

OceanStor 9000 V300R006C00

# **Product Description**

Issue 01

Date 2016-09-30



#### Copyright © Huawei Technologies Co., Ltd. 2016. 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.

#### **Trademarks and Permissions**

All other trademarks and trade names mentioned in this document are the property of their respective holders.

#### **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

# Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base

Bantian, Longgang Shenzhen 518129

People's Republic of China

Website: <a href="http://e.huawei.com">http://e.huawei.com</a>

# **About This Document**

This document describes the appearance, typical configuration, networking, application scenarios, functions, hardware structure, software structure, technical specifications, standards compliance, and certifications of the OceanStor 9000 big data storage system.

# **Intended Audience**

This document is intended for:

- Marketing engineers
- Planning engineers
- Maintenance engineers
- Training engineers

# **Symbol Conventions**

The symbols that may be found in this document are defined as follows:

Symbol	Description
<b>DANGER</b>	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
<b>WARNING</b>	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
<b>A</b> CAUTION	Indicates a potentially hazardous situation which, if not avoided, may result in minor or moderate injury.
NOTICE	Indicates a potentially hazardous situation which, if not avoided, could result in equipment damage, data loss, performance deterioration, or unanticipated results.
	NOTICE is used to address practices not related to personal injury.

Symbol	Description
NOTE	Calls attention to important information, best practices and tips.
	NOTE is used to address information not related to personal injury, equipment damage, and environment deterioration.

# **Change History**

Changes between document issues are cumulative. The latest document issue contains all the changes made in earlier issues.

# Issue 01 (2016-09-30)

This issue is the first official release.

# **Contents**

About This Document	ii
1 Positioning	1
2 Product Highlights	3
3 Product Architecture	
3.1 Hardware Architecture	8
3.1.1 Hardware Components.	8
3.1.2 Storage Nodes.	
3.1.3 Switches.	64
3.1.4 FR42612L Cabinets	74
3.1.5 Optional Hardware	
3.2 Software Architecture	78
3.3 Network Architecture.	79
4 Product Features	81
4.1 Service Function.	82
4.1.1 CIFS/NFS/FTP Sharing	82
4.1.2 CIFS-NFS Cross-Protocol Access Sharing.	82
4.1.3 Object Storage Service (Compatible with Amazon S3 APIs)	84
4.1.4 Object Storage Service (Compatible with OpenStack Swift APIs)	86
4.1.5 HDFS Interface	87
4.1.6 NDMP Backup	88
4.1.7 High Data Reliability (InfoProtector)	89
4.1.8 Permission Control.	90
4.1.9 Windows Hard Link	91
4.1.10 Unified Namespace	92
4.2 Management Function.	93
4.2.1 Data Consistency Scanning.	93
4.2.2 High System Security	93
4.2.3 Flexible Expansion.	94
4.2.4 Centralized Management	
4.2.5 Data Balancing	
4.2.6 Unified Network Management	95
4.2.7 Unified View	96

4.3 File System Basic Features	
4.3.1 InfoTier	
4.3.2 InfoAllocator	98
4.3.3 InfoEqualizer	99
4.3.4 InfoLocker	
4.3.5 InfoStamper	101
4.3.6 InfoReplicator.	102
4.3.7 InfoScanner.	
4.3.8 InfoMigrator	105
4.4 File System Enhanced Features	106
4.4.1 InfoRevive	106
4.4.2 InfoTurbo	107
4.4.3 Rapid internal file replication	
5 Typical Applications	115
5.1 Application Scenarios	116
5.2 Typical System Configuration	121
6 Operation and Maintenance	124
7 Technical Specifications	127
7.1 Performance and Capacity Specifications	128
7.2 Overall System Hardware Specifications	129
7.3 Reliability Specifications.	131
7.4 Environmental Specifications	
8 Standards Compliance	134
A Glossary	138

# 1 Positioning

As a big data storage product, the OceanStor 9000 provides advantages such as large capacity, high performance, and flexible scalability as well as a range of value-added features. It can be applied to various fields such as broadcasting, media and entertainment (M&E), high-performance computing (HPC), Internet operation, data centers, and large-sized enterprises.

Big Data is high-volume, high-velocity, and high-variety information assets that demand an effective information processing mode for enhanced insight, decision making, and process optimization. Features of Big Data are: Volume, Velocity, and Variety. The arrival of the Big Data era places the following requirements on storage systems:

- Volume includes capacity and quantity and requires flexible storage scalability.
- Velocity requests distributed processing and storage, a scale-out architecture, common hardware, and high performance with controllable costs.
- Variety requires storage to provide a variety of interfaces to address different storage requirements.

OceanStor 9000 is a Big Data storage platform that uses the general x86 architecture hardware and a scale-out NAS system. OceanStor 9000 adopts a fully symmetric architecture and supports seamless expansion of 3 to 288 nodes. The following functions are supported:

- Huawei self-developed distributed file system OceanStor DFS can store massive unstructured data and supports global unified namespace.
- Provides object storage service compatible with Amazon S3 APIs and OpenStack Swift APIs. Supports migration from the media cloud to the cloud computing architecture, and access to mainstream cloud computing systems.
- Interconnects with FusionInsight Hadoop/Cloudera Hadoop consisting of open source Hadoop components to help users build enterprise-class Big Data analysis platforms easily.

The OceanStor 9000 employs a networking mechanism that features full mesh and redundancy, allowing any nodes to concurrently access any data. Multiple nodes can concurrently access different areas of the same data, achieving high-performance data read and write. The OceanStor 9000 provides standard NFS/CIFS/FTP/HDFS/object storage service compatible with Amazon S3 and OpenStack Swift APIs

Figure 1-1 shows the system architecture of the OceanStor 9000.

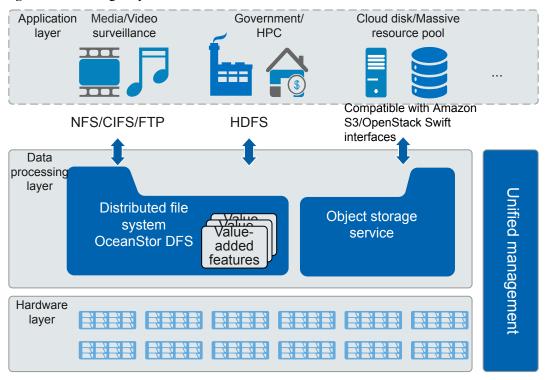


Figure 1-1 Converged system architecture

# Product Highlights

The OceanStor 9000 features outstanding performance, flexible scalability, openness and convergence, robust reliability, and simplified management, helping enterprises reduce the capital expenditure (CAPEX) and operating expenditure (OPEX).

#### **Outstanding Performance**

The OceanStor 9000 delivers industry-leading read/write performance that fully addresses the performance requirements of different fields such as media editing and high-performance computing (HPC).

The OceanStor 9000 provides the following key technologies for file system applications:

#### Concurrent access

In a traditional NAS system consisting of an engine and storage units, concurrent data access requests must be processed by the engine first. Therefore, engines are prone to a performance bottleneck.

The OceanStor 9000 employs a fully symmetric and logical architecture where all nodes directly and externally provide services and workloads are balanced among nodes in a cluster, improving the processing capability of concurrent requests.

#### Global cache

Accessing data in the cache is much faster than doing so on a disk. However, cache space on nodes is limited and independent from each other, resulting in a low cache hit ratio and complex data consistency protection.

The OceanStor 9000 provides the global cache mechanism to resolve these problems. A virtual cache pool is provided to incorporate all node cache. The cache pool stores only one copy of a file and is accessible to any node. The global cache mechanism elevates system performance and improves data hit ratio while shortening the latency.

#### High-speed node interconnection

The OceanStor 9000 supports 10GE high-speed networks. Each OceanStor 9000 node provides four 10GE ports. The OceanStor 9000 also supports node interconnection over an InfiniBand network to meet more demanding performance requirements.

With high-speed node interconnection, the OceanStor 9000 meets different networking requirements and delivers lower internal latency and better performance.

• SSD acceleration and dynamic storage tiering (DST)

High-performance solid state drives (SSDs) help eliminate I/O performance bottlenecks, lower latency, and increase throughput. The OceanStor 9000 can automatically identify

hotspot data and small files, and migrate data between different storage tiers based on data access frequency to improve resource utilization.

#### SMB3 Multichannel

The OceanStor 9000 supports the SMB3 Multichannel function. By bringing multi-core CPUs and bandwidth resources of clients into full play, SMB3 Multichannel greatly improves service performance and reliability. In addition, if one channel fails, SMB3 Multichannel transmits data over another channel to prevent services from being affected.

#### • Rapid internal file replication

Internal file replication is rapidly implemented within the OceanStor 9000 without involving clients.

#### NFS protocol enhancement

The traditional NFS protocol is limited by bandwidth of a single network port and block size. A client supports a maximum bandwidth of 1 GB/s. The OceanStor 9000 provides enhanced features of the NFS protocol. Multiple network ports and the NFS protocol plug-in DFSClient are configured on the NFS client for concurrent connections of multiple network ports and cache optimization to improve performance of the client.

- Bandwidth of a Mac OS X client is 1.5 GB/s.
- Bandwidth of an RHEL client is 2.5 GB/s.

#### Disk defragmentation

The OceanStor 9000 provides the disk defragmentation function. In a specified period of time, periodically executing disk defragmentation tasks ensures satisfactory system performance for a long time.

The OceanStor 9000 provides the following key technologies for object storage applications:

#### Concurrent access

The OceanStor 9000 employs a symmetric and logical architecture where each node can provide services. The load balancing design enables data access requests to be evenly distributed in a cluster, greatly improving concurrent access capability.

#### High-speed interconnectivity

The OceanStor 9000 supports 10GE high-speed networks, and each node provides four 10GE ports.

By interconnecting nodes at a high speed, OceanStor 9000 reduces the internal latency and provides outstanding performance, adapting to different networking requirements.

#### SSD acceleration

SSDs provide low latency and high throughput to eliminate I/O performance bottlenecks.

#### • Disk defragmentation

The OceanStor 9000 provides the disk defragmentation function. In a specified period of time, periodically executing disk defragmentation tasks ensures satisfactory system performance for a long time.

#### Flexible Scalability

With a distributed architecture, the OceanStor 9000 supports the flexible and seamless scaleout from 3 to 288 nodes. All nodes are symmetrically arranged in the OceanStor 9000 cluster, and all data and metadata are evenly distributed across each node to eliminate system resource contention and system bottlenecks. The OceanStor 9000 can be smoothly scaled out in line with service growth, enabling linear increase of capacity and performance. The OceanStor 9000 offers up to 100 PB capacity for a single file system. The flexible scalability of the OceanStor 9000 helps minimize the CAPEX and OPEX while adapting to changing storage requirements.

The OceanStor 9000 provides the following node types to fit into different application scenarios:

- P12/P12E, P25/P25E, and P36/P36E are high-performance storage nodes and oriented to OPS-intensive and high-bandwidth applications.
- C36/C36E and C72 are large-capacity archive nodes applicable to large-capacity application scenarios.

#### **Openness and Convergence**

The OceanStor 9000 is an open and converged storage system. It supports multiple interfaces.

- Supports multiple interfaces for object storage, HDFS, NFS, CIFS, NDMP, and FTP to support files and data.
- Interconnects with public cloud systems such as Amazon and OpenStack.
- Supports multiple nodes such as OPS-intensive, bandwidth-intensive, and large-capacity nodes to carry various applications.
- Both the front-end service network and management network support IPv4 and IPv6.
   However, IPv4 and IPv6 cannot be used together for the front-end service network but can be used together for the management network.

#### **Robust Reliability**

The OceanStor 9000 employs a distributed cluster architecture where all components are fully redundant, eliminating single points of failure. Key components of storage nodes, such as fan modules, power modules, and interface modules, adopt a redundant design to improve hardware reliability. Other reliability functions include:

• N+M data protection level

Erasure Code enables inter-node RAID, ensuring data integrity of a node even though the entire node fails. Multiple protection levels are supported, such as N+1, N+2, N+3, N +4, N+2:1, and N+3:1. You can specify a protection level for any empty directory in a file system.

Mirroring

Mirroring helps OceanStor DFS store multiple copies of metadata onto different storage nodes for improved data protection.

Data reconstruction

Automatic node fault detection and automatic data reconstruction are supported. When a disk or node fails, data can be retrieved, preventing data loss upon node or disk failures.

Physical domain division

Physical domain division is an efficient way to isolate faulty domains from functioning ones. When a disk or node in a domain (node pool or node pool group) fails, physical domain division ensures that reliability and security of data in other domains are not affected, so is the I/O performance of other domains during data recovery.

Data consistency scanning

Data consistency scanning, also called data scrub, periodically checks data in the background without affecting services, thereby discovering disk errors in advance to prevent silent corruption from causing data loss.

The OceanStor 9000 also provides a variety of advanced data protection technologies that ensure data availability even in a catastrophe, thereby protecting service continuity.

- InfoReplicator: It is a remote replication feature that allows you to set an automatic replication cycle or manually start a data replication task, so that the specified directories or files are replicated from one storage system to other storage systems over IP links (either a LAN or a WAN), thereby maximizing the disaster recovery capability.
- Network Data Management Protocol (NDMP): The OceanStor DFS acts as a data source, and data is backed up to backup media such as tape libraries for data backup and recovery purposes.
- InfoStamper: It is a directory-level snapshot feature that allows you to create snapshots of any directory (except the root directory) in a file system, so that important data can be easily backed up.
- InfoRevive: In video surveillance scenarios, if the number of failed nodes or disks exceeds the redundancy upper limit, a certain percentage of damaged video file data can still be read and data can be written to the damaged video file, thereby enhancing the availability, security, and continuity of system services.

#### Simplified Management

The OceanStor 9000 supports graphical automatic deployment, flexible space quotas, comprehensive performance monitoring, unified network topology management, and unified namespace. Management and maintenance are efficient, centralized, simple, and easy.

- Administrators can use DeviceManager to manage hardware, software, clusters, and services in a unified manner.
- The OceanStor 9000 adopts flexible quota space to dispatch storage resources at three levels: directory, user, and user group.
- If multiple OceanStor 9000 storage systems are used in application scenarios such as video surveillance and remote replication, eSight can be used to centrally manage and maintain these storage systems.
- SystemReporter, which is the performance analysis tool of the OceanStor 9000, uses data collection, archiving, analysis, and prediction methods to provide real-time monitoring, historical trend analysis, comparison-based analysis, diagnostic analysis, and capacity prediction for you to better use resources.
- System software is automatically deployed and configured. One-click rapid capacity expansion is supported.

# 3 Product Architecture

# **About This Chapter**

This chapter describes the hardware, software, and network architecture of the OceanStor 9000.

#### 3.1 Hardware Architecture

This section describes the hardware architecture and hardware components of the OceanStor 9000.

#### 3.2 Software Architecture

This section describes the software architecture of the OceanStor 9000.

#### 3.3 Network Architecture

The OceanStor 9000 network architecture contains the management network, front-end service network, and back-end storage network.

# 3.1 Hardware Architecture

This section describes the hardware architecture and hardware components of the OceanStor 9000.

# 3.1.1 Hardware Components

The OceanStor 9000 hardware contains storage nodes, network devices, keyboard, video, and mouse (KVM) devices, and short message service (SMS) modems.

Table 3-1 lists the OceanStor 9000 hardware components.

**Table 3-1** Hardware components

Hardware Type	Model	Description
(Recommended) Cabinet	FR42612L	Providing 42 U installation space
Storage node	P12E	2 U 12-slot node with higher performance (configuration example: 12 x SATA disks or 1 x SSD+11 x SATA disks)
	P25E	2 U 25-slot node with higher performance (configuration example: 1 x SSD+24 x SAS disks)
	P36E	4 U 36-slot node with higher performance (configuration example: 1 x SSD+35 x SATA disks)
	C36E	4 U 36-slot node with higher performance (configuration example: 36 x SATA disks)
	P12	2 U 12-slot node (configuration example: 12 SATA disks or 1 x SSD+11 SATA disks)
	P25	2 U 25-slot node (configuration example: 1 x SSD+24 SAS disks)
	P36	4 U 36-slot node (configuration example: 1 x SSD+35 SATA disks)
	C36	4 U 36-slot node (configuration example: 36 x SATA disks)
	C72	4 U 72-slot node (configuration example: 72 x SATA disks)
Network device	CE6810-48S4Q-EI/ CE6810-24S2Q-LI	10GE switch

Hardware Type	Model	Description
	S5700-52C-SI/ S5352C-SI	GE switch
	SX6018	InfiniBand switch
8-port KVM controller		Providing eight KVM ports
Modem		Providing SMS-based alarm notification

# 3.1.2 Storage Nodes

The OceanStor 9000 storage nodes include P12E, P25E, P36E, C36E, P12, P25, P36, C36, and C72.

#### **P12E**

#### **Function**

Based on requirements of the Internet, Internet Data Center (IDC), cloud computing, enterprise markets, and telecom service applications, Huawei launches the P12E nodes. With advantages such as high-performance computing, large-capacity storage, low power consumption, powerful expansion capabilities, robust reliability, easy deployment and management, and support for virtualization, P12E nodes are suitable for storage services including distributed storage, data mining, electronic photo albums, videos, basic applications of enterprises, and telecom service applications.

#### **Exterior**

#### **Entire device**

Figure 3-1 shows the P12E overall appearance.

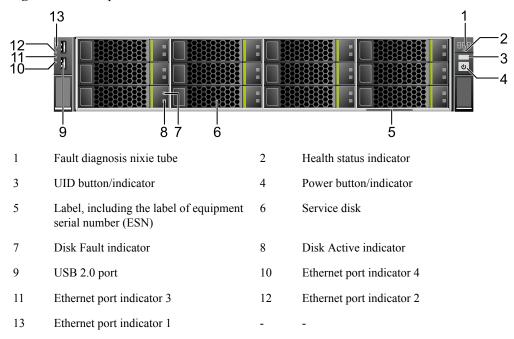
Figure 3-1 P12E overall appearance



#### Front view

Figure 3-2 shows the front panel of a P12E node.

Figure 3-2 Front panel



The front panel of a P12E node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in Figure 3-3.

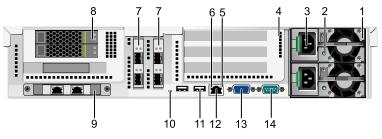
Figure 3-3 Disk slots on the front panel of a P12E node



#### Rear view

Figure 3-4 shows the rear panel of a P12E node.

Figure 3-4 Rear panel of a P12E node



1 Power module

- 2 Power module indicator
- 3 Power socket of the power module
- I/O module (SLOT6, SLOT7, and SLOT8 from up to bottom)

5 Link indicator

- 6 Data transfer indicator
- 7 Onboard PCIe slot (SLOT4 and SLOT5 from left to right)
- 8 System disk

4

- 9 Onboard network interface card (NIC)
- 10 UID indicator

11 USB 3.0 port

12 Management network port

13 VGA port

14 Serial port

#### **APIs**

**Table 3-2** and **Table 3-3** show ports provided by a P12E node.

**Table 3-2** Ports on the front panel

Name	Туре	Quantity	Description
USB port	USB 2.0	2	The USB port is connected to a USB device.  NOTICE  Before connecting to an external USB device, ensure that the USB device functions properly. Otherwise, the server may operate abnormally.

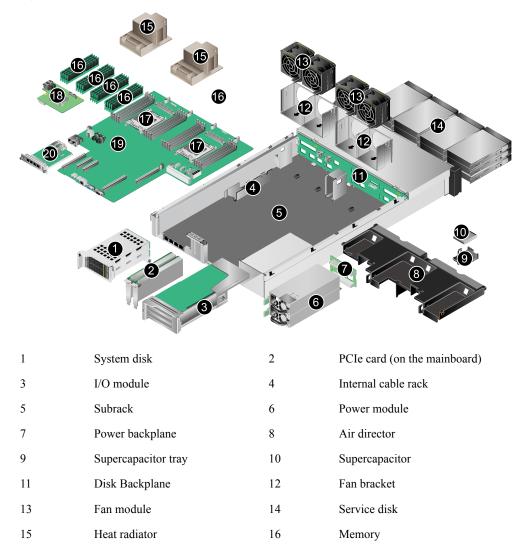
**Table 3-3** Ports on the rear panel

Port	Туре	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor or KVM.
USB port	USB3.0	2	The USB port is connected to a USB device.  NOTICE  Before connecting to an external USB device, check that the USB device operates properly. A server may fail if it is connected to an abnormal USB device.  The USB 3.0 is disabled by default. You can enable it in the BIOS. When you use the VMware OS, do not enable the USB 3.0. Otherwise, the keyboard and mouse are unavailable.
Mgmt management network port	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the server.
Serial port	DB9	1	This port complies with RS232. The baud rate is 115200 bit/s.  The serial port is used as the system serial port by default. You can set it to the iMana 200 serial port by using the iMana 200 command. The port is used for debugging.
Network port	-	-	The port types and quantity vary according to the configured NIC.

# **Physical Structure**

**Figure 3-5** shows the basic structure of a P12E node.

Figure 3-5 Basic structure



# **Technical Specifications**

17

19

**Table 3-4** shows technical specifications of a P12E node.

Table 3-4 P12E technical specifications

CPU

Main board

Category	Item	Specifications
Physical specifications	Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul><li>Net weight: 30 kg</li><li>Weight of packing materials: 5 kg</li></ul>

18

20

RAID controller card

NIC

Category	Item	Specifications
Environmental	Operating temperature	5°C to 35°C
specifications	Operating humidity	20% RH to 80% RH
Power supply specifications	Input voltage	1200 W high-voltage AC power supply module: 1200 W (input: 100 V AC to 240 V AC)
Power consumption specifications	Maximum power consumption	350 W
Configuration examples	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 1
	Memory	DDR4 RDIMM: 32 GB NVDIMM: 8 GB
	Disk	• 2 x 300 GB 2.5-inch back-end operating system disks (SAS)
		<ul> <li>The following two configuration modes are available for service disks:</li> </ul>
		- 12 x 2 TB/4 TB/6 TB/10 TB 3.5-inch service disks (SATA) in slots 0 to 11
		- 1 x 600 GB or 900 GB 3.5- inch service disk (SSD) in slot 0 and 11 x 2 TB/4 TB/6 TB/10 TB 3.5-inch service disks (SATA) in slots 1 to 11
		NOTE  The mixing of 4KN and 512N/512e <sup>d</sup> system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

#### **P25E**

#### **Function**

Based on requirements of the Internet, Internet Data Center (IDC), cloud computing, enterprise markets, and telecom service applications, Huawei launches the P25E nodes. With advantages such as high-performance computing, large-capacity storage, low power consumption, powerful expansion capabilities, robust reliability, easy deployment and management, and support for virtualization, P12E nodes are suitable for storage services

including distributed storage, data mining, electronic photo albums, videos, basic applications of enterprises, and telecom service applications.

#### **Exterior**

#### **Entire device**

Figure 3-6 shows the P25E overall appearance.

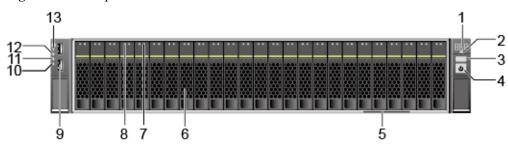
Figure 3-6 P25E overall appearance



#### Front view

Figure 3-7 shows the front panel of a P25E node.

Figure 3-7 Front panel



- 1 Fault diagnosis nixie tube
- 3 UID button/indicator
- 5 Label, including the ESN
- 7 Disk Fault indicator
- 9 USB 2.0 port
- 11 Ethernet port indicator 3
- 13 Ethernet port indicator 1

- 2 Health status indicator
- 4 Power button/indicator
- 6 Service disk
- 8 Disk Active indicator
- 10 Ethernet port indicator 4
- 12 Ethernet port indicator 2
- \_

The front panel of a P25E provides 25 service disk slots numbered from 0 to 24 from left to right, as shown in **Figure 3-8**.

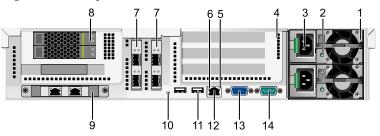
Figure 3-8 Disk slots on the front panel of a P25E node



#### Rear view

Figure 3-9 shows the rear panel of a P25E node.

Figure 3-9 Rear panel



- 1 Power module 2 Power module indicator
- 3 Power socket of the power module 4 I/O module (SLOT6, SLOT7, and SLOT8 from up to bottom)
- 5 Link indicator 6 Data transfer indicator
- 7 Onboard PCIe slot (SLOT4 and SLOT5 8 System disk from left to right)
- 9 Onboard network interface card (NIC) 10 UID indicator
- 11 USB 3.0 port 12 Management network port
- 13 VGA port 14 Serial port

#### **APIs**

**Table 3-5** and **Table 3-6** show ports provided by a P25E node.

**Table 3-5** Ports on the front panel

Name	Туре	Quantity	Description
USB port	USB 2.0	2	The USB port is connected to a USB device.
			NOTICE  Before connecting to an external USB device, ensure that the USB device functions properly. Otherwise, the server may operate abnormally.

**Table 3-6** Ports on the rear panel

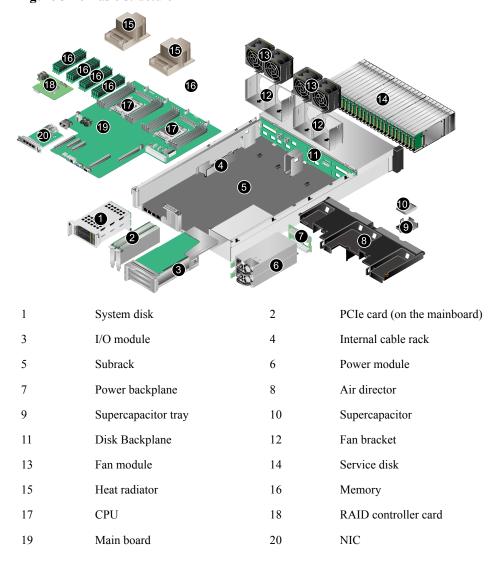
Port	Туре	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor or KVM.
USB port	USB3.0	2	The USB port is connected to a USB device.  NOTICE  Before connecting to an external USB device, check that the USB device operates properly. A server may fail if it is connected to an abnormal USB device.  The USB 3.0 is disabled by default. You can enable it in the BIOS. When you use the VMware OS, do not enable the USB 3.0. Otherwise, the keyboard and mouse are unavailable.
Mgmt management network port	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the server.
Serial port	DB9	1	This port complies with RS232. The baud rate is 115200 bit/s.
			The serial port is used as the system serial port by default. You can set it to the iMana 200 serial port by using the iMana 200 command. The port is used for debugging.

Port	Туре	Quantity	Description
Network port	-	-	The port types and quantity vary according to the configured NIC.

#### **Physical Structure**

Figure 3-10 shows the basic structure of a P25E node.

Figure 3-10 Basic structure



### **Technical Specifications**

**Table 3-7** shows technical specifications of a P25E node.

**Table 3-7** P25E technical specifications

Category	Item	Specifications		
Physical specifications	Dimensions (H x W x D)	86.1 mm x 447 mm x 708 mm, 2 U (1 U = 44.45 mm)		
	Weight	Net weight: 30 kg		
		Weight of packing materials: 5 kg		
Environmental	Operating temperature	5°C to 35°C		
specifications	Operating humidity	20% RH to 80% RH		
Power supply specifications	Input voltage	1200 W high-voltage AC power supply module: 1200 W (input: 100 V AC to 240 V AC)		
Power consumption specifications	Maximum power consumption	500 W		
Configuration	Processor	Model: Intel Xeon 8-core 2.1 GHz		
examples		Quantity: 2		
	Memory	DDR4 RDIMM: 48 GB		
		NVDIMM: 8 GB		
	Disk	<ul> <li>2 x 300 GB 2.5-inch back-end operating system disks (SAS)</li> <li>1 x 600 GB or 900 GB 2.5-inch</li> </ul>		
		service disk (SSD) in slot 0		
		• 24 x 2.5-inch 900 GB, 1.2 TB, or 1.8 TB service disks in slots 1 to 24		
		NOTE  The system does not support a combination of 4KN disks and 512N/ 512e disks <sup>d</sup> or a combination of 300 GB disks and 600 GB disks.		

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

#### P36E

#### **Function**

A P36E node is the Huawei new-generation 4 U dual-socket or single-socket rack server that adopts efficient design to ensure excellent computing performance as well as flexible and large-capacity local storage. It is suitable for storage-intensive applications, such as cold data storage, video surveillance, cloud storage, and Big Data applications, widely applied in media, financing, public security, and other industries.

#### **Exterior**

#### **Entire device**

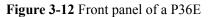
Figure 3-11 shows the entire appearance of a P36E.

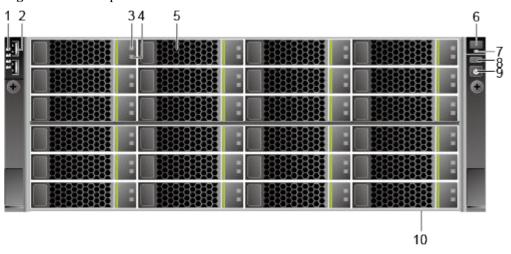
Figure 3-11 Appearance



#### Front view

Figure 3-12 shows the front view of a P36E.





- 1 Ethernet port indicator
- 3 Disk Fault indicator
- 5 Service disk
- 7 Health status indicator
- 9 Power button/indicator

- 2 USB 2.0 port
- 4 Disk Active indicator
- 6 Fault diagnosis nixie tube
- 8 UID button/indicator
- 10 Label, including the label of ESN

The front panel of a P36E provides 24 service disk slots numbered from 0 to 23 from left to right, as shown in **Figure 3-13**.

Figure 3-13 Disk slots on the front panel of a P36E

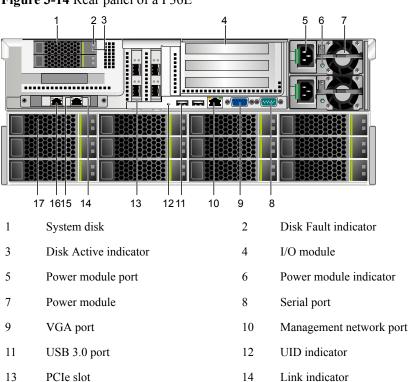


#### Rear view

Figure 3-14 shows the rear view of a P36E.

#### NOTE

This section uses the 2-port TOE 10 GE card as an example.

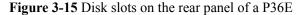


16

Service network port

Figure 3-14 Rear panel of a P36E

Figure 3-15 shows disk slots on the front panel of a P36E.

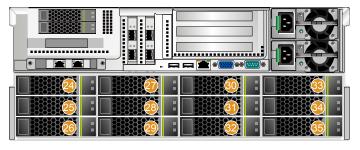


Data transfer indicator

Service disk

15

17



#### **APIs**

**Table 3-8** and **Table 3-9** show ports provided by a P36E.

**Table 3-8** Ports on the front panel

Name	Туре	Quantity	Description
USB port	USB 2.0	2	The USB port is connected to a USB device.
			NOTICE  Before connecting to an external USB device, ensure that the USB device functions properly. Otherwise, the server may operate abnormally.

 Table 3-9 Ports on the rear panel

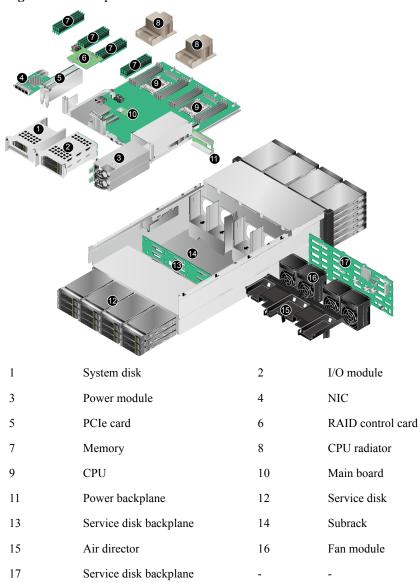
Port	Туре	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor or KVM.
USB port	USB3.0	2	The USB port is connected to a USB device.  NOTICE  Before connecting to an external USB device, check that the USB device operates properly. A server may fail if it is connected to an abnormal USB device.  The USB 3.0 is disabled by default. You can enable it in the BIOS. When you use the VMware OS, do not enable the USB 3.0. Otherwise, the keyboard and mouse are unavailable.
Mgmt management network port	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the server.
Serial port	DB9	1	This port complies with RS232. The baud rate is 115200 bit/s.
			The serial port is used as the system serial port by default. You can set it to the iMana 200 serial port by using the iMana 200 command. The port is used for debugging.

Port	Туре	Quantity	Description
Network port	-	-	The port types and quantity vary according to the configured NIC.

# **Physical Structure**

Figure 3-16 shows components of a P36E.

Figure 3-16 Components of a P36E



#### **Technical Specifications**

**Table 3-10** shows technical specifications of a P36E node.

**Table 3-10** P36E technical specifications

Category	Item	Specifications	
Physical specifications	Dimensions (H x W x D)	175 mm x 447 mm x 748 mm, 4 U (1 U = 44.45 mm)	
	Weight	<ul> <li>Net weight: 57 kg</li> <li>Weight of packing materials: 15 kg</li> </ul>	
Environmental	Operating temperature	5°C to 35°C	
specifications	Operating humidity	20% RH to 80% RH	
Power supply specifications	Input voltage	1200 W high-voltage AC power supply module: 1200 W (input: 200 V AC to 240 V AC)	
Power consumption specifications	Maximum power consumption	750 W	
Configuration examples	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2	
	Memory	DDR4 RDIMM: 48 GB	
		NVDIMM: 8 GB	
	Disk	• 2 x 300 GB 2.5-inch back-end operating system disks (SAS)	
		• 1 x 600 GB or 900 GB 3.5- inch service disk (SSD) in slot 0	
		• 35 x 2 TB/4 TB/6 TB/10 TB 3.5-inch service disks (SATA) in slots 1 to 35	
		NOTE	
		The mixing of 4KN and 512N/512ed system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.	

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

#### **C36E**

#### **Function**

C36E is the Huawei new-generation 4 U dual-socket or single-socket rack server that adopts efficient design to ensure excellent computing performance as well as flexible and large-capacity local storage. It is suitable for storage-intensive applications, such as cold data storage, video surveillance, cloud storage, and Big Data applications, widely applied in media, financing, public security, and other industries.

#### **Exterior**

#### **Entire device**

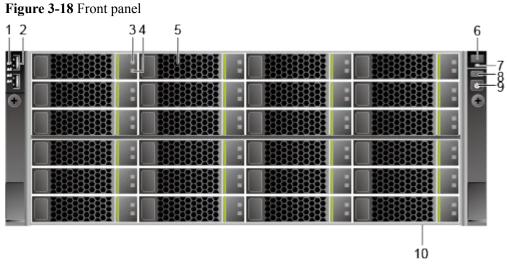
Figure 3-17 shows the overall appearance.

Figure 3-17 Appearance



#### Front view

Figure 3-18 shows the front panel of a C36E node.





The front panel of a C36E node provides 24 service disk slots numbered from 0 to 23 from left to right, as shown in **Figure 3-19**.

Figure 3-19 Disk slots on the front panel of a C36E node



#### Rear view

Figure 3-20 shows the rear panel of a C36E node.

#### NOTE

This section uses the 2-port TOE 10 GE card as an example.

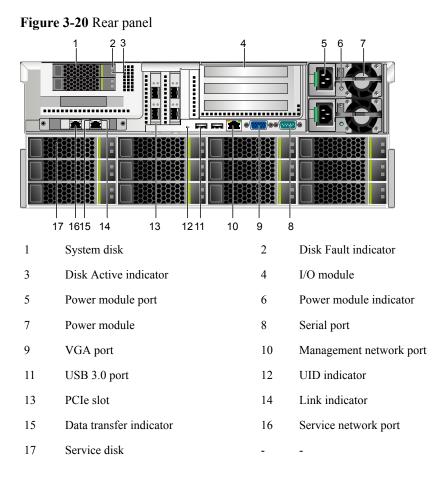
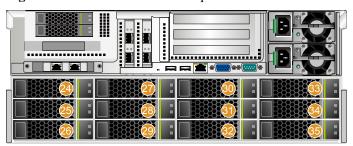


Figure 3-21 shows disk slots on the rear panel of a C36E node.

Figure 3-21 Disk slots on the rear panel



#### **APIs**

**Table 3-11** and **Table 3-12** show ports provided by a C36E node.

Table 3-11 Ports on the front panel

Name	Туре	Quantity	Description
USB port	USB 2.0	2	The USB port is connected to a USB device.  NOTICE
			Before connecting to an external USB device, ensure that the USB device functions properly. Otherwise, the server may operate abnormally.

Table 3-12 Ports on the rear panel

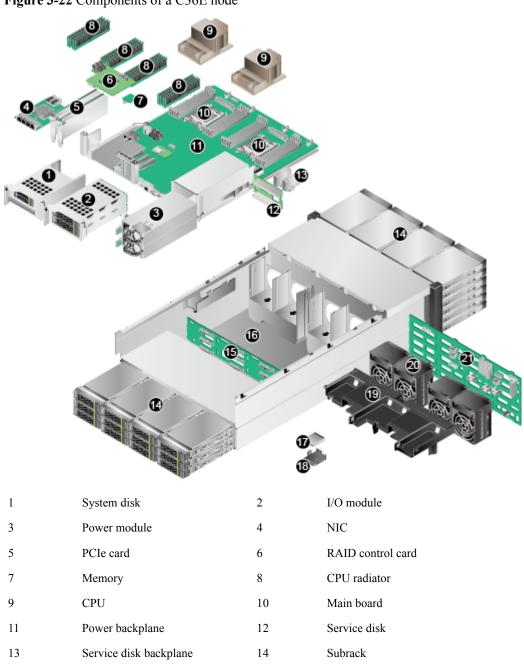
Port	Туре	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor or KVM.
USB port	USB3.0	2	The USB port is connected to a USB device.  NOTICE  Before connecting to an external USB device, check that the USB device operates properly. A server may fail if it is connected to an abnormal USB device.  The USB 3.0 is disabled by default. You can enable it in the BIOS. When you use the VMware OS, do not enable the USB 3.0. Otherwise, the keyboard and mouse are unavailable.
Mgmt management network port	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the server.
Serial port	DB9	1	This port complies with RS232. The baud rate is 115200 bit/s.
			The serial port is used as the system serial port by default. You can set it to the iMana 200 serial port by using the iMana 200 command. The port is used for debugging.

Port	Туре	Quantity	Description
Network port	-	-	The port types and quantity vary according to the configured NIC.

# **Physical Structure**

Figure 3-22 describes the components of a C36E node.

Figure 3-22 Components of a C36E node



15	Air director	16	Fan module
17	Service disk backplane	_	-

# **Technical Specifications**

Table 3-13 shows technical specifications of a C36E node.

**Table 3-13** C36E technical specifications

Category	Item	Specifications
Physical specifications	Dimensions (H x W x D)	175 mm x 447 mm x 748 mm, 4 U (1 U = 44.45 mm)
	Weight	• Net weight: 57 kg
		<ul> <li>Weight of packing materials:</li> <li>15 kg</li> </ul>
Environmental	Operating temperature	5°C to 35°C
specifications	Operating humidity	20% RH to 80% RH
Power supply specifications	Input voltage	1200 W high-voltage AC power supply module: 1200 W (input: 200 V AC to 240 V AC)
Power consumption specifications	Maximum power consumption	700 W
Configuration examples	Processor	Model: Intel Xeon 8-core 2.1 GHz
		Quantity: 1
	Memory	DDR4 RDIMM: 32 GB
		NVDIMM: 8 GB
	Disk	• 2 x 300 GB 2.5-inch back-end operating system disks (SAS)
		• 36 x 2 TB/4 TB/6 TB/10 TB 3.5-inch service disks (SATA) in slots 0 to 35
		NOTE
		The mixing of 4KN and 512N/512e <sup>d</sup> system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.

d: Disk sector format. 4KN means that the sector size is 4KB. 512e means that the sector size is 4KB and allows an integral multiple of 512B data to be written. 512N means that the sector size is 512B.

### P12

## **Function**

A P12 node is a high-performance storage node applicable to OPS-intensive applications, equipped with 12 front service disk slots and 2 rear 2.5-inch system disk slots.

### **Exterior**

#### Overview

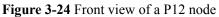
Figure 3-23 shows the exterior of a P12 node.

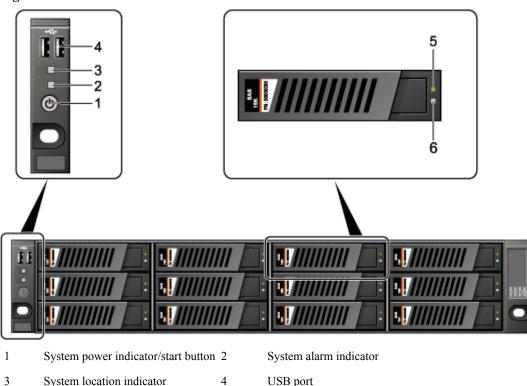
Figure 3-23 Exterior of a P12 node



### Front View

Figure 3-24 shows the front view of a P12 node.





System location indicator USB port

5 Disk activity indicator 6 Disk alarm/location indicator

The front panel of a P12 node provides 12 service disk slots, numbered from 0 to 11 from left to right and top to bottom, as shown in Figure 3-25.

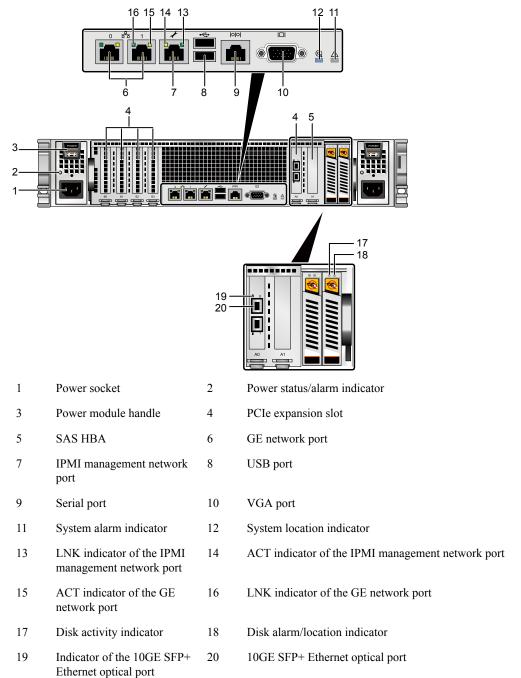
Figure 3-25 Disk slots on the front panel of a P12



### **Rear View**

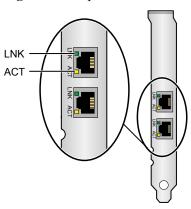
Figure 3-26 shows the rear panel of the P12 node.

**Figure 3-26** Rear panel of a P12 node with a SAS HBA and 2-port TOE 10GE network adapter



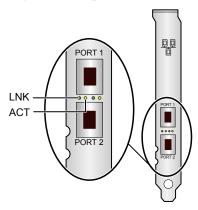
PCIe expansion slots of a P12 node can house 2-port GE network adapters and 2-port IB network adapters. Each 2-port GE network adapter provides two GE network ports, as shown in **Figure 3-27**.

Figure 3-27 2-port GE network adapter



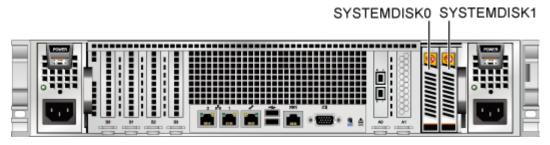
Each 2-port IB network adapter provides two 40 QSFP+ Ethernet optical ports, as shown in **Figure 3-28**.

Figure 3-28 2-port IB network adapter



The rear panel of a P12 node has two system disk slots, numbered as SYSTEMDISK 0 (left) and SYSTEMDISK 1 (right). **Figure 3-29** shows disk slots on the rear panel of a P12 node with a SAS HBA and 2-port TOE 10GE network adapter.

Figure 3-29 Disk slots on the rear panel of a P12 node



### **Ports**

Table 3-14 describes ports provided by the P12 node.

 Table 3-14 Description of ports

Port	Description
GE network port	If the front-end and rear-end devices are networked using GE switches, GE ports are used to connect those devices to the GE switches.
	The mainboard externally provides two GE network ports.
	Each node can be configured with two 2-port network adapters, each of wich externally provides two GE network ports.
IPMI	Provides the Web UI for the management and maintenance of the server.
USB port	Connects the mouse, keyboard, and other USB devices.
Serial port	Extends functions of the operating system and third-party software.
VGA port	Connects monitors.
10GE SFP+ Ethernet optical port	When the OceanStor 9000 is networked using 10GE switches, each node is configured with 2-port TOE 10GE network adapters, each of which provides two 10GE small form-factor pluggable plus (SFP+) Ethernet optical ports to connect to the 10GE switches.
40GE QSFP+ Ethernet optical port	When the OceanStor 9000 is networked using InfiniBand (IB) switches, each node is configured with 2-port IB network adapters, each of which provides two 40GE quad small-factor pluggable plus (QSFP+) to connect to the IB switches.

# **Physical Structure**

Figure 3-30 shows the physical structure of a P12 node.

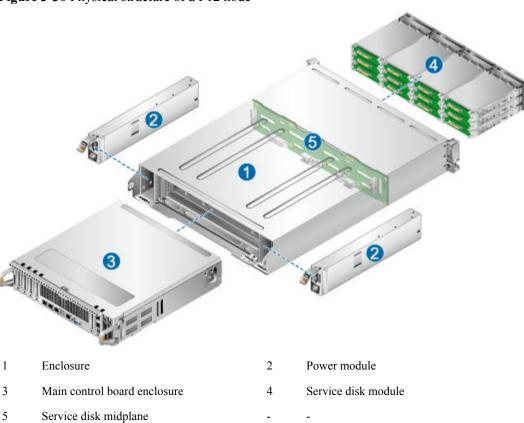


Figure 3-30 Physical structure of a P12 node

## **Technical Parameters**

**Table 3-15** lists the technical parameters of the P12.

**Table 3-15** Technical parameters of the P12

Category	Item	Parameter
Physical parameters	Dimensions (H x W x D)	86.1 mm x 446 mm x 582 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul> <li>Maximum weight<sup>a</sup>: 32 kg</li> <li>Unloaded weight<sup>b</sup>: 14.2 kg</li> <li>Shipment weight<sup>c</sup>: 45 kg</li> </ul>
Ambient parameters	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power parameter	AC input voltage	Rated input voltage: 100 V AC to 127 V AC or 200 V AC to 240 V AC
		Rated frequency: 50 Hz or 60 Hz

Category	Item	Parameter
Power consumption parameter	Maximum power consumption	470 W
Configuration example	Processor	Model: Intel Xeon 6-core 2.2 GHz Quantity: 2
	Memory	DDR3 RDIMM: 32 GB NVDIMM: 2 GB
	Disk	• Support for 2 x 300 GB 2.5-inch operating system disks (SAS) inserted in the rear panel
		• Support for 12 x 2 TB or 4 TB or 6 TB or 10 TB 3.5-inch service disks (SATA) inserted in slots 0 to 11
		Or support for 1 x 400 GB or 900 GB 3.5-inch service disk (SSD) inserted in slot 0, and 11 x 2 TB or 4 TB or 6 TB or 10 TB 3.5-inch service disks (SATA) inserted in slots 1 to 11
		NOTE  The mixing of 4KN and 512N/512e <sup>d</sup> system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.

a: The maximum weight refers to the weight of the P12 equipped with 12 service disks, 2 operating system disks, 2 power modules, and 3 memory modules.

### P25

### **Function**

A P25 node is a high-performance storage node applicable to OPS-intensive applications, equipped with 25 front service disk slots and 2 rear 2.5-inch system disk slots.

#### **Exterior**

#### Overview

b: The unloaded weight refers to the weight of the P12 without any components.

c: The shipment weight refers to the sum of the maximum weight of the P12 and the maximum weight of the shipping materials.

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

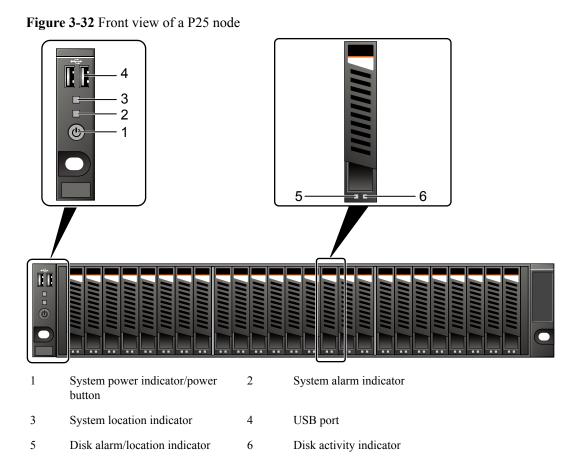
Figure 3-31 shows the exterior of a P25 node.

Figure 3-31 Exterior of a P25 node



### **Front View**

Figure 3-32 shows the front view of a P25 node.



The front panel of a P25 provides 25 service disk slots, numbered from 0 to 24 from left to right, as shown in **Figure 3-33**.

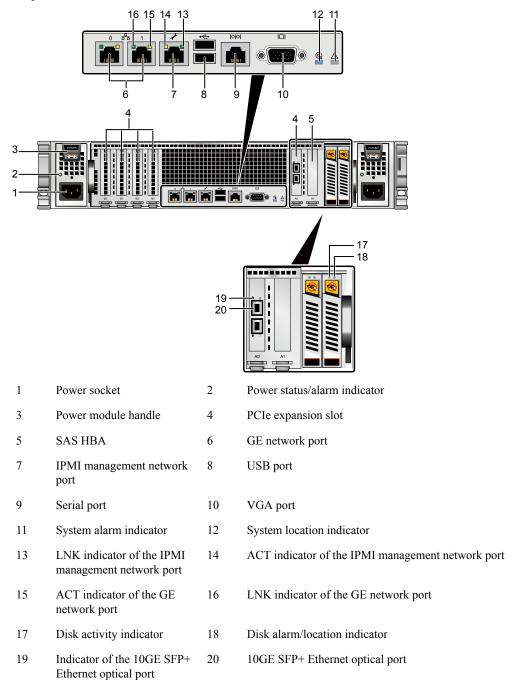
Figure 3-33 Disk slots on the front panel of a P25 node



#### **Rear View**

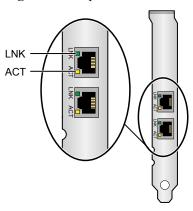
Figure 3-34 shows the rear panel of a P25 node.

**Figure 3-34** Rear panel of a P25 node with a SAS HBA and 2-port TOE 10GE network adapter



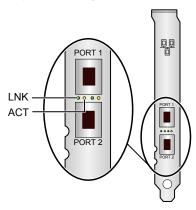
PCIe expansion slots of a P25 can house 2-port GE network adapters and 2-port IB network adapters. Each 2-port GE network adapter provides two GE network ports, as shown in **Figure 3-35**.

Figure 3-35 2-port GE network adapter



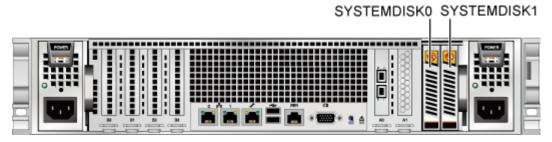
Each 2-port IB network adapter provides two 40 QSFP+ Ethernet optical ports, as shown in **Figure 3-36**.

Figure 3-36 2-port IB network adapter



The rear panel of a P25 node has two system disk slots, numbered as SYSTEMDISK 0 (left) and SYSTEMDISK 1 (right). **Figure 3-37** shows disk slots on the rear panel of a P25 node with a SAS HBA and 2-port TOE 10GE network adapter.

Figure 3-37 Disk slots on the rear panel of a P25 node



### **Ports**

Table 3-16 describes ports provided by a P25 node.

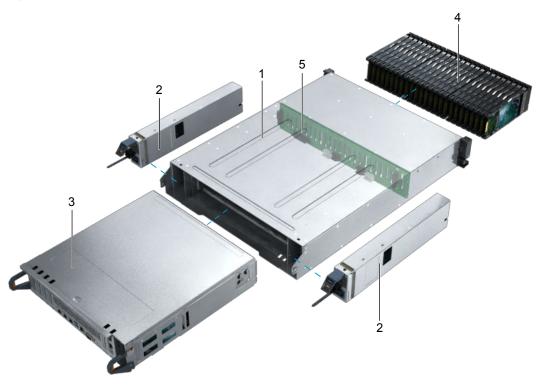
Table 3-16 Description of ports

Port	Description
GE network port	If the front-end and rear-end devices are networked using GE switches, GE ports are used to connect those devices to the GE switches.
	The mainboard externally provides two GE network ports.
	Each node can be configured with two 2-port network adapters, each of wich externally provides two GE network ports.
IPMI	Provides the Web UI for the management and maintenance of the server.
USB port	Connects the mouse, keyboard, and other USB devices.
Serial port	Extends functions of the operating system and third-party software.
VGA port	Connects monitors.
10GE SFP+ Ethernet optical port	When the OceanStor 9000 is networked using 10GE switches, each node is configured with 2-port TOE 10GE network adapters, each of which provides two 10GE small form-factor pluggable plus (SFP+) Ethernet optical ports to connect to the 10GE switches.
40GE QSFP+ Ethernet optical port	When the OceanStor 9000 is networked using InfiniBand (IB) switches, each node is configured with 2-port IB network adapters, each of which provides two 40GE quad small-factor pluggable plus (QSFP+) to connect to the IB switches.

# **Physical Structure**

Figure 3-38 shows the physical structure of a P25 node.

Figure 3-38 Physical structure of a P25 node



1 Enclosure 2 Power module

3 Main control board enclosure 4 Service disk module

5 Service disk midplane - -

## **Technical Parameters**

**Table 3-17** lists the technical parameters of the P25.

**Table 3-17** Technical parameters of the P25

Category	Item	Parameter
Physical parameters	Dimensions (H x W x D)	86.1 mm x 446 mm x 582 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul> <li>Maximum weight<sup>a</sup>: 35 kg</li> <li>Unloaded weight<sup>b</sup>: 10.65 kg</li> <li>Shipment weight<sup>c</sup>: 43 kg</li> </ul>
Ambient parameters	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH

Category	Item	Parameter
Power parameter	AC input voltage	<ul> <li>Rated input voltage: 100 V AC to 127 V AC or 200 V AC to 240 V AC</li> </ul>
		• Rated frequency: 50 Hz or 60 Hz
Power consumption parameter	Maximum power consumption	560 W
Configuration example	Processor	Model: Intel Xeon 6-core 2.2 GHz Quantity: 2
	Memory	DDR3 RDIMM: 48 GB
		NVDIMM: 2 GB
	Disk	• Support for 2 x 300 GB 2.5-inch operating system disks (SAS) inserted in the rear panel
		• Support for 1 x 400 GB or 900 GB 2.5-inch service disk (SSD) inserted in slot 0
		• Support for 24 x 900 GB or 1.2 TB or 1.8 TB 2.5-inch service disks (SAS) inserted in slots 1 to 24
		NOTE  4KN and 512N/512e disks <sup>d</sup> cannot be jointly used as operating system disks. In addition, 300 GB and 600 GB disks cannot be jointly used as operating system disks.

a: The maximum weight refers to the weight of the P25 equipped with 25 service disks, 2 operating system disks, 2 power modules, and 4 memory modules.

### P36

### **Function**

A P36 node is a high-performance storage node applicable to OPS-intensive applications, equipped with 36 service disk slots (including 24 front service disk slots and 12 rear service disk slots) and 2 rear 2.5-inch system disk slots.

b: The unloaded weight refers to the weight of the P25 without any components.

c: The shipment weight refers to the sum of the maximum weight of the P25 and the maximum weight of the shipping materials.

d: 4KN and 512N/512e refer to disk sector formats. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB but data can be written to the disk based on an integral of 512 bytes. 512N means that the sector size is 512 bytes.

### **Exterior**

### Overview

Figure 3-39 shows exterior of a P36 node.

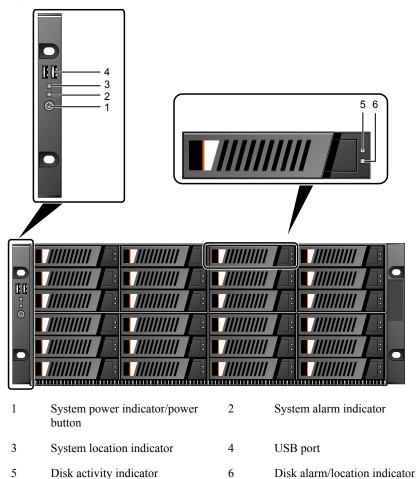
Figure 3-39 Exterior of a P36 node



### **Front View**

Figure 3-40 shows the front view of a P36 node.

Figure 3-40 Front view of a P36 node



The front panel of a P36 node provides 24 service disk slots, numbered from 0 to 23 from left to right and top to bottom, as shown in **Figure 3-41**. Slot 0 is located in the upper left corner of the front panel.

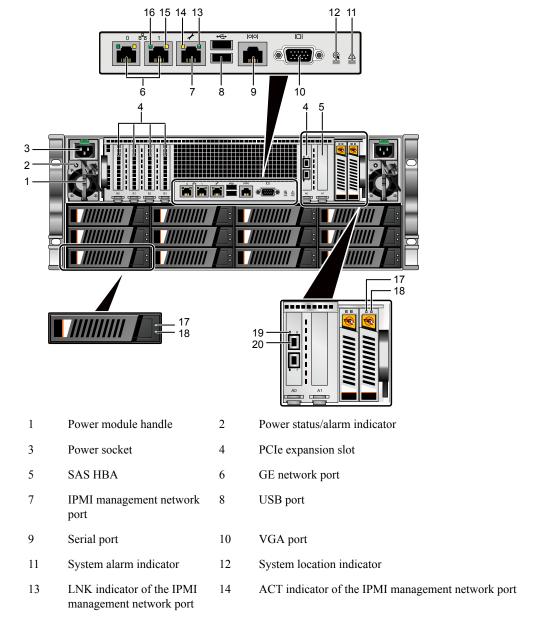
Figure 3-41 Disk slots on the front panel of a P36 node



### **Rear View**

Figure 3-42 shows the rear panel of a P36 node.

**Figure 3-42** Rear panel of a P36 node with a SAS HBA and 2-port TOE 10GE network adapter



PCIe expansion slots of a P36 node can house 2-port GE network adapters and 2-port IB network adapters. Each 2-port GE network adapter provides two GE network ports, as shown in **Figure 3-43**.

LNK indicator of the GE network port

Disk alarm/location indicator

10GE SFP+ Ethernet optical port

15

17

19

ACT indicator of the GE

Disk activity indicator

Ethernet optical port

Indicator of the 10GE SFP+

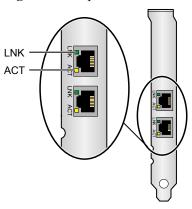
network port

16

18

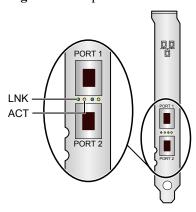
20

Figure 3-43 2-port GE network adapter



Each 2-port IB network adapter provides two 40 QSFP+ Ethernet optical ports, as shown in **Figure 3-44**.

Figure 3-44 2-port IB network adapter



The rear panel of a P36 node provides 2 system disk slots and 12 service disk slots.

- The two system disk slots are SYSTEMDISK 0 (left) and SYSTEMDISK 1 (right).
- Service disk slots in rear continue the numbering from those on the front panel from 24 to 35, from left to right and top to bottom. Slot 35 is located in the lower right corner of the rear panel.

**Figure 3-45** shows disk slots on the rear panel of a P36 node with a SAS HBA and 2-port TOE 10GE network adapter.

SYSTEMDISKO SYSTEMDISK1

Figure 3-45 Disk slots on the rear panel of a P36 node

### **Ports**

Table 3-18 describes ports provided by a P36 node.

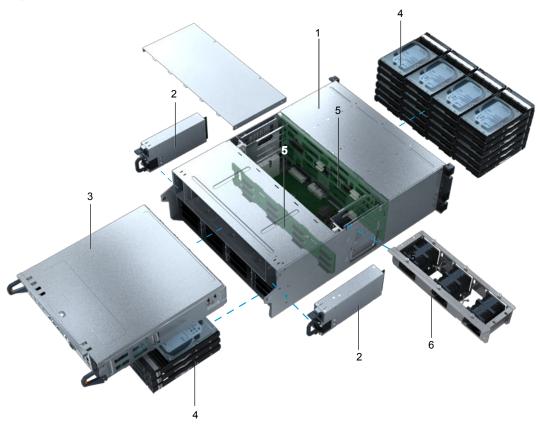
Table 3-18 Description of ports

Port	Description
GE network port	If the front-end and rear-end devices are networked using GE switches, GE ports are used to connect those devices to the GE switches.
	The mainboard externally provides two GE network ports.
	• Each node can be configured with two 2-port network adapters, each of wich externally provides two GE network ports.
IPMI	Provides the Web UI for the management and maintenance of the server.
USB port	Connects the mouse, keyboard, and other USB devices.
Serial port	Extends functions of the operating system and third-party software.
VGA port	Connects monitors.
10GE SFP+ Ethernet optical port	When the OceanStor 9000 is networked using 10GE switches, each node is configured with 2-port TOE 10GE network adapters, each of which provides two 10GE small form-factor pluggable plus (SFP+) Ethernet optical ports to connect to the 10GE switches.
40GE QSFP+ Ethernet optical port	When the OceanStor 9000 is networked using InfiniBand (IB) switches, each node is configured with 2-port IB network adapters, each of which provides two 40GE quad small-factor pluggable plus (QSFP+) to connect to the IB switches.

### **Physical Structure**

Figure 3-46 shows the physical structure of a P36 node.

Figure 3-46 Physical structure of a P36 node



- 1 Enclosure
- 3 Main control board enclosure
- 5 Service disk midplane

- 2 Power module
- 4 Service disk module
- 6 Fan module

### **Technical Parameters**

**Table 3-19** lists the technical parameters of the P36.

Table 3-19 Technical parameters of the P36

Category	Item	Parameter
Physical parameters	Dimensions (H x W x D)	175 mm x 446 mm x 582 mm, 4 U (1 U = 44.45 mm)
	Weight	Maximum weight <sup>a</sup> : 70 kg
		<ul> <li>Unloaded weight<sup>b</sup>: 20.1 kg</li> <li>Shipment weight<sup>c</sup>: 80 kg</li> </ul>

Category	Item	Parameter
Ambient parameters	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power parameter	AC input voltage	Rated input voltage range: 200     V AC to 240 V AC
		• Rated frequency: 50 Hz/60 Hz
Power consumption parameter	Maximum power consumption	880 W
Configuration example	Processor	Model: Intel Xeon 6-core 2.2 GHz Quantity: 2
	Memory	DDR3 RDIMM: 48 GB NVDIMM: 2 GB
	Disk	• Support for 2 x 300 GB 2.5- inch operating system disks (SAS) inserted in the rear panel
		• Support for 1 x 400 GB or 900 GB 3.5-inch service disk (SSD) inserted in slot 0
		• Support for 35 x 2 TB or 4 TB or 6 TB or 10 TB 3.5-inch service disks (SATA) inserted in slots 1 to 35
		NOTE  The mixing of 4KN and 512N/512e <sup>d</sup> system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.

a: The maximum weight refers to the weight of the P36 equipped with 36 service disks, 2 operating system disks, 2 power modules, and 4 memory modules.

C36

b: The unloaded weight refers to the weight of the P36 without any components.

c: The shipment weight refers to the sum of the maximum weight of the P36 and the maximum weight of the shipping materials.

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

### **Function**

A C36 node is a large-capacity archiving node applicable to capacity applications, equipped with 36 service disk slots (including 24 front service disk slots and 12 rear service disk slots) and 2 rear 2.5-inch system disk slots.

### **Exterior**

#### Overview

Figure 3-47 shows exterior of a C36 node.

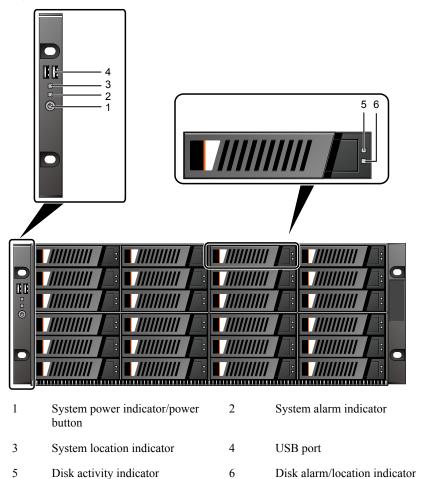
Figure 3-47 Exterior of a C36 node



#### Front View

Figure 3-48 shows the front view of a C36 node.

Figure 3-48 Front view of a C36 node



The front panel of a C36 node provides 24 service disk slots, numbered from 0 to 23 from left to right and top to bottom, as shown in **Figure 3-49**. Slot 0 is located in the upper left corner of the front panel.

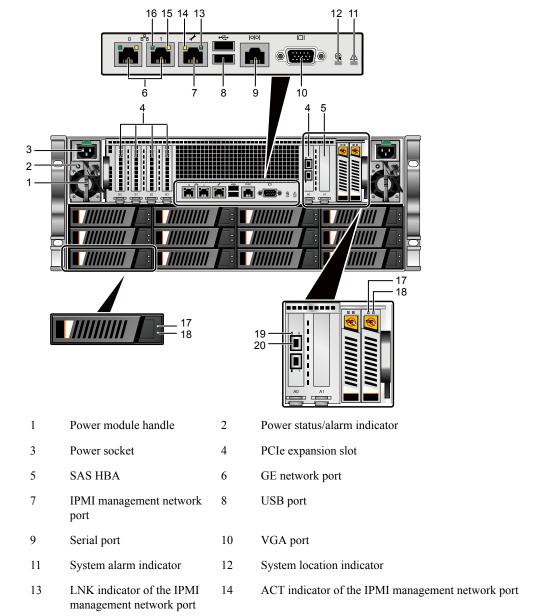
Figure 3-49 Disk slots on the front panel of a C36 node



### **Rear View**

Figure 3-50 shows the rear panel of a C36 node.

**Figure 3-50** Rear panel of a C36 node with a SAS HBA and 2-port TOE 10GE network adapter



PCIe expansion slots of a C36 node can house 2-port GE network adapters and 2-port IB network adapters. Each 2-port GE network adapter provides two GE network ports, as shown in **Figure 3-51**.

LNK indicator of the GE network port

Disk alarm/location indicator

10GE SFP+ Ethernet optical port

15

17

19

ACT indicator of the GE

Disk activity indicator

Ethernet optical port

Indicator of the 10GE SFP+

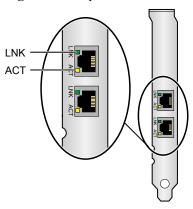
network port

16

18

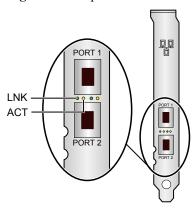
20

Figure 3-51 2-port GE network adapter



Each 2-port IB network adapter provides two 40 QSFP+ Ethernet optical ports, as shown in **Figure 3-52**.

Figure 3-52 2-port IB network adapter



The rear panel of a C36 provides 2 system disk slots and 12 service disk slots.

- The two system disk slots are SYSTEMDISK 0 (left) and SYSTEMDISK 1 (right).
- Service disk slots in rear continue the numbering from those on the front panel from 24 to 35, from left to right and top to bottom. Slot 35 is located in the lower right corner of the rear panel.

**Figure 3-53** shows disk slots on the rear panel of a C36 node with a SAS HBA and 2-port TOE 10GE network adapter.

SYSTEMDISKO SYSTEMDISK1

Figure 3-53 Disk slots on the rear panel of a C36 node

### **Ports**

**Table 3-20** describes ports provided by the C36 node.

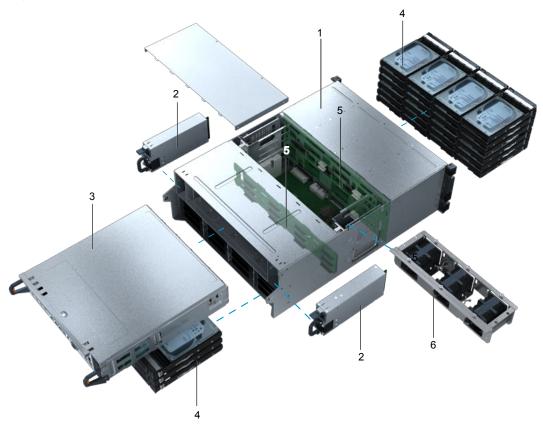
Table 3-20 Description of ports

Port	Description
GE network port	If the front-end and rear-end devices are networked using GE switches, GE ports are used to connect those devices to the GE switches.
	The mainboard externally provides two GE network ports.
	• Each node can be configured with two 2-port network adapters, each of wich externally provides two GE network ports.
IPMI	Provides the Web UI for the management and maintenance of the server.
USB port	Connects the mouse, keyboard, and other USB devices.
Serial port	Extends functions of the operating system and third-party software.
VGA port	Connects monitors.
10GE SFP+ Ethernet optical port	When the OceanStor 9000 is networked using 10GE switches, each node is configured with 2-port TOE 10GE network adapters, each of which provides two 10GE small form-factor pluggable plus (SFP+) Ethernet optical ports to connect to the 10GE switches.
40GE QSFP+ Ethernet optical port	When the OceanStor 9000 is networked using InfiniBand (IB) switches, each node is configured with 2-port IB network adapters, each of which provides two 40GE quad small-factor pluggable plus (QSFP+) to connect to the IB switches.

### **Physical Structure**

Figure 3-54 shows the physical structure of a C36 node.

Figure 3-54 Physical structure of a C36 node



- 1 Enclosure
- 3 Main control board enclosure
- 5 Service disk midplane

- 2 Power module
- 4 Service disk module
- 6 Fan module

### **Technical Parameters**

**Table 3-21** lists the technical parameters of the C36.

 Table 3-21 Technical parameters of the C36

Category	Item	Parameter
Physical parameters	Dimensions (H x W x D)	175 mm x 446 mm x 582 mm, 4 U (1 U = 44.45 mm)
	Weight	<ul> <li>Maximum weight<sup>a</sup>: 70 kg</li> <li>Unloaded weight<sup>b</sup>: 20.1 kg</li> <li>Shipment weight<sup>c</sup>: 80 kg</li> </ul>

Category	Item	Parameter
Ambient parameters	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power parameter	AC input voltage	Rated input voltage range: 200     V AC to 240 V AC
		• Rated frequency: 50 Hz/60 Hz
Power consumption parameter	Maximum power consumption	880 W
		Model: Intel Xeon 6-core 2.2 GHz Quantity: 2
	Memory	DDR3 RDIMM: 32 GB NVDIMM: 2 GB
	Disk	• Support for 2 x 300 GB 2.5-inch operating system disks (SAS) inserted in the rear panel
		Support for 36 x 2 TB or 4 TB or 6 TB or 10TB 3.5-inch service disks (SATA) inserted in slots 0 to 35
		NOTE  The mixing of 4KN and 512N/512e <sup>d</sup> system disks is not supported. The mixing of 300 GB and 600 GB system disks is not supported either.

a: The maximum weight refers to the weight of the C36 equipped with 36 service disks, 2 operating system disks, 2 power modules, and 3 memory modules.

### C72

### **Function**

C72 node is a large-capacity archiving node applicable to capacity applications. It has 72 service disk slots and four operating system disk slots.

### **Exterior**

#### **Device**

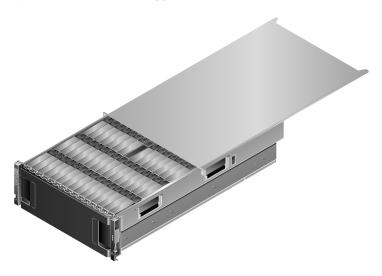
b: The unloaded weight refers to the weight of the C36 without any components.

c: The shipment weight refers to the sum of the maximum weight of the C36 and the maximum weight of the shipping materials.

d: Disk sector format. 4KN means that the sector size is 4 KB. 512e means that the sector size is 4 KB and allows an integral multiple of 512 B data to be written. 512N means that the sector size is 512 B.

Figure 3-55 shows the C72 overall appearance.

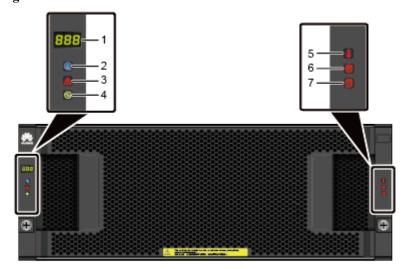
Figure 3-55 C72 overall appearance



### Front view

Figure 3-56 shows the front view of C72.

Figure 3-56 C72 front view

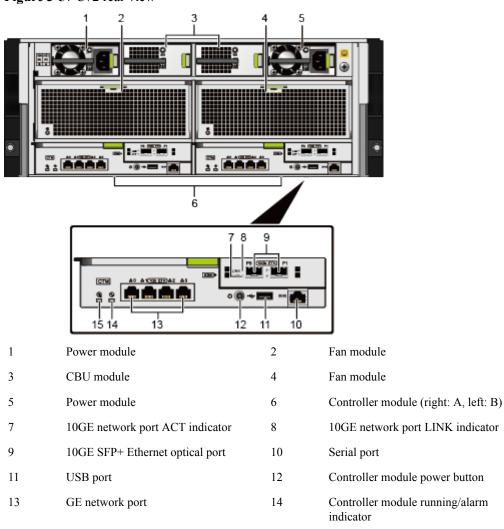


1	Enclosure ID	2	Location indicator
3	Alarm indicator	4	Running indicator
5	Overtemperature alarm indicator	6	Disk alarm indicator
7	Rear FRU alarm indicator	-	-

#### Rear view

Figure 3-57 shows the rear view of C72.

Figure 3-57 C72 rear view



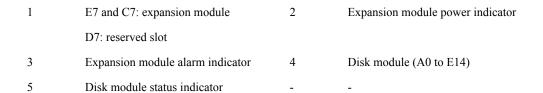
### Top view

15

Figure 3-58 shows the top view of C72.

Controller module location indicator

Figure 3-58 C72 top view



### Port

Table 3-22 describes ports on C72.

Table 3-22 Port description

Port	Description
GE network port	Connects to GE switches. C72 provides four onboard GE network ports.
USB port	Connects to a mouse, a keyboard, or other USB devices.
Serial port	Extends the functions of the operating system and third-party software.
10GE SFP+ Ethernet optical port	When OceanStor 9000 is networked using 10GE switches, each node is configured with 2-port TOE 10GE network adapters, each of which provides two 10GE small form-factor pluggable plus (SFP+) Ethernet optical ports to connect to the 10GE switches.

# **Physical Structure**

Figure 3-59 shows the basic structure of C72.

**Figure 3-59** basic structure of C72

1	Controller module	2	Fan module
3	Power module	4	CBU module
5	System enclosure	6	Disk module
7	Expansion module		

### **Technical Parameters**

**Table 3-23** lists the technical parameters of the C72.

**Table 3-23** Technical parameters of the C72

Category	Item	Parameter
Physical parameters	Dimensions (H x W x D)	175 mm x 446 mm x 790 mm, 4 U (1 U = 44.45 mm)
	Weight	<ul> <li>Maximum weight<sup>a</sup>: 96 kg</li> <li>Unloaded weight<sup>b</sup>: 26.1 kg</li> <li>Shipment weight (without service disks): 80 kg</li> <li>Shipment weight (only with 15 service disks per carton): 15 kg/carton</li> </ul>
Ambient parameters	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power parameter	AC input voltage	<ul> <li>Rated input voltage range: 200 V AC to 240 V AC</li> <li>Rated frequency: 50 Hz/60 Hz</li> </ul>
Power consumption parameter	Maximum power consumption	1100 W
Configuration example	Processor	Model: Intel Atom 8-core 2.0 GHz Quantity: 2
	Memory	DDR3 RDIMM: 64 GB (using the CBU module to endow 4 GB of memory with the power failure protection capability)
	Disk	<ul> <li>Support for 4 x 64 GB operating system disks (mini-SATA disks that reside in the controller module)</li> <li>Support for 72 x 4 TB or 6 TB or 10 TB 3.5-inch service disks (SATA) that are inserted in slots A0 to E14</li> </ul>

a: The maximum weight refers to the weight of the C72 equipped with 72 service disks, 2 controller modules, 2 power modules, 2 fan modules, 2 CBU modules, 8 memory modules, and 2 expansion modules.

# 3.1.3 Switches

The OceanStor 9000 supports multiple models of switches such as CE6810-48S4Q-EI (CE6810 for short), S5700-52C-SI (S5752 for short)/S5352-SI (S5352 for short) and SX6018.

b: The unloaded weight refers to the weight of the C72 without any components.

### CE6810-24S2Q-LI

CE6810-24S2Q-LI switches provide 10GE access.

### **Function**

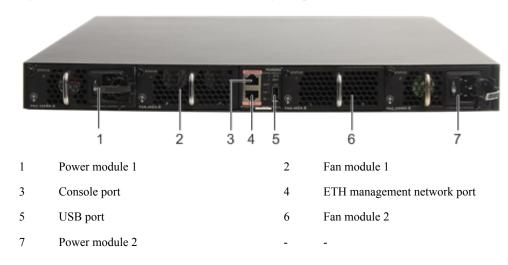
A CE6810-24S2Q-LI switch provides 10GE ports for node interconnection and communication, and supports 40GE upstream ports connection to the front-end network. CE6810-24S2Q-LI can provide twenty-four 10GE SFP+ Ethernet optical ports and two 40GE QSFP+ Ethernet optical ports, each of which can be divided into four 10GE ports.

### **Exterior**

#### **Front View**

Figure 3-60 shows the front view of the switch.

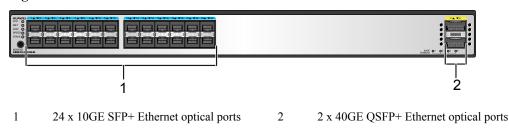
Figure 3-60 Front view of the switch (housing AC power modules)



#### **Rear View**

Figure 3-61 shows the rear view of the switch.

Figure 3-61 Rear view of the switch



### **Ports**

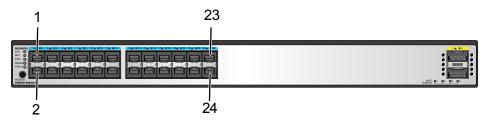
**Table 3-24** lists the ports of the switch.

Table 3-24 Ports of the switch

Port Name	Description
10GE SFP+ Ethernet optical port	Used for sending and receiving GE or 10GE services
40GE QSFP+ Ethernet optical port	Used for switch stacking or as an upstream port
Console port	Connects to the console for onsite device configuration
ETH management network port	Connects to the configuration terminal or network management workstation for building an onsite or remote configuration environment

Figure 3-62 shows the numbering of ethernet ports of a CE6810-24S2Q-LI switch

Figure 3-62 Ethernet ports of a CE6810-24S2Q-LI switch



# **Technical Specifications**

**Table 3-25** lists the technical specifications of the switch.

Table 3-25 Technical specifications of the switch

Item		Value	
Physical Dimensions (H x W x D) specifications		43.6 mm x 442.0 mm x 600.0 mm (1.72 in. x 17.4 in. x 23.62 in.)	
	Fan quantity	2	
	Max. power consumption	171W	
	Weight	10.1kg	
System	Processor	Main frequency: 1.2 GHz	
specifications	DDR memory	2 GB	

Item		Value
	NOR flash memory	16 MB
Power	Rated voltage	100 V AC to 240 V AC; 50 Hz or 60 Hz
specifications	Input voltage	90 V AC to 290 V AC; 45 Hz or 65 Hz
Environmenta l specifications	Temperature	<ul> <li>Operating temperature: 0°C to 40°C (32°F to 104°F)</li> <li>Storage temperature: -40°C to +70°C (-104°F to +158°F)</li> </ul>
	Relative humidity	5% RH to 95% RH
	Altitude	< 1800 m

#### CE6810-48S4Q-EI

CE6810-48S4Q-EI switches provide 10GE access.

#### **Function**

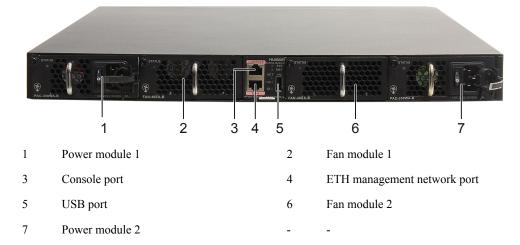
A CE6810-48S4Q-EI switch provides 10GE ports for node interconnection and communication, and supports 40GE upstream ports connection to the front-end network. CE6810-48S4Q-EI can provide forty-eight 10GE SFP+ Ethernet optical ports and four 40GE QSFP+ Ethernet optical ports, each of which can be divided into four 10GE ports.

#### **Exterior**

#### Front View

Figure 3-63 shows the front view of the switch.

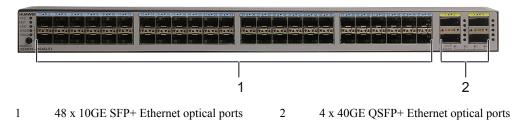
Figure 3-63 Front view of the switch (housing AC power modules)



#### **Rear View**

Figure 3-64 shows the rear view of the switch.

Figure 3-64 Rear view of the switch



#### **Ports**

**Table 3-26** lists the ports of the switch.

Table 3-26 Ports of the switch

Port Name	Description
10GE SFP+ Ethernet optical port	Used for sending and receiving GE or 10GE services
40GE QSFP+ Ethernet optical port	Used for switch stacking or as an upstream port
Console port	Connects to the console for onsite device configuration
ETH management network port	Connects to the configuration terminal or network management workstation for building an onsite or remote configuration environment

Figure 3-65 shows the numbering of ethernet ports of a CE6810-48S4Q-EI switch

Figure 3-65 Ethernet ports of a CE6810-48S4Q-EI switch



## **Technical Specifications**

**Table 3-27** lists the technical specifications of the switch.

Table 3-27 Technical specifications of the switch

Item		Value
Physical specifications	Dimensions (H x W x D)	43.6 mm x 442.0 mm x 600.0 mm (1.72 in. x 17.4 in. x 23.62 in.)
	Fan quantity	2
	Max. power consumption	238 W
	Weight	10.4 kg
System	Processor	Main frequency: 1.5 GHz
specifications	Switching capacity	1280 Gbit/s
	Packet forwarding capability	952.32 Mpps
	DDR memory	2 GB
	NOR flash memory	8 MB
Power	Rated voltage	100 V AC to 240 V AC; 50 Hz or 60 Hz
specifications	Input voltage	90 V AC to 290 V AC; 45 Hz or 65 Hz
Environmenta 1	Temperature	• Operating temperature: 0°C to 40°C (32°F to 104°F)
specifications		• Storage temperature: -40°C to +70°C (-104°F to +158°F)
	Relative humidity	5% RH to 95% RH
	Altitude	0 m to 1800 m (5905.44 ft)

## S5700-52C-SI/S5352C-SI

S5700-52C-SI/S5352C-SI switches provide GE access.

#### **Function**

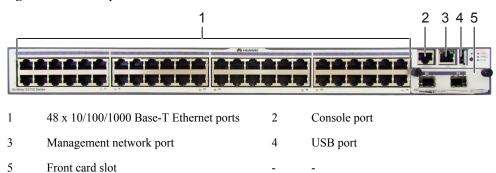
S5700-52C-SI/S5352C-SI (S5752/S5352 for short) is a device that connects servers and enables data communication and active/standby configuration for connected components.

#### **Exterior**

#### Front View

Figure 3-66 shows the front panel of a S5752/S5352 switch.

Figure 3-66 Front panel of a S5752/S5352 switch



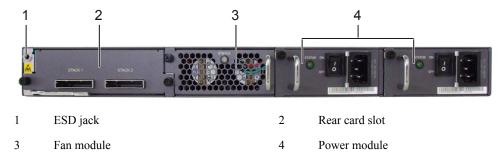
#### NOTE

- An optical card with two 10GE XFP optical ports is installed for an S5352C-SI switch.
- An optical card with two 10GE SFP+ optical ports is installed for an S5700-52C-SI switch.

#### **Rear View**

Figure 3-67 shows the rear panel of a S5752/S5352 switch.

Figure 3-67 Rear panel of a S5752/S5352 switch equipped with AC power modules



#### **Ports**

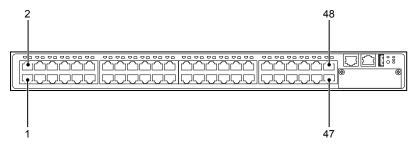
Table 3-28 lists the ports of a S5752/S5352 switch.

**Table 3-28** Ports of a S5752/S5352 switch

Port Name	Description
10/100/1000 Base-T Ethernet ports	Used for sending and receiving Ethernet services
Console port	Connects to the console for onsite device configuration
Management network port	Used for accessing and managing Ethernet switches
10 Gbit/s optical port	Used for receiving and sending 10GE services

**Figure 3-68** shows the numbering of the forty-eight 10/100/1000 Base-T Ethernet ports of a S5752/S5352 switch.

Figure 3-68 Ethernet ports of a S5752/S5352 switch



## **Technical Specifications**

**Table 3-29** lists the technical specifications of a S5752/S5352 switch.

Table 3-29 Technical specifications of a S5752/S5352 switch

Item		Value	
Physical specifications	Dimensions (H x W x D)	43.6 mm x 442.0 mm x 420.0 mm (1.72 in. x 17.4 in. x 16.54 in.)	
	Fan quantity	2	
	Max. power	78 W	
	Max. Weight	8.5 kg	
System	Processor	Frequency: 800 MHz	
specifications	Switching capacity	176 Gbit/s	
	Packet forwarding capability	130.94 Mpps	
	DDR memory	256 MB	
	Flash memory	32 MB	
Power specifications	Input voltage	AC input voltage:  Rated input voltage: 100 V AC to 240 V AC; 50 Hz or 60 Hz  Maximum input voltage: 90 V AC to 264 V AC; 47 Hz to 63 Hz	
Environmenta l specifications	Temperature	<ul> <li>Working temperature: 0°C to 50°C (32°F to 122°F)</li> <li>Storage temperature: -40°C to +70°C (-104°F to +158°F)</li> </ul>	

Item		Value
	Relative humidity	10% RH to 90% RH
	Altitude	0 m to 2000 m (0 to 6561.6 ft)

#### SX6018

SX6018 switches provide InfiniBand access.

#### **Function**

An SX6018 switch provides eighteen 40 Gbit/s QSFP+ ports for node interconnection, communication, and access to the front-end network.

#### **Exterior**

Figure 3-69 shows the exterior of an SX6018 switch.

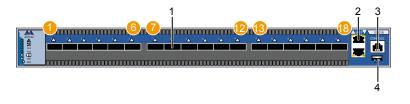
Figure 3-69 Exterior of an SX6018 switch



#### **Ports**

Figure 3-70 shows ports on the front panel of an SX6018 switch.

Figure 3-70 Ports on SX6018's front panel



- 1 QSFP+ ports 2 Management network port
- 3 Console 4 USB

Figure 3-71 shows the ports on the rear panel of an SX6018 switch.

Figure 3-71 Ports on SX6018's rear panel



**Table 3-30** lists the ports of an SX6018 switch.

Table 3-30 Ports of an SX6018 switch

Name	Used For
QSFP+ port	User access
Console	Initializing and debugging devices.
Management network port	Remote management
USB	Connecting to external disks. The MLNX-OS software is used to implement software update or file management, and USB 2.0 is supported.
I2C port	Upgrading firewall software and other advanced debugging operations
Power socket	Connecting to power cables

## **Technical Specifications**

**Table 3-31** lists the technical specifications of an SX6018 switch.

Table 3-31 Technical Specifications of an SX6018 switch

Item		Value
Physical	Dimensions (H x W x D)	43.6 mm×427.7 mm×627.4 mm
specifications	Max. power	152 W
	Weight	9 kg
Power specifications	Input voltage	100 V AC to 240 V AC; 50 Hz or 60 Hz
Environmenta	Temperature	0°C to 45°C
specifications	Relative humidity	10% RH to 90% RH (non-condensing)
	Altitude	-60 m to 3200 m

#### **3.1.4 FR42612L Cabinets**

It is recommended that the OceanStor 9000 be installed in FR42612L cabinets.

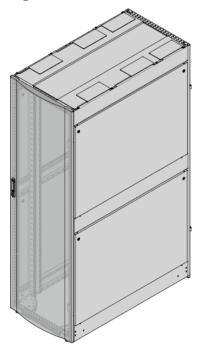
#### **Function**

An FR42612L cabinet provides 42 U (1866.9 mm or 73.5 in.) of internal space for accommodating devices and protects the devices from electromagnetic interference.

#### **Exterior**

The exterior of an FR42612L cabinet is in sand texture black. **Figure 3-72** shows an FR42612L cabinet.

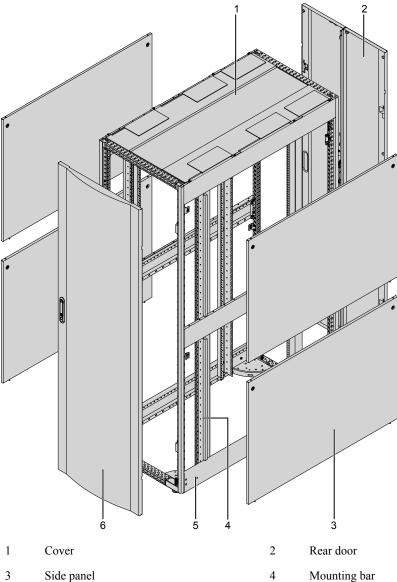
Figure 3-72 Exterior of an FR42612L cabinet



#### Structure

The FR42612L cabinet consists of the racks, front doors, back doors, side panels, cover, and mounting bars, as shown in **Figure 3-73**.

Figure 3-73 Cabinet structure



Rack

5

- Front door

## **Technical Specifications**

**Table 3-32** lists the technical specifications of an FR42612L cabinet.

Table 3-32 Technical specifications of an FR42612L cabinet

Parameter	Value
Dimensions (H x W x D)	2000 mm x 600 mm x 1200 mm
Capacity	42 U (1866.9 mm or 73.5 in.) of internal space

Parameter	Value
Weight	• 120 kg (only with the front and rear doors)
	• 170 kg (with front doors, back doors, guide rails, and cables)
Cabling mode	Overhead and underfloor cabling
Installation mode	Fastening installation and non-fastening installation
	The two modes are applicable both to the concrete ground and ESD floor.
Material	High-intensity G-A quality carbon cold-rolled steel plates and galvanized sheets that comply with Restriction of the Use of Certain Hazardous Substances (RoHS) and Underwriter Laboratories (UL)
Heat dissipation	Perforated doors, front-to-rear cooling, and underfloor air intake
Operating	• Long term: 0°C to 50°C
temperature	• Short term: - 5°C to +50°C
Operating humidity	Long term: 5% RH to 85% RH
	• Short term: 5% RH to 95% RH

#### **NOTE**

- Operating temperature and humidity are measured 1.5 m above the floor and 0.4 m in front of the device in the equipment room.
- Short-term operation means that the consecutive operating time does not exceed 48 hours, and the annual accumulative operating time does not exceed 15 days.

## 3.1.5 Optional Hardware

SMS modems and keyboard, video, and mouse (KVM) devices can be configured for the OceanStor 9000 for easy device maintenance.

#### **SMS Modems**

SMS modems are used for alarm notification by SMS. When maintenance personnel are not on site, you can set the SMS notification and send the alarm through SMS to the maintenance personnel.

Figure 3-74 shows the exterior of an SMS modem.

Figure 3-74 Exterior of an SMS modem



#### **KVM**

An 8-port KVM can be configured for access and control of the OceanStor 9000.

Figure 3-75 shows the exterior of a KVM.

Figure 3-75 Exterior of a KVM



Table 3-33 lists the technical specifications of a KVM.

Table 3-33 Technical specifications of a KVM

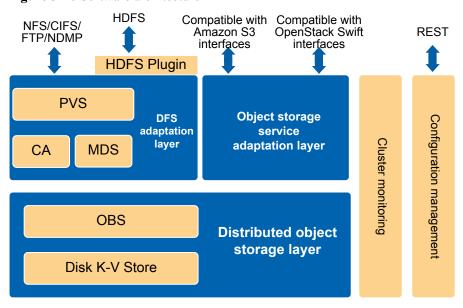
Parameter	Value
Height	1 U
Power input	90 V AC to 264 V AC
Power consumption	46 W
Weight	15 kg
Туре	8-port KVM
Switchover mode	Shortcut keys
Operating temperature	0°C to 50°C (32°F to 122°F)
Operating humidity	0% RH to 90% RH

## 3.2 Software Architecture

This section describes the software architecture of the OceanStor 9000.

Figure 3-76 shows the software architecture of the OceanStor 9000.

Figure 3-76 Software architecture



**Table 3-34** describes the major software modules shown in **Figure 3-76**.

Table 3-34 Software module description

Type		Description
DFS HDFS adaptatio Plugin		The HDFS Plugin converts the file access requests of the HDFS protocol to the requests of the NFS protocol.
n layer	Protocol and Value- added Service (PVS)	The PVS responds to connection requests initiated by external NFS, CIFS, and FTP applications. In addition, the PVS provides value-added features such as snapshot, storage tiering, and remote replication.
	Client Agent (CA)	The CA implements semantic analysis of application protocols such as NFS, CIFS, and FTP and sends the analysis results to underlying modules.
	MetaData Service (MDS)	The MDS manages file system metadata and acts as the cache for reading and writing file system metadata. The MDS cluster organizes file system metadata into a tree data structure. Each node only stores one sub-tree.

Туре		Description
Object-based storage service adaptation layer		The object-based storage (OBS) service adaptation layer provides OBS services, and is compatible with Amazon S3 and OpenStack Swift interfaces.
Distribut ed object- based storage layer	OBS (Object- Based Store)	The OBS provides reliable object-based storage services for the metadata and data of file systems and the OBS. It supports multi-level data protection, load balancing, and fast data reconstruction.
	Disk K-V Store	The Disk K-V Store supports disk-level Key-Value OBS.

## 3.3 Network Architecture

The OceanStor 9000 network architecture contains the management network, front-end service network, and back-end storage network.

#### Overview

The OceanStor 9000 network architecture consists of the following:

- A front-end service network that connects the OceanStor 9000 to the customer's network
- A back-end storage network that interconnects OceanStor 9000 nodes
- A management network that connects the OceanStor 9000 to the customer's maintenance network

The OceanStor 9000 supports multiple networking modes, including 10GE, InfiniBand, and GE networking to meet different network requirements.

## **Typical Network**

This section provides an example network for operations per second (OPS)-intensive applications. The OceanStor 9000 in this example adopts a 10GE network where all nodes provide IPMI ports and are managed and maintained through S5752 switches. **Figure 3-77** shows the 10GE network.

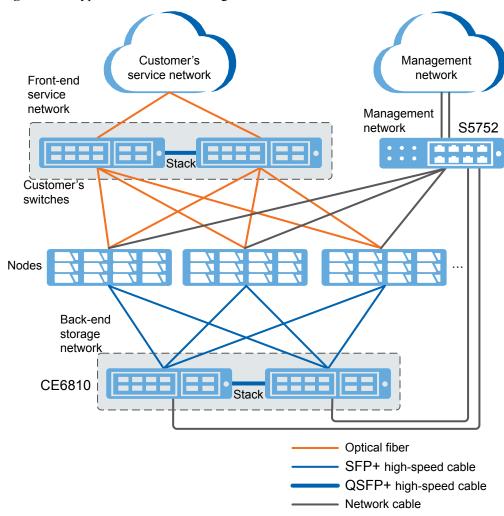


Figure 3-77 Typical 10GE network diagram

The network is described as follows:

- Front-end service network: Two 10GE ports on each node are connected to two user switches through optical fibers.
- Back-end storage network: Two 10GE ports on each node are connected to two CE6810 switches through high-speed cables.
- Management network: The IPMI port of a P12, P25, P36, or C36 node and the Mgmt port of a P12E, P25E, P36E, or C36E node connect to an S5752 switch with one GE network cable. A GE port on each of the first three nodes is connected to the S5752 switch. A GE port on each CE6810 switch is connected to the S5752 switch through a GE network cable.

#### NOTE

The two CE6810 switches are stacked through 40GE QSFP+ high-speed cables. The number of cables varies with the number of OceanStor 9000 nodes.

# 4 Product Features

## **About This Chapter**

This section introduces the functions and features of OceanStor 9000.

#### 4.1 Service Function

The OceanStor 9000 provides service functions such as CIFS/NFS/FTP share, object-based storage, and NDMP backup.

#### 4.2 Management Function

The OceanStor 9000 provides management functions such as data consistency scan, high system security, and unified namespace.

#### 4.3 File System Basic Features

The OceanStor 9000 provides basic features such as InfoTier, InfoAllocator, and InfoEqualizer.

#### 4.4 File System Enhanced Features

OceanStor 9000 provides enhanced features such as InfoRevive, InfoCopy, and InfoTurbo.

### 4.1 Service Function

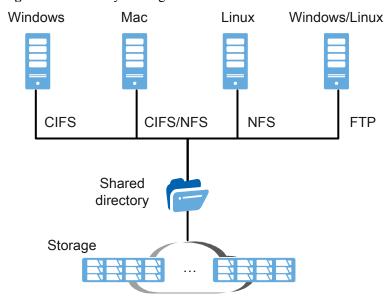
The OceanStor 9000 provides service functions such as CIFS/NFS/FTP share, object-based storage, and NDMP backup.

## 4.1.1 CIFS/NFS/FTP Sharing

The OceanStor 9000 provides a reliable file sharing service.

File sharing is a basic feature of a NAS storage system. With support for CIFS, NFS, and FTP as well as compatibility with multiple types of operating systems, the OceanStor 9000 can implement file resource sharing in a heterogeneous network environment and allows a directory to be shared based on CIFS, NFS and FTP at the same time, as shown in **Figure 4-1**.

Figure 4-1 Directory sharing



The CIFS/NFS/FTP sharing feature shares data by directory. A shared directory can be accessed by multiple clients. All nodes of the OceanStor 9000 are clustered to provide sharing services and monitor service status of each other.

The OceanStor 9000 balances data access requests among nodes based on service loads and node status, ensuring high performance in high concurrency scenarios. When a node is faulty, the OceanStor 9000 fails over services from the faulty node to a functioning one to ensure business continuity.

## 4.1.2 CIFS-NFS Cross-Protocol Access Sharing

NFS and CIFS sharing can be set simultaneously for a directory of the OceanStor 9000. The user mapping function ensures refined permission control for CIFS-NFS cross-protocol file access sharing on different clients.

#### Cross-Protocol (CIFS-NFS) Share Access

The OceanStor 9000 allows users to share a directory using NFS and CIFS at the same time. Different clients, such as Windows-based clients (CIFS), Linux clients (NFS), and Mac OS X

clients (CIFS or NFS) can access one shared directory or file simultaneously. Since Windows, Linux, and Unix adopt different mechanisms to authenticate users and control access, the OceanStor 9000 uses a mechanism to centrally map users and control access, protecting the security of cross-protocol (CIFS-NFS) access.

- If a CIFS user attempts to access a file or directory on the OceanStor 9000, the OceanStor 9000 authenticates local or AD domain users in the first place. If the UNIX permission (UNIX Mode bits, POSIX ACL, or NFSv4 ACL) has been configured the file or directory to be accessed, the CIFS user is mapped as an NFS user based on preset user mapping rules during authentication. Then the OceanStor 9000 implements UNIX permission authentication for the user.
- If an NFS user attempts to access a file or directory that has NT ACL on the OceanStor 9000, the OceanStor 9000 maps the NFS user as a CIFS user based on the preset mapping rules. Then the OceanStor 9000 implements NT ACL permission authentication for the user.

#### **CIFS-NFS Cross-Protocol User Mapping**

Windows systems (CIFS) and Linux systems (NFS) use different mechanisms to identify and authenticate users:

- Windows systems use security identifiers (SIDs) to identify users. SIDs apply to all
  users, user groups, services, and computers in the systems. Regarding authentication,
  CIFS support NT ACL.
- Linux systems use user identities (UIDs) and one or more group identities (GIDs) to identify users. One user belongs to one user group at least. Regarding authentication, NFS supports diversified security control mechanisms such as UNIX Mode bits and NFSv4 ACL.

During CIFS-NFS cross-protocol share access, users on different protocols must be mapped based on user mapping rules for user authentication and precise permission control.

The timing of user mapping is as follows:

- When a CIFS client accesses files or directories with the POSIX ACL/NFSv4 ACL or UNIX Mode bits permission, a user mapping occurs. The user will have both the permissions before and after the user is mapped.
- When a NFS client accesses files or directories with the NT ACL permission, a user mapping occurs. The user will have both the permissions before and after the user is mapped.
- Cross-protocol permission editing changes permission types. For example, users are mapped when an NFS client accesses a file or directory that has the NT ACL. If the NFS client runs the **chmod** command or sets the NFS ACL to change the permission of the file or directory, users are not mapped when the NFS client accesses the file or directory after the change. Then users do not have the permission of mapped users.

#### NOTE

You are advised not to cross protocol edit permission, avoid to change permissions type.

• When the parent directory has the inheritable NT ACL permission, the files or directories created no matter on a NFS client or a CIFS client will have the NT ACL permission by default. When the NFS client accesses files, a user mapping will always occur. That is, the user will have both the permissions before and after the user is mapped; When the parent directory does not have the inheritable NT ACL permission, the files or directories created no matter on a NFS client or a CIFS client will have the UNIX Mode

- bits permission. When the NFS client accesses files, no user mapping occurs. That is, the user does not have the permission after the user is mapped.
- On CIFS clients, users are mapped upon authentication. If mappings are changed, the change takes effect after next authentication.
- User mappings on NFS clients are cached and expire after four hours by default. New user mappings and user information changes take effect after the cached data expires.

User mapping rules specify the mappings among different user accounts. The user mapping rules can be saved in a local database or centrally managed in the AD domain. A user mapping rule includes the mapping type, original user, mapped user, and mapping priority. If a user matches multiple mapping rules, it is mapped based on the rule with a higher priority. If the rules have the same priority, the user is mapped based on the rule that is configured the earliest.

A user mapping process is as follows: (Local mapping is used as an example.)

- NFS-CIFS user mapping: A NFS user is authenticated by UID on the service end. When
  a user mapping occurs, the user name to which the UID corresponds will be queried in
  the sequence of the local host, LDAP domain, and NIS domain. Based on the queried
  user name and the local mapping, the user name, SID, and the owning group of the
  mapped user will be queried.
- CIFS-NFS user mapping: A CIFS user is authenticated by SID on the service end. When
  a user mapping occurs, the mapped user will be queried based on the user name to which
  the SID corresponds and the local mapping. Then the UID to which the mapped user
  name corresponds and its owing group will be queried in the sequence of the local host,
  LDAP domain, and NIS domain.

#### NOTE

You are advised not to configure the same UIDs or user names in the local host, LDAP domain, or NIS domain. If the same UIDs or user names exist, the user mapping result will not be the expected result.

## 4.1.3 Object Storage Service (Compatible with Amazon S3 APIs)

The OceanStor 9000 provides object storage service (compatible with Amazon S3 APIs). Object storage service (compatible with Amazon S3 APIs) offers secure, enduring, and highly scalable storage devices to R&D personnel and IT teams. This allows end users to use object storage at ease. At the same time, they can also store and search for data of any amount at any location on a web.

The Object Storage Service (compatible with Amazon S3 APIs) provides RESTful APIs, enabling end users to use Uniform Resource Identifier (URI) to store and search for their own data. The Object Storage Service (compatible with Amazon S3 APIs) discards the directory tree structure and simplifies read and write semantics. It also supports web-based storage and management of massive data.

Figure 4-2 shows Amazon S3 APIs with which the object storage service is compatible.

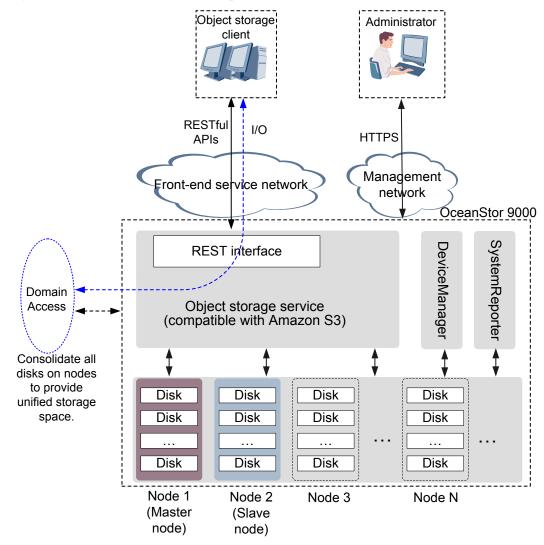


Figure 4-2 Object storage service (compatible with Amazon S3 APIs)

- Object storage service (compatible with Amazon S3 APIs) runs on all the nodes that
  provide the service. Metadata is evenly distributed so that performance bottlenecks
  caused by a certain overloaded node can be avoided. Elastic and seamless scale-out is
  supported, allowing a cluster to scale up to 100 PB.
- An OceanStor 9000 cluster that provides object storage service (compatible with Amazon S3 APIs) has a primary node and a secondary node as management nodes to manage accounts and users of the object storage service.
- Object storage service (compatible with Amazon S3 APIs) provides RESTful APIs externally. Clients can access the OceanStor 9000 to store data through the RESTful APIs.

#### NOTE

Clients can be interpreted as Internet-based applications of the object storage service and are provided, planned, or developed by the OceanStor 9000 customers. Clients can be application software and visual websites.

 Administrators can securely access DeviceManager and SystemReporter (optional performance analysis tool) to manage system resources and monitor system status and performance through HTTPS. DeviceManager supports RESTful APIs. Third-party developers can use RESTful APIs to access DeviceManager open resources and perform operations such as querying alarms, monitoring performance, and allocating resources.

## 4.1.4 Object Storage Service (Compatible with OpenStack Swift APIs)

The OceanStor 9000 provides object storage service (compatible with OpenStack Swift APIs). Storage systems that provide object storage service (compatible with OpenStack Swift APIs) are highly available and distributed systems used to create scalable and redundant object storage. If standardized servers are used, PB-level data can be stored.

The object storage service that is compatible with OpenStack Swift has a distributed and scalable architecture and provides RESTful APIs. All the resources can be located by URLs. Data is organized in an Account/Container/Object non-mesh structure. Massive data can be stored.

Figure 4-3 shows OpenStack Swift APIs with which the object storage service is compatible.

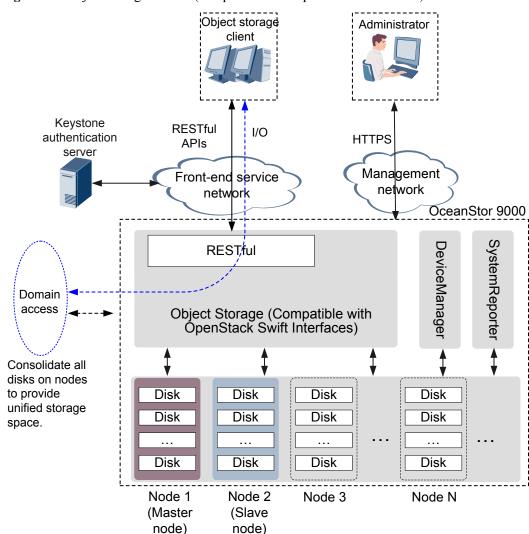


Figure 4-3 Object storage service (compatible with OpenStack Swift APIs)

- The object storage service (compatible with OpenStack Swift APIs) runs on all the nodes
  that provide the service. Metadata is evenly distributed so that performance bottleneck
  caused by a certain overloaded node can be avoided. Elastic and seamless scale-out of
  three to 288 nodes is supported. A cluster can scale to 100 PB.
- The OceanStor 9000 cluster employs one master node and one slave node as management nodes to manage Accounts of the object storage service.
- The object storage service provides RESTful APIs (compatible with OpenStack Swift), through which clients can access OceanStor 9000 to implement data storage.

#### NOTE

These clients can be regarded as the Internet-based object storage applications and are provided or developed by users of OceanStor 9000. The clients can be application software or visual websites.

- Users' access to the object storage service needs to be authenticated by the Keystone server.
- Administrators can securely access DeviceManager and SystemReporter through HTTPS to manage system resources and monitor system status and performance. DeviceManager also supports RESTful APIs.
- The OceanStor 9000 can be connected to external NMSs that are based on the SNMP protocol and the SMI-S protocol.

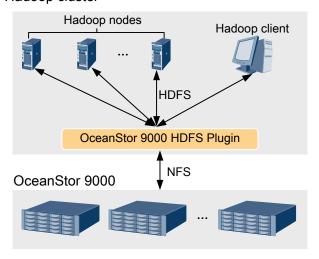
#### 4.1.5 HDFS Interface

By interconnecting with FusionInsight Hadoop or Cloudera Hadoop (Hadoop offerings released by Huawei and Cloudera respectively), the HDFS feature of the OceanStor 9000 supports distributed file storage with increased capacity utilization, reliability, and performance while supporting value-added services such as snapshot, remote replication, and quota management.

The OceanStor 9000 HDFS feature is also named the HDFS interface feature. Huawei HDFS Plugin deployed on Hadoop nodes and the Hadoop client can convert the HDFS-based access requests to NFS-based access requests. In this way, Hadoop service data can be directly stored to OceanStor 9000. **Figure 4-4** shows the HDFS feature.

Figure 4-4 HDFS feature

#### Hadoop cluster



OceanStor 9000 supports the following Hadoop products and versions:

- FusionInsight Hadoop V100R002C50
- Cloudera Hadoop 5.4.1

NOTE

The OceanStor 9000 cannot interconnect with Apache Hadoop that belongs to the open-source community.

## 4.1.6 NDMP Backup

The OceanStor 9000 provides the Network Data Management Protocol (NDMP) backup function, enabling backup data to be transferred from the storage system to the backup media.

NDMP is an open protocol for enterprise-level data management. It is used to control backup, recovery, and data transfer between the primary memories and secondary memories. NDMP works based on the client/server model. The OceanStor 9000 serves as the NDMP server and the data management application (DMA) serves as the client.

The OceanStor 9000 supports NDMP V4 and adopts 3-way networking mode, as shown in **Figure 4-5**.

Figure 4-5 3-way networking solution

OceanStor 9000
(NDMP server, primary storage)

LAN

Primary server Media server

Backup server
(NDMP client)

Tape library
(Secondary storage)

---- NDMP control flow
Backup data flow

In this mode, the OceanStor 9000 is the primary storage. The backup server uses the NDMP-based management interface to instruct the OceanStor 9000 to back up data to a tape library on the network. During the backup process, information about the files and directories written to tapes is transferred to the backup server. Then the backup server generates indexes to maintain the information.

The OceanStor 9000 provides snapshot-based backup and source directory-based backup, and supports following backup functions:

Full folder backup

Data in a directory is all backed up.

Differential Incremental Backup

Only the file changed since last backup is backed up.

Only the file changed since last full backup is backed up.

Cumulative Incremental Backup

#### NOTE

The definition of incremental backup may vary depending on different backup software.

The OceanStor 9000 provides the following recovery functions:

• Full Recovery

All data is recovered.

• Direct Access Recovery (DAR)

Backup software locates and accesses specific data based on user requirements and then recovers it.

Non-DAR

Backup software finds data to be recovered in tape libraries and then recover it.

## 4.1.7 High Data Reliability (InfoProtector)

The OceanStor 9000 provides the unique InfoProtector feature to ensure high data reliability.

The OceanStor 9000 can ensure access to and automatically recover data from physical devices that have failed. Data protection levels (data redundancy ratio) can be configured based on site requirements to strike an optimal balance between data reliability and storage space utilization.

Traditional RAID-based protection stores data onto different disks of one node. Data cannot be restored when the node fails. The OceanStor 9000 stores data on different nodes for redundancy to prevent data loss.

The OceanStor 9000 protects file data in original data count (ODC)+redundant data count (RDC) mode (N+M mode for short). ODC or N represents the number of source data slices, and RDC or M represents the number of redundant data slices. In N+M mode, the OceanStor 9000 allows a maximum of M disks or nodes to fail at one time. During this period, data can be accessed normally and data recovery will be automatically initiated. **Figure 4-6** shows the data protection process for a 6-node OceanStor 9000 cluster protected in 4+2 mode.

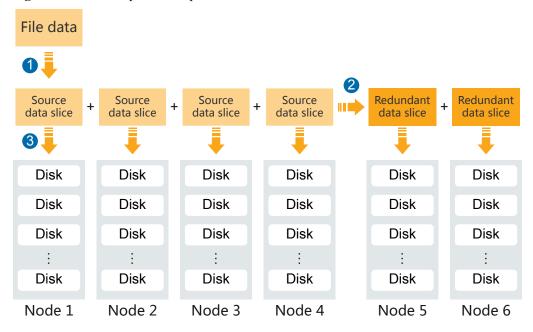


Figure 4-6 File data protection process

- 1. File data is divided into multiple source data slices. ODC ranges from 2 to 16. ODC is automatically calculated by the OceanStor 9000 and cannot be manually configured. In this example, ODC is 4.
- 2. Redundant slices are generated for source data slices.
  - File system storage: RDC ranges from 1 to 4 and can be configured for a specific empty directory. In this example, RDC is set to 2 (default value).
  - Object storage: RDC ranges from 1 to 3 and can be configured for a specific object storage account. In this example, RDC is set to 2 (default value).
- 3. Data slices are evenly stored onto different disks across all nodes.

ODC+RDC:RDN, also known as N+M:B, is a data protection mode used only in special circumstances. RDN or B represents the number of redundant data nodes. In this mode, the OceanStor 9000 allows the failure of up to M disks or B nodes at one time. N+M:B protection mode helps reduce the number of required nodes.

Services on a faulty node or disk are failed over to another node for continuous system operation. When the node or disk is recovered, the OceanStor 9000 quickly recovers the lost data by reconstructing data across nodes, ensuring high data reliability.

- If a faulty disk is recovered within 10 minutes, the system automatically begins data recovery on data changes during the fault period.
- If a faulty disk is recovered after more than 10 minutes, data on the disk is automatically recovered and data stored before the fault and data changed during the fault are recovered.
- Data recovery is automatically started for a node immediately after the node is recovered. Only the data changed during the fault is recovered. You can also perform full data recovery using the network management system to recover data stored before the fault and data changed during the fault.

#### 4.1.8 Permission Control

The OceanStor 9000 employs permission control to improve service access security.

Permission control includes system access permission control and directory access permission control. With permission control, the OceanStor 9000 allows only authorized user access and provides storage space for only specified users, protecting users' private data against unauthorized access and facilitating the establishment of upper-layer services.

During system access permission control, the OceanStor 9000 or authentication server authenticates users accessing the NAS services. The OceanStor 9000 supports multiple authentication modes, as described in **Table 4-1**.

Table 4-1 Authentication modes

Authentication Mode	Description	Application Scenario
Anonymous access	Direct access without identity authentication	CIFS sharing
Local user authentication	Authenticated by the OceanStor 9000	
AD domain authentication	Authenticated by the AD domain server	
Client IP address or the IP address segment of clients authentication	Authenticated using the client IP address or the IP address segment of clients	NFS sharing
Client host name authentication	Authenticated using the host name of access clients	
LDAP authentication	Authenticated by the LDAP domain server	
NIS authentication	Authenticated by the NIS domain server	
Anonymous access	Access without identity authentication	FTP sharing
Local user authentication	Authenticated by the OceanStor 9000	

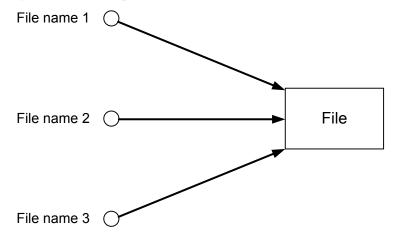
During directory access permission control, users' access permission for shared directories are controlled. System administrators can grant other users read-only or read/write permission for directories to control users' access to and operation on shared directories.

#### 4.1.9 Windows Hard Link

The OceanStor 9000 supports Windows hard links. Hard links of the shared directory can be created on the Windows client for media asset management.

A hard link represents one or multiple names of a file. Users can link multiple file names in the same or different directories to the same file. The number of a file's hard links indicates the number of file names. In **Figure 4-7**, the file has three file names, indicating that the number of hard links is three.

Figure 4-7 An example of hard links



Hard links enable the clients to modify the same file simultaneously using multiple file names under different or the same directory, so that one copy of data can support multiple service systems.

Many media asset service systems are based on the Windows platform. The systems need to frequently access the same materials but cannot directly access the files due to restrictions of directories. When hard links are not supported, one copy of data must be duplicated for several times, affecting performance and wasting space.

When hard links are supported, one copy of data can be sent to multiple service systems, saving storage space and improving competitiveness in the area of media assets.

You can only create hard links for local file systems by running the **mklink** and **fsutil hardlink** commands on the Windows command-line interface. On the OceanStor 9000, you can only create hard links by using the third-party open source suite **UnxUtils** or by calling WINDOWS API interfaces through application software.

The OceanStor 9000 has the following restrictions on Windows hard links:

- Hard links can only be created using SMB 2.0/SMB 2.1/SMB 3.0.
- Hard links and source files must be created under the same mounted directory.
- Hard links of a directory are not supported.
- Hard links of files cannot be created across quota directories.
- Hard links cannot be created for files in the WORM directory.
- Hard links cannot be created for snapshot files.
- Hard links are protected by the snapshots in the source files' directory.

## 4.1.10 Unified Namespace

The OceanStor 9000 provides a unified namespace to facilitate the allocation, maintenance, and expansion of storage resources.

The OceanStor 9000 provides a unified file system namespace. This namespace contains all operation objects such as files, share settings, and management permissions. You can use the unified namespace to manage all available capacity of the storage system.

The namespace is automatically created upon startup of the OceanStor 9000 and is named after the OceanStor 9000. All shared directories are subdirectories in the unified namespace.

As the root directory of the storage system, the unified namespace can be shared and accessed.

The unified namespace provides a unified storage pool where storage space can be centrally managed and allocated. New storage space is added to the storage pool after capacity expansion, without the need to maintain multiple separate file systems.

## 4.2 Management Function

The OceanStor 9000 provides management functions such as data consistency scan, high system security, and unified namespace.

## 4.2.1 Data Consistency Scanning

The OceanStor 9000 provides data consistency scanning that periodically checks data in the background without affecting services, thereby discovering disk errors in advance to prevent silent corruption from causing data loss.

If a data error is found, the OceanStor 9000 automatically attempts to retrieve the lost data using redundancy algorithms. If the amount of lost data exceeds the upper limit allowed by the data protection level, an alarm is reported.

The OceanStor 9000 supports both quick scanning and in-depth scanning.

Quick scanning

Only the size and attribute of a data object are checked. Quick scanning consumes relatively a small amount of overhead and time.

In-depth scanning

In addition to the size and attribute, the data object is also read and checked. In-depth scanning can detect more types of errors but consumes relatively a large amount of overhead and time.

You can set an automatic startup cycle for a data consistency scanning task, so that the task is implemented during off-peak hours. In addition, you can specify the implementation ratio of quick scanning to in-depth scanning.

## 4.2.2 High System Security

The OceanStor 9000 security framework ensures system security based on the following layers: storage device security, storage network security, and storage service security. In addition, the management plane security design ensures security among these layers.

- Storage device security
  - Operating system hardening
  - Patch management
  - System antivirus
  - Web security
- Storage network security
  - Plane isolation
  - Channel security

- Storage service security
  - Storage resource access control
  - Certification and authentication
- Storage management security
  - Management system user security
  - Password security
  - Certification and authentication
  - Log and alarm management

## 4.2.3 Flexible Expansion

The OceanStor 9000 can be flexibly expanded to enable easy deployment and capacity expansion.

The OceanStor 9000 adopts a distributed cluster architecture where all nodes are symmetrically deployed and provide services.

Metadata and data are evenly distributed across all nodes of the OceanStor 9000 cluster, eliminating metadata or data node single points of failure. The clustered management of metadata improves the reliability and processing capabilities of the OceanStor 9000, enabling linear performance growth with capacity.

The OceanStor 9000 supports the seamless scale out from 3 to 288 nodes. You can deploy three nodes in the initial setup phase and scale the system out later on demand. Capacity expansion requires no service interruption or migration. The flexible expansion of the OceanStor 9000 protects your device investment.

## 4.2.4 Centralized Management

The OceanStor 9000 can be centrally and remotely managed and maintained with ease.

The OceanStor 9000 provides both a graphical user interface (GUI) and a command-line interface (CLI).

An administrator can log in to the management GUI using the management IP address to monitor and configure the OceanStor 9000, allocate resources, collect performance statistics, and view alarms. There is no need to deploy an independent network management server. The failure of any node does not affect user login and use of the network management system.

The OceanStor 9000 also provides various remote management and maintenance functions. You can manage all devices using the network management system and only need to install and replace hardware in equipment rooms.

In addition, when multiple OceanStor 9000 clusters are deployed, you can deploy a unified network management system (eSight) to provide a unified device management GUI for these clusters.

## 4.2.5 Data Balancing

The OceanStor 9000 provides the data balancing function to ensure that data is evenly distributed across nodes when a node or disk is added, disk pre-failure occurs, or a node is enabled or disabled.

The OceanStor 9000 supports data balancing in the following scenarios:

#### Node addition

After a node is added, the disk space of the node is idle. The system migrates some data from other nodes to the node to balance workloads.

#### Disk addition

After a disk is added, the disk space of the disk is idle. The system migrates some data from other disks to the disk to balance workloads.

#### Node deletion

Before a node is deleted, the system migrates data from the node to other nodes to balance workloads.

#### Disk pre-failure

When a disk is set to the pre-failure state due to a fault, the system evenly migrates its data to other operating disks, and then the disk is allowed to stop functioning.

#### Slow disk

When a slow disk (low-performance disk) is detected, the system migrates its data to other operating disks.

#### Node enabling/disabling

After a node is disabled, the system migrates data from the node to other nodes to balance workloads.

After the node is enabled, the disk space of the node is idle. The system migrates some data from other nodes to the node to balance workloads.

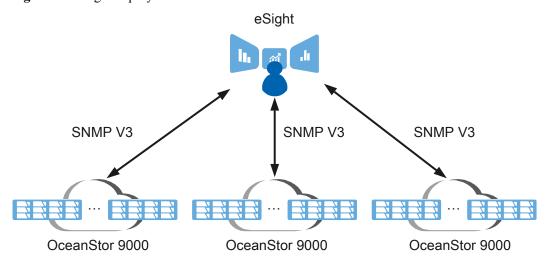
The OceanStor 9000 supports manual and automatic data balancing and uses the flow control mechanism to minimize the impact of data balancing on ongoing services.

## 4.2.6 Unified Network Management

The eSight ICT management system (eSight for short) is used to manage multiple sets of OceanStor 9000 clustered systems and other storage systems.

eSight is deployed independently and can monitor multiple sets of OceanStor 9000 systems, as shown in **Figure 4-8**.

Figure 4-8 eSight deployment



eSight provides the following functions:

- In system monitoring scenarios, eSight allows you to check the basic parameters, system status, and space usage of the registered OceanStor 9000 systems and other storage systems.
- In alarm management scenarios, eSight allows you to check, confirm, and clear the alarms generated on the OceanStor 9000 systems or other storage systems.
- In configuration scenarios, eSight allows you to use the single sign-on function to log in to the DeviceManager and perform system settings.

#### 4.2.7 Unified View

The OceanStor 9000 provides unified view client software. The software can consolidate multiple OceanStor 9000 clusters in different areas and provides a unified view (global namespace).

In video surveillance scenarios, to prevent unexpected events, it is required to check video surveillance data scattered across different areas. The unified view client consolidates multiple underlying OceanStor 9000 clusters into a storage resource pool and conceals the physical distribution of storage resources because the upper-layer applications do not care the distribution information of the underlying storage systems. The unified view client uses the network file system (NFS) to access OceanStor 9000 clusters that are deployed in different areas.

The unified view client provides the standard POSIX file system interface. By the interface, the upper-layer video surveillance management platform can access the files on the unified view

Figure 4-9 shows the unified view client applications.

Video surveillance management platform

POSIX

Unified view client

NFS

NFS

OceanStor 9000

OceanStor 9000

OceanStor 9000

Figure 4-9 Unified view client applications

The unified view client provides the following functions:

- The unified view client allows you to access files in multiple storage domains.
- When a storage system is added, the unified view client can retrieve and access new storage resources.

• When a storage system is offline or deleted, the unified view client deletes the offline storage resources from the global namespace.

## 4.3 File System Basic Features

The OceanStor 9000 provides basic features such as InfoTier, InfoAllocator, and InfoEqualizer.

#### 4.3.1 InfoTier

The OceanStor 9000's InfoTier feature provides automatic storage tiering.

InfoTier divides storage space into different storage tiers. Based on such file attributes: file name, file path, file size, create time, change time, last access time, owning users/groups, I/O count, I/O popularity, storage type, and SSD acceleration, InfoTier stores files to different storage tiers, meeting user requirements for file processing speed, transmission bandwidth, and storage capacity, bringing storage space into full play, improving the storage access performance, and reducing the total cost of deployment.

Transparent to upper-layer services, InfoTier automatically identifies cold and hot data and migrates data between storage media, as shown in **Figure 4-10**.

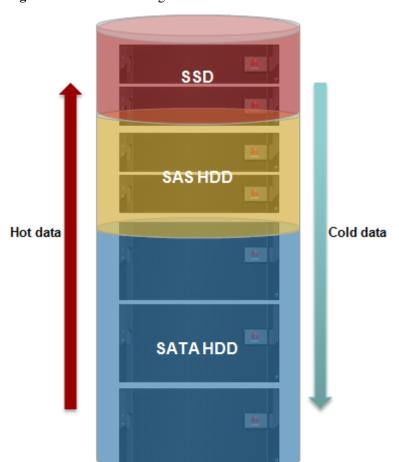


Figure 4-10 Schematic diagram of InfoTier

InfoTier provides the following functions:

- Node pool: Nodes of the same physical type can form a node pool. If a node pool consists of P12, P25, P36, and C36 nodes, the node pool can contain 3 to 20 nodes. If a node pool consists of C72 nodes, the node pool can contain 2 to 10 nodes. If more than maximum nodes of the same physical type exist, multiple node pools need to be formed.
- Disk pool: Disks of all nodes in each node pool form disk pools by disk type. Disk pools are used to store metadata and data.
- Tiering: Based on service requirements, you can set storage tiers and specify the mapping relationship between a node pool and a storage tier.
- File pool policy: You can specify the mapping relationship between file attributes and each storage tier to define a policy for storing and migrating data between different storage tiers.
- Watermark policy: You can set the high watermark and read-only watermark for each node pool to control data storage and migration.

The OceanStor 9000 allows you set node pools based on the physical types of storage servers and specify the mapping relationship between a node pool and a storage tier.

- For newly written data, the OceanStor 9000 allows you to define the initial storage locations by specifying the mapping relationship between the file storage path, file name, user name, user group name, and storage tiers, thereby properly distributing data in the storage system.
- For stored data, the OceanStor 9000 monitors changes in file attributes and automatically
  migrates data that meets the file migration policy to the corresponding storage tiers,
  thereby providing the optimal storage performance for different applications.

As a value-added feature of the OceanStor 9000, a license is required for using InfoTier. If the license is unavailable, the OceanStor 9000 will provide only tier management. File pool policy and automatic file migration will not be available.

## 4.3.2 InfoAllocator

The OceanStor 9000 provides the InfoAllocator feature to manage quotas.

InfoAllocator enables the OceanStor 9000 to:

- Manage storage space for users or user groups to prevent disproportionate storage space occupation.
- Collect statistics on and check storage space for users or user groups.
- Plan storage space for users or user groups in an appropriate way.
- Create default quota. Default quota is the default available quota of a designated directory for any user or any user in a user group. Default quota enables the system to automatically limit usage and collect usage information of any user for a designated directory. Quota configuration operations are simplified because you do not need to create quota for a newly created user.

Quota types supported by the OceanStor 9000 include:

- Capacity quota: used to manage and monitor storage space usage.
- File number quota: used to manage and monitor file number and usage.

Quotas are categorized based on their configuration method as follows:

- Calculating quota: monitors only storage capacity and file quantity but not storage resource usage.
- Mandatory quota: monitors storage capacity usage or file quantity as well as storage resource usage based on the preset thresholds.

**Table 4-2** describes the thresholds for mandatory quotas.

Table 4-2 Thresholds for mandatory quotas

Threshold Type	Description
Recommended threshold	When storage space usage or file quantity reaches a recommended threshold, the OceanStor 9000 allows data writes and reports an alarm.
Soft threshold	When storage space or file number reaches a soft threshold, the OceanStor 9000 still allows data writes for a grace period. However, after the grace period expires, the OceanStor 9000 forbids data writes and reports an alarm.  You need to configure a grace period when configuring a soft threshold.
Hard threshold	When storage space usage or file number reaches a hard threshold, the OceanStor 9000 forbids data writes and reports a run log and an alarm.

InfoAllocator is a license-controlled value-added feature of the OceanStor 9000.

## 4.3.3 InfoEqualizer

The OceanStor 9000 provides the InfoEqualizer feature that manages client connections to the OceanStor 9000 and implements load balancing and failover.

InfoEqualizer provides the following functions to improve availability, performance, and reliability of the OceanStor 9000:

- Enables a client to access the OceanStor 9000 using only a domain name, facilitating client connection.
- Intelligently allocates client requests to idle nodes for load balancing, ensuring high client access efficiency.
- Enables failover of services from faulty nodes to other normal nodes, improving system reliability.

Figure 4-11 shows a schematic diagram of InfoEqualizer.

In domain-name-based access mode, client access requests are intelligently distributed based on loads.

In a failure scenario such as a node failure, the failure is automatically detected and services are automatically migrated to other functional nodes.

Figure 4-11 Schematic diagram of InfoEqualizer

OceanStor 9000 supports following load balancing functions:

- Round-robin: Nodes process client connection requests in turn.
- CPU usage: The node with the lowest CPU usage processes client connection requests.
- Number of node connections: The node with the least clients processes client connection requests.
- Node throughput: The node with the lowest throughput processes client connection requests.
- Node load processing capability: The node with a high load processing capability is selected to process client connection requests. The OceanStor 9000 calculates the load processing capability of a node based on the node's CPU and bandwidth configurations. If the node is overloaded, the OceanStor 9000 decreases its load processing capability. For example, in a zone that contains node A whose capability value is 30 and node B whose capability value is 70, node A is selected at a 30% probability, and node B is selected at a 70% probability.

#### 4.3.4 InfoLocker

The OceanStor 9000 provides the InfoLocker function, that is, enterprise write only read many (WORM) to archive data.

Write only read many (WORM) is a technology that allows data to be only read multiple times once being written. Specifically, after data is written, you cannot change, delete, or move it within the file protection period but you can read it as you wish.

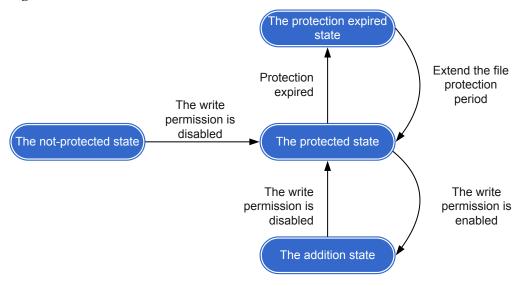
The enterprise WORM function provided by the OceanStor 9000 can be used to protect archived data in banks or governments from being changed by CIFS/NFS/FTP users except that the **root** administrator can delete it.

An administrator that has the WORM management permission can set the WORM reference clock on the DeviceManager, preventing files from expiring in advance because the local

clock is mistakenly changed. Before storing WORM files, the administrator needs to create a WORM directory and set the expiration time and automatic lock time for the WORM directory.

Figure 4-12 shows the WORM file states.

Figure 4-12 WORM file status



- A file in the not-protected state can be changed or deleted.
- When the write permission is disabled from a file, the file is in the protected state and can be accessed but cannot be deleted or changed.
- When a file is in the protection expired state, the file can be deleted or accessed but cannot be changed. After extending the file protection period, the file can enter the protection state again.
- If the write permission is granted to an empty file in the protection state, that file enters the addition state. New data can be written to a non-empty file in the addition state. However, existing data in such a file cannot be modified.
- After disabling the write permission for the file in the addition state, the file enters the protection state again.

The license-required InfoLocker is an optional function of the OceanStor 9000.

## 4.3.5 InfoStamper

The OceanStor 9000 provides the InfoStamper function, that is, directory-level snapshot.

A snapshot is a fully usable copy of a specific data collection. Such a copy contains a static image of the source data at the copy point in time. A snapshot can be used for manufacturing tests and data backup and recovery.

The OceanStor 9000 provides the directory-level snapshot function that enables snapshots to be created for any directory (except the root directory) in a file system, facilitating backup of mission-critical data. The OceanStor 9000 allows a maximum of 2048 snapshots to be created for a directory and allows a maximum of 8192 snapshots (less than 4096 is recommended) to be created for the system.

The OceanStor 9000 adopts two snapshot technologies, Copy On Write (COW) specific to metadata and Redirect On Write (ROW) specific to data, greatly increasing system space utilization.

- COW: Before new data is written, the original data is copied to another location or object. Then the new data overwrites the original data.
- ROW: New data is written to a new location or object without overwriting the original data.

The OceanStor 9000 supports manual snapshot and timing snapshot.

- Manual snapshot: A user creates snapshots for a specific directory at the current point in time. Manual snapshot is applicable to a manual backup scenario.
- Timing snapshot: Snapshots can be automatically created based on the timing snapshot policy preset by a user. Timing snapshot is applicable to an automatic backup scenario.

The license-required InfoStamper function is an optional function of the OceanStor 9000.

## 4.3.6 InfoReplicator

The OceanStor 9000 provides the InfoReplicator function, that is, asynchronous remote replication to create data mirrors between the production site and the remote site, implement remote data disaster recovery (DR), and ensure service continuity.

Remote replication is a core technology for DR and backup, as well as a foundation of remote data synchronization and DR. Data at the primary site can be backed up to the remote DR site to ensure that the data can be still accessible when the primary site is damaged due to some unexpected reasons. If the primary system at the production site fails or the primary system must be taken offline due to system maintenance or upgrade, the secondary system at the remote site can provide services for the upper-layer applications, minimizing the downtime.

The OceanStor 9000 supports directory-level asynchronous remote replication. Its specific functions include:

#### Data Replication

In a remote replication task, one primary directory and one secondary directory compose a pair. The primary and secondary replication directories reside in different OceanStor 9000 clusters.

After a remote replication task is created, the system creates a snapshot for the primary directory as the base snapshot prepared for initial synchronization. Based on a preset synchronization mode (manual or automatic), the system replicates data from the primary directory to the secondary directory. After the replication, data on the secondary directory is the same as the base snapshot.

At some time, for example when the bandwidth is insufficient to support critical services, you probably do not want to copy data from the primary directory to the secondary directory. In such cases, you can split the secondary file system from the primary file system to suspend data synchronization.

Users can effectively control a remote replication process by performing synchronization and splitting.

#### Service Switchover

When the primary device in the production center needs to be upgraded or maintained, or fails due to disasters, the secondary device in the DR center takes over services. After disabling the write protection for the secondary directory, the client can mount the

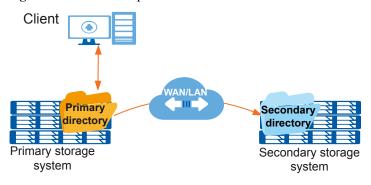
secondary directory. After the mounting, I/O services are stopped in the production center and are taken over by the secondary center.

#### • Service Recovery

When the primary device fails, the secondary device takes over the services. When the primary device is recovered, a remote replication pair can be rebuilt between the primary and secondary storage systems and data on the primary device can be recovered from the secondary device. After data is recovered, the primary device takes over the services again.

For example, in a DR scenario, data at the production site is periodically replicated to the DR site. When the primary site fails due to fires, floods, or earthquakes, services are switched over from the primary site to the secondary site, as shown in **Figure 4-13**.

Figure 4-13 Remote replication in a 1:1 DR scenario



The license-required InfoReplicator is an optional function of the OceanStor 9000.

## 4.3.7 InfoScanner

OceanStor 9000 provides InfoScanner, the antivirus function, to perform real-time and periodic scanning for files in the shared directory by connecting to the antivirus software to improve the overall security protection capability.

Figure 4-14 shows the network structure of InfoScanner.

Antivirus console Client

Antivirus console Client

Antivirus server Storage system

Scan policy Antivirus software

Scan directory Huawei Antivirus Agent

Figure 4-14 InfoScanner network structure

**Table 4-3** provides details about the network structure. The front-end service network of OceanStor 9000, antivirus servers, and AD domain servers are interconnected.

Table 4-3 Network architecture description

Device	Description	
AD domain server	<ul> <li>Scan authentication users must be AD domain users and need to be added to AntivirusGroup (default local authentication user group in OceanStor 9000).</li> </ul>	
	NOTE When the antivirus switch is enabled and antivirus servers are configured, scan authentication users can access storage resources in OceanStor 9000 without being authenticated, namely, they have full control.	
	• To ensure that Scan authentication users can access the OceanStor 9000 by using antivirus servers, antivirus servers and OceanStor 9000 need to be added to the AD domain.	
Antivirus console	The antivirus console is provided by Rising to collect feedback from antivirus software. The antivirus console can be deployed on one antivirus server or another kind of server.  NOTE	
	The antivirus console needs to be deployed only after Rising virtualization system security software is installed on the antivirus server.	
Client	A user accesses the OceanStor 9000 through a client and writes data to the OceanStor 9000.	

Device	Description
Antivirus server	<ul> <li>Antivirus servers access storage resources in scan directories in OceanStor 9000 using the CIFS protocol. You are advised to configure multiple antivirus servers to prevent antivirus function failure when the only server fails and to accelerate service processing.</li> </ul>
	<ul> <li>Antivirus Agent Software is downloaded using DeviceManager and installed on every antivirus server to trigger antivirus software to scan files in scan directories.</li> </ul>
	The following types of antivirus software are supported:
	- Rising virtualization system security software 1.0.0.89 or later
	- Symantec Protection Engine 7.5 or later
	- Symantec Endpoint Protection 12.1 or later
	- Trend Micro ServerProtect 5.8 or later
	Rising virtualization system security software mainly scans files in scan directories, backs up infected files to isolation directories, kills viruses of files in the scan directories, and restores files from isolation directories to scan directories. Symantec Protection Engine, Symantec Endpoint Protection, and Trend Micro ServerProtect scan files in scan directories for viruses and kill viruses if any. Antivirus software must be installed on each antivirus server.
Storage system (OceanStor	<ul> <li>A scan directory refers to the directory to be scanned in OceanStor 9000.</li> <li>A directory can be configured as a real-time or periodical scan directory.</li> <li>Antivirus servers access a scan directory using the CIFS protocol.</li> </ul>
9000)	<ul> <li>A scan policy specifies the non-scan period, non-scan file types, and the maximum size of files that can be scanned for a scan directory.</li> </ul>
	The OceanStor 9000 interconnects with the antivirus software deployed on an antivirus server to scan the scan directories for viruses.

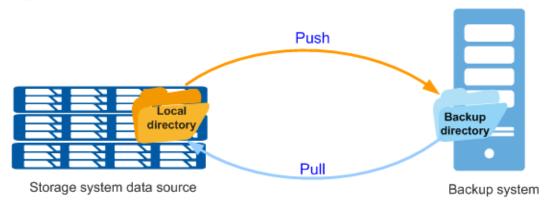
## 4.3.8 InfoMigrator

The OceanStor 9000 provides asynchronous InfoMigrator in full and incremental manners, allowing you to synchronize data from the local OceanStor 9000 to a remote backup system (an OceanStor 9000 or another type of storage system) that provides an NFS shared directory. InfoMigrator helps you easily and remotely back up and recover data, improving data reliability.

The OceanStor 9000 supports Push and Pull file synchronization scenarios, as shown in **Figure 4-15**.

- Push: The OceanStor 9000 replicates data from a local directory to a backup directory in Push mode for data backup.
- Pull: The OceanStor 9000 takes the initiative to replicate data from a backup directory to a local directory in Pull mode for local data recovery.

Figure 4-15 File synchronization scenarios



In the Push scenario, the OceanStor 9000 data source acts as an NFS client, and the backup system acts as an NFS server. The data source is mounted to the backup system through NFS. The local directory acts as the source directory, and files are synchronized to the backup directory in the backup system. Extended attributes of files and sub-directories such as sharing, quota, and ACL are not synchronized. In the Pull scenario, files are synchronized from the backup directory to the local directory in the OceanStor 9000.

## 4.4 File System Enhanced Features

OceanStor 9000 provides enhanced features such as InfoRevive, InfoCopy, and InfoTurbo.

## 4.4.1 InfoRevive

The OceanStor 9000 provides the InfoRevive function to ensure that video surveillance data remains accessible in extreme scenarios.

In video surveillance scenarios, it is unacceptable that video surveillance data becomes inaccessible after being slightly damaged. In this case, users expect that videos can be normally played even though the videos may be unclear occasionally.

To protect user data and provide better video surveillance service, the OceanStor 9000 delivers the InfoRevive function. InfoRevive maintains data to be accessible even though the number of failed nodes or disks exceeds the maximum allowed limit, ensuring continuity of the video surveillance services.

InfoRevive supports two operating modes:

#### • Read error tolerance

With read error tolerance enabled, when the number of failed nodes or disks exceeds the maximum allowed limit, you can still read a proportion of damaged video file data. The system security is enhanced.

#### Read/write error tolerance

With read/write error tolerance enabled, when the number of failed nodes or disks exceeds the maximum allowed limit, you can still read a proportion of damaged video file data and continue writing data to the file. The service continuity and security are enhanced.

The license-required InfoRevive function is an optional function of the OceanStor 9000 and can be only used in the video surveillance scenario.

#### 4.4.2 InfoTurbo

The OceanStor 9000 provides the InfoTurbo feature (namely the performance acceleration feature) that includes intelligent prefetch, SMB3 Multichannel, and NFS protocol enhancement.

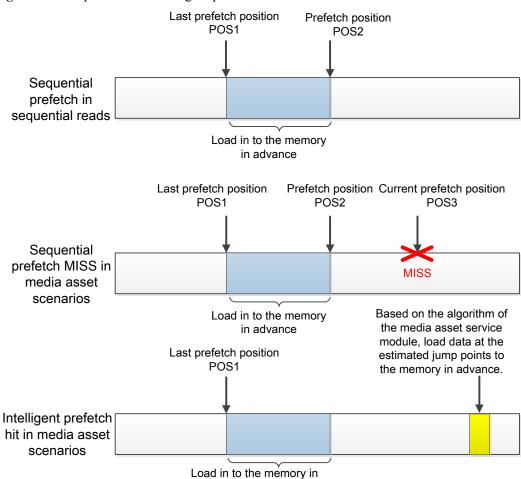
## **Intelligent Prefetch**

The OceanStor 9000 supports intelligent prefetch to provide a higher cache hit ratio for users in the media asset scenarios. In delay-sensitive scenarios, performance can be greatly improved.

Mainstream operating systems divide read modes into random reads and sequential reads and only prefetch sequential reads. This guarantees a high prefetch hit ratio and ensures data reading efficiency and an excellent coverage rate. Sequential reads are simple and common while random reads are unpredictable.

For media asset scenarios, the OceanStor 9000 provides the intelligent prefetch technology based on sequential prefetch. **Figure 4-16** shows sequential and intelligent prefetch.

Figure 4-16 Sequential and intelligent prefetch



advance

Sequential prefetch in sequential reads

If the prefetch thread determines that service I/Os are sequential, the thread will load the data between POS1 and POS2 to the cache when the upper-layer application reads data from POS1. If the next I/O needs to read data between POS1 and POS2, data will be read directly from the cache instead of the disk.

• Sequential prefetch MISS in media asset scenarios

The non-linear editing (NLE) software used in media asset scenarios needs to process multiple types of media files. Many media files are prefetched depending on their encoding formats instead of in sequence. For example, files of certain formats do not store video and audio data successively and have their own video and audio storage areas.

The play software needs to read video and audio data. Frequently, the video and audio data are not in the cache. The software has to read the data from the disk. This will cause great latency. As shown in **Figure 4-16**, the prefetch thread will load the data between POS1 and POS2 to the cache in advance based on the previous prefetch rules. However, the I/O operations do not read data between POS1 and POS2, and jump to other positions to read data. In this case, the cache cannot be hit.

Intelligent prefetch hit in media asset scenarios

To improve storage cache hit ratio in media asset scenarios, the OceanStor 9000 compares encoding formats of multiple media asset files on various mainstream NLE software and provides a series of service read modules. Based on the previous service modules, the OceanStor 9000 uses sequential prefetch and some forecast algorithms to make a forecast and loads the data that may be needed for later services to the cache in advance.

This feature supports the following NLE software: Final Cut Pro, Adobe CC, and Avid Media Composer.

This feature supports the following video encoding formats: mxf and mov.

#### SMB3 Multichannel

The OceanStor 9000 provides the SMB3 Multichannel function to greatly improve service performance and reliability. In addition, if one channel fails, the OceanStor 9000 transmits data over another channel to prevent services from being affected.

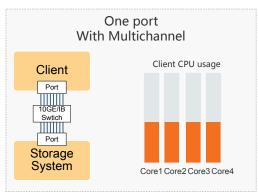
In CIFS file sharing scenarios, if a client that uses SMB3.0 (which is delivered with Windows 8, Windows 2012, and their later versions) is equipped with two or more GE/10GE/IB network ports of the same type or with one GE/10GE/IB network port that supports Receive-Side Scaling (RSS), the client will set up multiple channels with the OceanStor 9000.

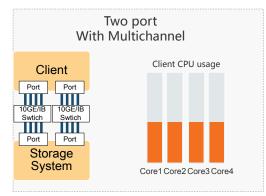
**Figure 4-17** shows a schematic diagram of SMB3 Multichannel. By bringing multi-core CPUs and bandwidth resources of clients into full play, SMB3 Multichannel greatly improves service performance. In addition, if one channel fails, SMB3 Multichannel transmits data over another channel, thereby improving service reliability.

Without Multichannel

Client
Port
TIGGE/IB
Swlich
Storage
System
Core1Core2Core3 Core4

Figure 4-17 Schematic diagram of SMB3 Multichannel





#### **NFS Protocol Enhancement**

OceanStor 9000 supports enhanced functions of the NFS protocol to greatly improve service performance on the client.

The NFS protocol enhancement feature is a performance acceleration feature provided by OceanStor 9000. By configuring multiple network ports and installing NFS protocol optimization plug-in DFSClient on MAC OS X and RHEL clients, concurrent connections can be established between a client and the storage system, thereby increasing the access bandwidth. Cache optimization is enabled on the Mac OS X client for improved access performance to adapt to 4K video editing in the media assets scenario.

#### Multi-Connection Mode on the Mac OS X Client

This part describes the establishment of multiple connections and the data transmission process on the Mac OS X client.

- 1. The Mac OS X client sends a connection request to OceanStor 9000.
- 2. OceanStor 9000 returns multiple dynamic front-end service IP addresses. Each of the IP addresses is connected with an IP address on the client, as shown in **Figure 4-18**.

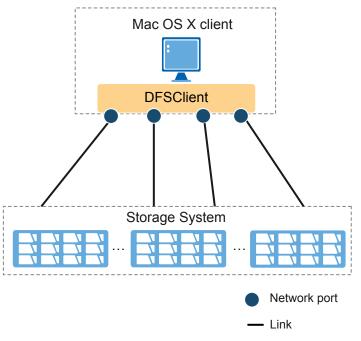


Figure 4-18 Establishment of multiple connections on the Mac OS X client

The number of connections that can be established using DFSClient is the number of IP addresses to be connected on the client (M) or the dynamic front-end service IP addresses of OceanStor 9000 (S), specifically:

- When  $M \ge S$ :
  - OceanStor 9000 returns all the dynamic front-end service IP addresses. The number of connections that can be established is S.
- When M < S:
  - OceanStor 9000 randomly returns M dynamic front-end service IP addresses. The number of connections that can be established is M.
- 3. When upper-layer services read and write video files concurrently, the system selects one connection to transmit one file. When all the connections are polled, the system starts a new round of polling, as shown in **Figure 4-19**.

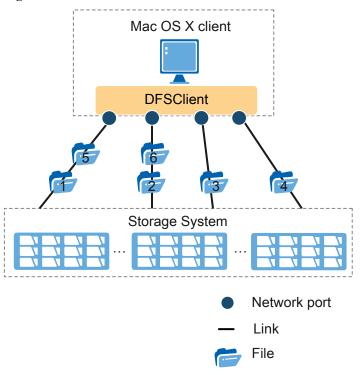


Figure 4-19 Data transmission mode on the Mac OS X client

Compared with the single-connection mode, the multi-connection mode increases the bandwidth provided for the Mac OS X client accessing OceanStor 9000.

During multi-layer video editing, video files are accessed concurrently. DFSClient can distribute files to connections evenly, thereby accelerating data transmission.

## Multi-Connection Mode on the RHEL Client

This part describes the establishment of multiple connections and the data transmission process on the RHEL client.

- 1. The RHEL client sends a connection request to OceanStor 9000.
- 2. OceanStor 9000 returns multiple dynamic front-end service IP addresses. Each of the IP addresses is connected with all the IP addresses on the client, as shown in **Figure 4-20**.

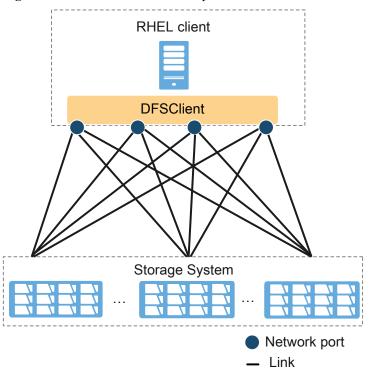


Figure 4-20 Establishment of multiple connections on the RHEL client

The number of connections that can be established =  $R \times S$ , where:

- R indicates the number of IP addresses to be connected on the RHEL client.
- S indicates the number of dynamic front-end service IP addresses returned by OceanStor 9000.
- 3. When upper-layer services read and write files concurrently, the system selects the connection with the minimum workload for data transmission, achieving link load balancing, as shown in **Figure 4-21**.

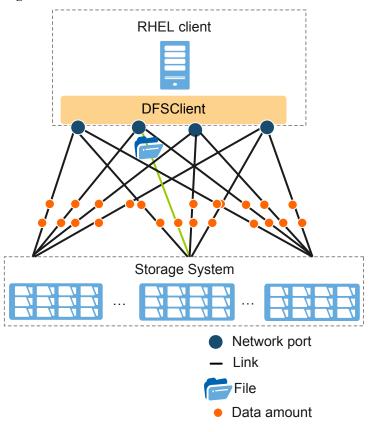


Figure 4-21 Data transmission mode on the RHEL client

Compared with the single-connection mode, the multi-connection mode increases the bandwidth provided for the RHEL client accessing OceanStor 9000.

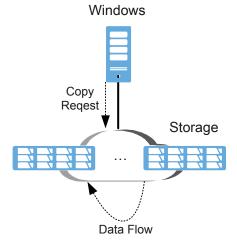
When there are a great amount of data reading and writing, DFSClient can balance data amount on links and accelerate data transmission.

## 4.4.3 Rapid internal file replication

If a client employs SMB 2.0/2.1/3.0, the OceanStor 9000 can rapidly implement internal file replication without involving the client.

In the OceanStor 9000, this function is enabled. Whether the function takes effect is automatically determined based on the client environment (when sending FSCTL\_SRV\_COPYCHUNK requests is in effect). Typically, the function takes effect when the copy command is executed in the command-line interface of Windows 7 or copy and paste operations are performed in the resource manager of Windows 8 or Windows Server 2012. In these scenarios, the network port on the client does not generate file data traffic. Figure 4-22 shows a schematic diagram of rapid internal file replication.

Figure 4-22 Schematic diagram of rapid internal file replication



## 5 Typical Applications

## **About This Chapter**

This chapter describes typical application scenarios and typical configurations of the OceanStor 9000.

#### 5.1 Application Scenarios

Typical application scenarios of the OceanStor 9000 include media asset management, high-performance computing, video surveillance, backup and archive, and individual and enterprise web disks.

#### 5.2 Typical System Configuration

This section describes the typical system configuration of the OceanStor 9000.

## **5.1 Application Scenarios**

Typical application scenarios of the OceanStor 9000 include media asset management, high-performance computing, video surveillance, backup and archive, and individual and enterprise web disks.

#### **MAM**

A MAM system is the core of a television program production system and is used to upload, download, catalog, retrieve, transcode, archive, store, invoke, and manage media materials. Its characteristics and storage requirements are as follows:

- Media materials have a high bit rate and are large in size. Large-capacity and highly-scalable storage systems are required.
- Ingesting, editing, and composing of audio and video require stable and low-latency storage systems.
- Concurrent editing by multiple users requires a storage system to deliver reliable and easy-to-use data sharing.
- Small files are frequently processed in video rendering and special effect processing. This requires storage systems to deliver high I/O performance.

**Figure 5-1** shows the MAM system application scenario of the OceanStor 9000. The OceanStor 9000 provides the following benefits:

- High-speed data sharing that improves program production efficiency
- On-demand and hitless in-service capacity expansion
- Dynamic storage tiering that meets high performance requirements while reducing device purchase costs
- Energy saving that reduces operational expenditure (OPEX)

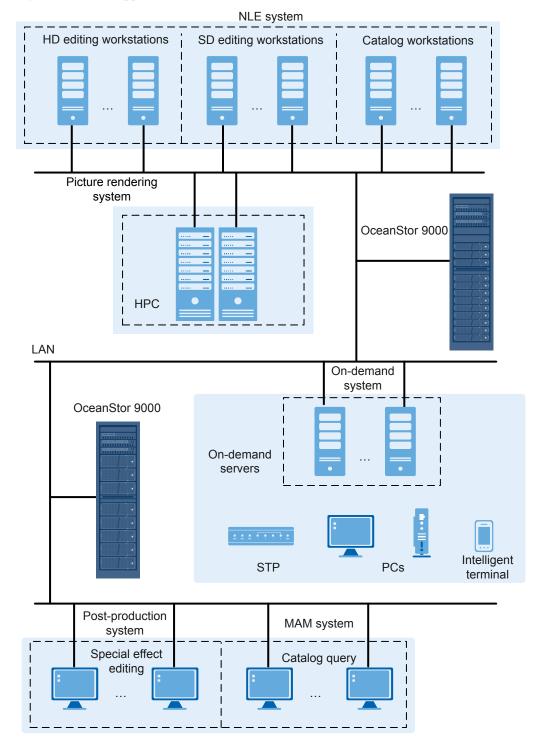


Figure 5-1 MAM application scenario

The OceanStor 9000 is a centralized storage solution for ingesting, producing, and distributing media materials. It meets the media assets industry's demanding requirements for storage capacity, performance, stability, and scalability.

#### **HPC**

HPC is widely applied in sectors such as life science, research & education, energy and geology, and space exploration. HPC involves massive data computing and has demanding requirements on storage performance, capacity, and stability.

**Figure 5-2** shows an HPC application scenario of the OceanStor 9000. The OceanStor 9000 provides the following benefits:

- High performance: Excellent I/O performance meets demanding service requirements.
- High availability: Nodes are clustered to eliminate single points of failure, ensuring data security and reliability.
- Flexible scalability: A distributed architecture with the fully symmetrical node design enables on-demand expansion of capacity and performance.
- Ease of maintenance: A unified namespace is provided to reduce maintenance costs.

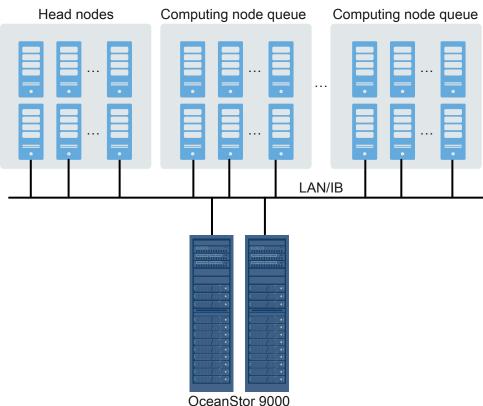


Figure 5-2 HPC application scenario

#### Video Surveillance

As cities are constantly extended, city security is facing a variety of threats posed by social security, public health, natural disasters, and accidents. The video surveillance system plays a crucial role in building a safe city.

The OceanStor 9000 provides a client software application that supports unified view display. The software can consolidate storage clusters in different areas into a storage resource pool. When unexpected events occur, cross-area video surveillance data can be easily checked.

Video surveillance management platform

Figure 5-3 shows applications of the OceanStor 9000 in the video surveillance scenario.

Figure 5-3 Video surveillance scenario

**Backup and Archiving** 

The OceanStor 9000 provides a variety of data protection functions, including snapshot, InfoReplicator, and NDMP backup, to help multiple data centers build backup, archiving, data protection, and disaster recovery capabilities.

OceanStor 9000

Figure 5-4 shows the use of the OceanStor 9000 in the backup and archiving scenario.

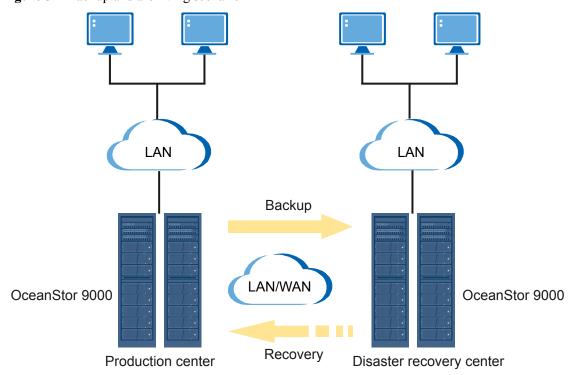


Figure 5-4 Backup and archiving scenario

## Individual and Enterprise Web Disks

The system provides the object storage service through RESTful APIs. The object storage service helps construct Internet applications including web hosting and online videos, albums, and backups. Object storage service (compatible with Amazon S3 APIs) provides the RESTful API access interface for web applications to access object storage service (compatible with Amazon S3 APIs).

**Figure 5-5** shows typical application scenarios of OceanStor 9000 in individual and enterprise web disks.

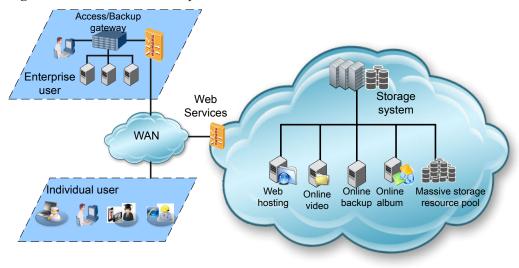


Figure 5-5 Individual and Enterprise Web Disks

## 5.2 Typical System Configuration

This section describes the typical system configuration of the OceanStor 9000.

The number of nodes in a fully configured cabinet depends on power supply capability of the cabinet. The following shows typical configurations of a cabinet when a power supply of 220V/63A is adopted.

- A fully configured P12/P25/P12E/P25E cabinet can house 12 storage nodes, two 10GE switches, one GE switch, and one KVM, as shown in **Figure 5-6** (where the storage nodes are P25.)
- A fully configured P36/C36/P36E/C36E cabinet can house eight storage nodes, two 10GE switches, one GE switch, and one KVM, as shown in **Figure 5-7**.
- A fully configured C72 cabinet can house five storage nodes, two 10GE switches, and one GE switch, as shown in Figure 5-8. A C72 cabinet equipped with a Huawei power distribution unit (PDU) whose ID is 02120643 can house six storage nodes, two 10GE switches, and one GE switch.

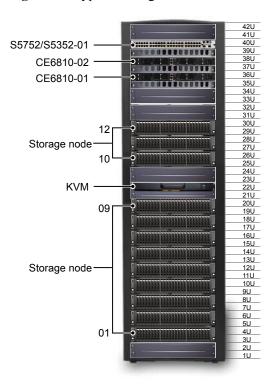
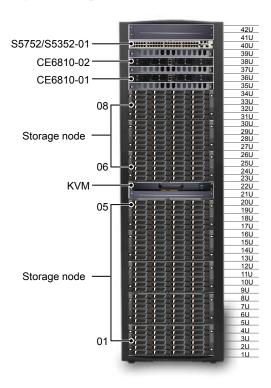


Figure 5-6 Typical configuration of a P12/P25/P12E/P25E cabinet

Figure 5-7 Typical configuration of a P36/C36/P36E/C36E cabinet



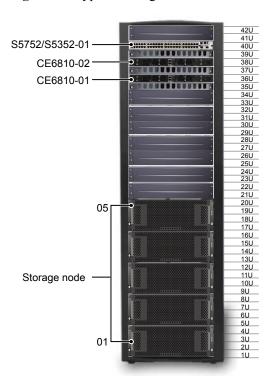


Figure 5-8 Typical configuration of a C72 cabinet

# 6 Operation and Maintenance

The OceanStor 9000 provides easy to use centralized management functionality and can run remote operations, monitoring, and maintenance through a unified network management.

## **Cluster System Management**

The OceanStor 9000 allows you to perform the following cluster management operations using the unified network management software DeviceManager:

- Monitor the topology and running status of the cluster system.
- Monitor node operating status.
- Manage system time, for example, configure the time zone and NTP settings.
- Monitor system performance, for example, view key indicators such as capacity usage and throughput.
- Maintains and manages devices. For example, you can enable, disable, freeze, unfreeze, stop, start, and restart devices.
- Maintains and manages the overall system. For example, you can start, restart, and shut down the system. This eliminates the need to perform these operations on each device and helps enhance system manageability.

The eSight ICT management system can help centrally manage multiple sets of OceanStor 9000 clustered systems and other storage systems.

## File System Management

The OceanStor 9000 enables you to perform the following file system management operations using DeviceManager:

- Configure an authentication mode.
- Configure service sharing.
- Manage quotas.
- Configure a redundancy ratio.

The OceanStor 9000 provides a resource manager that allows you to create, delete, and browse directories and configure directory properties with ease.

## Object Storage Service (Compatible with Amazon S3 APIs) Management

The OceanStor 9000 enables you to perform the following object storage service (compatible with Amazon S3 APIs) management using DeviceManager:

Access setting

An object storage service (compatible with Amazon S3 APIs) domain name can be set for a cluster.

Account management

An administrator can manage accounts of the object storage service (compatible with Amazon S3 APIs) using DeviceManager.

Charging service management

Charging service management helps safeguard charging accounts. You can determine whether to enable the charging service.

#### Object Storage Service (Compatible with OpenStack Swift APIs) Management

The OceanStor 9000 enables you to perform the following object storage service (compatible with OpenStack Swift APIs) management using DeviceManager:

Access setting

An object storage service (compatible with OpenStack Swift APIs) domain name can be set for a cluster.

Account management

An administrator can manage accounts of the object storage service (compatible with OpenStack Swift APIs) using DeviceManager.

Cross-domain access management

You can configure the cross-region resource sharing (CORS) and cross-domain policy file

• TempURL management

You can configure TempURL parameters to manage a TempURL, including Deleted Message Headers, Reserved Message Headers, Deleted Response Headers, Reserved Response Headers, and Use Method.

## **Software Management**

The OceanStor 9000 provides a software deployment tool that can automate software installation, software upgrade, and capacity expansion for the OceanStor 9000.

The OceanStor 9000 supports both online and offline upgrade.

- An online upgrade does not interrupt services. During an upgrade, the OceanStor 9000 can provide services to you. Perform the upgrade in off-peak hours (when running services are less than 10% of the maximum service volume supported by the system). Otherwise, the upgrade may fail.
- An offline upgrade interrupts services. During an offline upgrade, the OceanStor 9000 cannot provide services.

Based on the time, limitations, and reliability of an upgrade, the offline upgrade mode is advised to be used.

## **User Management**

The OceanStor 9000 supports user management. You can create, delete, lock, and unlock users to ensure secure user access.

## **Fault Management**

The OceanStor 9000 supports fault management. Once an exception occurs, the OceanStor 9000 reports an alarm in real time and provides handling suggestions. It can also inform the maintenance personnel of faults by SMS or email. You can configure alarm thresholds, export alarms, clear alarms, and dump alarms.

## **7**Technical Specifications

## **About This Chapter**

Technical specifications of the OceanStor 9000 include performance specifications, capacity specifications, overall system specifications, reliability specifications, and environmental specifications.

- 7.1 Performance and Capacity Specifications
- 7.2 Overall System Hardware Specifications
- 7.3 Reliability Specifications
- 7.4 Environmental Specifications

## 7.1 Performance and Capacity Specifications

, , Table 7-1, and Table 7-2 list the performance specifications of the OceanStor 9000.

Table 7-1 Performance indicators of the OceanStor 9000 file system

Item	Indicator	
OPS	A maximum of 25,000 (CIFS)/20,000 (NFS) per node	
CIFS stable read	A maximum of 800 MB/s bandwidth per node	
CIFS stable write	A maximum of 800 MB/s bandwidth per node	
CIFS ultimate read	A maximum of 1280 MB/s bandwidth per node	
CIFS ultimate write	A maximum of 1300 MB/s bandwidth per node	
NFS ultimate read	A maximum of 1600 MB/s bandwidth per node	
NFS ultimate write	A maximum of 1100 MB/s bandwidth per node	

Table 7-2 Performance indicators of the OceanStor 9000 object storage

Item	Indicator
Large object read of object storage	A maximum of 900 MB/s bandwidth per node
Large object write of object storage	A maximum of 800 MB/s bandwidth per node
Small object read of object storage	A maximum of 4000 transactions per second (TPS) per node
Small object write of object storage	A maximum of 500 TPS per node

**Table 7-3** lists the capacity specifications of the OceanStor 9000.

Table 7-3 Capacity specifications of the OceanStor 9000

Item	Description
Maximum number of nodes in a cluster	288
Maximum capacity of a file system	100 PB
Maximum size of a file	240 TB

Item	Description
Maximum number of files or sub- directories in a directory	10 million
Maximum number of files (including directories)	10 billion
Maximum number of object storage buckets supported by a cluster	100 million
Maximum number of objects supported by a cluster	10 billion
Maximum number of accounts supported by a cluster	C72 node: 100,000 P12/P25/P36/C36 node: 1,000,000 P12E/P25E/P36E/C36E node: 1,000,000

## 7.2 Overall System Hardware Specifications

**Table 7-4**, **Table 7-5**, **Table 7-6**, **Table 7-7**, **Table 7-8**, and **Table 7-9** list the overall system hardware specifications of the OceanStor 9000 that adopts the typical configuration.

Table 7-4 Overall system hardware specifications of a P12E or P25E node cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	12
Power supply	AC input voltage	90 V AC to 264 V AC
	AC input current	63 A/32 A
Weight	Weight of an AC cabinet	593.5 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	> 600 kg/m <sup>2</sup>

Table 7-5 Overall system hardware specifications of a P36E or C36E cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	8
Power supply	AC input voltage	90 V AC to 264 V AC
	AC input current	63 A/32 A

Category	Item	Value
Weight	Weight of an AC cabinet	689.5 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	$> 600 \text{ kg/m}^2$

Table 7-6 Hardware specifications of the P12 cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	12
Power supply	AC input voltage	100 V to 127 V or 200 V to 240 V
	AC input current	63 A/32 A
Weight	Weight of an AC cabinet	617.5 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	> 600 kg/m <sup>2</sup>

Table 7-7 Hardware specifications of the P25 cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	12
Power supply	AC input voltage	100 V to 127 V or 200 V to 240 V
	AC input current	63 A/32 A
Weight	Weight of an AC cabinet	653.5 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	> 600 kg/m <sup>2</sup>

Table 7-8 Hardware specifications of the P36/C36 cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	8

Category	Item	Value
Power supply	AC input voltage	200 V to 240 V
	AC input current	63 A/32 A
Weight	Weight of an AC cabinet	725.5 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	$>600 \text{ kg/m}^2$

Table 7-9 Hardware specifications of the C72 cabinet

Category	Item	Value
Capacity	Number of storage nodes in a cabinet	6
Power supply	AC input voltage	200 V to 240 V
	AC input current	63 A/32 A
Weight	Weight of an AC cabinet	920 kg
Load bearing capacity of the equipment room	Average load bearing capacity per square meter with fully-configured cabinets	> 600 kg/m <sup>2</sup>

## 7.3 Reliability Specifications

**Table 7-10** lists the reliability specifications of the OceanStor 9000.

Table 7-10 Reliability specifications

Item	Value
Service reliability	The file system supports +1, +2, +3, +4, +2:1, and +3:1 redundancy ratios, which can be set for any empty directory. Object storage supports +1, +2, +3, +2:1, and +3:1 redundancy ratios which can be set for an object storage account.
Data reconstruction speed	1 TB/h

## 7.4 Environmental Specifications

**Table 7-11** lists the environmental specifications of the OceanStor 9000.

Table 7-11 Environmental specifications of the OceanStor 9000

Parameter	Value
Operating temperature <sup>a</sup>	5°C to 35°C (41°F to 95°F)
Storage temperature	-40°C to +70°C (-40°F to +158°F)
Transportation temperature	-40°C to +70°C (-40°F to +158°F)
Temperature gradient	10°C/h (18°F/h)
Operating humidity	20% RH to 80% RH
Storage humidity	5% RH to 95% RH
Transportation humidity	5% RH to 95% RH
Particle contaminants	ISO <sup>b</sup> 14664-1 Class8
Corrosive airborne contaminants	• The corrosion rate of the copper test wafer is lower than 300 Å <sup>b</sup> /month per ANSI <sup>c</sup> /ISA <sup>d</sup> -71.04 severity level G1.
	• The corrosion rate of the silver test wafer is lower than 200 Å <sup>b</sup> /month per ANSI <sup>c</sup> /ISA <sup>d</sup> -71.04 severity level G1.
	<ul> <li>You are advised to ask a professional organization to monitor the corrosive airborne contaminants in the equipment room.</li> </ul>
	<ul> <li>The following is an example of initial evaluation for environmental corrosion risks.</li> </ul>
	<ul> <li>Whether the equipment room is near any sulfurous gas emission sources, for example, porcelain factories, rubber plants, tire factories, chemical factories, sewage plants, power stations, paper mills, smelters, automobile factories, coal mines, electroplating factories, food factories, and tanneries? If yes, environmental corrosion risks may exist.</li> </ul>
	<ul> <li>Whether the equipment room is near the sea, saline, sewer outlets, sewage treatment tanks, and industrial/heating boilers? If yes, environmental corrosion risks may exist.</li> </ul>
	<ul> <li>Whether the equipment room was decorated in the latest six months? If yes, environmental corrosion risks may exist.</li> </ul>
	<ul> <li>Whether batteries are stocked in the equipment room? Whether the storage battery leaks? If yes, environmental corrosion risks may exist.</li> </ul>
	<ul> <li>Whether the equipment room is closed tightly?</li> <li>Whether the room windows and doors are always closed? If no, environmental corrosion risks may exist.</li> </ul>

Parameter	Value
Operating altitude	-30.5 m to +3000 m (-100.06 ft to +9842.4 ft)
Storage altitude	-30.5 m to +3000 m (-100.06 ft to +9842.4 ft)
Transportation altitude	-30.5 m to +3000 m (-100.06 ft to +9842.4 ft)
Noise <sup>f</sup>	< 72 dBA

- a: The operating temperature ranges from  $5^{\circ}$ C to  $35^{\circ}$ C ( $41^{\circ}$ F to  $95^{\circ}$ F) when the system runs at an altitude lower than 1800 m (5905.44 ft), and the upper threshold for operating temperature decreases by  $1^{\circ}$ C ( $33.8^{\circ}$ F) every 300 m (984.24 ft) altitude increment when the system runs at an altitude higher than 1800 m (5905.44 ft).
- b: International Organization for Standardization
- c: a unit of length equal to 1/10000000000 (one ten billionth) of a meter
- d: American National Standards Institute
- e: Instrument Society of America
- f: the maximum noise generated by a storage node when the ambient temperature is lower than 25°C (77°F).

## 8 Standards Compliance

The OceanStor 9000 complies with protocol standards, safety and electromagnetic compatibility (EMC) standards, industry standards, and authentication standards.

## **Protocol Standards**

Table 8-1 lists the protocol standards with which the OceanStor 9000 complies.

Table 8-1 Protocol standards met by the OceanStor 9000

Name	Standards No.
IPMI2.0	Intelligent Platform Management Interface Specification Second Generation v2.0, Document Revision 1.0, February 12, 2004
SMBIOS	System Management BIOS (SMBIOS) Reference Specification, Version 2.5, Status: Final Standard, September 5, 2006
SAS2.1	Information technology — Serial Attached SCSI - 2
ACPI	Advanced Configuration and Power Interface Specification, Revision 3.0, September 2, 2004
IP	RFC0791: Internet Protocol
Swift	OpenStack Swift 2.5.0, April 30 2015
S3	Amazon Simple Storage Service  NOTE  For details about compatibility, see the OceanStor9000 VX00R00XCXX Object Storage Service (Compatible with Amazon S3 APIs) API Reference.
HDFS	Hadoop Distributed File System Version 2.5.0, November 13 2014

## Safety and EMC Standards

Table 8-2 lists the safety and EMC standards with which the OceanStor 9000 complies.

Table 8-2 Safety and EMC standards met by the OceanStor 9000

Name	Standards No.
IT equipment safety standard	GB4943-2001
ICE standard	IEC 60950-1
North America compliance and safety standards	UL 60950-1
US EMC standard	FCC, 47 CFR Part 15, Subpart B
European safety standard	EN 60950-1
European EMC standard	EN 55024: 1998+A1+A2

## **Industry Standards**

Table 8-3 lists the industry standards with which the OceanStor 9000 complies.

Table 8-3 Industry standards met by the OceanStor 9000

Name	Standards No.
Ethernet-based standard	IEEE 802.3
Standard for fast Ethernet	IEEE 802.3u
Standard for GB Ethernet	IEEE 802.3z
IEEE standard test interface and boundary scan structure	IEEE 1149.1-2001
Fault mode effect analysis (FMEA)	IEC 812
Standard for reliability, maintainability, and availability	IEC 863
Environmental protection	REACH/ROHS/WEEE
Clean room and related controlled environments	ISO 14664-1 Class8
Airborne contaminants and environment standards	ANSI/ISA-71.04-1985 G1 gas corrosion level

## Certifications

Table 8-4 lists the certifications that have been granted to the OceanStor 9000.

Table 8-4 Certifications granted to the OceanStor 9000

Certificate	Description
СВ	The CB Scheme is managed by the Worldwide System for Conformity Testing and Certification of the Electrotechnical Equipment and Components (IECEE) System. The IECEE member agencies in each country test the safety of electrical equipment based on IEC standards. The test results (the CB test report and the CB test certificate) are recognized by the IECEE member bodies. The CB Scheme aims at reducing trade barriers caused by different approval or certification standards in different countries.
CCC	China Compulsory Certification (CCC) enforces compulsory certification on products relating to human health, animal life and health, environmental protection, and public safety.
CE	Conformite Europeenne (CE) refers to the certification required for products to be sold in Europe. The products marked CE comply with the electromagnetic compatibility (EMC) regulations (2004/108/EEC) and low voltage regulations (2006/95/EEC) released by the European Commission.
C-Tick	C-Tick is an identification trademark registered to the Australian Communications and Media Authority (ACMA). Currently, the C-Tick is only applicable to products with electromagnetic interference and became effective from on January 1, 1999.
FCC	According to Chapter 15 in Federal Communications Commission (FCC) Rules and Regulations, the device conforms to the standard for level A digital device.
GOST	Russia national standard certification requiring CE or CB.
IC	Industry Canada (IC) sets up the test standards for analog and digital terminal devices and specifies corresponding EMC certificates that all import electronic products must obtain.
REACH	REACH is the European Community Regulation on chemicals and their safe use (EC 1907/2006). It deals with the Registration, Evaluation, Authorization and Restriction of Chemical substances.
RoHS	Restriction of the Use of Certain Hazardous Substances (RoHS) is published by EU in 2003. It applies to the management of environmental impact created by electric and electronic products when produced and scrapped. RoHS stipulates the maximum volume of hazardous substances arising from production.
SASO	Saudi Arabian Standards Organization (SASO) is a partner in maintaining the safety and public health protecting the consumer and environment and enhancing the competitive of national products through the application of the best international practices in different standardization fields in accordance with Islamic law. The International Conformity Certification Programme (ICCP) was implemented by MoCI of the Kingdom of Saudi Arabia to ensure the compliance of imported products, with applicable Saudi Arabian technical regulations.

Certificate	Description
SONCAP	According to the Standards Organization of Nigeria Conformity Assessment Programme (SONCAP), all goods imported into Nigeria conform to Nigeria Industrial Standards (NIS) or equivalent acceptable international standards.
UL	Underwriters Laboratories (UL) is a non-profit agency engaged in product safety testing in the USA.
VCCI	Voluntary Control Council for Interference (VCCI) is an identification trademark registered to Japan.
WEEE	Electrical and electronic products sold in the EU market must comply with the EU Directive on Waste of Electric and Electronic Equipment (WEEE) and possesses their logo/symbol.
Country Commodity Inspection Certification	The certification is applicable to Saudi Arabia, Nigeria, Tajikistan, Uganda, Kuwait, Algeria, Botswana, Qatar, and Egypt.
NOM	The Norma Oficial Mexicana (NOM) is the name of each of a series of official compulsory standards and regulations for diverse activities in Mexico. According to the List of Mandatory Mexican Standards Enforced at the Border of NAFTA Facts Document 9012, general electrical and electronic equipment (EEE) must obtain the NOM certificate except that the equipment is proved to be highly specialized equipment (HSE) or the input power of the equipment does not exceed 24 V (rms). Three organizations can grant such a certificate, UL, NYCE, and ANCE.
CCEL	To obtain Certificate of China Environmental Labeling (CCEL), the product must meet the requirements of energy saving, security, and EMC during the whole process through product design, product manufacture, production waste disposal, reclamation, and recycling (from cradle to cradle).
Air and Sea Transport Certification	-



A

AC alternating current
AD active directory

В

BIOS See basic input/output system.

back-end storage

network

In system networking, the network that interconnects storage servers for data

interaction is called the back-end storage network.

basic input/output system (BIOS)

Firmware stored on the computer motherboard that contains basic input/output control programs, power-on self test (POST) programs, bootstraps, and system setting information. The BIOS provides hardware setting and control functions for the

computer.

C

**CAPEX** capital expenditure

CIFS See common Internet file system.

CLI command-line interface

Capacity node C nodes apply to Capacity and high-throughput.

See central processing unit.

central processing unit

(CPU)

**CPU** 

The computational and control unit of a computer. The CPU is the device that interprets and executes instructions. The CPU has the ability to fetch, decode, and

execute instructions and to transfer information to and from other resources over the

computer's main data-transfer path, the bus.

**cluster** Minimum management unit in the capacity management system. It is a collection of

applications that support automatic load balancing or near load balancing through manual adjustment. It supports cloud-based scale-in and scale-out and non-cloud—

based capacity expansion and reduction.

common Internet file system (CIFS)

A network file system access protocol primarily used by Windows clients to communicate file access requests to Windows servers. CIFS was originally called Server Message Block (SMB). Today,other implementations of the CIFS protocol allow other clients and servers to use it for intercommunication and interoperation with Microsoft operating systems.

D

**DAR** dynamic address resolution

**DC** direct current

**domain** A logical subscriber group based on which the subscriber rights are controlled.

E

**EMC** See electromagnetic compatibility.

**ESD** electrostatic discharge

**Ethernet** A LAN technology that uses the carrier sense multiple access with collision detection

(CSMA/CD) media access control method. The Ethernet network is highly reliable and easy to maintain. The speed of an Ethernet interface can be 10 Mbit/s, 100 Mbit/s,

1000 Mbit/s, or 10,000 Mbit/s.

electromagnetic compatibility (EMC)

A condition which prevails when telecommunications equipment is performing its individually designed function in a common electromagnetic environment without causing or suffering unacceptable degradation due to unintentional electromagnetic

interference to or from other equipment in the same environment.

F

**FMEA** failure mode and effects analysis

**FTP** File Transfer Protocol

front-end service

network

In system networking, the network that connects to the customer's service network for

data interaction is called the front-end service network.

G

**GE** Gigabit Ethernet

GUI graphical user interface

**global cache** A cache management function that consolidates caches of all nodes in the cluster into

a global cache. It makes data cached on any node directly accessible to the other

nodes.

global namespace A type of namespace. Users can access all space of a storage domain through such a

namespace.

I

I/O input/output

IB See InfiniBand.

**IEC** International Electrotechnical Commission

IEEE See Institute of Electrical and Electronics Engineers.

IP address A 32-bit (4-byte) binary number that uniquely identifies a host connected to the

> Internet. An IP address is expressed in dotted decimal notation, consisting of the decimal values of its 4 bytes, separated with periods; for example, 127.0.0.1. The first three bytes of the IP address identify the network to which the host is connected, and

the last byte identifies the host itself.

**IPMI** See Intelligent Platform Management Interface.

ISA See instruction set architecture.

ISO International Organization for Standardization

InfiniBand (IB) A switched fabric I/O technology that uses a central InfiniBand switch to set up single

> connections between remote storage devices, network devices, and servers and to redirect traffic. With a tight structure, InfiniBand greatly improves system performance, reliability, and efficiency and alleviates data congestion between

hardware devices.

InfoAllocator A piece of quota management software that poses space usage or file quantity limits

on directories or users

**InfoEqualizer** A piece of load balancing software for client links. It uses unified domain name

access, load balancing, and failover technologies to simplify the process of setting up

a connection with a client, improving system performance and reliability.

InfoTier A piece of automatic storage tiering software that enables policy-driven automatic file

migration among multiple levels of storage nodes.

**Institute of Electrical** 

and Electronics **Engineers (IEEE)**  A professional association of electrical and electronics engineers based in the United States, but with membership from numerous other countries. The IEEE focuses on electrical, electronics, and computer engineering, and produces many important

technology standards.

**Intelligent Platform** 

(IPMI)

A specification, developed by Dell, HP, Intel and NEC, for the purpose of improving Management Interface serviceability of servers. The IPMI provides servers with device management, sensor and event management, user management, fan box and power supply management,

and remote maintenance.

instruction set architecture (ISA) A part of computer architecture relating to program design that comprises local data types, instructions, registers, address modes, memory architectures, interrupts,

exception handles, and external I/Os.

K

**KVM** See keyboard, video, and mouse.

keyboard, video, and mouse (KVM)

A hardware device installed in the integrated configuration cabinet. KVM serves as the input and output device for the components inside the cabinet. It consists of a

screen, a keyboard, and a mouse.

L

LDAP See Lightweight Directory Access Protocol. OceanStor 9000 Product Description

A Glossary

Lightweight Directory Access Protocol

(LDAP)

A network protocol based on TCP/IP, which allows access to a directory system agent (DSA). It involves some reduced functionality from X.500 Directory Access Protocol

(DAP) specifications.

**load balancing** The distribution of activity across two or more servers or components in order to

avoid overloading any one with too many requests or too much traffic.

M

Mini-Capacity node

M nodes apply to Mini-Capacity applications.

metadata

Data about data. Metadata defines data properties and is used to specify data storage locations and historical data, retrieve resource data, and record information.

N

NAS network attached storage

NDMP network data management protocol

**NE Interface Server** 

(NIS)

A call center platform component that manages registration information of other

components.

NFS See Network File System.

NIS See NE Interface Server.

NTP Network Time Protocol

**NVDIMM** See non-volatile dual in-line memory module.

**Network File System** 

(NFS)

A method of implementing file sharing between UNIX systems. Through the NFS and the network, application programs on the client can save and obtain data on the disk of a server. The NFS was first developed by Sun in 1984, and designed for the sharing of

materials among different PCs and operating systems through networks.

non-volatile dual inline memory module (NVDIMM) A type of dual in-line memory module. Such a memory module can save data within it

upon a system power failure.

O

**OPEX** operating expense

**OPS** See operation per second.

operation per second

(OPS)

The number of operations within a second supported by a storage device. It is a critical

storage performance indicator.

P

PCI See peripheral component interconnect.

**Performance node** P nodes apply to OPS-intensive applications.

peripheral component interconnect (PCI)

A standard designed for the bus connecting the computer main board to peripheral devices. The PCI1.0 standard was released by Intel in 1992 and related standards have been released by PCI-SIG since 1993. Peripheral component interconnect (PCI) delivers I/O functionality for computers ranging from servers to workstations, PCs, laptop PCs and mobile devices.

Q

**quota** A limit that restricts the amount of a resource, such as disk space, that a user, group or

directory structure may consume.

R

**RAID** redundant array of independent disks

**RoHS** restriction of the use of certain hazardous substances

 $\mathbf{S}$ 

SAS serial attached SCSI

SATA Serial Advanced Technology Attachment

SCSI Small Computer System Interface

**SFP** small form-factor pluggable

**SMS** short message service

**snapshot** A data backup technology through which a fully usable copy of a data object can be

quickly generated. The duplicate contains the image of the data object at a point in

time.

T

threshold A limitation on an amount, scale, or level. Changes will occur when a threshold is

reached.

U

UL Underwriters Laboratory
USB See Universal Serial Bus.

**Universal Serial Bus** 

(USB)

A serial bus standard to interface devices. It was designed for computers such as PCs

and the Apple Macintosh, but its popularity has prompted it to also become

commonplace on video game consoles and PDAs.

W

**WORM** See write once read many.

write once read many

(WORM)

A type of data protection method. After being written, the data can be read but cannot

be altered.