

# Alibaba Cloud MaxCompute

## Release Note

Issue: 20190819

# 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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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 contact Alibaba Cloud directly if you discover any errors in this document.



## Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 <b>Danger:</b> Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 <b>Warning:</b> Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 <b>Notice:</b> Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 <b>Note:</b> You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
<b>Bold</b>	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[ ] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand   slave}</code>



# Contents

---

Legal disclaimer.....	I
Generic conventions.....	I
1 Announcements.....	1
2 History.....	5



# 1 Announcements

---

This topic lists updates to MaxCompute in descending chronological order.

August 21, 2019 (Beijing time): External table is upgraded.

An upgrade will be carried out on MaxCompute at 14:00-23:00, 21 August 2019 (Beijing time). During this upgrade, when you use external table custom storage handler to implement the outputter interface, the job may fail if column data is obtained through column names rather than digital subscriptions.

Regions: Asia Pacific NE 1, Central Europe 1, China(Hong Kong), Asia Pacific SE 2

The details are as follows: The record passed in `Outputter . output ( Record record )` is the record generated by the last operator of Outputter. The column name may change, and the system cannot guarantee the fixed column name. For example, the column name generated by the expression `some_function ( column_a )` is a temporary column name, so the use of `record . get ( column name )` to obtain column content may be affected. You may use `record . get ( index )` instead. To get the column name of the table in Outputter, please call `DataAttributes . getFullTableColumns ()`.

For any questions, you can submit the ticket to contact us.

July 24, 2019(Beijing time): MaxCompute Spark is released.

Region: West USA 1, China(Hong Kong), Central Europe 1, Asia Pacific SE 1, Asia Pacific SOU 1.

March 26, 2019: MaxCompute SQL is upgraded.

- The `GROUPING SETS` clause (together with the `CUBE` and `ROLLUP` subclauses) can be used to aggregate and analyze data of multiple dimensions. For more information, see [#unique\\_4](#).
- The `INTERSECT`, `MINUS`, and `EXCEPT` clauses gained support. For more information, see [#unique\\_5](#).
- When the system reads files in ORC or Parquet format by using external tables, it can crop the columns in the files, which can reduce I/O and resource usage, thereby lowering overall computing costs.

- Systems that run in the Java UDX framework are enhanced to support writable parameters. For more information, see [#unique\\_6](#).

March 1, 2019: External tables of MaxCompute begin to incur charges.

Starting from March 1, 2019, SQL external tables (which are used to process OSS data and Table Store data) of MaxCompute begin to incur charges.

The charging policy is as follows:

$$\text{One-time SQL computing fee} = \text{Input data volume} \times \text{SQL complexity} \times \text{SQL price}$$

The SQL price is 0.0044 USD/GB/Complexity. The complexity coefficient is 1. All the fees are charged on the next day, and you will receive an account bill.

For more information, see [#unique\\_7](#).

If you have any questions, open a ticket.

January 15, 2019: The underlying structure of MaxCompute in China (Hong Kong) is optimized from 16:00 to 20:00.

The underlying metadata warehouse of MaxCompute in China (Hong Kong) is optimized from 16:00 to 20:00 on January 15, 2019 to improve the performance and stability of MaxCompute. During the release window, users in the Hong Kong region may encounter submission delays or failures for tasks, which may last about one minute. In the worst cases, the application may be unavailable for 30 minutes (or half an hour). Therefore, we recommend that you do not submit any tasks during the release window. If you have any questions, contact us through DingTalk or by opening a ticket. Users in other regions are not affected.

December 24, 2018: MaxCompute supports time zone configuration.

MaxCompute project uses the UTC+8 time zone by default. The time-related built-in functions and datetime, timestamp, date fields are calculated based on UTC+8. From December 24, 2018, users can configure time zones in MaxCompute using either of the following methods:

- Session level: Submit the `set odps . sql . timezone =< timezoneid >;` SQL statement and a calculation statement. The following is an example:

```
set odps . sql . timezone = Asia / Tokyo ;
select getdate ();
-- Results :
output :
```

```

+-----+
|  _c0  |
+-----+
| 2018 - 10 - 30    23 : 49 : 50  |
+-----+

```

- **Project level:** The project owner runs the `setProject odps . sql . timezone =< timezoneid >;` SQL statement using a CLI. After a project is configured, the corresponding time zone is used automatically, which will affect the data of existing tasks. Therefore, we recommend that you do not configure the existing projects. Instead, you can configure new projects as needed.

**Note:**

- The time zone configuration supports SQL built-in date functions, UDF, UDT, UDJ, and select transform.
- The time zone format such as Asia/Shanghai (daylight saving is considered) is supported. The GMT+9 format is not supported.
- If the time in the SDK time zone differs from that in the project time zone, you need to configure the GMT time zone so as to convert the date data to a string.
- After the time zone is configured, there might be difference between the real time and the output time when you run the related SQL statements through MaxCompute. Between the years of 1900 and 1928, the time difference is 352 seconds. Before the year of 1900, the time difference is 9 seconds.
- To ensure the datetime data accuracy in different regions, we will upgrade MaxCompute and the Java SDK and related console versions with the `-oversea` suffix. After the upgrade, the display of existing datetime data (before the year of 1928) stored in MaxCompute might be affected.
- When you upgrade MaxComput, we recommend that you upgrade the Java SDK and console versions if the local time zone is not UTC+8, so as to ensure the accuracy and consistency between the SQL computing result and the Tunnel transferred data after '1900-01-01'. For the datetime data before '1900-01-01', the SQL computing output and the Tunnel transferred data might differ by 343 seconds. For the existing datetime data before '1928-01-01', the time difference is 352 seconds.
- If you continue using the SDK and console versions without the `-oversea` suffix, you might encounter time difference between the SQL output and the Tunnel

transferred data. The time difference before '1900-01-01' is 9 seconds, and the time difference between '1900-01-01' and '1928-01-01' is 352 seconds.



**Note:**

When you update or configure the time zone of Java SDK and console versions, the time zone of DataWorks remains unchanged. Therefore, there might be time difference and you need to evaluate the impact of task scheduling in DataWorks. In the Japan region, the time zone of DataWorks is GMT+9. In the Singapore region, the time zone of DataWorks is GMT+8.

- If you are using a third-party client connected by JDBC, you need to set the time zone on the client to ensure the time consistency.
- MapReduce supports time zone configuration.
- Spark supports time zone configuration.
  1. For tasks that are submitted to ODPS computing clusters, the project time zone can be automatically obtained.
  2. For settings that are made through the yarn-client method, such as spark-shell, spark-sql, and pyspark, you need to configure the spark-defaults.conf parameter of the driver and add `spark.driver.extraJavaOptions -Duser.timezone = America / Los_Angeles`. The 'timezone' in the preceding statement is the time zone to be used.
- PAI supports time zone configuration.
- Graph supports time zone configuration.

## 2 History

---

When Alibaba Cloud was founded in September 2009, our vision was to be the first platform for data processing and sharing. In April 2010, that vision came a little closer to reality with the Open Data Processing Service (ODPS) supporting the newly launched loan business area of Ant Financial. From there, things went from strength to strength. In 2012, a Unified Data Platform was established. By the end of 2013, we had the ability to process massive volumes of data on a large scale. From 2014 to 2015, the big data platform was refined, and in 2016, MaxCompute 2.0 was born, providing massive data warehousing solutions and big data modeling capabilities.

### Key milestones

Time	Development
2010.04	Named ODPS, the service is released as an operational component of Alibaba Group' s Ant Financial
2013.05	ODPS is released for beta testing
2013.07	ODPS v1.0 is released as a commercially available service. A single cluster contains 5,000 servers, with support for multi-level clusters available.
2016.09	Renamed MaxCompute, v2.0 is released as a commercially available service.