

## Technical Review

# Huawei OceanStor Pacific Next-gen HPDA Storage

**Date:** June 2021 **Author:** Tony Palmer, Senior Validation Analyst

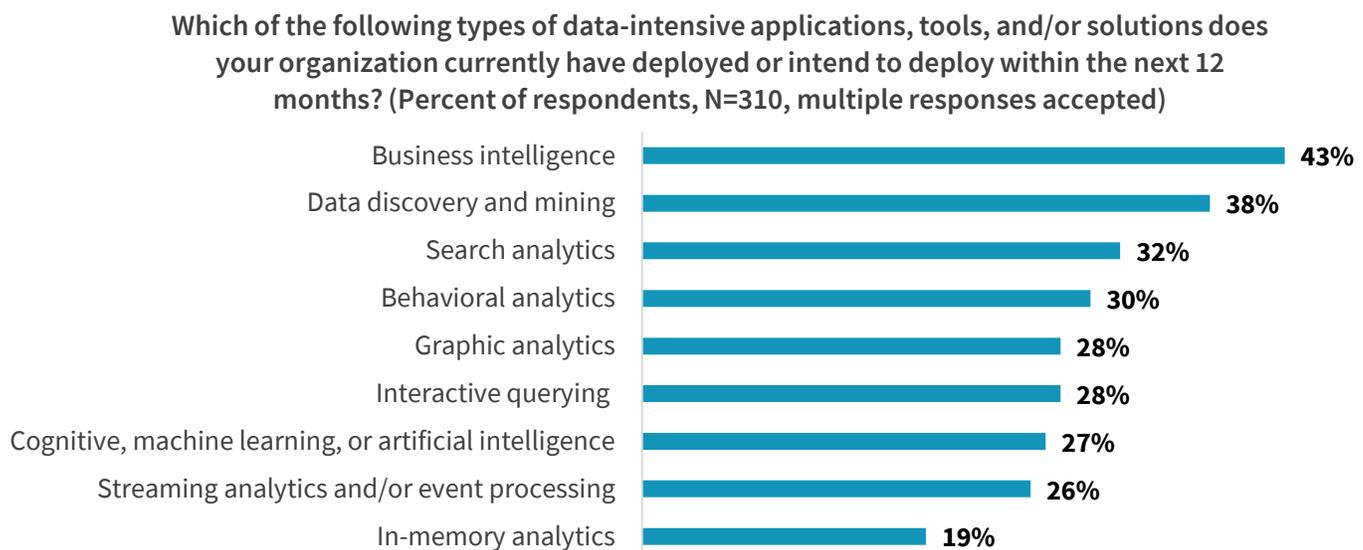
## Abstract

This ESG Technical Review documents hands-on performance and function testing of Huawei OceanStor Pacific storage for high performance data analytics (HPDA) and presents the findings of a five-year TCO analysis highlighting the economic benefits of Huawei OceanStor Pacific when compared with storage systems from other major vendors.

## The Challenges

High-performance computing (HPC) was originally conceived to perform computation-intensive actions on relatively small amounts of data. New technologies and applications are driving rapid growth of data volumes. Level 4 autonomous driving data sets can grow to exabyte levels and a single modern genome sequencer can generate 6TB of data per day, for example. Iterative analyses of these data sets are used to continuously develop and refine the algorithms that autonomous vehicles depend on and provide lifesaving medical information about patients and diseases. In data-intensive analysis like this, storage systems can become the bottleneck, restricting the efficiency of data analysis. All these data sources present the opportunity to add significant business value using HPDA (see Figure 1). This does not come without challenges; the volume of data has been increasing at an accelerating pace for a long time. In a recent survey, one in five (21%) organizations reported that they are managing 10PB of data or more, and 5% are managing more than 50PB. This explosion in the volume of data makes it difficult to manage, safely store, securely analyze, and generate robust insights.<sup>1</sup>

**Figure 1. Data-intensive Applications, Tools, and Solutions Deployed**



Source: Enterprise Strategy Group

Further, 63% of respondents to a separate survey indicated that spending on AI/ML in 2021 would increase over the prior year, which will create and use even more data.<sup>2</sup>

<sup>1</sup> Source: ESG Master Survey Results, [The State of Data Analytics](#), August 2019.

<sup>2</sup> Source: ESG Research Report, [2021 Technology Spending Intentions Survey](#), January 2021.

The demand for reliable availability of data is even more urgent. Organizations are telling ESG that data has become a core asset and data storage technology has become strategic: nearly half of organizations polled by ESG stated that data either is their business (23%), or both is their business and supports their business (26%).<sup>3</sup> In this same survey, 71% of organizations tell ESG that data storage technology is strategic—it’s critical to their core applications and business processes and can provide competitive advantage.<sup>4</sup>

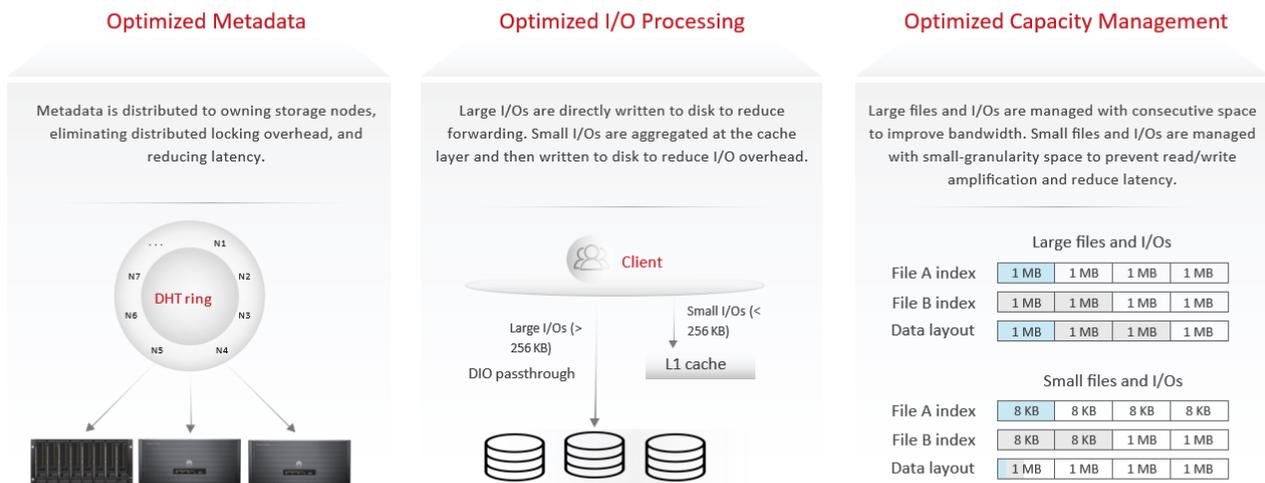
Data-intensive HPC brings new challenges to storage. It requires more economical and reliable storage that can effectively cope with diverse workloads. From the traditional single application/workload to complex mixed workloads driven by the refinement of HPC-business process links and the integration of multiple application scenarios, business workloads have become more complex. Looking at seismic exploration, for example, the reservoir simulation stage is dominated by large files and requires high bandwidth, while the reservoir interpretation stage is dominated by smaller—but still fairly large—files and requires high IOPS.

As larger amounts of both structured and unstructured data are generated, collected, and analyzed, organizations have striven to build out their storage infrastructure more efficiently. In addition to the examples listed above, organizations are employing data-intensive applications such as those used for AI/ML, financial modeling, business data analytics, post-production editing, and the internet of things (IoT), which all utilize unstructured data using multiple protocols. Organizations seek out solutions that will provide fast and consistent storage performance for reads, writes, and metadata operations for multiple applications. These solutions are key to application performance at scale, as their storage grows and supports the large amount of application processing required so that they can process data and extract value without unnecessary delay.

### The Solution: Huawei OceanStor Pacific Next-gen HPDA Storage

Huawei OceanStor Pacific is a distributed scale-out storage system designed to support organizations’ business- and mission-critical HPDA workloads. The Huawei OceanStor Pacific parallel file system is designed to optimize I/O metadata placement, keeping it close to the data on the nodes that own it. I/O processing and capacity management are likewise optimized using large I/O passthrough to bypass cache when appropriate and a granularity-specific layout that increases bandwidth for large-block I/O streams while decreasing latency and I/O amplification for small-block I/O.

Figure 2. Huawei OceanStor Pacific Next-gen Parallel File System



Source: Enterprise Strategy Group

<sup>3</sup> Source: ESG Research Report, [Data Storage Trends in an Increasingly Hybrid Cloud World](#), March 2020.

<sup>4</sup> Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019.

OceanStor Pacific is engineered to provide the performance and flexible access required by a variety of data-intensive scenarios with disparate access requirements, including HPC, AI/ML, big data analytics, large-scale virtualization, content repositories, seismic analysis, life sciences, finance, and any app that requires the ability to store huge volumes of data and provide high-performance, multi-protocol access.

To handle the massive amounts of data produced by HPC applications, Huawei has designed next-gen, high-density hardware architectures for OceanStor Pacific. Extremely high-density design has always been a challenge in the industry as a whole; Huawei is addressing this challenge with innovative design choices, including Huawei-developed half-palm-size NVMe SSDs to decrease the cross section by 65%. Huawei’s design includes advanced heat dissipation materials, strategic location of fans, and a new structural design designed to improve cooling efficiency by 30%. Huawei has implemented elastic erasure coding (EC) and end-to-end data integrity field (DIF), which they state helps them maintain a disk utilization rate of up to 91.6%.

Huawei has also made changes to field replaceable units (FRUs), bidirectional drawer slides, and tank chain techniques to increase the system maintenance efficiency.

Huawei OceanStor Pacific offers two high-density hardware architectures: the OceanStor Pacific 9950 for high-density performance and the OceanStor Pacific 9550 for high-density capacity (Figure 3).

**Figure 3. Huawei OceanStor Pacific Next-gen Models**



Source: Enterprise Strategy Group

The OceanStor Pacific 9950 is a 5U chassis that supports up to 8 storage nodes and 80 NVMe SSDs, delivering a maximum of 160 GB/s bandwidth and 2 million IOPS. The OceanStor Pacific 9550 is a dual-node 5U chassis that supports up to 120 3.5-inch SATA disks, delivering over 1.6PB of raw capacity in just 5U.

### ESG Tested

ESG performed hands-on testing and validation of Huawei OceanStor Pacific. Testing was designed to validate the performance, reliability, data management, and TCO of the OceanStor Pacific storage platform with a focus on delivering high levels of predictable performance across multiple protocols for data-intensive applications. Finally, a five-year TCO analysis was performed. It is important to note that the performance tests were not designed to obtain maximum performance of an OceanStor Pacific configuration. All test results were obtained with small clusters dedicated to the workloads that were being tested. Most organizations deploying OceanStor Pacific leverage much larger clusters to support multiple applications and workloads and can achieve much higher performance.

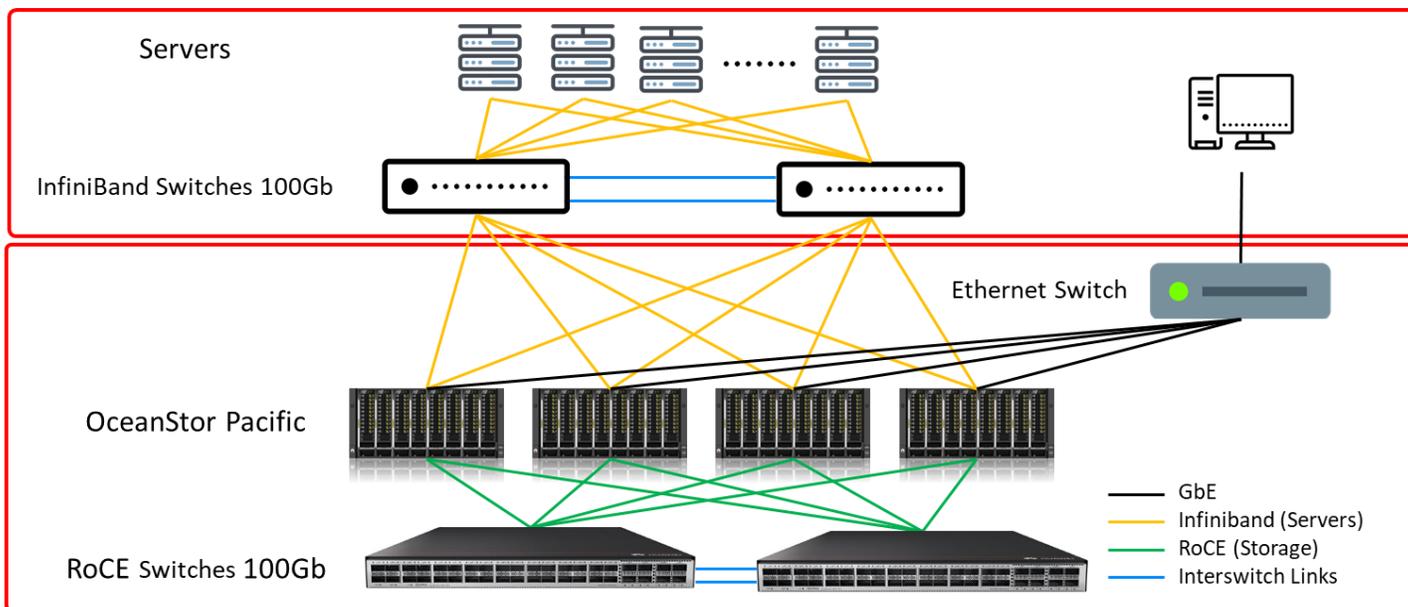
## Multi-protocol Support

File, object, and Hadoop distributed file system (HDFS) services are all commonly used by organizations in different phases of their data pipelines for HPDA scenarios like autonomous driving, precision medicine, and smart manufacturing. Traditionally, these services have either been provided by separate storage platforms—where multiple copies of data are required—or by using gateways in front of a central storage platform. Both of these scenarios are suboptimal; making multiple copies consumes time, increases complexity, and wastes storage space, while using a NAS or object gateway in front of a block storage array will compromise performance.

In contrast, the multi-protocol capability of OceanStor Pacific allows one copy of data to be shared using multiple protocols. OceanStor Pacific supports NFS, CIFS, HDFS, and S3 protocols. This is designed to improve analytical efficiency because data written using one protocol can be read over multiple protocols without data migration while preserving protocol semantics and providing consistent performance.

ESG analyzed OceanStor Pacific in a multi-protocol test environment to validate semantic integrity, performance, and advanced functionality like snapshots, quotas, QoS, object storage versioning, and object versioning. Figure 4 shows a simplified version of the test environment.

**Figure 4. The Test Environment**



Source: Enterprise Strategy Group

Table 1 shows a detailed list of the components used in testing.

**Table 1. Test Environment Specifications**

Components	Description	Quantity
OceanStor Pacific 9950	Disks: 10x3.84TB NVMe SSD, RAM: 128 GB Network: 100 Gigabit InfiniBand	4
Hosts	Huawei RH2288H V5 x86, CPU: 2* Xeon Gold 6151, RAM: 256GB, OS: CentOS 7.6	11
Front-end (Server) Switches	Mellanox SB7800 Managed EDR InfiniBand Switch	2
Back-end (Storage) Switches	Huawei CloudEngine 8850-64CQ-EI	2

Source: Enterprise Strategy Group

Multi-protocol testing began with the creation of four files in a shared directory—one each using a CIFS, NFS, HDFS, and S3 client. When each file was created, the MD5 hash was recorded for verification purposes using the client that created it. The files were then examined with all four clients. Figure 5 shows the four files on the NFS client.

**Figure 5. Multi-protocol Access**

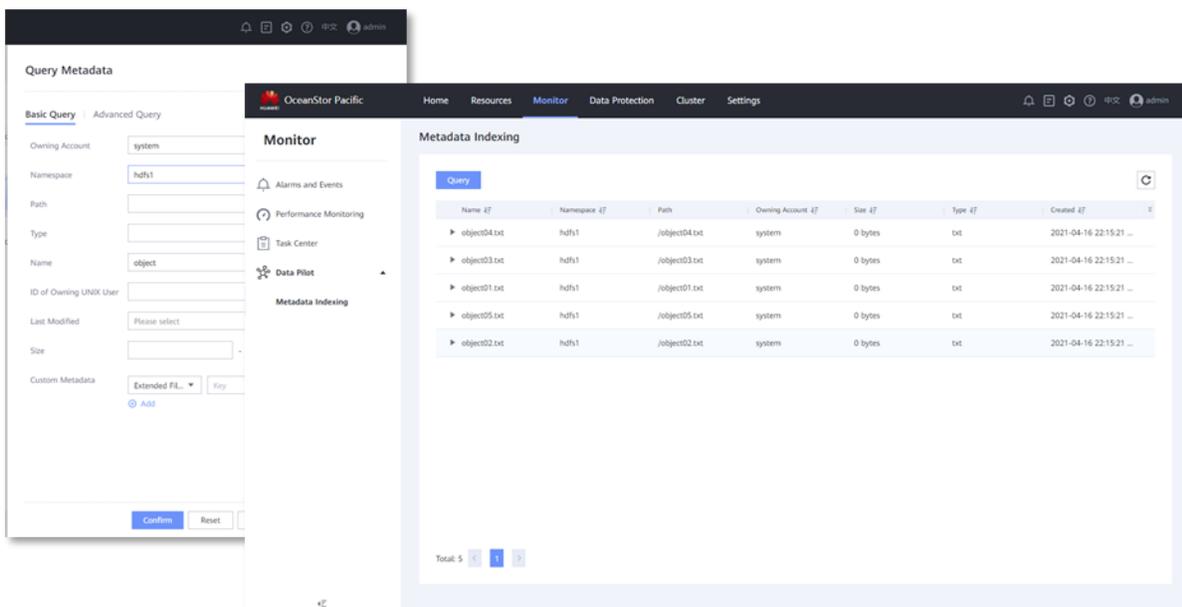
```

hdfs.huawei:/hdfs 75T 6.1T 69T 9% /mnt/hdfs
[root@host100 hdfs]# ll
total 60
drwxrwxrwx 2 kms      kafka 4096 Apr 12 15:23 app-logs
drwxr-xr-x 5 root     root  4096 Apr 12 17:21 benchmarks
drwxr-xr-x 3 kms      kafka 4096 Apr 12 15:23 datasets
drwxr-xr-x 2 kms      kafka 4096 Apr 12 15:23 datastore
drwxr-xr-x 12        970    root  4096 Apr 12 17:14 hbase
drwxrwxrwx 4 impala   kafka 4096 Apr 12 15:21 mr-history
drwxrwxrwt 2 polkitd   kafka 4096 Apr 12 15:23 spark2xJobHistory2x
drwxrwxrwt 2          968    kafka 4096 Apr 12 15:23 sparkJobHistory
drwx--x--x 2 20000 10002 4096 Apr 12 15:21 tenant
-rwxr--r-- 1 root     root   21 Apr 13 15:27 testcifs.txt
-rw-r--r-- 1 root     root   22 Apr 13 15:59 testhdfs
-rw-r--r-- 1 root     root   21 Apr 13 15:18 testnfs
-rw-r--r-- 1 root     root   23 Apr 13 16:03 testobject.txt
drwxrwxrwx 7 kms      kafka 4096 Apr 12 16:12 tmp
drwxrwxrwx 8 kms      kafka 4096 Apr 12 16:12 user
  
```

Source: Enterprise Strategy Group

ESG confirmed that the files were identical after accessing and downloading to each of the four clients, then checking the MD5 hashes against the one recorded at the source system. Next, more files were created in the shared directory by each of the four clients using a prefix matching their protocol. Search queries were able to quickly find all the files that were created across all protocols ( Figure 6).

**Figure 6. Metadata Search**

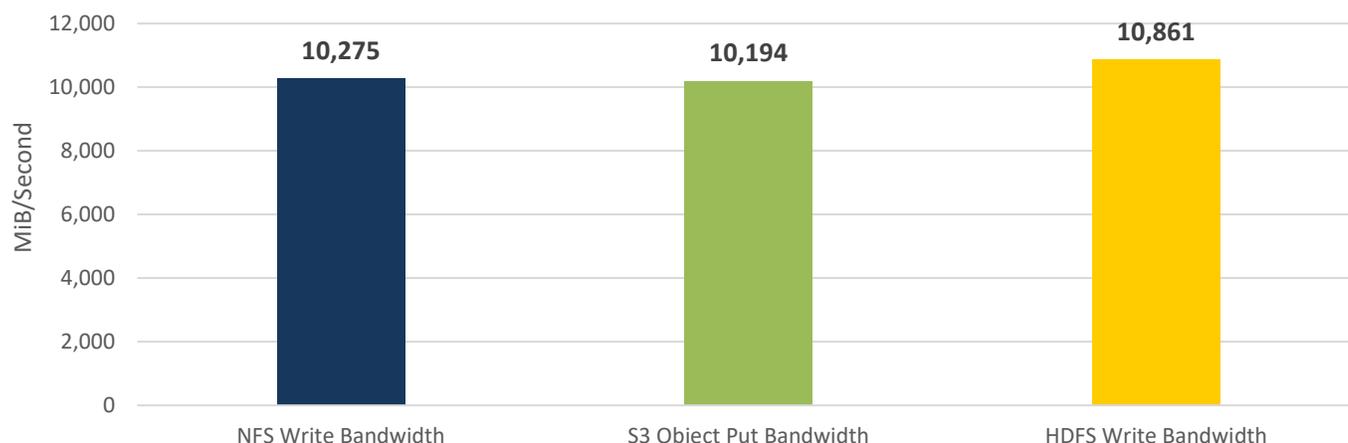


Source: Enterprise Strategy Group

ESG also looked at a number of advanced features, including snapshots, quotas, QoS, data encryption, S3 bucket policy control, and object versioning. Each of these features functioned perfectly across all protocols in our tests.

Finally, ESG evaluated Huawei OceanStor Pacific native multi-protocol performance. In these tests, hosts were configured to access the system using NFS, S3, and HDFS protocols. Each client generated read and write workloads using 4GB files and sequential I/O. As seen in Figure 7, multiple clients were able to drive more than 10,000 MiB/sec (10GB/sec) of writes accessing the same namespace and using the same data.

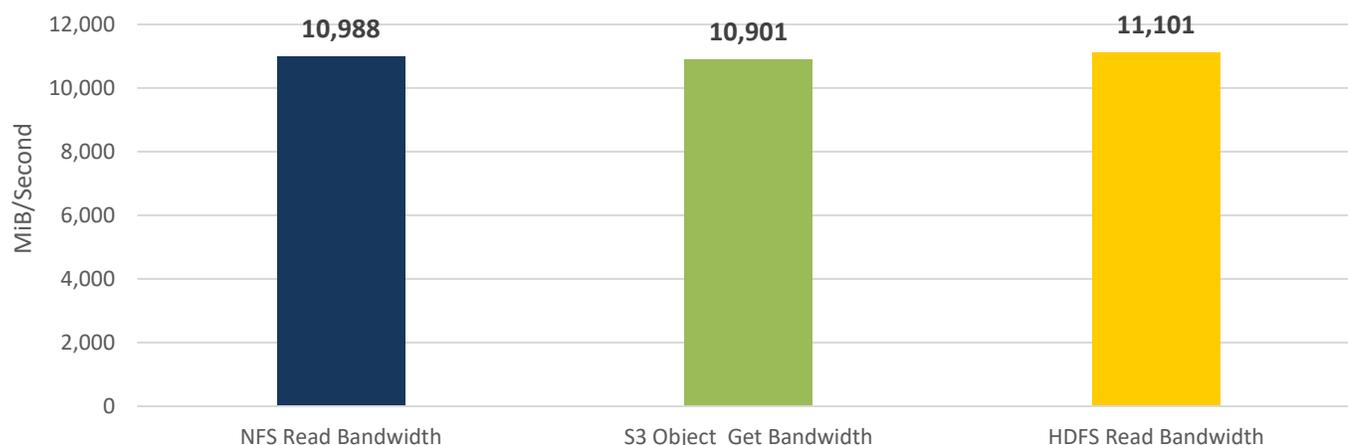
**Figure 7. Huawei OceanStor Pacific Write Bandwidth**



Source: Enterprise Strategy Group

Figure 8 shows the results from the same hosts testing read performance. In this case, the multiple clients were able to support nearly 11,000 MiB/sec (11GiB/sec).

**Figure 8. Huawei OceanStor Pacific Read Bandwidth**



Source: Enterprise Strategy Group

ESG validated that the Huawei OceanStor Pacific platform can drive consistently high performance across file and object protocols with zero loss.

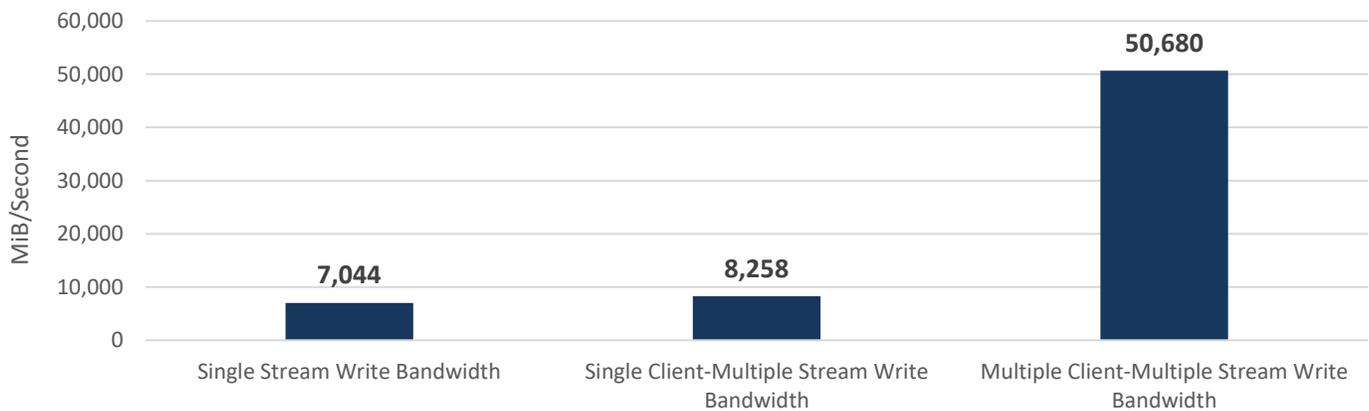
### Hybrid Workload Testing

HPC workloads are diverse, even in the same application. Seismic data processing, for example, requires high bandwidth, while interpretation of the processed data drives high IOPS. Big data and AI technologies further intensify this challenge. This means that performance bottlenecks are also diverse. Bandwidth bottlenecks can be caused by deficiencies in the network, disk, or memory. IOPS bottlenecks can be caused by insufficient CPU power or software issues like call stack depth.

The OceanStor Pacific file system uses features like metadata distribution, targeted processing of large and small I/O, and disk indexing to satisfy both high bandwidth and high IOPS requirements.

Applications that require extremely high bandwidth and use the message passing interface (MPI-IO) to support parallel I/O present a serious challenge. ESG tested the performance of the Huawei OceanStor Pacific parallel filesystem with the Huawei Distributed Parallel Client (DPC). Unlike traditional NFS clients, DPC enables a single client to concurrently access multiple storage nodes, eliminating single-client and single-stream performance bottlenecks. DPC supports MPI-IO and RDMA networks to better adapt to application ecosystems and reduce response time. DPC implements I/O-level load balancing to fully leverage storage cluster capabilities. Write bandwidth tests were run from a single client with a single thread and from a single client with multiple threads. The test was then repeated with 11 clients running multiple threads each. As seen in Figure 9, a single client was able to drive 7,044 MiB/sec with a single stream and 8,258 MiB/sec running eight streams. Eleven hosts were able to drive 50,680 Mib/sec of write throughput.

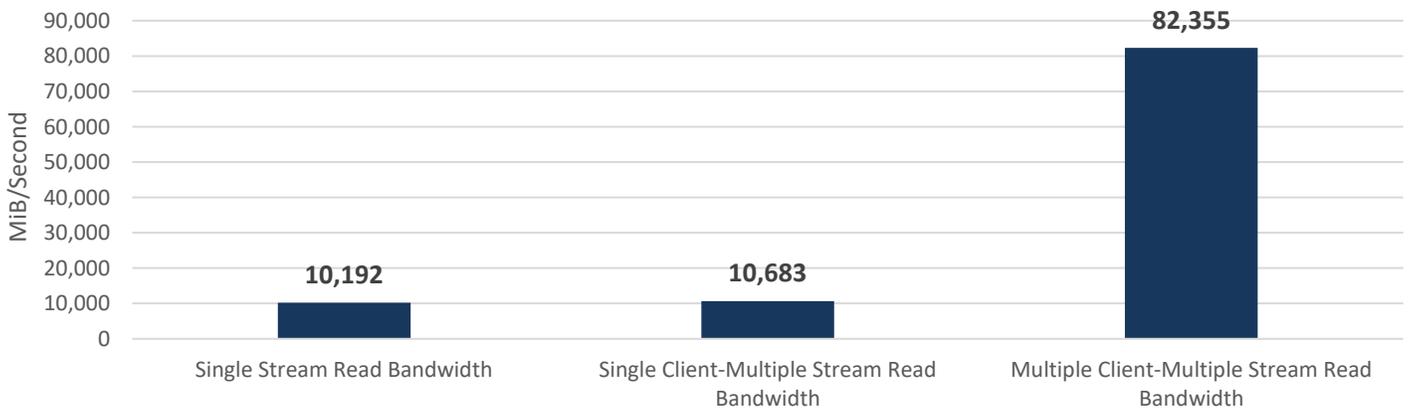
**Figure 9. Huawei OceanStor Pacific Write Bandwidth**



Source: Enterprise Strategy Group

These tests were repeated with a read workload, and the results are shown in Figure 10. Again, tests were run from a single client with a single thread and then from a single client with multiple threads. Next, the test was repeated with 11 clients running multiple threads each.

**Figure 10. Huawei OceanStor Pacific Read Bandwidth**



Source: Enterprise Strategy Group

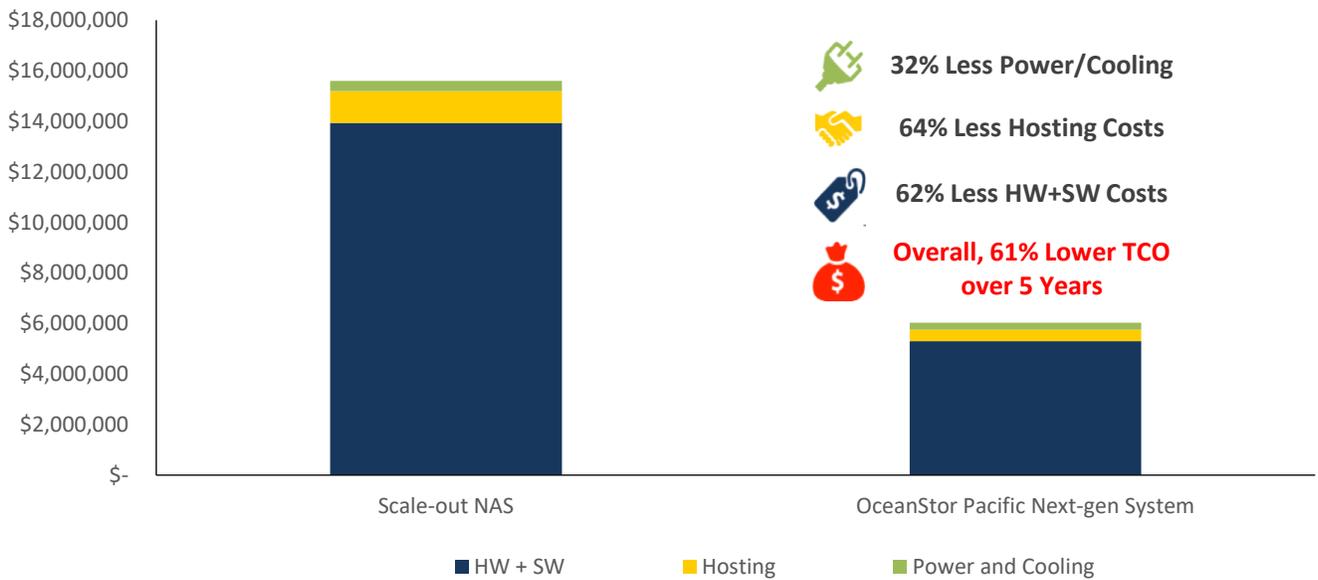
As seen in Figure 10, a single client was able to drive 10,192 MiB/sec with a single stream and 10,683 MiB/sec running eight streams. Eleven hosts were able to drive 82,355 Mib/sec of read throughput.

### High-density Design and TCO

ESG modeled and compared the storage-related costs that could be expected when deploying a scale-out NAS system and a Huawei OceanStor Pacific 9550 high-density system. The costs associated with purchasing, maintaining, powering, and cooling the storage systems were calculated in US Dollars, and the average cost for electricity in the United States as reported by the US Energy Information Administration<sup>5</sup> was used to calculate power and cooling costs. ESG modeled the expected storage total cost of ownership (TCO) for a company that needed to support a high availability, mixed-protocol production HPDA environment with 16.5PiB of usable capacity. Competing solutions were configured as similarly as possible. The largest disk drives available in each solution were used to build the solution.

TCO was calculated using a simplified model based on costs that would be incurred over a five-year period without taking into consideration capacity and performance growth requirements or IT operational costs. List prices for Huawei OceanStor Pacific were provided to ESG by Huawei. Costs for other solutions were obtained from publicly available sources. Maintenance and support contracts, along with typical customer discounts for hardware, software, and maintenance were factored into the estimated costs. Figure 11 shows the TCO cost comparison between scale-out NAS and next-gen Huawei OceanStor Pacific.

**Figure 11. TCO Analysis of the OceanStor Pacific Next-gen Storage System**



Source: Enterprise Strategy Group

As Figure 11 shows, Huawei OceanStor Pacific demonstrates a 61% overall TCO advantage over five years compared with a high-density scale-out NAS system. The largest savings (64%) come from hosting costs, thanks to the extremely high-density platform. CapEx savings are 62%, while power and cooling show a 32% advantage in this comparison.

<sup>5</sup> <https://www.eia.gov/>

## Why This Matters

With the number of tools and technologies that exist in a traditional enterprise environment, the cost and complexity related to maintaining the infrastructure, ensuring constant uptime, and guaranteeing performance levels can easily get out of hand. When asked to name their biggest challenges in terms of their on-premises file storage environments, data protection (30%), data migration (27%), hardware costs (26%), and rapid data growth rates (25%) were the most commonly cited responses.<sup>6</sup> A storage system must address all these challenges without compromising performance.

ESG validated that the Huawei OceanStor Pacific effectively addresses these issues. A single Huawei OceanStor Pacific storage system was able to deliver high performance with low latency for multiple workloads and consistent semantics across multiple protocols. While running our test workloads, we saw consistent performance across multiple file and object protocols with no discernable loss on any protocol.

In our tests, the Huawei Distributed Parallel Client enabled a single client to concurrently access multiple storage nodes, eliminating single-client and single-stream performance bottlenecks. A single-client, single-stream workload was able to drive 10.1 GiB/sec in reads and 7 GiB/sec in writes, while running multiple streams drove 10.68 GiB/sec reads and 8.25 GiB/sec in writes. The system also offers impressive density, scaling to 1.68 PB in just 5 RU.

ESG was impressed with the Huawei OceanStor Pacific platform's single- and multi-client performance, multi-protocol support, and the value provided by the platform in our 5-year cost of ownership analysis.

## The Bigger Truth

Organizations are continuing to generate and store exceptionally large amounts of unstructured data. ESG uncovered that more than half (56%) of organizations expect their on-premises data to grow by at least 21% annually over the next three years.<sup>7</sup> With the increasing adoption and use of data-intensive applications—life sciences, financial analysis, autonomous driving, and AI/ML, to name just a few—organizations require a solution that can efficiently store and process exceptionally large volumes of data with consistently high performance. The solution should also scale in a manner that enables organizations to increase performance and capacity independently.

Data growth is accelerating, and the resulting infrastructure required to store and protect that data is costly and complex. Organizations are tasked with providing a high-quality computing environment for an ever-growing number of HPDA applications while enterprise environments have become increasingly unpredictable as their underlying IT infrastructure grows in complexity and size. Mission-critical HPDA application performance is sensitive to storage performance and latency and highly dependent on the resilience of the IT environment.

The OceanStor Pacific storage system is designed to handle business- and mission-critical data analytics applications and workloads across multiple protocols simultaneously. The OceanStor Pacific enterprise-class availability features are implemented in software to provide a platform engineered for consolidating mission- and business-critical workloads with massive data sets at extremely low latencies.

ESG testing validated OceanStor Pacific's ability to consolidate heterogeneous data-intensive workloads on a single, high-performance, high-availability platform. The environment ESG tested serviced multiple workloads simultaneously, using multiple protocols to access the same data.

<sup>6</sup> Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019.

<sup>7</sup> Ibid.

The results that are presented in this document are based on testing in a controlled environment. Due to the many variables in each production data center, it is important to perform planning and testing in your own environment to validate the viability and efficacy of any solution.

ESG is pleased to validate that the Huawei OceanStor Pacific delivers consistently high performance for extremely large data sets and is clearly well suited to support demanding real-world, data-centric applications running in a performance-critical environment. Next-generation HPDA storage systems are designed with a goal of providing the best possible performance and capacity density while avoiding many of the limiting factors of traditional storage systems. Huawei designed the Pacific series around the dual goal of solving very large-scale analytics problems quickly and efficiently.

It is no surprise that ESG's five-year analysis demonstrated that, by deploying Huawei OceanStor Pacific rather than a traditional scale-out NAS system, organizations can lower their storage TCO by up to 61% while improving availability and reducing operational effort. If your organization is looking to lower storage TCO while increasing capacity and performance, ESG recommends investing in a next-generation system designed for HPDA, and [Huawei OceanStor Pacific](#) is worth a closer look.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.