



- Efficient Processing for Diverse Data
- Everlasting Operations for Services
- Economical Storage for Mass Data

Huawei OceanStor Pacific is an intelligent distributed storage series with scale-out capability designed to support the business needs of today and tomorrow. It features a wide range of storage systems that provide the high performance of traditional parallel storage and meet the needs of mission-critical and emerging workloads.

Product Overview

◀ Performance model

- OceanStor Pacific 9950 is a high-density, all-flash storage product that offers outstanding performance, capacity, and scalability. Each 5 U chassis houses a maximum of 8 storage nodes using all NVMe SSDs. Each chassis provides a raw capacity ranging from 128 TB to 614.4 TB, a bandwidth of up to 160 GB/s, and 2 million IOPS for ultra-fast data access. It is the perfect choice for mass unstructured data storage¹.
- OceanStor Pacific 9920 is Huawei's brand-new all-flash (SSD) distributed storage system with each 2 U chassis housing 1 node. It delivers excellent performance and features flexible component configurations to meet the access requirements of a variety of structured² and unstructured workloads.

◀ Capacity model

- OceanStor Pacific 9550 is a hybrid storage product that features high capacity density to deliver optimal cost effectiveness. Each 5 U chassis integrates 2 storage nodes and uses HDDs as large-capacity primary storage. Each chassis provides a raw capacity ranging from 720 TB to 1680 TB, reducing the cabinet footprint by 62.5% compared to general-purpose storage servers. It is suitable for mass unstructured data storage and can work with OceanStor Pacific 9950 to form a tiered data storage solution.
- OceanStor Pacific 9540 is a brand-new hybrid distributed storage system with high capacity density. Each 4 U chassis accommodates 1 node. It enables high capacity density and flexible component configurations to fulfill the access requirements of a wide range of structured and unstructured workloads.
- OceanStor Pacific 9520 is a brand-new hybrid distributed storage system. Each 2 U chassis houses 1 node, and it provides flexible component configurations to meet the access requirements of various structured and unstructured workloads.

Note: 1. Unstructured data storage includes distributed file, object, and HDFS storage. 2. Structured data storage includes distributed block storage.

The OceanStor Pacific series uses storage system software to integrate local storage resources from hardware nodes into a fully distributed storage pool. It provides upper-layer applications with file, HDFS, and object storage services, or block storage services. Its diverse and adaptable range of features provide efficient processing for diverse data, everlasting operations for services, and economical storage for mass data.

- File storage is applicable to high-performance computing (HPC) and other performance-intensive scenarios. It is fully compatible with native NFS and CIFS protocols, as well as parallel interfaces like POSIX and MPI-IO.

- HDFS storage provides a decoupled storage-compute big data solution using native HDFS semantics. It does not require plug-ins to be installed on compute nodes. Its intelligent tiering and coexistence of legacy coupled Hadoop clusters and new storage reduce the total cost of ownership (TCO) and offer a consistent user experience.
- Object storage provides superb small object processing performance and comprehensive disaster recovery capabilities. It is fully compatible with Amazon S3.
- Block storage allows for SCSI or iSCSI access modes, and delivers HA solutions such as distributed active-active and cabinet-level redundancy. It is widely applicable to virtualization, cloud resource pools, and databases.

Each product and storage service is available for individual or group purchase and on-demand deployment in a storage pool. You can access the same piece of data using multiple protocols, such as file, HDFS, and object storage, to implement efficient data access without migration and reduce storage space occupation.

The OceanStor Pacific series has been widely adopted in a variety of scenarios, including HPC, big data analytics, intelligent applications, virtualization, clouds, content repository, backup, and archiving. We can provide superior performance for oil & gas exploration, life sciences, financial institutions, carriers, Smart City projects, and Internet companies.

Features

◀ Efficient Processing for Diverse Data

The OceanStor Pacific series uses innovative performance acceleration technologies, such as directory DHT partitioning, intelligent disk granularity management, large I/O pass-through, small I/O aggregation, and multi-tier intelligent cache, to meet the access requirements of bandwidth- and OPS-intensive workloads with just one storage system. The new-generation parallel file system can be accessed using either standard protocols such as NFS and CIFS or Distributed Parallel Client (DPC) compatible with POSIX and MPI-IO. DPC can simultaneously connect multiple storage nodes to implement I/O load balancing. Additionally, it supports remote direct memory access (RDMA), achieving higher single-thread and single-client performance.

The OceanStor Pacific series uses the computing power provided by high-performance processors to offload some storage functions to the processor layer. It offers a 20% improvement in storage computing power compared to products with similar configurations. It adapts to any customer need for I/O, bandwidth, latency, or capacity. The OceanStor Pacific series supercharges the technology of today for the business of tomorrow.

◀ Everlasting Operations for Services

The OceanStor Pacific series uses E2E reliability assurance at the I/O, system, and data center levels to offer 99.9999% availability for different levels of data protection solutions. Specifically, it supports E2E data integrity checks and self-healing, and uses online verification and periodical background verification to handle silent data corruptions, such as bit changes and incorrect positions of read/write data. This helps ensure high data availability. The powerful elastic Erasure Coding (EC) provides data redundancy protection so that a single storage system can tolerate the simultaneous failure of up to four nodes. The OceanStor Pacific series supports dynamic EC. Faulty nodes trigger automatic adjustments to EC ratios without compromising system reliability. Technologies such as multi-module concurrent service takeover enable services on a fault node to be taken over within 10 seconds. The OceanStor Pacific series monitors device status in real-time to provide comprehensive sub-health detection and self-healing for disks, nodes, and networks. It also builds a disaster recovery system using replication or distributed active-active functions for data center reliability.

◀ Economical Storage for Mass Data

The OceanStor Pacific series organizes storage media, including HDDs and SSDs, into large-scale resource pools using distributed technologies and provides industry standard interfaces for upper-layer applications and clients. This eliminates the bottlenecks of traditional data centers and overcomes obstacles to system performance, such as unbalanced utilization of hardware resources by silo storage systems. It can start small and can scale-out to up to 4096 nodes in a storage cluster. This allows for linear performance growth as capacity expands, optimizing investment.

The series protects storage through the data redundancy of powerful elastic EC. EC nearly triples the disk space utilization compared to traditional multi-copy protection, offering a variety of EC ratios for flexible on-demand deployment. An EC ratio as high as 22+2 and disk space utilization rate of 91.6% help you reduce hardware investments. It uses abundant primary storage tiering policies to support automatic data migration between high-performance and large-capacity pools, enabling efficient storage and management of hot and cold data.

The OceanStor Pacific series also provides scenario-specific solutions to help you further optimize IT investments. For example, the decoupled storage-compute big data solution greatly improves data analysis efficiency and intelligently takes over services from third-party HDFS, maximizing the return on investment (ROI) and reducing the TCO.

◀ Simplified Lifecycle Management

The OceanStor Pacific series provides unified, converged management and intelligent O&M for multiple products and storage services. Intelligent resource prediction identifies the service risks of storage resources in advance. This enables you to make informed decisions regarding capacity expansion, procurement, and service changes. In addition, intelligent fault location provides comprehensive sub-health detection and processing across the system for precise fault location and troubleshooting.

Application Scenarios

◀ HPC

The OceanStor Pacific series provides high-performance file storage services for HPC platforms. It supports DPC that is compatible with POSIX and MPI-IO to better adapt to HPC services. Automatic storage tiering optimizes storage architecture to increase both efficiency and cost effectiveness.

Typical industry scenarios: Oil & gas exploration, genome sequencing, electron cryomicroscopy, weather forecasting, and autonomous driving.

◀ Big data analytics

The OceanStor Pacific series provides a decoupled storage-compute big data solution. This enables on-demand configuration and flexible expansion of storage and compute resources, and reduces TCO. It is fully compatible with native HDFS semantics for a consistent user experience and supports the coexistence of coupled and decoupled storage-compute architectures for seamless infrastructure evolution.

Industry scenarios: Big data for offline finance analytics, Internet log retention, operational carrier analytics, governments, and Smart City projects.

◀ Content repository and backup archiving

The object storage resource pools of the OceanStor Pacific series provide enterprises with high performance and reliability for large throughput, frequent access to hot data, long-term storage, and online access. The OceanStor Pacific series is ideal for real-time online services, such as Internet data, online audiovisual data, and enterprise web disks.

Industry scenarios: Production, storage, backup, and archiving of banking check images, audio and video recordings, medical images, electronic documents of governments and enterprises, Smart City projects, and the Internet of Vehicles (IoV).

◀ Virtualization/Cloud resource pool

The OceanStor Pacific series provides a storage resource pool to host mass data for on-demand resource provisioning and elastic capacity expansion in virtualization and cloud environments. It dramatically improves storage resource deployment, expansion, and O&M efficiency.

Industry scenarios: Channel access clouds for Internet finance, development and testing clouds, carrier cloud services for BOM domains and B2B cloud resource pools, smart governments, and Smart City clouds.

Specifications

◀ Key specifications of distributed file, object, and HDFS storage products

| Model | OceanStor Pacific 9520 | OceanStor Pacific 9540 | OceanStor Pacific 9550 | OceanStor Pacific 9920 | OceanStor Pacific 9950 |
|---|---|--|--|---|---|
| System Architecture | Fully symmetric distributed architecture | | | | |
| Max. Raw Capacity per Chassis | 168TB | 504TB | 1680TB | 92TB | 614.4TB |
| Height per Chassis | 2U | 4U | 5U | 2U | 5U |
| Number of Nodes per Chassis | 1 | 1 | 2 | 1 | 8 |
| Max. Number of Primary Storage Disks per Node | 12 | 36 | 60 | 12 | 10 |
| Number of Processors per Node | 2 x Huawei Kunpeng 920 Processors or 2 x Intel® Xeon® Scalable Processors | 2 x Huawei Kunpeng 920 Processors or 2 x Intel® Xeon® Scalable Processors | Huawei Kunpeng 920 Processor | 2 x Huawei Kunpeng 920 Processors | Huawei Kunpeng 920 Processor |
| Max. Memory per Node | 256 GB, 512 GB | 256 GB, 512 GB | 256 GB | 512 GB | 256 GB |
| Max. Cache per Node | 4 x NVMe SSDs | 4 x NVMe SSDs | 4 x Half-palm NVMe SSDs | N/A | N/A |
| Number of System Disks per Node | 2 x 600 GB SAS HDDs | 2 x 600 GB SAS HDDs | 2 x 480 GB SSDs | 2 x 600 GB SAS HDDs | 2 x 480 GB SSDs |
| Data Disk Types | 3.5-inch HDDs | 3.5-inch HDDs | 3.5-inch HDDs | 3.5-inch SAS SSDs | Half-palm NVMe SSDs |
| Front-End Service Networks ¹ | 10GE or 25GE TCP/IP | <ul style="list-style-type: none"> 10GE, 25GE, or 100GE TCP/IP 25GE or 100GE RoCE 100 Gb/s EDR InfiniBand | <ul style="list-style-type: none"> 10GE, 25GE, or 100GE TCP/IP 100GE RoCE 100 Gb/s EDR/HDR InfiniBand | <ul style="list-style-type: none"> 25GE or 100GE TCP/IP 100GE RoCE 100 Gb/s EDR InfiniBand | <ul style="list-style-type: none"> 25GE or 100GE TCP/IP 100GE RoCE 100 Gb/s EDR/HDR InfiniBand |
| Storage Interconnection Networks | 25GE RoCE | <ul style="list-style-type: none"> 25GE RoCE 100 Gb/s EDR InfiniBand | 25 GE RoCE | <ul style="list-style-type: none"> 100GE RoCE 100 Gb/s EDR InfiniBand | 100 GE RoCE |
| Data Redundancy Protection Mechanism | EC: N + M (M is 2, 3, or 4) | | | | |
| Storage Access Protocols | NFS, CIFS, POSIX, MPI-IO, HDFS, and Amazon S3 | | | | |
| Key Features | SmartQuota (quotas), SmartTier (storage tiering), SmartQoS (service quality), SmartEqualizer (load balancing), SmartMulti-Tenant (multi-tenancy), SmartEncryption (data encryption), SmartAuditlog (audit logs), HyperSnap (snapshots), HyperReplication (asynchronous replication), SmartIndexing (metadata indexing), Recycle Bin (recycle bin), SmartInterworking (multi-protocol interworking), DIF (end-to-end data integrity verification), Object Versioning (versioning) ² , and SmartTakeover (intelligent takeover) ³ | | | | |
| Data Self-Healing | Automatic concurrent data reconstruction at 2 TB per hour | | | | |
| Chassis Dimensions (H x W x D) | <ul style="list-style-type: none"> Kunpeng model: 86.1 mm x 447 mm x 790 mm x86 model: 86.1 mm x 447 mm x 748 mm | <ul style="list-style-type: none"> Kunpeng model: 175 mm x 447 mm x 790 mm x86 model: 175 mm x 447 mm x 748 mm | 219.5 mm x 447 mm x 1030 mm | 86.1 mm x 447 mm x 790 mm | 219.5 mm x 447 mm x 926 mm |
| Max. Weight per Chassis (with Disks) | <ul style="list-style-type: none"> Kunpeng model: ≤ 32 kg x86 model: ≤ 34.1 kg | ≤ 65 kg | ≤ 164 kg | ≤ 32 kg | ≤ 115 kg |
| Operating Temperature | 5°C to 35°C | 5°C to 35°C | 5°C to 35°C | 5°C to 35°C | 5°C to 35°C |
| Operating Humidity | 8% to 90% RH (non-condensing) | 8% to 90% RH (non-condensing) | 5% to 90% RH (non-condensing) | 8% to 90% RH (non-condensing) | 5% to 90% RH (non-condensing) |

◀ Key specifications of distributed block storage products

| Model | OceanStor Pacific 9520 | OceanStor Pacific 9540 | OceanStor Pacific 9920 |
|---|---|--|--|
| System Architecture | Fully symmetric distributed architecture | | |
| Max. Raw Capacity per Chassis | 168TB | 504TB | 92TB |
| Height per Chassis | 2U | 4U | 2U |
| Number of Nodes per Chassis | 1 | 1 | 1 |
| Max. Number of Primary Storage Disks per Node | 12 or 25 | 36 | 12, 22, or 25 |
| Number of Processors per Node | 2 x Huawei Kunpeng 920 Processors or 2 x Intel® Xeon® Scalable Processors | 2 x Huawei Kunpeng 920 Processors or 2 x Intel® Xeon® Scalable Processors | 2 x Huawei Kunpeng 920 Processors or 2 x Intel® Xeon® Scalable Processors |
| Max. Memory per Node | 512 GB, 768 GB, or 1 TB | 512 GB, 768 GB, or 1 TB | 768 GB or 1 TB |
| Max. Cache per Node | 4 x NVMe SSDs or SAS SSDs | 4 x NVMe SSDs or SAS SSDs | N/A |
| Number of System Disks per Node | 2 x 600 GB SAS HDDs | 2 x 600 GB SAS HDDs | 2 x 600 GB SAS HDDs |
| Data Disk Types | 2.5-inch or 3.5-inch HDDs | 3.5-inch HDDs | 2.5-inch or 3.5-inch SAS SSD or NVMe SSD |
| Front-End Service Networks | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE RoCE | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE RoCE | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE or 100GE RoCE 100 Gb/s InfiniBand |
| Storage Interconnection Networks | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE RoCE | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE RoCE | <ul style="list-style-type: none"> 10GE or 25GE TCP/IP 25GE or 100GE RoCE 100 Gb/s InfiniBand |
| Data Redundancy Protection Mechanism | <ul style="list-style-type: none"> EC: N + M (M is 2, 3, or 4), applicable to SSDs or HDDs used as primary storage Multi-copy: 3-copy mode | | |
| Storage Access Protocols | iSCSI, SCSI, and OpenStack Cinder | | |
| System Security Policies | Disk, node, and cabinet levels | | |
| Key Features | SmartThin (thin provisioning), SmartDedupe & SmartCompression (data reduction), SmartQoS (service quality), SmartAuditlog (audit logs), HyperSnap (snapshots), HyperClone (linked clone), HyperMetro (distributed active-active), HyperReplication (asynchronous replication) MultiPool (multiple resource pools), and DIF (end-to-end data integrity verification) | | |
| Data Self-Healing | Automatic concurrent data reconstruction at 4 TB per hour | | |
| Deployment Scheme | Decoupled and coupled deployment of compute and storage resources | | |
| Compatible Platforms | Huawei FusionSphere, VMware vSphere, Microsoft Windows Hyper-V, OpenStack, and containers ⁴ | | |
| Chassis Dimensions (H x W x D) | <ul style="list-style-type: none"> Kunpeng model: 86.1 mm x 447 mm x 790 mm x86 model: Chassis with 3.5-inch disks: 86.1 mm x 447 mm x 748 mm Chassis with 2.5-inch disks: 86.1 mm x 447 mm x 708 mm | <ul style="list-style-type: none"> Kunpeng model: 175 mm x 447 mm x 790 mm x86 model: 175 mm x 447 mm x 748 mm | <ul style="list-style-type: none"> Kunpeng model: 86.1 mm x 447 mm x 790 mm x86 model: Chassis with 3.5-inch disks: 86.1 mm x 447 mm x 748 mm Chassis with 2.5-inch disks: 86.1 mm x 447 mm x 708 mm |
| Max. Weight per Chassis (with Disks) | <ul style="list-style-type: none"> Kunpeng model: ≤ 32 kg x86 model: ≤ 34.1 kg | ≤ 65 kg | ≤ 32 kg |
| Operating Temperature | 5°C to 35°C | 5°C to 35°C | 5°C to 35°C |
| Operating Humidity | 8% to 90% RH (non-condensing) | 8% to 90% RH (non-condensing) | 8% to 90% RH (non-condensing) |

Note: 1. NFS, CIFS, HDFS, and S3 support TCP/IP networks. DPC supports RoCE and InfiniBand networks. 2. Object Versioning is applicable to object storage. 3. SmartTakeover is applicable to HDFS storage. 4. iSCSI can be used to connect containers.

For More Information

To learn more about Huawei storage, please contact the local office or visit Huawei Enterprise website <http://e.huawei.com>.



Huawei Enterprise APP





Huawei IT



Copyright © Huawei Technologies Co., Ltd. 2021 . All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademark Notice

 , HUAWEI , and  are trademarks or registered trademarks of Huawei Technologies Co., Ltd.
Other trademarks, product, service and company names mentioned are the property of their respective owners.

NO WARRANTY

THE CONTENTS OF THIS MANUAL ARE PROVIDED "AS IS". EXCEPT AS REQUIRED BY APPLICABLE LAWS, NO WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ARE MADE IN RELATION TO THE ACCURACY, RELIABILITY OR CONTENTS OF THIS MANUAL.

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO CASE SHALL HUAWEI TECHNOLOGIES CO., LTD BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES, OR LOST PROFITS, BUSINESS, REVENUE, DATA, GOODWILL OR ANTICIPATED SAVINGS ARISING OUT OF OR IN CONNECTION WITH THE USE OF THIS MANUAL.

HUAWEI TECHNOLOGIES CO., LTD.
Bantian, Longgang District
Shenzhen 518129, P. R. China
Tel: +86-755-28780808

www.huawei.com