

Ceph-Based Software Defined Storage Solution Optimization Practice In Financial Industry

The head of computing infrastructure at the bank indicated that: "By analyzing the bottleneck of the Ceph storage system together with the Intel team, we discovered an effective way to improve performance. Judging by the current test results, the performance improvements brought about by Intel® SSD cache has been instant. Therefore, we are now more confident in moving more of our essential services to the Ceph storage system. This will help provide a flexible and robust IT service platform for our upper layer applications and service innovations."

Against the general trend of digital transformation, the finance industry is currently restructuring its IT infrastructure to aid innovation in upper layer applications such as online financial services. In order to further optimize the storage system after choosing OpenStack* to build a private cloud and using Ceph* as the software-defined platform software, a department in a certain large bank created a high performance Ceph storage system by deploying an Inspur NF5280M5 server which has an Intel® Xeon® Gold 5120 processor, Intel® SSD DC P4600, and Intel® SSD D3-s4510, which replaces mechanical hard drives with high performance SSDs to overcome the performance bottlenecks of mechanical hard drives.

Industry background: Software-defined Storage Becomes a Major Trend in the Finance Industry

Out of the traditional industries, the finance industry has always been one of the industries with the highest degree of informatization, highest reliance on IT systems, as well as the fastest rate of technological development. As cloud computing technologies have become more widely used in recent years, and as the digital transformation of the finance industry takes form, the major developmental focus of IT in the finance industry is on the restructuring of IT infrastructures and building distributed system architectures to provide a flexible and robust IT service platform.

Whether it is mobile banking, online loans, other innovative internet financial services, or artificial intelligence, IoT and other innovative integrations of technologies with the finance industry, all have presented additional pressure for storage systems in the finance industry to evolve. The rapid growth in financial data, along with the challenges in performance, operation and maintenance, costs, and so on, regarding scenario-specific, mobile, and fragmented non-structured financial data, has made it difficult for traditional Storage Area Network (SAN) systems to handle. Therefore, many financial institutions have gradually moved towards software-defined storage systems. Compared to SAN storage, software-defined storage has a wide range of advantages in implementing storage service platforms with shared resource pools, high utilization, high availability, low cost, and low power consumption.

Ceph is an open source software-defined horizontally scalable storage system that can be used at large scales. It can provide object-, block-, and file-level storage for financial institutions on a single platform. Ceph storage can be run on standard large capacity servers which helps reduce costs while also providing outstanding flexibility and scaling on demand. It has huge advantages when storing a large amount of non-structured data when compared to SAN storage. For companies in the finance industry that are moving their services to the cloud, Ceph has become a major component in storage system restructuring.

Challenge: Building a High Performance Ceph Storage System

This large bank has always been at the forefront of IT systems in China's finance industry, and began to move from the traditional centralized architecture to a distributed cloud computing platform a few years ago. It has chosen to use OpenStack to build its private cloud. The bank has also implemented a distributed architecture in a highly effective way and used Ceph as the software-defined storage platform software to fully unleash its advantages in areas such as costs, agility, and so on, to improve product innovation and service capabilities.

Ceph was mainly used for cold and warm data storage in the initial design phase and was not advantageous in terms of I/O performance when compared to SAN storage. This bank then used Ceph as the data storage system for non-essential services while essential services were still handled by SAN storage. However, as the requirements of the PaaS platform and the online database became more demanding, the disadvantages of SAN in flexibility and deployment time became a technical bottleneck that hindered the development of the cloud platform. Therefore, the bank hoped to move more essential services onto Ceph storage. But this led to an unavoidable problem: how to improve the I/O performance of Ceph so that it could handle essential services?

One of the methods to improve the performance of Ceph storage is to deploy storage clusters. Nodes can be added to implement agile scaling of I/O performance, but in consideration of costs, power consumption, space usage, and other factors, the deployment of a large number of nodes is not suitable. Therefore, this bank hoped to

control the Ceph cluster while optimizing the hardware and software of single storage nodes in order to greatly increase the I/O capabilities of the Ceph cluster and satisfy the storage performance demands of essential applications.

Solution: Intel and Inspur Work Together to Provide a Ceph Optimization Solution

In order to help the bank resolve the bottlenecks in the Ceph system, Intel started a technical partnership with the bank and both put in technical resources to analyze and study the bank's Ceph system. By collecting system information, Intel engineers discovered that while in operation, there were no bottlenecks in the processor, memory, or network. However, the system's iostat command shows that although the average disk utilization was not high, every sample point had hard drives running at 100% utilization. Further investigation revealed that each sample point had at least 6 hard drives running at 90% utilization or above. This clearly shows where the bottlenecks in Ceph systems are.

Engineers from both parties tried using software optimization methods to overcome this bottleneck at first. However, they found that due to the special conditions of the bank's internal IT environment, many of the optimization methods could not be applied. Further analysis revealed that hard drives in this bank had a large amount of useless read operation. These read operations were caused by lost metadata in FileStore. To address this problem, Intel recommended using the Intel® SSD DCP4600 cache to cache the metadata. The advantage of this is the complete separation of data and metadata, where only data operations and no metadata operations are performed on mechanical hard drives. Since the Intel® SSD DCP4600 cache performs much



Inspur NF5280M5 Server

better than mechanical hard drives, the overall performance will not be affected even if metadata is lost.

The bank approved this upgrade solution and began replacing the original mechanical hard drives with Intel® SSDs. In terms of server node selection, Inspur recommended using the Intel® Xeon® Gold 5120 processor-based Inspur NF5280M5 server with Intel® SSD DCP4600 and Intel® SSD D3-s4510. This server uses all-new smart computing acceleration technology and is able to handle the huge amount of data in the finance cloud in an effective and flexible way.

The Inspur NF5280M5 server is a 2U dual-channel server which supports a maximum of 31 2.5" hard drives and has overall performance which is 1.3 times better than previous generation systems. It has 8 standard PCIe slots and 1 OCP slots for flexible scalability. Furthermore, the NF5280M5 uses different layers of high availability technology. For example, it has a wide range of RAID configurations, hot backup memory, imaging, error correction, as well as redundant hot swap design for power supply fans and has complete hardware and software error tolerance. The high availability ensures the stability of financial services.

As an important component of this upgraded solution, the Intel® SSD DCP4600 is an SSD based on the NVMe* (Non-Volatile Memory Express) specification and provides the industry's leading throughput rates and latency performances with high service quality and ultra-high durability. It is optimized for read-intensive workloads and can effectively lower costs, enhance server effectiveness, and improve data reliability.

- Intel® SSD DCP4600 uses an all-new NVMe controller which is optimized for mixed workloads commonly seen in data caching and significantly improves server flexibility and utilization.
- Thanks to the controller's support of a maximum of 128 arrays, the DCP4600 assists in minimizing the idle CPU time and achieves the maximum efficiency on Intel® platforms.
- With the help of DCP4600, cache in data centers can be accelerated, and every server can provide more services and handle heavier workloads for more users.
- The new firmware management feature and updating process of the DCP4600 and its scaling and monitoring feature helps in reducing server downtime.
- Intel has built an industry leading end-to-end data protection function in the DCP4600 that can reduce the risk of adverse effects from data loss, and reduce catastrophic downtime and major errors.

Furthermore, the solution also uses Intel® SSD D3-s4510 as the operating system disk to improve operating system performance.

In terms of processors, the Intel® Xeon® Gold 5120 processor is used to provide a solid foundation for the Ceph storage solution which enables a leap in performance improvement in terms of agility and scalability. This type of innovative processor utilizes revolutionary design and establishes new standards in the areas of platform convergence, computing, storage, memory, networking, service continuity, and so on.



Intel® Xeon® Scalable processors

Intel® Xeon® Scalable processors:

- Offer high scalability to enterprise data centers
- Provide higher performance in virtual infrastructures when compared to previous generation processors
- Implement outstanding resource utilization rates and agility
- Improve data and work load completeness and ensure regulatory compliance in data center solutions

This bank uses Intel® Xeon® Gold processors in its Ceph-based storage solution.



Result: Provides High Speed I/O Capacity Support for Essential Finance Applications

Thanks to the performance advantages of Intel® SSDs and software, this bank significantly improved its I/O capabilities in its Ceph clusters, and achieved much better throughput performance for large and small data blocks when compared with solutions with mechanical hard drives. Since the Ceph cluster performance bottleneck was overcome, this bank could move more essential services to the Ceph storage system to fully unleash the advantages of Ceph storage in terms of costs, flexibility, and so on.

Prospect: Building a Ceph Storage Platform for the Future

The experience with working with the bank has shown that a Ceph-based open source storage system, along with Intel® Xeon® processors, Intel® SSDs and other such products, can be used to realize an improvement in performance. In particular, after using Intel® SSD cache, the overall performance of the Ceph cluster increased significantly. It also verifies that Ceph can become the storage back-end for essential finance industry applications.

As China's leading financial service provider, the bank used Ceph to build this breakthrough storage solution and verified that it could be used to replace traditional SAN storage in a wider range of applications. The bank plans to use Ceph storage in more of its essential financial services, and continue to work with Intel to create high performance storage solutions to provide more efficient and reliable services to internal and external users.



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. For more complete information about performance and benchmark results, visit <http://www.intel.cn/content/www/cn/zh/benchmarks/benchmark.html>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark® and MobileMark®, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information please visit <http://www.intel.cn/content/www/cn/zh/benchmarks/benchmark.html>

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

Intel has not made any explicit or implicit guarantees, including not limited to matters related to marketability, suitability for intended purpose, and infringement, as well as any guarantees related to compliance processes, transaction processes, or trade practices. Intel, and the Intel logo, are registered trademarks of the Intel Corporation in the United States and/or other countries.

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

Intel Corporation © 2019 all rights reserved. All rights reserved.