

Alibaba Cloud

Tablestore Pricing

Document Version: 20200828

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 <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1. Billing overview	05
2. Storage usage	07
3. Search index	10
4. Billing rules	14
5. Overdue payments, renewals, and upgrades	19
6. Billing examples	21
7. FAQ	22

1. Billing overview

Tablestore charges resources for each instance. This topic describes the pricing, billing methods, and billing items of Tablestore.

Pricing



For more information about pricing, see [Tablestore Pricing](#).

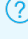
Billing methods


Tablestore charges resources based on the pay-as-you-go billing method. This method indicates that Tablestore charges resources you have used on an hourly basis.

Billing items

The following table lists the billing items of Tablestore.

Billing item		Billing method
Storage usage		The total storage usage charged based on the average storage usage per hour in a billing cycle.
Read throughput	Reserved read throughput	<p>The total reserved read CUs of all tables charged based on the average reserved read CUs per hour in a billing cycle.</p> <p> Note Reserved read CUs are available only for high-performance instances. For more information about instances, see Instance.</p>
	Additional read throughput	The total additional read CUs of all tables charged based on the average additional read CUs per second in a billing cycle.
Write throughput	Reserved write throughput	<p>The total reserved write CUs of all tables charged based on the average reserved write CUs per hour in a billing cycle.</p> <p> Note Reserved write CUs are available only for high-performance instances. For more information about instances, see Instance.</p>
	Additional write throughput	The total additional write CUs of all tables charged based on the average additional write CUs per second in a billing cycle.

Billing item		Billing method
Internet outbound traffic		<p>The total Internet outbound traffic generated when applications access Tablestore over HTTP.</p> <div data-bbox="598 392 1385 705"><p> Note</p><ul style="list-style-type: none">• Tablestore charges Internet outbound traffic instead of inbound traffic and internal network traffic.• When access fails, Tablestore returns access failure information, which also generates outbound traffic.• Access between different regions is also billed as Internet access.</div>

 **Note**

- You are also charged when you use search indexes. For more information about the billing, see [Search index](#).
- You are also charged when you use global secondary indexes. For more information about the billing, see [Billing rules](#).

2.Storage usage

Tablestore charges storage usage on an hourly basis based on the total size of data. Tablestore calculates the total size of data based on predetermined intervals and calculates the average data size per hour.

 **Note** For more information about pricing, see [Tablestore pricing](#).

The following section describes how to calculate the size of a single row and tables.

Calculate the size of a single row

Each row of data in Tablestore uses storage space. After max versions or time to live (TTL) is configured, data of each version includes the version number (8 bytes), column name, and data value.

The size of a single row is calculated based on the formula: $\text{Size of a single row} = \text{Size of primary key columns} + \text{Size of all attribute columns}$. For more information, see [Primary keys and attributes](#).

The sizes of primary key columns and all attribute columns are calculated based on the formulas:

- $\text{Size of primary key columns} = \text{Length of all primary key column names} + \text{Size of values in primary key columns}$
- For more information about the size of all attribute columns, see the examples on how to calculate the size of a single row and a table in this topic.

The following table describes how to calculate the size of values.

Date type	Number of bytes
STRING	The number of bytes used by characters encoded in UTF-8. Tablestore allows empty strings. The size of empty strings is 0.
INTEGER	8.
DOUBLE	8.
BOOLEAN	1.
BINARY	The number of bytes used by binary data.

Example: Calculate the size of a row.


The primary column name is ID. Other columns are attribute columns.

ID	Name	Length	Comments
1.	timestamp = 1466676354000, value = 'zhangsan'	timestamp = 1466676354000, value = 20	timestamp = 1466676354000, value = String (100 Bytes) timestamp = 1466679954000, value = String (150 Bytes)

In the row, two versions are available for the value in the Comments column.

- When MaxVersions is set to 2 and TTL is set to 2592000:


The size of an attribute column is calculated based on the formula: $\text{Size of an attribute column} = (\text{Length of the attribute column name} + 8) \times \text{Number of valid versions} + \text{Total size of values in the attribute column}$.

 **Note** When max versions is set to a value greater than 1 or TTL is set to -1, each version number uses 8 bytes. The version number is represented by a timestamp. For more information, see [Data versions and time to live](#).

The size of the row is the sum of 10, 20, 22, and 282. Details:

- Size of the primary key column: $\text{len}('ID') + \text{len}(1) = 10$ bytes
 - Size of the Name attribute column: $(\text{len}('Name') + 8) * 1 + \text{len}('zhangsan') = 20$ bytes
 - Size of the Length column: $(\text{len}('Length') + 8) * 1 + \text{len}(20) = 22$ bytes
 - Size of the Comments column: $(\text{len}('Comments') + 8) * 2 + 100 + 150 = 282$ bytes
- When MaxVersions is set to 1 and TTL is set to -1:

The size of an attribute column is calculated based on the formula: $\text{Size of an attribute column} = \text{Length of the attribute column name} + \text{Total size of the value in the attribute column}$.

 **Note**

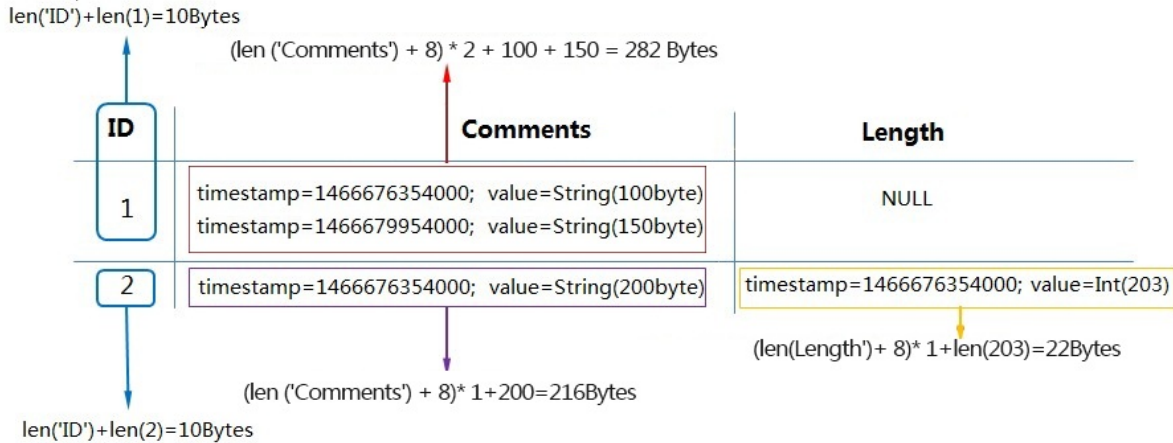
- When MaxVersions is set to 1 and TTL is set to -1, the version number consumes no bytes.
- Although the value in the Comments column has two versions, only the latest version is calculated because MaxVersions is set to 1.

The size of the row is the sum of 10, 12, 14, and 158, which equals 194 bytes. Details:

- Size of the primary key column: $\text{len}('ID') + \text{len}(1) = 10$ bytes
- Size of the Name attribute column: $\text{len}('Name') + \text{len}('zhangsan') = 12$ bytes
- Size of the Length attribute column: $\text{len}('Length') + \text{len}(20) = 14$ bytes
- Size of the Comments attribute column: $\text{len}('Comments') + 150$ (bytes) = 158 bytes

Calculate the size of a table

The size of a table consists of sizes of all rows of data in the table. Example: A table contains the ID primary key column and attribute columns. When MaxVersions is set to 2 and TTL is set to -1: The following figure shows how to calculate the size of the table.



- The size of the row whose value in the ID column is 1 = 10 (size of the primary key column) + 282 (size of the values of the two versions in the Comments column). In total, the size of the row whose value in the ID column is 1 is 292 bytes.
- The size of the row whose value in the ID column is 2 = 10 (size of the primary key column) + 216 (size of the value in the Comments column) + 22 (size of the value in the Length column). In total, the size of the row whose value in the ID column is 2 is 248 bytes.
- The size of the table is the sum of 292 and 248, which equals 540 bytes.

If the size of the data in the table remains unchanged within one hour, Tablestore charges storage fees based on 540 bytes. Tablestore does not impose limits on the size of data that can be stored in a table. Tablestore charges storage fees for actual data stored in the table.

Note Tablestore asynchronously deletes expired data in each partition and the additional versions when the max versions value is exceeded. The time used to delete data is related to the total size of the data, which is within 24 hours. Data written to a partition after a data delete operation is performed is charged after the next data delete operation is performed.

3. Search index

Search index-based queries require extra space to store indexed data and consume read throughput. This topic describes the billing items and billing formulas of search indexes.

Note

- The metering and billing of indexes are independent of base tables.
- The price of each billing item of a search index is the same as that of a high-performance instance.

Billing items

Billing item	Billing method	Description
Data size	Pay-as-you-go	<p>Unit: GB. Billed size is rounded up to the next GB.</p> <p>Tablestore charges you for the total volume of indexed data on an hourly basis. The utilization of system resources varies with field types and index types. Therefore, indexed data is billed based on the volume of data compressed after you create indexes. The raw data volume in a base table does not affect this billing.</p>

Billing item		Billing method	Description
Read throughput	Reserved read throughput	Pay-as-you-go	<p>Unit: CU</p> <p>Tablestore specifies a reserved read throughput based on the indexed data size. The charges of reserved read throughput are calculated based on the following operations:</p> <ul style="list-style-type: none"> • When you create a search index, Tablestore reads data from a base table, which consumes read throughput. • Tokenization during the creation of a search index also consumes read throughput. Fees incurred from tokenization are included in the fees for the reserved read throughput. • To ensure index and query performance, some indexed data is loaded to the memory in advance and remains in memory. Such data consumes read throughput. Fees incurred from these operations are also included in the fees for the reserved read throughput. <p>The minimum fees are based on the reserved read throughput value. For example, assume that the reserved read throughput is 10,000 CUs for an index. Each index query reads 10 rows, and the size of each row is less than 4 KB. When the number of queries per second (QPS) is less than 1,000, the actual consumed read throughput is less than the reserved read throughput, and no extra fees are charged.</p> <p>Calculation of reserved read throughput: The reserved read throughput is proportional to the size and number of rows of the indexed data. For example, 1 GB or 2 million rows of indexed data corresponds to a reserved read throughput of 10 CUs. When the reserved read throughput values corresponding to the data size and the number of rows are different, the system uses the larger one as the reserved read throughput.</p> <div style="background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p>Note</p> <ul style="list-style-type: none"> • The maximum reserved read throughput for Tablestore is 100,000 CUs. • When the data size is less than 200 MB and the number of rows is less than 400,000, the reserved read throughput for Tablestore must be 20 CUs, which is suitable for tests that involve a small amount of data. When the data size is greater than or equal to 200 MB or the number of rows is greater than or equal to 400,000, the reserved read throughput must be greater than or equal to 100 CUs. </div>

Billing item		Billing method	Description
	Additional throughput	Pay-as-you-go	Unit: CU The portion of the actual consumed read throughput that exceeds the reserved read throughput is charged as additional read throughput.
Internet outbound traffic		Pay-as-you-go	Unit: GB Fees for Internet outbound traffic.

Billing formulas


The following table describes how to calculate the data size and read throughput of a search index.

Billing item	Formula	Description
Data size	$f(Size) = \frac{Size}{1GB} * USD\ 0.00030/GB/Hour$	The Size parameter indicates the size of a compressed index.
Read throughput	<p>Reserved read CUs per index:</p> $ReservedCuPerIndex(Size, Rows) = \min(\max(\frac{Size}{0.2GB}, \frac{Rows}{0.4m}, 10) * 2, 100000)$ <p>Read CUs per query:</p> $CuPerQuery = Ceil(\frac{ReturnRowSize}{4KB}) * ReturnRowCount$	<ul style="list-style-type: none"> The Size parameter indicates the size of a compressed index. The Rows parameter indicates the total number of rows in an index, except for child rows in a nested query. The ReturnRowSize parameter indicates the size of returned rows. The ReturnRowCount parameter indicates the number of returned rows.

Billing examples

Storage	Number of rows	Billing

Storage	Number of rows	Billing
8 GB	9 million rows	<ul style="list-style-type: none"> Storage fees: $8 \text{ GB} \times \text{USD } 0.00030/\text{GB}/\text{hour} = \text{USD } 0.0024/\text{hour}$ Calculation of reserved read throughput: 8 GB of data corresponds to 80 CUs and 9 million rows correspond to 45 CUs. Fees of reserved read throughput: $80 \text{ CUs} \times \text{USD } 0.0002/\text{CU}/\text{hour} = \text{USD } 0.016/\text{hour}$ Total fees: $\text{USD } 0.0024/\text{hour} + \text{USD } 0.016/\text{hour} = \text{USD } 0.0184/\text{hour}$ <p>The portion of the actual consumed read throughput that exceeds the reserved read throughput is charged as additional read throughput. The billing method of the Internet downstream traffic is the same as that of the primary table.</p>
100 GB	300 million rows	<ul style="list-style-type: none"> Storage fees: $100 \text{ GB} \times \text{USD } 0.00030/\text{GB}/\text{hour} = \text{USD } 0.03/\text{hour}$ Calculation of reserved read throughput: 100 GB of data corresponds to 1,000 CUs and 300 million rows correspond to 1,500 CUs. The reserved read throughput is 1,500 CUs. Fees of reserved read throughput: $1,500 \text{ CUs} \times \text{USD } 0.0002/\text{CU}/\text{hour} = \text{USD } 0.3/\text{hour}$ Total fees: $\text{USD } 0.03/\text{hour} + \text{USD } 0.3/\text{hour} = \text{USD } 0.33/\text{hour}$ <p>The portion of the actual consumed read throughput that exceeds the reserved read throughput is charged as additional read throughput. The billing method of the Internet downstream traffic is the same as that of the primary table.</p>
30 TB	10 billion rows	<ul style="list-style-type: none"> Storage fees: $30,000 \text{ GB} \times \text{USD } 0.00030/\text{GB}/\text{hour} = \text{USD } 9/\text{hour}$ Calculation of reserved read throughput: 30 TB of data corresponds to 300,000 CUs and 10 billion rows correspond to 50,000 CUs. Although the former CU value is larger (300,000 CUs), it exceeds the maximum reserved read throughput of 100,000 CUs. Therefore, the reserved read throughput is 100,000 CUs. Fees of reserved read throughput: $100,000 \text{ CUs} \times \text{USD } 0.0002/\text{CU}/\text{hour} = \text{USD } 20/\text{hour}$ Total fees: $\text{USD } 9/\text{hour} + \text{USD } 20/\text{hour} = \text{USD } 29/\text{hour}$ <p>The portion of the actual consumed read throughput that exceeds the reserved read throughput is charged as additional read throughput. The billing method of the Internet downstream traffic is the same as that of the primary table.</p>

 **Note** The price in the preceding table is for reference only. For more information, visit the Tablestore console.

4. Billing rules

To use secondary indexes, index tables are needed. Therefore, additional storage space is required to store index tables. When the system inserts data to a primary table, it may also need to write the index tables created on the primary table at the same time. During this process, read and write CUs are consumed. This topic describes the billing rules for secondary indexes.

Note Capacity units (CUs) are read and write throughput units. They are the smallest units used to measure the costs of read and write operations. For example, when the system reads 4 KB from one row per second, one read CU is consumed.

To use secondary indexes, index tables are needed. Therefore, additional storage space is required to store index tables. When the system inserts data to a primary table, it may also need to write the index tables created on the primary table at the same time. During this process, read and write CUs are consumed.

Secondary index billing includes the following parts: the number of read and write CUs consumed to write index tables, the amount of data stored in the index tables, and the amount of data that is read from the index tables.

Billing item	Description
Data storage	The storage space used to store a primary table and its index tables.
Read CUs consumed to write index tables	The number of CUs that are consumed by read operations to delete, insert, or update index rows.
Write CUs consumed to write index tables	The number of CUs that are consumed to insert or update index rows.
CUs consumed by regular read operations	The number of CUs that are consumed to read data from a primary table or index tables using an API.
CUs consumed by regular write operations	The number of CUs that are consumed to insert data to a primary table using an API.

Billing rules for storing, writing, and reading an index table:

- The billing rules for storing and reading an index table are the same as those of a primary table. For more information, see [Billing items and pricing](#).
- CUs are consumed based on the following rules when the system writes an index table:
 - Write CUs are consumed only when an index row is inserted or updated.
 - Read CUs are consumed when an index row is deleted, updated, or inserted. The number of read CUs equals the amount of data read from the corresponding indexed columns in the primary table.

Calculate the number of read CUs consumed to write index tables

When you create secondary indexes on the primary table, read CUs are consumed based on the following rules:


- When you use the PUT operation to insert a data row to the primary table:

- The PUT operation does not insert data to the indexed attribute columns in the primary table, which means that no index row is inserted. In this case, one read CU is consumed.
- The PUT operation inserts data to the indexed attribute columns in the primary table, which means that new index rows are inserted. In this case, one read CU is consumed.
- When you use the PUT operation to overwrite a row in the primary table:
 - The PUT operation does not update the indexed attribute columns in the primary table. In this case, one read CU is consumed.
 - The PUT operation updates the indexed attribute columns in the primary table. In this case, the read CUs are consumed as follows:

Divide the total amount of data read from the indexed attribute columns by four, excluding primary key columns. The number of consumed CUs equals the calculated value rounded up to the nearest integer. If the total amount is 0 KB, one CU is consumed.
- When you use the UPDATE operation to insert a data row to the primary table:
 - If the UPDATE operation does not insert data to the indexed columns in the primary table, no read CU is consumed.
 - If the UPDATE operation inserts data to the indexed columns in the primary table, one read CU is consumed.
- When you use the UPDATE operation to update a row in the primary table:
 - If the UPDATE operation does not insert data to the indexed attribute columns in the primary table, no read CU is consumed.
 - If the UPDATE operation inserts data to the indexed attribute columns in the primary table, read CUs are consumed based on the following rules:

Divide the total amount of data read from the indexed columns by four, excluding the primary key columns. The number of consumed CUs equals the calculated value rounded up to the nearest integer. If the total amount is 0 KB, one CU is consumed.
- When you use the Delete operation to delete a row in the primary table, read CUs are consumed based on the following rules:

Divide the total amount of data read from the indexed columns by four, excluding the primary key columns. The number of consumed CUs equals the calculated value rounded up to the nearest integer. If the total amount is 0 KB, one CU is consumed.
- If the primary table uses primary key auto increment, inserting data to the primary table does not consume any read CUs. Updating a row in a primary table that uses primary key auto increment consumes read CUs. CUs are calculated based on the same rules as those of the UPDATE operation.

 **Note** We recommend that you use primary key auto increment to insert data to a primary table to decrease the number of CUs that are consumed by index tables.

For primary tables that do not use primary key auto increment, one read CU is consumed if a read operation is performed on the indexed columns, even if no data is retrieved. For primary tables that use primary key auto increment, no read operation is performed on the indexed columns when you insert data. Therefore, no read CU is consumed.

Calculate the number of write CUs

When you insert data to the primary table and create secondary indexes, write CUs are consumed. Write CUs are consumed based on the following rules:

- If you insert a row to the primary table and no data in the index table is updated, no write CUs are consumed.
- If you insert a row to the primary table and a new index row is inserted to the index table, write CUs are consumed. The number of the write CUs is determined by the size of the inserted index row.
- If you insert a row to the primary table and an index row is deleted from the index table, write CUs are consumed. The number of the write CUs is determined by the size of the deleted index row.
- If you insert a row to the primary table and an index row in the index table is updated, write CUs are consumed. The number of the write CUs is determined by the size of the updated index row.
- If you insert a row to the primary table, an index row is deleted from the index table, and another index row is inserted to the index table, write CUs are consumed. The number of the write CUs is determined by the total size of the deleted and inserted index rows.

The detailed rules are as follows:

- When you use the PUT operation to insert a data row to a primary table:
 - The PUT operation does not insert data to the indexed attribute columns in the primary table, which means that no index row is inserted. In this case, no read CU is consumed.
 - The PUT operation inserts data to the indexed attribute columns in the primary table, which means that new index rows are inserted. The write CUs consumed for each index table are calculated as follows:

Divide the total amount of data in the inserted index row by four. The number of consumed CUs equals the calculated value rounded up to the nearest integer.
- When you use the PUT operation to overwrite a row in the primary table:
 - The PUT operation only updates the indexed primary key columns in the primary table. In this case, no write CUs are consumed.
 - The PUT operation updates the indexed columns in the primary table. The write CUs are consumed based on the following rules:

All indexes updated by the PUT operation consume a certain number of write CUs, except sparse indexes.
- When you use the UPDATE operation to insert a data row to the primary table:
 - If the UPDATE operation does not insert data to the indexed columns in the primary table, no write CUs are consumed.
 - If the UPDATE operation inserts data to the indexed columns in the primary table, the write CUs consumed for each index table are calculated as follows:
 - If the UPDATE operation inserts a new index row, write CUs are consumed. Divide the total size of the data in the index row by four. The number of consumed CUs equals the calculated value rounded up to the nearest integer.
 - If no index row is inserted, no write CUs are consumed.
- When you use the UPDATE operation to update a row in the primary table:
 - If the UPDATE operation does not update the indexed attribute columns, no write CUs are

consumed.

- If the UPDATE operation updates the indexed attribute columns, write CUs consumed for each index table are calculated based on the following rules:
 - If the index table already contains an index row created based on the row to be updated, delete CUs are consumed. The number of the delete CUs is determined by the size of the indexed primary keys in the deleted index row.
 - If a new index row is inserted based on the updated row, write CUs are consumed. The number of the write CUs is determined by the size of the indexed primary keys in the inserted index row.
 - If the UPDATE operation only updates the attribute data in the existing index row but no new index row is inserted, update CUs are consumed.

Divide the total amount of data in the index row by four. The number of consumed CUs equals the calculated value rounded up to the nearest integer.

- When you use the DELETE operation to delete a row in the primary table, write CUs are consumed based on the following rules:

If an index table already contains an index row created based on the row to be deleted, write CUs are consumed. Divide the total amount of the data in the corresponding indexed columns by four, excluding the primary key columns. The consumed write CUs equal the calculated value rounded up to the nearest integer.

- If you insert data to a primary table that uses primary key auto increment, write CUs are consumed. The write CUs are calculated based on the same rules as those of the PUT operation. If you update a row in a primary table that uses primary key auto increment, write CUs are consumed. The write CUs are calculated based on the same rules as those of the UPDATE operation.

Measure index table size

The size of an index table is measured based on the same rule as that of a primary table. The size of an index table equals the total size of all rows. The total size of the rows equals the total size of primary keys and attribute data. For more information, see [Data storage](#).

Calculate the number of CUs consumed to read an index table

When you use an SDK, the console, or other methods, such as a DLA, to read an index table, read CUs are consumed. The number of read CUs are calculated based on the same rules as those of reading a primary table.

Examples


The following example uses a primary table that has two index tables to describe how CUs are consumed under different conditions.

The primary table Table contains two primary key columns PK0 and PK1, and three predefined columns Col0, Col1, and Col2. Two index tables, Index0 and Index1, are created on the primary table. Index0 contains three primary keys Col0, PK0, and PK1 and one attribute column Col2. Index1 contains four primary keys Col1, Col0, PK0, and PK1, and no attribute columns. Use the UPDATE operation to update PK0 and PK1.

- If the target row does not exist in the primary table:
 - Updating Col3 does not consume read or write CUs.
 - Updating Col1 consumes the following CUs:
 - One read CU
 - No write CUs
 - Updating Col0 and Col1 consumes the following CUs:
 - One read CU
 - Index0 consumes write CUs. The number of the write CUs is determined by the total amount of data inserted to Col0, PK0, and PK1. Index1 consumes write CUs. The number of the write CUs is determined by the total amount of data inserted to Col0, Col1, PK0, and PK1.
- If the target row already exists in the primary table:
 - Updating Col3 does not consume read or write CUs.
 - Updating Col2 consumes the following CUs:
 - Read CUs are consumed. The number of the read CUs is determined by the amount of data read from Col0. If the UPDATE operation inserts data to Col0, one CU is consumed.
 - For Index0, if the UPDATE operation inserts data to Col0, Index0 does not consume write CUs. If the UPDATE operation updates the data in Col0, Index0 consumed write CUs. The number of the write CUs is determined by the total amount of data inserted to Col0, PK0, PK1, and Col2. Index1 does not consume write CUs.
 - Updating Col1 consumes the following CUs:
 - Read CUs are consumed. The number of the read CUs is determined by the amount of data read from Col0 and Col1. If the total amount is 0 KB, one CU is consumed.
 - Index0 does not consume write CUs. For Index1, if an index row is inserted, write CUs are consumed. The number of the write CUs is determined by the amount of data read from Col0 and inserted to Col1, PK0, and PK1. For Index1, if no data in Col0 is updated, no index row is inserted and no write CUs are consumed. If the data in Col0 and Col1 is updated, write CUs are consumed to delete the corresponding index row. The number of write CUs is determined by the total amount of data read from Col0, Col1, PK0, and PK1.

5. Overdue payments, renewals, and upgrades

This topic describes the overdue payment, renewal, and upgrade policies of Tablestore instances.

 **Notice** You may receive notifications if you have overdue payments. When this occurs, pay off all overdue bills to avoid instances being released. Note that your instances may be released at a system-selected time after the payment due date.

Billing method	Overdue payment	Renewal and upgrade
----------------	-----------------	---------------------

Billing method	Overdue payment	Renewal and upgrade
<p>Pay-as-you-go</p>	<p>Fees are calculated on an hourly basis. When your account balance is insufficient to cover the charges of the last billing cycle, you will have an overdue payment.</p> <p>When an overdue payment is generated, the system sends notifications to you based on the following situations:</p> <ul style="list-style-type: none"> You will not be affected by the service suspension if you top up your balance within 24 hours. Your Tablestore is automatically suspended if you fail to pay off all overdue bills within 24 hours. However, you will still be charged for the storage space resources are using. Consequently, the overdue amount will continue to increase. <p>Your Tablestore is automatically started if you top up your balance to pay off all overdue bills within 15 days after the overdue payment is generated.</p> <ul style="list-style-type: none"> If you fail to pay off all overdue bills within 15 days, you will be regarded as voluntarily discarding Tablestore. Data in your Tablestore instance may be deleted and deleted data cannot be recovered. 	<p>Pay-as-you-go instances are charged based on the service duration. You do not need to renew the instances. Instead, you need only to top up your balance in the Alibaba Cloud Management console.</p>


6. Billing examples

This topic provides billing examples to describes how Tablestore calculates fees.

Background information

A user in the United States creates a high-performance instance after they activate Tablestore. Data in the table of the instance supports queries per second (QPS) of 10000 and throughput of smaller than 4 KB (1 CU). The user wants to learn about how Tablestore calculates fees related to the table within one day.


Case analysis

 **Note** The unit price for Tablestore was released on the Alibaba Cloud official website on August 1, 2018. To remain updated on the unit price, log on to the Alibaba Cloud official website.

Billing item	Unit price for high-performance instances
Additional read throughput	USD 0.0030 per 10,000 CUs

Additional read throughput within the day is charged based on the formula:

Read throughput fees = $10000 \times 86400 / 10000 \times 0.0030$. A total of USD 259.2 is charged.

 **Note** When Tablestore calculates additional read/write throughput fees, the total number of consumed CUs is calculated. In this example, the total number of consumed CUs is 864 million CUs (10000×86400).

7.FAQ