

TDSQL-A for PostgreSQL

FAQs

Product Documentation



Copyright Notice

©2013-2019 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

FAQs

Last updated : 2021-07-02 16:46:16

What are the differences between TDSQL-A for PostgreSQL and TDSQL for PostgreSQL?

Both of them are distributed databases developed by Tencent Cloud based on PostgreSQL and belong to the formerly TBase series.

The main difference is that TDSQL-A for PostgreSQL is an OLAP-enhanced edition. It has proprietary column storage and vectorized execution engines and focuses on OLAP scenarios, while TDSQL for PostgreSQL focuses on high-concurrency OLTP scenarios and hybrid HTAP scenarios.

The two services have many similar features that are inherited from each other, such as compatibility with Oracle, multi-level security features, and hot/cold data separation capabilities, and they will inherit more features from each other in the future.

What are the differences between TDSQL-A for PostgreSQL and TDSQL for ClickHouse?

Both of them belong to the TDSQL-A series. They are developed based on open-source PostgreSQL and ClickHouse respectively and compatible with PostgreSQL and MySQL syntax protocols respectively.

TDSQL-A for PostgreSQL has better complex SQL query capabilities such as multi-table complex query than TDSQL-A for ClickHouse and also has better transaction logic to ensure the ACID of transactions.

What are the differences between TDSQL-A for PostgreSQL and Tencent Cloud Data Warehouse PostgreSQL (CDWPG)?

TDSQL-A for PostgreSQL is a distributed analytical database developed by Tencent Cloud. It is compatible with PostgreSQL protocols and has proprietary column storage and vectorized execution engines as well as deeply optimized kernels. It shares the master code branch with TDSQL for PostgreSQL and inherits and adapts to the full set of enterprise-grade features, such as high availability, security capability, management and control capabilities, and kernel performance. It can be purchased in both public and private clouds.

CDWPG (formerly Snova) is Tencent Cloud's in-cloud data warehouse based on Greenplum.

Can an Oracle database be migrated to TDSQL-A for PostgreSQL?

TDSQL-A for PostgreSQL is one of the best choices to replace Oracle databases. It supports complex subqueries, stored procedures, window functions, and many other enterprise-grade features. It also supports common Oracle syntax and functions at the kernel level, making it possible to implement

rapid business migration to TDSQL-A for PostgreSQL with few or no business code modification required.

Feature	Description
Analytic function (window function)	Common analytic functions such as Rank, Lead, and Lag are supported, and analytic function customization is also supported.
NoSQL data type	Multiple NoSQL data types such as JSON, JSONB, XML, and array are supported.
Complex data analysis	Powerful multi-table join queries, subqueries, materialized views, and partition tables are supported.
Iterated WITH (CTE)	Iterated <code>WITH</code> statements are supported, which can be used for tree-structure or graphic hierarchy processing. <code>WITH</code> also supports <code>update</code> and <code>delete</code> .
Custom function and aggregate function	Functions, aggregate functions, types, etc., can be customized.
DDL rollback	Dangerous operations such as <code>drop table</code> , <code>drop index</code> , and <code>truncate table</code> can be rolled back before the transaction ends.
External data source management through SQL/MED	Interactive operations can be performed on external data sources such as Oracle, SQL Server, MySQL, MongoDB, and Redis through the FDW mechanism.
Extension	PostgreSQL extensions contributed by the community can be downloaded here .

How compatible is TDSQL-A for PostgreSQL with PostgreSQL?

- In terms of functionality, TDSQL-A for PostgreSQL implements most of the features specified in the SQL:2011 standard as PostgreSQL does. It supports basic features such as SELECT, UPDATE, INSERT, JOIN, and GROUP BY as well as advanced features such as foreign key constraint, check, primary key, trigger, view, stored procedure, multi-level transaction, and multi-version concurrency control.
- In terms of data types, TDSQL-A for PostgreSQL supports common data types such as numeric, string, and time as well as special types such as self-incrementing sequence, currency, geometry, UUID, array, JSON, and range.
- In terms of API, TDSQL-A for PostgreSQL supports APIs for most programming languages such as JDBC, ODBC, Shell, C, Python, PHP, and .NET.

You can use TDSQL-A for PostgreSQL exactly like standalone PostgreSQL with no additional learning costs incurred.

Does TDSQL-A for PostgreSQL support strong sync replication?

Yes. Ensuring that the primary and standby nodes have the identical data is the basis of the entire disaster recovery system. If the primary node fails, the database service can be switched to the standby node without any data loss.

Does TDSQL-A for PostgreSQL support audit?

Yes. TDSQL-A for PostgreSQL supports all-round auditing in multiple dimensions. Bypass detection is used for auditing, which has little impact on database operations.

How do I perform efficient distributed joins?

- In terms of the execution method, the coordinator receives your SQL requests, generates the optimal cluster-level distributed query plan based on the collected cluster statistics, and distributes it to the datanodes participating in the computation for execution, that is, the coordinator sends the execution plan, and the datanodes are responsible for executing the plan.
- In terms of data interaction, efficient data exchange tunnels are established between datanodes, enabling efficient data exchange. The process of data exchange is called data redistribution in TDSQL-A for PostgreSQL.

Based on the efficient global query plan and data redistribution technology, TDSQL-A for PostgreSQL can take full advantage of parallel computing and complete the joins efficiently.

How do I use partition tables?

For more information, please see [Partition Table](#).

How do I use the column storage and compression features?

For more information, please see [Table Operations](#).