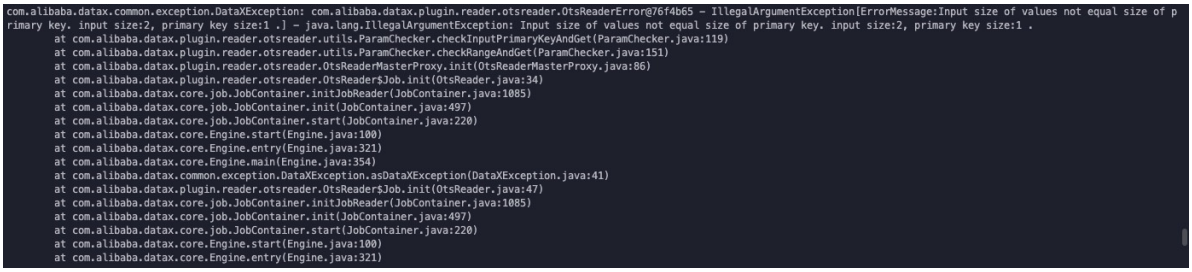

- Possible cause

  The number of primary key columns that are specified in the Tablestore Reader script is different from that in the Tablestore data table.

- Solution

  Check the schema of the Tablestore data table, obtain the number of primary key columns in the data table, and then set the obtained value as the number of primary key columns in the Tablestore Reader script.

  For example, a Tablestore data table contains two primary key columns and you want to synchronize all data in the table. The following sample code shows how to specify the number of primary key columns in the Tablestore Reader script:

```
"range": {
  "begin": [
    {
      "type": "INF_MIN"
    },
    {
      "type": "INF_MIN"
    }
  ],
  "end": [
    {
      "type": "INF_MAX"
    },
    {
      "type": "INF_MAX"
    }
  ]
}
```

Error message: "Input type of 'range-split' not match partition key. Item of 'range-split' type:xx, Partition type:xx"

- Problem description

The "Input type of 'range-split' not match partition key. Item of 'range-split' type:xx, Partition type:xx" error message appears when Tablestore Reader is running. The following figure shows the error message.

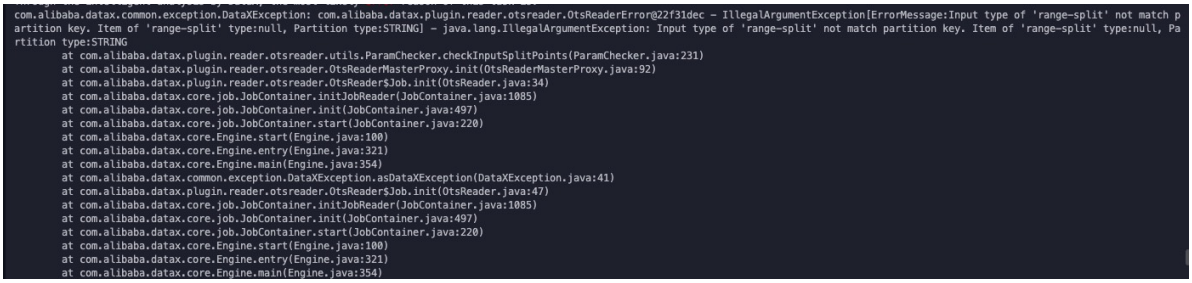

- Possible cause

  The split point is incorrectly specified in the Tablestore Reader script. The following sample code provides an example on how the split point is incorrectly specified:

```
"split": [
  {
    "type": "INF_MIN"
  },
  {
  "type": "INF_MAX"
  }
]
```

- Solution
  - Delete the split point setting from the Tablestore Reader script.
  - Specify the split point again by following the instructions provided in How do I configure split points when I use Tablestore Reader to synchronize full data?.

Error message: "Invalid 'column', Can not parse Object to 'OTSColumn', item of list is not a map"

- Problem description

  The "Invalid 'column', Can not parse Object to 'OTSColumn', item of list is not a map" error message appears when Tablestore Reader is running. The following figure shows the error message.

- Possible cause

  The column parameter is incorrectly specified in the Tablestore Reader script. The following sample code provides an example on how the column parameter is incorrectly specified:

```
"column": [
  "id",
  "name",
  "age"
]
```

- Solution

  Make sure that the column parameter is correctly specified in the Tablestore Reader script. The following sample code provides an example on how to specify the column parameter:

```
"column": [
  {
    "name": "id"
  },
  {
    "name": "name"
  },
  {
    "name": "age"
  }
]
```

1.22. What do I do if the NoPermission error message appears when I view the monitoring data in the console?

If the NoPermission error message appears when a RAM user views the monitoring data in the console, grant the RAM user the required permissions on Cloud Monitor.

Problem description

The NoPermission error message appears when I view the instance monitoring data or table monitoring data in the Tablestore console as a RAM user.

Possible cause

The RAM user is not granted the required permissions to view Cloud Monitor.

Solution

Log on to the RAM console by using your Alibaba Cloud account, and grant the RAM user the AliyunCloudMonitorFullAccess or AliyunCloudMonitorReadOnlyAccess permission. The AliyunCloudMonitorFullAccess permission allows the RAM user to manage Cloud Monitor, and the AliyunCloudMonitorReadOnlyAccess permission allows the RAM user to access Cloud Monitor in the read-only mode. For more information, see Grant permissions to a RAM user.

1.23. What do I do if data whose primary keys are out of the specified range is returned when the GetRange operation is called to read data?

If data whose primary keys are out of the specified range is returned when the GetRange operation is called to read data, you can create a secondary index to read data whose primary keys are only within the specified range.

Problem description

When the GetRange operation is called to perform range query in the Tablestore console or by using Tablestore SDKs, query the external tables of Tablestore, or synchronize data from Tablestore by using Tablestore Reader, data whose primary keys are out of the specified range is returned.

Possible cause

The specified query conditions do not comply with the leftmost matching principle. For example, if a data table contains two primary key columns, `beginPrimaryKey = INF_MIN` and `endPrimaryKey = INF_MAX` are specified for the first primary key column, and `beginPrimaryKey = 10` and `endPrimaryKey = 10` are specified for the second primary key column, all data in the table, instead of only the rows in which the second primary key column value is 10, is returned.

> **Note** The leftmost matching principle indicates that the range specified by the start primary key and the end primary key in the right primary key column only takes effect when the start primary key and the end primary key are the same within each leftmost primary key column. If the start primary key and the end primary key in a leftmost primary key column are different, the range specified by the start primary key and the end primary key in the right primary key column does not take effect. The leftmost primary key columns are defined earlier than the right primary key column in the table schema.

Solution

Create a secondary index with the sequence of primary key columns specified, and choose a proper query method as required.

- To query data within a specified range in the Tablestore console or by using specific SDKs, use the secondary index. For more information about specific operations, see Read data in the console and Read data by using SDKs.
- To query data within a specified range in a Tablestore external table, set the table name to the secondary index name in the SQL statement that is executed to create the external table. For more information, see Access Tablestore data.
- To synchronize data within a specified range from Tablestore by using Tablestore Reader, set the table name to the secondary index name in the data synchronization script. For more information about specific operations, see Tablestore Reader.

1.24. What do I do if the "Wait future timeout" error message appears when I use an SDK to send an API request?

If the "Wait future timeout" error message appears, make sure that FullGC is not triggered on the client side and the client.shutdown() method is not called during the request. You must also check whether the latency on the server side is high.

Problem description

When I use an SDK to send an API request, the "Wait future timeout" error message appears on the server side.

Possible causes

The duration from the time when the request is sent to the time when the response is returned exceeds the value of syncClientWaitFutureTimeoutInMillis. The default value of syncClientWaitFutureTimeoutInMillis is 60000, in milliseconds.

> **Note** We recommend that you do not set syncClientWaitFutureTimeoutInMillis to an excessively small value. You can specify a value for syncClientWaitFutureTimeoutInMillis in ClientConfiguration.

Possible causes:

- FullGC is triggered on the client side or you call the client.shutdown() method on the client side during the request.
- The latency on the server side is greater than the value of syncClientWaitFutureTimeoutInMillis.

Solution

- Check whether the callback thread is shut down because you call the client.shutdown() method on the client side during the request or whether FullGC is triggered on the client side.
- Log on to the Tablestore console and perform the following steps to check whether the latency on the server side is high:

  On the **Manage Table** page of the data table, click the **Monitoring Indicators** tab, select a table or an index, specify a time range, and then set Categories to **Average Latency** to view the average latency of different operation types.

  If the latency on the server side is greater than the value of syncClientWaitFutureTimeoutInMillis, submit a ticket to contact technical support.

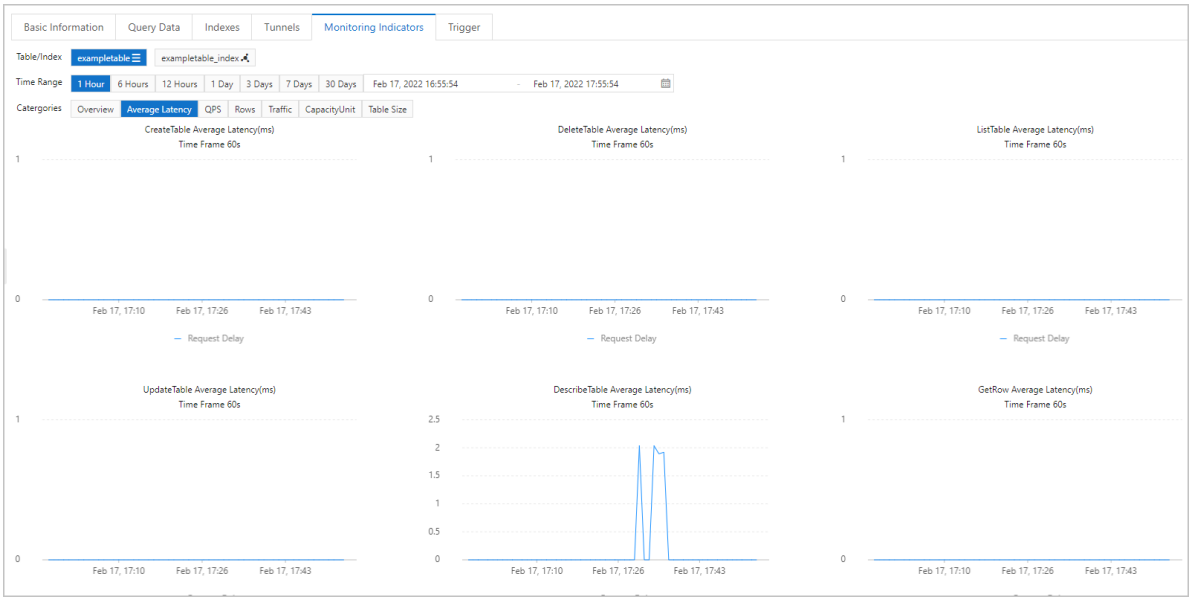
Basic Information
Query Data
Indexes
Tunnels
Monitoring Indicators
Trigger
Table/Index
exampletable
exampletable_index
Time Range
1 Hour
6 Hours
12 Hours
1 Day
3 Days
7 Days
30 Days
Feb 17, 2022 16:55:54
Feb 17, 2022 17:55:54
Catergories
Overview
Average Latency
QPS
Rows
Traffic
CapacityUnit
Table Size
CreateTable Average Latency(ms)
Time Frame 60s
DeleteTable Average Latency(ms)
Time Frame 60s
ListTable Average Latency(ms)
Time Frame 60s
UpdateTable Average Latency(ms)
Time Frame 60s
DescribeTable Average Latency(ms)
Time Frame 60s
GetRow Average Latency(ms)
Time Frame 60s
Feb 17, 17:10
Feb 17, 17:26
Feb 17, 17:43
Request Delay

2.FAQ about API/SDK

2.1. What do I do if SocketTimeoutException is reported when I use Tablestore SDK for Java?

Problem description

When the difference between the time Tablestore SDK for Java receives data and the time Tablestore SDK for Java sends data exceeds the socket timeout value, Tablestore SDK for Java returns SocketTimeoutException. The socket timeout period includes the period for network transmission between the application and server, which is calculated from the time the application sends a request to when server processes the request and when the application receives the response. You can customize the socket timeout value when you create an OTSClient instance. By default, the socket timeout value is 15s.

Solution

Solutions are provided based on the following possible causes:

- Network connection failure

  If SocketTimeoutException is returned for all requests, the possible cause is network connection failures. You can run the ping or curl command to test the connectivity of the network.

  ```
  ping aaaa.cn-hangzhou.ots.aliyuncs.com
  curl aaaa.cn-hangzhou.ots.aliyuncs.com
  ```

  In most cases, the curl command returns a similar output:

  ```
  <? xml version="1.0" encoding="UTF-8"? >
  <Error><Code>OTSUnsupportOperation</Code><Message>Unsupported operation: ''. </Message><R
  equestID>00054ec5-822c-8964-adaf-990a07a4d0c9</RequestID><HostID>MTAuMTUzLjE3NS4xNzM=</Ho
  stID></Error>
  ```

  If the connection fails, a possible cause is that an internal endpoint is used when the request is not sent by an ECS instance.

- Server-side latency that exceeds the socket timeout value specified in Tablestore SDK for Java

  The period for the Tablestore server to process a request does not exceed 15 seconds because the default maximum overtime period for the server is about 10 seconds. If the actual period exceeds 10 seconds, OTSTimeout is returned to the client.

  However, if a socket timeout value such as 2 is specified in Tablestore SDK for Java, Tablestore SDK for Java returns SocketTimeoutException.

- Network transmission latency

SocketTimeoutException may also be returned if the overall latency is extended by network transmission instead of server-side processing. Check whether traffic and bandwidth usage as well as the packet retransmission rate are high.

- Frequent full GC

  SocketTimeoutException may be returned when the program runs under heavy load due to frequent full GC. The possible cause is that when full GC occurs, the request cannot be sent or the response cannot be received. SocketTimeoutException is returned when the actual timeout period exceeds the socket timeout value specified in Tablestore SDK for Java.

  If this error occurs, you must use a tool to analyze the GC status of the process to resolve frequent GC.

2.2. FAQ about the Tablestore SDK for Java logging framework

Which logging framework does Tablestore SDK for Java use?

Tablestore SDK for Java uses Simple Logging Facade for Java (SLF4J) that depends on **Log4j 2** as the dependency to implement the logging framework.

How can I replace a logging framework?

You can remove the declaration of the dependency on **Log4j 2** from the dependency of Tablestore SDK for Java. SLF4J automatically searches your applications for a logging framework to use the framework as the dependency.

```
<dependency>
  <groupId>com.aliyun.openservices</groupId>
  <artifactId>ots-public</artifactId>
  <version>2.2.4</version>
  <exclusions>
           <exclusion>
               <groupId>org.apache.logging.log4j</groupId>
               <artifactId>log4j-api</artifactId>
          </exclusion>
           <exclusion>
               <groupId>org.apache.logging.log4j</groupId>
               <artifactId>log4j-core</artifactId>
          </exclusion>
           <exclusion>
               <groupId>org.apache.logging.log4j</groupId>
               <artifactId>log4j-slf4j-impl</artifactId>
          </exclusion>
  </exclusions>
</dependency>
```

2.3. What can I do if errors related to data types supported by primary key columns occur?

Problem description

```
Caused by: [ErrorCode]:OTSInvalidPK, [Message]:Validate PK type fail. Input: VT_STRING, Meta: VT_BLOB. [RequestId]:00055f43-3d31-012b-62c3-980a3eefe39e, [TraceId]:02822839-3b5b-af35-409a-cf68841239fa, [HttpStatus:]400
```

Cause

When the table is created, the data type supported by the primary key columns is BINARY. However, the data type of your data is STRING.

Solution

Ensure that your data is of the correct type.

2.4. What do I do if the Validate PK size fail exception occurs when I use SDKs?

This topic describes the problem description, cause, and solution of the Validate PK size fail exception when you use SDKs.

Problem description

The following exception occurs when you use SDKs to update data:

```
Caused by: [ErrorCode]:OTSInvalidPK, [Message]:Validate PK size fail
```

Cause

The number of primary keys configured is inconsistent with the actual number of primary keys of the table.

Solution

The number of primary keys configured must be consistent with the actual number of primary keys of the table.

2.5. What do I do if the Validate PK name fail exception occurs when I use SDKs?

This topic describes the problem, cause, and solution of the Validate PK name fail exception that may occur when you use SDKs.

Problem description

The following exception occurs when you use SDKs to query data:

```
Caused by: [ErrorCode]:OTSInvalidPK, [Message]:Validate PK name fail
```

Cause

The specified names of primary key columns are different from those in the data table, or the sequence of primary key columns is different from that in the data table.

Solution

The specified names and sequence of the primary key columns must be the same as those of the table.

2.6. What do I do if "java.lang.IllegalStateException: Request cannot be executed; I/O reactor status: STOPPED" is displayed when I use Tablestore SDK for Java?

Problem description

A similar output is displayed when I use Tablestore SDK for Java:

```
java.lang.IllegalStateException: Request cannot be executed; I/O reactor status: STOPPED
```

Cause

shutDown is called to shut down the OTSClient instance and the internal I/O reactor. If the OTSClient instance is called to read and write data, the error occurs.

Solution

Check whether the OTSClient instance is shut down.

2.7. What do I do if Invalid date format is reported when I use Tablestore SDK for Java?

Problem description

Runtime environment: JDK 8

A similar output is displayed when I use Tablestore SDK for Java:

```
 [Error Code]:OTSParameterInvalid, [Message]:Invalid date format: Wed, 18 May 2016 08:32:51
+00:00.
```

Cause

The version of Joda-Time on which Classpath depends is earlier than the required version. Similar errors occur when you run Joda-Time of an early version in JDK 8.

Solution

Update the ots-public version to 2.2.4. If your project also depends on the Joda-Time library, update the version of ots-public to 2.9.

2.8. What do I do if PB library conflicts occur when I use Tablestore SDK for Java?

This topic describes the problem description, cause, and solution of PB library conflicts that occur when you use Tablestore SDK for Java.

Problem description

A similar output is displayed when I use Tablestore SDK for Java:

```
Caused by: java.lang.UnsupportedOperationException: This is supposed to be overridden by su
bclassed
```

Cause

Tablestore SDK for Java V2.4.1 depends on Protobuf library 2.4.1 and httpasyncclient 4.0.2, which may conflict with the same built-in library of your application.

Solution

Add the following dependencies to the pom.xml file in the Maven project:

> **Note** classifier is jar-with-dependencies, which packages the HttpClient and Protobuf dependencies by using rename package to remove the dependency on HttpClient and Protobuf.

```
<dependency>
        <groupId>com.aliyun.openservices</groupId>
        <artifactId>tablestore</artifactId>
        <version>Your current version</version>
        <classifier>jar-with-dependencies</classifier>
        <exclusions>
            <exclusion>
                <groupId>com.google.protobuf</groupId>
                <artifactId>protobuf-java</artifactId>
            </exclusion>
            <exclusion>
                <groupId>org.apache.httpcomponents</groupId>
                <artifactId>httpasyncclient</artifactId>
            </exclusion>
        </exclusions>
    </dependency>
```

2.9. What do I do if OTSUnsupportOperation is reported when I use Tablestore SDK?

Problem description

A similar output is displayed when I call `syncClient.createTable(request)`:

```
Caused by: [ErrorCode]:OTSUnsupportOperation, [Message]:Unsupported operation: 'CreateTable
'.
```

Cause

Tablestore SDK later than V4.0.0 is used to access tables created by using Tablestore SDK V2.x.x or earlier.

Solution

Use Tablestore SDK V2.x.x.

```
<dependency>
    <groupId>com.aliyun.openservices</groupId>
    <artifactId>ots-public</artifactId>
    <version>2.2.5</version>
</dependency>
```

2.10. What do I do if OTSParameterInvalid is reported when I use BatchWriteRow to submit 100 data entries at a time?

Problem description

A similar output is displayed when I use BatchWriteRow to submit 100 data entries at a time:

```
ErrorCode: OTSParameterInvalid, ErrorMessage: The input parameter is invalid.
```

Cause

The possible cause is that one batch operation cannot be repeatedly performed on the same row. If the operation is repeatedly performed on the same row, an error occurs.

Solution

Change 100 to 1 to submit one data entry each time. Keep other code unchanged.

2.11. Why do 5xx-related errors occur when I use Tablestore?

Some users may encounter 5xx-related errors when they use Tablestore. The following table lists the HTTP status codes.

HTTP status code	Error code	Error message
503	OTSPartitionUnavailable	The partition is not available.
503	OTSServerUnavailable	Server is not available.
503	OTSServerBusy	Server is busy.
503	OTSTimeout	Operation timeout.

The possible cause is that Tablestore is a distributed NoSQL service. The server implements load balancing based on the amount of and access to data in each partition. This way, data and access can be seamlessly scaled and distributed to multiple servers.

The following figure shows that Tablestore distributes data to different partitions base on the order of the first primary key column. Different partitions are distributed to different racks across servers to provide data reading and writing services.

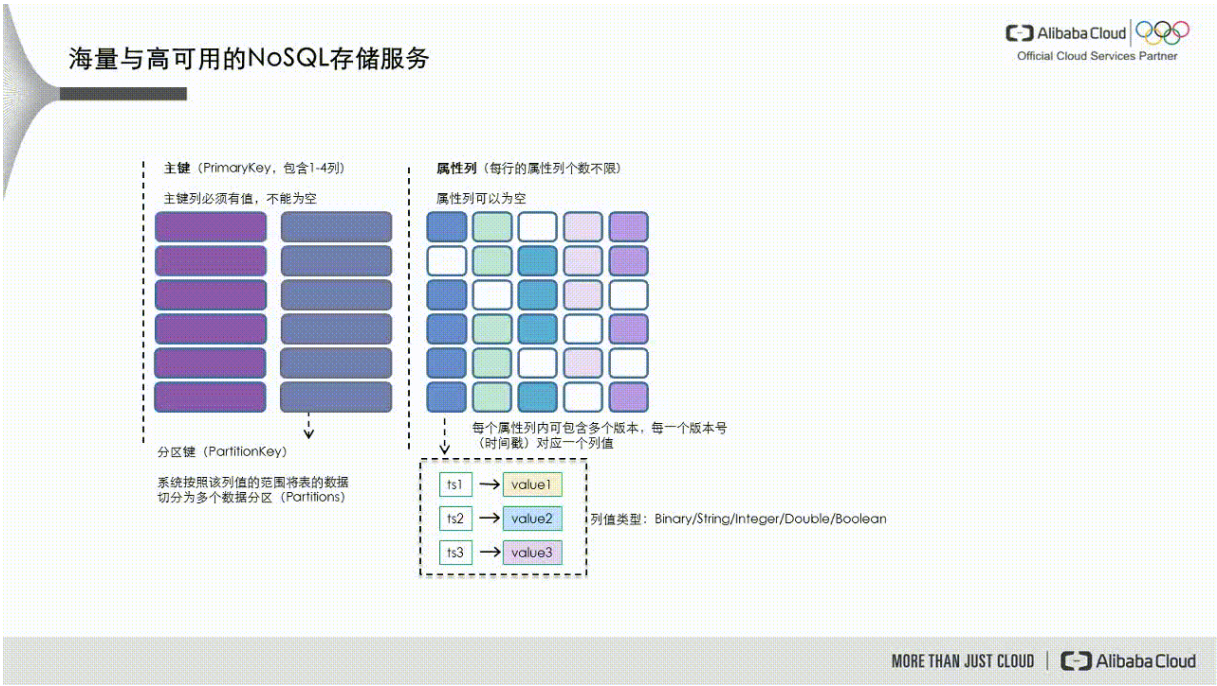

The dynamic load balancing mechanism of Tablestore can detect whether a partition such as P1 in the following figure stores large amounts of data or processes too frequent access. In this case, the partition is split into two partitions P1 and P5. The two partitions are distributed to the node that supports lower loads.

Tablestore uses the preceding automatic load balancing mechanism to store large amounts of data, process highly concurrent access, and implements automatic scalability without human intervention during the whole process. Only one partition is created when a table is created, which provides limited concurrent read and write capabilities. The automatic load balancing mechanism delivers latencies. To address these issues, contact Alibaba Cloud engineers to split a table into multiple partitions in advance.

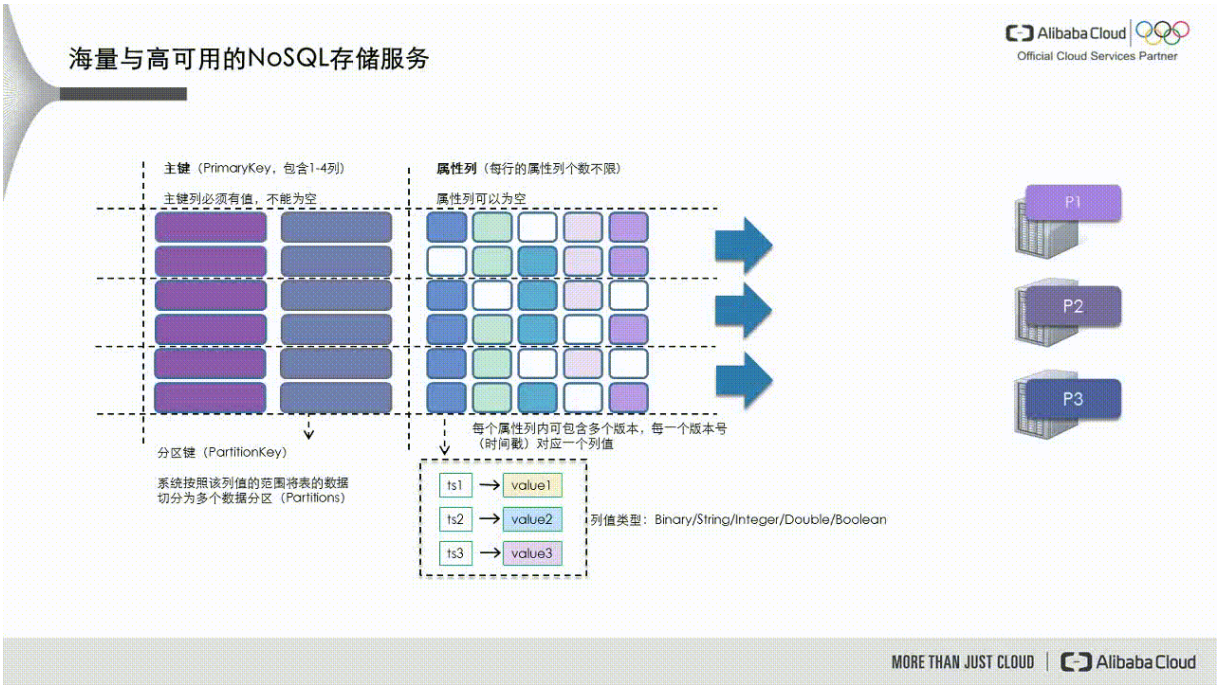

Tablestore uses a shared storage mechanism. Partitions are the logical storage unit of Tablestore. When Tablestore implements load balancing, data is not migrated and only metadata of tables is modified. When the metadata changes, partitions may become unavailable for a short period of time to ensure data consistency. **In most cases, this period ranges from several hundred milliseconds to several seconds when the partitions are under heavy load.** The preceding error may occur when read and write operations are performed on the partitions during this period of time. If this error occurs, retry the operations. Tablestore SDK provides default retry policies. You can specify retry policies when you initialize an OTSClient instance.

Tablestore uses the API that complies with the standard Restful protocol. Due to the uncontrollability of the network environment, we recommend that you add a retry policy for all reading and writing operations to respond to network errors for fault tolerance.

> **Note** Data read by calling BatchWriteRow and BatchGetRow may be distributed to a thousand tables or multiple partitions of a table. A partition may be being split when operations are called on the table. Therefore, the operations as a whole is not atomic. To ensure that each single-row operation is atomic, check getFailedRows() in the response for failed single-row operations when HTTP status code 200 is returned.

2.12. How can I obtain multiple rows of data if only one primary key column is set?

For more information about how to use GetRange to query multiple rows of data, see GetRange. For more information about specific code examples, visit GitHub.

2.13. How can I paginate query results?

Tablestore is a distributed storage system. You can use multiple methods to paginate query results. This topic describes how to paginate query results in detail.

Paginate query results for tables

If you create only Tablestore tables without creating search indexes, you can use one of the following methods to paginate query results:

- Use next tokens. The response to each request of the GetRange operation contains a next token. The next token can be used in the next request to read remaining data by page.
- Use GetRangeIterator. Use the iterator.next() method to obtain next piece of data.
- offset is not supported to paginate query results.
- The total number of rows and the total number of pages cannot be obtained.

Paginate query results for search indexes

If you create a search index for a table, use the following method to paginate query results:

- Use offset and limit. The sum of the offset and limit values cannot exceed 10000. If the sum of the offset and limit values exceed 10000, use next tokens to paginate query results.
- Use next tokens. The response to each request of the GetRange operation contains a next token. The next token can be used in the next request to read remaining data by page.
- Use SearchIterator. Use the iterator.next() method to obtain next piece of data.

- You can obtain the total number of rows and the total number of pages (Total number of pages = Total number of rows/Value of limit). To obtain the total number of rows, set getTotalCount in the request to true. If getTotalCount is set to true, more resources are consumed and performance may deteriorate.

Examples for tables

The following code provides an example on how to call an operation to paginate query results by setting offset and limit.

```
/**
     * Specify the range of data to query and return specified rows of data by setting the
page size and offset.
     */
    private static Pair<List<Row>, RowPrimaryKey> readByPage(OTSClient client, String table
Name,
            RowPrimaryKey startKey, RowPrimaryKey endKey, int offset, int pageSize) {
        Preconditions.checkArgument(offset >= 0, "Offset should not be negative.") ;
        Preconditions.checkArgument(pageSize > 0, "Page size should be greater than 0.") ;
        List<Row> rows = new ArrayList<Row>(pageSize);
        int limit = pageSize;
        int skip = offset;
        RowPrimaryKey nextStart = startKey;
        // If the amount of data to query is large, only part of data is returned for one r
equest. To query all required data, use multiple range query requests.
        while (limit > 0 && nextStart ! = null) {
            // Create query parameters for GetRange.
            // Note that startPrimaryKey must be set to the position where last reading sto
ps. This way, you can send multiple range query requests to read remaining data.
            RangeRowQueryCriteria criteria = new RangeRowQueryCriteria(tableName);
            criteria.setInclusiveStartPrimaryKey(nextStart);
            criteria.setExclusiveEndPrimaryKey(endKey);
            // Set an appropriate limit value. limit specifies the total number of rows to
return (a maximum of all rows of data on a page) from the offset value.
            criteria.setLimit(skip + limit);
            GetRangeRequest request = new GetRangeRequest();
            request.setRangeRowQueryCriteria(criteria);
            GetRangeResult response = client.getRange(request);
            for (Row row : response.getRows()) {
                if (skip > 0) {
                    skip--; // The number of rows of data to filter before the offset value
. The data is filtered on the client after the data is read.
                } else {
                    rows.add(row);
                    limit--;
                }
            }
            // Set the position from which to read data next time.
            nextStart = response.getNextStartPrimaryKey();
        }
        return new Pair<List<Row>, RowPrimaryKey>(rows, nextStart);
    }
```

The following code provides an example on how to use the preceding operation to sequentially read all data within a specified range by page:

```
  private static void readByPage(OTSClient client, String tableName) {
        int pageSize = 8;
        int offset = 33;
        RowPrimaryKey startKey = new RowPrimaryKey();
        startKey.addPrimaryKeyColumn(COLUMN_GID_NAME, PrimaryKeyValue.INF_MIN);
        startKey.addPrimaryKeyColumn(COLUMN_UID_NAME, PrimaryKeyValue.INF_MIN);
        RowPrimaryKey endKey = new RowPrimaryKey();
        endKey.addPrimaryKeyColumn(COLUMN_GID_NAME, PrimaryKeyValue.INF_MAX);
        endKey.addPrimaryKeyColumn(COLUMN_UID_NAME, PrimaryKeyValue.INF_MAX);
        // Read the first page from the 33th line.
        Pair<List<Row>, RowPrimaryKey> result = readByPage(client, tableName, startKey, end
Key, offset, pageSize);
        for (Row row : result.getKey()) {
            System.out.println(row.getColumns());
        }
        System.out.println("Total rows count: " + result.getKey().size());
        // Read all data within the range and paginate query results in sequence.
        startKey = result.getValue();
        while (startKey ! = null) {
            System.out.println("============= start read next page ==============");
            result = readByPage(client, tableName, startKey, endKey, 0, pageSize);
            for (Row row : result.getKey()) {
                System.out.println(row.getColumns());
            }
            startKey = result.getValue();
            System.out.println("Total rows count: " + result.getKey().size());
        }
    }
```

Examples for search indexes

For more information, see Sorting and pagination.

2.14. How can I add one to the values in a specific column?

The following code provides an example on how to add one to the values in a specified column:

```
row = getRow(primary_key, 'col') // Read the values in the column.
old_value = row['col'] // Record the original values in the column.
row['col'] = old_value + 1 // Calculate new values.
updateRow(row, condition: row['col'] == old_value) // Use conditional update when you write
new values. Condition: The data is written when the original values remain unchanged.
```

2.15. What are the code examples used to perform ListTable in Tablestore SDK for Python?

Examples

```
import time
import logging
import unittest
from ots2 import *
ENDPOINT = "https://xxx.cn-hangzhou.ots.aliyuncs.com";
ACCESSID = "xxx";
ACCESSKEY = "xxx";
INSTANCENAME = "xxx";
ots_client = OTSClient(ENDPOINT, ACCESSID, ACCESSKEY, INSTANCENAME)
list_response = ots_client.list_table()
print u'instance table:'
for table_name in list_response:
    print table_name
```

Note For more information about how to install and use Tablestore SDK for Python, see Preface in Tablestore SDK for Python.

Tablestore SDK for Python does not provide import at the beginning of each snippet for you to run code. If you do not add import to the beginning of each snippet when you run the code, a similar output is displayed:

```
Traceback (most recent call last):
  File "listtable.py", line 6, in
    ots_client = OTSClient(ENDPOINT, ACCESSID, ACCESSKEY, INSTANCENAME)
NameError: name 'OTSClient' is not defined
```

Add import to run the code.

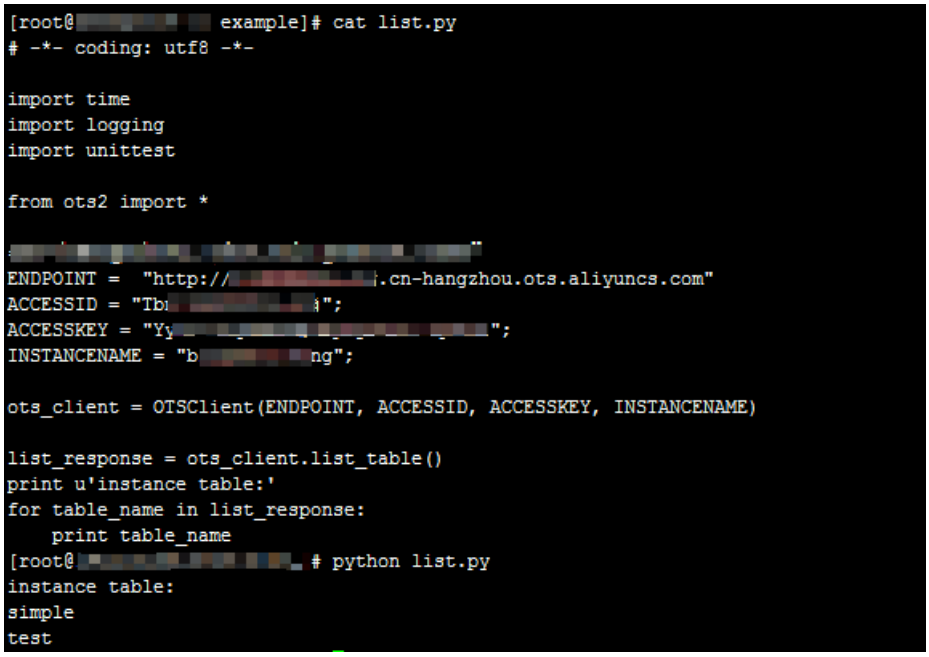

2.16. What do I do if Checksum mismatch is reported when I use Tablestore SDK for PHP?

This topic describes the problem description, cause, and solution of the Checksum mismatch exception when you use Tablestore SDK for PHP.

Problem description

The following exception occurs when you use Tablestore SDK for PHP by using PHP 5.6 in Windows:

```
Fatal error: Uncaught exception 'AliyunOTSOTSClientException' with message 'Checksum mismat
ch. expected:120,actual:-48'
```

Cause

Tablestore uses 64-bit integers. However, PHP versions earlier than PHP 7 in Windows do not fully support 64-bit integers and can only process 64 bit integers as strings. Therefore, Tablestore SDK for PHP does not support PHP versions earlier than PHP 7 in Windows.

Solution

When you use Tablestore SDK for PHP in Windows, ensure that the PHP version is PHP 7 or later. We recommend that you use PHP 7 to obtain the optimal performance.

You can use phpinfo() to check the value of Architecture in PHP configuration information to determine whether your PHP version supports 64-bit integers.

- If the value of Architecture is X86, the current PHP version supports only 32-bit integers. You must upgrade your PHP version to PHP 7 or later.
- If the value of Architecture is X64, the current PHP version supports 64-bit integers and can meet your

requirement.

2.17. What do I do if the "Request denied because this instance can only be accessed from the binded VPC" exception occurs when I use Tablestore SDK?

This topic describes the problem description, cause, and solution of the "Request denied because this instance can only be accessed from the binded VPC" exception when you use Tablestore SDK.

Problem description

The following exception occurs when you use SDKs to access an instance:

```
[ErrorCode]:OTSAuthFailed, [Message]:Request denied because this instance can only be acces
sed from the binded VPC.
```

Cause

After you bind a VPC to a Tablestore instance, you can access the instance only from the address of the bound VPC.

Solution

After you bind a VPC to a Tablestore instance, you must use the address of the bound VPC to access the instance by using SDKs.

Log on to the Tablestore console. Go to the **Network Management** tab to obtain the address of the bound VPC.