

NetEase Games Reduces Costs and Improves Efficiency with Intel® VROC



"The performance, cost, and availability of storage systems have a huge impact on system capabilities, end-user experience, and return on investment in gaming scenarios. Through PostgreSQL master-slave architecture, hardware storage solution optimization, and the replacement of RAID controller cards with Intel® VROC, we have laid a solid foundation for growth by significantly improving the storage performance of our PostgreSQL databases and reducing costs."

– **Chen Xiaosheng**

Deputy Technical Director,
Infrastructure Department,
NetEase Games

Overview

Redundant array of independent disks (RAID) has been widely used in server systems to improve the storage throughput and fault tolerance of data. However, the application of RAID technology also creates challenges. For example, the use of RAID controller cards may add to system deployment costs, create another point of failure, introduce bandwidth limitations from the upstream PCIe link, and create software stack overhead to realize SCSI and NVMe compatibility. In certain conditions, these issues may affect cloud and data center performance.

As one of the world's leading game companies, NetEase has always been at the forefront of independent industry research and development, and has launched and operated a large number of globally popular games. As NetEase's business continues to grow rapidly, the scale of its cloud data center resources is also expanding. In order to improve the stability of its cloud data center servers, take full advantage of the performance capabilities of NVMe SSDs, and mitigate the complexity, cost, and supply challenges of RAID controller cards, NetEase has adopted Intel® Virtual RAID on CPU (Intel® VROC) technology enabled by Intel® Xeon® Scalable processors. Intel® VROC provides performant and highly reliable local NVMe storage for game database applications through NVMe RAID, along with reliable OS data protection through SATA RAID. The solution both reduces costs and improves efficiency in cloud computing and database scenarios.

Challenge: Traditional RAID Solutions Face Cost and Performance Challenges

As enterprise data grows rapidly, high-level applications such as games, e-commerce, transactions, and search have placed higher demands on the performance and latency of storage systems. Data security and reliability have also become an increasing concern. The RAID solution has been traditionally used to meet these demands as it improves storage performance by writing data to multiple disks in an array and achieves data redundancy through mirroring, parity information, and other methods.

NetEase primarily uses the SATA RAID and NVMe RAID storage configurations. SATA RAID is primarily used for boot disks, while NVMe RAID is primarily used for data disks that store PostgreSQL and other critical database systems. PostgreSQL is a stable and scalable open-source relational database management system that supports both relational and non-relational storage, and is often used in web, financial system, and gaming scenarios. It is one of the databases used by NetEase.

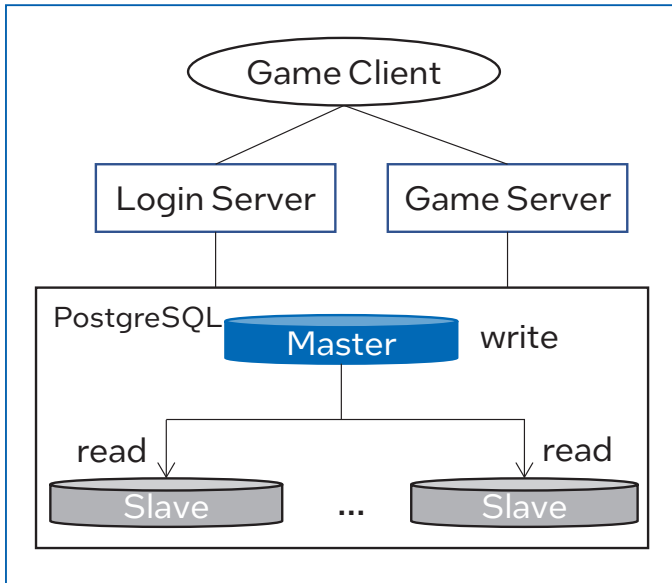


Figure 1. PostgreSQL architecture of a popular NetEase game

Due to the design of game architectures, PostgreSQL relies on the performance of the master database rather than the horizontal scaling of slave databases, meaning that improved hardware performance provides more direct results than optimization at the software level. As such, NetEase has been upgrading its data disks from SATA SSDs to NVMe SSDs and optimizing the selection of RAID solutions.

Server RAID solutions typically rely on RAID controller cards, which provide support for multiple RAID modes, offer key advantages in the simplicity and maturity of application and design, and meet the majority of RAID requirements. But despite these advantages, NetEase found that solutions based on RAID controller cards also face a number of challenges.

- **Cost:** As they are separate server components, RAID controller cards increase hardware costs. They also place certain specification

requirements on the hard disk backplane of servers, further driving up server costs.

- **Reliability:** The deployment of RAID controller cards requires corresponding hardware, along with components such as cables, firmware, and drivers, all of which introduce new points of failure and reliability risks.
- **Performance Loss:** Unlike NVMe SSDs directly connected to the CPU, RAID controller cards are limited by the upstream PCIe bandwidth (x8 or x16 PCIe lanes). NVMe RAID controller cards also experience performance loss in the host software stack as additional I/O protocol translation from NVMe to SCSI is required when handling read/write operations on NVMe SSDs.
- **Supply:** During special conditions such as chip shortages, RAID controller card chips will be difficult to procure in a timely manner, affecting the stability of server supply chains.

Solution: NetEase RAID Solution Based on Intel® VROC

In order to meet the challenges of traditional RAID solutions, NetEase explored the use of Intel® VROC in SATA RAID and NVMe RAID storage systems to reduce costs and improve efficiency.

Designed for NVMe SSDs, Intel® VROC is an enterprise-grade integrated RAID solution that delivers outstanding reliability and unlocks the full performance potential of NVMe SSDs. Intel® VROC is enabled by a hardware feature in Intel Xeon Scalable processors called Intel® Volume Management Device (Intel® VMD), an integrated controller inside the CPU PCIe* root complex. As NVMe SSDs are directly connected to the CPU, Intel® VROC fully unlocks performance potential for lower latency and higher bandwidth. Enabled by Intel VMD, Intel® VROC eliminates the costs and power consumption requirements of deploying a RAID host bus adapter (HBA) between the SSD and CPU.

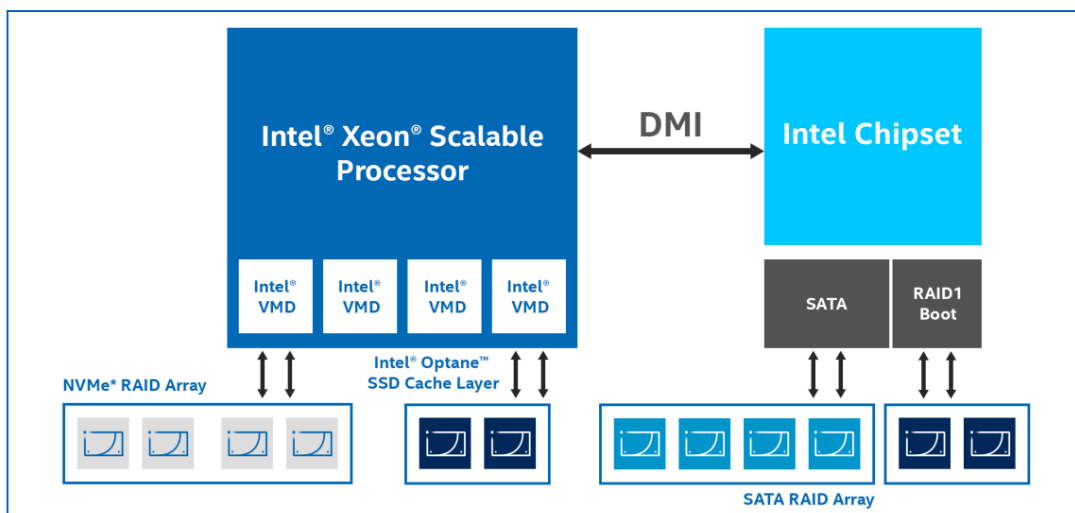


Figure 2. Intel® VROC saves costs and power by eliminating the need to deploy a RAID HBA

Along with high-performance NVMe SSD RAID, Intel® VROC also supports traditional SATA RAID storage. The technology provides data protection for both application data and the system boot disk to deliver performant data storage and data redundancy.

As shown in Fig. 2, Intel® VROC NVMe RAID supports the direct connection of NVMe SSDs to the VMD controller on the CPU, along with the flexible configuration of RAID arrays across VMD controllers to meet the storage capacity and performance demands of different business scenarios. Intel® VROC NVMe RAID also supports the direct connection of NVMe SSDs to the PCH, targeting for NVMe RAID1 boot disks. Intel® VROC SATA RAID supports the direct connection of SATA disks to SATA controllers on the PCH, and provides RAID support for both boot and data disks.

Compared with traditional RAID solutions, Intel® VROC NVMe RAID takes full advantage of the performance of NVMe storage while reducing hardware costs and eliminating points of failure. Tests performed by NetEase show that compared with the RAID controller card solution, the Intel® VROC NVMe RAID solution achieved a leading edge in most areas. Sequential read/write and random read/write test data are shown in Fig. 3 and Fig. 4 respectively.

To further optimize performance on the basis of Intel® VROC NVMe RAID, NetEase also optimized PostgreSQL by optimizing the database architecture and adding a cache layer between master and slave databases to store logs and reduce the communication delay between the master and each slave.

NetEase has also adopted Intel® VROC SATA RAID technology in its boot disk RAID solution. Enabled by the storage controller in the on-board Platform Controller Hub (PCH), Intel® VROC SATA RAID enables RAID for SATA drives, supports RAID 0/1/10/5, and does not require RAID controller cards.

NetEase has carried out a large amount of work on BIOS/UEFI compatibility, OS compatibility, disaster recovery drills, internal tool adaptation, and automated O&M. These efforts have ensured the stable implementation of Intel® VROC SATA RAID in the production environment, reduced hardware procurement costs, and reduced points of failure related to RAID components.

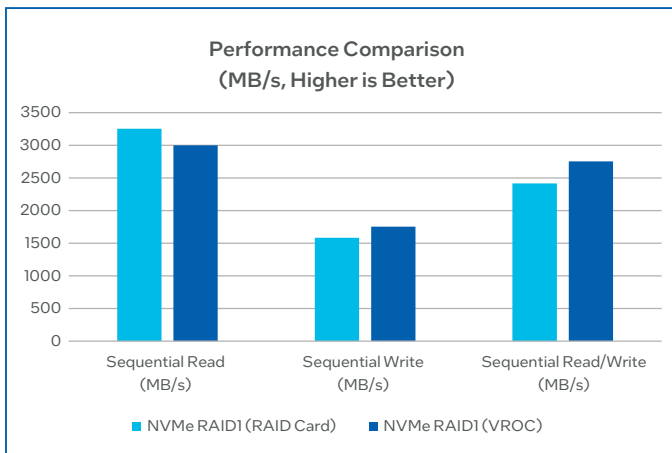


Figure 3. Sequential read/write performance comparison of the RAID controller card and Intel® VROC NVMe RAID solutions¹

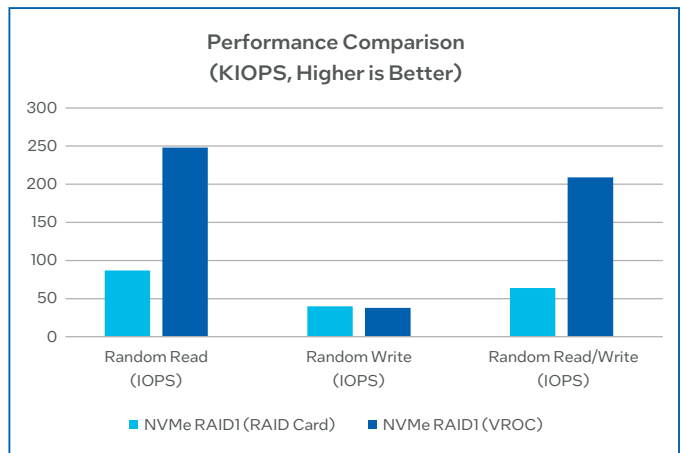


Figure 4. Random read/write performance comparison of the RAID controller card and Intel® VROC NVMe RAID solutions²

^{1,2} Data from internal test results of NetEase Games (conducted on November 2022). 2 x Intel® Xeon® Silver 4214R Processors U @ 2.40GHz, 12 cores, 384 GB RAM (12 x 32 GB DDR4 2933MT/s RDIMMS); 2 x 1.6TB PCIe 3.0 U.2 SSD, Debian10.3, Kernel: 4.19.0-8-amd64, mdadm - v4.1 - 2018-10-01, Intel VROC Pre-OS Version: 6.0.2.1002, Intel-SSD-only Intel® VROC HW Key, fio-3.15, iodepth=128, numjobs=1, bs=4k (random)/128k (sequential), rwmixread=50, fs=ext4. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Benefits

Through optimization with the Intel® VROC solution, NetEase has improved the performance, cost efficiency, and reliability of RAID solutions, and obtained a variety of benefits.

Enhanced Performance

When combined with optimizations to the PostgreSQL master-slave architecture and hardware storage solution, Intel® VROC has improved the QPS of NetEase's PostgreSQL database by around threefold³.

Improved Stability and Reliability

The Intel® VROC solution eliminates points of failure related to RAID controller cards, such as the card itself, cables, firmware, and drivers. In addition to improving stability and reliability, this also facilitates O&M.

Reduced Total Cost of Ownership (TCO)

The Intel® VROC solution eliminates the cost of RAID controller card hardware, cables, and hard disk backplane selection (accounting for 10% of total server costs), and reduces power consumption by around 15W/unit⁴.

Outlook

When compared with traditional solutions, Intel® VROC provides the ability to use NVMe drives to their full potential, fewer hardware dependency, bootable RAID, and hot insert/surprise removal. Affordable, efficient, and simple, Intel® VROC is ideal for building RAID solutions in modern cloud and data centers. Adoption by NetEase has shown that RAID solutions based on Intel® VROC provide highly reliable and cost-effective local storage for database applications, and realize reduced costs and improved efficiency in cloud computing and database scenarios.

As it continues to drive innovation, Intel will work closely with NetEase and other partners to accelerate the implementation and validation of Intel® VROC in actual application scenarios. Intel will also continue to integrate next-gen Intel Xeon Scalable processors and other hardware to help users improve storage performance, accelerate data processing in AI, analysis, cloud and microservices, network, database, storage, and other workloads, and facilitate digital transformation.

Acknowledgement

NetEase Games:
Chen Bo, Jin Rong

^{3,4} Data from internal test results of NetEase Games (conducted on November 2022). Test Configuration: 2 x Intel® Xeon® Silver 4214R Processors U @ 2.40GHz, 12 cores, 384 GB RAM (12 x 32 GB DDR4 2933MT/s RDIMMS); 2 x 1.6TB PCIe 3.0 U.2 SSD, Debian10.3, Kernel: 4.19.0-8-amd64, mdadm - v4.1 - 2018-10-01, Intel VROC Pre-OS Version: 6.0.2.1002, Intel-SSD-only Intel® VROC HW Key, fio-3.15, iodepth=128, numjobs=1, bs=4k (random)/128k (sequential), rwmixread=50, fs=ext4. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

About NetEase Games

Established in 2001, NetEase Games has continued to grow alongside the gaming community. After 20 years of rapid development, the company has become one of the largest in the gaming industry. As a leading game developer in China, NetEase Games has always been at the forefront of independent industry research and development. The company positions itself not only as a gaming platform and service provider, but also as an integral part of the gaming community.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.



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.