



Kuaishou Unleashes the Potential of Intel® Optane™ Persistent Memory With KVDK, Achieving More Than 20x Acceleration of GraphDB Performance¹

Summary

Increasing business volume and rapid growth of data represent a major challenge for businesses in processing data. More and more businesses are increasing their commitment to Graph Database (GraphDB) innovation and investment in order to effectively process large volumes of data, and provide support for applications such as user profiles, personalized recommendations, risk management, knowledge graphs, and social networks. While providing high-speed data processing and analysis performance, GraphDB is also extremely demanding on storage performance. Traditional NVMe SSDs, facing performance bottlenecks, are unable to keep up with the demands of large-scale GraphDB applications.

In order to support rapidly growing business with high-performance database capabilities, Kuaishou has implemented the Kuaishou GraphDB application based on open-source technology, and built high-performance, large-capacity storage pools with Intel® Optane™ persistent memory (PMem). Kuaishou Graph DB also utilizes Key-Value Development Kits (KVDK) designed specifically for Intel® Optane™ persistent memory, in order to fully leverage the performance advantage of Intel® Optane™ persistent memory, and provide advanced functionality for database workloads.

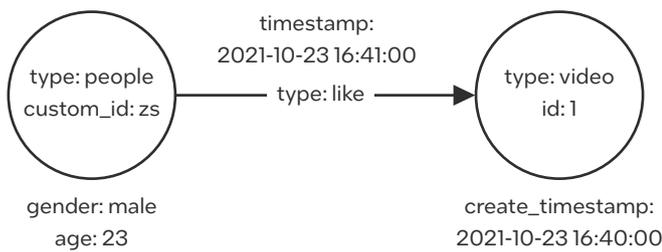
Background: Kuaishou's GraphDB application and performance bottleneck

While traditional relational databases are able to meet the needs of most scenarios, the increase in volume and depth of business data has led to an increase in demand for processing complex data relationships. In operations such as social networks, recommendation systems and data analysis, user properties often form massive entity diagrams with tremendous data volumes. Relational databases cannot compute results within a reasonable span of time. In order to achieve better connection between data, businesses need database technology that stores relational information as entities, flexibly expanding data models - GraphDB.

GraphDB is generally described by key data such as vertex, edge, and property. A vertex represents an entity, such as a person, a video or a device. Edge represents the relationship between vertexes, such as one user following another, or a device's ownership by one person. Property represents additional information beyond which is represented by the vertex or edge.

GraphDB smoothly integrates the explicit modelling and query of graphics into a standard database environment. It provides key attributes of object-oriented modeling for standard applications, such as the object class, object identifiers and object properties that comprise the hierarchy. At the same time, the model also allows graphs to be displayed by categorizing object classes into basic classes, link classes and path classes. The object of these classes may be considered the graphic's node, edge, and explicit storage path.

Compared to traditional relational databases, GraphDB excels at representing and processing relationship between data. Excellent bottom-up compatibility with the data structure of graphical models makes it extremely quick at data query and analysis. GraphDB also boasts a highly flexible data model, allowing the customer to adjust their graphical data structure model based on their needs, adding



KGraph	Graph	Type	Vertex/Edge	Property
Relational	Database	Table	Row	Column

Figure 1. Kuaishou GraphDB graphical data characteristics

and removing vertexes and edges, and scaling their graphical model as needed. GraphDB also meets agile development needs, enabling quick and simple updates.²

In spite of so many advantages, GraphDB faces a number of bottlenecks limiting its performance, including: In-memory Graph solutions have long recovery time and are unsuitable for large data sets, and linear scalability needs to be improved in order to meet business needs dynamically. In terms of performance, GraphDB physical nodes typically have performance requirements in excess of 200 million QPS and low single-vertex query latency. Meanwhile, due to the performance bottleneck of traditional NVMe SSDs, the performance improvement of GraphDB is limited. Because of this single-node performance limitation, extension nodes are often needed for traditional GraphDB to support large-scale data, resulting in increased server procurement, deployment and operation costs, a less-than-ideal choice from an economic perspective.

That is why Kuaishou began to look for innovative storage media and storage engines targeted at GraphDB applications, in order to maximize single-node throughput, reduce latency, and resolve bottlenecks in extensibility and recovery time, while achieving better control of the total cost of ownership (TCO) of GraphDB applications.

Solution: Optimizing GraphDB through Intel® Optane™ Persistent Memory (PMem) and KVDK

In order to resolve GraphDB's performance and latency bottlenecks, Kuaishou has worked with Intel to replace the bottom-level storage media of the GraphDB system with Intel® Optane™ persistent memory (PMem), and employing KVDK as the storage engine of GraphDB, replacing the RocksDB that was previously in use. Figure 2 shows a comparison of the old and new frameworks.

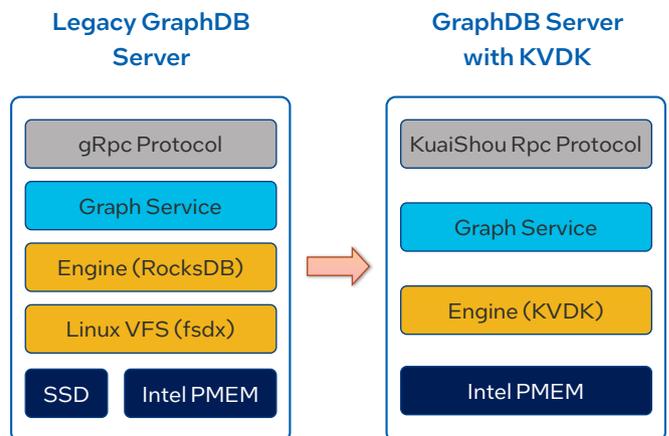


Figure 2. Comparison of Kuaishou's old and new frameworks

² RH Güting, Iv P I . GraphDB: A Data Model and Query Language for Graphs in Databases[J]. proc.int.conf.on very large data bases, 1995.

Intel® Optane™ persistent memory offers tremendous performance advantages over traditional NVMe SSDs. Intel® Optane™ persistent memory is a revolutionary technology that combines high-speed, high cost-to-performance ratio, large capacity, persistent data protection, and high-level encryption, offering two operating modes: Memory Mode and App Direct Mode. In Memory Mode, PMem offers 128 GB, 256 GB and 512 GB solutions that work in the same way as standard volatile memory, but at a lower cost, achieving higher memory capacity while maintaining the same system budget.

In App Direct Mode, applications and operating systems have access to two different types of memory. The system decides whether to read from or write to DRAM or PMem. Operations with stringent latency requirements but not permanent data storage can be executed in DRAM, while large data structures and data that require long-term storage are stored in PMem. In this mode, data is retained in memory even when the system is powered off.

Kuaishou replaced RocksDB with KVDB, in order to fully leverage the performance advantage of Intel® Optane™ persistent memory, and provide advanced functionality for database workloads. KVDB is a key-value repository written in C++, designed specifically for Intel® Optane™ persistent memory, providing a unified API for volatile and persistent scenarios, boosting the performance of Intel® Optane™ persistent memory. In addition to providing basic API for key-value storage, it also provides advanced functionality such as transactions and snapshots.

KVDB is equipped with a persistent memory allocator to prevent thread allocation racing and improve core extensibility. The allocator optimizes sequential access of data based on the characteristics of Intel® Optane™ persistent memory, improving the performance of RocksDB. The allocator also adjusts the access mechanism of Intel® Optane™ persistent memory, saving persistent memory bandwidth.

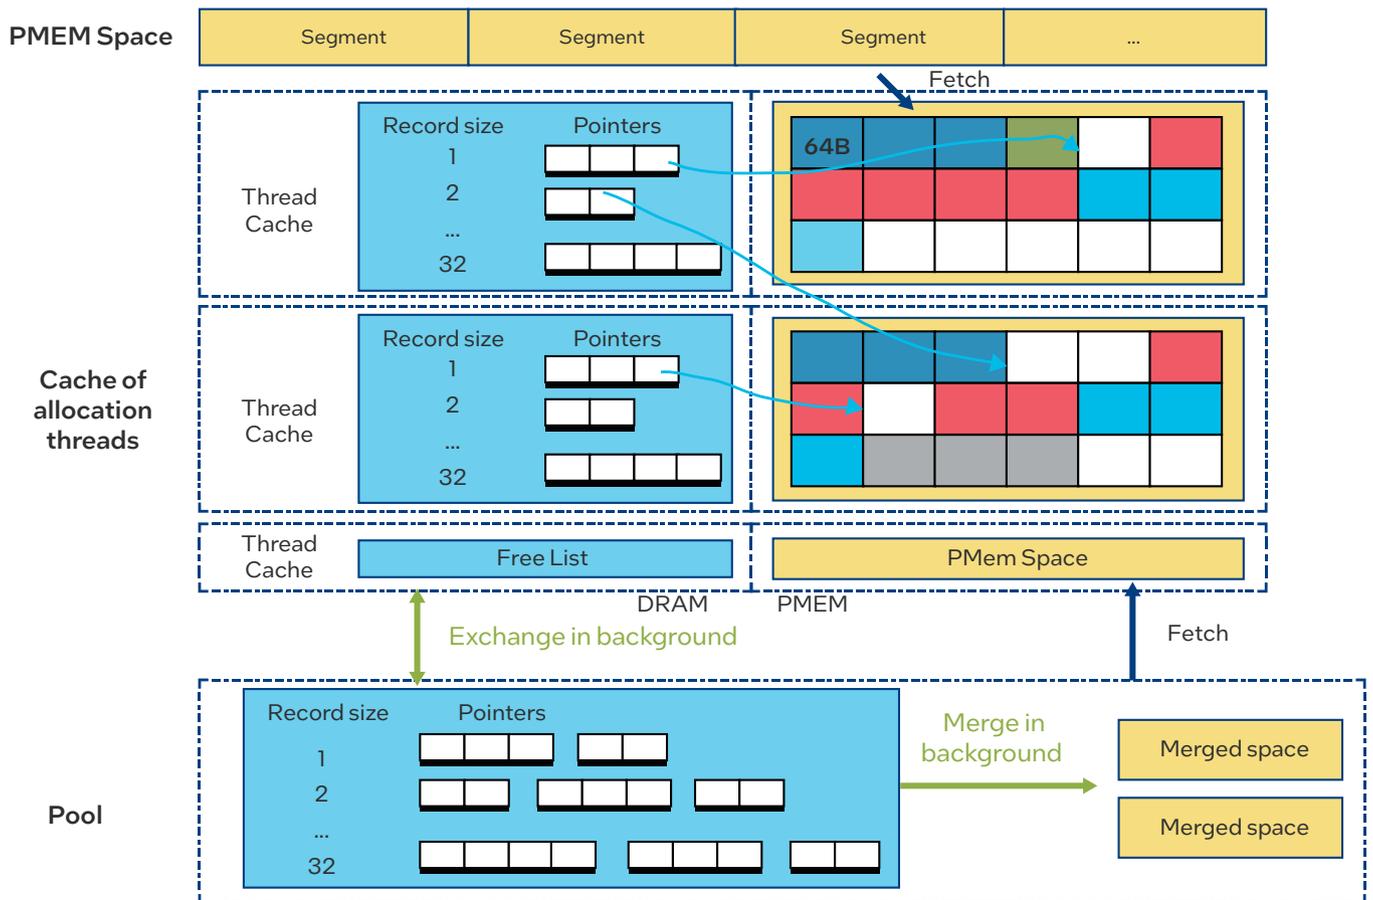


Figure 3. Structure of KVDB's persistent memory allocator

KVDK performs unsorted indexing through hash algorithms, and achieves sorted indexing through hash + skiplist. This has many advantages, including cache friendliness, lock-free read/write, fine-grained write locking, and fewer and predictable persistent memory costs, improving performance while reducing resource usage.

To further leverage the advantages of KVDK, Kuaishou has removed the redundant Linux system from the GraphDB build, and developed a new, faster RPC framework to meet its needs.

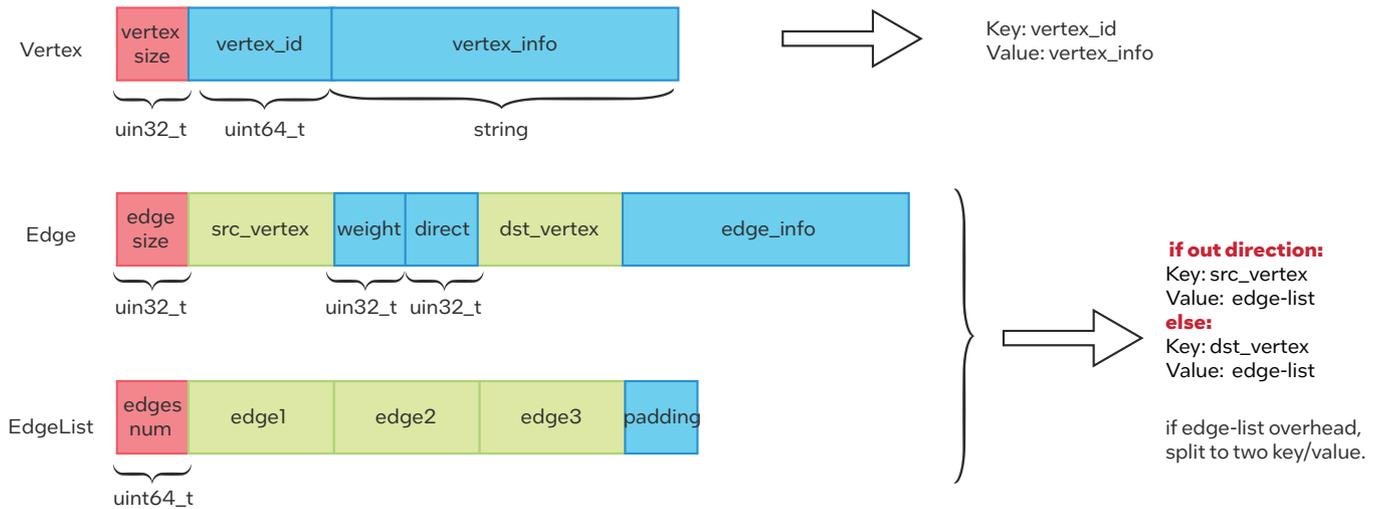


Figure 4. Kuaishou's targeted optimizations in GraphDB

Benefit: More than 20x performance gains¹

Kuaishou performed three tests to validate the performance gains from KVDK's Kuaishou GraphDB solution. The first test, which used a data set of 30,000,000 vertexes + 1200,000,000 edges, was designed to test the time needed to randomly search the build data. The second test used a data set of 1,000,000 vertexes + 100,000,000 edges, testing latency through building popular search data + 4 degree bfs. The third test involved computing the top 10 out of top N results.

The KVDK-based Kuaishou GraphDB significantly outperformed the legacy RocksDB solution, leading the latter by up to 26x in some cases.³

Kuaishou has significantly optimized GraphDB performance by utilizing Intel® Optane™ persistent memory and KVDK, meeting the performance demands of a database with tens of billions of entries, accelerating search and computing, providing complex data relational computing capability for applications such as user profiles, personalized recommendations, knowledge graphs, and social networks, supporting the effective operation of Kuaishou's innovative applications.

Increase in single-node performance means that under the same performance demands, Kuaishou needs less server resources to reach GraphDB performance goals, saving costs in terms of hardware procurement, space, electricity, and maintenance.

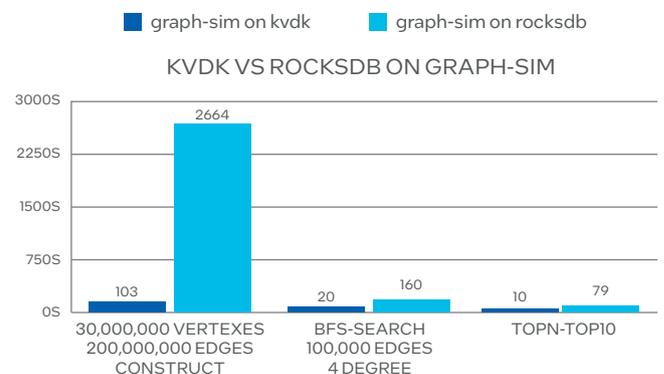


Figure 5. Kuaishou GraphDB test data before and after switching to KVDK

“High-performance storage media and a storage engine with targeted optimization is key to the optimization of GraphDB performance. By working with Intel and utilizing the KVDK repository, we are able to fully leverage the performance advantages of Intel® Optane™ persistent memory, resolving bottlenecks in extensibility and recovery time, building a cutting-edge GraphDB system that provides complex data relational computing capability for various applications, and building the foundation for deep applications.”

– Liu Xin
Kuaishou

Prospect: Empowering Innovative Digital Applications with Intel® Optane™ Persistent Memory

As an innovative storage media that bridges the gap between DRAM and storage, Intel® Optane™ persistent memory has been proven to be highly valuable in I/O-intensive workloads. It can replace the role of DRAM for in-memory databases, providing high-capacity, affordable memory, or provide a high-performance storage media for databases such as GraphDB, providing performance that far surpasses traditional SSDs. Repositories such as KVDK further leverage the advantages of Intel® Optane™ persistent memory, providing powerful foundational resources for databases.

Kuaishou's practical experience with GraphDB demonstrate that coordinated hardware and software optimization for Intel® Optane™ persistent memory can effectively improve database performance. Kuaishou and Intel are working together to promote innovation in IT infrastructure, improving performance, extensibility, stability and availability, building a firm foundation for Kuaishou's rapidly growing innovative internet business.

About Kuaishou

Kuaishou started out in Mar 2011 as a mobile app for creating and sharing GIF images. In Nov 2012, Kuaishou transitioned to offering a community for users to share video clips, recording their everyday life and creative work. Today, Kuaishou is a popular social platform with millions of original videos uploaded daily. Kuaishou's applications had an average of 346 million daily active users as of Q1 2022.

About Intel

Intel (NASDAQ: INTC) is an industry leader that focuses on developing technologies that change the world, drive global progress, and enrich human lives. Driven by Moore's Law, Intel is dedicated to the continuous innovation of semiconductor design and manufacturing to provide customers with solutions to major challenges. By integrating intelligence with the cloud, network, edge, and various computing devices, Intel unlocks the potential of data and helps improve both business and society. For more information on Intel innovation, refer to our China News Center at newsroom.intel.com and our official website at intel.com.



^{1,3}Results provided by Kuaishou and were based on internal tests as of June 2022. Please contact Kuaishou for further details.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.