GaussDB Evolution and Research Directions

Ronen Grosman, Yao Wu

Huawei Canada
ronen.grosman@huawei.com,wuyao18@huawei.com

1 Introduction to GaussDB

GaussDB is a distributed relational database from Huawei. It supports intra-city cross-AZ deployment [1] with zero data loss. With a distributed architecture, GaussDB supports petabytes of storage and contains more than 1,000 nodes per DB instance. It is highly available, secure, and scalable and provides services including quick deployment, backup, restoration, monitoring, and alarm reporting for enterprises. The overall architecture of a distributed instance of GaussDB is as Figure 1, which is a share-nothing distributed cluster with log replication HA. GaussDB is built on top of openGauss [1], which is a multi-core-oriented open-source relational database that provides ultimate performance, full-link service and data security, AI-based tuning [2], and efficient O&M capabilities. This leading database at enterprise level is developed in collaboration with global partners and is released under the Mulan Permissive Software License v2. openGauss deeply integrates Huawei's years of R&D experience in the database field and continuously builds competitive features based on enterprise-level scenario requirements.

2 Challenges and Opportunities of GaussDB

Huawei is actively working with industry and ecosystem partners worldwide, promoting joint innovation among businesses, academia, research institutes, and users to create practical value based on the needs of different industries. We promote cross-domain and cross-technology collaboration in various forms, in order to tackle real-world problems that different industries face. The openGauss kernel was derived from PostgreSQL and focused on building advance features for architecture, transactions and storage engines with performance optimization. It is deeply optimized for ARM architecture and retains compatibility with x86 architecture. In this presentation, we will cover the main challenges and opportunities for openGauss and GaussDB, e.g., disaggregated architecture, hybrid data processing, unified storage engine [3]. Huawei will continue to work with developers and partners in building the openGauss community for a prosperous global ecosystem.

[1] https://docs.opengauss.org/en

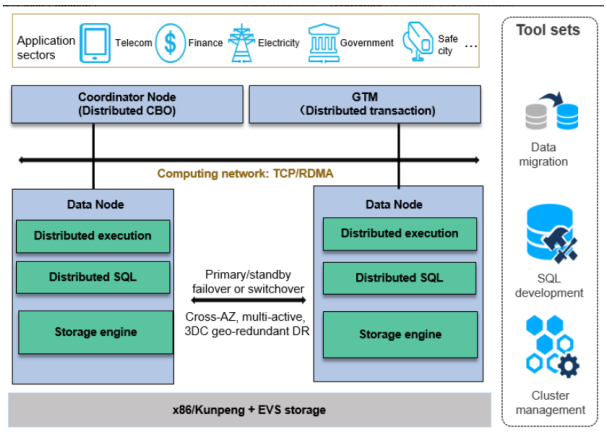

Figure 1: GaussDB Architecture

3 Open Discussions

In this section, we will discuss a few open challenges from an industry perspective towards the database community [4].

References

[1] Weixing Zhou, Qi Peng, Zijie Zhang, et al. Geogauss: Strongly consistent and light-coordinated oltp for geo-replicated sql database. *Proceedings of the ACM on Management of Data*, 1(1):1–27, 2023.

[2] Guoliang Li, Xuanhe Zhou, Ji Sun, et al. opengauss: An autonomous database system. *Proceedings of the VLDB Endowment*, 14(12):3028–3042, 2021.

[3] Hillel Avni, Alisher Aliev, Oren Amor, et al. Industrial-strength oltp using main memory and many cores. *Proceedings of the VLDB Endowment*, 13(12):3099–3111, 2020.

[4] Guoliang Li, Haowen Dong, and Chao Zhang. Cloud databases: New techniques, challenges, and opportunities. *Proceedings of the VLDB Endowment*, 15(12):3758–3761, 2022.