


Cloud Data Warehouse

Product Introduction

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

Product Introduction

Overview

Basic Concepts

Cluster Architecture

Strengths

Use Cases

Product Introduction

Overview

Last updated : 2024-01-19 16:46:19

Cloud Data Warehouse offers easy-to-use, flexible, and stable ClickHouse hosting services in the cloud. A data warehouse can be created in minutes for massive real-time data query and analysis, thereby improving the overall efficiency of data value mining. By leveraging the massively parallel processing (MPP) architecture, it can query data several times faster than traditional data warehouses.

Features

Cloud Hosting

Cloud Data Warehouse enables you to build a data warehouse at the terabyte or even petabyte scale in just a few minutes. With the mature cluster management functionality, it saves the heavy workload of human OPS to implement expense management, cluster configuration adjustment, and parameter configuration in the console.

OPS and Monitoring

Cloud Data Warehouse visualizes monitoring data so that you can view and stay up to date with cluster operations in a timely manner.

High-Performance Analysis

Cloud Data Warehouse allows you to linearly scale the storage and computing capabilities to process queries several times faster than traditional data warehouses, even up to terabytes of data per second per query, with the excellent query performance of ClickHouse and the distributed massively parallel processing (MPP) framework.

Storage-Compute Separation

Within a Cloud Data Warehouse cluster, configurations are performed separately for storage and compute nodes, so that you can scale them independently as needed to reduce hardware costs.

Security and Reliability

Your clusters can be independently deployed in isolated VPCs for more secure data access. Moreover, Cloud Data Warehouse comprehensively supports the data replica mechanism to implement service disaster recovery and failover imperceptible to users.

Basic Concepts

Last updated : 2024-01-19 16:46:19

Region

A region is a geographical area where a physical Cloud Data Warehouse server is located. Select a Tencent Cloud region with caution as you cannot change it after service purchase and networks are completely isolated between regions. For a lower access latency and higher read and write speeds, we recommend you select a region closest to your users.

AZ

An availability zone (AZ) is a physical IDC with isolated resources within a region.

Cloud Data Warehouse Cluster

A Cloud Data Warehouse cluster is a distributed cluster of multiple ClickHouse nodes, each of which may have one or two replicas depending on the specifications you purchase. A cluster may contain one or more shards.

Shard

Cloud Data Warehouse stores massive amounts of data on multiple nodes, each of which only stores and processes a part of the data. When there is one replica, a shard corresponds to a node; when there are two replicas, a shard corresponds to two nodes.

Replica

A replica is used to guarantee data security and service availability in case of exceptions. Cloud Data Warehouse stores data on two nodes to generate two replicas.

HA

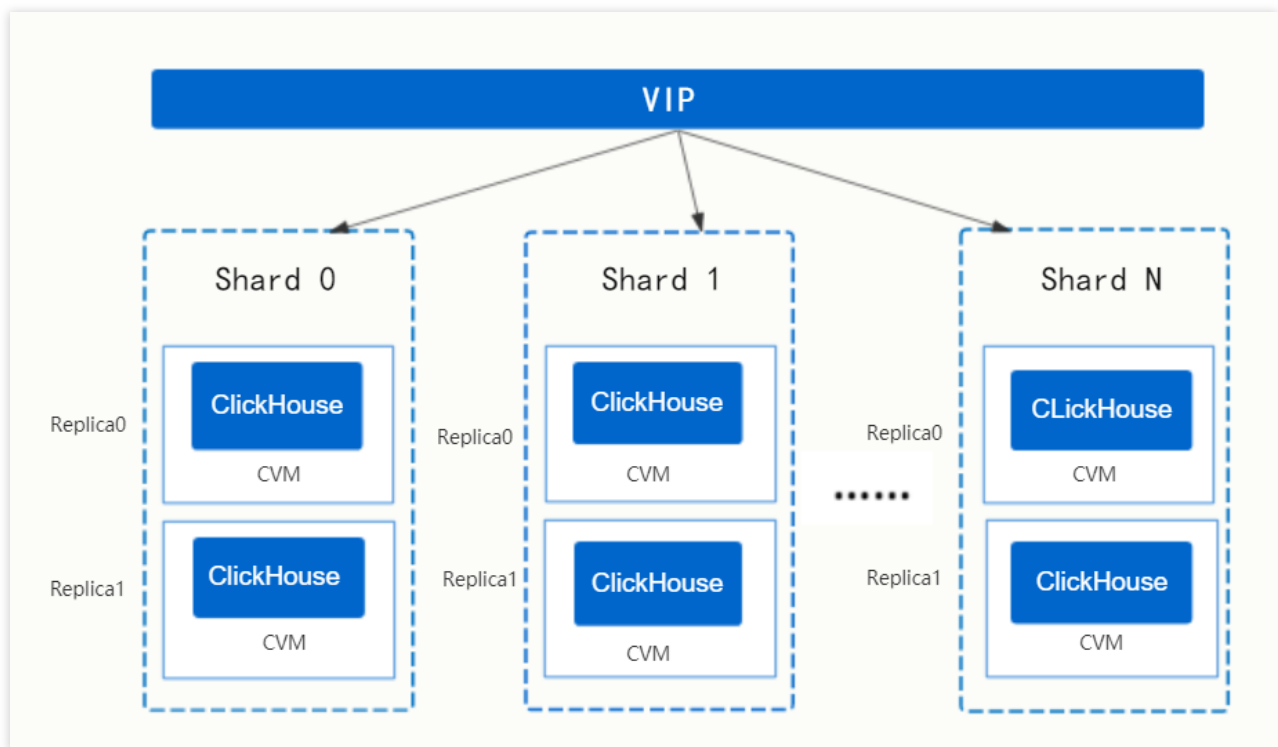
In HA mode, each shard has two replicas to ensure the cluster availability in case one of them fails.

Cluster Architecture

Last updated : 2024-01-19 16:46:19

Cloud Data Warehouse is a fully managed cloud product based on the open-source ClickHouse database management system to deliver easier use, better security, and higher reliability. The two are consistent in terms of architecture.

The general architecture of Cloud Data Warehouse is as shown below:



Strengths

Last updated : 2024-01-19 16:46:19

Ease of Use

You can create a ClickHouse analysis cluster in minutes in the console. A diversity of features such as OPS management, monitoring, and alarming free you from underlying infrastructure management and let you focus on analyzing data value with complete SQL statements.

Superior Performance

Cloud Data Warehouse leverages distributed massively parallel processing (MPP) architecture, allowing you to process queries several times faster than traditional data warehouses, even up to terabytes of data per second per query.

Elastic Scalability

The Cloud Data Warehouse console simplifies and accelerates cluster scaling and node configuration adjustment, dynamically supporting your business growth with high scalability.

Security and Reliability

Your clusters are independently deployed in isolated VPCs for more secure data access. Moreover, Cloud Data Warehouse comprehensively supports high cluster availability to implement service disaster recovery and failover imperceptible to users.

High Cost-Effectiveness

Cloud Data Warehouse enables you to build highly cost-effective managed ClickHouse clusters with devices in the cloud. This includes leveraging the 10x data compression algorithm of ClickHouse to efficiently reduce disk usage and costs compared with traditional data warehouses.

Use Cases

Last updated : 2024-01-19 16:46:19

User Behavior Analysis

You can import data such as clicks and dwells in websites, apps, and games into Cloud Data Warehouse to build a large wide table for user characteristics analysis. Cloud Data Warehouse excels in query and responds to queries in milliseconds during multidimensional and multimodal analysis. This facilitates the analysis of user behavior characteristics and patterns, providing strong support for targeted marketing and membership conversion.

Enterprise Operations Analysis

You can import massive amounts of enterprise operations data to Cloud Data Warehouse, which allows for hundreds of millions of queries in a large wide table in hundreds of dimensions in milliseconds. In this way, you can implement personalized statistics collection and uninterrupted analysis at any time to make informed business decisions. Cloud Data Warehouse can query data several times faster than traditional data warehouses. Additionally, it can be elastically scaled as needed, perfectly satisfying your need for a high-performance, cost-effective, and scalable data warehouse in the era of big data.