

THREE KEYS FOR A BRIGHT ALL-FLASH ARRAY FUTURE

AFA READY PLAYER ONE

Gartner:
Moving Toward the
All-Flash Storage Data Centers

Moving Your Mission-Critical
Applications to Cloud Just Got
Easier!

3 Things You Should Consider
When Buying Flash Storage

INSIGHT INTO THE FUTURE
OF FLASH TECHNOLOGIES

Moore's Law and Huawei Storage's
Continuous Improvement

4 Questions to Understand
All-Flash Performance

CRYSTAL KEY: CUSTOMERS'
VIEWS & SUCCESS STORIES

Can You Imagine That the
Nasdaq Market Shutdown Was
Caused
by a False Fire Alarm?

Vostochny Bank Quickly Expands Businesses
with High-Performance Storage Devices

The All-Flash "Renaissance" of Telecom Italia

BYD Auto Becomes a Leader in New Energy
with All-Flash Storage

#ALL-FLASH ARRAY
BUILT FOR YOUR
APPLICATION SCENARIOS



Contents

04 / Gartner: Moving Toward the All-Flash Storage Data Centers

08 / Moore's Law and Huawei Storage's Continuous Improvement

12 / Huawei's All-Flash Storage Roadmap

Trend: Insights Into The Industry

Copper Key: Devices & Black Technologies

State-of-the-art Devices

15 / DCIG All-Flash Array Buyer's Guide: An Authoritative Report

17 / Three Things You Should Consider When Buying Flash Storage

22 / Four Questions to Understand All-Flash Performance

27 / Overview of Huawei OceanStor All-Flash Storage Products

28 / Huawei's All-Flash Arrays Offer Record-Breaking SAN and NAS Storage Performance

Fabulous Skills

34 / Huawei All-Flash Reliability Assurances for Your Mission-Critical Businesses

40 / Huawei OceanStor Dorado V3 All-Flash Storage 3 to 1 Data Reduction Guarantee

41 / Huawei Accomplishes 3:1 Data Reduction Rate in All-Flash Offerings

44 / Chips Purpose-Built for All-Flash Storage: the Core of Competitiveness

50 / OceanStor OS Enables Fast and Stable All-Flash Storage

58 / FlashLink: The Secret to Huawei's High-Performance All-Flash Storage

Native AFAs

64 / Lightning-Fast All-Flash Storage Ushers in an NVMe Era

70 / Huawei All-Flash Storage OceanStor Dorado V3 Wins Best of Show Award at Interop

Beyond Imagination

72 / ESG Test Report: Huawei OceanStor Dorado V3 Delivers Excellent Performance for Mainstream All-Flash Applications

76 / Huawei OceanStor Dorado V3 All-Flash Solutions, Built for Your Application Scenarios

87 / Can You Imagine That the Nasdaq Market Shutdown Was Caused by a False Fire Alarm?

Across Boundaries

91 / Moving Your Mission-Critical Applications to Cloud Just Got Easier!

93 / Huawei's Business Integration Solution: A Revolutionary Optimal Data Migration Method

Jade Key: Scenarios & Solutions

Crystal Key: Customers' Views & Success Stories

Industry Feedback

99 / Voice of Industry User

105 / Voice of Industry User

Success Cases

100 / Vostochny Bank Quickly Expands Businesses with High-Performance Storage Devices

103 / China Life: Transformation to Technical Guidance for Business Innovation

106 / BYD Auto Becomes a Leader in New Energy with All-Flash Storage

109 / One-Picture-Say-It-All: Yahoo Japan

110 / One-Picture-Say-It-All: Telecom Italia

111 / The All-Flash "Renaissance" of Telecom Italia

114 / Insight into the Future of
Flash Technologies

Plan: To The AFA Future

Get Ready For The Fantastic AFA Journey



Email: OceanStor@huawei.com

Gartner: Moving Toward the **All-Flash Storage** Data Centers


By Gartner



Market Overview and Forecast

According to the data analysis gathered by Gartner, for the storage industry in 2017, the overall market volume of all-flash storage increased by 36%, reaching US\$6.3 billion, whereas the procurement price continued to decline by at least 40% in the past 18 months.


In response to increasing competition among different brands, vendors launched software packaging, performance, and reliability commitments, further emphasizing the value of all-flash storage in data centers. It is predicted that 50% of all traditional disk storage in data centers will be replaced by all-flash storage by 2020.



Source: Gartner

The application of NVMe and 3D XPoint will help improve the all-flash storage performance by ten folds with a doubled density and 50% lower unit capacity price in the next year, reducing the costs of all-flash storage and making it more accessible than ever before.

Average price (US\$)



Source: Gartner

Analysis


Most storage vendors capitalize on enhanced product functions to compete with each other. Driven by the advancements in different technologies, the storage infrastructure will be fundamentally different in the future. Solid-state Disks (SSDs) and all-flash storage are transformational technologies that increase performance and storage utilization by up to two orders of magnitude.

SSDs and all-flash storage delivers high performance, low cost, and high reliability, which are regarded as unobtainable in traditional storage scenarios. Many commonly accepted rules will be no longer applicable. For example, the previous unwritten requirements such as matching storage performance, application performance, and service importance to maximize resource utilization are no longer major issues because all-flash storage data centers provide a lower TCO than traditional disk storage. All-flash storage helps customers save on O&M costs while improving administration productivity and infrastructure efficiency thanks to its data reduction mechanism. This results in a flat storage infrastructure with only one storage tier, consisting of SSDs.

Competition and progress in all-flash storage has created a level of maturity, reliability, and cost-effectiveness that exceeds hybrid disk arrays in all characteristics except raw capacity, providing data centers with higher agility and service capability.

Reliability Comparison

The advanced disk technologies have prolonged the life cycle of traditional disk storage from 3 years to 5 years, but it is difficult to improve on that length. Alternatively, the reliability of all-flash storage extends the storage life cycle to 7 years or even longer, which is why many enterprises are leading towards all-flash for their storage needs.



Source: Gartner

TCO Comparison

Take an American medical analysis organization as an example. It replaced its hybrid storage with all-flash storage at 3.6:1 data reduction ratio. The all-flash storage helped decrease the procurement cost and operational cost by 75% and 83.5% respectively.

	Old Environment	New Environment
Capacity	155TB	80TB ¹
HDD	75%	0%
SSD	25%	100%
Capex — Total Cost	75% higher than new environment	
Opex — Administration FTE	Two FTEs, including zoning	0.5 FTEs, including zoning
Opex — Power	4 times more power draw About 12,000 watts	Approximately 3,200 watts
Opex — Cooling	3.6 times more Btu required About 40,000 Btu per hour	About 11,000 Btu per hour
Opex — Space	Seven system/storage bays Seven bays by H: 76" x W: 30" x D: 42"	(23U) 1 standard rack: H: 79.00" x W: 44.30" x D: 23.54"
Opex — Maintenance	Annual maintenance cost: \$1 million	Annual maintenance cost: \$165,000
Consolidation — Software License	\$150,000 paid for value-add replication and cloning software Oracle charges by processor count; it had two times the processor count and thus two times the Oracle licenses in the old software.	Replication, cloning and data deduplication additional cost is \$0 Reduction in Oracle software cost driven by core count reduction
Productivity	Oracle RAC storage topology deployment for 38 LUNs takes one to two hours per deployment	Oracle RAC storage topology deployment for 38 LUNs takes five minutes per deployment
¹ Data reduction ratio is 3.6:1. Btu = British thermal unit D = diameter H = height LUN = logical unit number W = width		

Impacts and Recommendations

Impacts

Many vendors are on their second, third, or fourth generation of product development. They successfully provide reliability, availability and service ability features that enable them to be used in the most demanding environments. Dynamic software upgrades, hardware hot swapping, remote support, and proven service functions are equivalent to or better than those of existing disk storage or hybrid storage arrays.

Recommendations

Select all-flash storage for primary applications, but do not limit the application for other workloads.
Select the most mature all-flash storage for mission-critical and core applications.

Impacts

The value proposition of higher-capacity SSDs in denser storage systems will enable I&O leaders to better support the digital business by enabling new types of applications.

Many different types of SSDs are used in all-flash storage, servers, integration systems, and JBOFs. Large-capacity and low-cost flash storage based on 3D NAND is being used to create 4 TB to 15 TB SSDs. The increase in SSDs capacity and relative reduction in price per TB will make high-density JBOFs and all-flash storage possible. JBOFs Or SSAs that have not been designed with performance as the primary feature to become cost-effective for non-primary data workloads such as active archive and big data.

Recommendations

Implement high-density, high-capacity JBOFs or all-flash storage for active archives, data lakes, big data infrastructures and IT operations and analytics (machine-learning applications) to shorten the search time.

Create new services from previously time-constrained analytics workloads and algorithms.

Use all-flash storage to improve the performance of large data sets, such as archive, auditing trails, and security searches. This type of applications could not be processed as they took too long to run.

Impacts

Implementing SSDs throughout the data center can reduce operational costs in terms of storage array administration, physical floor or rack space, plus power and cooling reductions.

In addition to performance, all-flash storage boasts many advantages over hard disk storage. The productivity of storage administrators will improve as they spend less time configuring and tuning storage arrays. Unlike hard disk storage, all-flash storage provides the data compression function, which helps reduce footprint and power consumption. There are many advantages of all-flash storage that should be mentioned: less electromagnetic interference and lower equipment room noise.

Recommendations

Quantify non-performance-related advantages, such as increased staff productivity and lower operational costs like power, cooling and rack space costs.

To move to a solid-state data center, use all types of SSDs found in servers, JBOF, integrated or converged systems, and solid-state arrays.

Moore's Law and Huawei Storage's Continuous Improvement

By Meng Guangbin, President of Huawei Storage Product Line



Charles Dickens described the French Revolution era as "It was the best of times, it was the worst of times..." a sentiment that also applies to today's world. We are at a historic turning point full of opportunities and challenges in a complex global environment.

The Age of Steam and Age of Electricity are a thing of the past. In today's intelligent information age with rapid advancements of science and technology, it's winner-take-all. That is, companies grabbing advantageous opportunities can survive. In the storage field, an increasing number of real-time interaction applications, especially in cloud computing, online payment, and mobile social networking arenas, are generating a soaring data amount. Storage performance of data centers has become paramount.

Over the past 20 years, flash storage was the largest technological revolution in the industry. Flash storage supports faster data processing than

traditional disk storage and is essential for enterprises to cope with explosive data growth, accelerating mission-critical applications. Gartner predicts that flash storage will become more prevalent than traditional disk storage within two years, becoming the mainstream storage medium of data centers, and applying to mission-critical enterprise workloads.

In this fast-evolving era, Huawei is never a blind follower. We continue to innovate to stay ahead of the technological curve. The past few years witnessed the emergence of a variety of concepts in the storage industry – optical storage, quantum storage, and biological storage, to name a few – but few of them have been commercialized. Such innovations are only abstract ideas bringing no real benefits to customers. This goes against Huawei's philosophy. At Huawei, we plan and design products with the goal of improving customer experience, delivering innovative, better, and faster offerings. To that end, we develop distinctive technological features, engineering deployment, and intelligent O&M, suited to the core requirements of customers.

R&D Innovation: Huawei's Core Strength, Following Moore's Law

Cloud and storage are the most demanding areas in the ICT industry. Users demand that their core business processes and data are always stable and secure. Huawei invests more funds and personnel in storage R&D than any peer company. We began researching storage technologies in 2002 and now have R&D offices around the world. Our Silicon Valley office is the bridgehead of storage technologies, and the storage algorithm research center in Russia and delivery competence centers in Shenzhen, Chengdu, and Beijing all contribute to technology innovation and core competitiveness. Huawei has already invested more than US\$2 billion and 3200 R&D engineers in this field, obtaining more than 800 storage patents.


In addition to the continuous accumulation of hardware strengths, Huawei is also an industry leader in data applications. In the future, the core competitiveness in the storage industry lies in the organic combination of software and hardware capabilities, integrating chips and solid-state drives (SSDs) with algorithms, compute, and software capabilities. Huawei is the only vendor with a proprietary storage operating system (OceanStor OS), controllers, and SSDs (Huawei's SSDs). We implement end-to-end, in-depth, flash-oriented optimization on OceanStor OS and SSDs to maximize the advantages of all-flash storage. The dedicated flash-optimized OceanStor OS, unique flash-optimized algorithms, and dedicated performance-acceleration chips enable Huawei's all-flash storage to deliver the industry's highest performance and reliability.

As the core for Huawei's all-flash storage software and hardware, OceanStor OS and chips are of strategic significance for us.

The core of a storage system is its operating system. An optimal operating system ensures smooth storage experience of end users. Currently, the configuration, service logics, and hardware deployment of Huawei's storage rely on OceanStor OS, the soul of Huawei's storage products, fully developed by Huawei's employees.

Powered by OceanStor OS, Huawei's storage products enjoy continual innovation in storage management. Traditional storage vendors such as Dell EMC and HPE usually obtain the latest popular technologies through M&As. Over the past few years, HPE has acquired storage start-ups, such as 3PAR in 2010 and Nimble Storage in 2017, helping them offer new products.


However, M&A brings unique problems. Users find that storage products of multiple sub-brands in a data center fail to collaborate properly with each other. It is difficult to migrate services from an old storage device to a new one because sub-brand products adopt different storage operating systems at the underlying layer, making flexible data replication and migration impossible.



Like cash for enterprises, data is valuable for customers only when it is mobilized. The faster data mobility, the greater data value. When deploying a storage system in medium- and large-sized data centers, Huawei uses OceanStor OS to enable hitless data migration from disk storage to all-flash storage thanks to a unified software architecture. This allows users to store core data on flash storage and cold data on more cost-effective traditional disk storage. The data can freely mobilize between two types of storage systems according to preset policies, minimizing cost and management complexity.

As for hardware, Huawei launches chips supporting comprehensive capabilities pertaining to computing, storage, network, and management, to improve overall storage performance and reliability. Take a flash controller chip as an example. It is an essential component of an SSD. In the industry, only Huawei, Intel, and Samsung are capable of developing SSDs and SSD controller chips.

Through in-depth understanding of the underlying layer in a data center, we have designed proprietary high-performance, highly reliable SSDs. Competitors like Dell EMC and NetApp use SSDs from Toshiba, SanDisk, Micron, and Memblaze, which are powered by third-party controller chips. Huawei's HSSD controller chips adopt hard logics to implement dynamic RAID, data inspection, and LDPC algorithms, providing higher efficiency and better performance than SSD controller chips of other vendors using built-in CPU algorithms.



Moore's law, named after Gordon Moore, the co-founder of Fairchild Semiconductor and Intel, proposed the, is the observation that the number of transistors in a dense integrated circuit approximately doubles about every two years. It describes Intel's success in crossing the industrial development gap and has accurately described the development of the IT industry for decades. Moore's law not only describes technological development, but also represents a physical or natural law. Each successful product has its own "Moore's law" describing its exponential growth trend. We must have an open mind and insist on learning and self-iteration to adapt to the changing times.

Putting User Experience First

It's easier said than done when innovating services specific to users' demands. User experience includes two aspects: Product experience (pertaining to product quality) and service experience (related to online upgrade, cloud, and maintenance services).

In terms of product experience, we envision of "Data on Demand," aiming to build converged storage resource pools including all-flash storage for cloud data centers, consolidate storage infrastructure, and remove differences of underlying storage products. This enables on-demand data migration and sharing between workloads by providing automated, on-demand, and consistent data services on- and off-premises. As a result, customers enjoy simpler, more agile cloud transformation.


Adhering to its vision, Huawei's storage has been applied in many industries. In the big data cloud project of China Merchants Bank in 2017, Huawei's storage helped the bank reduce TCO by 40% and speed up the resource provisioning tenfold compared with traditional storage. The rollout time of a new service was shortened from several weeks to several hours, facilitating efficient marketing decision-making.

In terms of service experience, Huawei's storage products adopt artificial intelligence (AI) technologies and big data analysis to transform storage management, supporting intelligent device O&M and resource scheduling, and enabling data mining. This is one of the pillar fields of Huawei Storage in 2018. Huawei will release its intelligent storage management systems one after another. The soon-to-be-released eService is such an example. As an intelligent management system for remote access and control, eService will simplify storage device management and O&M and substantially increase overall O&M efficiency, cost-effectiveness, and reliability.

For example, faults may occur on running IT devices. To rapidly identify faults and recover from faults is a core reliability indicator. Huawei's unique intelligent management chips serve as the management heart of storage. The chips have built-in fault diagnosis and pre-warning expert databases to improve fault diagnosis accuracy. Fast fault diagnosis is the prerequisite for fast data recovery. The chips provide 5x higher computing capabilities per second than peer products. Data switchovers can be completed within seconds in scenarios involving controller faults, front-end interface card faults, and management board faults. The chips ensure no data loss and user-unaware non-disruptive services during the switchover.

Intrinsically Motivated to March Forward

To date, Huawei's storage products have witnessed an increasingly wider presence across the globe. The overall growth rate in sales revenue has long topped the industry. Remarkable achievements have been made in markets outside China, in particular: In 2017, the sales revenue of Huawei's storage products increased by 147% in Europe and by 79% in the Asia-Pacific (excluding China). Meanwhile, Huawei continues to improve a stable sales pattern within China. In all-flash storage alone, the sales revenue grew at nearly twice the rate of 2017, topping in the industry.




Since its launch in 2011, Huawei OceanStor Dorado All-Flash Storage series has maintained a "zero accident" record. It has been trusted by key industry customers worldwide thanks to its proven product quality and technical capabilities. A large number of industry leaders in key industries around the world have chosen Huawei's storage. Examples include Volkswagen, Vodafone, Saudi Aramco, Caixa Econômica Federal, Russian Post, Industrial and Commercial Bank of China, China Mobile, and China Pacific Insurance Group.

Many third-party analysts recognize Huawei's strengths in storage. Huawei's storage has been named a Leader in Gartner's Magic Quadrant for General-Purpose Disk Arrays for two years in a row. American technical consulting firm DCIG ranked Huawei's all-flash storage in the Recommended Ranking in the 2018-19 All-flash Array Buyer's Guide, a higher rank than Dell EMC and HPE. In standard performance tests conducted by the Storage Performance Council (SPC) and Standard Performance Evaluation Corporation (SPEC), Huawei has always broken its own records. Among the Top 10 SPC-1 By Performance products, Huawei storage occupies six positions.

In today's fast-paced world, the storage industry is undergoing profound changes accompanied by emerging hybrid clouds, intelligence, data lakes, storage media, data protection technologies, and new technical opportunities. Keeping up with Moore's law and customer demands, Huawei will continue to release industry-leading products and technologies, such as faster SCM flash storage and software-defined distributed all-flash storage powered by end-to-end NVMe architecture and better proprietary storage chips. A 1 U all-flash storage system houses over 1 PB data. Huawei strives to become the data leader in the industry.

Internally, all product lines in Huawei strive to offer the best products. It is our constant pursuit to put customers first as we continue to innovate. Huawei's storage products aim to top the world, offering unique and optimal customer experience.

We believe that our efforts and persistence allow Huawei's storage product line to shoulder more responsibilities. Let's embrace the future together. 

Huawei's All-Flash Storage Roadmap

In 2005, Huawei began to research and develop flash technologies.


In 2008, Huawei released the first SATA SSD disk, marking a milestone.

In 2011, Huawei released the first generation of all-flash storage Dorado2100 and Dorado5100, configured with Huawei-proprietary 6 Gbit/s SAS SSD disks.

In 2012, Huawei released the second generation of all-flash storage Dorado2100 G2.

In May 2014, Huawei released OceanStor V3 series converged flash storage, beginning an era of all-converged data architecture.

In May 2015, Huawei released its high-end converged storage system, OceanStor 18000 V3, setting a benchmark for high-end storage with superb stability, excellent performance, and onstop innovation.



In September 2016, Huawei released all-flash storage OceanStor Dorado5000 V3 and OceanStor Dorado6000 V3, with a gateway-free, active-active storage solution and Huawei-proprietary 12 Gbit/s SAS SSD disks, providing excellent performance.

In March 2017, Huawei released the NVMe-based all-flash storage OceanStor Dorado5000 V3.

In June 2017, Huawei's all-flash OceanStor Dorado V3 won Best of Show Award at the Tokyo Interop Expo.

In September 2017, the ESG lab (a third-party authoritative organization) released an evaluation report for Huawei OceanStor Dorado, proving it is an ideal choice for mission-critical workloads.



In January 2018, an authoritative third-party report showed that the revenue growth rate of Huawei's all-flash storage ranked first globally.

In March 2018, Huawei launched intelligent flash storage for enterprise mission-critical applications, including the OceanStor F V5 all-flash series and OceanStor V5 hybrid-flash series.

In March 2018, Huawei high-end all-flash storage, OceanStor 18800F V5, passed the most authoritative SPC-1 performance benchmark test in existence. In addition, it broke the performance record of OceanStor 18500 V3 with 6 million IOPS, indicating the highest performance ever recorded in SAN storage.

In May 2018, Huawei's eight all-flash storage products were given a Recommended ranking in DCIG 2018-2019 Enterprise General Purpose All-flash Array Buyer's Guide.


Huawei's all-flash storage does not stand out from its peers by chance, but thanks to 5000 days of continuous efforts by Huawei's 3000 storage engineers. At present, Huawei storage has been used in more than 150 countries and regions worldwide.

- China Life use Huawei OceanStor 18000F high-end hybrid-flash storage, improving the core insurance claim efficiency by three times.

- Telecom Italia (TIM) replaced an existing 55 cabinets with nine Huawei all-flash storage OceanStor Dorado cabinets, reducing the OPEX by 72%.

- Cliniques Universitaires Saint-Luc (a non-profit hospital in Belgium) uses Huawei all-flash file storage, improving the concurrent access scale by five times and has constructed an active-active solution, ensuring 24/7 service continuity.

Huawei flash storage will continue to innovate for mission-critical enterprise applications and take the lead in the all-flash, all-cloud, and all-intelligence data storage era. ⚡



COPPER KEY: DEVICES & BLACK TECHNOLOGIES

DCIG All-Flash Array Buyer's Guide: An Authoritative Report

Three Things You Should Consider When Buying Flash Storage

Huawei All-Flash Reliability Assurances for Your Mission-Critical Businesses

Huawei's All-Flash Arrays Offer Record-Breaking SAN and NAS Storage Performance

OceanStor OS Enables Fast and Stable All-Flash Storage

Huawei Accomplishes 3:1 Data Reduction Rate in All-Flash Offerings

Four Questions to Understand All-Flash Performance

FlashLink: The Secret to Huawei's High-Performance All-Flash Storage

Chips Purpose-Built for All-Flash Storage: the Core of Competitiveness

DCIG All-Flash Array Buyer's Guide: An Authoritative Report

By Wang Zhen from Huawei

All-flash storage has been positively affecting the digital transformation of enterprises. The first sentence of the DCIG 2018-19 Enterprise General Purpose All-Flash Array Buyer's Guide indicates the significance of all-flash storage: "Any organization that has yet to adopt an all-flash storage infrastructure for all active workloads is operating at a competitive disadvantage". This shows that the use of all-flash applications is supported and recommended by DCIG.



Given that some readers are not familiar with the DCIG Enterprise General Purpose All-Flash Array Buyer's Guide, this article extracts some of the key points.

The DCIG guide provides users with comprehensive and in-depth analyses and suggestions for product and technology procurement, helping buyers save time and money. The evaluations in the guide cover business value, integration efficiency, and data reliability. Many CIOs regard the guide as an important reference when planning for and purchasing data center infrastructures.

In this guide, DCIG scores and ranks 32 all-flash storage products from 7 mainstream storage vendors (Dell EMC, Hitachi Vantara, HPE, Huawei, NetApp, Pure Storage, and Tegile) based on 57 evaluation items in terms of performance, management and software, hardware, virtualization, and technical support. Eight analysis steps are performed to classify and analyze the preceding evaluation items. The products of different vendors are given Recommended or Excellent rankings based on their scores. Only all-flash storage products with the best overall performance and highest scores are awarded a Recommended ranking.

Three models of OceanStor Dorado V3 series and five models of OceanStor F V5 series are listed with a Recommended ranking. Following the Gold Medal received at Interop in Tokyo 2017, Huawei's all-flash storage has once again receives high recognition in the market and demonstrated its technical strength.

RANKING	PRODUCT
RECOMMENDED	Huawei OceanStor Dorado6000 V3
	Huawei OceanStor Dorado5000 V3 (NVMe)
	Huawei OceanStor Dorado5000 V3 (SAS)
	Huawei OceanStor 6800F V5
	Huawei OceanStor 5800F V5
	Huawei OceanStor 5600F V5
	Huawei OceanStor 5500F V5
	Huawei OceanStor 5300F V5
	Hitachi Vantara VSP F1500
	Hitachi Vantara VSP F800
	Hitachi Vantara VSP F600
	Hitachi Vantara VSP F400
	HPE 3PAR StoreServ 20450
	HPE 3PAR StoreServ 9450
NetApp AFF A300	
EXCELLENT	Dell EMC VMAX 450FX AFA
	Dell EMC VMAX 250FX AFA
	Dell EMC Unity 650F All-Flash
	Dell EMC Unity 550F All-Flash
	HPE Nimble Storage AF9000 AllFlash Array




DCIG awarded Huawei OceanStor Dorado a Recommended ranking for the following reasons:

1. OceanStor Dorado5000 V3 is among the first batch of storage arrays that commercially used NVMe.
2. Huawei-proprietary SSDs and the disk controller collaboration of OceanStor Dorado arrays ensure high performance, low latency, and a long service life.
3. Huawei introduced the scale-out architecture of SmartMatrix to OceanStor Dorado, allowing it to support up to 16 controllers. In addition, Huawei increased the density and capacity by 1 to 4 times thanks to 15.36 TB SAS and NVMe SSDs. This allows a single cluster to provide up to 36.9 PB of raw flash capacity.

DCIG also includes OceanStor 6800F V5, 5800F V5, 5600F V5, 5500F V5, and 5300F V5 in the Recommended ranking for the following reasons:

1. These flash products provide unified storage, powerful specifications, and high density. The mature and multi-functional OceanStor OS has endured continuous tests by last-generation products and been widely used in applications and data center environments.
2. These arrays use LUN-level SmartQoS to optimize the resource utilization of core applications, meeting IOPS, bandwidth, and latency requirements.
3. OceanStor F V5 disk arrays use SmartVirtualization to streamline storage infrastructure and integrate heterogeneous storage systems. SmartVirtualization eliminates data silos, simplifies storage management and enables data migration and disaster recovery processes.

In addition, DICG summarizes the development technologies of all-flash storage in the industry:



1. DCIG believes that inline compression and deduplication, supported by most all-flash arrays, improve the value, service life, and performance of all-flash storage and maximize the available storage capacity. In contrast, post-process deduplication adversely affects the performance and service life of all-flash storage. This is because it increases the number of writes in underlying flash media, shortens the service life of media, and can reduce production efficiency during peak hours.

2. 12 GB SAS interfaces are still the mainstream interfaces of flash media and other non-volatile memory. However, NVMe enables all-flash storage to respond with sub-millisecond speeds, and supports millions of IOPS, real-time analysis, and transactional processing. Therefore, NVMe has gained the favor of enterprise storage customers the world over.



Get More Tips


Report: DCIG 2018-19 All- Flash Array Buyer's Guide

DCIG released the 2018-19 All-Flash Array Buyer's Guide, in which Huawei's all-flash array earned the Recommended Ranking.

Three Things You Should Consider When Buying Flash Storage

By Chen Yong from Huawei

Applications of flash storage are poised for explosive growth, thanks to increasingly lower-priced and higher-capacity SSDs. In fact, the newest generation of 3D NAND flash technology has now reduced the cost of a single SSD to less than US\$1 per GB.



As a result, more and more users are using all-flash arrays to build their own data centers. All-flash storage is expected to account for nearly 30% of storage in data centers in 2018, according to third-party authorities.

All Solid-State Storage Data Centers Accelerate Digital Transformation for Enterprises

Gartner



Solid-state arrays have reached a level of maturity that allows them to surpass hybrid disk arrays in reliability, affordability, and agility—all respects except raw capacity.




Using SSDs throughout a data center can reduce operational costs in terms of storage array administration, physical floor and/or rack space, power, and cooling.



Higher-capacity SSDs in denser storage systems will enable I&O leaders to better support the digital business by enabling new types of applications.

Source: Gartner's Moving Toward the All Solid-State Storage Data Center in 2016



All-flash storage benefits IT systems in many ways:

1. It improves user experience. All-flash storage can quickly respond to service requests. For example, in the securities industry, all-flash storage enables traders to smoothly process online transactions. In the medical industry, registration, diagnosis, and treatment can be quickly coordinated, and patient information can be swiftly retrieved from the hospital IT system.

2. It reduces O&M costs. Traditional storage systems need dozens of disk cabinets and tens of thousands of HDDs to meet performance requirements. Now, the required performance can be delivered using just a few all-flash cabinets. This significantly reduces procurement costs, power consumption, space requirements, and maintenance manpower, all while slashing OPEX by more than 75%.


3. It supports discovery of new business opportunities. For example, consider a cloud service provider that uses the original IT architecture and provides only gold, silver, and copper services. All-flash storage could help that service provider provide a platinum cloud space leasing service with higher performance and lower latency. Deduplication and compression make the provider's services even more commercially competitive.

Considering the massive potential impact on your business, the decision to build an all-flash data center should not be taken lightly. All-flash storage vendors will often push for you to make the switch without giving you the full story. When the time comes to make the switch to all-flash, here are three crucial points to keep in mind:



1. Hardware architecture should offer large capacity and efficient data exchange

In the past 10 years, SSD performance has seen a 4000-fold improvement, but CPU performance is only about 20 times better. The bottleneck at data centers lies in CPU usage, because it reaches 90% or higher when only a dozen disks are used. To get the most out of SSD performance, the latest CPU must be used to meet business requirements. This is especially true for front-end financial transactions, carrier CRM, billing, and any core business with high storage performance demands. Only with the right all-flash hardware architecture can these needs be met.



The use of SSDs drives demand for larger cache capacity. Data on SSDs is scattered all over the place. As a result, corresponding metadata increases as storage data grows. In particular, after deduplication and compression are enabled, metadata must be frequently invoked and fingerprints must be compared before data is flushed. In the cache, more and more metadata is stored. Therefore, every all-flash array must offer a cache with large capacity.

High-speed interconnection between controllers is also necessary. Data exchange channels must be efficient enough for all-flash storage to be fully utilized in the multi-controller layout. Therefore, most storage controllers are interconnected using PCIe 3.0 to improve switching bandwidth.



2. All-flash data centers need software architecture designed for all-flash storage

Your software architecture cannot be wasting SSDs at places such as protocol parsing, data flushing, and space reclamation. In addition, to maximize SSD performance, read and write I/Os must have the highest priority in obtaining resources.

All-flash storage requires end-to-end all-flash architecture design, including flash-oriented CPU multi-core optimization, cache-adaptive algorithms, link transmission optimization, SSD scheduling algorithms, and drive optimization. Without these elements, in-depth software-hardware integration is out of the question. Proper all-flash software architecture allows storage controllers to intelligently sense each SSD's data layout, enables both controllers and SSDs to adjust data layouts and I/O priorities as needed, and leverages disk and controller algorithms to deliver optimal performance.

In addition to SSD performance, the cost effectiveness of deduplication and compression must be considered. Deduplication and compression greatly reduce the amount of physical storage space needed, and alleviate investment, environment, and maintenance pressures. Deduplication and compression can also effectively reduce the number of write times and amounts of data written to hosts, prolonging the service life of SSDs and protecting customer investments.



3. Strong security and reliability are critical to centralized data access

With all-flash storage, performance is 10 times better than with HDD-based storage. This means that a single all-flash storage system can host many more services, such as CRM systems, billing systems, B2B reselling, server virtualization, and office systems. If you're going to put all your eggs in one basket like this, that basket had better be extremely reliable.

Vendors have their own ways of ensuring high system reliability. Some vendors improve their underlying algorithms, some use multi-controller redundancy architecture, and some leverage active-active or geo-redundant DR solutions. It can be challenging to evaluate a vendor's approach to ensuring reliability, but the key factors are as follows.

First, check whether underlying virtualization technology is used and whether physical disks can be logically scattered and effectively organized. These features prevent the permanent data loss that can be caused by time-consuming data reconstruction using traditional RAID groups. Underlying virtualization technology enables all disks in the disk group to simultaneously reconstruct data if one disk is faulty,



instead of reconstructing data on a single hot spare disk. This dramatically shortens reconstruction time. Currently, only HPE and Huawei are capable of offering this technology. Huawei's all-flash storage can reconstruct 1 TB of capacity in just 30 minutes. That's 10 times faster than traditional storage.


Next, check whether the architecture can provide multi-controller redundancy. In other words, if multiple controllers on a single engine fail, can service continuity be ensured? Only HDS and Huawei currently offer this ability.

Last but not least, check whether the active-active solution is simple and reliable:

- Is the gateway-free SAN and NAS integration mode used? This mode can reduce faults and latency while cutting procurement costs.
- Does the solution enable access in active-active mode, rather than active-passive mode? This mode can implement rapid site failovers and achieve $RTO \approx 0$, and $RPO = 0$.
- Is a simple networking mode adopted? IP or FC networking can be used to meet the requirements of different users. NetApp and Huawei are the only vendors in the industry that can implement this architecture.

The Ideal Choice for Mission-Critical Business

Considering hardware architecture, software architecture, and reliability, there is only one solution available that checks all the boxes: Huawei OceanStor all-flash storage.


A storage system specially designed for flash storage, OceanStor boasts both industry-leading reliability and superior performance. Equipped with SmartMatrix 2.0 (enhanced intelligent switching matrix), this storage system features full redundancy with 4 controllers in a single engine. In addition, the gateway-free HyperMetro-based (active-active) solution ensures 24/7 uninterrupted services for users and ensures a latency of less than 1 millisecond. With OceanStor, you can rest assured that your data center is supported by the best in the industry. 

Four Questions to Understand All-Flash Performance

By Qin Xuan from Huawei

In recent years, flash has become a hot topic. All-flash arrays are critical for enterprises to cope with explosive data growth and to accelerate mission-critical applications. According to data from IT Brand Pulse, the market share of enterprise-level HDDs has been declining since 2015. In 2017, the market share of SSDs exceeded that of traditional HDDs. It is expected that all-flash arrays will soon become mainstream data center storage and be used for key enterprise business.

Today, let's talk about the performance of all-flash arrays. We all know that the all-flash arrays have high performance, and we often see storage vendors painstakingly promoting all-flash performance via mainstream media in the market. All-flash arrays can offer a performance ranging from hundreds of thousands of IOPS to millions of IOPS or even higher. Therefore, we must ask four questions to learn more about the performance of all-flash arrays.



1. What IOPS and latency can all-flash arrays provide?

Vendors must provide their IOPS and latency during promotion.

Generally, random read/write performance of an enterprise-class 10,000 rpm SAS disk offers about 200 to 300 IOPS. In comparison, the performance of a single SSD can reach thousands or even tens of thousands of IOPS, depending on the SSD type (SLC, MLC, or TLC). When the performance of an all-flash array exceeds a certain threshold, the latency increases accordingly. Therefore, when vendors claim that they can provide very high levels of performance, we must ask one question: how long is the latency? 0.5 ms? 1 ms? or 5 ms?

Huawei OceanStor Dorado V3 excels in both performance and latency. The following figure shows the performance data of OceanStor Dorado6000 V3, which was tested by the Storage Performance Council. The average latency is just 0.38 milliseconds, and when performance reaches the maximum IOPS of 1 million, its latency is less than 0.5 ms.

In addition, we must check whether the IOPS is stable or maximum performance.

2. What are application scenarios and workloads?

The performance of all-flash arrays varies with application scenarios, application models, and workloads. Workloads we are discussing here include:


- I/O size (8 KB, 32 KB, 64 KB, or 128 KB)
- Ratio of reads to writes (that is, the ratio of read I/Os or write I/Os in a host I/O request)
- Sequence and randomness of I/Os
- Other factors (such as cache hit ratio)

In this case, any change to a parameter (for example, 8 KB I/Os and 64 KB I/Os, 100% random reads/writes and sequential reads/writes, 7:3 mixed reads/writes, and 100% reads) can greatly affect performance. Therefore, when seeing very high levels of performance, you must ask the vendor what the system workloads like in order to achieve this performance.

In most real-world production environments, database OLTP is a typical application scenario. Therefore, when evaluating the performance of all-flash arrays, the commonly used workload models are 8 KB I/Os, 100% random I/Os, and read/write ratio 7:3.




3. What is the performance of all-flash arrays when value-added functions such as deduplication and compression are enabled?



For all-flash arrays, inline deduplication and compression not only improve the service life of SSDs, but also help you reduce space and power consumption, further slashing costs. IDC reports show that data reduction functions such as inline deduplication and compression are key factors and technologies for reducing the costs of all-flash arrays.

To keep up with this trend, most traditional storage vendors have launched all-flash array models or versions. However, due to inherent weakness in the architecture, even though deduplication or compression is supported, all-flash array performance is seriously affected when deduplication, compression, or snapshots are enabled.

Take vendor H's all-flash product as an example. The all-flash array is configured with two controllers, 48 x 480 GB SSDs, and RAID 10 groups, allowing the array to provide about 150,000 IOPS (workloads: random 8 KB I/Os, read/write ratio 7:3, and 1 ms latency). When RAID 5 groups are used, the array can only offer about 130,000 IOPS. However, after enabling the deduplication function, performance is halved, providing only 60,000 to 70,000 IOPS. If other value-added features (such as snapshots, replication, or active-active) are enabled, performance will see further massive decreases.



Therefore, we should focus on maintaining excellent performance while value-added functions are enabled. This is much closer to real-life application environments.

4. What is the long-term steady performance when capacity usage reaches 80%?



For example, at a customer's POC site, vendor E's all-flash arrays showed outstanding performance to begin with. However, as the amount of test data increased, performance deteriorated. The root cause of this was that the garbage collection mechanism started running in the background to handle increasing data volumes. The mechanism's running mode and efficiency have a major impact on all-flash array performance.


A simple way to establish or simulate a real-life application scenario is to embed around 80% of data (or perform random overwrites for one time) before starting an official performance test. By doing this, you can easily test the long-term steady performance of all-flash arrays.

The OceanStor Dorado V3 uses the global garbage collection mechanism, which is triggered as a background task when the ratio of idle blocks reaches a specified threshold. The OceanStor Dorado V3 uses system-level end-to-end I/O priority to ensure that I/Os generated by garbage collection have the lowest possible impact on host I/O performance, thereby ensuring the stable latency of host I/Os.

- On-demand start and stop of garbage collection: Controls matching between garbage collection bandwidth and front-end write I/O bandwidth, migrates data on demand, and reduces write amplification, minimizing the impact of garbage collection on host performance.

- Optimized cost-benefit policy used to select target data for garbage collection: Separately stores cold and hot data (separation of metadata and data, frequently accessed data blocks and seldom accessed data blocks), moves invalid data to idle blocks based on the cost of releasing space, and batch reclaims data written at the same point in time. This reduces unnecessary data migrations and minimizes the impact of garbage collection on application performance.

The following figure shows a performance comparison between Huawei OceanStor Dorado5000 V3 and other mainstream all-flash arrays when under strict test conditions (two controllers, random 8 KB I/Os, read/write ratio 7:3, 80% occupation rate, inline deduplication and compression enabled, and 0.5 ms latency). We can see that Huawei's product provides superior performance compared to the other storage arrays.



Test conditions: dual - controller, 80% capacity occupation rate, inline dedupe & compression, 0.5ms latency

Summary

Equipped with native, flash-architecture design capabilities, Huawei is the only vendor that develops SSD chips, SSDs, and all-flash systems. Huawei enables disk and system collaboration to make the most of all-flash features, and Huawei all-flash arrays deliver the highest performance and reliability in the industry by using operating systems specially designed for all-flash arrays, unique FlashLink technology, and dedicated performance acceleration chips.

We hope that this article helps you to look at dazzling promotions and understand the true performance capabilities of all-flash arrays.

Huawei has demonstrated its amazing performance in multiple all-flash competition tests. Thanks to its strong product competitiveness, Huawei's all-flash arrays are growing rapidly. In the first three quarters of 2017, Huawei's all-flash storage revenue enjoyed the fastest growth speed in the world, and started to serve many famous enterprises worldwide such as Vodafone, Saudi Arabia Aramco, Russian Post, Brazil's Caixa, and China Pacific Insurance. ⚡



Overview of Huawei OceanStor All-Flash Storage Products

OceanStor
18500/18800 V5

OceanStor V5

Provides converged block and file data services with high performance and unwavering reliability to perfectly host core enterprise services.

OceanStor Dorado V3

Easily copes with enterprises' diverse data integration needs with ultimate performance experience and powerful functions.

OceanStor
6800 V5


OceanStor
Dorado6000 V3

OceanStor
Dorado5000 V3

OceanStor
5600/5800 V5

OceanStor
5300/5500 V5

Lightning Fast
Rock Solid



Huawei's All-Flash Arrays Offer Record-Breaking SAN and NAS Storage Performance

By Wang Jiaxin from Huawei

Excellent performance is an everlasting pursuit for storage system vendors. Vendors use different performance parameters to promote product performance, such as 100% cache hit ratio (maximum IOPS), 256 KB I/O bandwidth, and 100% write latency. All vendors want you to believe that they can provide the highest performance and can easily handle business pressure during peak hours. You may want to know how to identify all-flash arrays that truly have high performance from amongst a huge number of potentially misleading promotions.


Track 1: 100% cache hit ratio

When promoting storage performance, many vendors often boast they can achieve "100% cache hit". This means that all data is stored in memory, instead of being stored permanently. A criterion to determine whether IT architecture is high performance is to check whether it stores the most frequently accessed data in the location that has the fastest response times.

This essentially means that it is almost impossible to have a 100% cache hit ratio in customer environments. This is because all-flash storage applies to not only different industries but also to varying business models within industries. Therefore, such a claim must be taken with a grain of salt, and customers should thoroughly check their needs against what the product can offer.

Trick 2: 100% read performance

Mainstream SSDs offer high performance in data reads, but their write performance is lacking because each time data is written, SSDs erase data from a NAND before writing new data. This process is called "Program/Erase", which takes 1 ms to 2 ms in MLC/TLC NAND flash, and leads to huge performance differences in data reads and writes.



Program/Erase in NAND flash

SSDs often reserve a large amount of over-provisioning space and the "Erase/Program" process is not required every time data is written onto SSDs. Generally speaking, SSD performance suffers more in mixed read and write scenarios than in read-only scenarios. Therefore, real-life performance values must reflect customer business models (such as 8 KB I/O data blocks, read/write ratio 7:3, and RAID 5/6 groups), instead of boasting the value of 100% read performance.

Application	Seek Type	I/O Request Sizes	Percentage of I/O as Writes
Microsoft Exchange	Random	32 KB	Moderate to high
SAP/Oracle applications	Random	~8 KB	Depends on application
RDBMS:Data entry/OLTP	Random	Data base or file system page size	Moderate to high
RDBMS:Online (transaction)logs	Sequential	512 byte+	High,except for archiving process
RDBMS:temp space	Random	Data base or file system page size	Very high

Data models in typical scenarios

Trick 3: Performance stability

Many vendors only promote reference values that are achieved using a specific model and under certain conditions, which does not show their products' stability.



The stability of storage performance is critical, especially during service pressure changes and in complex service environments. When choosing storage SLA, customers regard stability as one of the most important indicators, encouraging vendors to produce limited results, which mislead their customers. Stability indicates that response times of 99% of services must remain stable under certain conditions. If services fluctuate, user experience suffers. However, most customers are, unfortunately, unaware of this issue.

SPC, the most commonly used third-party platform for evaluating SAN storage performance

Test standards (especially SPC-1 tests for random IOPS effectiveness) proposed by the Storage Performance Council (SPC) are commonly used as a reference for measuring SAN storage product performance in the industry. Mainstream vendors such as EMC, NetApp, Dell, HPE, HDS, Huawei, IBM, and Fujitsu have participated in SPC tests and used the results to promote their product performance and cost-effective features. This makes SPC-1 a well-known indicator for measuring the performance of products from various brands.

The test specifications of SPC-1 are both strict and fair. Its test program prevents deliberate data value manipulation and its test models are similar to mainstream transaction workloads, meaning SPC-1 can be reliably used as a reference to reflect performance in real-life environments.

Service type: For production services, the SPC provides the SPC-1 benchmark test guide. For data analysis services, the SPC provides the SPC-2 benchmark test guide.

Mixed service workload: In the SPC-1 test, multiple service workloads are used for concurrency testing to verify response times of storage systems in mixed service scenarios. The main indicators include different data types (text, binary, and sparse), different block sizes (8 KB to 128 KB), and different read/write ratios, such as sequential reads and writes or random reads and writes.

Parameter Type	I/O Stream 1	I/O Stream 2	I/O Stream 3	I/O Stream 4
ASU	1	1	1	1
Transfer alignment (512 byte blocks)	16	16	16	16
Stimulus group	1	100	1	100
Intensity multiplier	0.035	0.281	0.070	0.210
Memory alignment	16	16	16	16
Model type	Open	Open	Open	Open
Population	1	1	1	1
Read fraction	0.5	0.5	1.0	0.5
Stream identifier	"ASU 1-1"	"ASU 1-2"	"ASU 1-3"	"ASU 1-4"
Transfer address	Uniform: 0.0 – 1.0	R1/W1: 0.15 – 0.2	Incremental (0.4, 0.4, 0, 0.1)	R1/W1: 0.7 – 0.75
Transfer size (512 byte blocks)	16	16	SMIX	16
Workload identifier	"SPC-1.30"	"SPC-1.30"	"SPC-1.30"	"SPC-1.30"

SPC- 1 testing model

Varying service pressure: The SPC increases and decreases dynamic workload behaviors to test the storage performance of varying service pressure. Specifically, the SPC starts tests from the peak value before gradually decreasing and then increasing service pressure. The SPC then simulates two service peaks as shown in the following figure.

SPEC, the most commonly used third-party platform for evaluating NAS storage performance

Founded in 1988, the Standard Performance Evaluation Corporation (SPEC) is a global and authoritative application performance testing organization, designed to formulate a series of application performance evaluation standards. SPEC is a third-party organization that consists of dozens of prestigious and influential universities, research institutes, and IT enterprises.

SPEC uses the latest SPEC SFS®2014 standard performance evaluation benchmark to test file server throughput and response times, and

provides a standard measurement method for comparing the storage performance of devices from different vendors in specific application scenarios. This benchmark uses software to establish swbuild scenarios and tests NAS performance when handling metadata-intensive workloads. NAS products from multiple brands, including Huawei and NetApp, have released the performance evaluation results obtained in this scenario.



Huawei products with record-breaking performance

With 13 years of experience, Huawei has designed intelligent chips, hardware platforms, and operating systems for its all-flash arrays, and has innovatively used FlashLink technology to fully utilize the capabilities of all-flash arrays.


Three intelligent chips implement end-to-end acceleration for transmission, storage, and management. Specifically, the multi-protocol interface chip accelerates data transmission, the SSD controller chip accelerates data access, and intelligent management chip speeds up fault management. At the hardware layer, the latest-generation Purley CPU, along with dual-uplink acceleration are used to provide a stable chassis for all-flash arrays.

At the operating system layer, a unique disk-controller collaboration algorithm is developed based on the all-flash array features. This algorithm enables SSDs and storage controllers to work together in order to make the most of all-flash systems.

In March 2018, Huawei's next-generation high-end all-flash storage OceanStor 18800F V5 went through the most authoritative SPC-1 performance benchmark testing currently available and set a new SAN performance record, with 6,000,572 IOPS, breaking that previously set by the OceanStor 18500 V3.




Response Time and Throughput Graph



The OceanStor 18800F V5 has a maximum performance of 6 million IOPS and a latency of less than 1 ms. Statistics show that 99.984% I/O responses are within 1 ms.


OceanStor 18800F V5 IOPS in SPC-1 testing



Business Metric (Builds)	Average Latency (msec)	Builds Ops/Sec	Builds MB/Sec
100	0.152	50002	646
200	0.148	100004	1293
300	0.168	150006	1941
400	0.567	200009	2586
500	0.411	250008	3235
600	0.559	300012	3881
700	0.439	350015	4528
800	0.658	400015	5175
900	1.279	150011	5822
1000	2.073	500009	6469

Unrivaled performance both the SAN and NAS fields reflects the advantages of Huawei's next-generation intelligent all-flash storage systems in parallel architecture, and helps Huawei to continuously provide state-of-the-art storage products, ensuring a better data management experience for customers.

This is also not the first time that Huawei storage has reached the top. Huawei OceanStor has been participating in SPC certifications since 2010, and has broken the SPC-1 performance testing records multiple times, surpassing IBM, Hitachi and NetApp. ⚡



OceanStor 6800F V5 performance in SPEC testing



Get More Tips

SPC-1 Report: Huawei all-flash Storage OceanStor 18800F V5

Huawei OceanStor 18800F V5 achieved 6,000,572 IOPS in SPC-1 benchmark testing.



Get More Tips

SPEC Report: Huawei All-Flash Storage OceanStor 6800F V5

During the SPEC SFS2014_swbuild benchmark test, Huawei OceanStor 6800F V5 provides 500,009 operations per second (OPS).

Huawei All-Flash Reliability Assurances for Your Mission-Critical Businesses

By Wang Jiaxin from Huawei

High-performance flash storage carries mission-critical business systems for various industries. If a problem occurred, enterprises would be hit hard. Qualix Group released figures to show the impacts of business interruption. In transportation, a one-minute stoppage would result in average losses of 150,000 USD, while that for banks would be 270,000 USD. The same one-minute stoppage for a telecommunications company would cost an average of 350,000 USD, while manufacturing would be hit with a 420,000 dollar loss, and security traders would top the list, losing 450,000 USD.

Therefore, ensuring mission-critical business continuity is a top priority for all-flash storage systems. Reliability, designed in an end-to-end way, is no easy task. Looking from media to systems and solutions, let's examine how Huawei OceanStor Dorado all-flash storage can provide high performance and reliability for customers.

Disk-level reliability

SSD reliability is measured by examining the mean time between failure (MTBF) and annualized failure rate (AFR). The industry MTBF benchmark is between 2 and 2.5 million hours. Huawei raises the bar well beyond this, reaching 3 million hours between failures on its homegrown disks.

How does Huawei accomplish this feat and extend the life of its SSDs? Huawei has maintained long-standing cooperation with its vendors, such as Samsung, Micron and Toshiba, to ensure that components are manufactured according to Huawei's solution design objectives. Another reason is the extensive cooperation achieved between arrays and disks, which combines a series of reliability designs (such as optimization in graphene dissipation technology (GDT), global wear leveling, and global anti-wear leveling).

1. Global wear leveling design

At the beginning of the SSD lifecycle, service loads are spread in a balanced manner throughout SSDs to avoid overloading specific disks. This leads to the idleness of some disks and premature retirement of others.

2. Global anti-wear leveling design (Huawei-patented)

At the end of the SSD lifecycle, when the wear on an SSD exceeds 80%, the anti-leveling mechanism takes the gradient wear to more than a 2% difference to avoid simultaneous disk failure. The service life of the system can be prolonged by gradually replacing the disks, ensuring sufficient time for system upgrade.

3. System software optimization

First, at the algorithm layer, Huawei is the first vendor to commercialize the LDPC algorithm in SSDs. After years of optimization, Huawei now supports a 4 K ultra-long code algorithm. This brings the error correction granularity to twice of other SSD providers in the industry.

Second, at the flash chip layer, the number of erase cycles is limited in an SSD. The service life of an SSD can be prolonged if the number of erase cycles can be increased through algorithms. Huawei's innovative adaptive program & erase (APE) technology automatically controls the erase strength and frequency of flash chips based on the amount of read and write data. In this way, the number of erase cycles can be effectively extended without changing costs or media granules, prolonging the SSD service life.

Third, at the data protection layer, while the storage controller system has RAID protection, SSDs also support two-dimension RAID groups with interleaving parity at channel and CE levels, ensuring chip-level failure data protection. The disk RAID and system RAID groups work together to conduct automatic data recovery if multiple chips of a single disk are faulty. Then, after being recovered, SSDs will be operational again.



System-level reliability

Achieving reliability is complex. In addition to the hardware structure design and software fault tolerance mechanism, the storage system must tolerate physical and logical faults and support quick recovery. This will prevent data loss caused by system faults and ensure businesses continue running stably.

1. Magnitude 9.0 earthquake-resistant design

The irregular seismic waves and intensified shaking caused by huge earthquakes will affect the stability and service life of electronic equipment. Huawei OceanStor Dorado all-flash storage has passed the magnitude 9 earthquake-resistant test run by China Telecommunication Technology Labs (TTL). This makes Huawei the only company to have done so and satisfy the TTL's IT standards. Once an exception is detected, the system can also diagnose and rectify the fault quickly enough to prevent business interruption.

2. Tolerance of three-disk failures

Disk capacity increases linearly with disk reconstruction time. Traditional RAID 5 or RAID 6 technologies allow 5 hours for the reconstruction of 1 TB of data, and 80 hours for 16 TB. However if one or two more disks become faulty during reconstruction, systems running RAID 5 or RAID 6 are unable to cope, severely disrupting business. Therefore, traditional RAID technologies cannot ensure system reliability, causing data loss and business interruption.


Huawei's innovative RAID-TP software technology is based on the Erasure Code (EC) algorithm. Parity bits support 1-, 2-, 3-dimensions and can tolerate 1 to 3 simultaneous disk failures. This means that in the case of three disk failures, the system will not suffer from data loss or service interruption. Currently, only products from Huawei, NetApp, and Nimble can tolerate the simultaneous failure of three disks.

Although NetApp and Nimble can tolerate simultaneous failures of three disks, they both use traditional RAID architecture with fixed data disks and hot spare disks. For these companies, hot spare disk reconstruction for 1 TB of data takes 5 hours. OceanStor Dorado employs a global virtualization system able to reconstruct the data in just 30 minutes, fulfilling the requirements of ultra-large capacity profiles.

3. End-to-end data integrity protection and tolerance of silent data corruption

In data access, any errors that occur can cause issues for data integrity when data is transferred through multiple components, channels, and complex software. However, such errors can only be detected in subsequent data checks and access. This phenomenon is called silent data corruption.

Often overlooked, silent data corruption has greatly impacted services, such as databases, that require absolute data integrity. Launched by Huawei, Emulex, and Oracle, the data integrity solution changes the traditional condition where hosts and storage systems protect data independently. This has been achieved by implementing end-to-end protection across applications, hosts, storage systems, and disks. As a result, this solution prevents silent data corruption for mission-critical businesses and eliminates potential down times.



4. Intelligent prefetch

When a disk detects block faults or even severe die failures, the storage system receives failure reports from SSDs and uses redundant data in RAID groups to rapidly reconstruct and repair damaged data, reducing data loss risks and ensuring system reliability.

Huawei all-flash storage systems can accurately query internal data, such as SSD data, and use innovative prediction algorithms to monitor and predict the service life of disks. The personnel in charge of customer businesses will be told that their disks need replacing before the disks become faulty or one month before the service life is exhausted.



Solution-level reliability

Huawei OceanStor Dorado all-flash storage supports multiple data protection technologies, such as snapshot, clone, remote replication, and active-active data protection. This allows it to implement data protection solutions from local or intra-city to remote disaster recovery. This solution provides high availability and non-disruptive storage data services for customers, preventing data loss caused by logical or physical disasters.

1. Lossless snapshot

Traditionally, COW-based snapshot technology requires data to be written to a location after being read and migrated to a new location. Therefore, such snapshot processes involve one read, two writes, and one metadata update. COW-based snapshot affects system performance due to performance loss during each data migration.

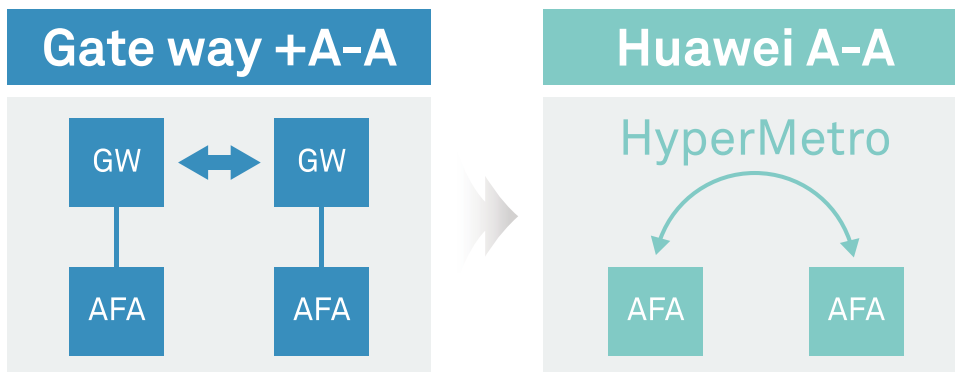
Huawei OceanStor Dorado all-flash storage implements lossless snapshot using ROW. When a snapshot is activated, data is written to the new location and the pointer of the mapping table is modified. Only one data write and one metadata update are involved, with data operation complexity being only 1/3 of that seen for COW-based snapshot. In addition, no extra data migration is required when the ROW snapshot is activated, resulting in no compromises in performance regarding production businesses.



In addition, OceanStor Dorado storage supports second-level periodic snapshot, which is superior to the minute- or hour-level snapshots used by competitors' all-flash storage. OceanStor Dorado snapshot provides users with a more intensive and powerful continuous data management (CDM) solution, enabling real-time data protection.

2. Gateway-free active-active architecture

Huawei OceanStor Dorado storage adopts a gateway-free active-active layout, removing the gateways on both sides. This immediately reduces customer procurement costs and lowers possible failures, achieving reduced latency, improved reliability, and accelerated performance. In addition, the overall networking is greatly simplified, with the number of deployment steps halved, thereby shortening the delivery cycle.



Active-Active architecture

HyperMetro is deployed on two arrays in an active-active profile. Data on the active-active LUNs at both ends is synchronized in real time, and both ends process read and write I/Os from application servers to provide the servers with parallel active-active access. Should either array encounter a fault, services are seamlessly switched to the other end without interrupting service access, achieving RPO = 0 and RTO ≈ 0.

In remote data protection scenarios, the active-active solution can be effortlessly upgraded to the data center solution in geo-redundant mode, requiring no extra gateways and causing no business interruptions. This allows it to deliver a huge reliability protection rate of 99.9999% for customers. Third-party sites can even use Huawei OceanStor converged storage systems to provide cost-effective DR solutions for remote DR centers that require only ordinary response times.

Summary

All-flash storage focuses on performance and efficiency. Like a giant container ship on the sea, all-flash storage continuously pursues higher speeds and a larger capacity. With continuous and stable operations, Huawei OceanStor Dorado all-flash systems can deliver 99.9999% reliability, providing the public with a lightning-fast, rock-solid platform. ⚡



Get More Tips

IDC White Paper: Huawei HyperMetro-Active-Active Data Centers

HyperMetro-Active-Active Data Centers Help Enterprises to Safeguard Business Continuity in the Era of Digital Transformation.

Huawei OceanStor Dorado V3 All Flash Storage 3:1 Data Reduction Guarantee



The transition to all flash storage is around the corner. The biggest challenge facing data centers today is reducing TCO to make all flash storage more affordable for their users.

OceanStor Dorado V3, a Next-Gen All Flash Storage System with Excellent Performance, Reliability, and Efficiency

OceanStor Dorado V3 leverages FlashLink technology, specifically designed for flash storage, to provide supreme performance while maintaining a consistent latency of 0.5 ms. This allows enterprises' core applications to respond more quickly without compromising their reliability and data processing efficiency.

The OceanStor Dorado V3 supports a gateway-free active-active design, guaranteeing 99.9999% availability. It also includes cutting-edge inline deduplication and compression technologies to provide high performance at lower procurement costs and TCO.

Huawei 3:1 Data Reduction Guarantee

Huawei OceanStor Dorado V3 All Flash Storage Guarantees 3:1 Data Reduction for Higher ROI at a Lower Investment

If the 3:1 data reduction set forth in this guarantee is not met, Huawei is liable for providing additional effective capacity^[1] or exempting the customer from future costs on the capacity.^[2]

Prerequisites

1. Applicable to OLTP/OLAP database, virtual desktop infrastructure (VDI), and virtual server infrastructure (VSI) scenarios
2. Excludes compressed or encrypted data, or other types of data that cannot be reduced in size
3. Inline deduplication and compression must be enabled

3:1 Data Reduction Guarantee, Accelerating the All Flash Storage Transformation of Data Centers

[1]: Effective capacity refers to writable capacity available to hosts after data reduction functions are enabled.

[2]: For specific conditions and procedures, please refer to supporting documents.

[3]: Refer to best practices for deploying the OceanStor Dorado V3 in OLTP, OLAP, VDI, and VSI scenarios for detailed configurations.

Huawei Data Reduction Guarantee

- 3:1 data reduction
- A written guarantee for your business

How Can You Join?

- Contact Huawei sales or a Huawei-authorized distributor/channel partner
- Deploy all flash storage following Huawei's directions^[3]

* Huawei reserves the right of final interpretation for this program.

Huawei Accomplishes 3:1 Data Reduction Rate in All-Flash Offerings

By Wang Zhen from Huawei

Demands on data processing performance and storage efficiency continue to escalate as data mushrooms in the digital age. IT is gradually transforming from being a cost center into a value center to enhance the competitiveness of enterprises, shorten time to market, and even reduce operating costs. Part of this transformation involves all flash, which is becoming a focus for storage vendors and attracting users.

All-flash attraction

All flash continues to gain market share and vendors are fighting for position. In 2012, EMC spent 430 million USD acquiring XtremIO. In 2017, HPE spent 1 billion USD on the Nimble Storage acquisition to enrich its all-flash product line. HDS, Dell EMC and many other manufacturers have also augmented their portfolios with proprietary flash offerings. The threshold is high and Huawei is taking center stage with the most comprehensive portfolio in flash.

Silicon is in and magnetic tape is moving out as all flash replaces the aging medium. And, the pace of that supersession is accelerating. Major manufacturers are all chasing after their piece of the pie.

The faster, the better

All flash presents a compelling case with its service attributes and improved efficiencies. I/O latency and sluggish performance in ERP systems, OLTP databases, and real-time performance analysis equate to time and money losses, not to mention degradation to user experience. The need for speed is intensifying as society and business are all becoming networked.

In this backdrop, which level of latency is acceptable? Is 1-second latency enough or must it be reduced to 10 ms, 1 ms, or even 0.5 ms? For CIOs, latency can never be low enough. Latency in each all-flash offering differs from vendor to vendor.

Most architectures are still geared to HDD with some optimized algorithms and modules undergoing SSD grafting. The vast majority of offering can only ensure 1-ms latency and that rises to 3 ms under heavy service loads or when deduplication and compression utilities are running. In contrast, native SSD is purpose-built to leverage the innate advantages of all flash layouts. Native SSD design entails much more than simply replacing HDD with all flash drives. The new architectures adopt algorithms specifically designed for SSD and can achieve system-wide latency of 0.5 ms with the right end-to-end optimizations. Enterprises are finally enjoying the long awaited sub-millisecond latency in their latency-sensitive applications.



Stable latency and speed must-haves for critical services

Users demand more than just speed. They also place heavy weight on service continuity. If business is not sustainable, then all the speed in the world won't make a difference. All flash is aimed at accelerating the entire stock of critical applications running on enterprise systems rather than on a single service. Increasingly more users are opting for active-active layouts to ensure high availability of mission-critical services in their high-performance lineup.

Which type of active-active design is most suitable for all flash?

Two active-active designs apply to the storage layer: solution with an external gateway and solution with arrays by using built-in active-active functions.

Solution with external gateways means adding more nodes in IT systems, increasing system latency by 1 ms to 1.6 ms. Moving into all flash generally requires storage latency be controlled to one millisecond or less. Having a gateway in the mix robs the layout of the performance advantage in SSD.

Rock-solid, lightning-fast storage arrays with built-in active-active functions are becoming the compelling choice.

Striving for the best performance and stability

All flash vendor offerings differ significantly in speed and stability.

Flash-accelerated traditional storage layouts can deliver stability and implement snapshot, active-active data protections, and other utilities. However, the layout optimizations are no match for native SSD design in terms of speed.

At the same time, native SSD design often lacks the needed stability and predictability in latency when replication utilities, active-active protections, and other features are running. Huawei OceanStor Dorado V3 sets itself apart from the pack because it does not require any external gateways. Latency is kept to within 1 ms when systems are in active-active mode, achieving rock-solid stability and lightning-fast performance. Enterprises no longer have to choose between speed and stability!

Flash speed combined with the reliability of Huawei storage technology accelerates and solidifies the matrix, raising the bar in the all flash arena.

Safeguarding performance while making data reduction claims

Users recognize the performance and reliability of all flash. However, compared with that of traditional storage, the higher cost is the last hurdle keeping the accelerated medium from fully entering the data center.

How to lower TCO? Obviously, deduplication and compression are two most common technologies for enterprises to reduce the purchase amount of raw capacity, hence reducing the IT footprint and improving green credentials.

Various vendors offer guarantee rates from 2:1 to 4:1, and free additional capacity are provided for users if the data reduction levels are not reached. HDS claimed of 2:1 data reduction rate on its all flash systems back in 2016. Huawei OceanStor Dorado V3 guarantees 3:1 data reduction rate in its long list of performance and stability commitments.

However, storage performance in other vendor offerings takes a hit when deduplication and compression are enabled. This point is often overlooked but critical to the purchase decision. Vendors do not offer compensation for these degradations in performance, defeating the original intent in the all flash procurement. This is precisely why some enterprises are unable to see the full value of their SSD procurements because they forget to check the loss in performance when deduplication and compression are enabled.

When deduplication and compression are enabled, all-flash layouts optimized with SSDs are no match for the native SSD designs in performance. Therefore, Huawei OceanStor Dorado V3 can make solid claims of delivering 3:1 data reduction rate with zero impact on performance! ⚡




Chips Purpose-Built for All-Flash Storage: the Core of Competitiveness

By Wang Jiaxin from Huawei

Data is the core of intelligent management. Intelligent management, scheduling, mining, and analysis of data help applications deliver personalized experience and companies implement precision control over the design, production and logistics process to reduce management costs. Data is defining new business models and creating new commercial opportunities.

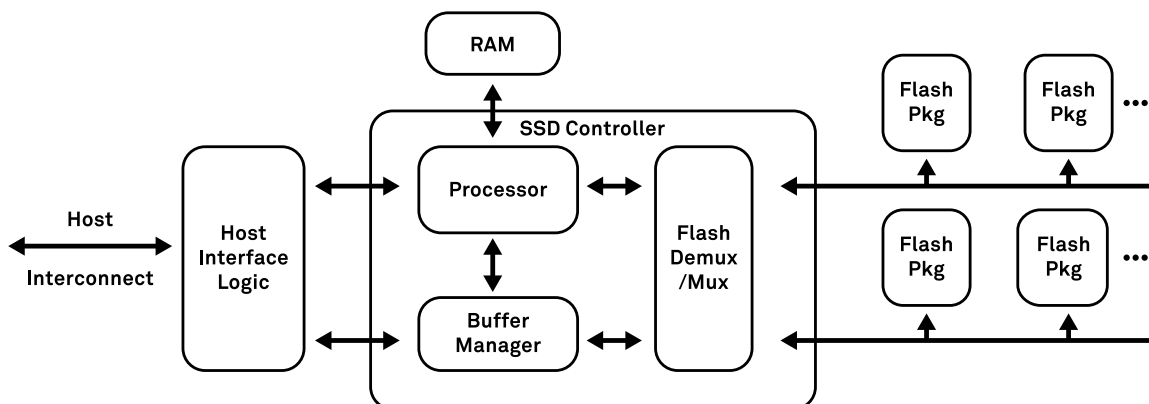
The digital and intelligent transformation demands high-performance infrastructures, of which all-flash storage is an essential part. Born with high performance, all-flash storage is the new engine for mission-critical services. By vertically integrating its capabilities in chips, networks, and management, Huawei releases the OceanStor Dorado all-flash storage series, which adopts an intelligent multi-protocol processing chip, intelligent SSD controller chip, and intelligent device management chip to build an end-to-end service acceleration platform, delivering triple performance and maximizing resource utilization.





Intelligent SSD controller chip, accelerating data reads and writes

Since data is stored on SSDs, the performance and stability of SSDs are critical. An SSD comprises of a control unit (SSD controller + DRAM) and a NAND flash storage unit. The control unit is responsible for data reads and writes. The Flash Translation Layer (FTL) is used to store the mapping between the user LBA and the physical page in the SSD. When the storage controller reads data from the SSD, it provides an LBA address. The control software in the SSD flash finds the physical address corresponding to the LBA address, reads the data from the flash, and returns the data to the host. In the case of writing, the control software writes data and then updates the FTL mapping table.




To achieve the ultimate in storage speed, Huawei leverages an innovative SSD controller chip to accelerate data reads and writes. The FTL algorithm is moved from the control software layer to the SSD controller chip, so that all FTL reads and writes are performed by the chip, significantly reducing the number of software interactions and hence the I/O response latency.

An example can help you better understand the reason. A computer using Windows 95 needs to load a long string of codes after being powered on, which takes two to three minutes. After switching to Windows 2010, the computer has a stronger CPU to load the codes. We can enter the startup interface instantaneously. According to tests by Huawei's performance and interoperability lab, the read latency of Huawei SSDs is as low as 80 μ s in low-load scenarios, which is only 60% of the same type of SSDs in the industry. The performance is two times higher than that of competitors.



Intelligens SSD controller chip with FTL algorithm to accelerate data read and write, read latency much lower than than of competitors

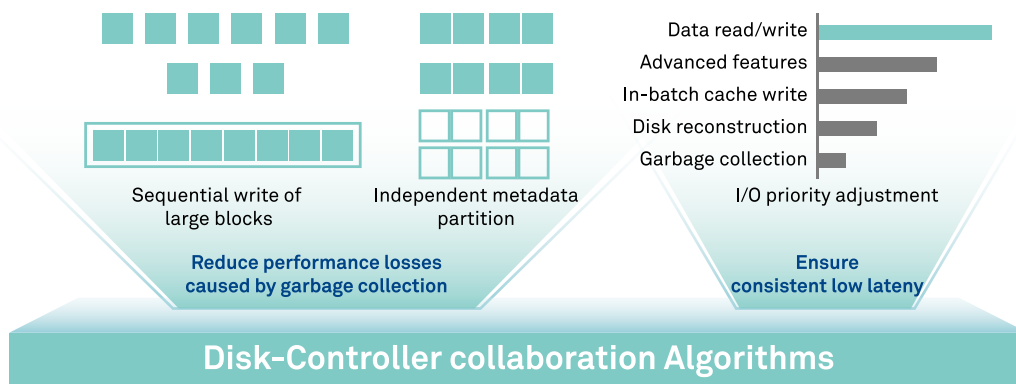


However, that's not all of it. To achieve end-to-end acceleration, Huawei develops FlashLink™ technology by combining proprietary SSD controller chip, SSDs, NVMe architecture, and a storage operating system designed for flash from the ground up. Huawei's Dorado storage can deliver 3x higher performance after enabling of value-added features such as deduplication, compression, and snapshot, while maintaining 0.5 ms latency, eliminating performance shortage issues in peak hours.

Based on the flash-oriented storage operating system, Huawei Storage has developed an innovative disk-controller collaboration algorithm. As is known to all, garbage collection is one of the main factors that affects the performance of SSDs.

The innovative disk-controller collaboration algorithm enables a storage controller to learn about the data layout in SSDs in real time and adjust the data layout to ensure consistent data

layout between the storage controller and SSDs. Data in the controller is written into SSDs in the format required by the SSDs, avoiding subsequent data migration and garbage collection and ensuring consistently high performance of the flash storage system. This is the basic principle of the disk-controller collaboration algorithm. The detailed technologies involved include the sequential write of large blocks, independent partitioning of metadata, and end-to-end I/O priority adjustment.



Based on the intelligent SSD controller chip and FlashLink™ technology, Huawei's lightning-fast and rock-solid OceanStor Dorado all-flash storage can deliver high performance while keeping the response time within 0.5 ms.



Intelligent multi-protocol interface chip, accelerating data reads and writes of front-end network interfaces

Front-end interface modules are essential to a storage system. Application data is transmitted from servers to storage arrays through these modules. Currently, mainstream front-end interface modules include 8G/16G/32G FC, 1/10/25/40/100 GE, and 10G FCoE. Each front-end interface module supports only one protocol, which is a kind of waste of resources. To improve the efficiency, Huawei develops a multi-protocol interface chip that integrates GE/10GE/FC/FCoE protocol interfaces. Customers can use one interface chip to transmit data carried over IP and FC protocols. On a 10GE or 8/16G FC network, only optical module components need to be replaced, instead of the entire module. The flexible conversion between protocols greatly improves network flexibility and reduces the network construction and maintenance costs of data centers.

More importantly, the internal hardware logic module of the intelligent multi-protocol interface chip supports many protocol stack functions such as checksum and FC. The processing flow, logic, and functions are switched over from the CPU software to the chip, helping ensuring high I/O concurrences and low latency of storage services.

In other words, the intelligent multi-protocol interface chip offloads the checksum and FC workloads from the CPU, which can improve network processing performance, release the x86 CPU processor resources, accelerate network access and data exchange, and improve the overall storage performance. Huawei's performance tests show that when configured with the same front-end interface module (16G FC) and test model (7:3 read/write, 8K I/O blocks), OceanStor Dorado can deliver 3x performance over competitors' products.

Finally, let's look at the FCoE protocol. It supports both FC and Ethernet-based transmission, thereby protecting FC-based software and hardware investments. The intelligent multi-protocol interface chip integrates the functions of the Ethernet NIC and FC network HBA card, that is, one interface can support two types of network I/Os, eliminating extra cables and switches and simplifying network management. Previously, the FCoE protocol runs on the host CPU, which consumes plenty of CPU resources. The host CPU cannot parse other network protocols at the same time, failing to provide the required high performance. OceanStor Dorado adopts the intelligent multi-protocol interface chip to process and parse the FCoE protocol, offloading the host CPU and improving the overall network performance and server availability.

Intelligent device management chip, accelerating fault management and rectification

Whether faults can be quickly identified and eliminated is a core indicator of the reliability of IT devices. The intelligent device management chip is the "heart" of OceanStor Dorado. It is built in with fault diagnosis and fault prewarning libraries to improve fault diagnosis accuracy. Fast diagnosis is the prerequisite for fast recovery. The intelligent management chip provides a management computing capability of 2000 Dhrystone Millions of Instructions Per Second (DMIPS), which is five times higher than that of competitors. If a controller, front-end interface module, or management module fails, the switchover can be performed in seconds. No data is lost, services are not interrupted, and users are completely unaware of the switchover.

The intelligent device management chip excels in energy saving. It monitors the health, power consumption, and temperature of each module in a refined manner, and combines the static and dynamic power consumption control technologies to adjust the heat dissipation of the system and reduce the chip temperature, thereby controlling static power consumption of the chip. On the other hand, the embedded Dynamic Energy Management Technology (DEMT) enables the chip to monitor and analyze the status of applications on the system, and to identify and disable the clock and voltage of idle modules. Hence the dynamic power consumption of the chip can be reduced. According to field tests, under the same CPU workload, the power efficiency ratio of the entire system can be improved by 16%. Besides, power consumption is obviously reduced.



Whether for new applications such as artificial intelligence, Big Data, autonomous driving, and block chain, or traditional healthcare, manufacturing, and finance industries, innovative chip technologies have become the main driving force of enterprises. Huawei Storage continuously invests in innovation of chips and technologies to help enterprises cope with digital flood and accelerate digital transformation. The intelligent SSD controller chip introduces optimizations purpose-built for the flash media, exploring full potential of flash to improve user experience. The multi-protocol intelligent processing chip simplifies networking, reduces management costs, and offloads network protocols to accelerate data reads/writes. The intelligent device management chip accelerates fault locating, implements failover in seconds, and excels in energy saving, helping users lease IT devices to reduce costs and explore new business opportunities.

Through technological innovation and vertical optimizations of software, hardware, and chips, Huawei Storage is committed to eliminating the gap caused by unbalanced CPU, media, and network development, providing faster, better, and more cost-effective products and solutions. Explore more about Huawei Storage, and win business success together with Huawei. ⚡



OceanStor OS Enables Fast and Stable All-flash Storage

By Wang Jiaxin from Huawei

However, this is not the case. Not all products can fully offer all the benefits it claims to provide. For example, nearly all cars have the scale of 200 km/h on its dashboard, but few can actually operate freely when driving at this speed. Without solid chassis, precise steering, and good suspension, a car is taking risks when it runs at fast speeds. For storage, these components are all present in all-flash. The all-flash operating system can easily operate the all-flash super sports car (storage), making it run steadily under extreme road conditions.

IT equipment is configured with a large number of integrated circuits and PBC boards. The equipment design and production require sophisticated technologies and strict process control. Cutting-edge technology, sufficient capital, and decades of experience are all crucial for manufacturing reliable IT hardware devices. Storage is one of the most important aspects for all enterprises. It adopts the integrated design of hardware and software. All hardware needs to work with software to maximize its capability, which further enhances difficulties in storage engineering.




Hardware is the carrier of storage, and software is the soul of storage. Neither of them can survive independently. It requires great efforts to manufacture hardware, and even greater efforts for top-quality software. Manufacturing the operating system, considered the core of software, is the most demanding and time-consuming part. A common storage operating system contains at least 17 million lines of codes, requires the diligent work of 3000 first-class engineers, and at least five years of development and tests before it is completed. Over 65% of all development tests are made to test the fault-tolerance designs to ensure the high availability of the system. Only a handful of vendors in the industry have achieved such a feat. Most storage startup vendors' are incompetent in designing a complete storage operating system. Therefore, a storage leader's strength lies in an independent storage operating system.

With years of industry experience accumulated from varying industries, Huawei has developed its own state of the art storage operating system – OceanStor OS. For over eight years, this OS has been running on more than 50,000 sets of devices on the live network of enterprises, ensuring the stability and reliability of enterprises' key data. By focusing on the most popular SSD technology in the industry, Huawei is one of the vendors holding a large number of SSD-level patents. OceanStor OS is based on a stable software platform and many SSD-level patents. Recently, Huawei launched a new OceanStor OS version dedicated to flash storage to fully utilize the capabilities of SSDs and flash systems.

Another aspect of Huawei's cutting edge OceanStor products is the use of SSDs and garbage collection. For SSD read/write, SSD flash cells can be re-written only after being erased. Generally, SSDs write data by page (size: 16 KB) and erase data by block (size: 8 MB). Each block consists of multiple pages comprising valid pages, invalid pages, and empty pages. In an SSD, if pages within a logical location are newly written, the pages of original physical location become invalid. To avoid erasing valid pages, valid pages in a block must be migrated to another space. When all valid pages in the block are converted to invalid pages, the entire block can be erased at one time. The process of migrating valid data is known as garbage collection. Garbage collection improves the space re-utilization of an SSD, but each migration undermines the performance of the storage system. In addition, larger amounts of migrated valid data and shorter periods lasting from when each page is written to the SSD to when the page becomes invalid imposes greater impacts on the system performance.

Innovative disk-control collaboration lays a solid foundation for high performance

The key to effectively controlling garbage collection is to maximize the performance of SSDs and flash storage systems. Powered by proprietary SSDs and the flash operating system, OceanStor OS, Huawei's OceanStor Dorado adopts an innovative disk-controller collaboration technology. By optimizing internal software algorithms, Dorado enables storage controllers to detect the data layouts in SSDs in real time and make adjustments accordingly. In this way, the data layout in the storage controller and SSDs can maintain consistent with each other. This prevents data migration and garbage collection after the data is written to SSDs, ensuring consistently high performance for flash storage systems.



Huawei-patented global wear/anti-wear leveling technology prolongs SSD service life

Different from HDDs, SSDs can only withstand a limited number of read and write operations. The service life of an SSD is an inverse proportion to the amount of data written to the SSD. Therefore, an all-flash storage system requires load balancing between multiple SSDs to prevent over-used disks from failing.

OceanStor OS adopts a global wear leveling technology. Based on the collaboration between the controller software and SSD drives, all data is evenly distributed to multiple SSDs to share the service pressure. In addition, OceanStor OS periodically queries the SSD controller for the disk wear degree, and uses the wear degree as the basis for space allocation, thereby ensuring the reliability of the entire system.




Global wear leveling and global anti-wear leveling

However, when an SSD disk enters the end of its service life, for example, the disk wear degree reaches 80% or higher, multiple disks may be faulty at the same time, resulting in data loss. Huawei has developed a patented global anti-wear leveling technology to prevent SSDs from being faulty in batches.


OceanStor OS selects the most severely worn SSD and writes new data onto it as long as it has idle space. This reduces that SSD's life faster and users are advised to replace it sooner, avoiding potential simultaneous failures and service interruptions. This technology is perfect for scenarios requiring IT devices replacement.

Purchasing a batch of new IT devices to replace the existing system is not an overnight task. It takes a long time to complete the procurement process approval, new device deployment, legacy service migration, and user acceptance. This is a very delicate process for users' core services and data, so the effects are major if all SSDs break down during this period. Global anti-wear leveling technology can gradually replace the faulty SSDs and prolong the service life of legacy devices on the live network until official rollout of new devices.



Global inline deduplication and compression improve efficiency and service life SSD service life

Based on its analysis on HDDs and SSDs, IDC predicts that the price of an SSD will drop from more than 3 times in 2016 to 2 times at the end of 2018 over a 10K SAS disk. To accelerate the commercial use of SSDs, the industry uses inline deduplication and compression technologies to reduce the data volume before data is written into SSDs and minimize the amount of data that actually moves into SSDs, without affecting user experience.




IDC predicts the price comparison between SSDs and HDDs





Huawei OceanStor OS developed global inline deduplication and compression technologies. To obtain the best data reduction ratio, different types of services require different deduplication and compression granularities. The weak hash algorithm and byte-by-byte comparison are used for deduplication. After data is divided into data blocks by service type, deduplication is started. Then the system uses the weak hash algorithm to calculate the fingerprints of the data blocks and compares the fingerprints with the existing ones. If the fingerprint of a data block exists in the system, the system does not write the data block but only increases the fingerprint count. If a fingerprint is unique, the system adds it to the fingerprint table and writes the data block to SSDs. Deduplication is performed in real time, not after data has been written to SSDs.



Working principle of deduplication

The byte-by-byte comparison technology for addendum compares deduplicated data by byte to prevent the fingerprint hash conflict and ensure 100% data reliability.

The compression algorithm is a compute-intensive program. Inline compression consumes significant CPU resources, affecting end-to-end performance of the system. Industry peers often use the open source compression algorithm with high performance and low compression rate, such as LZ4, LZ0, and Snappy. Huawei's OceanStor OS is optimized based on the open-source LZ4 compression algorithm. The unit for storing compressed data is 1 KB, which doubles the compression efficiency and saves the storage space for compressed data.

Most all-flash storage vendors in the industry claim that their operating systems support inline deduplication and compression technologies. However, there are technological differences between OceanStor Dorado V3 all-flash storage and other products. According to the actual project test, under the data model of dual-controller, 100% random 8 KB I/O blocks and 7:3 mixed read/write, by stimulating the most common database scenarios, OceanStor Dorado with inline deduplication and compression enabled can maintain a 0.5 ms low latency and high performance. In the same test environment, the performance of OceanStor Dorado is twice that of EMC VMAX 950F or HPE StorServ 20850.



Based on the efficient deduplication and compression algorithms, Huawei promises a 3:1 data reduction ratio to customers who purchase OceanStor Dorado V3 all-flash storage series, helping users save investment and achieve higher return on investment (ROI). What's more, if the guaranteed ratio is not met, Huawei is liable for providing additional storage capacity or exempting price on the capacity in future procurement. This function improves the storage system utilization, reduces user's effective capacity cost per GB, maximizes space occupation, and reduces power consumption, lowers air conditioning and maintenance cost, helping to the end-to-end OPEX. Less data written into SSDs reduces the wear of SSDs, and prolongs the service life of SSDs and the storage system.

Comprehensive data protection and efficient software maximize the advantages of all-flash high performance

Thanks to years of expertise accumulation, Huawei's new-generation OceanStor OS applies features of data protection software, including clone, remote replication, active-active, and 3DC, and inherits features from high efficiency software, including thin provisioning, QoS, and heterogeneous Virtualization. In addition, Huawei has fully optimized SSDs and developed more competitive features, such as lossless snapshot, RAID-TP that tolerates three-disk failure, and non-disruptive data migration, staying ahead of competitors in the all-flash era.

In addition to supporting gateway-free active-active mode and ensuring critical services with zero RPO and a close-to-zero RTO, OceanStor OS is ahead of its peers in terms of performance, reliability, and efficiency.

- For performance, OceanStor OS is based on the high performance of all-flash storage and the optimization of internal lock mechanism in the active-active software, enabling Dorado all-flash storage to reach 200,000 IOPS at 1 ms latency, topping in the industry.
- For reliability, OceanStor OS supports upgrading an active-active solution to a geo-redundant solution to ensure 99.9999% high availability for critical services.
- For efficiency, OceanStor OS enables Huawei's all-flash storage to support HyperMetro for both SAN and NAS. The integration of SAN and NAS changes the traditional active-active solution where extra gateways are added on arrays to provide active-active SAN and NAS services, decreasing the number of devices by more than two and reducing deployment complexity and costs. In addition, the advantages of SAN and NAS parallel architectures are fully utilized to improve service performance.



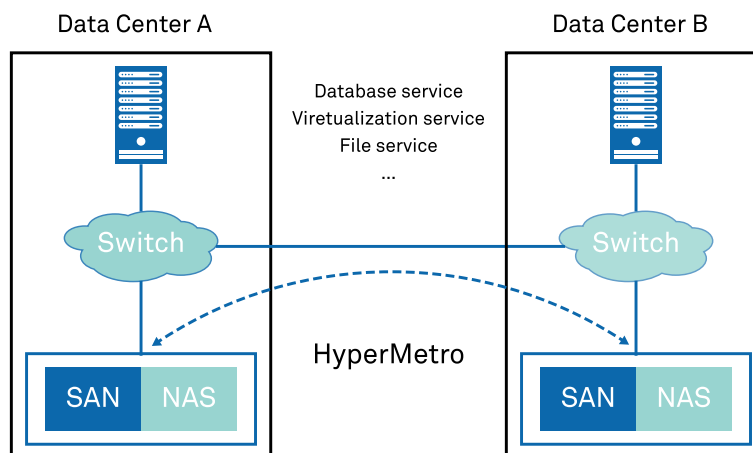
In addition, network integration changes the coexistence of multiple networks, such as Fibre Channel and IP networks. Between active-active sites, Fibre Channel or IP is used to deploy data replication networks, configure networks, and heartbeat networks in a unified manner, reducing deployment costs. In the traditional storage active-active solution, two arbitration mechanisms work separately, leading to inconsistent arbitration results of SAN and NAS services in the event a network fault occurs between

sites. OceanStor OS adopts unified arbitration to ensure that in all instances SAN and NAS services are deployed at the same site and share the same number of resources.

Currently, one of the industry's largest active-active clusters supports only eight nodes, and it cannot meet storage performance requirements in large-scale deployment scenarios. Huawei's integrated SAN + NAS active-active solution inherits the scale-out architecture of common clusters

and supports a maximum of 32 nodes in active-active mode, meeting customers' fast-growing requirements for storage performance.

Yahoo, the Japan's largest Internet company, uses Huawei's integrated SAN+NAS all-flash active-active solution to ensure real-time synchronization of inventory data in online stores, maintaining consistency between online and offline inventory. This solution completes fault switchovers (180 km far between two sites) within seconds, five times faster than a solution provided by NetApp.



Integrated SAN+NAS active-active solution

Manufacturing all-flash operating system is not an easy task. The preceding description of OceanStor OS is just the tip of the iceberg. When enterprise users actually put the complete set of OceanStor OS in use, they can find out the unique advantages by themselves. Ultimately Huawei is nothing without its customer base. Thankfully, after 20 years of successful deployment and intuitive research, Huawei has become of the world's leading brands with a customer base ranging across multiple countries and industries, and has since become a brand that you can trust. ⚡



Huawei's High-Performance All-Flash Storage

By Wang Jiaxin from Huawei

In the digital economy era, data is growing explosively, transforming people's understanding of data. Data used to serve business operations but has now become one of the driving forces enabling digital transformation for enterprises. As enterprises' data centers carry an increasing number of businesses, data is used and mobilizes more and more frequently, requiring enterprises to find data systems that provide lower latency and higher service levels.

According to a survey on hundreds of data system users, about 87% of system performance problems occur in the interaction between the storage subsystem and the application database. That is to say, the response latency and concurrent access traffic of the storage subsystem determine those of the application system. High latency and small concurrent traffic of a storage subsystem has become the performance bottleneck of the entire system, which is an infuriating reality for many enterprises.

Let's take a quick look at the history of latency requirements of storage systems needed for business expansion. In the HDD era, enterprise backup and web

disk applications require relatively low storage latency, and 10 ms can help users easily cope with application demands. These days most virtual desktops require a latency of 5 ms, and with the emergence of cloud and virtualization technologies, virtual desktop offices have become the mainstream in large enterprises. Big data has changed business models, and the surge in data volume has a huge impact on enterprises' customer relationship management (CRM) and enterprise resource planning (ERP) systems. The latency of 0.5 ms has become the ultimate goal of enterprises to ensure quality services.

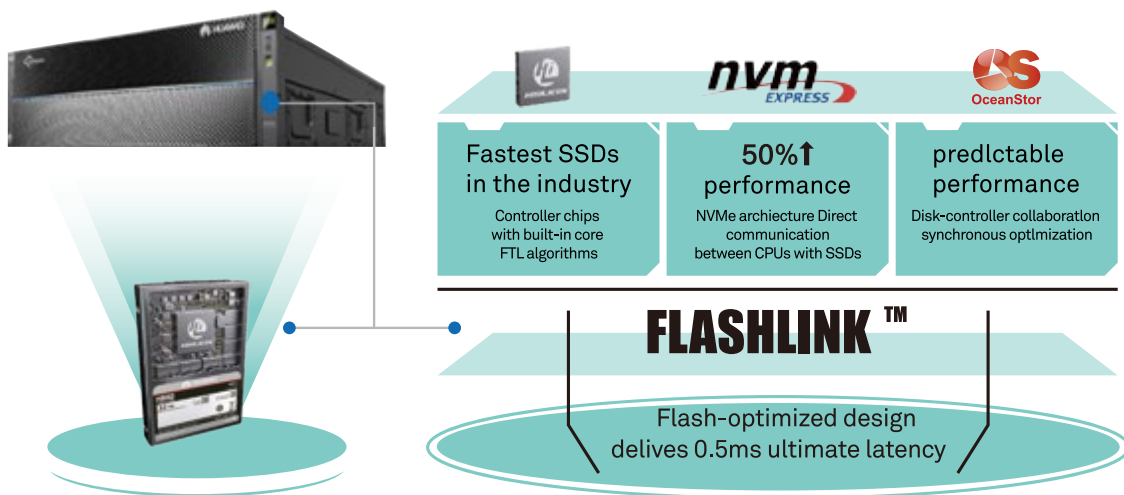
The 0.5 ms latency is a predicted value that can operate under real-world heavy-workload conditions. It is neither a manipulated peak number produced in carefully controlled test parameters nor a number produced when systems are under zero workload. Storage systems are required to maintain 0.5 ms predictable latency even during service peaks so to provide users with a consistent experience.



Innovative FlashLink Technology

After continuously accumulating technical experience over the past 20 years, Huawei proudly launched the lightning-fast and rock-solid OceanStor Dorado V3 all flash storage in 2016. Still a player in the field today, it delivers the industry-leading performance powered by the innovative FlashLink technology, and the high performance is maintained from three aspects: chip, architecture, and operating system.

OceanStor Dorado V3 adopts three intelligent chips to achieve end-to-end service acceleration and provides a performance of 45% higher than SAS all-flash storage. Huawei is a groundbreaking telecommunication provider in that it continuously keeps up with the latest architectural technology trends and even develops its own technologies. For example, OceanStor Dorado V3 is one of the first all flash storage systems to use NVMe in commercial use. Further, OceanStor Dorado V3 adopts a brand-new SSD-optimized design and disk-controller collaboration technology to enable storage controllers to detect data layouts in SSDs in real time and synchronize data in controllers and SSDs. This helps reduce performance losses caused by garbage collection and ensures rapid response to data read and write I/Os. While these are just the highlights of OceanStor Dorado V3's abundant back catalog, together they help maintain a predictable latency of 0.5 ms even under heavy workloads. The secret for such advancements is FlashLink, helping OceanStor Dorado to improve service performance by three times in comparison with traditional storage.



FlashLink overview

Innovative Disk-Controller Collaboration Ensures Predictable High Performance

The flash storage cells in an SSD can be re-written only after being erased. Generally, the basic writing unit of an SSD is a 16 KB page, and the basic erasing unit is an 8 MB block. To avoid erasing valid pages, valid pages in a block need to be migrated to another space for storage. The block space of valid pages is converted into invalid page space, after which the block can be erased at a time. The process of migrating valid pages is known as garbage collection.

Garbage collection improves the space re-utilization of an SSD, but each migration undermines the performance of the storage system. Large amounts of migrated valid data and shorter periods lasting from when each page is written to the SSD to when the page becomes invalid imposes greater impacts on the system performance.

To ensure that they are maximizing the performance of SSDs and flash storage systems, enterprises must effectively control garbage collection. Powered by proprietary SSDs and the flash operating system, OceanStor OS, Huawei's OceanStor Dorado adopts an innovative disk-controller collaboration technology to prevent a drop in performance caused by garbage collection. By optimizing internal software algorithms, OceanStor Dorado enables storage controllers to detect the data layouts in SSDs in real time and make adjustments accordingly. This helps prevent data migration after being written to SSDs and garbage collection, ensuring predictable high performance for flash storage systems.




Large Block Sequential Writes Reduce the Frequency of Garbage Collection

Take real-time ridesharing as an example. Such services allow multiple users with similar routes to share a ride, which helps save travel fees of each user and reduce overall energy consumption.

Using this analogy, controllers of OceanStor Dorado V3 detect the data layouts in SSDs in real time and aggregate data blocks to be written to SSDs in the controller cache. The formats of the data blocks are unified, and then into a larger data block that is written to SSDs at a time to improve the overall system performance. Detailed benefits include:

- The large block sequential write technology controls the frequency at which random small blocks (I/O) are written to SSDs for multiple times, which makes full use of the bandwidth of back-end SAS.
- RAID write penalty (extra reads and writes required during verification) has long been one of the factors hindering the performance in a storage system using RAID protection. OceanStor Dorado V3 writes data into SSDs once after data aggregation, effectively reducing the number of disk writes and the number of extra read and write requests required for verification. This feature ensures a predictable system performance when RAID 5, RAID 6, and RAID-TP are used. Even in such an unlikely scenario as three disks failing concurrently, RAID-TP, a unique technology of OceanStor Dorado V3, can still ensure that services are unaffected to.
- OceanStor Dorado supports global garbage collection. It monitors the system pressure in real time and controls the frequency of garbage collection in disks, mitigating the impact of garbage collection on system performance.



Large block sequential write technology

Technical principles

- Controllers detect the data layouts within proprietary SSDs
- In the controller memory, multiple discrete I/Os are aggregated into a sequential large data block.
- Large data blocks are written into SSDs in sequence.


Controllers automatically detect the data layouts within SSDs.

Discrete write of multiple small I/O blocks → Sequential write of one large block
 Maximizes SAS bandwidth
 Reduces garbage collection

Independent Metadata Partition Controls the Frequency of Garbage Collection

In a storage system, the frequency of updating user data differs from that of updating metadata in that metadata is often updated more frequently. In scenarios where metadata and user data are written into the same partition on a disk, more garbage collections are required than in user data-only scenarios. This is because when pages of metadata become invalid, pages of the user data may remain valid. Therefore, a large amount of user data needs to be migrated during garbage collection, resulting in excessively large write amplification on the disk as well as shorter service life and lower performance of the SSD.

OceanStor Dorado V3 all flash storage uses independent metadata partitioning. It frequently writes updates metadata to a partition and infrequently updates to a different partition in the storage system and SSD. This reduces migration of user data blocks when upgrading metadata, mitigating the impact of garbage collection on system performance. In simple terms, the independent metadata partitioning technology controls the number of garbage collections, ensuring a predictable high performance of the storage system.




Independent metadata partitioning

Prioritizing Data Read and Write I/Os

Large financial enterprises, such as big banks, often set a special counter for VIP customers to separate them from the regular customers. Similarly, when VIP counters are busy, VIP customers can jump the queue to common counters. This kind of model ensures that VIP customers enjoy the most effective services at a quicker speed than regular customers.

This kind of model is also adopted by OceanStor Dorado, which introduces an I/O priority scheduling mechanism to ensure predictably low latency of service requests. OceanStor Dorado prioritizes data read/write requests with IT resources including CPUs, memories, and concurrent disk access traffic in storage systems. Other requests such as data reconstruction, asynchronous cache flushing, and background requests within the system should compromise in the case of resource contention.

The OceanStor Dorado priority adjustment is performed synchronously in the storage controller and the SSD to ensure that the data read/write requests enjoy the top priority at all times. Other types of data I/O requests are suspended when the read/write request arrives, and resume after the read/write operation completes, guaranteeing an optimal response latency of data read and write in the storage system.



Controllers automatically detect the data layouts within SSDs.

All types of I/O are responded in time order with the same - level priority.



The read and write I/Os are always responded to with the highest priority. **Other types of I/Os pause and restart when read and write I/Os finished, ensuring lowest latency for read and write I/Os.**

I/O priority adjustment

3x Higher Performance of All Flash Storage

In the digital transformation era, replacing traditional storage with all flash opens a new chapter. Individuals and enterprises no longer need to wait for response of applications in life and work.


For the financial industry, especially those in the securities field providing frequent real-time transactions, time is money. Take Hundsun in China as an example. Before it cooperated with Huawei, its traditional IT architecture supported only 60,000 transactions per second (TPS) at business peaks, insufficient of the desired 100,000 TPS. OceanStor Dorado V3 can help enable Hundsun to process 150,000 transactions per second, and can scale for future business expansion.

In the manufacturing industry, batch processing capabilities of data warehouses are the basis for ERP business analysis. For example, BYD, the largest new energy vehicle manufacturer in China, needed at least 3.5 hours a day to batch process business requests. When the business volume was large, the system spent too long processing these requests, which caused great pressure on those at decision-making levels regarding service the next day. However once they implemented Huawei's OceanStor Dorado V3, the system batch processing takes only 1 hour and 12 minutes, reserving sufficient time for final decision making.

In the medical industry, the hospital information system

(HIS) is the core for hospital service management. It connects to multiple processes, such as registration, diagnosis, treatment, charging, and medication. Take a well-known tertiary hospital in China as an example. In its traditional IT architecture, each patient spent 3 seconds for registration on average and should wait in three to six queues during diagnosis, lasting for at least one hour. After Huawei OceanStor Dorado is used, the registration time of each patient takes only 0.5 seconds, improving diagnosis efficiency and improving the doctor-patient relationships.

Storage leasing is one of the main services of carriers and independent service providers (ISPs). Taking ACESI Group, the largest ISP in eastern France as an example. The speed of batch VM deployment was a major concern of ISPs because it was related to new business rollout. It took 30 minutes to deploy 100 VMs using traditional storage, making rapid service development impossible. OceanStor Dorado V3 shortened the deployment time of 100 VMs to just 10 minutes. In addition, Huawei enabled ACESI to develop new platinum leasing services based on the high-performance Dorado all flash storage, enhancing the overall competitiveness ACESI. This makes ACESI stay ahead of competitors in the industry.

In the future, Huawei's OceanStor Dorado V3 all flash storage will benefit more customers. 

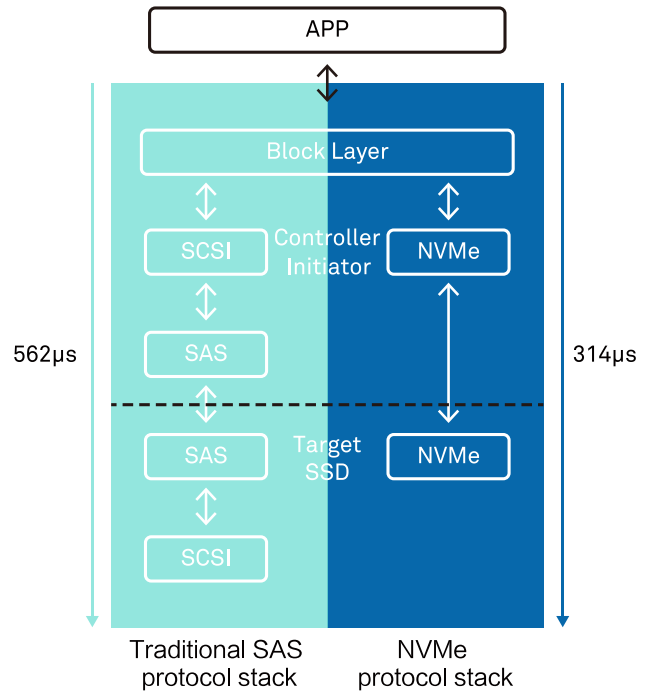
Lightning-Fast All-Flash Storage Ushers in an NVMe Era

By Wang Jiaxin from Huawei

Speed is the ultimate weapon to survive and thrive in fierce competition. Today, NVMe is making the fastest speed possible to help you achieve your goals.

Designed for HDDs, traditional SAS protocol hinders SSD performance due to complex system architectures, excessive numbers of protocol to be parsed, and limited queue concurrences. NVM Express formulated NVMe protocol standards, and replaced complex protocol layers such as I/O Scheduler and SCSI in the SAS system with lightweight NVMe protocol. As a result, NVMe is a quicker, smarter, and more intuitive option for enterprises, evident with its superb performance in all-flash arrays (AFAs).

Huawei has extensively researched the NVMe protocols, covering only end-end development of NVMe flash controllers, NVMe all-flash OSs, and NVMe SSDs in the industry. Besides, powered by NVMe architecture with disk-controller collaboration, NVMe all-flash storage eventually ensures a stable latency of 0.5 ms.



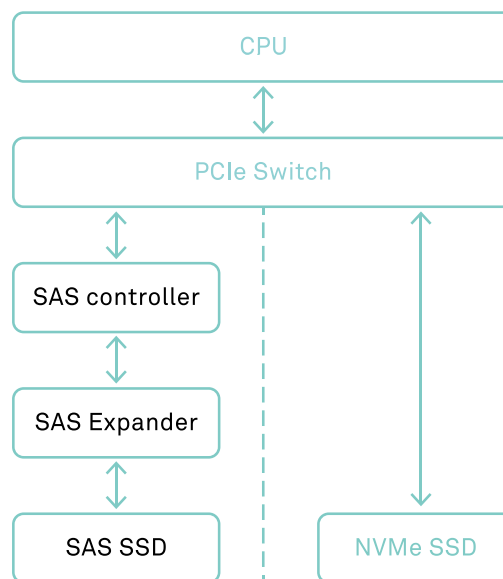
Architecture: SAS vs. NVMe



First, at the transmission layer, SAS-based AFAs deliver I/Os from CPUs to SSDs through the following paths:

Step 1: I/Os are transmitted from CPUs to SAS chips through the PCIe links and switches.

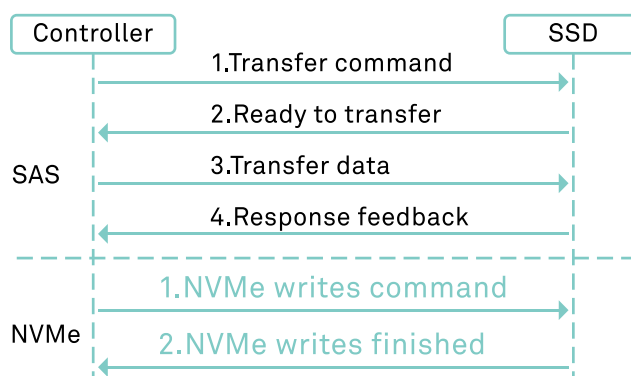
Step 2: I/Os are converted into SAS packets before arriving at SSDs through the SAS switching chips.



Transmission path: SAS vs. NVMe

For NVMe-based AFAs, I/Os are transmitted from CPUs to SSDs through the PCIe links and switches. CPUs of the NVMe-based AFAs directly communicate with NVMe SSDs over a shorter transmission path, resulting in higher transmission efficiency and less transmission latency.

Second, at the software protocol parsing layer, SAS- and NVMe-based AFAs differ greatly in protocol interaction technologies for data writes. A complete data write request requires 4 protocol interactions in SCSI protocol (connected through SAS back end); however, the NVMe protocol requires only 2 protocol interactions to complete a write.



Protocol parsing: SAS vs. NVMe

Third, at the protocol encapsulation layer, when SAS protocol stacks are used, I/O requests are sent from block devices and reach SSDs through SAS links after the encapsulation of two layers (SCSI and SAS protocols). However, when the NVMe protocol stacks are used, I/O requests, however, require the encapsulation of only one layer (NVMe protocol). Simplified NVMe protocol stacks lower encapsulation costs by 50%, thereby reducing CPU consumption and I/O transfer latency caused by each encapsulation.

Fourth, at the multi-queue concurrency layer, the SAS protocol supports single queue, while the NVMe protocol supports up to 64 K queues, with each queue

supporting a maximum of 64 K concurrent commands. Multi-queue NVMe protocol performs better thanks to higher concurrency processing and better cooperation with multiple channels and multiple dies in SSDs.

Fifth, at the lock mechanism layer, the single-queue SAS programming must be locked in the multi-core environment. Huawei designed an I/O scheduling mechanism for NVMe to completely cancel the disk-level mutex in original I/O paths and avoid I/O processing conflicts. This mechanism gives full play to the concurrent processing of multi-core processors, reduces software overheads, and improves back-end processing performance. In I/O scheduling, multiple threads work with multi-queues to achieve optimal performance.

Sixth, at the OS optimization layer, storage OSs designed for flash innovate the disk-controller coordination algorithms.



Be stable, be perfect

Because of the advantages mentioned previously, customers choosing or creating new storage protocol are likely to be attracted to the many benefits of NVMe protocol. Despite offering superb performance unseen so far in the industry, the NVMe protocol has new challenges for disk and system design.

- Non-native dual-PCIe ports. Dual-controller redundancy technology is used in enterprise storage to ensure system reliability. However, PCIe-based dual-port technology is a challenge for SSD designs. Currently, some vendors' SSDs do not use native dual-PCIe ports. For example, Intel's NVMe SSDs use dual-port PCIe ports implemented through the internal PCIe switches. Such technology shortcuts inevitably reduce system reliability.
- Traditional PCIe hot plug with poor performance. Commercial NVMe SSDs must support non-disruptive replacement. However, traditional PCIe hot plug, particularly surprise hot plug, causes poor user experience due to system exceptions and service interruption.
- Data reliability guarantee pressure. As large-capacity disks are used, large amounts of data demand for higher reliability.
- Difficult cross-site reliability design. Reliability design for single points of storage has already been challenged. The cross-site reliability is going to be the same.

As one of the leading brands to successfully develop stable AFAs, Huawei continues to showcase its capabilities, evident in over 1000 NVMe all-flash success stories.

First, at the interface layer, Huawei NVMe SSDs support native dual-port technology, with two independent PCIe 3.0 x2 links. This provides hardware basis for system recovery and exceptions, and ensures dual-controller redundancy, helping improve the system reliability for enterprises.

Second, at the hot plug layer, Huawei is an expert in the telecommunication industry, especially regarding comprehensive PCIe link management, PCIe troubleshooting technology, and hot plug technology. The PCIe driver is designed to support SSD removal anytime and anyway, providing end-to-end PCIe system reliability if a single disk is replaced or a fault occurs.

Third, at the data protection layer, Huawei's innovative RAID-TP software technology is based on the Erasure Code (EC) algorithm. Parity bits support 1-, 2-, 3-dimensions and can tolerate 1 to 3 simultaneous disk failures. This means that in the case of three disk failures, the system will not suffer from data loss or service interruption. Currently, only products from Huawei, NetApp, and Nimble can tolerate simultaneous failures of three disks; and other vendors (such as Dell EMC, HDS, and IBM) are unable to make this claim.

Although NetApp and Nimble can tolerate simultaneous failures of three disks, they both use traditional RAID architecture with fixed data disks and hot spare disks. For these companies, hot spare disk reconstruction for 1 TB of data takes 5 hours. OceanStor Dorado employs a global virtualization system able to reconstruct the data in a mere 30 minutes to fulfill requirements in ultra-large capacity profiles.

Fourth, at the cross-site data protection layer, Huawei NVMe all-flash storage provides comprehensive data protection technologies, such as snapshot, clone, and remote replication, to help customers build a hierarchical data protection solution from local or intra-city DCs to remote DCs. Huawei is a revolutionary member in implementing a gateway-free active-active solution in all-flash storage.



Get with the trend, prepare for the future


Traditional HDD storage has a latency of more than 10 ms due to the long seek time. However, SSDs reduce 50% of the storage system latency to about 5 ms by using electronic mapping tables.


Traditional storage controllers often provide the same OSs despite differences in HDDs disk form, meaning it is convenient to use even when the disk type changes. However, many of the HDDs and subsequent OSs have become redundant. That is why Huawei released AFAs, such as Pure Storage and OceanStor Dorado V3. Designed for SSDs, these AFAs effectively reduce the storage system latency to less than 1 ms.

In the future, faster storage media will undoubtedly be the next move for many enterprises looking to capitalize on innovative storage methods. Such is the benefits of using modern technologies. There is a large performance gap of 2 to 3 orders of magnitude between DRAMs and NAND SSDs, and even more between SCMs and DRAMs. All-flash storage using SCMs has the latency low to 250 μs, which ensures faster service response.

In addition to NVMe application in the local PCIe SSDs, NVM Express have also released the NVMe over Fabrics specification in June 2016. The new specification enables NVMe to be utilized over different fabric types, such as RDMA and FC, which can provide high-performance solutions for remote access to SSDs and remove resource sharing barriers among local SSDs.

Huawei uses NVMe over Fabrics to fully share SSD resources, and provides 32 Gbps FC and 100 GE full-IP networking design for front-end network connection, back-end disk enclosure connection, and scale-out controller interconnection. These functions decrease the storage latency and simplify storage network management by using one IP system to control the whole DC. This design avoids complex network protocol and planning, streamlines DC deployment, and reduces DC maintenance costs.





Huawei supports NVMe Over Fabric E2E solution

New SCM media are introduced to further improve system performance. With NVMe over Fabrics, SSD resources are fully shared, and front-end NVMe interfaces optimize hardware and software architectures. Ready to build the more competitive all-flash storage? Then think Huawei. ⚡



Huawei All-Flash Storage OceanStor Dorado V3 Wins Best of Show Award at Interop

The superior performance, stability, and efficiency of OceanStor Dorado V3, Huawei's all-flash array, make it the first all-flash array in the industry to meet the requirements on high performance and high reliability of mission-critical businesses. At Interop being held in Tokyo in 2017, it won over judges to grab the Best of Show award.




Huawei OceanStor Dorado V3

Meng Guangbin, President of the IT Storage Product Line at Huawei, said: "I am delighted our OceanStor Dorado V3 won the award. This is the second time since 2014 Huawei Storage has had this honor bestowed upon it, marking yet another nod to the viability of Huawei's IT products in fulfilling even the most exacting demands of the Japanese market. Huawei started to research flash technologies in 2005, having accumulated 13 years of technical expertise. OceanStor Dorado V3 is purpose built to handle heavily loaded mission-critical applications for organizations. Our products have some rather compelling features and our comprehensive portfolio is able to accelerate core applications while ensuring hallmarked reliability."

Huawei is the vendor with the end-to-end optimizations in flash from home growing its own operating systems, chips, and SSDs. Huawei OceanStor Dorado V3 adopt proprietary FlashLink technology that enables controllers to intelligently sense the data layout and make the global I/O priority adjustments. The on-board sequential write for large blocks, global garbage

collection, and other utilities help yield an impressive 4 million IOPS in performance while keeping latency stable at 0.5 ms. OceanStor Dorado is able to completely satisfy all processing capability requirements on critical services. In tandem deployments with the gateway-free Huawei HyperMetro active-active feature enabled, OceanStor Dorado V3 is able to achieve six-nines (99.9999%) availability. Huawei guarantees a 3:1 data reduction rate in many service scenarios when lossless inline deduplication and compression are enabled, reducing the TCO by over 50%.

Customers the world over have placed their trust in Huawei all-flash products. The long list includes the second-largest medical insurance in Switzerland (CSS), RPMI (British fund management), and BYD Automotive Group. Huawei is unleashing the power of flash to accelerate critical services for organizations big and small.

Huawei is fully committed to its Data on Demand strategy in building converged resource pools in cloud-based data center layouts and consolidating storage infrastructures to allow assets to be shared across different services. Find out how Huawei can empower you with the needed automations and many on-demand conveniences to effectuate an agile transformation to the cloud. 



JADE KEY: SCENARIOS & SOLUTIONS

ESG Test Report: Huawei OceanStor Dorado V3 Delivers Excellent
Performance for Mainstream All-Flash Applications

Huawei OceanStor Dorado V3 All-Flash Solutions, Built for Your Application Scenarios

Huawei's Business Integration Solution: A Revolutionary Optimal Data Migration Method

Can You Imagine That the Nasdaq
Market Shutdown Was Caused by a False Fire Alarm?


Moving Your Mission-Critical
Applications to Cloud Just Got Easier!

ESG Test Report: Huawei OceanStor Dorado V3 Delivers Excellent Performance for Mainstream All-Flash Applications

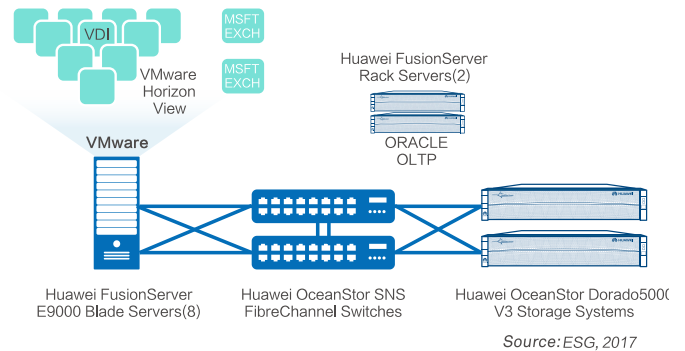
By Zhao Yanling from Huawei

The ESG lab is a well-known IT consulting organization based in the United States. It performs IT analyses, research, and verification, and is renowned for its independent, rigorous, and practical laboratory tests.

In September 2017, the ESG lab tested the performance, reliability, cost-effectiveness, and total cost of ownership (TCO) of Huawei OceanStor Dorado V3. The ESG set up a real service environment and tested the performance and reliability of Huawei all-flash storage for applications such as VMs, databases, and mailboxes.



The following figure shows the test environment used in the ESG lab. The VMware vSphere 6.0 virtualization infrastructure is deployed on 10 Huawei FusionServer RH series rack servers, with two dual-controller Huawei OceanStor Dorado5000 V3 storage systems connected to the back-end.



Test environment used in the ESG lab

According to the ESG lab test data, Huawei Dorado V3 all-flash storage delivers a consistent latency of 0.3 ms under the pressure of 5000 Exchange users, 1000 VMware Horizon virtual desktops, and hundreds of Oracle users. Even if storage devices are faulty, upper-layer applications can automatically recover within 3 seconds by using the active-active solution, ensuring service continuity. In addition, the data reduction technology reduces the five-year total cost of ownership (TCO) by 75%, making Huawei Dorado V3 an ideal choice for mission-critical workloads.

Stable High Performance

The ESG lab has verified that the OceanStor Dorado V3 can integrate the most challenging workloads and mission-critical tasks into a high-performance and high-availability platform. When 5000 Exchange users, 1000 VMware Horizon virtual desktops, and hundreds of Oracle users are integrated into one OceanStor Dorado5000 V3 system, the integration environment has an average response time of 0.3 ms for these services.

The following figure shows the test results of five thousand 1 GB mailboxes and 9.6 TB Oracle OLTP databases. In this case, the system supports about 105,000 continuous IOPS, while maintaining an average response time of just 0.3 milliseconds.



Source: ESG, 2017

Storage performance of the database and mailbox services

In the next test, personnel simulated a boot storm of 1000 virtual desktops. The I/O soared to about 160,000 IOPS in approximately one minute. At this time, the peak response time increased but remained lower than 0.5 ms. Shortly after this, the two indicators returned to the same level as before the boot storm.



Source: ESG, 2017

Storage performance of a VM boot storm

In addition, the ESG lab created multiple snapshots and then rolled these back, which did not affect system performance. Deduplication and compression are always in the enabled state, and the deduplication and compression ratio is 7:1 for the three types of hybrid workloads. Together with Thin Provisioning, this reduces the total space required by 14:1.

High Availability

Huawei OceanStor Dorado V3 all-flash platform ensures high availability and stable performance during planned maintenance and unexpected interruptions. The HyperMetro feature provides a 100-kilometer, gateway-free, active-active solution to ensure service continuity. RAID-TP technology ensures services continue to run stably even if three disks in a single disk array fail.

In the ESG lab, HyperMetro is first tested in a campus environment. Two OceanStor Dorado V3 storage systems are configured with HyperMetro and run the Oracle RAC database to simulate OLTP transaction database traffic. The total traffic generated by the two servers is 50000 IOPS, the processing volume of a single disk array is 25000 IOPS, and the average I/O response time is 0.25 ms.

The ESG simulated the failure of array 1 by removing its power supply, reducing the total IOPS of array 2 to 0. After about 3 seconds, the total IOPS of array 2 was restored to provide services for all requests in the cluster. A total of 50,000 IOPS are supported, and during the switchover the maximum latency on the controller of array 2 was 4 ms and 8 ms. The latency was reduced to 0.25 ms after 3 seconds.

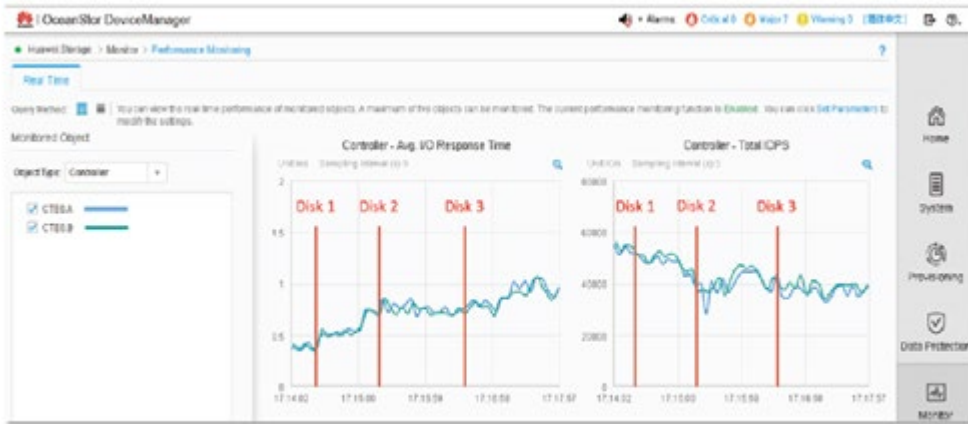


Source: ESG, 2017

Storage performance in the active-active environment in the case of a single-array failure

The ESG also tested the performance of the Dorado RAID TP technology when three disks are faulty. Three SSDs were removed from the disk array at intervals of one to two minutes to simulate the simultaneous failure of three SSDs. The disk array was still available after three SSDs were shown to be faulty without significantly reducing performance.

Storage performance when three disks fail




Source: ESG, 2017

The RAID-TP reconstruction time was also tested. The ESG lab removed an SSD from a disk array that was configured with nine SSDs and 7.4 TB volumes. It took just nine minutes for the RAID-TP to reconstruct the 7.4 TB volume. There is a 7.4 TB capacity and there are nine SSDs, while the reconstruction data volume of each SSD is about 822 GB. Therefore, the reconstruction time of 1 TB of data was only 11 minutes.

TCO

The ESG lab simulated the expected storage TCO of a company that required a high availability hybrid service environment. The OceanStor Dorado5000 V3 has an obvious advantage in this area when compared with Huawei's hybrid storage arrays and the first-generation all-flash storage devices of other mainstream vendors. The five-year TCO was shown to be reduced by more than 75%.

The ESG lab acknowledges that Huawei OceanStor Dorado V3 provides stable high performance and high reliability for its users. Huawei OceanStor V3 perfectly adapts to high-performance service applications and can run in a virtualization environment with high performance requirements. ⚡



Source: ESG, 2017

Five-year TCO comparison



Get More Tips

ESG Lab Review: Huawei OceanStor Dorado V3 All-Flash Storage ESG Lab

Huawei Oceanstor Dorado V3 can be shown to have a lower five year TCO when compared with first generation AFA and hybrid storage solutions.



Huawei OceanStor Dorado V3 All-Flash Solutions, Built for Your Application Scenarios

By Wang Taiping from Huawei

Purpose-built for mission-critical businesses, OceanStor Dorado V3 all flash storage system delivers high performance, reliability, and efficiency augmented by FlashLink technology.

- Innovative FlashLink : enables 0.5 ms stable latency and 3-fold better storage performance than that of traditional storage.
- HyperMetro gateway-free active-active : provides smooth business migration to DCs in geo-redundant mode and solution-level 99.9999% availability, with RPO = 0 and RTO ≈ 0.
- Inline deduplication and compression : Huawei storage claim 3:1 data reduction guarantee (depending on application and configuration) and reduce TCO by 75%.

Huawei OceanStor Dorado V3 fully satisfies the high performance, reliability, and efficiency requirements of mission-critical database services and virtualization efforts, smoothing the way for customers in finance, manufacturing, telecom, and other sectors in the all-flash era.



Server virtualization



Virtualization figure



Database acceleration

Database acceleration scenario (Oracle OLTP)

Oracle database hardware applications are the most widely used series by large enterprises. However, today's information explosion poses great challenges to the performance and costs of Oracle systems.

1. Stacking of devices

Large e-commerce and banking systems process more than 100 million transactions per day. Core business systems perform in excess of 200,000 IOPS while keeping the latency within only 1 to 3 ms. However, traditional storage systems simply stack more devices to achieve increases in performance.

2. Business continuity

Oracle OLTP carries mission-critical businesses for enterprises, remaining unchanged pursuit of business continuity.

3. Increasing costs for Oracle licensing

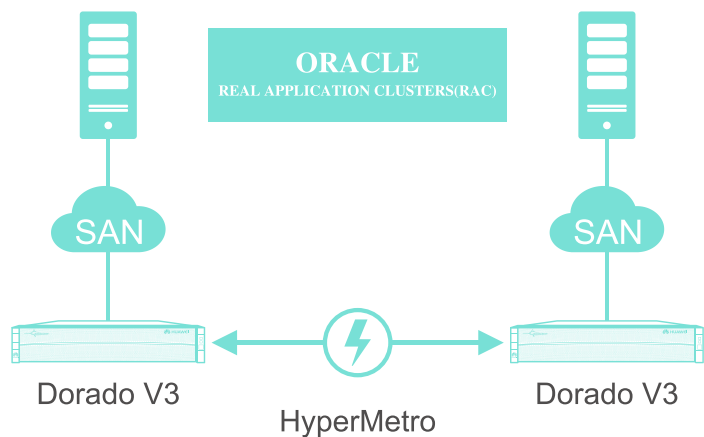
Oracle licensing costs account for a significant portion of the O&M costs in enterprise IT systems. For example, a database with 10 TB of capacity could cost an enterprise as much as 50% of its yearly maintenance budget. As data volumes increase, the annual licensing costs would also continue to rise. Reducing the licensing costs for Oracle applications is a top concern for enterprises.

Not having enough storage performance could stunt the growth of the enterprise, while not having disaster recovery capabilities in place could jeopardize the entire operation.

Huawei Dorado provides high performance, robust continuity assurances, and reduced TCO, ensuring the mission-critical OLTP services of the enterprise remain online while adding new degrees of efficiency in operations.

Huawei all-flash Oracle database (OLTP) solution

The intense performance of the Huawei all flash storage solution powered on OceanStor Dorado V3 doubles online transaction efficiency while keeping latency down to a stable 0.5 ms, minimizing wait times, and improving resource utilization of storage and server assets. Many enterprises have seen a 50% reduction in Oracle licensing costs. The scale-out layout allows for easy, linear expansion. Capacity increases are predictable and scale-outs can be implemented in nearly no time, allaying enterprise concerns over the unpredictability factors in service growth. The Huawei offerings going into the solution provide a rich catalog of data protection services, ensuring that mission-critical applications remain online 24/7.



Huawei all-flash Oracle database (OLTP) solution



Highlights

1. More than 1.5 times faster performance than that of other all-flash storage

After tested by using the simulated system for Oracle database online transaction, a dual-controller OceanStor Dorado V3 system can provide stable performance of 210,000 IOPS at 0.5 ms latency, 10 times higher than that of traditional storage and 1.5 times than that of XtremIO X2.

Storage Device	Workload Model	IOPS
OceanStor Dorado V3	7:3, 8 KB, random	210,000 at 0.5 ms latency 330,000 at 1 ms latency
XtremIO X2	7:3, 8 KB, random	140,000 at 0.5 ms latency 220,000 at 1 ms latency

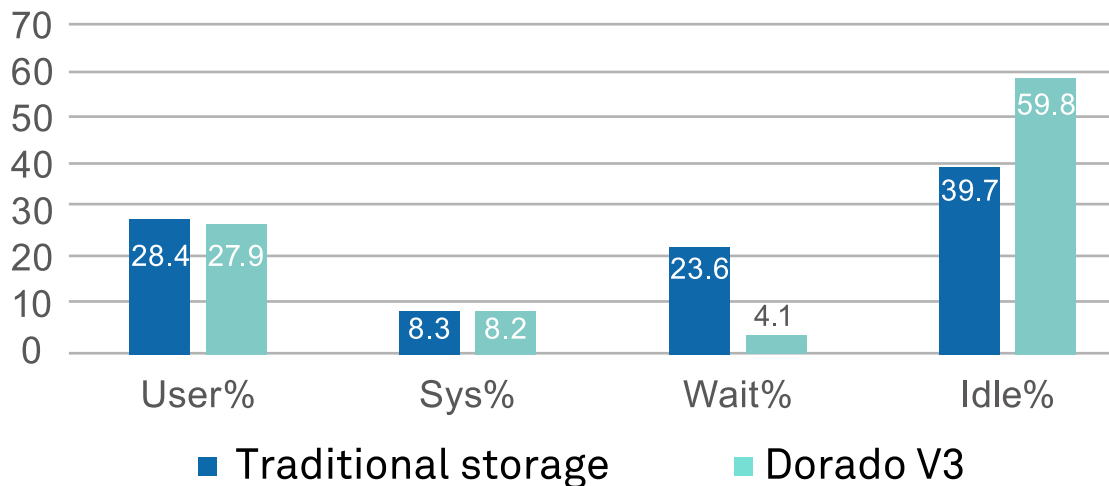
2. Complete DR solution

OceanStor Dorado V3 with the active-active HyperMetro gateway-free feature enabled achieves 99.9999% high availability for mission-critical businesses. The gateway-free attribute of the solution reduces the possible points of node failure, simplifies the networking, minimizes frequency of error from manual operations, and improves overall system reliability. The active-active architecture is able to balance loads and execute seamless service switching.

Vendor/Product	Gateway-free	Active-Active	DC Solution in Geo-Redundant Mode
Huawei OceanStor Dorado	✓	✓	✓
EMC XtremIO	×	–	–
Pure Storage	✓	✓	×
HPE	✓	×	✓
NetApp	✓	×	✓
IBM FlashSystem A9000	×	–	–

3. Lower Oracle licensing costs

The ultra-low latency on Dorado V3 reduces the wait time for CPU I/Os and effectively lowers CPU usage. Under the same workload, using less CPU resources helps increase processing speeds and lower Oracle licensing fees. As the figure shows, compared to traditional storage, OceanStor Dorado V3 storage reduces the CPU usage by about 50%.



Comparison of server CPU usage in Oracle scenarios between traditional storage and OceanStor Dorado V3

VDI scenario (VMware)

Terminal management and delivery services are rather time-intensive and costly in traditional office layouts dominated by PCs. Virtual desktop infrastructures (VDIs) have become widely adopted in enterprise IT systems thanks to the ability to deliver, protect, and manage Windows-based or Linux-based desktops and applications on premise or in the cloud all while controlling costs. Adding to the luster, users are able to work from any location and at any time they chose. However, the availability, performance, and maintenance requirements on centralized deployment of VDI systems are much higher than that of the traditional PC environment.

1. High costs

The VMware desktop solution claims the annual costs in typical use cases (30 GB system disks, 10 GB data disks) is \$699 USD for skilled users, which is higher than the traditional office solution. The investment into storage in the VDI solution accounts for 25 to 40% of the total, a big cost consideration to adopting such a layout.

2. Time-consuming O&M

Deploying a VDI environment for 2000 users takes over 20 hours with traditional storage and it takes considerable time to perform expansions. When emergency situations require patching, it can take over 10 hours to install the patches and perform the restarts for 2000 desktops, causing serious impact to daily operations.

3. User experience gap between virtual desktops and PCs

The traditional VDI solution delivers response time of about 5 to 10 seconds to open a slide, Word, Excel, and other files.

Compared to traditional storage devices, the Huawei OceanStor Dorado VDI solution provides a new balance among high performance, service continuity, and efficiency.

Huawei VDI solution

Huawei OceanStor Dorado V3 all-flash storage provides industry-leading performance and low latency, ensuring quick maintenance of VDI systems and good customer experience. The rich store of DR solutions available with the Dorado fully satisfies different reliability requirements. With the inline data reduction function enabled, TCO is significantly reduced in VDI environments.

Highlights

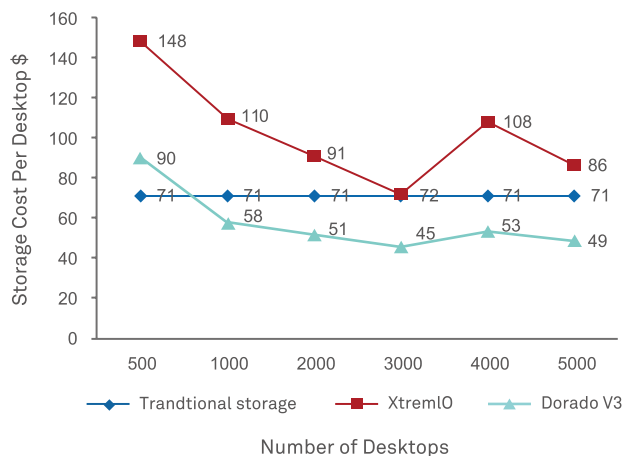
The following testing indicates the experience differences between Huawei OceanStor Dorado V3 and traditional storage arrays in common VDI scenarios.

Test Environment	Test Configuration
operating system	VMware vSphere ESXi 6.5 Windows Server 2012 R2 Windows 7 Enterprise x64 SUSE Linux Enterprise Server 11 x86_64
Virtualization suite	VMware vCenter Server 6.5 VMware horizon view 7
Test tool	View Planner 3.6
Office software	Mozilla Firefox 7.0/Adobe Reader version 10 /7-Zip /Microsoft Office Professional Plus 2010 32-bit
Storage 1	Dual-controller Dorado6000 V3, 25 x 1.8 TB SSDs
Storage 2	OceanStor 5500 V3 with 75 x 600 GB SAS disks
Server	2 x Huawei RH2288 V3 (2 x 2620 v4 CPUs, 128 GB memory) 3 x E9000 blade servers 16 x CH121 V3 (2 x 2600 v4 CPUs, 256 GB memory)
Switch	2 x SNS5120, 2 x CE5850

Comparison of the best practice data of mainstream products and test results of Huawei solution is shown as follows:

1. More than 30% lower cost per desktop

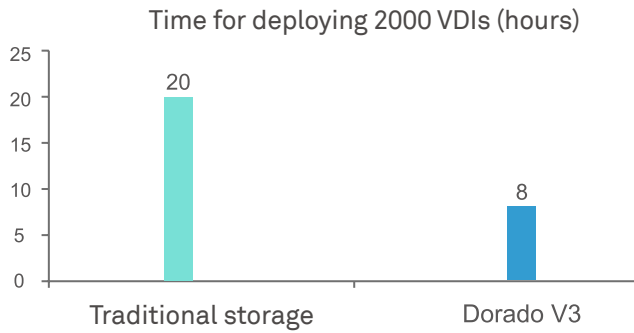
Comparing the single desktop cost among OceanStor Dorado V3, XtremIO, and traditional storage, it can be seen that OceanStor Dorado V3 outperforms XtremIO with the same 5:1 deduplication and compression ratio, while traditional storage did not enable this function. More than 30% of costs can be reduced when over 1000 desktops are operating.



Single desktop cost comparison (USD)


2. High O&M efficiency

OceanStor Dorado V3 storage reduces the time to deploy 2000 desktops from 20 hours to 8 hours compared to traditional storage, and lowers the startup storm time by 75% (60 minutes to 15 minutes).



3. Better user experience

Compared to traditional storage, the response time to open office applications is shortened by 50%.



Scalable VDI storage solution

Number of desktops	<500	500–3000	>3000
OS storage	OceanStor 5300 V5 with SSD&HDD	Dorado 5000 V3 SAS SSD	Dorado 6000 V3
User data storage		OceanStor 5500 V5 with HDD	OceanStor 5500 V5 with HDD

Server virtualization scenario (VMware)

Storage is the biggest bottleneck for a system to achieve centralized business deployment and maximum resource utilization. However, the server virtualization technology can dramatically improve the IT efficiency and elasticity for enterprises. Huawei all-flash storage can cope high performance challenges in server virtualization scenarios, enabling enterprises to avail themselves of the advantages.



1. Business integration

In terms of service requirements, everything from applications being carried, capacity, and performance are all on the sharp rise, making the consolidation and expansion capabilities of the virtualization platform all the more important. The common services carried on virtualized servers include Microsoft services, databases, Web applications, and OA. During peak usage hours, the loads on applications range from 2000 to 5000 IOPS/TB while that during moderate usage ranges from 500 to 1000 IOPS/TB. The performance density of traditional storage is 300 to 500 IOPS/TB, leaving much to be desired in terms of performance and capacity in the new environments.

2. Business continuity

A unified platform with multiple businesses has higher business continuity requirements than those of independent apps. Therefore, service continuity must be supported in virtualization scenarios. For VMware, its DR solution provides the SRM stretched active-active design.

3. Development and testing

The development and test system of virtualization platforms needs to quickly copy and access production data without adverse impact to production services. Use of LUN cloning in traditional storage is time-intensive, taking about 30 minutes to copy 10 TB of data and lowering processing capabilities in production services by 30% or more.

4. Cost reduction

Each VM in a virtualized environment takes up system resources. Resource utilization in storage needs to be improved and reduction in redundant data in the system helps lower TCO.

The Huawei server virtualization solution delivers both high performance and top levels of service continuity while reducing TCO for the enterprise.



Huawei server virtualization solution

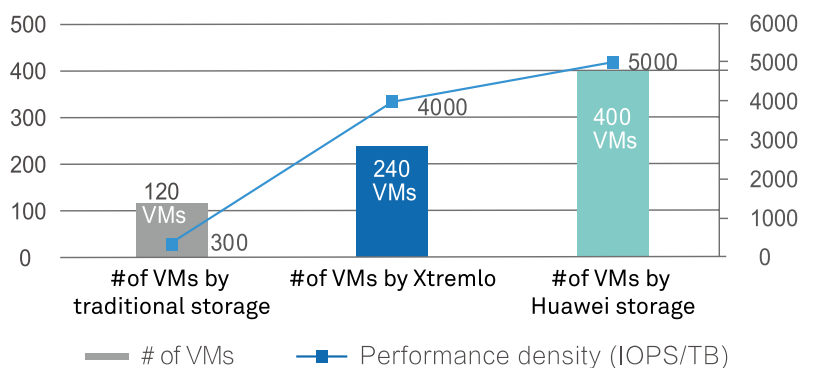
Huawei OceanStor Dorado V3 all-flash storage provides powerful performance, supports large-scale service integration and ultra-low latency, and handles high performance density in server virtualization scenarios and the business pressure in multi-business scenarios. At the same time, Huawei solution supports active-active design with VMware SRM stretched storage to build an end-to-end high-availability virtualization platform.

Based on the thin provisioning and inline data reduction technologies, OceanStor Dorado V3 reduces costs of virtualization infrastructure platform, saves equipment room footprint and power consumption, and lowers TCO.

Highlights

1. Three times better business integration capability, with VMs bearing mission-critical businesses

OceanStor Dorado V3 improves the performance density by three times compared to traditional storage, and by twice compared to XtremIO X2. OceanStor Dorado V3 delivers better performance in heavily-loaded virtualization scenarios and supports 500 IOPS/100 GB for a single VM.



2. Support for VMware SRM stretched storage active-active layouts

Service continuity: Huawei gateway-free active-active solutions have acquired certification from VMware SRM Stretched Storage.

Vendors	Asynchronou DR	SRM Stretched Active-Active	Gateway-Free
Huawei OceanStor Dorado	✓	✓	✓
EMC XtremIO	✓	✓	×
Pure Storage	✓	×	-
NetApp	✓	×	-

3. Three times higher development and testing efficiency based on lossless snapshots

Traditional development & test: three complete copies of data, slow speed, and high space occupation rate

Dorado V3 development & test: Lossless snapshot (no need for copying) occupies small volume of space and executes operations with zero impact to production services.

SAP HANA TDI solution

SAP HANA is a high-performance in-memory database platform launched by SAP SE. It aims to help enterprises analyze and process large amounts of data in real time and seize business opportunities. According to Gartner, a well-known third-party advisory firm, more than 100,000 SAP customers will be upgraded from traditional databases to SAP HANA databases by 2020, accounting for more than 35% of SAP users.

Although SAP HANA accelerates large-scale service innovation, SAP HANA has certain requirements for the legacy data center architecture.





1. Data center architecture

SAP HANA mainly employs the soiled construction model, bringing business island, low resource utilization, and complex business expansion. This is contrary to the trend of migrating businesses to the cloud for mainstream data centers.

2. DR and backup building

In SAP HANA appliance mode, the DR of SAP applications and database cannot be deployed in a unified manner. The database adopts software-based DR while the application adopts virtualization- and storage-based DR. The management, maintenance and DR operations are complex, which cannot ensure data consistency and cannot meet the requirements in large-scale data center scenarios.

3. Business maintenance process

Traditional databases use advanced functions of storage systems to provide data services. The SAP HANA platform supports internal enterprise businesses, involving dozens and hundreds of sub-service systems which demand for extensive development, testing, rollout, and positioning. Business innovation requires rapid release of development and testing, and sandbox systems. However, the SAP HANA system without snapshot function takes a long time to have required full copy testing for each new business rollout and upgrade. Each operation needs to be coordinated by multiple layers, such as application and hardware layers. Many operation steps make it difficult for hardware administrators to handle.

The preceding challenges are especially prominent in cloud-based data centers and in the evolution from traditional databases to the SAP HANA, making it a major challenge for enterprises to use SAP HANA system.



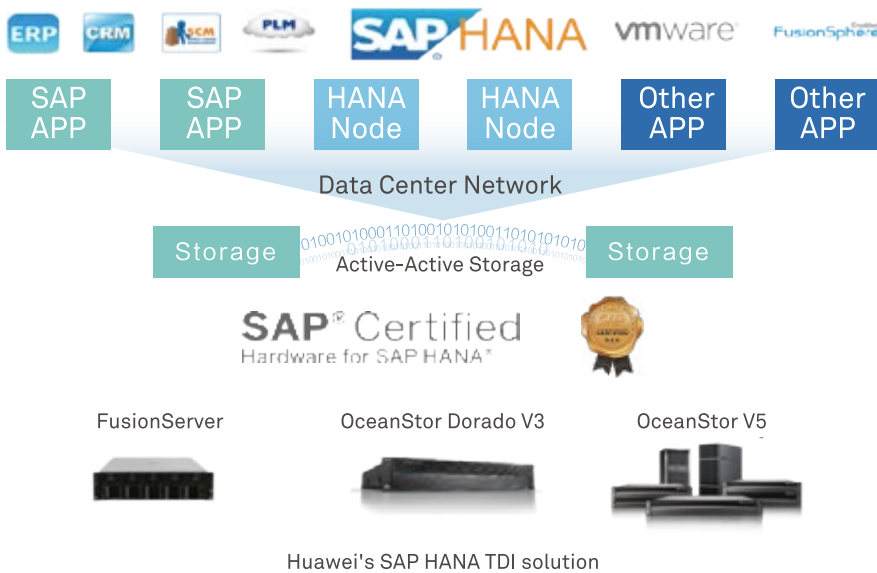
Huawei's SAP HANA TDI solution

With the maturity of SAP HANA, hardware faces fewer restrictions than before. To better integrate SAP HANA into existing data centers of enterprises, SAP launches Tailored Datacenter Integration (TDI) reference architecture to integrate SAP HANA database more flexibly into existing data centers.

Huawei's SAP HANA TDI is a solution designed based on Huawei's powerful hardware capabilities and integration capabilities of the TDI solution. Compared to other vendors' solutions, it has the following three features:

Scalable: SAP businesses require high hardware performance. The higher the performance is, the more HANA nodes are carried. In addition, the preceding performance data takes effect only after being certified by SAP. Huawei's outstanding performance in all-flash storage has been recognized by SAP. In the latest SAP certification, Huawei's high-end all-flash storage system supports a maximum of 200 SAP HANA nodes, ranking top in the industry.

Reliable: Huawei all-flash storage can still meet SAP's high performance (1 ms latency) requirements after enabling the active-active design. Therefore, Huawei's SAP HANA TDI solution can provide end-to-end SAP HA solutions that satisfy SAP performance requirements.



Efficient: With the lossless snapshot function, Huawei's SAP HANA TDI solution can quickly set up the development, testing, and sandbox systems. Response time to the business environment setup is reduced from 10 hours to 20 minutes compared to traditional solution.

Huawei is the first storage vendor that puts NVMe into commercial use. Designed for enterprises' heavily-loaded critical applications, Huawei's all-flash storage meets storage requirements of enterprise-class applications such as databases, virtual desktops, and server virtualization, perfectly matching solid-reliability and high-performance standards in SAP HANA scenarios to build an industry-leading SAP HANA TDI solution. ⚡



[Get More Tips](#)

SAP HANA TDI Solution
 Specifically designed for large enterprise data centers to deliver seamless business integration, solid reliability and efficient O&M.

Can You Imagine That the Nasdaq Market Shutdown Was Caused by a False Fire Alarm?

By Zhao Yanling from Huawei

Financial Disaster Caused by a False Fire Alarm

On April 18, 2018, the Northern European Nasdaq transaction market was frozen until 2:00 p.m. local time, according to The Register, a well-known British tech news website. A fire alarm was reported in the core equipment room of the data center, and high-pressure inert gas was vented to extinguish the fire. The resulting high pressure generated thundering noises, which caused the racks to shake violently. As a result, the core equipment of the stock trading system broke down. Although Nasdaq has enabled the backup data center in Väsby, Sweden, the transaction system could not be restored until 2:00 p.m.

Data Centre

Noise from blast of gas destroys Digiplex data depot disk drives

Nordic Nasdaq knocked as deleterious decibels crashed servers

By Chris Mellor 26 Apr 2018 at 13.57

34  SHARE 



Huawei Storage Survived a Similar Trial in a Top-class Carrier's Equipment Room

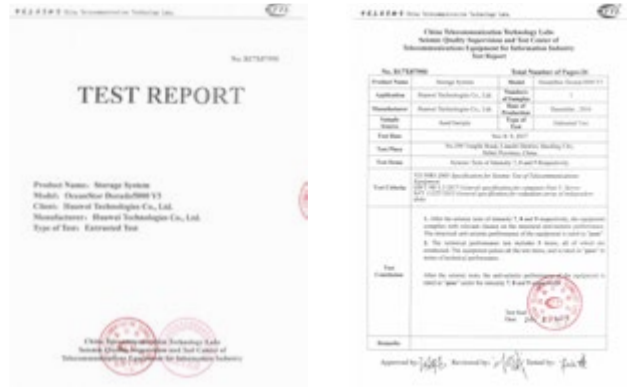
Data center breakdowns as a result of false fire alarms are not uncommon. The Register also mentioned that a large number of devices were damaged due to the same problem in a data center in Sydney four years ago. This resulted in large-scale complaints from tenants. In December 2017, the situation repeated itself once more in the core equipment room of a top-class carrier, where extensive equipment damage caused some services to stop for a week.

Storage devices of four vendors, including Huawei, were used in the core equipment room of the carrier. Huawei provided OceanStor Dorado V3 all-flash and OceanStor V3 hybrid-flash storage systems, with a total of 468 all-flash disk drives and 333 spinning disk drives. After the incident, 0 of Huawei's 468 SSDs and only 4 of the 333 NL-SAS disks required replacing, with no corruption of the service data running on Huawei storage devices. In contrast, the other three vendors in the same equipment room suffered damage to a large number of their spinning disk drives or even to the storage devices. This did not just result in data loss but also week-long service interruptions while all of the damaged devices were replaced.


Huawei Quality Proven by the National Magnitude-9 Earthquake-resistant Certification

So, is it luck that allowed Huawei’s storage devices to perform so well? The answer is certainly no.

Huawei has been dedicated to work in the ICT field for 30 years and has accrued profound experience in device vibration protection. Last year, Huawei’s all-flash storage system obtained the Magnitude-9 Earthquake-resistant Certification of the China Telecommunication Technology Labs (CTTL). It is the only IT product series that has been awarded this CTTL certification.



Magnitude-9 Earthquake-resistant Report of Huawei storage systems



Magnitude-9 Earthquake-resistant Test

Authority of the “National Magnitude-9 Earthquake-resistant Certification” of the CTTL

The CTTL was founded on September 21, 1981, with the approval of the former General Administration of National Standards and the former Ministry of Post and Telecommunications of the People’s Republic of China. In China, all public communication equipment must pass CTTL performance tests before obtaining a network access permit. In addition, the CTTL is an authorized lab for international admission systems of the U.S. FCC, EU CE, Canada IC, Japan JATE, and Egypt NTRA. The earthquake-resistant quality inspection center of the CTTL is the only authority in China that provides seismic proof that main equipment of the public telecommunication networks can resist earthquakes. It is fundamental to prove that the equipment can endure earthquakes that are magnitude-7 and above.

The CTTL has professional seismic testing equipment, including world-leading three-dimensional six-degree-of-freedom (DOF) seismic simulation hydraulic shaking table, vertical and horizontal shock wave tables, a variable capacitance triaxial accelerometer, and precision data acquisition analyzers. These devices can accurately simulate a magnitude-9 earthquake in X, Y, and Z dimensions. The following figure shows a Huawei storage system undergoing a simulated magnitude-9 earthquake at the CTTL inspection center.

Would you like to see the tests for yourself? If so, you are in luck! Thanks to the professional support of the CTTL, we have moved the three-dimensional six-DOF seismic simulation hydraulic shaking table to the exhibition hall of the Huawei Chengdu Research Center. As of now, our all-flash storage system has gone through simulated magnitude-9 earthquake tests for a combined total of 2578 hours, and it is still running smoothly.

Many earthquake-prone countries and regions, such as Japan, have very stringent requirements for seismic resistance of IT devices due to the serious threat posed by earthquakes. In addition to earthquakes, abnormal equipment room environments (such as high air pressure), transportation and installation of the equipment, and internal sources of vibration (such as fans) can threaten the operation of equipment. Therefore, an effective seismic-resistant design is critical for the proper operation of the equipment.



Highly Strengthened Cabinets

Since its research and development in 2002, Huawei storage has been refining its hardware architecture. In the early days, preventing vibration on mechanical hard disks was a core issue. In collaboration with the Data Storage Institute (DSI) of Singapore, Huawei has mastered the industry's most advanced methods of finite element modeling (FEM), simulation, and testing and optimization. In addition Huawei has calibrated the model based on the data of tens of thousands of observed cases, and has built the most stable and reliable storage device in the industry.

Based on the FEM, Huawei storage devices provide end-to-end anti-seismic protection from cabinets to subracks and then to internal core components, fans, and disk enclosures. Each storage device that is delivered to customers has passed 30 drop, shock, and vibration tests, as well as spectrum tests for over 170 types of roads. This is what makes Huawei storage systems so reliable in different environments.



E2E shockproof design

30 drop/shock/
vibratio tests




Road spectrum teste on
170+ roads



Highly Strengthened Cabinets

In terms of shockproof design, Huawei storage devices to be delivered in full cabinets have highly strengthened connections. This ensures that the cabinet can hold up to twice its weight without suffering irreversible deformation. In addition to this, if tension equal to 2.5 times the weight of the cabinet was applied from the front, back, left, and right of the cabinet, it would receive no permanent structural damage.



Highly strengthened cabinets

Three-dimensional Six-DOF Subracks

Subracks are fixed with three dimensions, six directions, and full degrees of freedom to ensure that the frame is incredibly strong. In addition, the design of guide pins and bayonet ensures that the boards and modules are tightly connected, and subracks, especially disk enclosures, are securely installed. The comprehensive anti-seismic capabilities of the cabinet meet rigorous magnitude-9 earthquake-resistant standards.



Subracks with 3D six-directional full Dof

Shockproof Design of Internal Core Components

Fans are a primary source of internal vibration. Huawei leverages vibration isolation and absorption technologies in the three-layer vibration reduction solution through the fan and bracket, bracket and module structure, and module structure and chassis layer according to layers of vibration reduction. This reduces the vibration level of fan modules, minimizing the impact on the system. For hard drive handles of vibration-sensitive components, the external rigid spring plate is used to overcome the low frequency vibration, and the internal damping pad is designed to mitigate the high frequency vibration. This ensures protection from vibration and prolongs the service life of hard disk drives. The return rate of Huawei hard disk drives is 70% lower than the average rate in the industry.

Huawei's Craftsmanship Spirit

In today's world, data transformation leads to the extensive use of IT infrastructure. This means that we need to consider a more rigorous and complex device application scenarios than ever before.

In Japan, the country that records the most earthquakes in the world, Huawei storage has helped Yahoo Japan to achieve 24/7 business continuity.

In rough seas, Huawei all-flash storage ensures steady services onboard the German cruise liners AIDA and Viking for dozens of days at sea.

In harsh mining environments, Huawei all-flash storage enables the Australian mining group south32 to achieve round-the-clock operation and data analysis of hundreds of mining machines.

We are ceaselessly working to ensure continuous improvement and to meet customer needs, and has endeavored to achieve excellent storage performance

and endurance in rigorous environments. This does not only refer to areas prone to disasters such as earthquakes, but also to environments such as those with extremely high/low temperatures, severe acid and alkali corrosion, and high and low pressure. We fully consider the environmental requirements of our customers and continue to develop and hone stable and reliable storage products. In addition, in terms of high performance, Huawei continually devotes R&D efforts to proprietary chips and the proprietary FlashLink™ algorithm, driving it to continue to surpass itself and deliver storage products that have premium performance.

As the saying goes "Little by little, one travels far", good products come from continuous efforts. Huawei has been forging ahead with the spirit of craftsmanship. ⚡

Moving Your Mission-Critical Applications to Cloud Just Got Easier!


By Liu Lin from Huawei

Public cloud services are developing rapidly all around the world. Enterprise data centers are quickly evolving to the cloud, leaving traditional storage methods obsolete in the 5G era. Reports have shown that 68% of enterprises in North America have started to use cloud services to various degrees. Companies in other parts of the world are reported to be trying out the new cloud-based service models.

Why are enterprises around the world so excited about the new service models?

Because they are easy to use and reduce TCO, and the ease of use is provided with on-demand and quick resource acquisition capabilities. Enterprises are always looking for ways to save on TCO through adoption of new applications and techs. According to statistics, the costs for storing data in self-purchased storage are 19 to 30 times higher than those for leasing public cloud storage. Considering the costs, it's no wonder more and more enterprises are using cloud services.

Large cloud service providers can handle just about any workload for organizations of all sizes. However, enterprises are still somewhat reserved about placing their mission-critical applications and data on the cloud due to the security requirements levied and high performance requirements. To address any reservations that enterprises may have, many cloud service providers deliver high security guarantees to prove to potential customers that migrating mission-critical services to the public cloud is safe, quick, and easy.



Dedicated Enterprise Storage Service (DESS) migrates mission-critical applications to the cloud without making any forklift upgrades

HUAWEI CLOUD has a proven track record of helping many move their applications to the cloud, even in the most demanding of scenarios. With HUAWEI CLOUD, security and extreme performance are no longer blocking mission-critical applications from moving to the cloud.

DESS uses Huawei all-flash and hybrid-flash storage to build dedicated storage services dedicated to enterprises. Oriented to enterprises' key applications, such as Oracle RAC and SAP HANA TDI, it provides the same superb performance and reliability as enterprise private cloud, simplifying service deployment on public cloud for enterprise users.

DESS users can exclusively use physical storage resources in HUAWEI CLOUD to ensure that applications are stably running on isolated storage devices. In a typical database performance model, databases can obtain a stable latency of

0.5 ms and a read/write capability of 9000 IOPS/TB. Thanks to the robust reliability of enterprise-class storage, users can use the same high-reliability solution as that used in self-owned data centers, with 99.99999% of data durability.

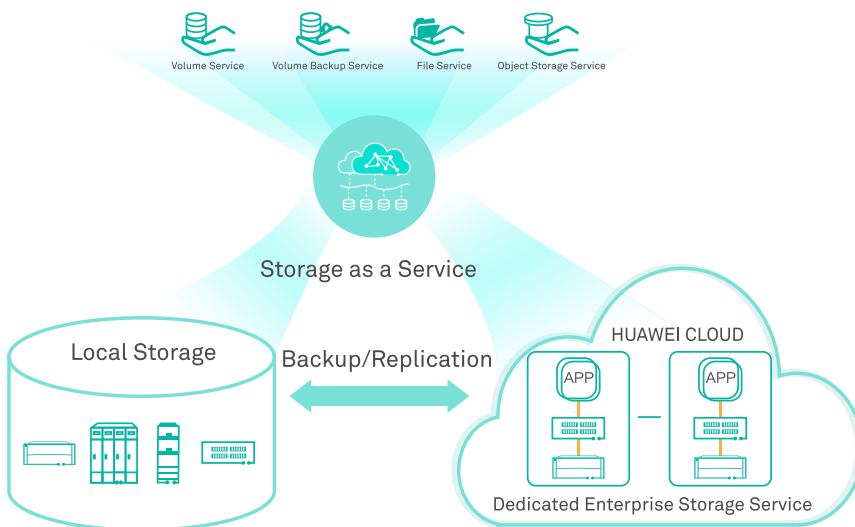
Since its release in September 2017, DESS has been used by many large enterprises and government agencies to successfully deploy their key applications on the cloud. For example, after migrating the ERP system to DESS, a large-sized private healthcare company has reduced the time needed to generate financial statements from several days to just several hours, helping to improve enterprise business analysis and decision-making efficiency by dozens of times.

STaaS solution centrally manages resources on- and off-premises

For customers who have some of their applications running on the cloud, Huawei also provides the Storage as a Service (STaaS) solution to facilitate storage resource management and data flow on- and off-premises. The STaaS solution can manage various storage devices in the on-site data center and access to HUAWEI CLOUD. Enterprises can manage on- and off-premises resources over a single pane of glass, eliminating the troubles in switching back and forth between interfaces.

The rich selection in the service catalog enables high degree of automation in provisioning of storage resources at both on-site and cloud-based data centers. Even such operations as volume creation, file sharing, creation of active-active relationships, volume backups, volume replications are all easily handled with time-saving automations.

The results speak for themselves. KIO, the largest ISP in Mexico, uses the STaaS solution to centrally manage and schedule storage resources in multiple data centers, reducing manual operations and improving work efficiency by over 50%



The impressive on- and off-premises resource management and automated service provisioning capabilities have won over customer after customer in finance, telecom, public services, retail, and other sectors. What's more, after intense vetting and proven track record with zero incidence after deployment, Huawei has since become a brand that you can trust. Therefore it's no surprise the Huawei STaaS solution has become the compelling choice in hybrid cloud environments able to carry the mission-critical applications for organizations of all sizes.

Find out how Huawei can smooth out moving your mission-critical applications to the cloud. ⚡



Get More Tips


Video: Huawei Storage-as-a-Service Solution
 Huawei Storage-as-a-Service Solution supports cloud-based storage consolidation and flexible orchestration of on-premises and off-premises resources.

Huawei's Business Integration Solution: A Revolutionary Optimal Data Migration Method


By Li Wuxian from Huawei

Enterprise data is stored in professional storage. For large enterprises comprising many data centers and massive volumes of data, storage is often a pain point. Storage devices need to both provide instant access and take up less floor space, all at a lower price than other alternatives. Storage device replacement is required when devices are out of warranty, end of life, or encounter insufficient performance and capacity.

Replacing storage devices involves business integration and data migration. Data migration does not mean simply moving data from one place to another. Instead it involves data security, compatibility, migration efficiency, outage time, as well as considerations of its impact on business. Currently, there are three main types of data migration: at the application layer, at the host layer, and at the storage layer. Because of the many different aspects surrounding storage and data migration, it is no wonder that it is the focus of many industry giants. From a customer standpoint, enterprises need to effectively judge how data migration will impact their business and services. Further, they need to find ways to both enable online data takeover and migration, and ensure a quick migration speed, as well as choose a solution that delivers an easy-to-use interface for future business expansion and integration.



Analysis of Mainstream Data Migration Solutions



Data migration at the application layer: A major migration method based on database and virtual machines (VMs) functions at the application layer, this solution is used in cross-platform and cross-media migration scenarios and scenarios with short downtime. It features simple access to the target storage and little impact on the source network topology. In some scenarios, migration and replacement tolerate short service interruption.

Because this method of storage requires host resources during data migration, it can create huge pressure on service systems. It is often associated with low migration speeds, sometimes even less than 200 MB/s. In addition, has application-layer migration solutions have restrictions on application configuration. For example, Oracle ASM Rebalance supports only Oracle ASM to manage disks, and VMware Storage vMotion does not support online migration in some scenarios (VM shared disks and RDM disks, etc.).

Data migration at the host layer: implemented by the volume management software. Common Veritas Volume Manager (VxVM) of a third party and Logical Volume Manager (LVM) of AIX, HP-UX, and Linux are used to migrate data online. Host-layer migration solutions ensure that services are not interrupted during migration, satisfying storage space consolidation and cross-media migration demands. However, the solutions may provide slow migration speeds of less than 100 MB/s, as well as large consumption of host CPU and memory resources. In addition, manual inputting of command lines is required during the migration, which is prone to errors.

Data migration at the storage layer: Utilizes the built-in functions of storage devices to migrate data at the storage layer. These functions are as follows:

LUN copy and remote replication: used for migration between homogeneous storage, where the migration speed of LUN copy can reach up to 500 MB/s.

Remote replication: used for long-distance migration, where its online migration speed is about 100 MB/s on average.

Heterogeneous virtualization: uses storage virtualization functions to take over data from heterogeneous devices and then migrate the data to the source end. Its migration speed is up to 500 MB/s and third-party storage devices are supported.

Data migration at the storage layer features fast migration speeds and does not impact host system performance. Likewise, there is no vendor lock-in, allowing users to migrate to and from third-party storage devices. However, LUN copy is offline, and the switchover of remote replication service takes a long time. Heterogeneous virtualization requires suspension when taking over services.


Storage-layer migration solutions need to provide non-disruptive data migration processes, and support online takeover migration.

Based on these demands, Huawei launched OceanStor V5 Hybrid Flash Storage, OceanStor F V5, and OceanStor Dorado All Flash Storage series. These are business-unaware migration solutions with zero downtime perfectly designed to optimize heterogeneous virtualization, implementing online takeover and migration of data at the storage layer.

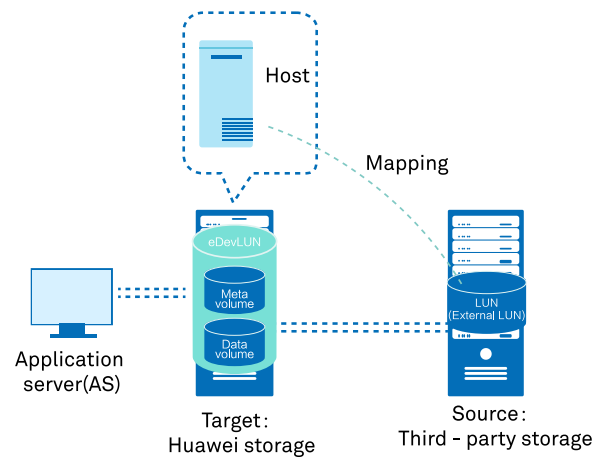


Huawei's Heterogeneous Virtualization Data Migration Solution

OceanStor V5 Hybrid Flash Storage, OceanStor V5, and OceanStor Dorado All Flash Storage series adopt SmartVirtualization and SmartMigration in heterogeneous virtualization migration scenarios. These functions migrate data in the source array to the target array at the storage layer. Users can use SmartVirtualization to take over the source array (third-party storage) online, and then migrate the data online by using SmartMigration. After the migration is complete, users can remove the source storage from the live network.



Implementation Principles of Huawei's Heterogeneous Virtualization Migration Solutions



Customers can move data from third-party storage to Huawei's storage External LUNs on third-party storage devices to store actual data. These Huawei devices are commonly known as eDevLUNs. An eDevLUN consists of a Meta Volume and a Data Volume.

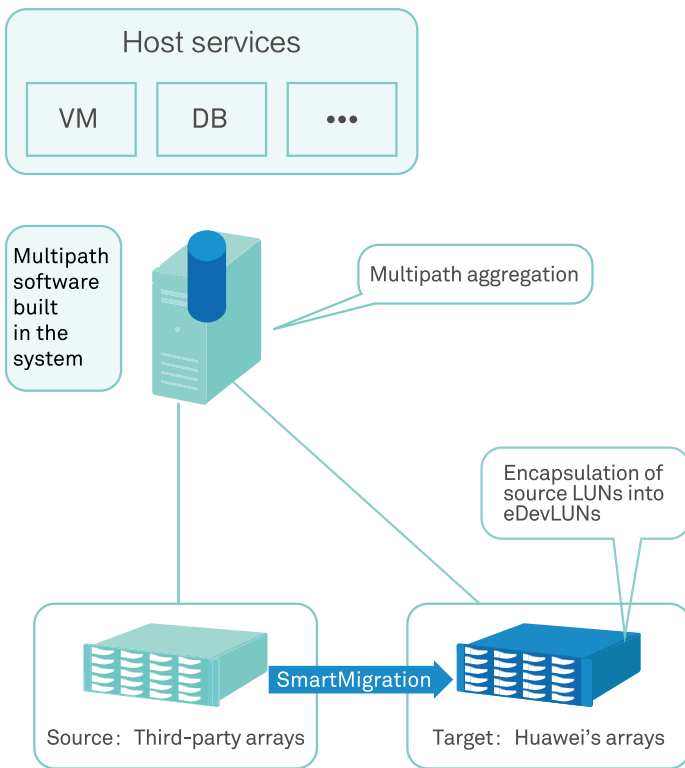
The Meta Volume is used to manage the data organization of eDevLUNs. The storage space of Meta Volumes is provided by Huawei's storage devices, and occupies little space.

A Data Volume is a logical abstract object of external LUN data. An external LUN provides all data storage space. eDevLUNs created on Huawei's storage systems have one-to-one mapping relationships with external LUNs. Application servers manage data on external LUNs through data reads and writes on eDevLUNs.

In a heterogeneous virtualization solution, LUNs mapped to source arrays are taken over during the initial migration phase, and LUNs of third-party storage are initialized into eDevLUNs that can be managed by the target arrays under the condition that no data is lost. LUNs are migrated using value-added features of heterogeneous virtualization, which enables data to be rapidly migrated between the source and target arrays. Heterogeneous takeover can be performed online, and the services are non-disruptive during data replication.



Huawei's Online Migration Technologies in Heterogeneous Virtualization Scenarios




Compared with takeovers in traditional heterogeneous virtualization scenarios, Huawei's takeover can encapsulate the source LUNs into eDevLUNs without interrupting host services. Huawei's storage maps the eDevLUNs to the host so that the hosts' multi-pathing software can automatically aggregate the physical paths of the source LUNs and eDevLUNs. The software runs host multi-pathing commands to delete the corresponding physical links of the source LUNs and switch all I/Os to the link to implement online takeover with zero downtime.



Process of Online Takeover and Migration

1. Connect the source arrays to the target arrays, and map the LUNs to be migrated to Huawei's disk arrays.
2. The Huawei's disk arrays encapsulate the LUNs mapped to the source arrays into eDevLUNs that are then mapped to the source host.
3. Use multi-pathing software on the host system to aggregate physical links between the source LUNs and eDevLUNs and delete the physical links of the source LUNs to complete online takeover.
4. After the online takeover, use SmartMigration to migrate data from the source LUNs to Huawei's disk arrays to implement online data migration.



Advantages of Huawei's Online Migration in Heterogeneous Virtualization Scenarios

- Mainstream operating systems can be used to migrate data between SAN storage arrays from different vendors.
- During migration, data is replicated on storage arrays, which does not occupy CPU or memory resources of the host, imposing zero impact on business.
- Services are not interrupted during array takeover and data migration
- Data migration speed can reach up to 500 MB/s.
- Data can be migrated conveniently and quickly without installing any extra software or agent in the existing environment.
- Migration Director is a professional migration tool that supports automatic operation on an efficient graphical user interface. ⚡



CRYSTAL KEY : CUSTOMERS' VIEWS & SUCCESS STORIES

Voice of Industry Users

Vostochny Bank Quickly Expands Businesses with High-Performance Storage Devices

The All-Flash "Renaissance" of Telecom Italia

China Life: Transformation to Technical Guidance for Business Innovation

BYD Auto Becomes a Leader in New Energy with All-Flash Storage

One-Picture-Say-It-All: Yahoo Japan

One-Picture-Say-It-All: Telecom Italia

Marco Geraci

CTO of ACESI Group

“

We plan to migrate progressively all our cloud storage to Huawei OceanStor Dorado6000 V3, and we expect to have a better deduplication rate in order to have lower costs for our storage. and then we will be able to propose more competitive cloud services to our customers.

We plan to add another Huawei OceanStor Dorado6000 V3 for disaster recovery plan, then we will have an emergency disaster recovery plan for our entire cloud.

”

Vostochny Bank Quickly Expands Businesses with High-Performance Storage Devices

By Zhao Yanling from Huawei

Founded in 1991, Vostochny Bank is the largest commercial bank in Russia's Far Eastern region. Following a series of successful mergers and acquisitions, Vostochny Bank has been expanding its business rapidly and now has more than 700 branches across 10 federal subjects in Russia.

High performance, robust stability, and 3x better business efficiency for the next five years of growth


A series of mergers and acquisitions has placed intensive pressure on Vostochny Bank, including its 2016 merger with Uniastrum Bank, which greatly expanded its user base and drove its maximum volume of business from 65% to 93%. The bank needed to consolidate its automatic banking system (ABS) with Uniastrum Bank's core systems, increasing the complexity of its business system management.

"This large-scale merger led to a sharp increase in our IT systems' workload," stated the IT director of Vostochny Bank. "We also worried that the performance and stability of our IT devices wouldn't cope when a large number of customers used the transaction system at the same time. Although we had configured DR protection for our two data centers, system switchovers could still have disrupted business for several hours and caused data losses if

a fault had occurred at our production data center, which would have impacted customer satisfaction and could even have led to attrition. What was particularly worrying was that at that time our storage systems also could not fulfill the regulatory requirements of the Central Bank of the Russian Federation. This meant we needed to upgrade our existing IT architecture to ensure high performance and provide the necessary stability and reliability for our IT business systems."

Customer information (user data and permissions), transaction SNs (e-banking transaction record lists), product libraries (peripheral product data and operation logs), and report databases (e-banking report review and reconciliation history data) are stored in data storage in the main e-banking business process. They are used to form an e-banking database.

Typical e-banking Business process and Data Storage Requirements





E-banking databases are the lifeline for banks, who often use the e-banking synchronous replication function to create local copies and intra-city copies for data in e-banking databases. E-banking services also archive data over one year old.

For many banks, the main issue surround storage is the amount of data generated. For example, every 10 million customers will create 10 TB of data, and data is often stored for 1 to 3 years before being archived. Even though data that comprises user information and transactions is not large in volume in some scenarios, high performance and business continuity must be ensured to deliver optimal customer experience. Vostochny Bank required data storage providing 400,000 IOPS at a latency of less than 3 ms.


CFS, the leading ISV in Russia's financial industry, is a business provider for Vostochny Bank's ABS. It acted as a consultant for this IT upgrade, proposing standards and IT planning requirements for the next 5 years based on the bank's status quo and its pace of development. They suggested a performance of 400,000 IOPS at 1 ms of latency and a target of zero disruption for the bank's core business transactions.

Vostochny Bank under the guidance of CFS invited vendors "E" and "H" (names not included for privacy reasons), and Huawei to participate in the POC testing. The bank designed 15 business models for the systems on the live network to test the performance of each vendor's devices. Originally the bank planned these tests to use two storage devices to achieve 400,000 IOPS at 1 ms of latency, but they were surprised to discover that a single OceanStor Dorado storage system can provide such a low latency as well as a performance of 410,000 IOPS. When two OceanStor Dorado storage systems were deployed at the production center, the performance improved 3fold.



OceanStor Dorado V3 Performance Test				
Scenarios	Read Ratio (%)	Random Ratio (%)	IOPS	Avg_rt(ms)
OLTP	30	100	390355	0.98
8K Random	75	100	415159	1.006

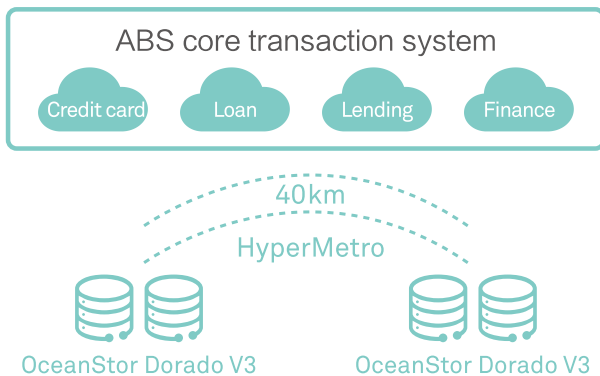
To handle its rapid expansion, Vostochny Bank attached great importance to the performance of the system with value-added features enabled. Huawei's OceanStor Dorado delivers stable performance even when deduplication and compression are enabled, snapshots are frequently created, and services are switched over between active-active storage systems. What's more OceanStor Dorado's secret weapon is that data is reconstructed if three disks fail simultaneously, possible even in snapshot-heavy scenarios in which snapshots are created several times each week for the testing and development of merger and acquisition business systems. In the bank's testing, Huawei's Dorado provided a stable performance 35% higher than that of H's new flash solution.



Gateway-Free Active-Active Design Ensures 24/7 Business Continuity

If a fault had occurred under Vostochny Bank's legacy dual-center architecture, data would have been lost and its services would have been disrupted for several hours. Losing core business data is unacceptable, and lengthy disruptions to the bank's business would have severely affected customer satisfaction and the trust they have in the bank.


Huawei's Dorado all-flash storage significantly improves performance and its industry-leading all-flash active-active architecture prevents disruption or data loss, ensuring 24/7 business continuity. Huawei's active-active data center architecture can even achieve 400,000 IOPS at 1 ms of latency when its two data centers are 40 km apart, which put Vostochny Bank's concerns to rest.



huawei's gateway-free active-active design provides best-of-breed performance at 1 ms of latency.

3:1 Data Reduction Guarantee Reduces TCO by 50% over 5 Years

Vostochny Bank's expansion strategy has promoted its rapid growth in recent years, but it has also consumed large cash flows. After its merger with Uniastrum Bank, Vostochny Bank needed to control costs. Huawei's OceanStor Dorado leverages inline deduplication and compression to deliver outstanding performance and improve its space utilization 3fold. The solution also reduces power consumption and cooling costs while greatly cutting down on construction and O&M expenses, which will lower the bank's TCO by 50% over the next 5 years.




Huawei's OceanStor Dorado all-flash active-active solution met all of Vostochny Bank's fast business development needs, providing superb performance, high reliability, and 24/7 business continuity. Even in the most demanding scenarios, Huawei's OceanStor Dorado can excel. ⚡

China Life: Transformation to Technical Guidance for Business Innovation

By Li Zhenjun from Huawei

China Life Insurance (Group) Company, founded in 1949, is the largest commercial insurance group in China and has been listed in the world's top 500 companies for 16 consecutive years with assets worth up to 3.35 trillion RMB.

Unit: 100,000,000RMB



As a leader in the insurance industry in China, China Life has formed an industry-leading overall IT layout based on multi-active data centers, including business models "big background, small front-end systems", and "network-wide interconnection". Following their three-step strategy of "Online China Life", "Intelligent China Life", and "Digital China Life", China Life actively lays a solid foundation for enterprise transformation and business development.

China Life has gained a deep insight into the development trends of massive data sets, perfect for development today and in the future. For many large-scale businesses, a common consensus is that SSDs will inevitably become the development direction of data storage, meeting the requirements of massive storage capacity and improving performance. In 2014, China Life was the first in the industry to explore innovative technologies in the all-flash field, making efforts to improve SSD usage in research, testing, trial runs, and commercial use. After three years of technological accumulation and production practices, Chinese Life successfully transformed from using HDDs at its data centers to all-flash media.

All-flash storage greatly improves the service processing efficiency of Chinese Life's data centers, and effectively meets the business requirements of large capacity and high concurrency from Internet users, providing infinite space for future business development and significantly enhancing customer experience.





Accelerating Technology Innovation and Achieving Advanced Productivity

Since 2015, Huawei OceanStor 18000 series is a high-end all-flash storage family that has been put into use in China Life's data centers. With six-nines stability, Huawei's high-availability solution powered by the high-end all-flash storage establishes a platform for China Life for stable and reliable data storage, analysis, and management. Equipped with 3 million IOPS, the storage helps China Life easily cope with business peaks. In addition to improving existing storage performance by 200%, it reduces space and power consumption by about 60% to 80% thanks to the use of 900 GB and 1.8 TB SSDs, and also reduces the number of standard cabinets from five to one.

In the China Life all-flash solution, Huawei HyperMetro technology provides the active-active function and integrated SAN and NAS, and has successfully deployed underwriting, claim settlement, ERP, actuary, office, and many other core business systems. During the all-flash transformation of China Life, the five convergence concepts of the Huawei OceanStor 18000, including the convergence of HDDs and SSDs, and that of SAN and NAS, have created a stable period for China Life's current and future business, as well as a step-based, continuous, and sustainable evolution towards all-flash storage.

Although the multi-active data centers of China Life have strict requirements on storage technologies, the Huawei OceanStor 18000 effectively meets the data synchronization requirements of China Life's remote multi-active data centers with its secure, efficient, and reliable technical capabilities. It can efficiently, safely, and reliably meet China Life's business production needs at any time.

Huawei's innovative technologies have also been tried and tested in many different scenarios under differing levels of extreme requirements. Huawei's all-flash technologies enable Huawei's high-end storage to withstand mounting pressure from China Life's services, demonstrating superb performance, high quality, and unwavering stability ever since its deployment in 2015. Huawei cooperated with China Life to streamline the equipment, system, development, and O&M departments to achieve an end-to-end solution design closer to the customer's real-life situation. This solution also helps customize roadmaps for future technical architectures, helping to accelerate the implementation of new products and technologies. Huawei will make dreams for new technologies a reality and will help productivity move at the fastest possible speed, paving the way for China Life to revel in their business success caused by technical innovation. ⚡



[Get More Tips](#)

Video: ACESI Group CTO Praises Huawei OceanStor Dorado V3 All-Flash Storage

Stefano Sandri CTO of MAPEI

“

We decided to go for a new all-flash solution two years ago, and we found Huawei is a good partner for implementing the project with good technology that was serving for us for leveraging our business operation.

We finished the project of full migration of infrastructure in two months by very quick and smooth operation without any major issues, and also during our follow-up of the project, we will get immediate support and strong technology from Huawei.

”

BYD Auto Becomes a Leader in New Energy with All-Flash Storage

By Wang Jiaxin from Huawei

On October 11 2017, the Low Carbon Vehicle Partnership (LowCVP) announced BYD Auto as the 2017 Low Carbon Champion. BYD Auto won the "Low Carbon Heavy Duty Vehicle of the Year Award" and, in partnership with Transport for London (TfL) and Go-Ahead London, won the "Grand Prix" - or winner-of-winners-award - for launching London's first all-electric bus routes (the 507 and 521). These awards specifically noted BYD's incredible efficiency in terms of delivering the vehicles,

operational capacity, and building infrastructure for the London project.

Founded in 1995, BYD engages in IT-related and automobile business for traditional fuel-powered vehicles and vehicles powered by new types of energy sources. With 30 production centers established around the globe, BYD is the only vendor worldwide that owns power battery and vehicle technologies in the new energy field.


Stretched ERP System of Unclear Causes

The ERP system is a core system that supports BYD's three product lines (IT, automobile, and new energy). The original legacy ERP system could not deliver the necessary stability and performance to meet growing business expansion requirements, for example, the number of employees more than doubling, factory orders growing by 10% every quarter (reaching 185,000), and concurrent connections to the ERP system exceeding 60,000. To make matters worse, the ERP system took more than 1 hour to query financial cost control-related reports and the system would occasionally breakdown during monthly settlements, severely affecting the operating of the production line.

Troubleshooting with Huawei

As a leading global ICT solutions provider, Huawei offers end-to-end solutions that cover networks, cloud computing, servers, databases, and storage. BYD decided to team up with Huawei, a company with extensive experience in IT solutions, to analyze server memory, CPU utilization, networks, and disk response times. The team discovered that over 200 GB of storage data that was bottlenecking service access performance was generated in the ERP's production system every month. In addition, the team found that the I/O latency was reaching 50 ms on the live network.


The original and outdated HDD-based storage devices in the ERP system delivered poor access performance. BYD wanted to improve its business performance by 100% without changing the existing ERP system architecture. Smoothly migrating applications without disrupting application systems was a challenge for Huawei.



Huawei OceanStor Dorado leads to a three-fold increase in BYD ERP performance

Based on the requirements and analysis results of BYD's live network, Huawei provided the OceanStor Dorado all-flash acceleration solution which leverages all-flash storage. This allows the seamless replacement of existing HDDs, accelerating read and write I/Os of the ERP and BW systems and improving system efficiency. This solution does not change the architecture, database, host, ERP software, or O&M mode of the live network. The minimal change translates to minimal risks.

To prevent system upgrades from affecting services, the Huawei data migration team provided a set of data mirroring migration solutions by using host volume management software. These solutions migrated data from legacy storage to OceanStor Dorado all-flash storage without affecting ERP services. In addition, Huawei recommended that BYD run development and test systems, such as sandbox, on the legacy storage system to ensure return on investment (ROI).



ERP system networking

After the ERP system switches to Huawei OceanStor Dorado, the latency is reduced to less than 1 ms. With the query condition design and upper-layer platform of ERP, the data extraction time for BW data warehouses is reduced by

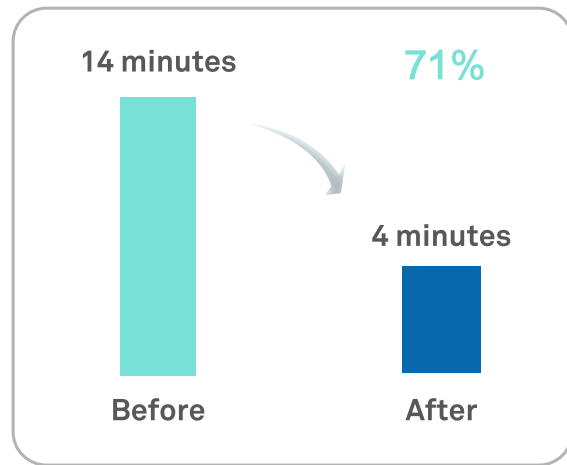

66%. Additionally, the report output speed of the MM, PP, SD, and FICO modules sees an increase of more than 200%. The time required to query material documents is reduced from 14 minutes to 4 minutes.



Time to extract BW data is reduced from 3.5 to 1.2 hours



Time to query material documents is reduced from 14 to 4 minutes



ERP system improvements

Legacy HDDs are configured with RAID 10 groups and both SAS disks and SSDs are deployed. To meet service capacity and performance requirements under these conditions, 40 U cabinets are required. Huawei OceanStor Dorado all-flash storage uses RAID 6 groups and all-SSD configurations, requiring only 20 U cabinets. This saves 50% on cabinet space and O&M manpower costs, while significantly reducing electricity and cooling costs. In addition, the estimated three-year OPEX can be slashed by more than 70%, ensuring competitive advantages for all-flash storage compared to traditional HDD-based storage.

Huawei OceanStor Dorado has now been running stably for more than half a year in the ERP system on BYD's live network. BYD has seen noticeable improvements to service performance and reductions in maintenance costs. The IT system reconstruction has surpassed all of BYD's expectations.

Inspired by the successful reconstruction of the ERP system, BYD once again chose Huawei's Dorado all-flash storage to upgrade and revamp its DMS system. Dorado's superb performance and robust stability have once again impressed BYD. It is expected that BYD and Huawei will further their cooperation through the smart manufacturing of "Made in China 2025". ⚡



[Get More Tips](#)

Video: The CTO of MAPEI discusses how Huawei helps to improve their performance



Yahoo! Japan Delivers More Enjoyable Shopping Experience with Huawei OceanStor



Active-Active in a distance as long as 180 KM



7x24

Real-time



5-fold

Increase in concurrent connections




40%

Lower TCO



Telecom Italia Uses Huawei All-Flash Storage to Consolidate Complicated Traditional Storage



72% OPEX Savings



Power Consumption



Footprint



O&M Personnel

The All-flash "Renaissance" of Telecom Italia

By Wang Jiaxin from Huawei

TIM, the largest carrier in Italy, has more than 50 million customers worldwide. As competition intensifies in the Internet and telecom market, TIM has been facing great challenges in regard to revenue growth in recent years. However, the operational cost of TIM's IT infrastructure has also continued increasing over the past few years, while the OPEX has increased by more than 10%. Therefore, it has become TIM's most urgent requirement to reduce the TCO of IT infrastructures.

Imminent Transformation

After more than 100 years of endeavor, TIM has become engaged in numerous areas of business such as landline phones, mobile phones, the Internet, multimedia content, and TV. TIM has more than 35.3 million mobile phone users, 29.2 of which are in Brazil, 22.8 million landline users in Italy, and 10.7 million broadband users in Europe.

With the use of new technologies such as 4G and 5G and the continuous development of the mobile Internet, service volume keeps increasing. According to TIM's statistics, the requirements for storage space increase by more than 10% each year. TIM is facing the pressing issue of increasing data storage requirements while continuously reducing the budget.

A Perfect Attempt in 2016

Compared with other elements of IT infrastructure, storage devices are more difficult to replace. Therefore, storage life cycle management and TCO management are the top priority of TIM.

At the beginning of 2016, 66 sets of storage devices in seven TIM data centers, distributed throughout Italy, were about to reach the EOM and needed to be replaced. After communicating with multiple IT vendors about its future storage evolution strategy, TIM became interested in Huawei's efficient solution.

An All-Flash Solution for Efficient Service Systems

The existing storage devices on the TIM live network are traditional storage devices. A large amount of over commitment capacity and system scale must be configured to meet the current service requirements, resulting in high TCO. Even so, some performance-sensitive services still fail to meet service requirements and are subject to the high maintenance costs associated with traditional storage. As a result, the TIM storage system cannot fulfill the Service Level Agreement (SLA).


With the maturity of flash storage, TIM wants to replace traditional storage devices with flash storage. However, most all-flash vendors lack comprehensive data service capabilities, making large-scale

deployment difficult on TIM's live networks.

Huawei's all-flash array, OceanStor Dorado V3, boasts high performance and provides high reliability and data service capabilities that are similar to traditional storage. This makes it perfect for meeting the requirements of the TIM service system.

In most of TIM's business scenarios (database, VDI, and server virtualization), OceanStor Dorado V3 provides a data reduction guarantee of 3:1. Compared with traditional storage, the OceanStor Dorado V3 requires much lower initial purchase costs. This is a key reason for TIM to choose Huawei.





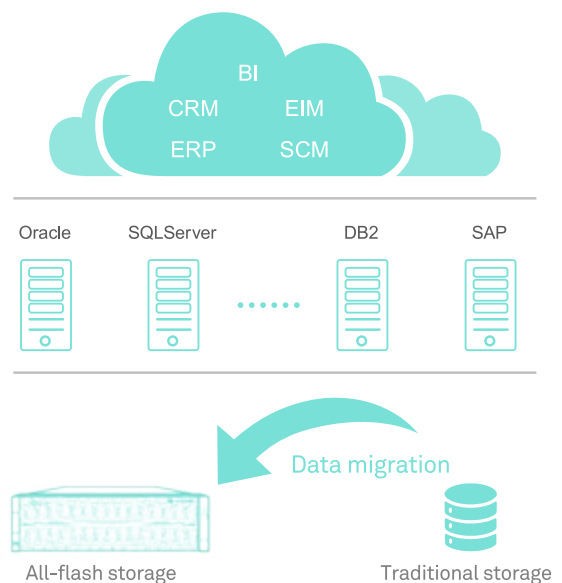
Energized Operations at Data Centers with Smooth Data Migration

Huawei has accrued rich experience in global markets over the past regarding how to aggregate and migrate storage. In addition, Huawei all-flash storage supports heterogeneous virtualization and can smoothly migrate services from traditional storage platforms to new all-flash storage platforms.

Doubled service processing capabilities than traditional storage

TIM performed strict tests on Huawei OceanStor Dorado V3 that were based on the service load on the live network. The test results showed that OceanStor Dorado V3 could improve the TIM live network performance by 54.7%, fully addressing the service challenges of the CRM system regarding business analysis and billing services.

Huawei OceanStor Dorado V3 all-flash storage was used to replace the live network storage in seven data centers distributed throughout Italy. OceanStor Dorado V3 provides excellent and stable performance and supports inline deduplication and compression. This helps TIM reduce cabinet occupation, save power, save space, and save labor costs, as well as reducing OPEX by 72%. When compared with traditional storage, the performance improves by 100%.




Huawei Storage for Rapid IT Transformation

Huawei OceanStor Dorado V3 is selected to carry the key services of TIM, and the services on the live network are integrated.

- Performance is improved by 100%: Service efficiency is doubled.
- Operation costs are reduced by 72%: The 3:1 data reduction guarantee saves the TCO in E2E mode.
- Unique gateway-free active-active all-flash storage: Stable running of services 24/7.

Huawei All-Flash Storage Ensures Stable Online Services

For carriers, the high availability of core services is the most basic requirement. TIM has been using traditional high-end storage to carry core services for many years.



Gateway-free all-flash active-active

Huawei OceanStor Dorado V3, unlike other all-flash storage devices, boasts array-based, gateway-free, active-active capabilities. In addition, Huawei OceanStor Dorado V3 can provide consistently low latency to meet carriers' core service requirements while ensuring high reliability and performance. ⚡



[Get More Tips](#)

Video: Huawei OceanStor Dorado All-Flash Storage Introduction

Insight into the Future of Flash Technologies

By Tu Weixin and Wang Xuesong from Huawei

The world is undergoing rapid and accelerating changes.

The information industry, driven by various innovative technologies, stands at the center of changes across the globe. The storage system, as a pillar of the information industry, is also experiencing tremendous changes. The past decade has seen the rapid rise of flash media, incredible capacity expansion of hard disk from hundreds of GB to tens of TB, and continuous iterations of new network interconnection technologies. In addition, we have seen numerous new concepts arise (such as data tiering, deduplication, RAID 2.0, distributed EC, FCoE, and NVMe), and new storage forms (converged storage, CI, HCI, public

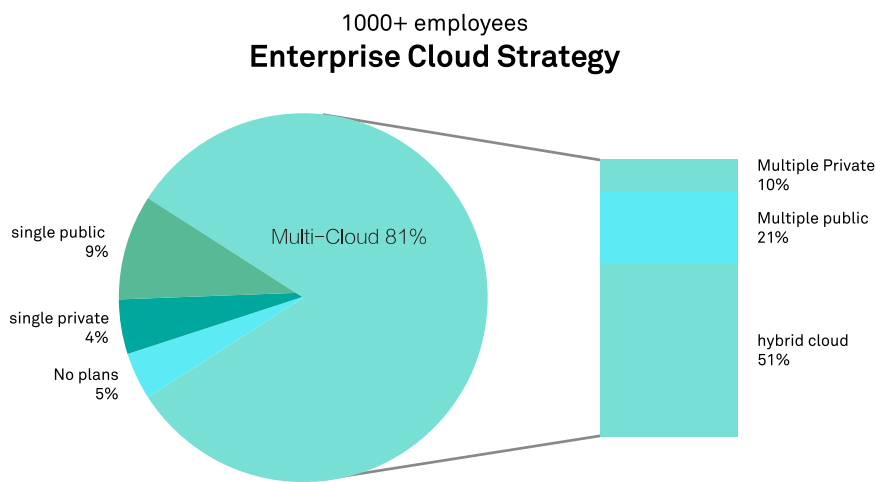
cloud storage, SDS, and object storage). Although some of these changes have been fruitful, such as flash memory; some have failed, such as FCoE; some are still developing, such as secondary storage; and some technologies are old yet remain strong, such as tapes. There is no doubt that changes are now emerging at a rate faster than they did in the last decade, not to mention in the previous era.

How can Huawei, as an industry-leading ICT infrastructure provider, find the best combination of technologies to allow customers to cope with these changes?

Changes in Data Center Architecture

Despite all these changes, the requirements for data centers have remained the same. Customers require lower TCO, higher performance, larger capacity, and more secure and reliable data centers. With the development of technologies, the architecture of data centers is constantly evolving to meet these requirements.

Huawei believes that cloud-enabled data centers will be the most significant trend over the next few years. More specifically, hybrid clouds will be widely deployed in data centers. According to the RightScale 2018 State of the Cloud Report, 81% enterprises adopt the multi-cloud strategy, among which the hybrid cloud accounts for 51%.




State of the Cloud Report by RightScale, 2018

There are plenty of discussions and analysis throughout the industry regarding the benefits of cloud-enabled data centers. We will focus on why the hybrid cloud is the most important change for these data centers. Huawei believes that the hybrid cloud architecture will be the most important infrastructure for medium- and large-sized enterprises over the next five to ten years. In hybrid cloud architecture, enterprises usually divide data center workloads into two types: Core workloads and non-core or innovative workloads. These two types of workloads adopt different IT strategies.

Core workloads include ERP, core transaction systems, such as category-A banking services, and core product design systems. These services require high performance and reliability, usually have requirements related to intellectual property rights and security compliance, and have their data stored in local data centers. Non-core workloads include backup, archiving, and innovative services such as IoT and AI, which require large capacity, low price, high scalability, flexibility, and wide accessibility. These services will be deployed in cloud data centers, or in one or more public clouds.

This hybrid deployment mode can better meet enterprises' ultimate requirements for data centers. For example, an automobile group holding company is mainly engaged in the R&D, production, and sales of auto components, and this company has more than 200 factories with an annual turnover of hundreds of billions of dollars, and the HQ provides cloud services and O&M management. Each edge serves as an OpenStack cluster and is managed by ManageOne at HQ, while the data centers are connected to the public cloud. This allows any data belonging to non-core workloads or the production centers to be backed up to the cloud. The hybrid cloud infrastructure, composed of Huawei public cloud and private cloud, then meets requirements for both scalability and flexibility.

Hybrid cloud architecture faces great challenges for storage. How can free data flow and unified management be implemented among multiple sites? How can vendors, like Huawei (Huawei Public Cloud), with both on-premises and off-premises cloud capabilities make full use of the advantages of collaboration? These issues will greatly affect the development of future storage systems. Huawei believes that there will be no independent storage systems in the future and instead all storage systems will be integrated into the cloud. Therefore, Huawei's future on-premise storage and off-premise storage will use the same architecture to implement free data flow and unified management without differentiating data locations, as well as unified bills. This will allow Huawei to provide customers with the most complete and smooth hybrid cloud architecture experience possible and help them to accelerate the deployment of cloud-enabled data centers.




Industry Innovative Applications

Fantastic innovations are occurring in all industries, but these innovations are also posing new challenges for the entire IT system, especially storage systems. Taking the security industry as an example, there are ongoing innovations based on AI technologies, such as facial recognition, trace tracking, and identification of abnormal objects and events. These innovations require high-speed data lake storage which can provide multiple access modes, such as video, image, MPPDB, and block, as well as ultra-large capacity and data lifecycle management (hot, warm, and cold data). The innovations also require storage to have local data processing capabilities (NDP-Near Data Processing) with the assistance of GPU acceleration. However, the traditional security storage system features a large capacity, low price, and large bandwidth, making the disparity obvious to see.

The rise of Internet finance in the financial industry has led to a requirement for massive access methods and has caused service uncertainties, posing new demands on storage flexibility. For example, during the 2017, Double 11 shopping festival in China, online transactions peaked at 256,000 per second, doubling the amount recorded in 2016. In addition, database processing peaked at 42 million times per second, which was greater than the processing volume recorded for the entire day before the festival. Therefore, the underlying IT infrastructure was under great pressure to cope with this massive increase.


The same situation can be seen in the mobile gaming industry. Before a game is launched, it is difficult to predict the expansion speed of users or how quickly game users will drop away. Take Flappy Bird as an example, which was released in May 2013, and went viral at the beginning of 2014. It quickly reached millions of daily downloads, and the increase in users could not be predicted.



Explosive growth of Flappy Bird in AppStore (Data source: <http://zachwill.com/>)



In social media, for example Weibo, a sudden trending event usually generates heavy traffic. For example, in October 2017, the breakdown of one of Sina Weibo's servers was caused by heavy traffic over a short period due to a celebrity declaring their love on Weibo. According to the officially released results, the heavy traffic recorded consisted of 462,884 reposts, 986,409 comments, and 2,566,617 likes.



The microblog was posted at 12:00 on October 8, and the reposts and comments reached peak at 12:10 on October 8. The peaks of the repost and comment reached 34,856 and 5,996 respectively. Later on, the spread speed declined gradually.

Microblog traffic data

In the manufacturing industry, with the innovations of Industry 4.0 and flexible manufacturing, massive product configurations and unpredictable product sales are being generated. These innovations require highly flexible systems, and Huawei believes that only distributed storage and cloud storage can meet the requirements.


Other recent important innovations are ultra-high-speed data analysis services, such as credit card fraud analysis in the financial industry. Multiple industries have the same innovations, such as the Internet of Vehicles (collision analysis and prevention), traffic control (real-time road condition analysis), and manufacturing (EDA). These services require storage that has unparalleled speed.

Generally speaking, we believe that industry innovations have introduced some new storage requirements, such as the intelligence trend of data lakes, distributed storage, cloud storage, and ultra-high-speed storage. Huawei storage defines the layout of future storage products based on these requirements and provides the best choices for customers.

Continuous Media Progress

There is no doubt that media remains the most important technology change in technology for storage systems.

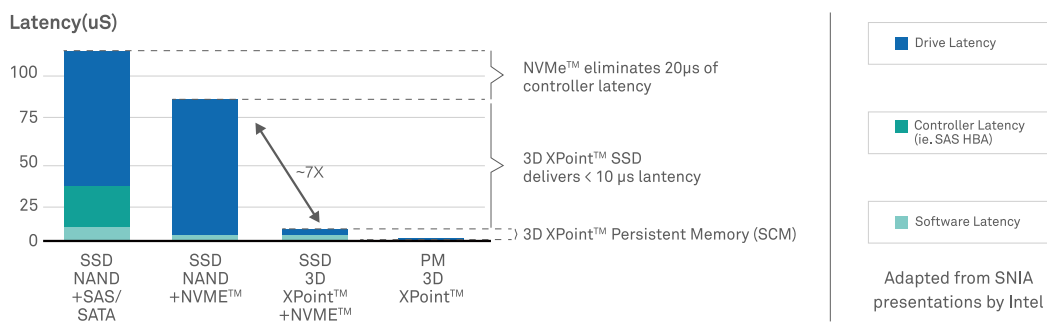
Over the past few years, we have witnessed the rise of all-flash array storage (AFA). As the price of flash storage continues to decline, we can expect AFA to continuously expand its share in data centers.



Price trend of different storage media

However, does the current AFA provide optimal flash capabilities? Huawei believes that storage vendors must still make long-term efforts in order to provide the best possible flash capabilities.

In addition, the computing capabilities of CPUs cannot meet the development speed of media. From HDDs to SSDs, performance is improved by hundreds or even thousands of times. The new-generation SCM has made a great leap in performance, reducing media latency to within μs .



Evolutionary and performance improvements of storage media

The CPU performance of each generation improves by only 20-30%, meaning the CPU becomes the bottleneck for AFA system performance.

One way to solve this problem in the industry is to use more CPUs in order to increase the Core/SSD configuration ratio. In modern AFAs, using OceanStor Dorado6000 as an example, the Core/SSD configuration ratio can reach 3.7:1. However, more CPUs cause higher power consumption, dissipation, and costs, therefore failing to achieve sustainable development. To tackle this, Huawei considers adopting CPUs with lower power consumption and costs and using more cores (multi-core). This further improves the Core/SSD ratio.

Another solution is to use a smart disk enclosure to distribute some media related functions to disks/enclosures (NDP). These functions include compression, media reliability scanning, reconstruction, and balancing. Offloading some computing capabilities from the CPU ensures reduced CPU occupation, lower network consumption and better performance.

Continuous Progress in Data Protection

Huawei will launch corresponding technologies and support the cascading and expansion of back-end JBoF through NVMeOF. Another solution is to optimize software and improve the parallelism and efficiency of computing through lock-free mode, and run to complete mode. Intel's SPDK represents an attempt at accomplishing this. Huawei OceanStor Dorado AFA was designed with the goal of enabling high concurrency and efficiency in mind, and will continue to be optimized.

The mismatch seen between performance and media requirements also occurs on the network side, the cause of which is the same as that described above. We have witnessed the rapid application of NVMe networks in storage, and Huawei was one of the first vendors to launch NVMe storage in the industry. Huawei will quickly promote the further application of the all-NVMe network (NVMe over Fabric) in storage systems and data centers.

Another major trend is that flash media is rapidly developing. One trend is to provide larger storage, while the other is to provide faster storage. In terms of larger storage,

Users have a constant need for improved data security, and many innovations have been made in this area. For example, Cloud Backup/DR ensures lower costs with the help of cloud, and the Copy Data Management (CDM) emphasizes the management, migration, and reuse of data copies, improving traditional backup system capabilities. Ultra-high density snapshots (completed in seconds) compress intervals, increase the number of snapshots, and provide similar

the emergence of QLC further extends the application scenario of flash memory and strengthens the justification for replacing HDDs. However, the reliability and performance challenges introduced by QLC determine that the storage system must optimize the QLC (similar to the SMR in storage) and control its application scenarios. Warm storage is one of the QLC application scenarios.

In terms of faster storage, SCM represents a major change. The byte addressable mode of SCM will fundamentally change the architecture design of storage systems. The byte converts a block-based design to a byte-based design, introducing new requirements for the new-generation data structure, transaction mechanism, and database design based on the SCM. Huawei has been conducting in-depth research and will release ultra-high-speed storage based on SCM in the future.

CDP capabilities (traditional CDP usually requires dedicated devices) on the storage system, improving the level of data protection. The 3DC DR can be upgraded to 4DC DR or even cross-cloud DR.

These innovations demonstrate the importance of user data. Huawei will continue to enhance its data protection capabilities and apply these to AFA, distributed storage, and cloud storage.

Broader Vision for New Technical Opportunities

In addition to the preceding changes and innovations, there are new technical developments that may affect future storage, such as AI and block chain.

AI has been a hot topic in the industry in recent years, and storage is also making use of AI. Multiple vendors, including Huawei, have deployed AI-based storage O&M systems to predict faults, locate problems, analyze performance and capacity trends, and provide the system health status, greatly reducing O&M costs. However, this is just the beginning. AI will be further applied to storage, such as for intelligent identification service workloads and the dynamic adjustment of storage system parameters (cache prefetch parameters, disk selection algorithms, and

deduplication algorithms), providing optimal system capabilities. In addition, the intelligent algorithm automatically optimizes the data distribution mode in distributed storage systems, reducing network consumption and balancing the system. Huawei will continue its progress in AI applications.

Compared with AI, the deployment of block chain in storage is still in the very early stages. Block chain is characterized by non-repudiation, and therefore a possible application mode is to use this technology to implement trusted storage systems, where all user logs, operation records, and even data are verified and traceable. Trusted storage systems, such as contract storage, have a competitive edge in areas that have high requirements for data security.

Conclusion

Storage is facing a number of tremendous changes, such as hybrid cloud, intelligence, data lakes, new media, new data protection technologies, and new technology opportunities. Huawei storage will continue to uphold its belief in "Data on Demand" and continuously innovate to provide customers with higher-quality, more advanced storage systems and better data management capabilities in this new era. ⚡



Huawei's Next-Gen Hybrid Flash Storage OceanStor V5

Multi-level Convergence

Gateway-free converged active-active solution for SAN & NAS, 99.9999% HA

Excellent Performance

In-depth flash-oriented optimizations, supporting scale-out to 6 million IOPS with 1 ms latency

Intelligent Services

Supporting intelligent cloud services throughout the entire lifecycle, supporting cloud transformation



AFA READY PLAYER ONE

AUGUST 2018

Creative Team

Sponsor	Huawei Storage Marketing Support Department
Consultant	Huang Jin, Meng Guangbin, Ge Qiancheng
Editor-in-Chief	Li Mingxiao
Editorial Reviewer	Wang Tao
Editor & Writer	Wang Zhen, Zhou Ling, Lan Bo, Chen Yong Liu Lin, Wang Xuesong, Tu Weixin Zhao Yanling, Qiao Lina, Wang Jiaxin Wang Taiping, Chi Haixia, Huang Lili Li Wuxian, Xu Yonghua, Zhang Weiguo
Foreign-Language Editor	Liu Yuhan
Graphic Designer	Xue Di

Get Ready For The Fantastic AFA Journey



Email: OceanStor@huawei.com



OceanStor Dorado V3 All Flash Storage Lightning Fast, Rock Solid

- 3x increase in application performance
- 99.9999% field-proven availability
- 75% OPEX savings

For more information, please visit e.huawei.com

Leading New ICT, The Road to Digital Transformation