HUAWEI

Building a Fully Connected, Intelligent World

OceanStor Pacific Series

Next-Gen Storage for High-Performance Data Analytics (HPDA)

HPC Is Benefiting Humankind

Mass data and HPC keep advancing technology and society

R&D of new drugs

5000 days → 100 days

R&D of new cars

60 months → 24 months

Genome sequencing time

13 years → 1 day

Weather forecast accuracy

21.8% → 80%

Larger data volumes caused by emerging services

Higher scalability and lower TCO are required.

6 TB
Daily data generated per device

Genome sequencing

1 PB
Annual data collected by a single satellite

Satellite remote sensing

10x
More data in 3D exploration than 2D

Oil & Gas exploration

60 TB
Daily training data generated per car

Autonomous driving

More complex workloads due to diversified computing power

One storage system must support different types of workloads.

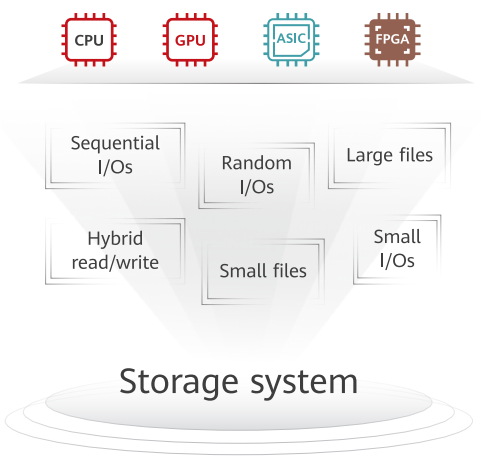

Efficient data mobility required by high-performance analytics

Multiple protocols share one copy of data, eliminating the need for data migration.

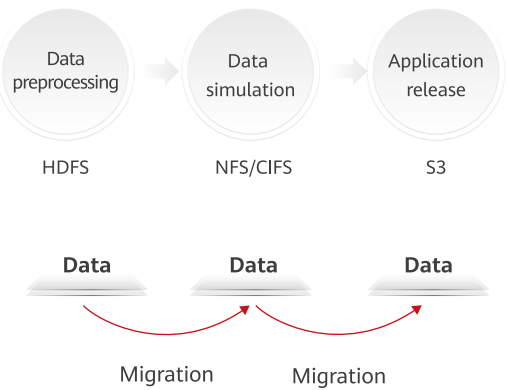

Big Data Becomes a New Driving Force

HPDA leads industry trends as HPC shifts from computing- to data-intensive


OceanStor Pacific Series

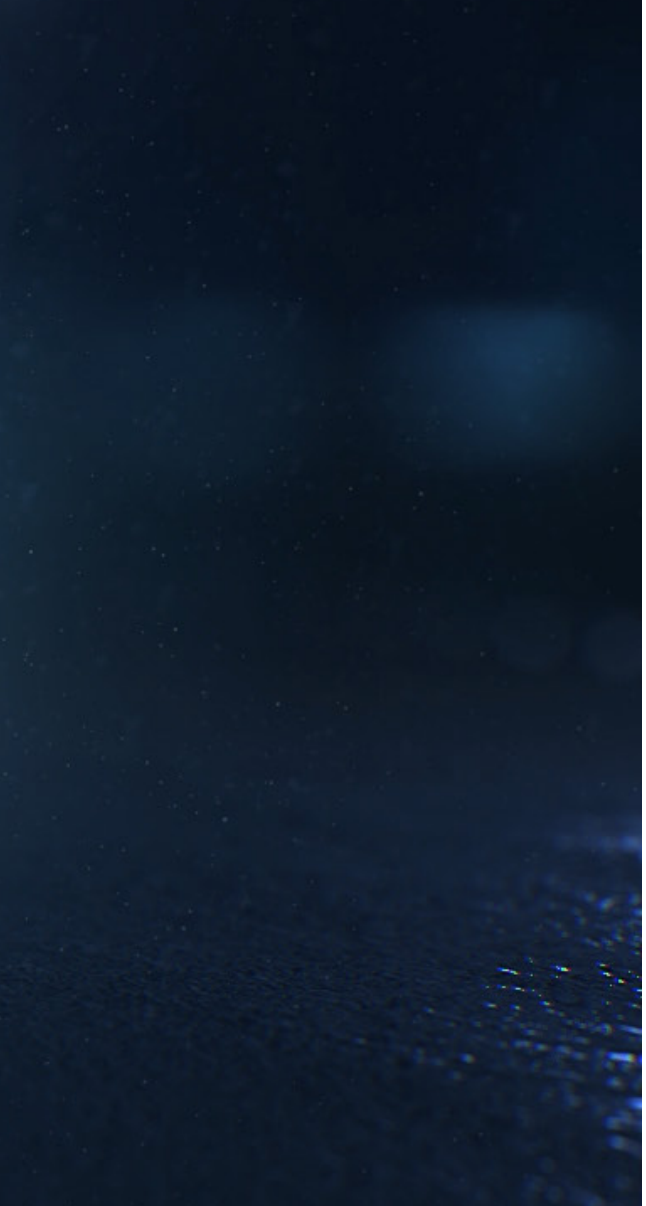

Next-Gen Storage for HPDA

High-density design

Hybrid workloads-oriented

Multi-protocol interworking

High-Density Design

Overall cost optimization

Performance models

2 U

OceanStor Pacific 9950

NEW

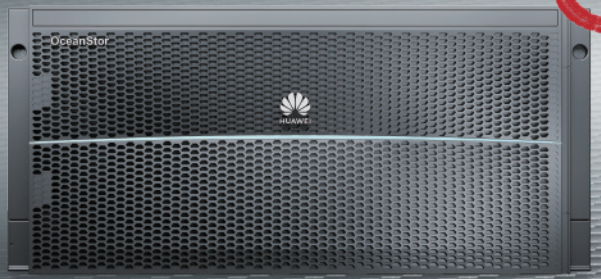

5 U high-density product

Capacity models

OceanStor Pacific 9550

NEW

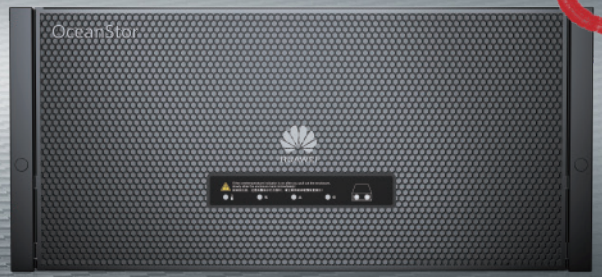

5 U high-density product

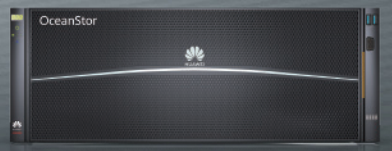

4 U

OceanStor Pacific 9950

High-density performance model

- 5 U, 8 nodes
- 80 disk slots, NVMe SSDs
- 160 GB/s bandwidth, 2 million IOPS

8 nodes/chassis, 10 half-palm NVMe SSDs/node

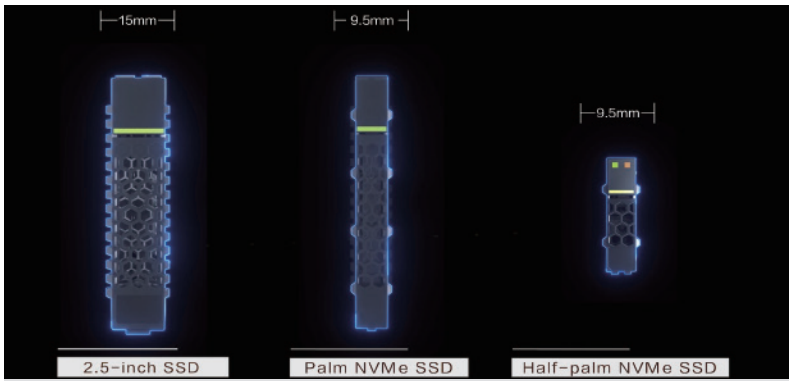

Half-palm NVMe SSD
Smaller cross-sectional area for higher density

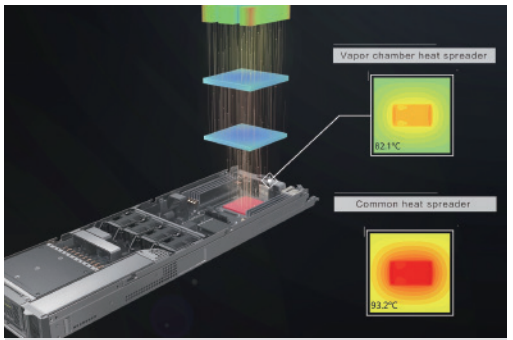

Carbon fiber pad + vapor chamber heat spreader for efficient heat dissipation

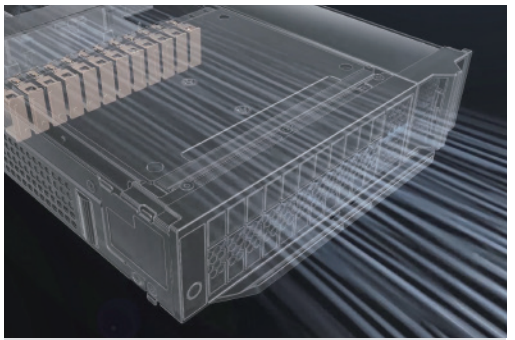

Horizontal backplane design for extremely low wind resistance

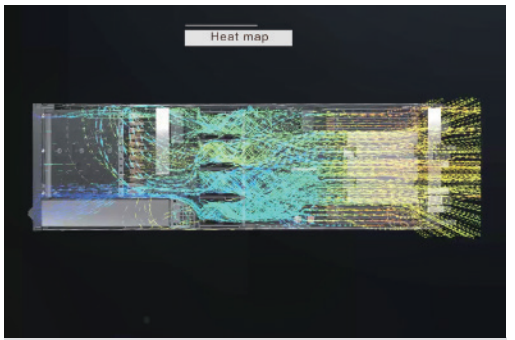

A well-ventilated channel for unobstructed airflow

OceanStor Pacific 9950

High-density performance model

- Back-end 100GE
- Front-end 100GE, 100 Gbit/s IB

- PCIe 4.0
- RDMA networking

- Passthrough of large I/Os, aggregation of small I/Os

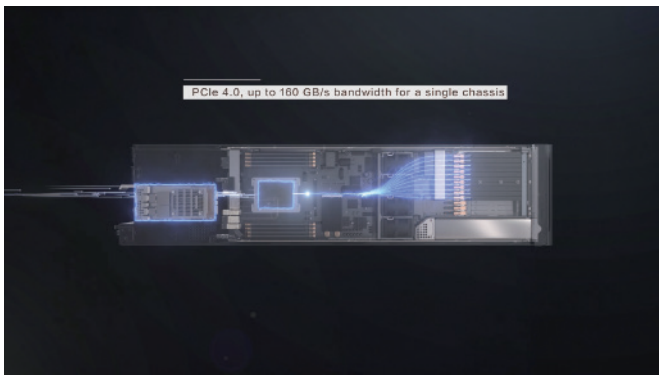

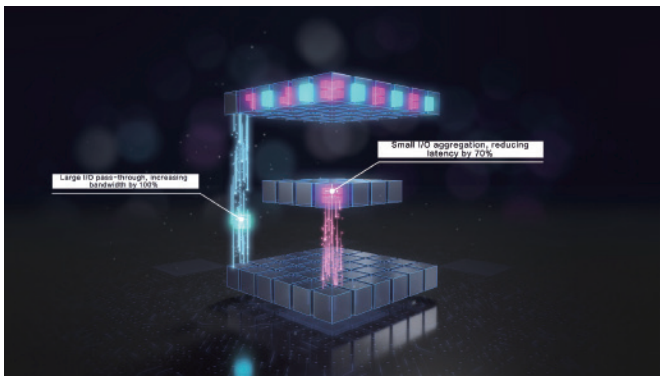

E2E NVMe Flash Design Delivers Unprecedentedly High Performance

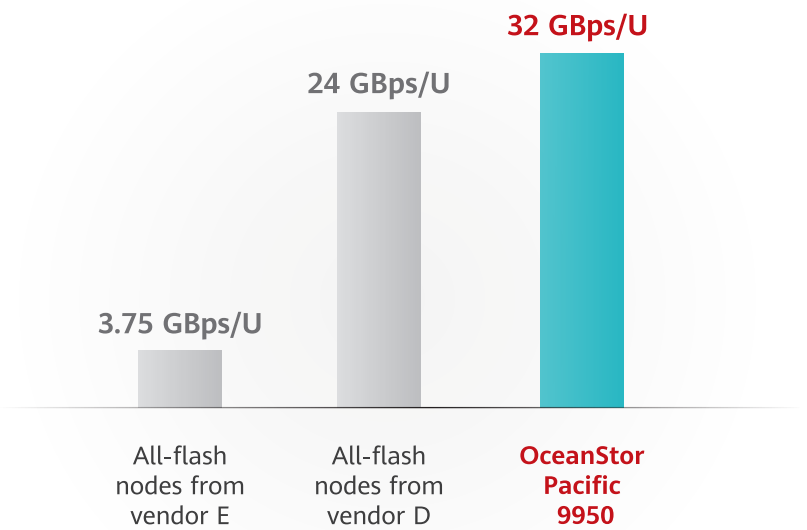

↑ 30% bandwidth/U

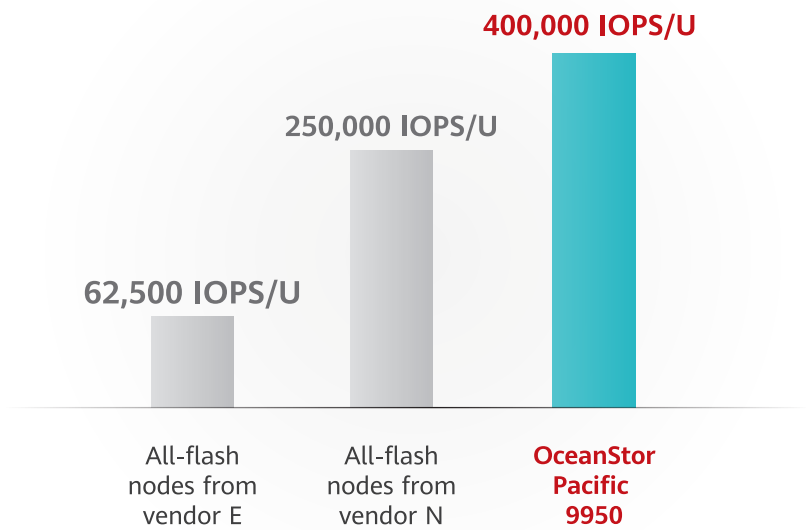

↑ 60% IOPS/U

OceanStor Pacific 9550

High-density capacity model

- 5 U, 120 disk slots, 3.5-inch SATA disks
- 1.68 PB raw capacity
- Elastic EC

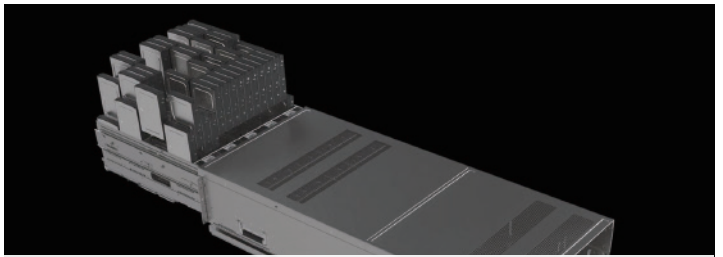

Bidirectional drawer slide
50% increase in space between disks, 30% increase in heat dissipation, and bidirectional maintainability

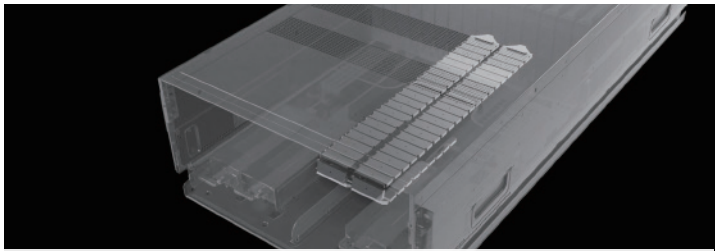

Tank chain + Flexible cable
Non-disruptive, independent disk maintenance

Higher Density for Better Space Utilization

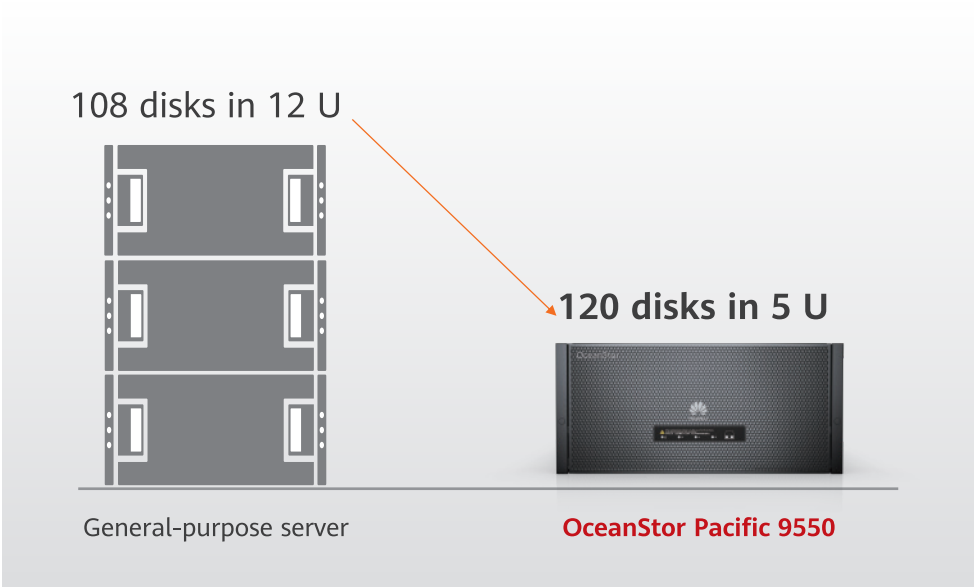

1 chassis = 3 nodes

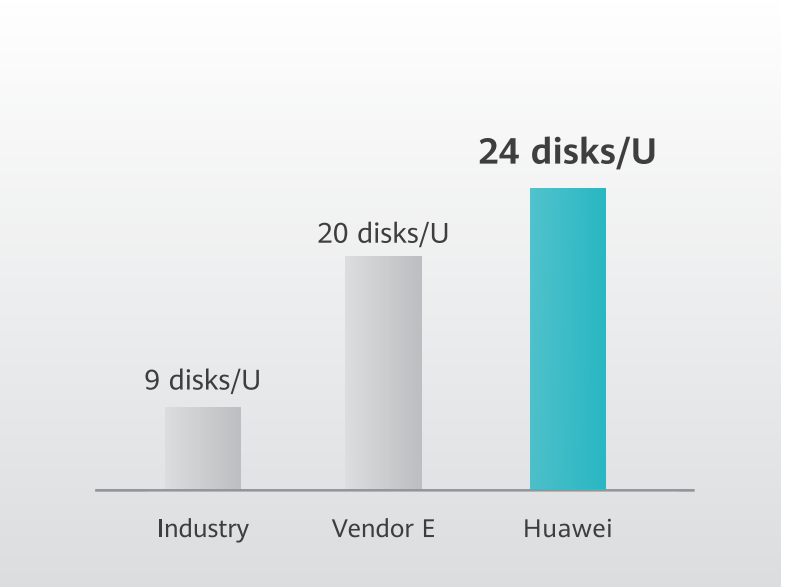

2.67x density | **62.5%** space reduction

OceanStor Pacific 9550

High-density capacity model

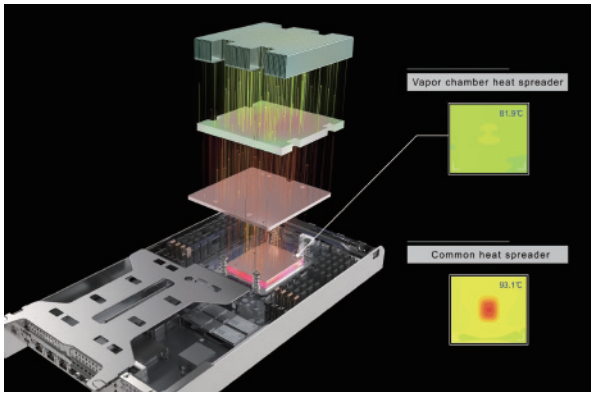

Carbon fiber pad + VC heat spreader

20% increase in heat dissipation, 11℃ decrease in CPU temperature

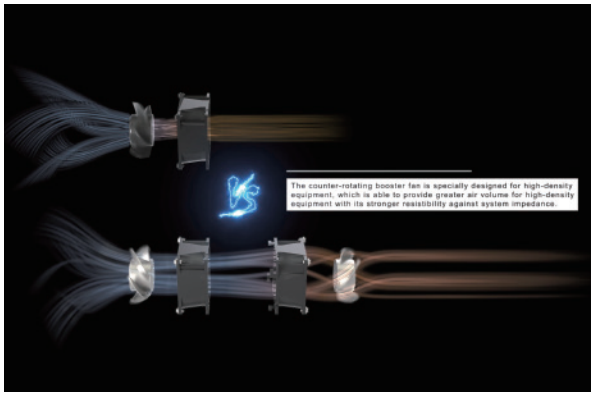

Aviation-level counter-rotating fans

Higher impedance resistance and larger air volume

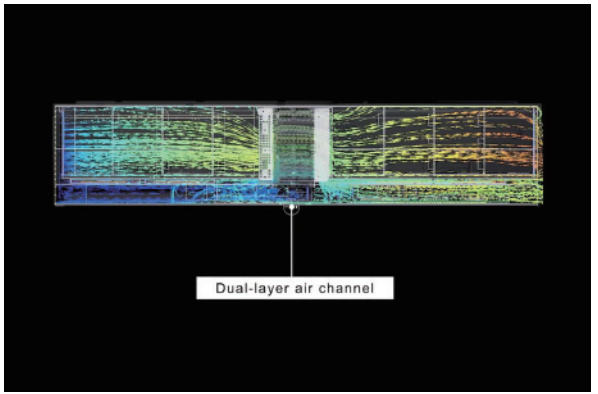

Dual-layer air channel

Higher impedance resistance and larger air volume

Larger Capacity for Mass Data

3,360
512 GB mobile phones

7 million
CDR entries

840
4K feature films

560,000
genome profiles

OceanStor Pacific 9550

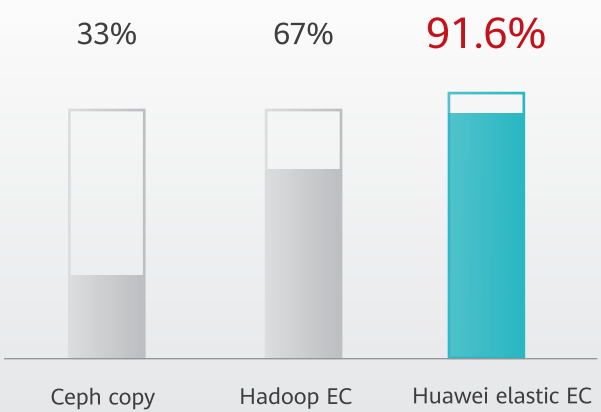

1.68 PB
raw capacity

2.7x
disk utilization

Hybrid Workloads-Oriented

All-scenario acceleration

Next-gen parallel file system

Support for bandwidth- and I/O-intensive workloads

Distributed parallel client (DPC)

Concurrent access to multiple service nodes for higher performance

Next-Gen Parallel File System Designed for Hybrid Workloads

Optimized metadata mechanism

Metadata is distributed to owning storage nodes, eliminating overheads of distributed locking and reducing latency.

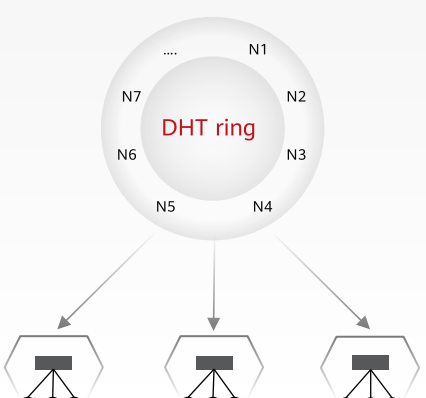

Separate processing of large and small I/Os

Large I/Os are directly written to disks to reduce forwarding. Small I/Os are aggregated at the cache layer and then written to disks to reduce interactions and I/O overheads.

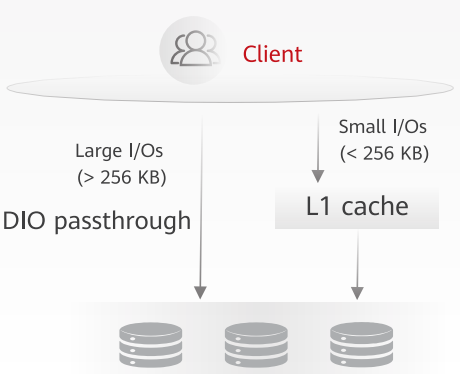

Optimized space management mechanism

Large files and I/Os are managed with consecutive spaces to improve bandwidth. Small files and I/Os are managed with small-granularity spaces to prevent read/write amplification and reduce latency.

One System Meets the Workload Requirements of Multiple Fields in **Oil and Gas Exploration**

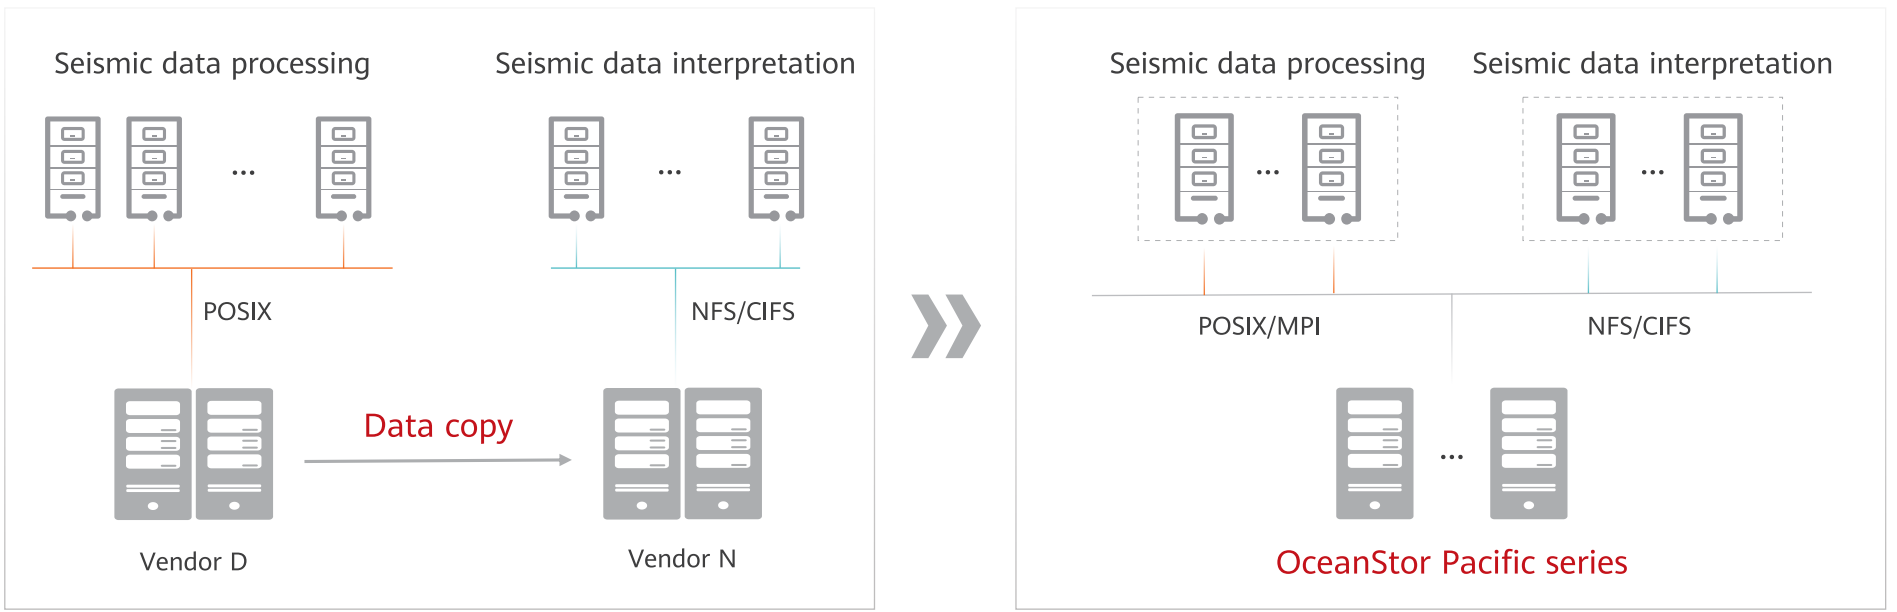

10+ PB capacity
50+ GB/s bandwidth

500+ TB capacity
600,000+ IOPS

EB-level scalability
160 GB/s bandwidth and 2 million IOPS in a 5 U space

DPC

NFS client access mode

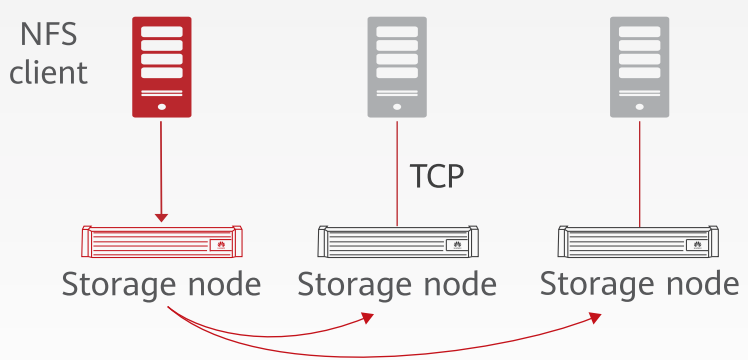

- A single client can only connect to one storage node, meaning that a single-point performance bottleneck is likely to occur.
- After data is written, cross-node forwarding is required, resulting in high latency.

DPC access mode

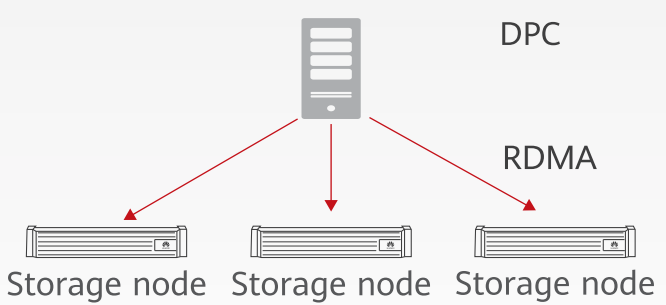

- A single client concurrently accesses multiple storage nodes. I/O-level load balancing ensures high performance per client and thread.
- Data is directly distributed to different nodes without cross-node forwarding, ensuring low latency.

DPC Improves Performance in Various Scenarios

Weather forecast

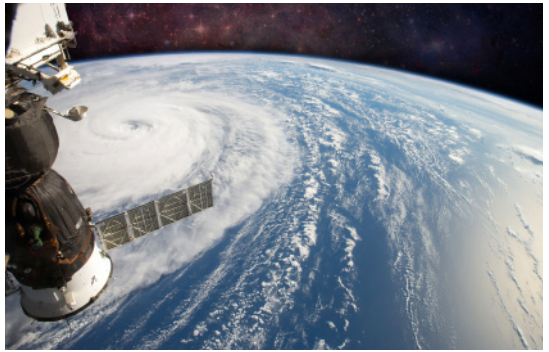

MPI-IO accelerates weather data analytics applications.

Satellite mapping

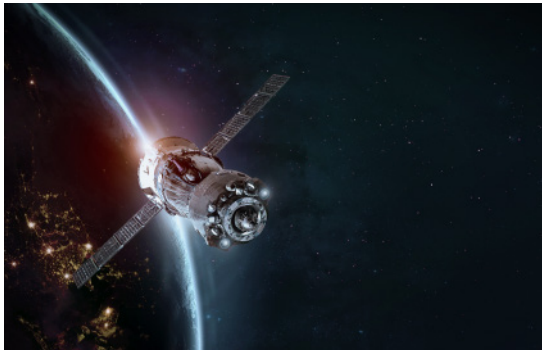

3+ GB/s single-thread performance ensures satellite data reception.

Supercomputing center

10+ GB/s single-client performance unleashes the potential of fat clients.

Multi-Protocol Interworking

Worry-free service evolution

One copy of data, migration-free

- Support for NFS, CIFS, HDFS, and S3
- Native interfaces without semantic and performance loss

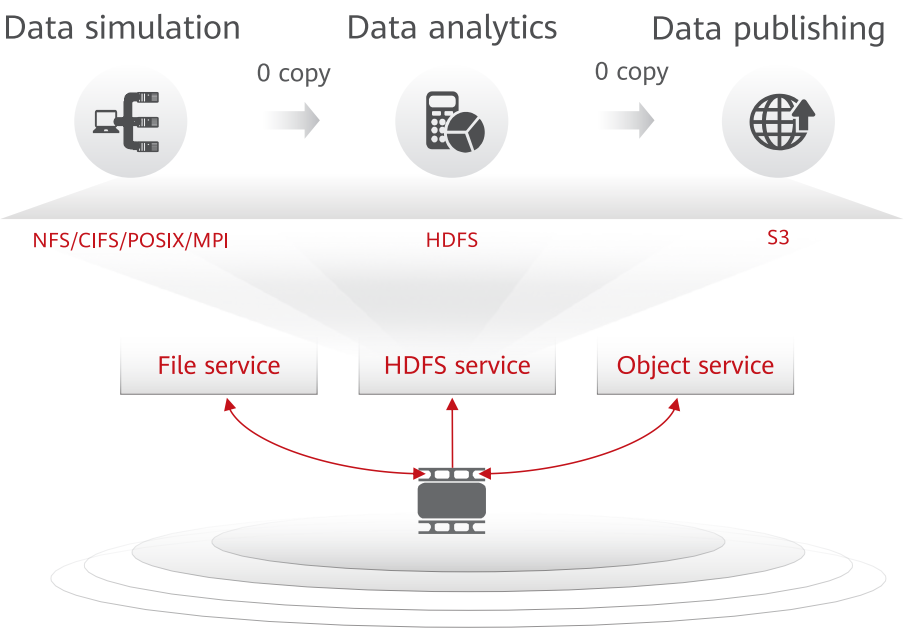

Multi-Protocol Interworking Improves the R&D Efficiency of **Autonomous Driving**

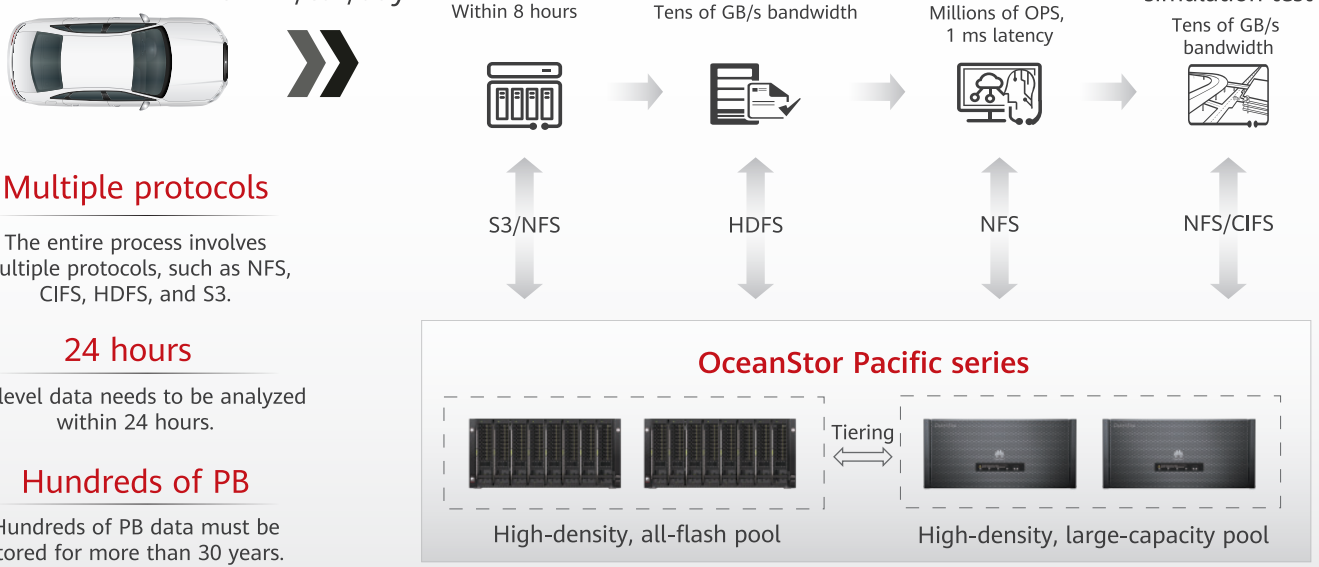

Multiple protocols

The entire process involves multiple protocols, such as NFS, CIFS, HDFS, and S3.

24 hours

PB-level data needs to be analyzed within 24 hours.

Hundreds of PB

Hundreds of PB data must be stored for more than 30 years.

Efficiency 25% ↑

Multi-protocol interworking enables one storage system to provide high bandwidth and OPS, without data copying.

TCO 20% ↓

Hot and cold data are intelligently tiered. High-density, large-capacity hardware reduces the footprint.

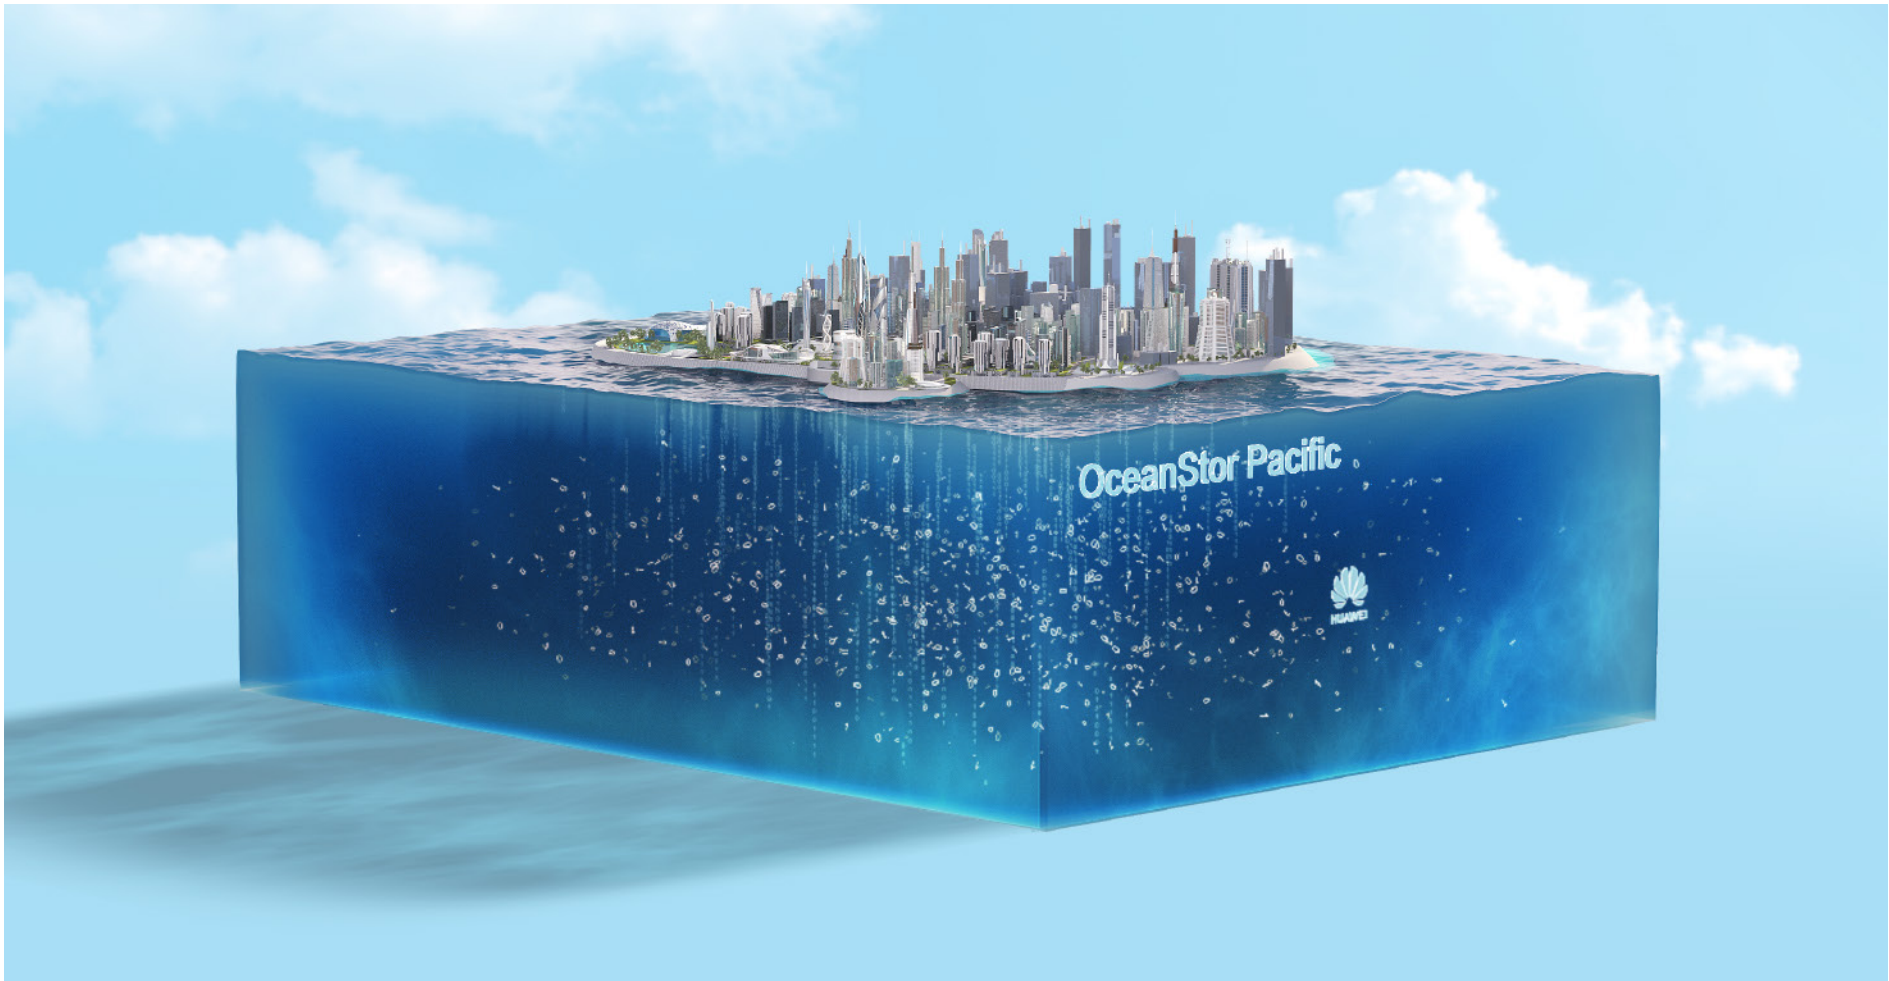
OceanStor Pacific
HUAWEI

OceanStor Pacific Series

Next-Gen Storage for HPDA

High-density design

Overall cost optimization

Hybrid workloads-oriented

All-scenario acceleration

Multi-protocol interworking

Worry-free service evolution

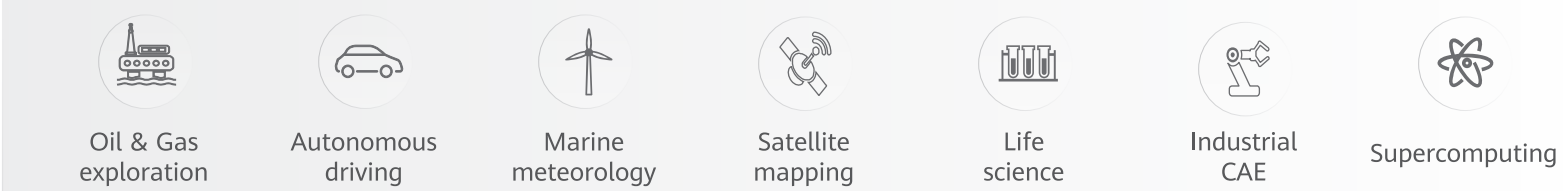

Oil & Gas Exploration

2D → 3D: 5x to 10x ↑
Grid density doubled each time: 3x to 4x ↑

10x to 20x data increase but incapable of non-disruptive capacity expansion

Raw data: 1 to 20 TB/day
Process data: 0.5 to 10 PB

Service bottleneck due to frequent data migrations

Multiple cross-cluster data migrations and multiple copies of the same data, resulting in low resource utilization

Large data volume, high concurrency, and frequent human-machine interaction

Aggregate bandwidth for seismic data processing: 2 to 20 GB/s per PB

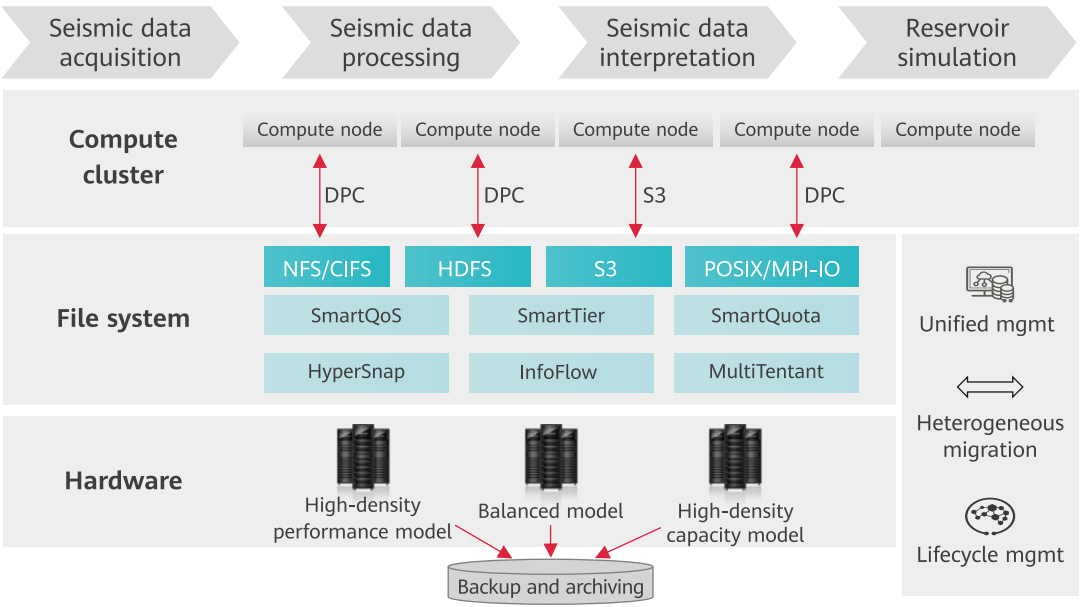

1 **One-time construction and any-time capacity expansion**

Fully symmetric architecture, small-scale initial deployment, and linear expansion to thousands of nodes along with data growth to optimize investment

2 **One storage device with automatic tiering based on data access frequency**

Automatic tiering of hot, warm, and cold data, with no impact on services and reduced TCO throughout the lifecycle

3 **N sets -> 1 set, migration-free and easy maintenance**

Seismic data processing/interpretation and reservoir simulation carried by the same architecture without data migration, improving exploration efficiency and reducing data copies

Genome Sequencing

Genome/person

3 GB

Parallel sequencing: 30+ times

Raw data (~100 GB) + Quality data (~300 GB)

-> Compressed data (~100 GB)

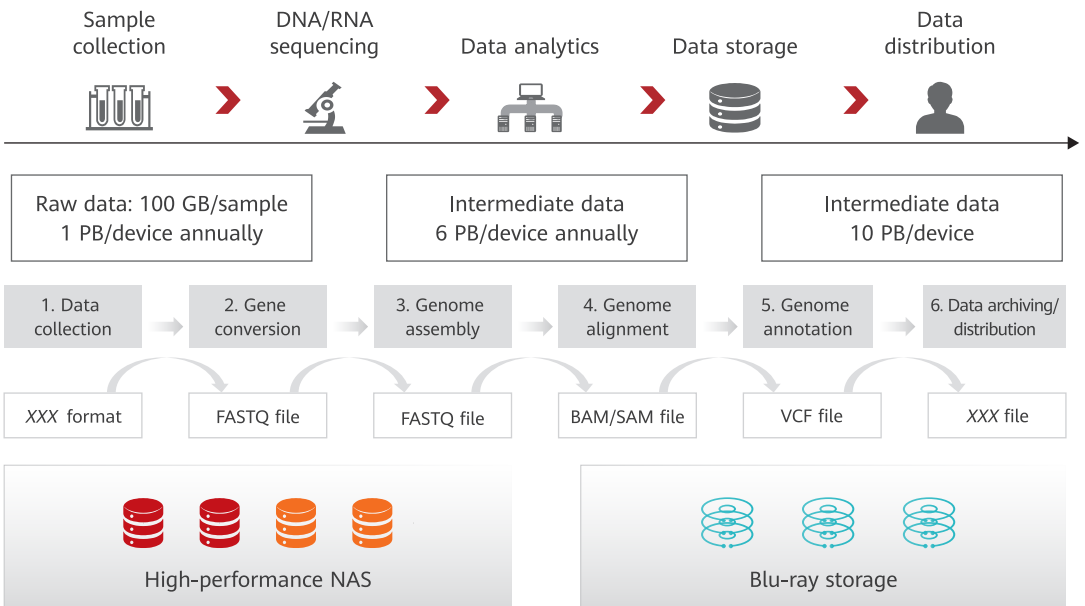

Long-term storage of mass genetic data

A sequencer produces 6 TB/40 hours. The intermediate files and results generated in the analysis process are around 5 times the raw data volume.

Large data calculation volume

Genome sequencing requires intensive computing instances for I/O, computing, and memory needs.

1

Storage for 50+ years

Automatic tiering and free mobility of hot, warm, and cold data

Cold data archived on Blu-ray discs for 50+ years

2

One data copy throughout the process, migration-free

The same data copy accessed over different protocols in sequencing, alignment, analysis, and display phases

3

Analytics performance improves with all-flash + private clients

NVMe all-flash storage with 160 GB/s bandwidth per chassis

3+ GB/s single-thread bandwidth for the DPC

Weather Forecast

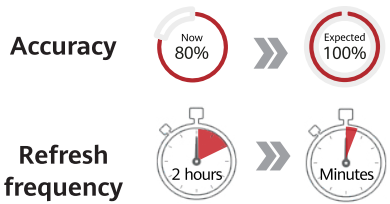

Mass data storage

100 TB observation data is collected per day. The production data reaches tens of PB with hundreds of PB historical data.

Timeliness and accuracy

8 rounds of calculation are performed per day with 80% accuracy. Multi-mode concurrent execution is required to improve accuracy.

Data mobility efficiency

Multiple data imports and exports across departments and processes result in silos and a serious waste of resources.

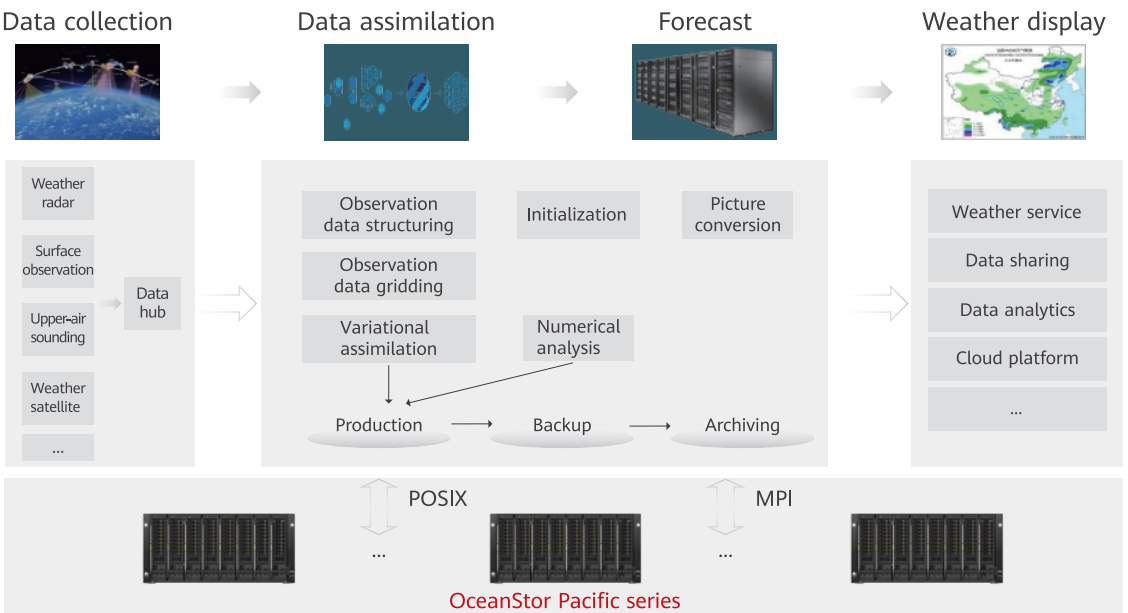

1 Multiple phases served by one storage system

Collection, preprocessing, mode calculation, post-processing, and data display served by one storage system, without data migration

2 Data tiering and archiving without affecting ongoing services

Automatic tiering of hot, warm, and cold data, with no impact on services, reduced TCO throughout the lifecycle, and data archived on Blu-ray discs for 100 years

3 Parallel analysis of mass data

Support for MPI-IO, enabling multiple compute nodes to access the same file and greatly improving analytical efficiency

Supercomputing

Traditional -> AI

Petascale -> Exascale

Hybrid workload acceleration

Multiple types of workloads and complex I/O models require the storage system to provide ideal performance in all scenarios.

Multi-tenant service

Services are provided for multiple industry customers at the same time, requiring QoS and resource isolation between tenants.

Convergence of traditional HPC, big data, and AI

Cross-protocol access capabilities improve data processing efficiency of big data and AI convergence services.

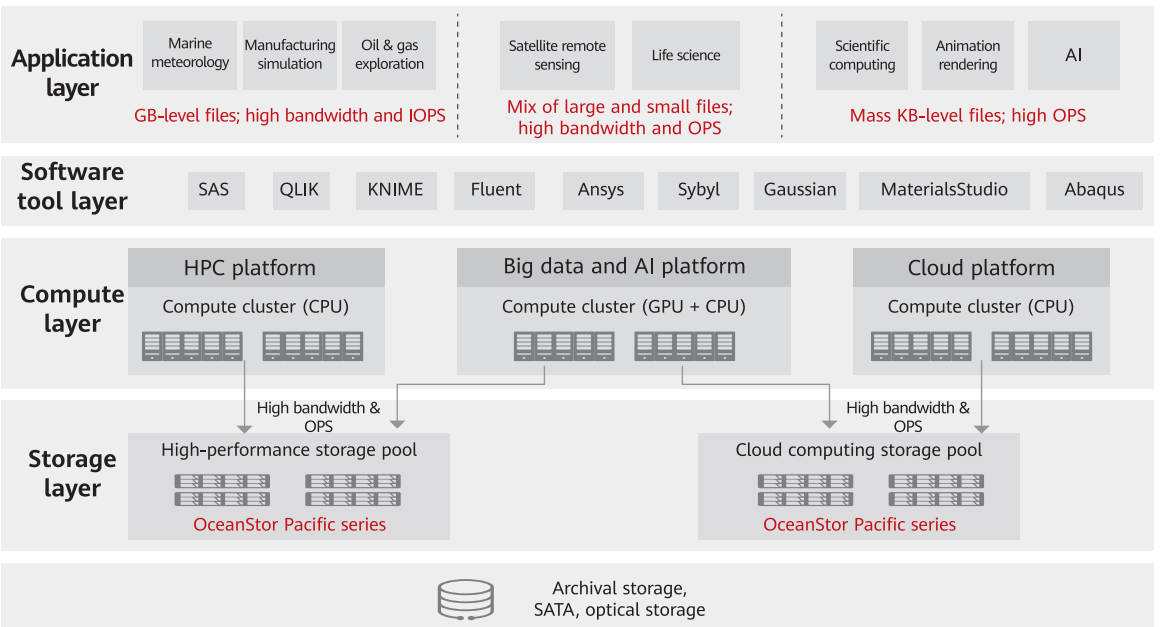

1 Adaptation to complex workload models of supercomputing

One storage system supports high bandwidth and IOPS, and DPCs support MPI-IO, providing high single-thread and single-client performance.

2 Flexible data access by multiple applications without data migration

Interworking between file, HDFS, and object protocols makes HPDA and AI data access more efficient.

3 Data and load isolation between tenants

Multi-tenant and QoS capabilities isolate service resources and performance.

Lundin Energy Norway Makes Oil Exploration More Efficient and Opens Up a New Era of Deep Sea Exploration with OceanStor Storage

A storage infrastructure that matches new reservoir simulation technologies is built for efficient data storage and analytics, higher exploration precision and efficiency, as well as faster business development of oil and gas reserves exploration.

Shanghai Astronomical Observatory Accelerates Exploration of the Unknown Universe with Huawei OceanStor Storage

A core storage system is built for the world's first SKA regional center prototype to aggregate mass scientific data and accelerate exploration of the formation and origin of the universe structure.

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

Scan for Mobile Reading