

TencentDB for MariaDB

Security White Paper

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

Security White Paper

Platform Security Features

Tenant Security Features

Security White Paper

Platform Security Features

Last updated : 2024-01-11 15:21:58

This document describes platform security features such as isolation, authentication, transfer encryption.

Isolation

The networks of different regions are fully isolated from each other, and Tencent Cloud services in different regions cannot communicate with each other over the private network by default. In addition, security groups and VPCs are also used for network isolation.

A **security group** is a stateful virtual firewall capable of filtering. As an important means for network security isolation provided by Tencent Cloud, it can be used to set network access controls for one or more Tencent Cloud services.

You can control the access permissions of a TencentDB for MariaDB instance in the following ways:

Create multiple security groups and specify different rules for them.

Assign one or multiple security groups to the TencentDB for MariaDB instance and use the security group rules to determine what traffic can access the instance and what resources can be accessed by the instance.

Configure a security group to allow only specified IP addresses to access the TencentDB for MariaDB instance.

VPC: it is a logically isolated network space defined in Tencent Cloud. Even in the same region, different VPCs cannot communicate with each other over the private network by default.

Authentication

Cloud Access Management (CAM) is a web-based Tencent Cloud service that helps you securely manage and control access to your Tencent Cloud resources. Using CAM, you can create, manage, and terminate users and user groups.

You can manage identities and policies to allow specific users to access your Tencent Cloud resources.

When using CAM, you can associate a policy to a user or user group. The policy can authorize or reject users' use of specified resources to finish specified jobs.

If you use multiple cloud services such as VPC, CVM, and TencentDB that are managed by different users sharing your cloud account key, the following problems may arise:

The risk of your key being compromised is high since multiple users are sharing it.

Your users might introduce security risks from misoperations due to the lack of user access control.

You can avoid the problems above by allowing different users to manage different services through sub-accounts. By default, a sub-account does not have permissions to use Tencent Cloud services or resources. Therefore, you need to create a policy to grant different permissions to the sub-accounts.

Transfer Encryption

The TencentDB for MariaDB console supports the HTTPS transfer protocol and standard network access protocols, guaranteeing your access security and meeting your needs for sensitive data encryption and transfer.

Tenant Security Features

Last updated : 2024-01-11 15:21:58

This document describes tenant security features such as MAR, auto failover, and data security encryption.

Multi-thread Async Replication (MAR)

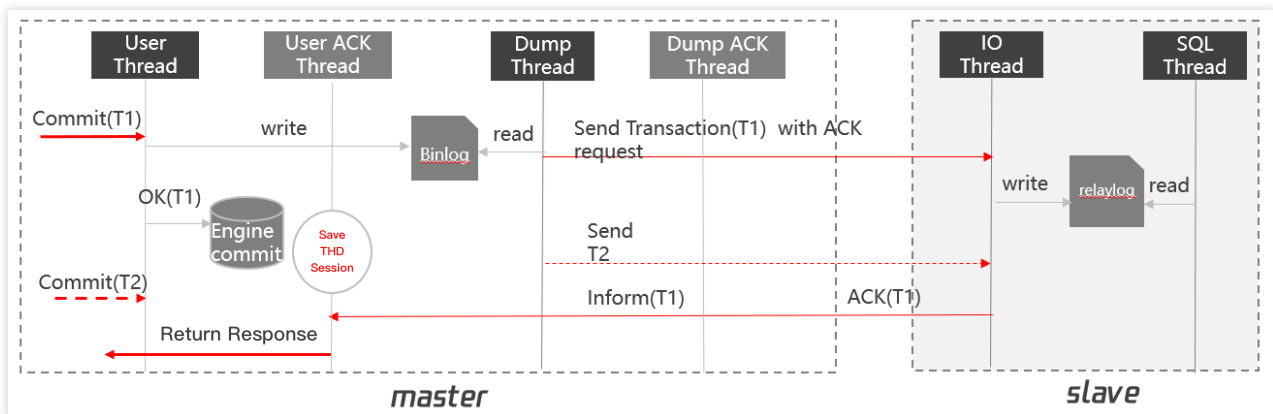
Background

As a database records data in it, to switch between multiple databases, the data in them must be in sync. Therefore, data sync is the foundation of the database high availability solution.

Currently, the open-source MySQL database supports the async and semi-sync data replication modes. However, in both modes, if a node failure occurs, the data may be lost or become incorrect or messy. In addition, the replication mode is serial, which has a low performance.

Solution

In Tencent Cloud's proprietary parallel multi-thread asynchronous replication (MAR, aka strong sync) solution based on the MySQL protocol, when a request is initiated at the application layer, only after a replica node successfully returns a message can the primary node respond to the application layer with a request success, ensuring that the primary and the replica nodes have the same exact data.

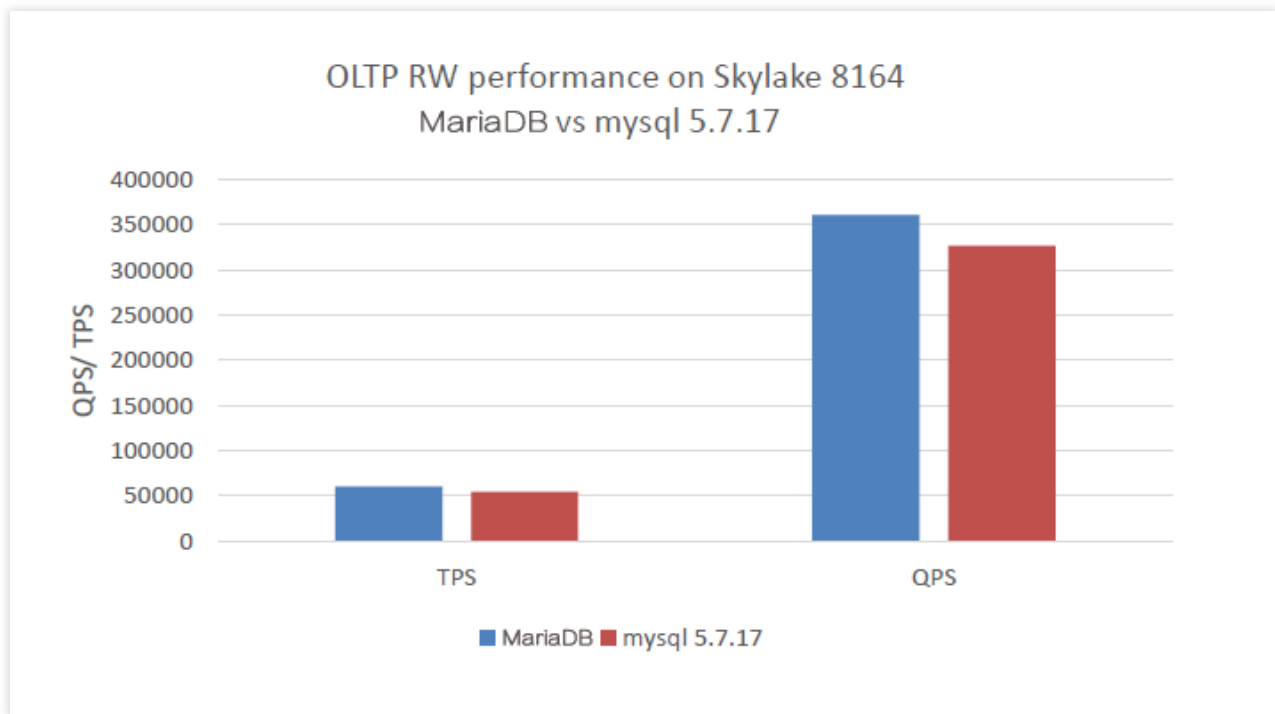


When you perform strong sync replication, the primary database will be hanged if it is disconnected from the replica database or if the replica database fails. If there is only one primary or replica database, the high-availability solution will be unavailable, because if only one single server is used, part of the data will be lost completely when a failure occurs, which does not meet the requirements for finance-level data security.

Therefore, based on MAR, MariaDB provides a downgradable strong sync solution, which is similar to the semi-sync

technology of Google but has a different implementation solution.

In addition, MariaDB MAR parallelizes the serial sync threads and introduces the worker thread capabilities, which greatly improve the performance. In the same cross-AZ (IDC with a latency of around 10–20 ms) test, the technical performance of MAR is around 5 times that of semi-sync replication on MySQL 5.6 and 1.5 times that of MariaDB Galera Cluster. In OLTP RW (mix read/write in primary/replica architecture), its performance is 1.2 times that of async replication on MySQL 5.7. The comparison of the specific performance tested by the Intel® technical team is as shown below:



Auto Failover and Recovery

In production systems, high availability solutions are often required to ensure uninterrupted system operations. As the core of system data storage and services, the availability requirement for the database is higher than that for computing service resources.

The high availability solution of MariaDB works by allowing the collaboration of multiple database services. In this way, if a database fails, another server will immediately take over its tasks, so the service will not be interrupted or be interrupted only for a very short period of time. This solution is also called primary/replica high availability.

Based on the general primary/replica high availability, MariaDB further supports the following advanced features: Auto failover, cluster member control, and node removal from the cluster are supported. For instance-level primary-replica switches, the virtual IP (VIP) will remain unchanged. The MAR policy ensures complete primary/replica data consistency in case of primary/replica failover, fully meeting the finance-level requirements for data consistency.

Auto recovery is supported. When a physical node carrying shards fails, the scheduling system will automatically try to recover the node. If the node cannot be recovered, it will be automatically replaced within 30 minutes. A new node will be rebuilt from backups and automatically added to the cluster, ensuring the high-availability architecture of instances for the long term.

Each node contains a complete replica of the data and can be switched according to the needs specified on the database management page.

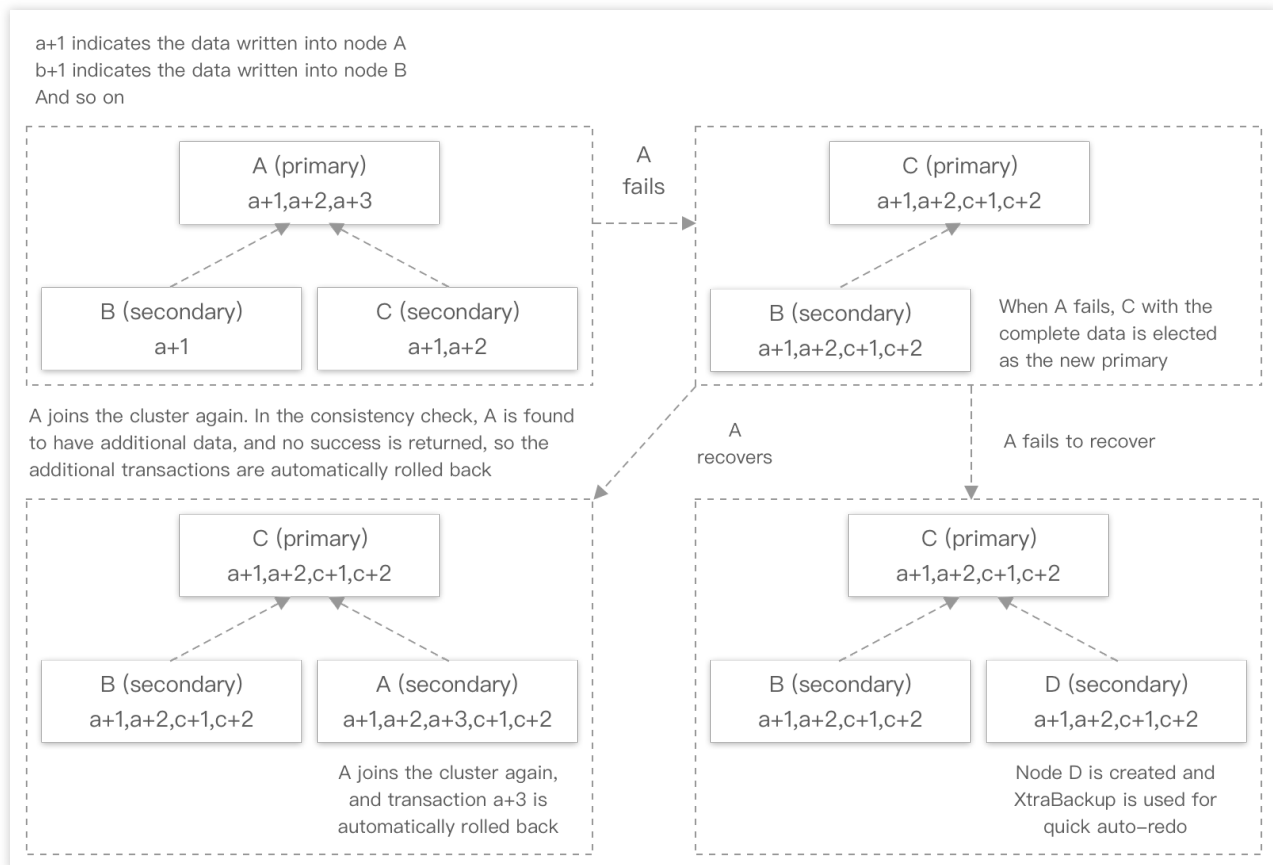
Do-not-switch configuration is supported. That is, failover will not be performed during the specified period of time. x86 PCs are supported, and there is no need to share storage devices.

Cross-AZ deployment is supported. Even if the primary and replica instances are in different data centers (regardless of whether they are in the same region), the data can be replicated through Direct Connect in real time. If the local node is the primary and the remote node is the replica, the local node will be accessed first, and if it fails or becomes unreachable, the remote node will be accessed. In addition, with the help of Tencent Cloud VPC, an intra-region active-active architecture can be implemented. That is, the business system can directly read/write the database in both data centers.

This feature provides MariaDB with multi-AZ disaster recovery capabilities, eliminating the operational risks with single-IDC deployment.

All MariaDB shards support the MAR-based high availability solution. If the primary database fails, the system will automatically select the optimal replica database immediately to take over the tasks. The switch process is imperceptible to users, the access IP remains unchanged, and 24/7 continuous monitoring is provided for the databases and underlying physical devices.

If a failure occurs, the system will automatically restart the database and relevant processes. If a node crashes and cannot be recovered, it will be automatically rebuilt from its backup files, as shown below:



Chinese and International Certifications

MariaDB complies with applicable Chinese information security standards and has earned many Chinese and international certifications on behalf of TencentDB.

MariaDB Platinum member

ACMUG and China Computer Industry Association - Open Source Database Committee (CCIA-ODC) Presidium member

ISO 27001

ISO 27001:2013

ISO 20000

ISO 20000-1:2011

ISO 22301

ISO 9001

ISO 27018

PCI DSS Level 1 Service Provider Qualification

SOC Audit

ITSS Cloud Service Advanced Certification

China's Cybersecurity Classified Protection Level 3 Filing and Evaluation for Public Cloud

China's Cybersecurity Classified Protection Level 4 Filing and Evaluation for Finance Cloud

Trusted Cloud Database Service Certification Issued by China Academy of Information and Communications Technology (CAICT)

Trusted Cloud User Data Security Protection Capability Assessment of China Academy of Information and Communications Technology (CAICT)

Trusted Cloud Gold Class Operations Special Assessment of China Academy of Information and Communications Technology (CAICT)

ITSS Certification

CSA STAR Gold certification and dual certifications for information security management system from CNAS and UKAS

Data Security Encryption

TencentDB for MariaDB supports connection encryption (SSL connection encryption) which ensures security of the traffic between the database and server.

SQL Firewall

SQL firewall is a security feature that filters out unauthorized SQL statements by analyzing the syntax of SQL statements sent by users. It works with SQL Engine to check whether an SQL statement is on the predefined list of unauthorized SQL statements so as to filter out and block it accordingly, which effectively prevents SQL injection attacks.

Note:

SQL firewall can be used together with Tencent Cloud services such as Web Application Firewall (WAF). Taking into account the business conditions and SQL complexity, there are currently no preset rules in the MariaDB SQL firewall.

Comprehensive Security Audit

Security audit is one of the most important tracing methods. Therefore, China's Cybersecurity Classified Protection Certification (Level 3) stipulates that an information system should support auditing. MariaDB provides audit capabilities at the following three layers to deliver complete security protection:

Security audit for the OPS system, which is implemented by the operation logs of the Chitu operation system.

Security audit for the database system, which is implemented by Tencent Cloud's proprietary database audit system.

Security audit for the server operating system, which is implemented by Tencent Cloud's proprietary Tiejiangjun system.

Note:

In public cloud, all security audit features are configured by default.

In private cloud, system operation logging (Chitu system) is configured by default, while database SQL audit and server operation audit features are optional.

Kernel-Level Security Policies

MariaDB provides various open-source security solutions at the database kernel level, some of which have earned the recognition of the community. The following are some kernel security measures:

Slow deletion

If you run the `drop table` or `alter table ... drop partition` command, the database will not delete the tablespace file immediately. Instead, it will rename the file, gradually shrink it on the backend, and finally delete it. This feature can avoid system performance fluctuation caused by I/O load surges in the server's file system when a large tablespace file is deleted in one single request.

Accidental metadata deletion prevention

Only authorized users can log in to the system and delete metadata tables, which helps prevent business unavailability due to faulty operations.

Banning of plugin installation by unauthorized users

The database service provides standard APIs for users to implement custom features, but hackers usually exploit this vulnerability to launch attacks. Therefore, only specified admin users can mount plugins.

Banning of unauthorized user access to the physical server file system

To prevent hackers from bypassing the security system by means such as file selection, file injection, and path detection, unauthorized users are blocked from accessing the directory structure and file system of the physical server.

Data Termination

When you terminate your TencentDB for MariaDB instance, all data (including backup data) stored in it will be destroyed. Tencent Cloud will not retain the data or actively recover your instance.

Suggestions on 1-DC Disaster Recovery Deployment

When deploying 1-DC disaster recovery, you should prevent the following failures for your database cluster:

Single points of failures on devices such as data center switch, load balancer, and ENI.

Single points of failures on devices such as rack power, fan, and cooler.

Single points of failures on database server hardware.

Therefore, we recommend you satisfy at least the following requirements for 1-DC disaster recovery deployment:

Deploy at least active-active disaster recovery for network devices such as switch and load balancer.

Deploy one primary and two replicas for the database server and management and scheduling system.

Deploy different devices of the same module across racks.

Deploy a data backup module.

Suggestions on 2-Region-3-DC Deployment

2-region-3-DC deployment is to add a disaster recovery center based on 1-region-2-DC deployment. The two disaster recovery instances are synced over a data communication network (DCN) to ensure data consistency.

