

Intel Technology Accelerates H3C ONEStor Distributed Storage System

H3C

“New technologies and corresponding new applications such as artificial intelligence, 5G networks, autonomous driving, and the Internet of Things come with an exponential explosion in data volume and rapidly growing demands for capacity, density, durability and reliability in data storage. Intel teams up with H3C in the storage field and applies new high-performance storage technologies such as non-volatile persistent memory and high-performance NVMe flash memory, in combination with optimized the end-to-end of distributed storage software solutions, which is a huge leap in high-performance distributed storage products.”

– Mr. Li Yadong

General Manager of Government and Enterprise Division, China, Data Center Group, Intel Corporation

“H3C, as a world-leading provider of intelligent connectivity, information security, big data, cloud computing, and distributed storage solutions, is also a strategic partner of Intel. In our close cooperation with Intel in the H3C all-flash distributed storage optimization project, we have utilized products like Intel® Optane™ persistent memory and SPDK ranging from software to hardware based on the 2nd Generation Intel® Xeon® Scalable processor, focusing on software optimization, software and hardware collaboration, and thus greatly enhance the performance of the distributed storage solution with the provision of higher performance and cost-effective all-flash storage solutions to H3C users as the core. In the future, we will continue our cooperation for providing users with the latest products that keep pace with the times.”

– Mr. Wang Youkang

Vice Head of H3C Distributed Storage Product Department

Fueled by artificial intelligence (AI), live streaming video and Internet transaction applications, the data to be stored or processed in real time by enterprises is increasing exponentially. The need for high-performance storage systems becomes critical for real-time read-write, random access to super large-scale data sets, and other use scenarios. As unit capacity cost of solid-state drive (SSD) continues to decline, the demand for high IO throughput and low latency services becomes more and more extensive. This offers development avenues for high-performance all-flash storage which is fully dominated by solid-state storage media. At present, all-flash storage has gradually become a preferable choice for enterprises while deploying or optimizing storage in the digital transformation process.

Given this background, New H3C Technologies Co., Ltd. (“H3C” for short) has enabled H3C ONEStor distributed storage system to be equipped with all-flash arrays. The all-flash array meets higher real-time and high-concurrency requirements in commercial use scenarios while providing higher performance and lower latency than traditional distributed storage environments. Furthermore, the H3C ONEStor distributed storage system also features high scalability, high flexibility, intelligence, and high availability. Its excellent performance density and capacity density are needed by most enterprise users in business-critical environments.

Challenge: How to enhance all-flash storage performance?

Thanks to its advantages such as high performance, reliability, durability and manageability, all-flash storage has become an important direction enterprise-grade storage step towards. The *Worldwide Quarterly Enterprise Storage Systems Tracker* released by IDC shows that in the third quarter of 2019, the global all-flash array market registered revenue of \$ 2.58 billion, up 11.7% from the previous year, significantly higher than the overall growth rate of the enterprise storage market in that quarter¹. On the other side, an increasing number of enterprises are looking to deploy scalable distributed storage based on all-flash arrays at a lower cost, while meeting the exacting requirements of critical business for real-time data access.

With the evolution of flash memory performance and the innovation of software optimization technology, the performance of all-flash storage has perpetuated robust growth in recent years. However, all-flash storage still has a long way to go. All-flash storage performance is being powered by higher-performance SSDs, innovative software optimization technology and exclusive performance acceleration chips. In order to improve the performance of the H3C ONEStor distributed storage system, H3C has made technology-exploration towards metadata access acceleration, lower processor (CPU) overhead and optimized distributed storage system.

First, as far as cloud data access is concerned, due to high performance and advanced features, read and write access to metadata in all-flash storage system is often of high frequency. In distributed storage

system, metadata and data caches are usually deployed on SSD which may somehow retard metadata reads and writes due to bottleneck in performance and then the overall performance of all-flash storage. Based on this, H3C are looking for metadata media storage that has higher performance than flash memory and can be persisted.

Second, along with enhanced flash media performance, the main bottlenecks of all-flash storage in performance have shifted from flash memory to CPU. In data center application scenarios, CPU performance overhead can reach more than 90% in case that multiple SSDs are read and written simultaneously. If the bottlenecks of the CPU in performance cannot be broken down, the result is lower scalability of all-flash storage and reduced number of SSDs, which will greatly limit the performance and capacity of all-flash storage.

Third, the H3C ONEStor distributed storage system is developed and optimized based on open source Ceph. However, no optimization has been made by Ceph's earlier version on all-flash storage. The said shortfalls should be identified, in parallel with corresponding code transplant for redevelopment, in a bid to further enhance the performance H3C ONEStor.

Solution: H3C ONEStor distributed storage system based on latest Intel technology

The H3C ONEStor distributed storage system builds on a distributed architecture, which underpins the reliability, availability, autonomous operation and high performance. The system boasts advanced converged distributed architecture, linear scalability, flexible

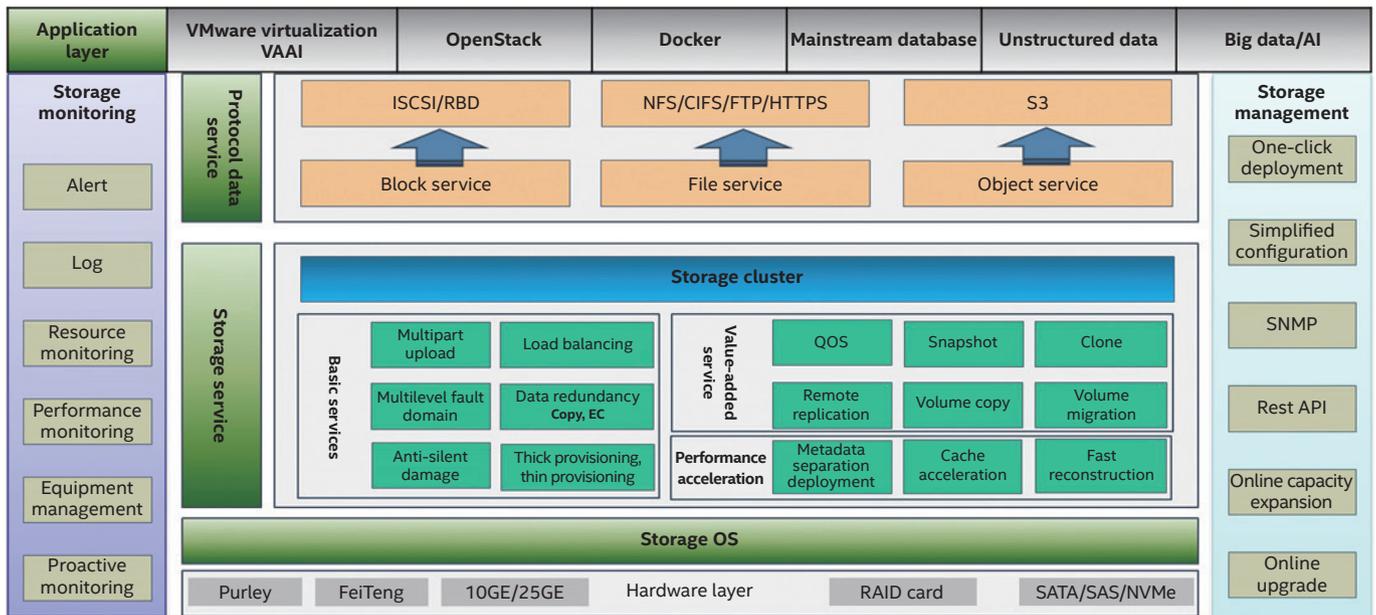


Figure 1: H3C ONEStor Distributed Storage System Architecture

2nd Generation Intel® Xeon® Scalable processor

2nd Generation Intel Xeon Scalable processor:

- Provides high scalability for enterprise data centers
- Provides higher performance for virtual infrastructure, compared to the prior generation
- Enable greater resource efficiency and agility
- Improve data and workload integrity and ensure regulatory compliance for data center solutions

H3C ONEStor distributed storage system based on latest Intel technology recommends using the 2nd Generation Intel Xeon Scalable processor.



partitioning storage, rich enterprise-grade features and autonomous operations, and can be applied to block storage, file storage and object storage among others to meet diverse storage needs of users in different industries.

In order to improve the performance of the H3C ONEStor distributed storage system, H3C partnered with Intel to incorporate Intel® Optane™ persistent memory, 2nd Generation Intel® Xeon® Scalable processor and Intel® SPDK for optimization of all-flash distributed storage. For storage server, 2nd Generation Intel Xeon Scalable processor is recommended to suit higher requirements for performance, agility, and scalability.

Accelerated metadata access

In an effort to enhance metadata read/write performance without compromising metadata non-volatility, H3C ONEStor distributed storage system leverages Intel Optane persistent memory as metadata storage media. Intel Optane persistent memory is a groundbreaking memory technology that delivers a unique combination of affordable large capacity and support for data persistence. It provides read and write bandwidth far ahead of NVMe devices and significantly decreases latency for storage. Intel Optane persistent memory is categorized into Memory Mode and App Direct Mode. In App Direct Mode, Intel Optane persistent memory offers the ability to accelerate data access and persistent data storage.

Intel® Optane™ persistent memory

Benefits of Intel Optane persistent memory:

- Lower total cost of ownership (TCO): It brings cost saving to set up a memory pool with parallel performance or capacity and helps enterprises to lower infrastructure costs and gain a greater competitive edge.
- Agile application modes: Enterprises can use the Memory Mode or App Direct Mode of Intel Optane persistent memory as required without the need to adjust hardware.
- Robust performance, far ahead of traditional SSDs.

H3C ONEStor distributed storage system based on latest Intel technology is recommended to adopt Intel Optane persistent memory.



In the optimization process, H3C capitalizes on Intel Optane persistent memory in App Direct Mode to power the metadata for H3C ONEStor distributed storage system, and access Intel Optane persistent memory by means of PMDK to expedite reading and writing of metadata.

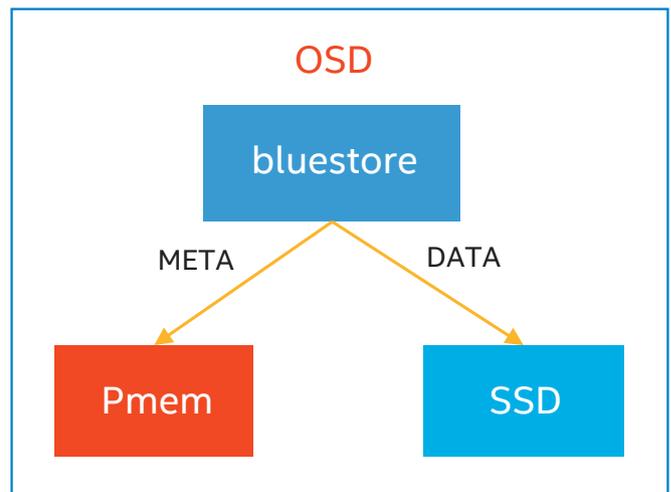


Figure 2: Accelerated Metadata Access for H3C ONEStor Distributed Storage System

SPDK driver lowers system overhead

The H3C ONESstor distributed storage system is powered by high-performance Intel Xeon Scalable processor which helps provide high-throughput, IOPS and low-latency storage performance for critical workloads. Additionally, H3C is also optimized with Intel SPDK which offers a portfolio of development kit based on NVMe drives. The goal of SPDK is to highlight the outstanding performance of flash storage media enabled by using Intel's networking, computing and storage technologies together.

Compared to traditional I/O model, Intel SPDK combines two key techniques: UIO and pooling. First, Intel SPDK can run the device driver code in User Mode to avoid the kernel context switches and interrupts. This saves a significant amount of processing overhead and allows more clock cycles to be spent doing the actual storing of the data. Second, Intel SPDK adopts the polling mode and changes the basic model for an I/O, avoiding the latency and overhead by interrupts, in parallel with allowing the application to improve I/O efficiency. In H3C ONESstor distributed storage system, H3C changes kernel driver into SPDK NVMe user mode driver so as to enhance performance.

Open source Ceph optimization

Due to a close cooperation with Intel at the code level, H3C has introduced the optimized implementation of flash storage media for BlueStore and RBD from the open source community, together with optimized part of the native Ceph configuration parameters, enabling optimal performance in all-flash scenarios.

Benefit: H3C ONESstor distributed storage system based on latest Intel technology for accelerated data transformation

Owing to above-mentioned optimizations, the H3C ONESstor distributed storage system has seen huge improvements in performance with 36% increase in SPDK driver, 34.6% in corresponding module integration from the latest open source community and 30% in total.

All-flash arrays optimized based on Intel Optane persistent memory and Intel SPDK secures a victory for H3C ONESstor distributed storage system to deliver higher performance and suit the requirements of government, education, enterprises, power energy, finance and operators for real-time data processing. Additionally, the solution also enjoys advantages in:

- **Better cost-effectiveness:** With optimized all-flash storage for Intel Optane persistent memory and Intel SPDK, H3C ONESstor reduces users' demands for the number of nodes and brings cost savings on procurement, energy consumption, space and operations, as well as total cost of ownership.
- **Stronger scalability and agility:** Intel Optane persistent memory offers the attributes of both storage and memory. It has a varied set of application scenarios and helps achieve performance-cost balance.

To sum up, the rollout of H3C ONESstor distributed storage system based on latest Intel technology will help enterprise and industry users to address the challenges data explosion brings and accelerate digital transformation.

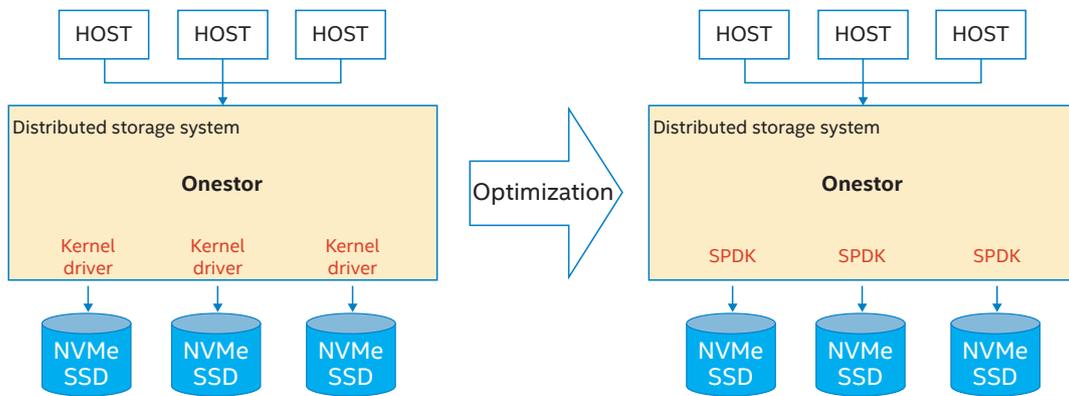


Figure 3: Optimized CPU Overhead for H3C ONESstor Distributed Storage System

Outlook: H3C will join hands with Intel to promote all-flash storage system innovation

As the result of cooperation with Intel, H3C has optimized all-flash storage performance utilizing Intel products and technologies and hence provides a solid underpinning for commercial deployment of H3C ONEStor distributed storage systems in key industries. The parties will continue their cooperation in terms of network technologies (RDMA), the latest Ceph version, Intel Optane

persistent memory (AD mode), etc. to further promote all-flash storage system innovation, advance the rollout of all-flash storage in more storage systems and application scenarios and maximize all-flash-enabled performance improvement, so as to bolster more data-critical businesses and spearhead digital transformation across sectors.

Appendix 1: List of recommended configurations for H3C ONEStor distributed storage solution based on latest Intel technology

Server: UIS 3000G3 (2U)	
Processor	2 x Intel® Xeon® Gold 5218 Processor @ 2.30GHz
RAM	8 x 32 GB DDR4-2666
Storage network	Management network: 1xGE Business network: 2x25GE Storage network: 2x25GE
HDD	4 x 1.8TB NVMe SSD
RAID	P460
Guest: H3C R4900 G2	
Processor	2 x Intel® Xeon® Processor E5-2690 @ 2.60GHz
Memory	16 * 32 GB 2400Ghz
Network	Management network: 1xGE Business network: 2x25GE (Bond4)

About H3C

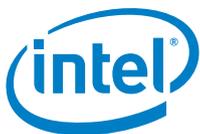
H3C is an industry leader in the provision of Digital Solutions, and is committed to becoming the most trusted partner of our customers in their quest for business innovation and digital transformation. H3C offers a full portfolio of Digital Infrastructure products, spanning across compute, storage, networking, security and related domains, and provide a comprehensive one-stop digital platform that includes cloud computing, big data, interconnectivity, information security, new safety, Internet of Things (IoT), edge computing, artificial intelligence (AI) and 5G solutions, as well as end-to-end technical services. H3C is also the exclusive provider of HPE® servers, storage and associated technical services in China.

About Intel

Intel Makes Possible the Most Amazing Experiences of the Future

You may know us for our processors. But we do so much more. Intel invents at the boundaries of technology to make amazing experiences possible for business and society, and for every person on Earth.

Harnessing the capability of the cloud, the ubiquity of the Internet of Things, the latest advances in memory and programmable solutions, and the promise of always-on 5G connectivity, Intel is disrupting industries and solving global challenges. Leading on policy, diversity, inclusion, education and sustainability, we create value for our stockholders, customers, and society.



¹ [Asia/Pacific grew the fastest of any region in the worldwide enterprise external storage systems market during the third quarter of 2019], <http://www.elecfans.com/consume/1140031.html>

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 www.intel.com/benchmarks

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 complete information visit intel.com/benchmarks

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. 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, 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.

Copyright © 2020 Intel Corporation. All rights reserved.