Model		Performance Node	Capacity Node	Mini-Capacity Node
Supported protocol		NFS, CIFS, HDFS, NIS, Microsoft Active Directory, LDAP, and SNMP		
Alarm notification		Email, SMS, SNMP, and Syslog		
Free from instant maintenance		Automatic bad disk detection, alarm notification, and centralized batch replacement of bad disks, avoiding instant replacement and reducing manual maintenance		
Physical Feature				
Power supply		AC 100 V to 127 V, AC 200 V to 240 V		
Dimensions (H x W x D)	Node	2 U, 86.1 mm x 446 mm x 582 mm (3.39 in. x 17.56 in. x 22.91 in.)	4 U, 175 mm x 446 mm x 582 mm (6.89 in. x 17.56 in. x 22.91 in.)	2 U, 86.1 mm x 446 mm x 582 mm (3.39 in. x 17.56 in. x 22.91 in.)
	Cabinet	Maximum dimensions: 2000 mm x 600 mm x 1100 mm (78.74 in. x 23.62 in. x 43.31 in.)		
Weight	Node	Fully loaded with 2.5-inch disks: ≤ 35 kg (77.18 lb)	Fully loaded with 3.5-inch disks: ≤ 70 kg (154.35 lb)	Fully loaded with 3.5-inch disks: ≤ 32 kg (70.56 lb)
	Cabinet	Maximum weight when fully loaded with 2 U nodes: 576.7 kg (1271.62 lb) Maximum weight when fully loaded with 4 U nodes: 725.5 kg (1599.73 lb)		
Typical power		420 W	580 W	260 W
Operating temperature		5°C to 35°C (41°F to 95°F) when the altitude ranges from –60 m to +1800 m (–196.85 ft. to +5905.44 ft.) When the altitude is higher than 1,800 m (5905.44 ft.) but lower than or equal to 3000 m (9842.4 ft.), the ambient temperature drops by 0.6°C (1.08°F) for every 100 m (328.08 ft.) altitude increment.		
Operating humidity		20% RH to 80% RH		

Technical Specifications of the Big Data Analysis Subsystem

Model		Insight Node	
Hardware Specifications			
System architecture		Fully symmetrical distributed architecture	
Number of nodes		3 to 32	
CPUs per node		2 x Intel E5 series	
Cache per node		Standard configuration: 96 GB, expandable to 192 GB	Standard configuration: 64 GB, expandable to 192 GB
Data disk type		2.5-inch 600 GB SAS	3.5-inch 2 TB SATA
Disks per node		Standard configuration: 25 x 2.5-inch 600 GB SAS disks or no disks	Standard configuration: 12 x 3.5-inch 2 TB SATA disks or no disks
Front-end network type		GE or 10GE	GE
Internal network type		10GE	GE
Application scenario		Unstructured and semi-structured data analysis and Hadoop	Structured data archiving
Software Feature			
Hadoop		FusionInsight Hadoop, supporting Sqoop, MapReduce, HBase, and Hive	None
Database		WushanSQL distributed database, supporting quick retrieval of a large amount of structured and unstructured data	
Database interface		JDBC and ODBC	
Data compression		The compression rate is automatically adjusted. The average compression ratio reaches 3:1.	
Quick retrieval		Quick retrieval of massive files (InfoExplorer)	
Data protection		Data written onto two physical nodes for redundancy	
Data recovery		Quick automated parallel data recovery at up to 1 TB per hour	
System expansion		Online expansion	
System management		Support for users of different management rights, and domain- and rights-based user management Alarm notification by email, SMS, SNMP, and Syslog	
Physical Feature			
Power supply		AC 100 V to 127 V, AC 200 V to 240 V	
Dimensions (H x W x D)	Node	2 U, 86.1 mm x 446 mm x 582 mm (3.39 in. x 17.56 in. x 22.91 in.)	
	Cabinet	Maximum size: 2000 mm x 600 mm x 1100 mm (78.74 in. x 23.62 in. x 43.31 in.)	
Node weight		Fully loaded with 2.5-inch disks: ≤ 35 kg (77.18 lb)	Fully loaded with 3.5-inch disks: ≤ 32 kg (70.56 lb)
Typical power		420 W	260 W
Operating temperature		5°C to 35°C (41°F to 95°F) when the altitude ranges from –60 m to +1800 m (–196.85 ft. to +5905.44 ft.) When the altitude is higher than 1800 m (5905.44 ft.) but lower than or equal to 3,000 m (9842.4 ft.), the ambient temperature drops by 0.6°C (1.08°F) for every 100 m (328.08 ft.) altitude increment.	
Operating humidity		20% RH to 80% RH	

HUAWEI TECHNOLOGIES CO., LTD.

Huawei Industrial Base
Bantian Longgang
Shenzhen 518129, P.R. China
Tel: +86-755-28780808

Version No.: M3-035255-20140803-C-3.0

www.huawei.com

OceanStor 9000
Big Data Storage System

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

Designed for Big Data, the HUAWEI OceanStor 9000 storage system employs a symmetric distributed architecture to consolidate data storage, archiving, and analysis. The OceanStor 9000 delivers industry-leading performance, extensive scale-out capabilities, and a super-large single file system to provide shared storage for structured data and unstructured data. The OceanStor 9000 meets many Big Data service application scenarios, including film and TV, satellite mapping, gene sequencing, energy exploration, scientific research, education, and telecom carriers. Based on leading processing performance and data lifecycle management, the OceanStor 9000 helps customers build industry's most efficient Big Data storage capabilities.

OceanStor 9000 Big Data Storage System

Integration of Storage, Analysis, Archive of Big Data

- Data type integration: Centrally stores structured and unstructured data, and quickly retrieves massive volumes of structured and unstructured data.
- Big data life cycle integration: includes life cycles of different subsystems such as the file storage subsystem WushanFS, analysis subsystem WushanSQL, and FusionInsight Hadoop, and supports data storage, analysis, and archiving for improved processing efficiency, simplified management, and reduced investment and cost.
- Management integration: Centrally manages IT devices, provides analysis reports, simplifies management, and improves operation efficiency.

Linear Scaling of Capacity and Performance with the Largest Ever Single File System

- Even data distribution: The shared-nothing symmetric distributed architecture evenly distributes data and metadata data to all nodes, eliminating system bottlenecks.
- Single file system: A single file system of up to 40 PB simplifies system management and maintenance and eliminates data silos caused by multiple namespaces.
- Ultra-high utilization: Ensures up to 95% disk utilization without compromising inter-node data reliability.

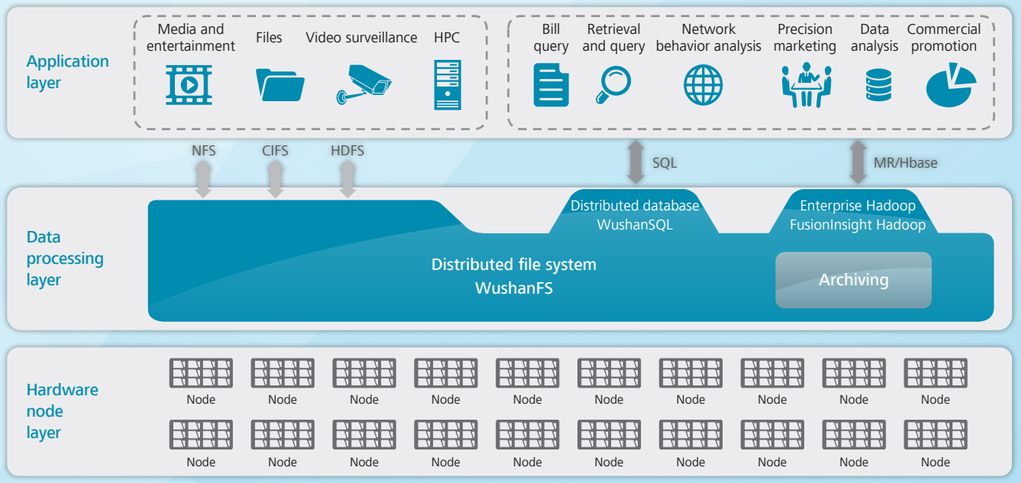

Advanced Symmetric Distributed Architecture for Industry-Leading Performance

- High-performance data access: The OceanStor 9000's performance ranks first in the SPEC sfs2008 benchmark test and exceeds 5 million Operations Per Second (OPS) in terms of Network File System (NFS), providing the ultimate fast-storage user experience.
- High-speed internal interworking: A 10 GE network or a 40 GE InfiniBand network is employed with support for TCP Offload Engine (TOE) and Remote Direct Memory Access (RDMA).
- SSD-based metadata access acceleration: SSDs are used to store metadata, improving metadata access efficiency.
- Global cache: A maximum of 55 TB global cache is provided, increasing the data access hit ratio.

Visualized and Unified Resource Management

- Flexible reliability configuration: Directory-based redundancy ratio policies provide a choice of data protection levels.
- Automatic statistics collection and analysis: Automatic performance statistics collection and analysis help customers use resources efficiently.
- Automatic deployment: A software platform is automatically deployed and configured. One-click capacity expansion is provided. A single node can be added within 60 seconds.
- Rights management: Access control of IP addresses, users, or user groups ensures that storage pools are secure and mutually isolated.

Info Series Software, Managing the OceanStor 9000 Intelligently

InfoEqualizer, implementing client connection load balancing

- Manages access IP addresses in a unified manner and supports automatic allocation, failover, and failback of node IP addresses.
- Implements load balancing based on domain names and supports various load-balancing polices.
- Nodes can be managed based on zones. An independent load-balancing policy and an independent domain name can be configured for each zone.

InfoTier, implementing DST

- DST is implemented between performance and capacity nodes, fully leveraging the advantages of different types of storage media and reducing Total Cost of Ownership (TCO).
- Various data migration policies and migration priorities are supported, adapting to changing service needs.

InfoAllocator, implementing space quota management

- Space quota management is implemented based on users, user groups, and directories, meeting different customer requirements.
- Quota nesting management allows flexible and easy access to storage space.

InfoProtector, implementing reliable data protection

- InfoProtector is Huawei''s proprietary N+M data protection technology that protects data against a concurrent failure of four nodes.
- Multiple nodes restore data concurrently at a speed of up to 1 TB/hour.

InfoExplorer for quick retrieval of massive volumes of data

- Retrieval time is shortened from dozens of hours to several seconds. The search efficiency is improved by several thousand times.
- The built-in full-text retrieval function is used to support fuzzy search based on the file name, path, user name, and user-defined tag.

Technical Specifications of the Big Data File Storage Subsystem

Model	Performance Node	Capacity Node	Mini-Capacity Node
Hardware Specifications			
System architecture	Symmetric distributed architecture		
Number of nodes	3 to 288		
CPUs per node	2 x Intel E5 series		
Cache per node	Standard configuration: 48 GB, expandable to 192 GB		Standard configuration: 32 GB, expandable to 192 GB
Disk type	SSD and SAS	SSD, SATA and NL-SAS	
Number of disks per node	Standard configuration: 4 x 2.5-inch 200 GB SSDs + 21 x 2.5-inch 600 GB SAS disks (Based on actual performance requirements, the SSD/HDD configuration ratio can be adjusted.)	Standard configuration: 1 x 3.5-inch 200 GB SSD + 35 x 3.5-inch 4 TB SATA disks (Based on actual performance requirements, the SSD/HDD configuration ratio can be adjusted.)	Standard configuration: 12 x 3.5-inch 2 TB SATA disks (Based on actual performance requirements, the SSD/HDD configuration ratio can be adjusted.)
Front-end network type	10GE, 40GE InfiniBand, and 1GE		
Internal network type	10GE and 40GE InfiniBand		
Application scenario	OPS-intensive	High-bandwidth	Small-capacity
Software Feature			
Data protection level	N+1, N+2, N+3, and N+4		
File system	Wushan distributed file system, which supports global namespace and can be dynamically expanded up to 40 PB		
Value-added feature	Dynamic storage tiering (InfoTier) Automatic client connection load balancing (InfoEqualizer) Space quota management (InfoAllocator)		
Thin provisioning	Support for thin provisioning, which does not need to be configured		
Data self-healing	Automatic, concurrent, and quick data restoration, with the maximum restoration speed of 1 TB/hour		
System expansion	One-click online expansion, with less than 60 seconds needed for expansion of a single node		
System management	Rights- and domain-based user management		
Global cache	Up to 55 TB		
Supported operating system	Windows, Linux, Mac OS		