

Table Store

Product Introduction

Product Introduction

Table Store is a NoSQL database service built on Alibaba Cloud's Apsara distributed file system that can store and access massive structured data in real time by allowing users to:

Organize data into instances and tables that can seamlessly scale using data partitioning and load balancing.

Shield applications from faults and errors that occur on the underlying hardware platform, providing fast recovery capability and high service availability.

Manage data with multiple backups using solid state disks (SSDs), enabling quick data access and high data reliability.

Furthermore, when using Table Store, you only pay for the resources you reserve and use. Factors such as cluster resizing, upgrades, and maintenance of database software and hardware are managed free of charge.

Table Store provides the following advantages:

Scalability

Dynamic adjustment of reserved read/write throughput

When creating a table, you can configure the reserved read/write throughput for an application based on business access conditions. Table Store schedules and reserves resources, based on the table's reserved read/write throughput, to minimize the resource usage costs. Then, Table Store can dynamically adjust the table's reserved read/write throughput based on the application.

Unlimited capacity

There is no limit to the amount of data stored in Table Store tables. If a table size increases, Table Store will adjust the data partitions for immediate storage space allocation to the increased table.

Data reliability

Table Store stores multiple data copies across different servers in different racks. If one backup fails, servers with copied data will immediately restore services, achieving zero data loss.

High availability

Through automatic failure detection and data migration, Table Store shields applications from both hardware and network-related faults to deliver high availability.

Ease of management

Table Store automatically manages complex tasks, such as the management of data partitions, software/hardware upgrades, configuration updates, and cluster resizing, allowing you to focus on growing your business.

Secure access platform

Table Store performs identity authentication for each application request, preventing unauthorized data access and ensuring data access security.

High consistency

Table Store ensures high consistency of writing data. Once a successful result is returned for a write operation, applications can read the latest data.

Flexible data models

Table Store tables do not require a fixed format. The column numbers of each row, and the value types in columns of the same name but different rows, can be different. Furthermore, Table Store supports multiple data types, such as Integer, Boolean, Double, String, and Binary.

Pay-As-You-Go


Table Store only charges fees based on the actual resources you have reserved and used.

Monitoring integration

The Table Store console provides real-time monitoring information, including the requests number per second and the average response latency.


Big data storage and analysis

Table Store provides low-cost, low-latency, and high-concurrency storage and online access of massive data, as well as incremental and full data tunnels, with support for SQL direct read and write on big data analysis platforms, such as MaxCompute. An efficient incremental streaming read interface is provided for easy computing of real-time data streams.




Social feed stream storage

Table Store can store massive volumes of social information produced by interactions between people, including IM chats, comments, threads, and likes. The elastic resources stored in Table Store are billed in Pay-As-You-Go method. At relatively low cost, Table Store can meet the needs of applications that feature significant traffic fluctuations and high concurrency when low latency is required. Table Store stores images and videos on OSS and, with CDN acceleration, provides optimal user experience.




Financial risk control

The advantages of Table Store such as low latency, high concurrency, and Pay-As-You-Go billing of elastic resources combine to optimize the risk control system, allowing you to strictly control transaction risks. Flexible data structures enable fast iteration of business models as market needs shift.




IoV data storage

A single table can store petabytes of data without distributing data in separate databases and tables, which simplifies the business logic. The schema-free data model enables easy access to the monitoring data of different vehicle-mounted devices. Table Store can be seamlessly integrated with multiple big data analysis platforms and real-time computing services for ease of real-time online query and business report analysis.




IoT time series data storage

With a single table capable of storing petabytes of data and processing thousands of queries per second (QPS), Table Store makes it easy to store the time series data of IoT devices and monitoring systems. The big data analysis SQL direct read function and the efficient incremental streaming read interface provide an easy way of offline data analysis and real-time streaming computing.



E-commerce recommendation

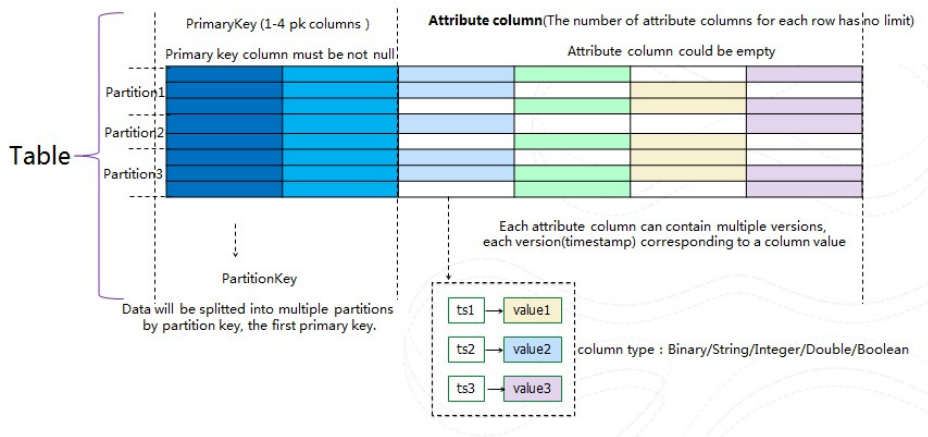
Table Store makes it possible for you to deal with data volumes and access performance with ease when handling a large number of historical transaction orders. Combined with MaxCompute, Table Store enables precision marketing, elastic resource storage, and Pay-As-You-Go billing, so that you can easily cope with peak hours when all customers go online.



Terms

Overview

The data model of Table Store is described by Table, Row, Primary Key, and Attribute, as shown in the following picture:



A table is a set of rows, and a row consists of the Primary Key and Attribute.

The Primary Key and Attribute columns consist of names and values.

All rows in a table must contain Primary Key columns with the same number and name. However, the number, name, and data type of Attribute columns within the rows can vary.

Each Attribute column can contain multiple versions, and each version (that is, the timestamp) corresponds to a column value, which is different from that of a Primary Key column.

Note: Timestamp is the sum of the milliseconds counted from 1970-01-01 00:00:00 UTC to the time when data is written.

Example

The following example illustrates two rows in a table. The **ID** column is the Primary Key column.

ID	Type	ISBN	PageCount	Length
----	------	------	-----------	--------

'4776'	timestamp = 146667635400 0, value = 'Book'	timestamp = 146667635400 0, value = '123*45678912345'	timestamp = 146667635400 0, value = 666	-
'6555'	timestamp = 146667635400 0, value = 'Music'	-	-	timestamp = 146667635400 0, value = 400; timestamp = 146676275400 0, value = 500

ID is the Primary Key of the given table. Rows with the ID of '4776' and '6555' have different attributes and can be stored in the same table.

The Attribute column **Type** of the row with ID '4776' only has one version. The version is 1466676354000 and the data is 'Book'.

The Attribute column **Length** of the row with ID '6555' has two versions. The data of version 1466676354000 is 400 and the data of version 1466762754000 is 500.

Max Versions is a data table attribute used to indicate the maximum number of data versions in each Attribute column of a data table. When the number of data versions in an Attribute column exceeds the value of Max Versions, the earliest versions will be deleted asynchronously.

After creating a data table, you can use the UpdateTable operation to dynamically change the value of Max Versions.

Data that exceeds the value of Max Versions will be cleared asynchronously. When dynamically modifying the value of Max Versions, the following restrictions may apply:

Data that exceeds the value of Max Versions is invalid and cannot be viewed or accessed, even though it is not deleted.

If you reduce the value of Max Versions, data that exceeds the new value will be deleted asynchronously.

If you increase the value of Max Versions, data that exceeds the old value, but has not been deleted, can be accessed again.

Time To Live (TTL) is a data table attribute measured in seconds that indicates the validity period of

data. To save data storage space and reduce storage costs, Table Store runs in the background and automatically clears any data that exceeds the TTL. TTL usage is described as follows:

Set TTL when creating a table. If you do not need a data expiration date, set TTL to **-1**.

After creating a data table, you can use the UpdateTable operation to dynamically change the value of TTL.

For a data table where the TTL is 86400 (one day), all attribute columns with versions earlier than 1468944000000 (divided by 1000 and converted to seconds to get 2016-07-20 00:00:00 UTC) will expire at 2016-07-21 00:00:00 UTC and be automatically cleared.

Expired data will be cleared asynchronously. When dynamically modifying the value of TTL, the following restrictions may apply:

Data with a date that exceeds the value of TTL is invalid and cannot be viewed or accessed, even though it is not deleted.

If you reduce the value of TTL, data with a date that exceeds the later value will be deleted asynchronously.

If you increase the value of TTL, data with a date that exceeds the earlier value, but has not been deleted, can be accessed again.

Max Version Offset is a data table attribute measured in seconds. To prevent the writing of unexpected data, the server checks the Attribute columns' versions when processing writing requests. Writing data to a specified row fails if the row has an Attribute column in which:

- Its version is earlier than the current writing time minus the value of Max Version Offset.
- Its version is later than or equal to the current writing time plus the value of Max Version Offset.

The valid version range for Attribute columns is [**Data written time – the value of Max Version Offset, Data written time + the value of Max Version Offset**). The data written time is the sum of the seconds counted from 1970-01-01 00:00:00 UTC to the time when data is written. Versions of the Attribute columns (expressed in milliseconds) must, after being divided by 1000 and converted to seconds, fall into the valid version range.

For example, if the valid version range of a data table is 86400 (one day), then at 2016-07-21 00:00:00 UTC, only data with versions later than 1468944000000 (converted to seconds to get 2016-07-20 00:00:00 UTC) but earlier than 1469116800000 (converted to seconds to get 2016-07-22 00:00:00 UTC) can be written to the data table. If a row has an Attribute column whose version is

1468943999000 (converted to seconds to get 2016-07-19 23:59:59 UTC, which is less than a day), the data cannot be written to the row.

Max Version Offset must be set to a non-zero value based on the number of seconds during the period from 1970-01-01 00:00:00 UTC to the current time.

If you do not set the value of Max Version Offset when creating a data table, the data table uses the default value **86400**.

After creating a data table, you can use the UpdateTable operation to dynamically change the value.

Primary Key

A Primary Key is the unique identifier of each row in a table. It consists of 1 to 4 columns.

When a table is created, the Primary Key must be defined. To define a Primary Key, the column name and data type of each Primary Key column, and the fixed sequence of the Primary Key columns, must be provided.

The data types of the Primary Key column can only be String, Integer, and Binary. For a Primary Key column of String or Binary data type, the size of the column value must not exceed 1 KB.

Partition key

The first Primary Key column that has a Primary Key is also called a partition key. Table Store checks the range where the partition key of each row is located, and automatically allocates the data in this row to the corresponding partition and machine to achieve load balancing.

Rows with the same partition key belong to the same partition. A partition may contain multiple partition keys. A partition key is the smallest partition unit. The data under a partition key cannot be split. To avoid a situation where the partition is too large to be split, we recommend that the total data volume of all rows under a single partition key does not exceed 1 GB.

To improve load balancing, Table Store splits and merges partitions according to specific rules. This process is automated without application intervention.

Attribute

Data of rows are stored in the Attribute columns. There is no limit to the number of attribute columns for each row.

Version

When writing data, you can specify the Attribute columns' versions. If you do not specify any

versions, the server generates the versions of the Attribute columns based on the current timestamp (expressed by the number of milliseconds that have elapsed since 1970-01-01 00:00:00 UTC). You can specify the maximum number of versions per column, or the version range, to limit the data that can be read from each row.

When the number of versions written to an Attribute column exceeds the value of Max Versions, the data of earlier versions is discarded so that the number of remaining versions is equal to the Max Versions value.

A version number also indicates the time when data is generated. The version is expressed by the number of milliseconds that have elapsed since 1970-01-01 00:00:00 UTC. For a data table where the TTL is set to 86400 (one day), the data in the Attribute column where the version is 1468944000000 (2016-07-20 00:00:00 UTC) will expire at 2016-07-21 00:00:00 UTC and be automatically deleted.

Note:

- During TTL comparison and Max Version Offset calculation, versions in milliseconds must be converted to seconds by dividing by 1000.
- If versions are determined by the server, data will be cleared after the time (in seconds) indicated when TTL has elapsed since the data was written.
- To prevent invalid written data, the system denies the writing of expired data. For example, at 2016-07-21 00:00:00, the data in which the version is earlier than 1468944000000 (2016-07-20 00:00:00 UTC) cannot be written to a data table where the TTL is 86400.
- To prevent writing errors, the system requires that the Attribute column version to which the data is written is within range of [Data written time – The value of Max Version Offset, Data written time + The value of Max Version Offset).

Column naming conventions

The Primary Key and Attribute columns follow the same naming conventions, in which:

Each name can be 1 to 255 characters in length.

Only uppercase and lowercase letters, digits, and underscores (_) are permitted.

The first character must be an uppercase or lowercase letter, or underscore (_).

All characters are case sensitive.

Data types of column values

Table Store supports five data types of column values, as shown in the following table.

Data type	Definition	Permitted as Primary Key	Size limitation
String	UTF-8, could be empty	Yes	For Primary Key column, not greater than 1 KB. For Attribute column, see restricted items.
Integer	64 bit Integer	Yes	8 Bytes.
Double	64 bit Double	No	8 Bytes.
Boolean	True/False	No	1 Byte.
Binary	Could be empty	Yes	For Primary Key column, not greater than 1 KB. For Attribute column, see restricted items.

The read/write throughput is measured by read/write capacity units (CUs), which is the smallest billing unit for the data read and write operations.

One read CU indicates that 4 KB data is read from the table, and one write CU indicates that 4 KB data is written into the table. Data smaller than 4 KB during the operation will be rounded up to an integer. For example, writing 7.6 KB data will consume two write CUs, and reading 0.1 KB data will consume one read CU.

When applications use an API to perform Table Store read/write operations, the corresponding amount of read/write CUs will be consumed.

Reserved throughput

The reserved read/write throughput:

Is an attribute of a table.

When creating a table, the application specifies the read/write throughput reserved for the table. Configuring the reserved read/write throughput does not affect the table's access performance and service capability.

Can be set to 0.

When the reserved read/write throughput is greater than 0, Table Store assigns and reserves enough resources for the table according to this configuration to ensure low resource costs. The reserved read/write throughput of a table can be dynamically changed using the UpdateTable operation.

For a non-zero reserved read/write throughput, your Table Store service is billed even if there are no read and write requests. To ensure billing accuracy, Table Store limits the maximum reserved read/write throughput to 5,000 CUs per table (neither read throughput nor write throughput can exceed 5,000 CUs). If you require more than 5,000 CUs of reserved read/write throughput for a single table, open a ticket to increase the throughput.

The reserved read/write throughput of a non-existent table is regarded as 0. To access a non-existent table, one additional read CU or one additional write CU will be consumed depending on the actual operation.

Additional throughput

The additional read/write throughput refers to the portion of the actual consumed read/write throughput that exceeds the reserved read/write throughput. Its statistical period is one second. For example, if the reserved read throughput of a table is set to 100 CUs and, within one second, the read operation actually consumes 120 CUs, then the additional read throughput consumed within the second is 20 CUs.

For the additional read/write throughput mode, it is difficult to estimate the amount of compute resources that need to be reserved for data tables. Table Store is required to provide sufficient service capability to effectively handle access traffic spikes. For this reason, the unit price of additional read/write throughput is higher than that of reserved read/write throughput. To ensure low costs are maintained, we recommend that a proper value of the reserved read/write throughput be set.

Note: Because it is difficult to accurately reserve resources based on the additional read/write throughput, in extreme situations, Table Store may return an error `OTSCapacityUnitExhausted` to an application when an access to a single partition key consumes 10,000 CUs per second. In this case, policies, such as backoff retry, are used to reduce the frequency of access to the table.

For more information, see [Table Store tables and billing methods](#).

Region refers to a service region of Alibaba Cloud. Table Store is deployed across many service regions, so that you can select the most suitable region according to your requirements.


The following table lists the regions supported by Table Store:

Region	Region name
China East 1 (Hangzhou)	cn-hangzhou
China East 2 (Shanghai)	cn-shanghai
China North 2 (Beijing)	cn-beijing
China North 5 (Huhehaote)	cn-huhehaote
China South 1 (Shenzhen)	cn-shenzhen
Hong Kong	cn-hongkong

Singapore	ap-southeast-1
US West 1 (Silicon Valley)	us-west-1
Asia Pacific NE 1 (Japan)	ap-northeast-1
Germany 1 (Frankfurt)	eu-central-1
Asia Pacific SE 2 (Sydney)	ap-southeast-2
Asia Pacific SE 3 (Kuala Lumpur)	ap-southeast-3

Overview

An instance is a logical entity in Table Store used to manage tables as database in a relational database management system (RDBMS). After activating Table Store, create an instance in the Table Store console and then create and manage tables within this instance. An instance is the basic unit of Table Store's resource management. Table Store implements access control and resource metering at the instance level.



Users can create different instances for different businesses to manage their respective tables. Users can also create multiple instances for one business based on different development, testing, and production purposes.

Table Store allows one Alibaba Cloud account to create up to 10 instances, and up to 64 tables can be created within each instance. If more instances or tables are needed, [open a ticket](#).

Naming conventions

The name of each instance must:

Be unique within each service region. If the instances are in different service regions, users can create instances of the same name.

Be 3 Bytes to 16 Bytes in length.

Contain only uppercase and lowercase letters, digits, and hyphens.

Start with either an uppercase or lowercase letter.

Not end with a hyphen.

Instance type

Table Store supports two instance types: high-performance instance and capacity instance.

Note: An instance type cannot be modified once the instance is created.

The two instance types have the same functions and support petabyte-sized data volumes for a single table, however, they differ in costs and application scenarios.

High-performance instance

High-performance instances support millions of read-write transactions per second (TPS) with 1 ms average latency of read and write operations per row. High-performance instances are suitable for scenarios requiring high read and write performance and concurrency, such as gaming, financial risk control, social networking apps, recommendation systems, and public opinion monitoring.

Capacity instance

Capacity instances provide write throughput and write performance comparable to that of the high-performance instances, but with lower costs. However, the capacity instances do not equal the read performance and concurrency of high-performance instances. The capacity instances are suitable for services with high write frequency, but low read frequency, as well as services with high affordability, but relatively low performance requirements. This includes access to log monitoring data, Internet of Vehicles data, device data, time sequence data, and logistic data.

Note: Capacity instances do not support reserved read/write throughput. All reads and writes are billed based on the additional read/write throughput.

Instance type supported by region

Region	High-performance instance	Capacity instance
China East 1 (Hangzhou)	Support	Support
China East 2 (Shanghai)	Support	Support
China North 2 (Beijing)	Support	Support
China North 5 (Huhehaote)	In development	Support
China South 1 (Shenzhen)	Support	Support
Hong Kong	In development	Support
Singapore	Support	In development
US West 1 (Silicon Valley)	Support	In development
Asia Pacific NE 1 (Japan)	In development	Support
Germany 1 (Frankfurt)	In development	Support
Asia Pacific SE 2 (Sydney)	In development	Support
Asia Pacific SE 3 (Kuala Lumpur)	In development	Support

Each instance corresponds to an endpoint that is also known as the connection URL. The endpoint needs to be specified before any operations on the tables and data of Table Store.

To access the data in Table Store from the Internet, the endpoint uses the following format:

```
https://instanceName.region.ots.aliyuncs.com
```

To access the data in Table Store from an Alibaba Cloud ECS instance of the same region through the intranet, the endpoint uses the following format:

```
https://instanceName.region.ots-internal.aliyuncs.com
```

For example, to access the Table Store instance in China East 1 (Hangzhou) region, with the instance name of myInstance:

```
Endpoint for Internet access: https://myInstance.cn-hangzhou.ots.aliyuncs.com
Endpoint for intranet access: https://myInstance.cn-hangzhou.ots-internal.aliyuncs.com
```

Better performance, such as lower response latency and no extra Internet traffic, can be

expected through the intranet.

If an application accesses Table Store from an ECS instance in VPC, the endpoint uses the following format:

```
https://vpcName-instanceName.region.vpc.ots.aliyuncs.com
```

For example, the service address used by an application in China East 1 (Hangzhou) region to access the instance named myInstance from a network named testVPC:

```
Endpoint of VPC access: https://testVPC-myInstance.cn-hangzhou.vpc.ots.aliyuncs.com
```

This VPC access address is only used for access initiated by servers in the testVPC network.

Item	Restriction	Description
Number of instances saved under an Alibaba Cloud user account	Up to 10	If you want to raise the limit, open a ticket.
Number of tables in an instance	Up to 64	If you want to raise the limit, open a ticket.
Instance name length	3-16 Bytes	Can contain uppercase and lowercase letters, digits, and hyphens. Must begin with a letter, and must not end with a hyphen.
Table name length	1-255 Bytes	Can contain uppercase and lowercase letters, digits, and underscores, and must begin with a letter or underscore.
Column name length	1-255 Bytes	Can contain uppercase and lowercase letters, digits, and underscores, and must begin with a letter or underscore.
Number of Primary Key columns	1-4	There must be at least one column.
Size of String type Primary Key column values	Up to 1 KB	A single Primary Key column's String type column value is limited to 1 KB.
Size of String type attribute column values	Up to 2 MB	A single attribute column's String type column value is limited to 2 MB.

Size of Binary type Primary Key column values	Up to 1 KB	A single Primary Key column's Binary type column value is limited to 1 KB.
Size of Binary type attribute column values	Up to 2 MB	A single attribute column's Binary type column value is limited to 2 MB.
Number of attribute columns in a single row	Unlimited	A single row can contain an unlimited amount of attribute columns.
Data size of a single row	Unlimited	The total size of all column names, and column value data, for a single row is unlimited.
Reserved read/write throughput for a single table	Up to 5000	If you want to raise the limit, open a ticket.
Number of columns in a read request's columns_to_get parameter	Up to 128	The maximum number of columns obtained in a row of data in the read request.
Table-level operation QPS	10	The QPS of a table-level operation on an instance must not exceed 10. For table-level operations, see Table operations .
Number of UpdateTable operations for a single table	Raise: Unlimited Lower: Unlimited	The reserved read/write throughput for each table can be raised or lowered unlimited times within a calendar day (from 00:00:00 to 00:00:00 of the next day in UTC time).
UpdateTable frequency for a single table	Maximum of one update every 2 minutes	The reserved read/write throughput for a single table cannot be adjusted beyond the frequency of once every 2 minutes.
The number of rows read by one BatchGetRow request	Up to 100	N/A
The number of rows written by one BatchWriteRow request	Up to 200	N/A
Data size of one BatchWriteRow request	Up to 4 MB	N/A
Data returned by one GetRange operation	Up to 5000 rows or 4 MB	The data returned by a single operation cannot exceed 5000 rows or 4 MB. Otherwise, you must read the excessive data with a returned token.

The data size of an HTTP Request Body	Up to 5 MB	N/A
---------------------------------------	------------	-----