

Cloud Data Warehouse for PostgreSQL

Performance Metrics

Product Documentation



Copyright Notice

©2013-2024 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.

Contents

Performance Metrics

TPC-B

Performance Metrics

TPC-B

Last updated : 2024-02-02 15:27:57

Test Description

The latest version of CDWPG encapsulates Greenplum 6.x. Compared to 5.x, it features greatly improved CRUD (i.e., OLTP) capabilities in multiple concurrent scenarios. These improvements include:

1. Upgraded the PostgreSQL kernel to v9.4 and introduced lock optimizations such as fastpath.
2. Provided global deadlock detection.
3. Optimized global transactions.

Greenplum officially uses [TPC-B](#) to test the OLTP capabilities of v6.x. Therefore, CDWPG also uses this benchmark for testing.

Test Environment

CDWPG offers **compute-intensive** and **storage-intensive** models. You can simply distinguish them from each other:

Compute-intensive: The underlying hardware is SSD with strong random read/write capabilities. It is suitable for hot data analysis and scenarios with mixed loads.

Storage-intensive: The underlying hardware is HDD with average random read/write capabilities but large disk capacity. It is suitable for storing and analyzing larger-scale historical data.

In conclusion, TPC-B requires the **compute-intensive** model. You can choose two nc2.large nodes and purchase them directly on the [purchase page](#).


Test Tool

Use pgbench, the standard tool of PostgreSQL, for testing.

1. If the test environment is CentOS 7.x, you can download the already compiled tool [here](#).
2. For other environments, you can compile PostgreSQL yourself or directly install the binary package. PostgreSQL 9.4 is recommended for better compatibility.

Test Steps

Creating test database



```
CREATE DATABASE pgbench;
```

Modifying query optimizer




```
ALTER DATABASE pgbench SET optimizer = off;
```

Note:

1. CDWPG is preconfigured with two query optimizers, GPORCA and Postgres Planner. The former is the default and suitable for complex query parsing, while the latter is required for OLTP queries.
2. This parameter can also be set at the session level. For convenience, it is set directly at the database level; that is, Postgres Planner is used to access this database, while other databases are still accessed via GPORCA.

Initializing test data



```
./pgbench -i pgbench -s 100 -p 5436 -h {host} -U {user} pgbench
```

Note:

10 million data records are tested here.

Running test script



```
./pgbench -h {host} -p 5436 -r -n -c 32 -j 32 -T 120 -P 1 -U {user} pgbench
```

Optimization

According to Greenplum's official data, v6.x has a peak TPS of about 5,000 in an ideal environment. You can modify the following parameters to reach a higher value in real tests in CDWPG:

Parameter	Value	Description
-----------	-------	-------------

log_statement	none	Disables the log output of the master node
gp_enable_fast_sri	on	Improves the efficiency of a single `INSERT`
gp_enable_gpperfmon	off	Disables monitoring sampling

Note:

1. In addition to the above configuration, you also need to disable the standby master node for optimal performance.
2. The above configuration is mainly used to test the extreme performance of CDWPG. Therefore, the above modifications are not recommended in actual production environments.
3. As the configuration modification feature is not available in CDWPG yet, you can [contact us](#) if you need it. You need to restart the cluster after the modification.