

Alibaba Cloud POLARDB*: An Innovative Cloud-Hosted Database

POLARDB uses software-defined storage powered by Intel® Optane™ Solid State Drives for data centers to improve performance and reduce latency.



Relational database management systems (RDBMSs) are the foundation of enterprises worldwide. They are also a prerequisite for company operations and digital transformation.

An on-premises RDBMS, though, can require significant investment in infrastructure, time, cost, and personnel to keep modernized, scale to meet business growth, and ensure it's capable of supporting digital transformation. One solution to these obstacles is cloud computing, which has evolved into a mainstream technology used by organizations of all sizes.

Relational databases are widely available as cloud-hosted services, known as database-as-a-service (DBaaS) solutions. But a cloud-based RDBMS can present its own challenges, which exposes a need for cloud-hosted relational database solutions to digitally transform. Popular databases include MySQL* and Oracle Database*; but when hosted on-premises or managed as a service in the cloud, they can be difficult to secure, manage, and scale, and they might not deliver the needed performance or quality of service (QoS).

Alibaba Cloud has developed an innovative, commercial-grade, cloud-hosted relational database, POLARDB*, that addresses these issues and offers enterprises a new DBaaS model. At the foundation of POLARDB is a distributed shared-storage architecture, which was engineered to use Intel® Optane™ DC Solid State Drives (SSDs) and Intel® 3D NAND SSDs.

Traditional Challenges of Cloud-Hosted Databases

In a public cloud environment, growth in the number of users, user services, and data volumes can result in issues with backups, performance, migration, upgrades, disk capacity, and delays caused by the binary log. The time needed for scaling, backing up, and migrating data increases as the data volume grows. It can take hours, or even days, to back up terabytes of data. Upgrades, scaling, backing up, and data migration can require that a database be taken offline and restarted, which can impact the business, its users, and its IT personnel. Migrating data can slow the performance of mission-critical applications. Maintaining everyday business operations can require significant CPU resources, which can cause delays and reduce price/performance when compounded over millions of transactions.

POLARDB Redefines Hosting a Database in the Cloud

POLARDB innovates on the existing DBaaS model. It uses an innovative software-defined scale-out system that incorporates fast, low latency Intel® Optane™ DC SSDs and low-cost, high-capacity Intel 3D NAND SSDs. Enterprises with large, growing volumes of data—and that require consistent low latency, high throughput, and a high QoS, combined with a good price/performance ratio—can benefit from POLARDB.

The storage paradigm of POLARDB, combined with other innovations, delivers the performance and availability of a commercial-grade database in a DBaaS model. Alibaba Cloud engineered POLARDB with a focus on business customers' on-cloud costs, online transaction processing (OLTP) performance, business continuity, business growth, and security. All of the hardware and software are engineered to work effectively together, including with the underlying Intel hardware, in a cloud-native design developed for high reliability and high availability. The architecture provides high performance and requires roughly one-tenth of the cost of a traditional on-premises relational database.²

The architecture of POLARDB enables it to deliver up to six times the throughput of standard MySQL 5.6 and 5.7 when MySQL is run on NAND-based Peripheral Component Interconnect Express* (PCIe*) and NVM Express* (NVMe*) SSDs. And POLARDB delivers throughput of up to one million input/output operations per second (IOPS) with latency of just milliseconds.² POLARDB can also scale up/in or out/down on demand in less than 10 minutes. It supports a capacity of up to 100 TB per database instance compared to just 64 TB for similar solutions. Storage capacity and clones can automatically scale. And POLARDB is 100-percent backward compatible with MySQL 5.6 and 5.7.

Intel® Optane™ DC SSDs Optimize Storage Performance

The POLARDB distributed storage design with Intel® Optane™ DC SSDs and Intel 3D NAND SSDs provides excellent storage efficiency, high QoS, high IOPS, high throughput, and consistent performance. Intel® SSDs also help keep costs and performance tuned. Intel® Optane™ DC SSDs help Alibaba Cloud eliminate data center storage bottlenecks, accommodate bigger, more affordable datasets, accelerate applications, reduce transaction costs for latency-sensitive workloads, and improve the overall data center total cost of ownership (TCO) for POLARDB.

By combining high-speed, high-capacity Intel® Optane™ DC SSDs with high-throughput Intel 3D NAND SSDs at the software layer, POLARDB implements a hybrid storage layer that can deliver consistent low latency, high throughput, high QoS, and a lower overall price/performance ratio compared to using only Intel 3D NAND SSDs for the storage layer.

With Intel® Optane™ DC SSDs and Intel 3D NAND SSDs, Alibaba Cloud has realized:

- A 30-percent improvement on queries per second (QPS), for an improved customer experience with POLARDB¹
- A 76-percent improvement on 95th-percentile latency with POLARDB¹

The POLARDB Architecture

The software-defined POLARDB architecture uses distributed storage, the Storage Performance Development Kit (SPDK), remote direct memory access (RDMA), and other innovations to create a high performance, scalable, cloud-based database. The POLARDB architecture separates the compute and storage pools. When CPU resources and memory are insufficient, the computing resource pool is expanded separately from the storage pool; and when the capacity or IOPS is low, the storage resource pool scales independent of the compute pool.

Intel® Optane™ DC SSDs power a

30% IMPROVEMENT

of queries per second (QPS), for an improved customer experience with POLARDB¹

76% IMPROVEMENT

on 95th-percentile latency with POLARDB¹

100 TB CAPACITY

for each POLARDB database instance

Intel® Optane™ DC SSDs and Intel® 3D NAND SSDs for POLARDB¹

POLARDB was originally developed to meet the needs of Alibaba Cloud's online shopping center. Alibaba Cloud initially made use of Intel® Optane™ DC SSDs and Intel 3D NAND SSDs for POLARDB for shopping center peak usage. Alibaba Cloud then realized a performance increase of up to six times for MySQL* compared to running MySQL only on NAND-based PCIe* NVMe* SSDs.¹

Today, by using the Intel® Optane™ SSD DC P4800X Series for the POLARDB offering, Alibaba Cloud makes use of SSDs that:

- Are the most responsive data center SSDs available with Intel® Optane™ technology³
- Offer a capacity of up to 750 GB
- Have enabled increased performance for POLARDB

Using Intel 3D NAND SSDs, in addition to Intel® Optane™ DC SSDs, brings economic data storage with outstanding quality, reliability, and advanced manageability and serviceability to minimize disruptions for POLARDB.

By working closely with customers like Alibaba Cloud to innovate databases, Intel is helping improve the enterprise database experience for customers worldwide.

The POLARDB distributed storage architecture utilizes three replicas or nodes instead of a local store. One replica is used for read and write, and the other two for read-only. When a new read-only instance is needed, no data needs to be copied; and no matter how large the new data volume is, it can be completed within five minutes compared to the 70 hours it can take to add an instance for a traditional database with more than 3 TB of data.¹ The primary instance and the read-only instances share the same storage resources, which increases read-only copy performance at no additional cost for new storage. Users pay only for CPU and memory consumed by the read-only instances. POLARDB uses this distributed storage configuration to locate data and scale faster than a traditional database.

Further performance benefits with POLARDB result from:

- Using the RDMA over Converged Ethernet (RoCE) network protocol to lower latency, reduce the CPU load, and increase bandwidth
- Bypassing the Linux* kernel with a software-defined stack that enables POLARDB to run in user mode, which helps decrease overhead
- Removing the need for the binary log, which reduces transaction times and input/output (I/O) overhead

Access a Commercial-Grade Database as a Service

A DBaaS implementation can free an enterprise IT team from managing databases in-house and from buying and maintaining the hardware that an on-premises database requires. The IT team can then focus on innovating the core business needs that help the enterprise digitally transform and improve its competitive opportunities. Where other DBaaS solutions can present challenges, POLARDB from Alibaba Cloud can help eliminate those challenges with a commercial-grade DBaaS solution. POLARDB empowers enterprises with an on-demand solution that can help maintain a good price/performance ratio. As a commercial-grade database, it helps enterprises manage large volumes of data while enabling low latency, high throughput, high QoS, and rapid scaling.

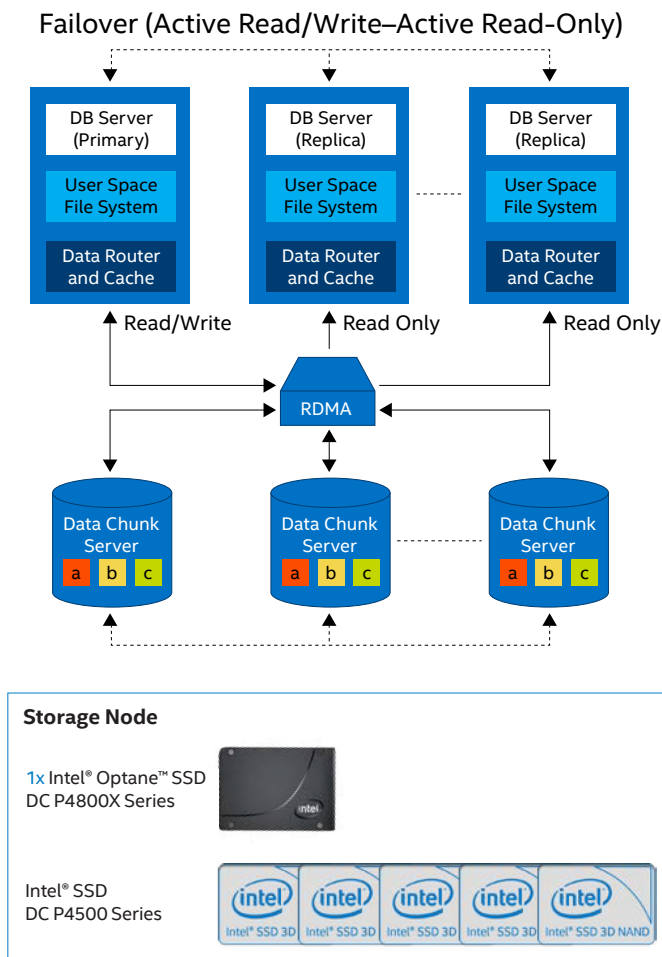


Figure 1. The Alibaba Cloud POLARDB* architecture with Intel® SSDs

Each POLARDB storage node uses one Intel® Optane™ SSD DC P4800X Series drive for fast journaling and indexing and to support Intel SSD DC P4500 Series drives, which are used for data storage. The SPDK writes data to each Intel 3D NAND SSDs' NVMe interface without having to access synchronized memory. Before implementing Intel Optane™ DC SSDs, journaling and data were collocated on NVMe-based Intel SSDs for the data center, which required accessing synchronized memory. Bypassing synchronized memory allows POLARDB to increase performance and reduce latency.

About Alibaba Cloud

Alibaba Cloud is the largest e-commerce business in China, and it is a rapidly growing cloud service provider. Alibaba Cloud delivers a range of global cloud computing services that power online businesses worldwide. The company also manages and maintains the e-commerce ecosystem for its parent company, Alibaba Group.

Alibaba Cloud's global operations are registered and headquartered in Singapore. It also maintains offices in Dubai, Frankfurt, Hong Kong, London, New York, Paris, San Mateo, Seoul, Singapore, Sydney, and Tokyo.


Learn More

Alibaba Cloud POLARDB: alibabacloud.com/product/apsaradb-for-rds

Intel® Optane™ technology: intel.com/optane

Intel® SSD Data Center Family: [intel.com/content/www/us/en/products/memory-storage/solid-state-drives/data-center-ssds.html](https://www.us/en/products/memory-storage/solid-state-drives/data-center-ssds.html)

SPDK: <https://software.intel.com/en-us/articles/introduction-to-the-storage-performance-development-kit-spdk>



¹ Testing completed by Alibaba Cloud on two 2-socket CPUs with 32 cores each, 256 GB DRAM, and dual 25 gigabit Ethernet (GbE) network interface controllers (NICs). 1 x Alibaba Cloud Elastic Compute Service* (ECS*), 1 x POLARDB* (1 x database and 3 x PolarStore* replica). Benchmark: Amazon Aurora* (sysbench*), thread = 1, iodepth = 1, 100 percent write, 4 TB database, request blocksize = 16 KB.

² Alibaba Cloud. "100TB Capacity and 6x Performance Improvement with Alibaba Cloud PolarDB." April 2018. alibabacloud.com/blog/100tb-capacity-and-6x-performance-improvement-with-alibaba-cloud-polardb_578294.

³ Responsiveness defined as average read latency measured at queue depth 1 during a 4K random write workload. Measured using FIO 2.15*. Common configuration: 2 x Intel® Xeon® processor E5-2699 v4 (2.20 GHz, 22 cores), CentOS 7.2*, kernel 3.10.0-327.el7.x86_64, 396 GB DDR RAM at 2,133 MHz. Intel® drives evaluated: 375 GB Intel® Optane™ SSD DC P4800X Series and 1,600 GB Intel® SSD DC P3700 Series. Samsung* drives evaluated: Samsung SSD PM1725a*, Samsung SSD PM1725*, Samsung SSD PM963*, Samsung SSD PM953*. Micron* drive evaluated: Micron 9100 PCIe NVMe SSD*. Toshiba* drive evaluated: Toshiba ZD6300*. Test: queue depth 1, random read, 4K latency; queue depth 1, random read/write, 4K 70-percent read latency; queue depth 1, random write, 4K latency using FIO 2.15.

Performance results are based on testing as of July 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

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, the Intel logo, Intel Optane, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2018 Intel Corporation. All rights reserved.

Printed in USA

0818/SH/PRW/PDF

Please Recycle

337450-002US