

Stream Compute Service

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

Strengths

Use Cases

Product Introduction

Overview

Last updated : 2023-11-07 15:26:47

Stream Compute Service is a powerful enterprise-grade real-time big data analytics platform based on Apache Flink. It features one-stop development, seamless connection, sub-second latency, low costs, high security, and high stability, helping businesses maximize the value of their data and accelerate their move to real-time data processing.

Comparison with open-source Apache Flink

Category	Item	Description		Stream Compute Service	Open-Source Apache Flink
Development and debugging	Data connection	Connection to various upstream and downstream data services through connectors, quickly aggregating real-time data and applying output data		Supported	Partially supported
		Cloud data products	Seamless integration with mainstream Tencent Cloud data products, including Cloud Data Warehouse, EMR, CDB, CKafka, TDMQ, and COS	Supported	Partially supported
		Self-built data services	Custom connectors for connecting to various external data systems, including those in the cloud and IDC	Supported	Supported
	Web IDE	Provides Web IDE as a one-stop development and debugging platform that integrates graphical development, multi-language development, metadata management, code debugging, and dependency management		Supported	Not supported
		Low-code	Real-time ETL graphical	Supported	Not

		development	development		supported
		Multiple languages	SQL, Java, and Scala	Supported	Not supported
		Metadata	Built-in metadata management that allows users to manage the variables of metadata parameters	Supported	Partially supported
		Automatic DDL	Intelligent detection of connected external data systems and automatic generation of table creation statements for data sources and data sinks	Supported	Not supported
		Custom functions	Custom functions such as UDF, UDTF, and UDAF	Supported	Supported
		Job debugging	Online SQL syntax check	Supported	Not supported
			Online SQL debugging that does not affect the production environment		
		Dependencies	Management of packages, and management and dynamic distribution of external dependencies (configuration files, dependent libraries, etc.)	Supported	Not supported
		Compute resources	Configuration of fine-grained resources and operator parallelism for JobManager and TaskManager	Supported	Partially supported
		Job versions	Historical version management of jobs and resources to facilitate business upgrade and development	Supported	Not supported
	Development	TencentCloud API provides all the APIs		Supported	Not

	using TencentCloud API	required for Web IDE development, through which business platforms and Stream Compute Service can be seamlessly connected.			supported
Deployment and Ops	Job deployment	Deployment management of jobs throughout the lifecycle (including start, stop, suspend, and resume) using Web IDE and TencentCloud API		Supported	Not supported
	Monitoring	Visualization of running parameters and status of jobs		Supported	Partially supported
		Running parameters	Visualization of the logic and parameters of running jobs, such as SQL statement details and parameter configurations	Supported	Not supported
		Running status events	Visualization of events such as unexpected job restarts, snapshot failures, and CPU or memory exceptions of JobManager/TaskManager	Supported	Not supported
	Metrics and alarms	Comprehensive job monitoring and alerting based on dynamic metrics defined at task granularity and over 65 monitoring metrics. The monitoring metrics are related to healthy running of cluster jobs and upstream and downstream systems and aggregated using functions such as sum, max, min, and avg.		Supported	Partially supported
		Monitoring metrics	Display of metrics at job, JobManager, TaskManager, and Task granularity, as well as accurate CPU/memory metrics for container Pods	Supported	Partially supported
		Monitoring alarms	Alarm rule (including AI-based dynamic threshold alarms) configuration for over 65 monitoring metrics using Tencent Cloud Observability Platform	Supported	Not supported

	Smart diagnosis	Intelligently diagnoses and analyzes the exception information by collecting and analyzing exception logs and provides recommended solutions		Supported	Partially supported
		Exception information collection	Ships the exception logs generated during job running to Tencent Cloud Log Service (CLS) in real time	Supported	Not supported
		Exception information search	Allows users to view job instance exception logs, which can be aggregated and searched by instance ID, process role, and process ID	Supported	Not supported
		Exception diagnosis	Allows users to quickly identify the causes of job exceptions using the exception log-based quick diagnosis feature and provides recommended solutions	Supported	Not supported
	Black window diagnosis	Automatically uploads the underlying OOM Dump, JFR, exception logs, and other exception information from actual jobs to the user's Object Storage Service (COS)		Supported	Not supported
		Information collection	Provides a list of files (including OOM Dump and JFR files) in the log directory for user analysis after the job process is ended	Supported	Not supported
	High availability	SLA guarantee	Distributed cluster design without any single point of failure, full-linkage monitoring, and rapid self-healing capabilities to ensure the availability of 99.9%	Supported	Not supported

Cost	Resource cost	Reduces computing resource costs by allocating and utilizing resources on demand based on actual business load		Supported	Not supported
		Auto scaling	Auto scaling based on the business load to ensure normal business operations and avoid resource waste	Supported	Not supported
		Fine-grained resources	Selection of fine-grained resources (e.g. 0.5 CU/process) for jobs to avoid resource waste	Supported	Partially supported
Security	Security isolation	Multi-dimensional isolation of the resource environment to ensure data security across tenants		Supported	Partially supported
		Space isolation	Exclusive tenant access to network space, computing resources, and storage resources, ensuring physical isolation between tenants	Supported	Partially supported
		Process isolation	Isolates job processes by container deployment to ensure process stability and security	Supported	Partially supported
	Access control	Ensures account security using the Tencent Cloud account system and SAML authentication mechanism, and achieves fine-grained access control		Supported	Not supported
		Account system	Connects to the Tencent Cloud account system and supports SAML authentication	Supported	Partially supported
		Fine-grained access control	Fine-grained permission management at cluster or job level to meet the needs of collaborative development by multiple persons	Supported	Not supported

Strengths

Last updated : 2023-11-07 15:27:34

One-stop development

Experience one-stop development and Ops with Web IDE or using TencentCloud API. Perform graphical development using multiple languages, manage metadata, dependencies, and full-lifecycle deployment, and conduct code debugging, operation monitoring, metric alerting, and smart diagnosis.

Seamless connection

Seamlessly connect to Tencent Cloud's popular data services, including message queue, database, data warehouse, and Elasticsearch Service, as well as open-source big data components. Expand connectors to access various external data systems (including self-built services in the cloud or IDC) as needed and enjoy full compatibility with Apache Flink for smooth migration to the cloud.

Sub-second latency

Process hundreds of thousands of records per second per CPU core with an end-to-end latency less than one second, and sustain tens of thousands of concurrent large-scale, real-time computing tasks.

Low costs

Reduce your per-core computing cost to 0.057 USD per hour by utilizing industry-leading proprietary server technology and assigning resources based on the actual business load, which is much more cost-effective than IDC-based self-built computing systems.

High security and stability

Isolate resource environments by space and process, ensure account security based on the Tencent Cloud account system and SAML authentication mechanism, ensure tenant data security with fine-grained access control, and provide full-linkage monitoring and alerting as well as self-healing in seconds to guarantee 99.9% service availability.

Expert service

Get end-to-end solutions and 24/7 Ops support services from the Tencent Cloud expert team.

Use Cases

Last updated : 2023-11-07 15:27:51

Stream Compute Service applies to the following scenarios:

Real-time extract, transform, and load (ETL) between heterogeneous data services

Real-time recommendation, real-time risk control, and database query acceleration

Real-time data warehouses, dashboards, and reports

Real-time monitoring, exception detection, and alerting

Real-time ETL

ETL is the process of loading data from a business system to the data sink after extraction, cleansing, and transformation. It aims to consolidate the scattered data complying with different standards in an enterprise to facilitate the decision-making of the enterprise.

Real-time financial risk management

Fast detection of financial risks can effectively reduce losses. By combining financial transaction big data with Stream Compute Service and introducing feature model algorithms, you can timely filter out abnormal transaction behaviors such as unauthorized transaction and implement comprehensive risk management to improve financial security.

IoT monitoring

Early detection of potential faults during the operations of industrial devices can greatly reduce production risks. Stream Compute Service can collect, aggregate, analyze, and filter device sensor data to implement monitoring and alerting for device exceptions in seconds, ensuring smooth operations of industrial production.

Targeted ecommerce recommendations

In ecommerce transactions, Stream Compute Service can extract feature variables in real time and promptly track users' desired categories to predict their consumption trends and provide basic capabilities for targeted recommendations. This helps enhance the shopping experience and drive purchases.

Database query acceleration

Relational databases face the challenges of weak performance and poor scalability in the case of massive data. Stream Compute Service can synchronize data from relational databases to ES in real time. ES provides capabilities of high concurrency and low latency query, SQL, and elastic scaling for massive data, meeting the demand for database query acceleration.