

FusionStorage

V100R006C20

Product Description

Issue 02

Date 2018-06-30

Copyright © Huawei Technologies Co., Ltd. 2018. 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



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

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: <http://e.huawei.com>

About This Document

This document describes the positioning, functions and features, hardware structure, software structure, networking, application scenarios, typical configuration, technical specifications, standards compliance and certifications of the FusionStorage.





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
 DANGER	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
 WARNING	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
 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 02 (2018-06-30)

This issue is the second official release.

Issue 01 (2018-04-30)

This issue is the first official release.

Contents

About This Document.....	ii
1 Product Positioning.....	1
2 Software Structure of FusionStorage.....	3
3 Product Networking Architecture.....	5
3.1 File Storage Service Network Architecture.....	5
3.2 Object Storage Service Network Structure.....	6
3.3 Block Storage Service Network Architecture.....	9
4 File Storage Service.....	11
4.1 Service Function of the File Storage Service.....	11
4.1.1 CIFS/NFS/FTP Sharing.....	11
4.1.2 CIFS-NFS Cross-Protocol Access Sharing.....	12
4.1.3 HDFS Interface.....	14
4.1.4 NDMP Backup.....	15
4.1.5 High Data Reliability.....	17
4.1.6 Permission Control.....	18
4.1.7 Windows Hard Link.....	19
4.1.8 Unified Namespace.....	20
4.1.9 Technical Specifications of File Storage Service.....	21
4.2 Features of the File Storage Service.....	21
4.2.1 Tiering.....	21
4.2.2 Quota.....	23
4.2.3 Load Balance.....	24
4.2.4 WORM.....	25
4.2.5 Snapshot.....	26
4.2.6 Replication.....	26
4.2.7 Anti-Virus.....	27
4.2.8 InfoMigrator.....	29
4.2.9 InfoRevive.....	30
4.2.10 InfoTurbo.....	30
4.2.10.1 Intelligent Prefetch.....	31
4.2.10.2 SMB3 Multichannel.....	32
4.2.10.3 NFS Protocol Enhancement.....	33

4.2.11 Rapid internal file replication.....	37
5 Object Storage Service.....	39
5.1 Object Storage Service (Compatible with Amazon S3 APIs).....	39
5.2 Object Storage Service (Compatible with OpenStack Swift APIs).....	41
5.3 Technical Specifications of Object Storage Service.....	42
6 Block Storage Service.....	43
6.1 Multi-Resource Pool.....	43
6.2 Distributed RAID.....	44
6.3 Snapshot.....	45
6.4 Thin Provisioning.....	46
6.5 Linked Cloning.....	47
6.6 High Data Reliability.....	48
6.7 Security Protection Mechanisms.....	53
7 Application Scenarios.....	58
8 Technical Specifications.....	63
8.1 Reliability Specifications.....	63
9 Recommended Hardware Configurations.....	64
9.1 Hardware Components.....	64
9.2 Storage Nodes.....	65
9.2.1 Nodes Providing the File Storage Service.....	65
9.2.1.1 RH2288 V3 12-slot Nodes.....	65
9.2.1.2 5288 V3 36-slot Nodes.....	70
9.2.2 Nodes Providing the Object Storage Service.....	76
9.2.2.1 RH2288 V3 12-slot Nodes.....	77
9.2.2.2 5288 V3 36-slot Nodes.....	82
9.2.3 Nodes Providing the Block Storage Service.....	88
9.2.3.1 2288H V5 12-slot Nodes.....	89
9.2.3.2 2288H V5 16-slot Nodes.....	95
9.2.3.3 Taishan 2280 12-slot Nodes.....	101
9.2.3.4 RH2288 V3 12-slot Nodes.....	106
9.2.3.5 5288 V3 36-slot Nodes.....	111
9.3 Switches.....	117
9.3.1 CE6851-48S6Q-HI.....	117
9.3.2 CE6855-48S6Q-HI.....	120
9.3.3 CE5855-48T4S2Q-EI.....	122
9.4 Standard IT Cabinets.....	125
9.5 Optional Hardware.....	127
9.6 Recommended Cabinet Configurations.....	129
9.7 Environmental Specifications.....	137
9.8 Standards Compliance.....	138

1 Product Positioning

This chapter describes the product positioning and features of FusionStorage.

FusionStorage is a highly scalable scale-out storage product. By using storage system software, FusionStorage consolidates all storage resources of x86 servers into a fully distributed storage pool. FusionStorage enables a single storage system to deliver any type of storage service for upper-layer applications, including block storage service, file storage service, or object storage service. This addresses the storage requirements of both structured and unstructured data while exempting you from planning storage service types in advance and purchasing diverse storage devices dedicated to different storage services.

FusionStorage supports a wide array of enterprise-class data service features, such as snapshot, thin provisioning, remote replication, and multi-tenant, ensuring flexible and efficient data access during constant service changes.

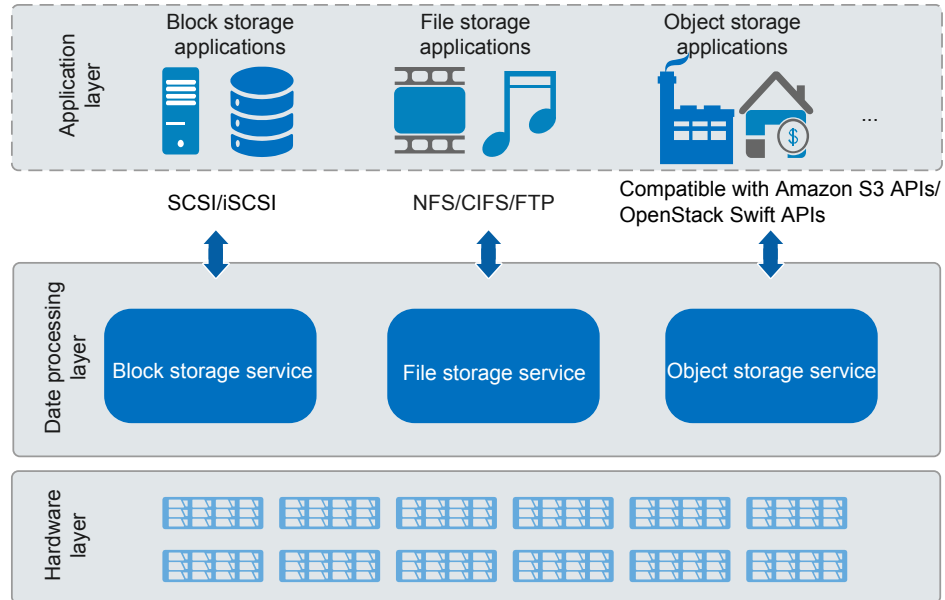
FusionStorage provides standard application programming interfaces (APIs) open to upper-layer applications and clients. With these APIs, FusionStorage can easily interconnect with OpenStack cloud infrastructure and Hadoop Big Data ecosystem. APIs provided by FusionStorage can be grouped into the following categories:

- **SCSI- and iSCSI-based block storage API:** With this type of API, FusionStorage can provide the distributed block storage service by virtualizing storage resources of local servers into SAN storage. Utilizing its high performance and extensive scalability, FusionStorage meets SAN storage requirements imposed by diverse database applications and virtualization platforms, such as cloud resource pooling, desktop cloud, and DevCloud.
- **NFS-, CIFS-, FTP-, and HDFS-based file storage API:** With this type of API, FusionStorage can provide the distributed file storage service. Benefiting from its excellent performance, strong scale-out capability, and super-large single file system, FusionStorage enables resource sharing of unstructured data and is applied in multi-service scenarios, such as massive video/audio storage and Big Data applications.
- **Object storage API, compatible with Amazon S3 or OpenStack Swift:** With this type of API, FusionStorage can provide the distributed object storage service and interconnect with mainstream cloud computing ecosystems, meeting the requirements of cloud backup, cloud archiving, and private cloud storage service operation.

FusionStorage object storage service provides OpenStack Swift APIs. We hope that you can develop your applications or modify your existing applications based on the APIs. To meet the requirements on running existing Amazon S3 applications, FusionStorage also provide a compatibility solution for you. However, due to differences in functions of platforms, we provide only some important compatibility functions.

Figure 1-1 shows the system architecture of FusionStorage.

Figure 1-1 System architecture of FusionStorage



2 Software Structure of FusionStorage

This chapter describes the software structure of FusionStorage.

Figure 2-1 shows the software structure of FusionStorage.

Figure 2-1 Software structure of FusionStorage

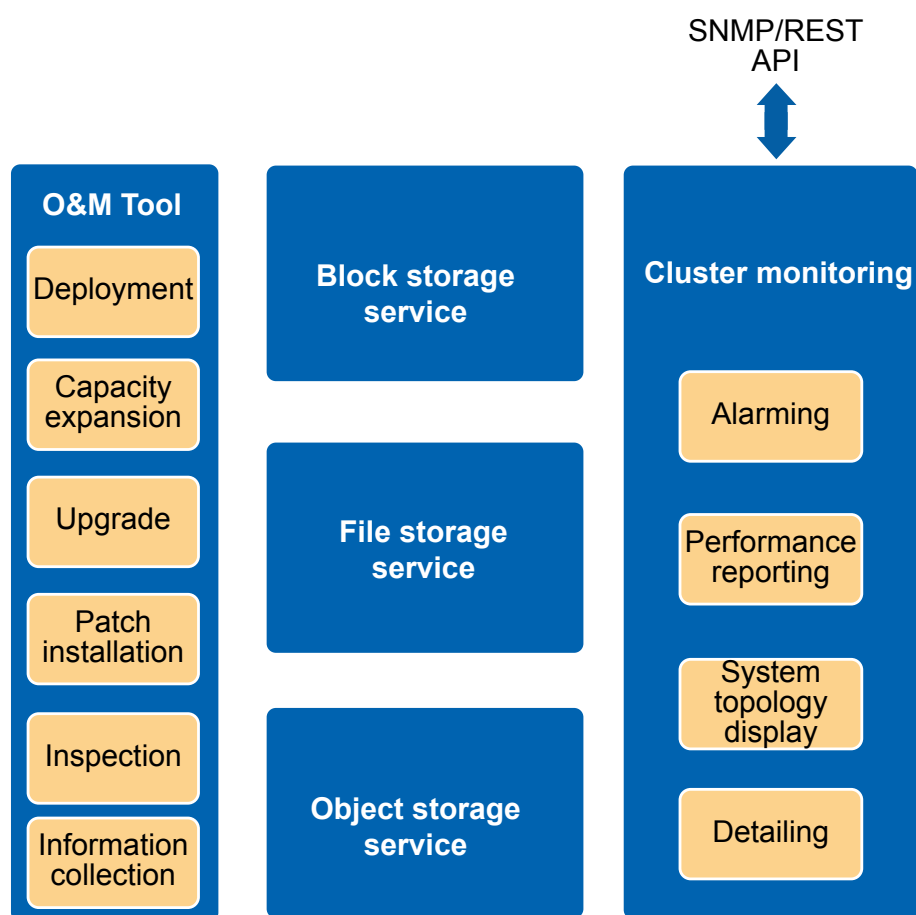


Table 2-1 describes each functional module of FusionStorage.

Table 2-1 Functional modules of FusionStorage

Module		Function
Cluster Monitoring	Alarming	Manages alarms of all service components centrally and sends them to the upper-layer NMS through a northbound interface.
	Performance reporting	Monitors system performance by key performance indexes, such as capacity usage and throughput, generates performance reports, and sends performance statistics to the upper-layer NMS through a northbound interface.
	System topology display	Displays device topology relationships.
	Device detailing	Provides device configuration details and monitoring information.
Block storage service		Provides distributed block storage services and supports SCSI and iSCSI protocols as well as value-added features, including snapshot and clone. NOTE iSCSI is supported only in VMware scenarios.
File storage service		Provides basic functions, such as NFS, CIFS, FTP, or HDFS-based file sharing, and advanced functions, such as InfoAllactor, InfoTier, InfoStamper, InfoReplicator, WORM, InfoScanner, NDMP, and CIFS/NFS cross-protocol access.
Object storage service		Provides object storage service compatible with Amazon S3 APIs and object storage service compatible with OpenStack Swift APIs.
Operation & Management Tool		Serves as an independent operation and maintenance tool used for system deployment, capacity expansion, upgrade, patch installation, and information collection. Toolkit can automatically identify hardware models and use different methods to deploy and maintain specific functional components.

3 Product Networking Architecture

About This Chapter

This chapter describes networking architecture of the file storage service, object storage service, and block storage service.

[3.1 File Storage Service Network Architecture](#)

The network architecture of file storage service contains the management network, front-end service network, and back-end storage network.

[3.2 Object Storage Service Network Structure](#)

This section describes network structures of the object storage service in single-cluster and multi-cluster scenarios.

[3.3 Block Storage Service Network Architecture](#)

3.1 File Storage Service Network Architecture

The network architecture of file storage service contains the management network, front-end service network, and back-end storage network.

Overview

The FusionStorage network architecture consists of the following:

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


The FusionStorage 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 FusionStorage in this example adopts a 10GE network where all nodes

provide Mgmt ports and are managed and maintained through management switch. **Figure 3-1** shows the 10GE network.

Figure 3-1 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.
- Management network: The Mgmt port of each node connects to an management switch with one GE network cable. A GE port on each of the first three nodes is connected to the management switch. A GE port on each back-end switch is connected to the management switch through a GE network cable.

3.2 Object Storage Service Network Structure

This section describes network structures of the object storage service in single-cluster and multi-cluster scenarios.

Networking Overview in Enterprise Scenarios

The network structure of the object storage service in enterprise scenarios consists of:


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

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

Typical Networking in Single-Cluster Scenarios

Use 10GE networking as an example. Each node provides a Mgmt port to connect to a management switch for node maintenance and management. **Figure 3-2** shows the network structure.

Figure 3-2 Typical networking in single-cluster scenarios



The networking is described as follows:

- Front-end service network: Two 10GE ports on each node are connected to two front-end switches through optical fibers.
- Back-end storage network: Two 10GE ports on each node are connected to two back-end switches through optical fibers.
- Management network: The Mgmt port of each node is connected to a management switch through one GE network cable. A GE port on each of the first three nodes is connected to the management switch. A GE port on each back-end switch is connected to the management switch through a GE network cable.

Networking Overview in Multi-Cluster Scenarios

The network structure of the object storage service in multi-cluster scenarios consists of:

- A front-end service network that communicates with the user's service network. The user can access the object storage service through object storage clients (such as OBS Console, OBS Browser, and PC client), REST API compatible with Amazon S3 APIs, and software development kit (SDK). By connecting to the external DNS server, Linux Virtual Server (LVS), Identity and Access Management (IAM), FusionStorage can resolve domain names, balance services, authenticate users, and encrypt data. In addition, FusionStorage can use the firewall to complete Network Address Translation (NAT) between networks.

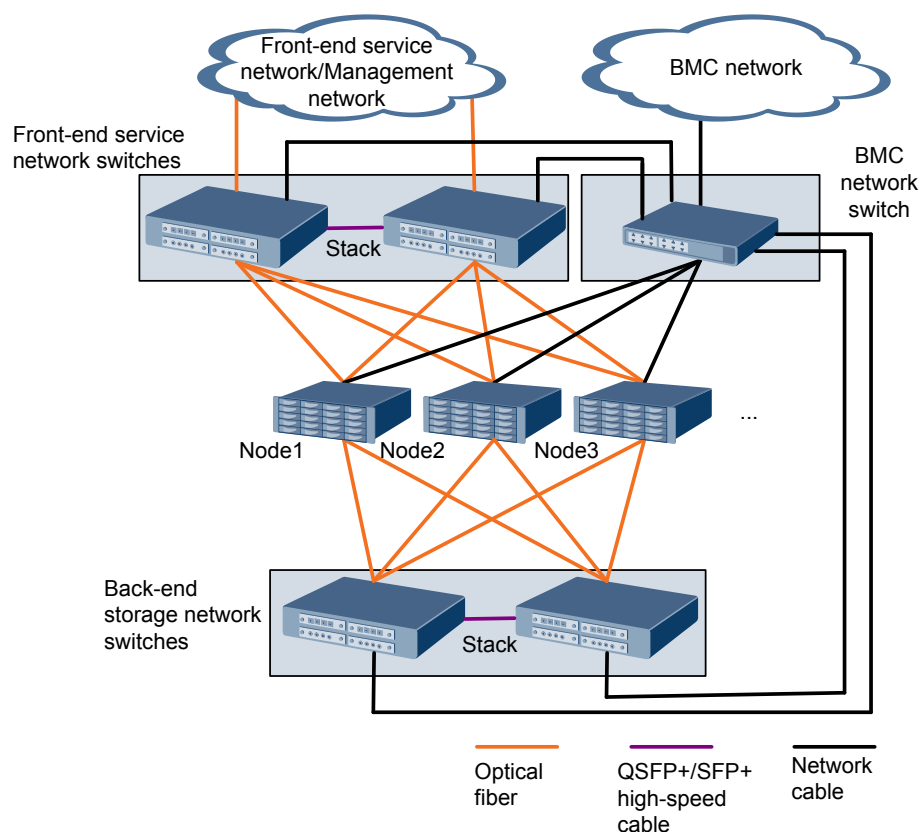
- A back-end storage network that transfers data and synchronizes status among all nodes of FusionStorage.
- A management network that communicates with the user's maintenance network. Administrators can use DeviceManager or Operation & Maintenance Tool to configure and maintain FusionStorage, eSight to collect, manage, and report alarms, OpsMonitor to monitor performance and report statistics, BSS to upload CDRs, and ELK to upload logs.
- A BMC network that remotely manages hardware by connecting to the Mgmt port of each node.

Only 10GE networking is supported in multi-cluster scenarios.

Typical Networking in Multi-Cluster Scenarios

Figure 3-3 shows the network structure.

Figure 3-3 Typical networking in multi-cluster scenarios



The networking is described as follows:

- Front-end service network and management network: Two 10GE ports (forming a logical bond interface) on each node are connected to two front-end switches through optical fibers.
- Back-end storage network: Two 10GE ports on each node are connected to two back-end switches through optical fibers.
- BMC network: The Mgmt port of each node is connected to a BMC switch through one GE network cable. A GE port on each back-end switch is connected to the BMC switch

through a GE network cable. A GE port on each front-end switch is connected to the BMC switch through a GE network cable.

3.3 Block Storage Service Network Architecture

Figure 3-4 shows the software architecture of FusionStorage.

Figure 3-4 Software architecture

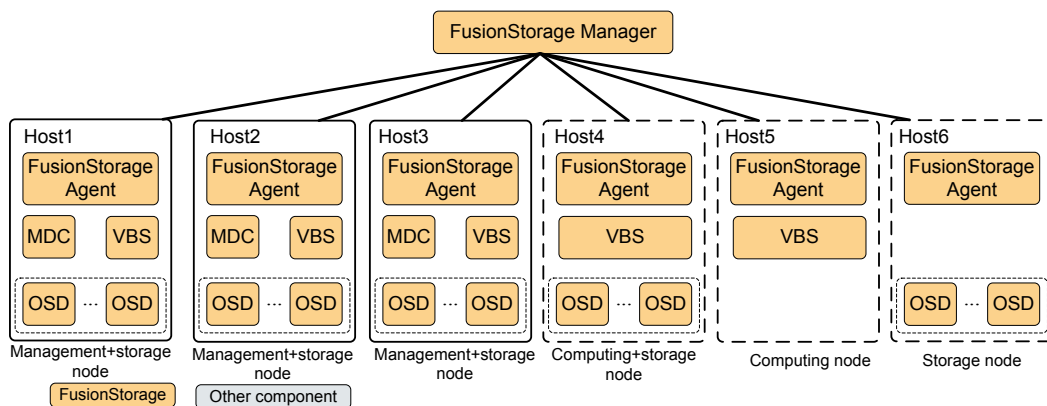


Table 3-1 describes the major components of FusionStorage.

Table 3-1 FusionStorage components

Component	Function
FusionStorage Manager (FSM)	A management process of the FusionStorage system. It supports O&M functions including alarm management, service monitoring, operation logging, and data configuration. It is best practice to deploy two FSM nodes working in active/standby mode.
FusionStorage Agent (FSA)	A management agent process of the FusionStorage system. It is deployed on each node (server) to communicate with the FSM node.
MDC	A service control process that controls status of distributed clusters and data distribution and reconstruction rules. When you create the control cluster, the metadata management service (ZooKeeper) is deployed on three, five, or seven nodes, and an MDC process is deployed on each ZooKeeper node. Nodes accommodating both the ZooKeeper and MDC processes form the control cluster. MDC can be deployed on a maximum of 96 nodes in the system. If the number of MDC nodes in the system is less than 96, the system automatically deploys MDC on a storage node upon a storage pool creation. Otherwise, the system does not deploy MDC.

Component	Function
VBS	A service input and output (I/O) process of the FusionStorage system. It manages metadata and provides an access service that enables computing resources to connect to distributed storage resources. A VBS process is deployed on each server to form a VBS cluster.
OSD	A service I/O process that performs I/O operations. Multiple OSD processes can be deployed on each server and one disk requires an OSD process.

4 File Storage Service

About This Chapter

This chapter introduces the functions and features supported by FusionStorage for the file storage service.

4.1 Service Function of the File Storage Service

The FusionStorage provides service functions such as CIFS/NFS/FTP share, HDFS, and NDMP backup.

4.2 Features of the File Storage Service

This section introduces the features supported by FusionStorage for the file storage service.

4.1 Service Function of the File Storage Service


The FusionStorage provides service functions such as CIFS/NFS/FTP share, HDFS, and NDMP backup.

4.1.1 CIFS/NFS/FTP Sharing

The FusionStorage 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 FusionStorage 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 FusionStorage are clustered to provide sharing services and monitor service status of each other.

The FusionStorage 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 FusionStorage fails over services from the faulty node to a functioning one to ensure business continuity.

NOTE

- FusionStorage cannot function as VM OS disks or data disks (VMDKs accessible through NFS) but can only function as VM web disks.
- FusionStorage does not support VAAI, a VM application optimization program designed for the VMware virtualization scenario.

4.1.2 CIFS-NFS Cross-Protocol Access Sharing

NFS and CIFS sharing can be set simultaneously for a directory of the FusionStorage. 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 FusionStorage 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 FusionStorage 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 FusionStorage, the FusionStorage 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 FusionStorage implements UNIX permission authentication for the user.

- If an NFS user attempts to access a file or directory that has NT ACL on the FusionStorage, the FusionStorage maps the NFS user as a CIFS user based on the preset mapping rules. Then the FusionStorage 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 an 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 changing the permissions type.

- When the parent directory has the inheritable NT ACL permission, the files or directories created no matter on an 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 an 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: An 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 owning 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.

After a user is mapped, the owner information of the files or directories owned by CIFS users (the files or directories that are created by CIFS users or the owner information of the files or directories are changed to CIFS users) is the information of the NFS users mapped from CIFS users on the NFS client. If no mapping rules have been configured for CIFS users, the owner information of the files or directories is about the IDs (calculated using IDMAP, a hash algorithm) of the CIFS users on the NFS client. If the client is an NFSv4 client, the owner information is displayed as **nobody**.

After a user is mapped, the owner information of the files or directories owned by NFS users (the files or directories that are created by NFS users or the owner information of the files or directories are changed to NFS users) is about NFS user names on the CIFS client:

- When NFS users are NIS domain users, the owner information is displayed as **NIS_DOMAIN\user name**.
- When NFS users are LDAP domain users, the owner information is displayed as **LDAP_DOMAIN\user name**.

When CIFS users are mapped to NFS users, quota statistics will be collected for the NFS users or owning user group.

If an NFS user creates a soft link using a full path on a Linux client, a mapped CIFS user cannot access the soft link on a Windows client.

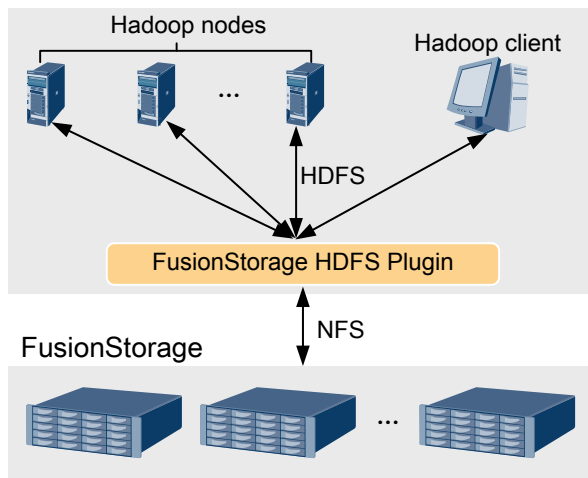
4.1.3 HDFS Interface

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

The FusionStorage 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 FusionStorage. [Figure 4-2](#) shows the HDFS feature.

Figure 4-2 HDFS feature

Hadoop cluster



FusionStorage supports Hadoop products and versions. For details, see the [OceanStor Interoperability Navigator](#).

NOTE

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


4.1.4 NDMP Backup

The FusionStorage 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 FusionStorage serves as the NDMP server and the data management application (DMA) serves as the client.

The FusionStorage supports NDMP V4 and adopts 3-way networking mode, as shown in [Figure 4-3](#).

Figure 4-3 3-way networking solution



In this mode, the FusionStorage is the primary storage. The backup server uses the NDMP-based management interface to instruct the FusionStorage 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 FusionStorage 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.
- Cumulative Incremental Backup
Only the file changed since last full backup is backed up.

NOTE

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

The FusionStorage 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.

4.1.5 High Data Reliability


The FusionStorage provides mirror and Erasure features to ensure high data reliability.

The FusionStorage 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 FusionStorage stores data on different nodes for redundancy to prevent data loss.

The FusionStorage 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 FusionStorage 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 4-4](#) shows the data protection process for a 6-node FusionStorage cluster protected in 4+2 mode.

Figure 4-4 File data protection process



1. File data is divided into multiple source data slices. ODC ranges from 2 to 18. ODC is automatically calculated by the FusionStorage. In this example, ODC is 4.

 **NOTE**

- The ODC value can be modified only in the data encryption scenario.
 - You can run the **change directory senior_info** command to change the ODC for a directory. After the change, files to be generated in the directory will be affected by the new ODC.
 - The possible values of the ODC are 0, 4, 6, 8, and 12. 0 indicates that no ODC is specified for the directory and each file to be created in the directory must comply with the following rules: If the ODC has been changed for some parent directories of the file, the ODC of the nearest parent directory is used. If the ODC of no parent directory has been changed, the ODC is automatically calculated by the system.
 - The automatically calculated value of the ODC may be 16 to 18 but this command does not support 16 to 18 to be a possible value of the ODC.
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 FusionStorage allows the failure of up to M disks on different nodes 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 FusionStorage 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.

4.1.6 Permission Control

The FusionStorage employs permission control to improve service access security.

Permission control includes system access permission control and directory access permission control. With permission control, the FusionStorage 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 FusionStorage or authentication server authenticates users accessing the NAS services. The FusionStorage 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 FusionStorage	
AD domain authentication	Authenticated by the AD domain server	
LDAP authentication	Authenticated by the LDAP 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 FusionStorage	


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.

4.1.7 Windows Hard Link

The FusionStorage 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 4-5](#), the file has three file names, indicating that the number of hard links is three.

Figure 4-5 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 FusionStorage, 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 FusionStorage 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.
- A compatibility issue may occur in Windows when you delete hard link files (created when NFS clients perform cross-protocol access to the OceanStor 9000) due to short name conversion required for Windows.

4.1.8 Unified Namespace

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

The FusionStorage 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 FusionStorage and is named after the FusionStorage. 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.

4.1.9 Technical Specifications of File Storage Service

This section describes the technical specifications of the file storage service.

Table 4-2 lists the capacity specifications of the file storage service.

Table 4-2 Capacity specifications of the file storage service

Specifications	Value
Max. number of nodes in a cluster	288
Max. capacity of a file system	100 PB
Max. size of a file	240 TB
Max. number of files/subdirectories in a directory	10 million
Max. number of files (including directories)	20 billion

4.2 Features of the File Storage Service

This section introduces the features supported by FusionStorage for the file storage service.


4.2.1 Tiering

The FusionStorage's Tiering feature provides automatic storage tiering.

Tiering 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, Tiering 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, Tiering automatically identifies cold and hot data and migrates data between storage media, as shown in **Figure 4-6**.

Figure 4-6 Schematic diagram of Tiering



Tiering provides the following functions:

- Node pool: Nodes of the same physical type can form a node pool. The node pool can contain 3 to 20 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 FusionStorage 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 FusionStorage 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 FusionStorage 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.

4.2.2 Quota

The FusionStorage provides the Quota feature to manage quotas.

Quota enables the FusionStorage 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 or any user in a user group 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 FusionStorage 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 4-3 describes the thresholds for mandatory quotas.

Table 4-3 Thresholds for mandatory quotas

Threshold Type	Description
Recommended threshold	When storage space usage or file quantity reaches a recommended threshold, the FusionStorage allows data writes and reports an alarm.
Soft threshold	When storage space or file number reaches a soft threshold, the FusionStorage still allows data writes for a grace period. However, after the grace period expires, the FusionStorage 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 FusionStorage forbids data writes and reports a run log and an alarm.

4.2.3 Load Balance

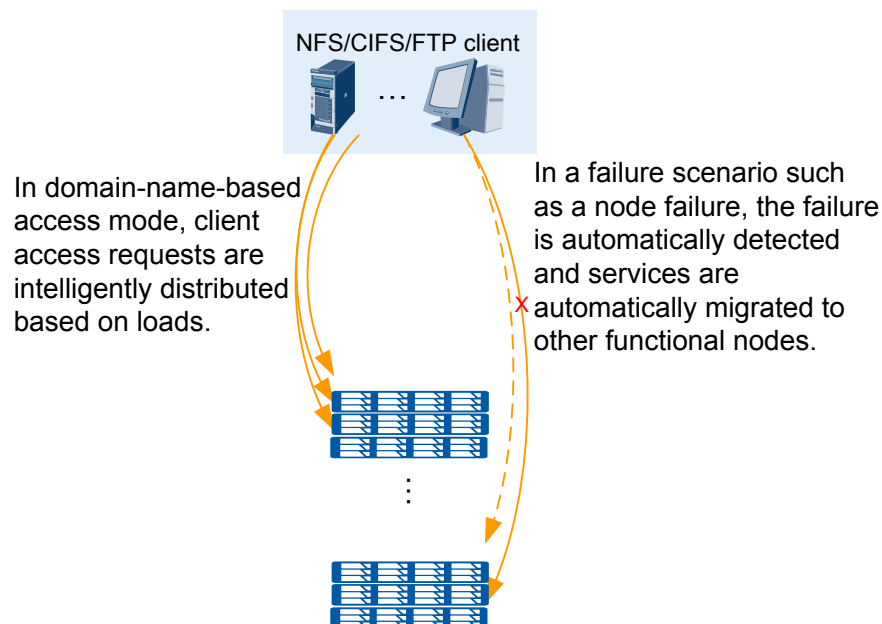
The FusionStorage provides the Load Balance feature that manages client connections to the FusionStorage and implements load balancing and failover.

Load Balance provides the following functions to improve availability, performance, and reliability of the FusionStorage:

- Enables a client to access the FusionStorage 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 4-7 shows a schematic diagram of Load Balance.

Figure 4-7 Schematic diagram of Load Balance



FusionStorage 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 FusionStorage calculates the load processing capability of a node based on the node's CPU and bandwidth configurations. If the node is overloaded, the FusionStorage 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.

4.2.4 WORM

The FusionStorage provides the WORM 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 FusionStorage 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 4-8 shows the WORM file states.

Figure 4-8 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.

4.2.5 Snapshot

The FusionStorage provides the Snapshot 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 FusionStorage 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 FusionStorage 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 FusionStorage 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 FusionStorage 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.

4.2.6 Replication

The FusionStorage provides the Replication 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 FusionStorage 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 FusionStorage 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

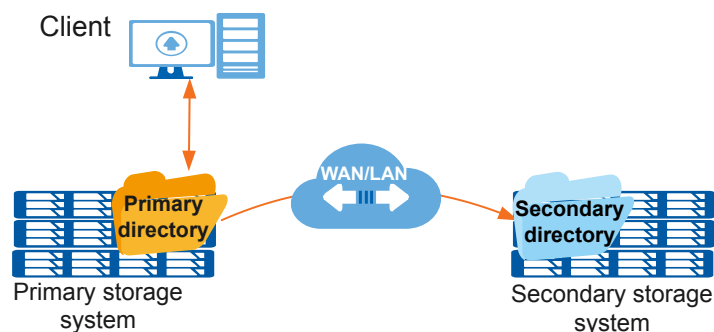
When the primary device in the production center needs to be upgraded or maintained, or fails due to disasters, the secondary device in the DR center takes over services. After disabling the write protection for the secondary directory, the client can mount the secondary directory. After the mounting, 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 4-9](#).

Figure 4-9 Remote replication in a 1:1 DR scenario



4.2.7 Anti-Virus

FusionStorage provides Anti-Virus, 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 4-10](#) shows the network structure of Anti-Virus.

Figure 4-10 Anti-Virus network structure

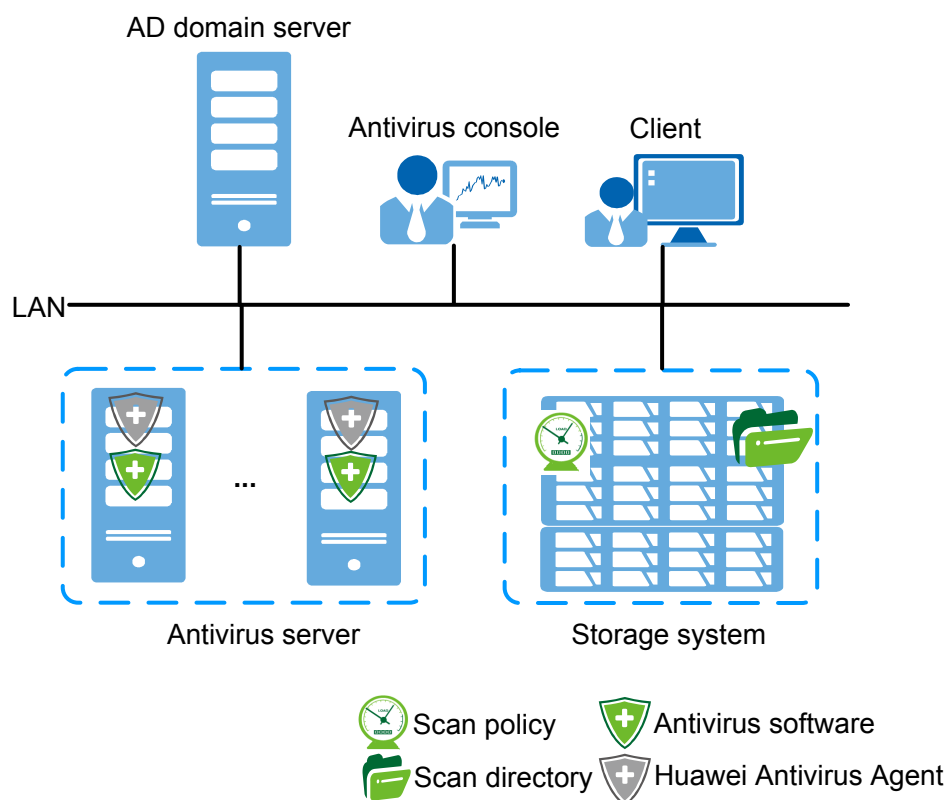


Table 4-4 provides details about the network structure. The front-end service network of FusionStorage, antivirus servers, and AD domain servers are interconnected.

Table 4-4 Network architecture description

Device	Description
AD domain server	<ul style="list-style-type: none"> Scan authentication users must be AD domain users and need to be added to AntivirusGroup (default local authentication user group in FusionStorage). <p>NOTE When the antivirus switch is enabled and antivirus servers are configured, scan authentication users can access storage resources in FusionStorage without being authenticated, namely, they have full control.</p> <ul style="list-style-type: none"> To ensure that Scan authentication users can access the FusionStorage by using antivirus servers, antivirus servers and FusionStorage need to be added to the AD domain.
Antivirus console	<p>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.</p> <p>NOTE The antivirus console needs to be deployed only after Rising virtualization system security software is installed on the antivirus server.</p>
Client	<p>A user accesses the FusionStorage through a client and writes data to the FusionStorage.</p>

Device	Description
Antivirus server	<ul style="list-style-type: none"> ● Antivirus servers access storage resources in scan directories in FusionStorage 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. ● Antivirus Agent Software is downloaded using DeviceManager and installed on every antivirus server to trigger antivirus software to scan files in scan directories. ● For details about supported antivirus software, see OceanStor Interoperability Navigator. 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 (FusionStorage)	<ul style="list-style-type: none"> ● A scan directory refers to the directory to be scanned in FusionStorage. A directory can be configured as a real-time or periodical scan directory. Antivirus servers access a scan directory using the CIFS protocol. ● 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. ● The FusionStorage interconnects with the antivirus software deployed on an antivirus server to scan the scan directories for viruses.

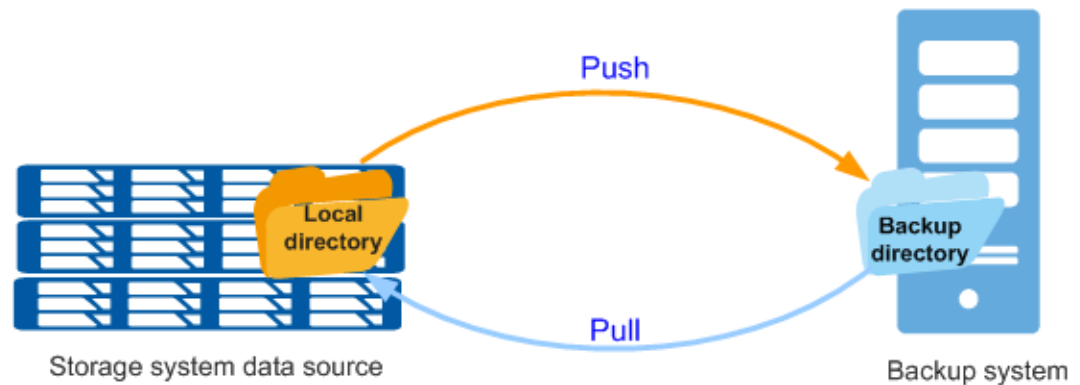
4.2.8 InfoMigrator

The FusionStorage provides asynchronous InfoMigrator in full and incremental manners, allowing you to synchronize data from the local FusionStorage to a remote backup system (an FusionStorage 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 FusionStorage supports Push and Pull file synchronization scenarios, as shown in [Figure 4-11](#).

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

Figure 4-11 File synchronization scenarios



In the Push scenario, the FusionStorage 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 FusionStorage.

4.2.9 InfoRevive

The FusionStorage 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 FusionStorage 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.

4.2.10 InfoTurbo

The FusionStorage provides the Performance Acceleration feature that includes intelligent prefetch, SMB3 Multichannel, and NFS protocol enhancement.

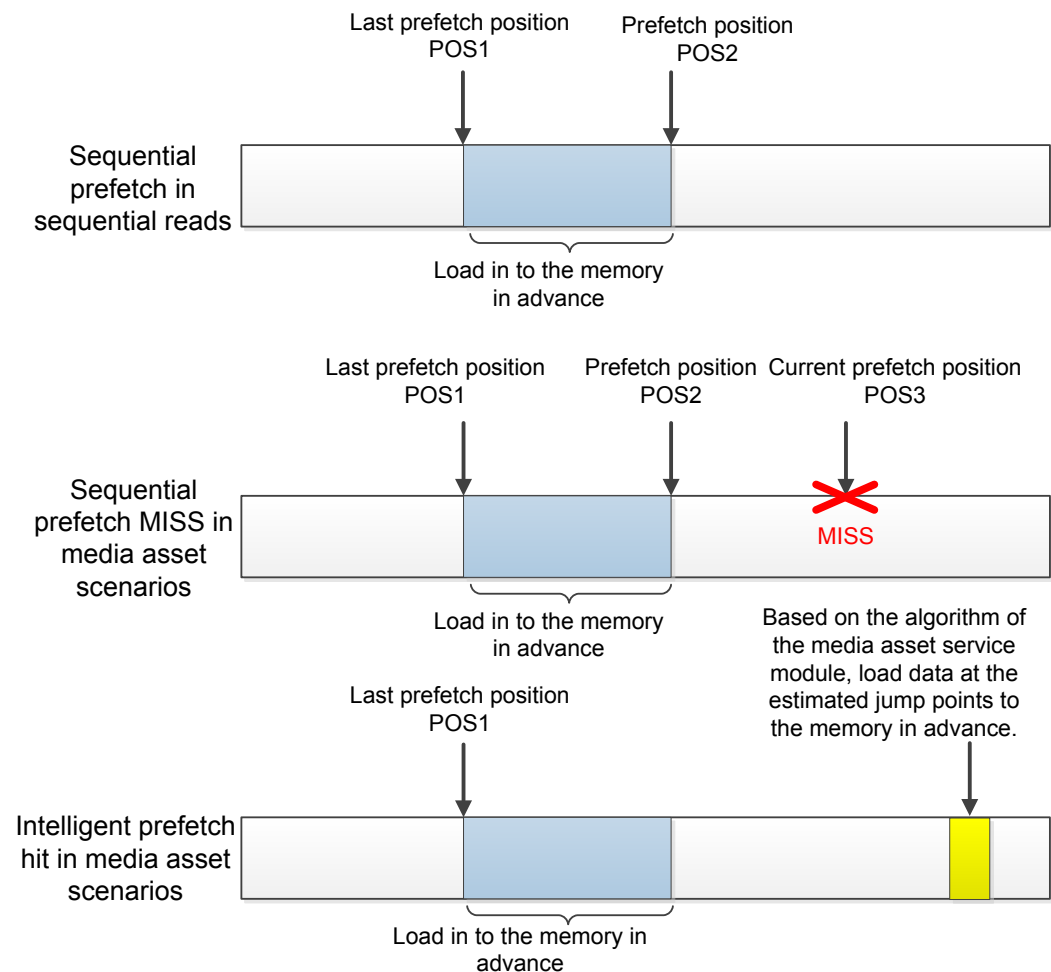
4.2.10.1 Intelligent Prefetch

The FusionStorage 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 FusionStorage provides the intelligent prefetch technology based on sequential prefetch. **Figure 4-12** shows sequential and intelligent prefetch.

Figure 4-12 Sequential and intelligent prefetch



- 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 4-12](#), 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 FusionStorage 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 FusionStorage 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.

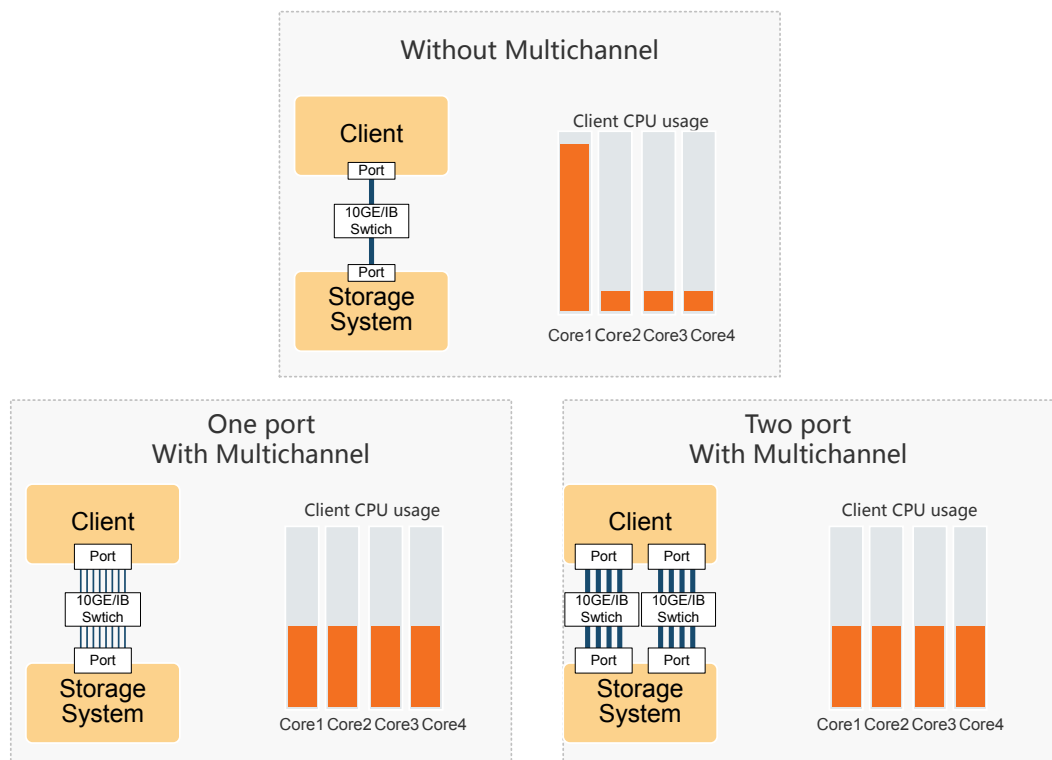
4.2.10.2 SMB3 Multichannel

The FusionStorage provides the SMB3 Multichannel function to greatly improve service performance and reliability. In addition, if one channel fails, the FusionStorage 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 network ports of the same type or with one GE/10GE network port that supports Receive-Side Scaling (RSS), the client will set up multiple channels with the FusionStorage.

[Figure 4-13](#) 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.

Figure 4-13 Schematic diagram of SMB3 Multichannel



4.2.10.3 NFS Protocol Enhancement

FusionStorage 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 FusionStorage. By configuring multiple network ports and installing NFS protocol optimization plug-in DFSCClient 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 FusionStorage.
2. FusionStorage 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 4-14](#).

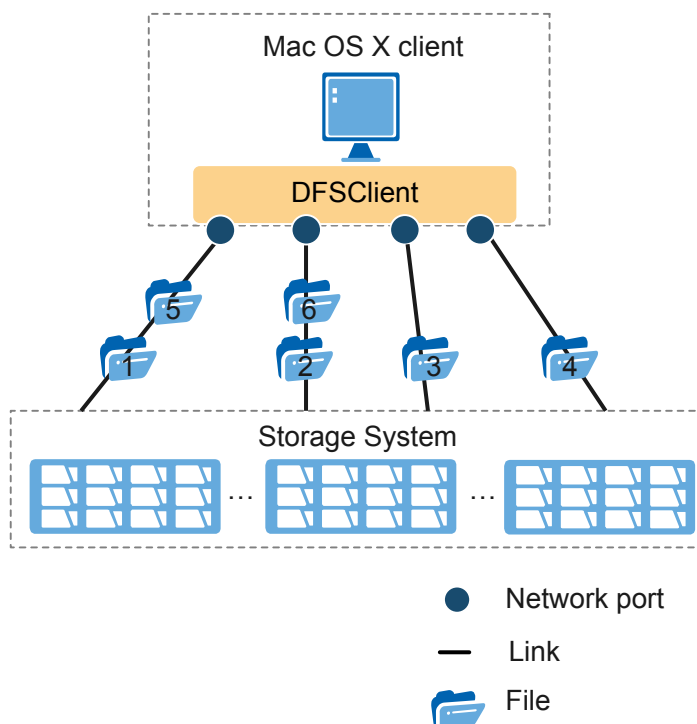
Figure 4-14 Establishment of multiple connections on the Mac OS X client



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

- When $M \geq S$:
FusionStorage returns all the dynamic front-end service IP addresses. The number of connections that can be established is S.
 - When $M < S$:
FusionStorage 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 4-15](#).

Figure 4-15 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 FusionStorage.

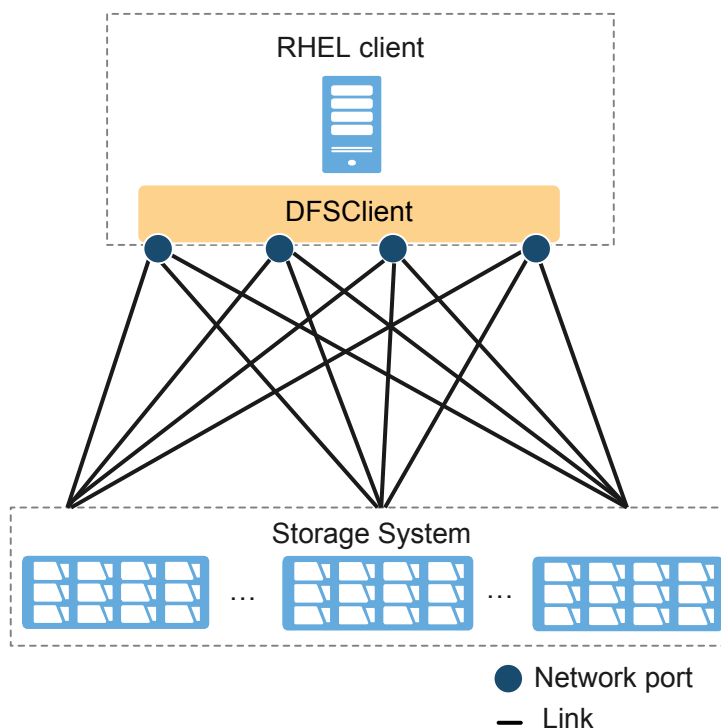
During multi-layer video editing, video files are accessed concurrently. DFSCient can distribute files to connections evenly, thereby accelerating data transmission.

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 FusionStorage.
2. FusionStorage 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 4-16](#).

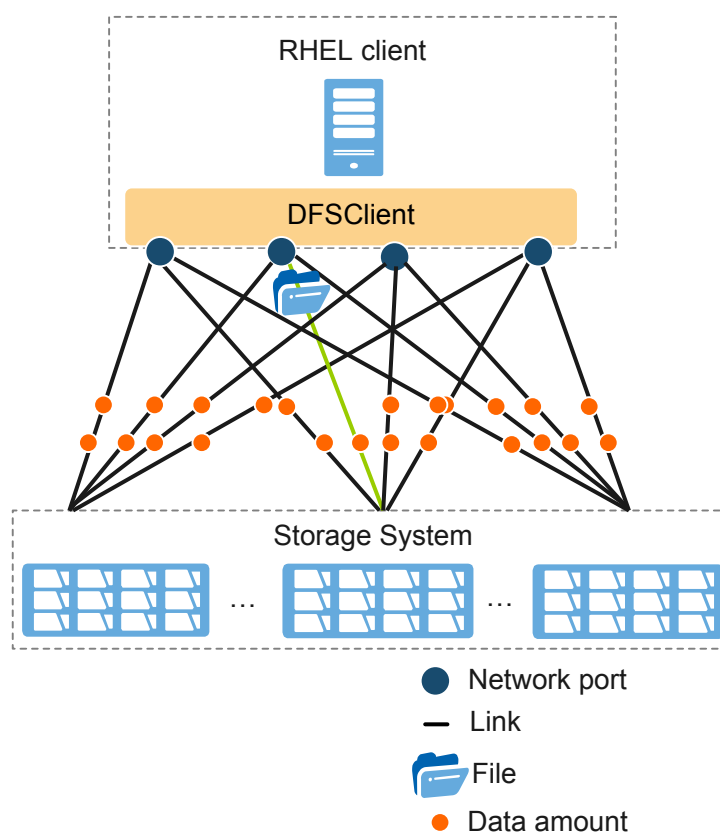
Figure 4-16 Establishment of multiple connections on the RHEL client



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 FusionStorage.
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 4-17](#).

Figure 4-17 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 FusionStorage.

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


4.2.11 Rapid internal file replication

FusionStorage supports CopyChunk. When CopyChunk is used, clients running on Windows 7, Windows Server 2008, Windows 8, Windows Server 2012, or later can complete access to shared files on server-side rather than traversed through the network. In this manner, the network interaction is reduced and the client performance can be maintained.

CopyChunk is an advanced feature of CIFS and is disabled by default. The Mac client does not support CopyChunk. Therefore, if you are using the Mac client, you have to contact Huawei technical support engineers before enabling CopyChunk.

Figure 4-18 shows a schematic diagram of rapid internal file replication.

Figure 4-18 Schematic diagram of rapid internal file replication



5 Object Storage Service

About This Chapter

This chapter introduces the functions and features supported by FusionStorage for the object storage service.

[5.1 Object Storage Service \(Compatible with Amazon S3 APIs\)](#)

The FusionStorage 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.

[5.2 Object Storage Service \(Compatible with OpenStack Swift APIs\)](#)

The FusionStorage 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.

[5.3 Technical Specifications of Object Storage Service](#)

This section describes the technical specifications of the object storage service.

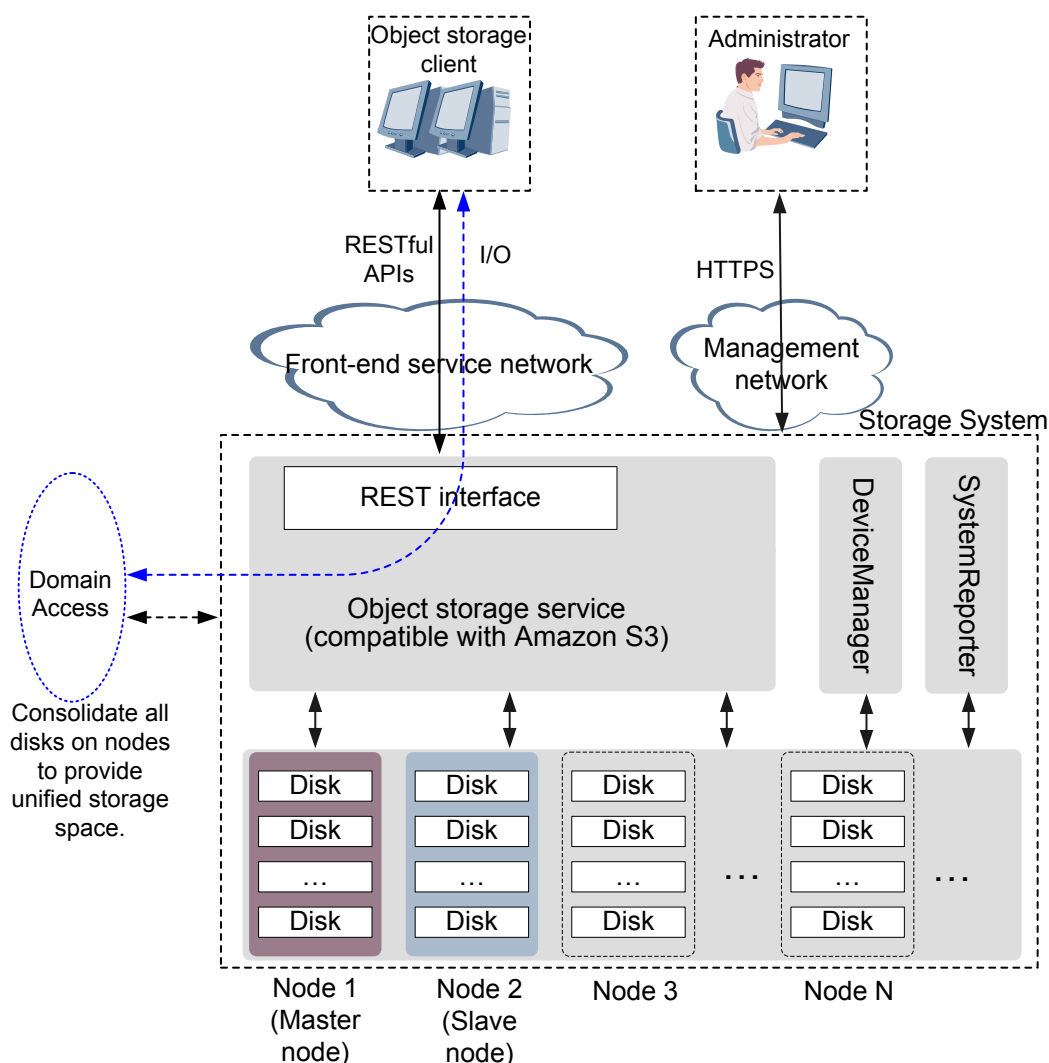
5.1 Object Storage Service (Compatible with Amazon S3 APIs)

The FusionStorage 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 5-1 shows Amazon S3 APIs with which the object storage service is compatible.

Figure 5-1 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 FusionStorage 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 FusionStorage 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 FusionStorage customers. Clients can be application software and visual websites.

- Users' access to the object storage service needs to be authenticated by AK/SK.

- Administrators can securely access the FusionStorage to manage system resources and monitor system status and performance through Hypertext Transfer Protocol Secure (HTTPS).

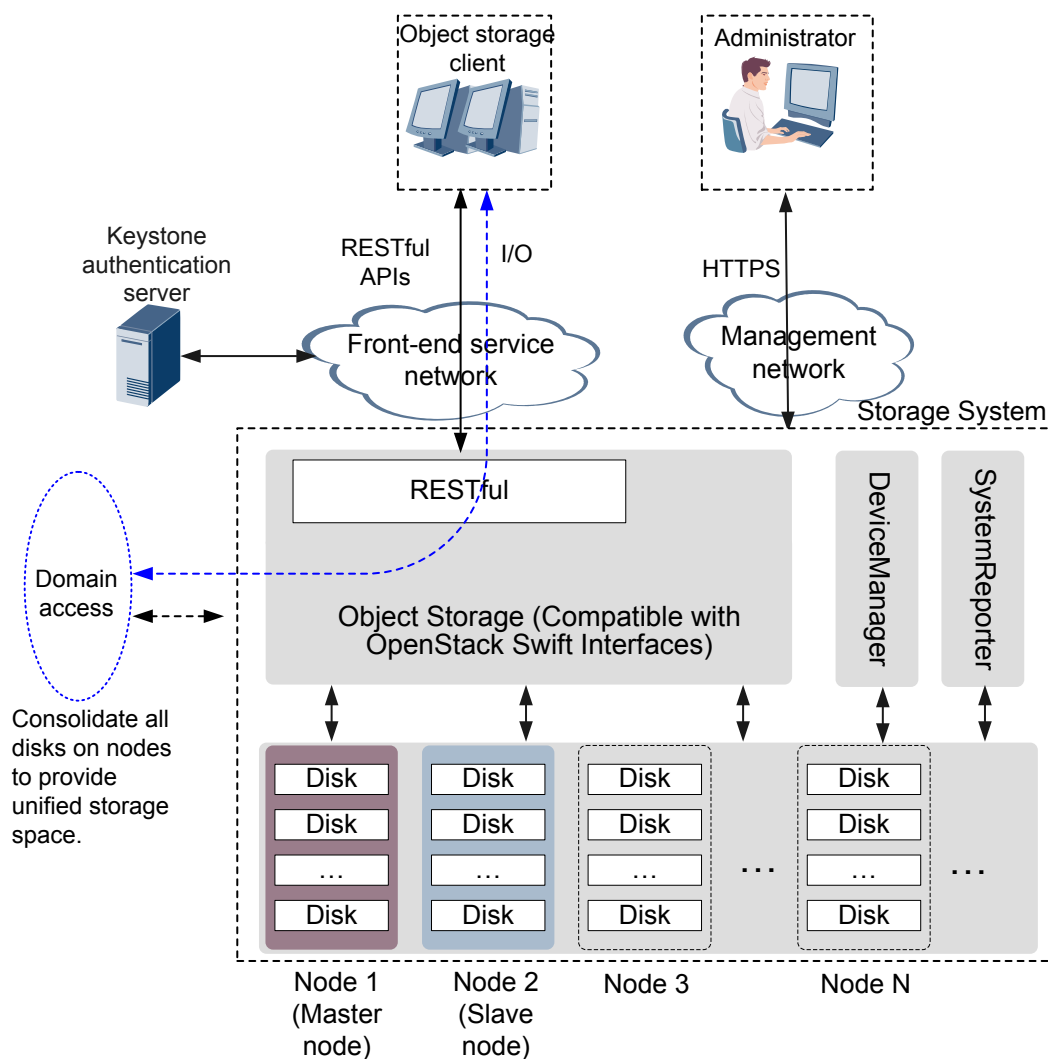
5.2 Object Storage Service (Compatible with OpenStack Swift APIs)

The FusionStorage 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 5-2 shows OpenStack Swift APIs with which the object storage service is compatible.

Figure 5-2 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 FusionStorage 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 API (compatible with OpenStack Swift), through which clients can access FusionStorage 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 FusionStorage. 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 FusionStorage through HTTPS to manage system resources and monitor system status and performance.

5.3 Technical Specifications of Object Storage Service

This section describes the technical specifications of the object storage service.

Table 5-1 lists the capacity specifications of the object storage service.

Table 5-1 Capacity specifications of the object storage service

Specifications	Value
Max. number of buckets in a cluster	100 million
Max. number of buckets created by a tenant	100
Max. number of containers created by a tenant	10000
Max. number of objects in a cluster	10 billion
Max. number of objects in a bucket	50 million
Max. number of accounts supported by a cluster	1 million

6 Block Storage Service

About This Chapter

This chapter introduces the functions and features supported by FusionStorage for the block storage service.

[6.1 Multi-Resource Pool](#)

[6.2 Distributed RAID](#)

[6.3 Snapshot](#)

[6.4 Thin Provisioning](#)

[6.5 Linked Cloning](#)

[6.6 High Data Reliability](#)

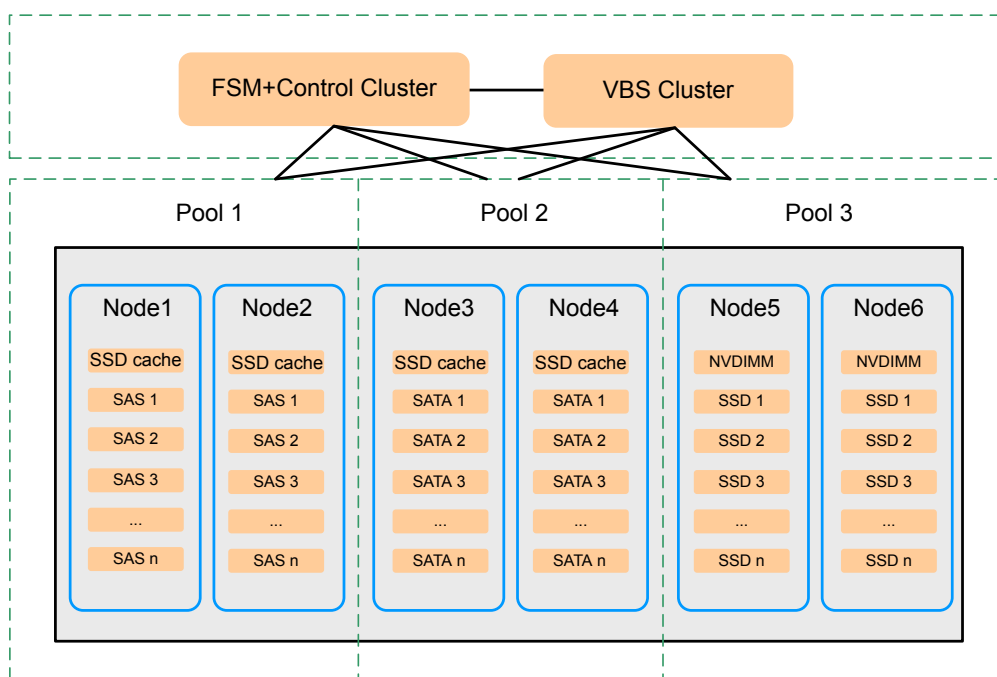
[6.7 Security Protection Mechanisms](#)

6.1 Multi-Resource Pool

FusionStorage simultaneously supports the demands of storage pools of different capacities, and isolates the faults to enhance reliability.

- Storage policy: Configures main storage media, cache media, copy redundancy type, and security policy (rack/server) for resource pools.
- High reliability: Uses faults isolation and multi-MDCs to enhance reliability.

Figure 6-1 Multi-Resource Pool



6.2 Distributed RAID

Cluster Management

FusionStorage manages the system in clusters. If a server or hard disk becomes faulty, it can be automatically isolated from the cluster and therefore has no adverse impact on system services. In a cluster, a controller process is elected among all the processes to process the data storage logic. When the controller process fails, the system automatically elects a new controller process.

Data Routing

Backed by the hash routing algorithm, FusionStorage has the following characteristics:

- Rapid load balancing: Only a minimal amount of data needs to be migrated to the new nodes to achieve load balancing.
- High data reliability: The flexible partition allocation algorithm effectively prevents identical data copies from being stored on the same server or disk, thereby avoiding single point of failure (SPOF).

Tight Consistency and Replication

When a user successfully writes application data into a disk, the system automatically duplicates the data into multiple copies and stores them in different disks. Then the user can read the data from any of the disks.

FusionStorage uses the tight consistency and replication technology to ensure data consistency between data copies. If the data written into the system is divided into multiple portions, the system also creates multiple copies for each data portion and stores them in

different disks. FusionStorage also ensures consistency between the copies of each data portion.

FusionStorage also supports the read repair mechanism. When failing to read data, FusionStorage automatically identifies the failure location. If the data fails to be read from a disk sector, FusionStorage retrieves the data from other copies of the data and writes the data back into the original disk sector. This ensures that the total number of data copies does not decrease.

6.3 Snapshot

FusionStorage provides a snapshot mechanism, which allows the system to create a snapshot for the status of the data written into a logical volume at a time point. The data snapshot then can be exported and used for restoring the volume data when required.

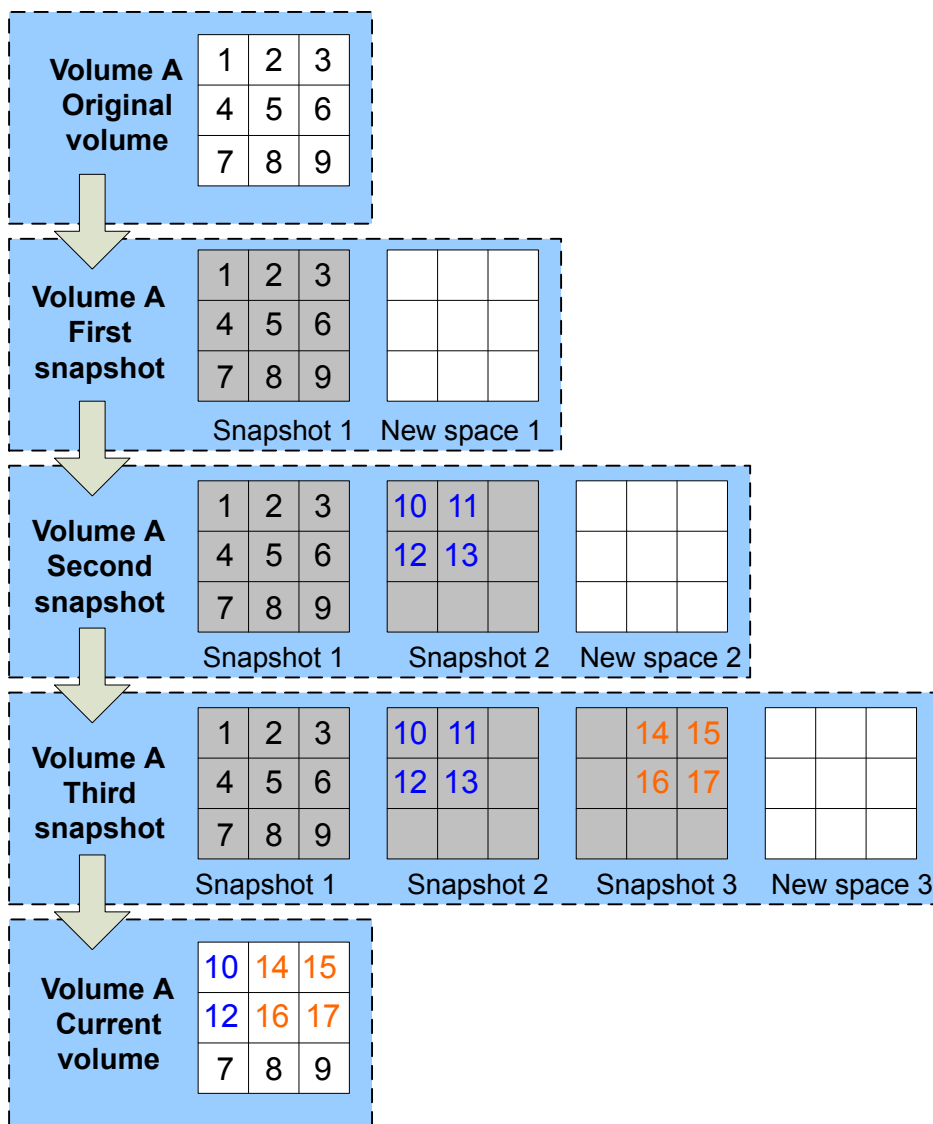
FusionStorage supports the following snapshot functions:

- Create a snapshot.
- Query a snapshot.
- Delete a snapshot.
- Back up snapshot data.
- Restore volume data using a snapshot.

FusionStorage creates snapshots based on the distributed hash table (DHT) technology. Creating snapshots does not have any adverse impact on the volumes. The DHT technology provides high query efficiency. For example, only one hash query can determine whether a snapshot has been created for a 2 TB hard disk. If a snapshot has been created, the hash query can also determine the storage location of the snapshot.

Figure 6-2 shows the principle of the FusionStorage snapshot mechanism.

Figure 6-2 Principle of the snapshot mechanism

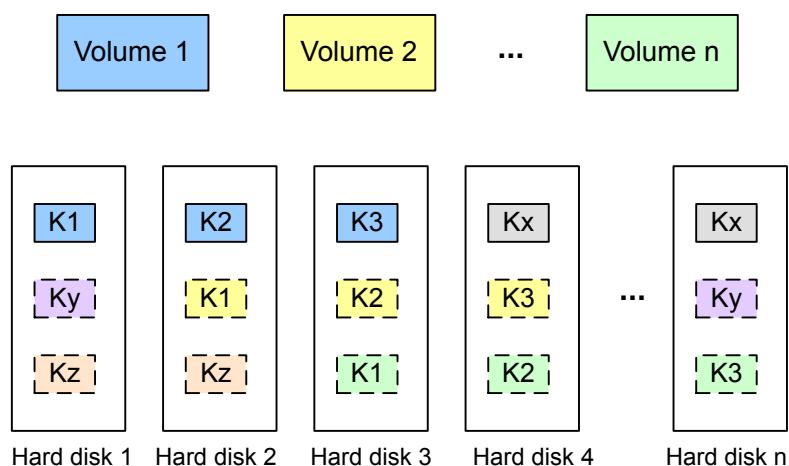


6.4 Thin Provisioning

FusionStorage uses the distributed architecture to support distributed thin provisioning.

As shown in [Figure 6-3](#), physical space is allocated for users only when users write data into volumes. FusionStorage only processes the mapping between virtual volume space and physical space, exerting no adverse impact on system performance.

Figure 6-3 Mapping between virtual volumes and physical hard disks



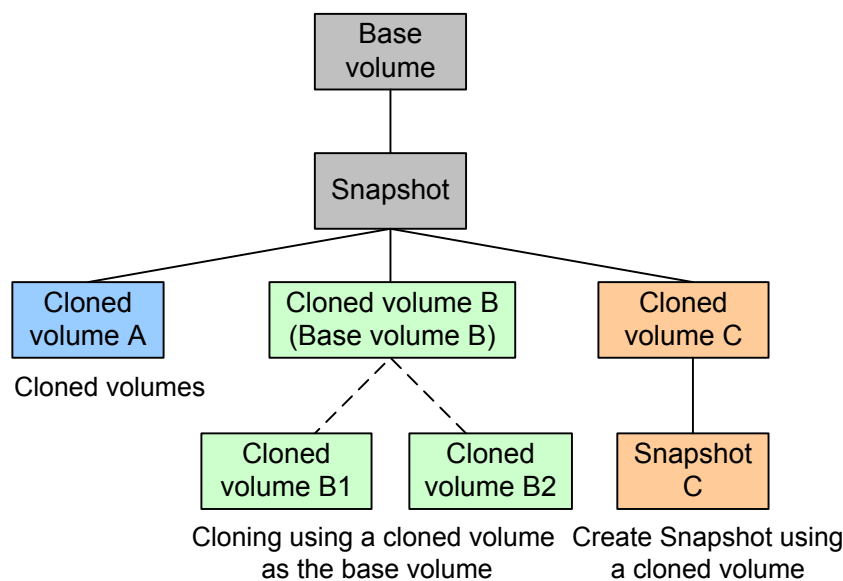
6.5 Linked Cloning

FusionStorage provides the linked clone mechanism for incremental snapshots so that multiple cloned volumes can be created for a snapshot. The data in the cloned volumes is the same as that in the snapshot. Subsequent modifications to a cloned volume do not affect the snapshot or other cloned volumes.

- FusionStorage supports a linked cloning ratio of 1:2048, which can greatly improve storage space utilization.
- FusionStorage supports batch volume creation for VMs and can create hundreds of volumes in a second.
- A cloned volume has all the functions of a common volume. You can create snapshots for a cloned volume, use the snapshot to restore the data in the cloned volume, and clone the data in the cloned volume.

Figure 6-4 shows the principle of the linked cloning mechanism.

Figure 6-4 Principle of the linked cloning mechanism



6.6 High Data Reliability

Data Protection

FusionStorage protects data using the multi-copy and Erasure Coding (EC) methods, enabling normal data access and automatic data recovery when a physical device is faulty. Users can use either of the two methods to optimally balance data reliability and storage space utilization. In traditional disk-level RAID mode, data resides on different disks in a single node and cannot be restored when the entire node fails. FusionStorage stores data across nodes in redundancy mode or creates replicas to prevent data loss.

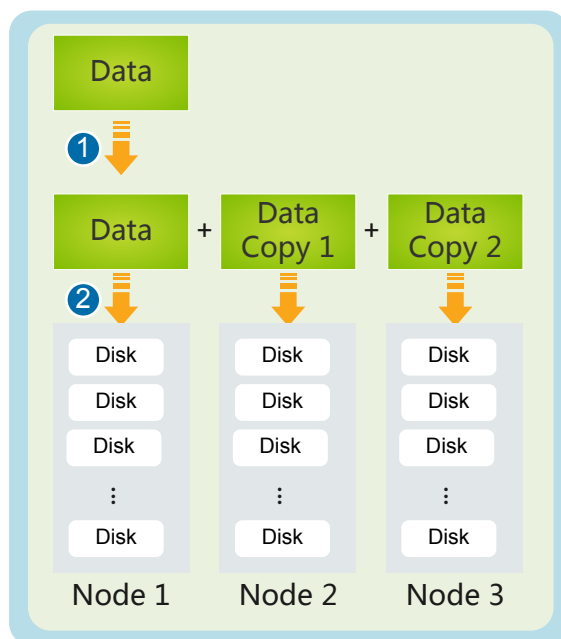
NOTE

The EC method with the same error tolerance as the multi-copy method has a higher space utilization but lower system performance than the multi-copy method. Users can decide which data protection method to use for a storage pool based on the actual site requirements.

Multi-copy

FusionStorage supports two-copy and three-copy storage pools. [Figure 6-5](#) shows an example of a three-copy storage pool.

Figure 6-5 Basic principles of Multi-copy



1. During data writing, two identical data copies are created for each piece of data.
2. Each piece of data and its two data copies are written into three storage nodes respectively.

Table 6-1 Multi-replica definition

Mode	Description	Disk Utilization	Min. Number of Nodes
Two-copy	One copy is created for each piece of data. When any storage node or any cabinet is faulty, data integrity is not affected.	50%	3
Three-copy	Two copies are created for each piece of data. When any two storage nodes or any two cabinets are faulty, data integrity is not affected.	33%	3

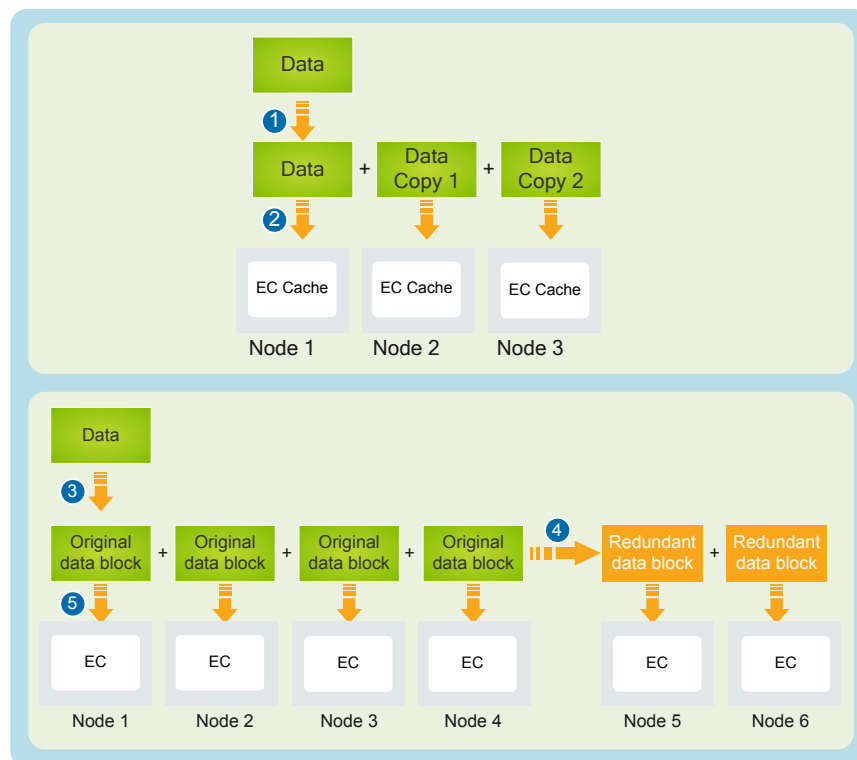
Erasur coding

FusionStorage uses the Erasure coding technology for storage pools, supporting the N+1, N+2, N+3 protection levels. **Figure 6-6** shows an example of the 4+2 protection level with the EC Cache function enabled.

NOTE

The EC Cache function is optional.

Figure 6-6 Basic principles of Erasure coding



1. After the EC Cache function is enabled, when data is written into the EC Cache, two copies are created for each data.
2. The data and its two copies are written into EC Caches of three storage nodes respectively.
3. When the system periodically delivers the original data in an EC Cache, each data is divided into four original data blocks.
4. The four original data blocks form a group and generate two redundant data blocks based on calculations.
5. The four original data blocks and two redundant data blocks are written into storage nodes at a redundancy ratio.

Table 6-2 Protection levels

Mode	Description	Disk Utilization	Min. Number of Nodes
N+1	Each piece of data contains <i>N</i> original data blocks and one redundant data block generated based on calculations. When any storage node or any cabinet is faulty, data integrity is not affected.	75%	4
N+2	Each piece of data contains <i>N</i> original data blocks and two redundant data block generated based on calculations. When any two storage nodes or any two cabinets are faulty, data integrity is not affected.	60% to 80%	5
N+3	Each piece of data contains <i>N</i> original data blocks and three redundant data block generated based on calculations. If any three storage nodes or any three cabinets are faulty, data integrity is not affected.	80%	15

Note 1: A protection level is expressed in the format of N+M, where N indicates the original data count (ODC), M indicates the redundant data count (RDC).

Data Recovery

Quick Data Rebuilding

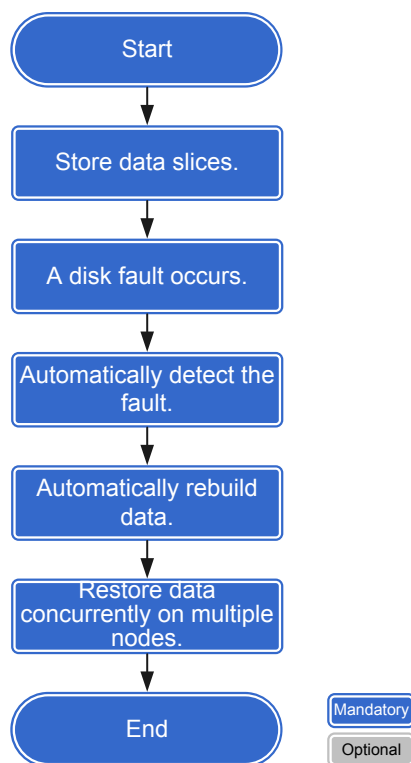
FusionStorage provides a powerful data protection mechanism, dividing data into slices and storing the slices on multiple nodes and cabinets using the multi-replica and Erasure coding technologies.

When a fault occurs on a node or cabinet, data stored on the node is automatically migrated to other nodes in the storage pool, ensuring service continuity. When the fault on the node or

cabinet is rectified, FusionStorage can quickly recover data stored on the faulty nodes by repairing and rebuilding a small amount of data concurrently on each faulty node where data slices are stored. Concurrently repairing a small amount of data on each node is free of performance bottlenecks and minimizes the impact on upper-layer services, delivering high data reliability.

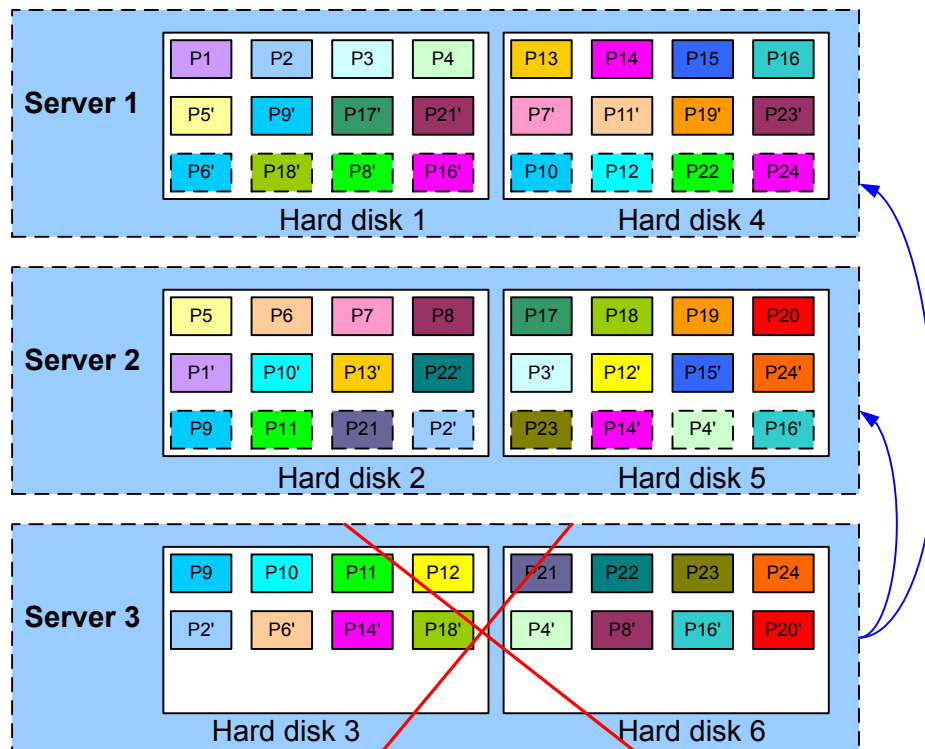
Figure 6-7 shows the automatic fault rectification flowchart.

Figure 6-7 Data rebuilding flowchart



In Figure 6-8, when server 3 is faulty, data on server 3 is automatically migrated to servers 1 and 2 in a balanced manner.

Figure 6-8 Data rebuilding principles



Power failure protection

A power failure may occur while FusionStorage is running. In this case, the metadata and cached written data stored in the memory may be lost due to the power failure. To prevent data loss in such cases, FusionStorage uses solid-state drive (SSD) cache to store and restore metadata and cached data.

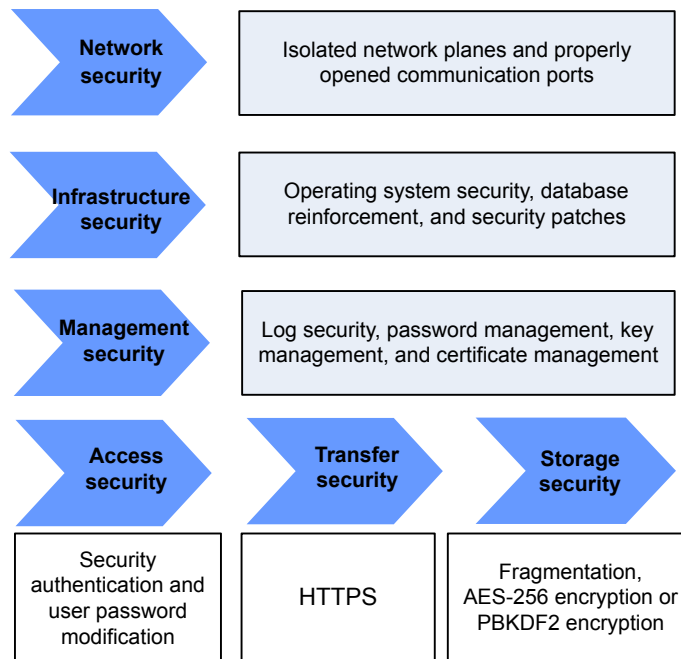
Each server running the FusionStorage software must be equipped with the SSD cache, so that the server can write the metadata and cached data into the flash of the SSD cache upon a power failure. After the power supply is resumed, FusionStorage restores the data stored in the flash back to the memory.

FusionStorage can identify the SSD cache in the system and store critical data in the SSD cache based on internal rules to prevent data loss caused by power failures.

6.7 Security Protection Mechanisms

FusionStorage employs multiple mechanisms to protect system and data security, including network, infrastructure, transmission, access, storage, and management security. [Figure 6-9](#) shows the application of security protection mechanisms.

Figure 6-9 Application of security protection mechanisms



Network Security

Communication ports open with security measures

FusionStorage opens only the communication ports necessary for implementing system functions, eliminating security risks caused by too many opened ports.

Infrastructure Security

Operating system security

To ensure device security, the FusionStorage operating system is configured as follows:

- Operating system account and password
Unnecessary users and user groups are deleted. The complexity and length of passwords must comply with specific requirements. The secure encryption algorithm is used to encrypt user passwords.
- System service
Insecure services such as Telnet and Network File System (NFS) are disabled. Needless and dangerous background processes and services are also disabled. Services can run only by using accounts that are granted specific permissions. Security protocols, such as Secure Shell (SSH) V2 and Secure File Transfer Protocol (SFTP), are used.
- Operating system kernel
Execution stacks are protected against buffer overflow attacks. IP spoofing prevention is enabled, and socket sequences are guarded against attacks.
- File and folder permission
The permissions for files and folders are strictly controlled.

Database reinforcement

The GaussDB database deployed on the FusionStorage Manager (FSM) node provides the following security measures:

- Access source control: allows only the access requests from the local host of the FSM node based on the actual service requirements and security standards. Access requests from other hosts are all denied, thereby protecting the system against external attacks.
- Minimum authorization: Roles are configured for the users other than the database administrator based on the minimum authorization rule.
- Directory protection: The installation user is the owner of the data installation directory and its data area directory and has the read, write, and execution permission for these directories and their subdirectories.
- Sensitive file protection: The installation user is the owner of the core configuration files in the database and has the read and write permission for these files.
- Connection number limit: The maximum number of connections is set to 200 to prevent malicious attacks.

Security patches

Security patches can prevent virus attacks and hacker invasion. FusionStorage has the latest security patches installed before delivery and periodically receives newly released security patches.

Access Security


Secure authentication mechanisms are adopted for invoking FusionStorage interfaces, which allows different accounts to perform different service operations. Users can change or reset their passwords.

Transfer Security

FusionStorage supports data transfer over HTTPS to prevent data leaks during transfer.

Figure 6-10 uses HTTPS access as an example to describe the technical principles of the data transfer security technology.

Figure 6-10 Technical principles of the data transfer security technology



1. A client sends a request for accessing an HTTPS resource.
2. A unique session key is generated at the client site. The public key generated by the server certificate is used to encrypt the session key, and then the session key is transferred to the server.

3. After receiving the session key, the server uses its own private key to decrypt it.
4. The connection is set up. The client can communicate with the server in secure mode.

Storage Security

FusionStorage uses diversified storage security technologies to ensure data security.

- **Fragmentation**
FusionStorage automatically stores each piece of data as multiple identical data copies into different storage nodes so that unauthorized users cannot obtain complete data from a single storage node or physical disk.
- **Sensitive user data encryption**
Before storage, the 256-bit Advanced Encryption Standard (AES) algorithm or PBKDF2 algorithm is used to encrypt sensitive data, such as authentication information.

Management Security

Log security

- In the FusionStorage system, modification operations, such as creating a storage pool or adding a server, are recorded in operation logs. Users can understand the system running status by viewing the logs. Operation logs are retained for 90 days by default, and the log retention policy is configurable.
- All the modification operations are recorded in audit logs. Audit logs are stored on local disks for 90 days by default, and access to these logs is controlled by permission.

Password management

- **Password change**
FusionStorage allows users to change passwords and administrators to change or reset passwords. In addition, FusionStorage automatically verifies password complexity.
- **Encryption and storage**
On all service nodes, user passwords are all encrypted using secure encryption algorithms before storage.
- **Operation logs**
All operations performed on passwords are recorded in logs. Those logs help users locate faults and implement audits.

Key management

- **Hierarchical key management**
The hierarchical key management mechanism consists of three layers of keys: root key, master key, and working key. Such mechanism lowers the key update costs and avoid the direct use of root keys to encrypt service data, significantly reducing the probability of attackers using various cryptanalysis methods to obtain root keys and securing the cryptography system.
- **Unique key usage**
To improve the cryptography system security, each key is used only for a single service. Such mechanism avoids the key management system crash induced by cracking the key as well as the ciphertext and data leakage.

- **Key update**
If the key leakage occurs, updating the key as well as the sensitive data encrypted using the key reduces the key leakage impact.
- **Distributed key storage**
The root key is distributed in different files, and access to these files are controlled by permission. Even if an attacker is breaking into the system, the attacker cannot obtain all the materials about the root key.
- **Audit logs**
Operations performed to keys are recorded in audit logs, and such records are used to locate faults and audit logs.

Certificate management

All FusionStorage certificates are centrally managed, and a user can import, update, and query a certificate. If a certificate expires or the certificate leakage occurs, update the certificate to improve the system security.

Web Security

- **Anti-brute force cracking supported on the login page**
On the FusionStorage login page, the system generates verification codes randomly. A user can log in to the system only when the username, password, and verification code they entered are correct.
- **Role-based user management**
System users can be managed based on three roles: system administrator, system operator, and system viewer. Users with different roles assigned have different permissions.
- **Session management**
Each user can only have one session in progress and manages the session. Administrators can view the session list and force specified sessions to exit the system.
- **Session timeout**
A user will be automatically logged out of the FusionStorage when no operations are performed within the timeout duration. To perform operations again, the user needs to enter the login username and password to log in to the system.
- **High-risk operation authentication**
High-risk operations, such as deleting a storage pool or unlocking a user, must be authenticated again.
- **Protection against Web application attacks**
The system strictly verifies input and output parameters to prevent cross-site scripting attacks, uses random tokens to prevent cross-site request forgery attacks, and uses PreparedStatement to execute SQL statements to prevent SQL injection attacks.
- **Restriction on file upload and download**
Only administrators have permissions to upload files, and the type and size of the uploaded files are strictly checked. Currently, only exported logs and templates used to import server information can be downloaded, and the directories containing the downloaded files are restricted.

7 Application Scenarios

FusionStorage is typically applied in such scenarios as DevCloud for banking, e-Government Cloud, virtualization for carriers, and FusionCloud private cloud.

DevCloud for Banking

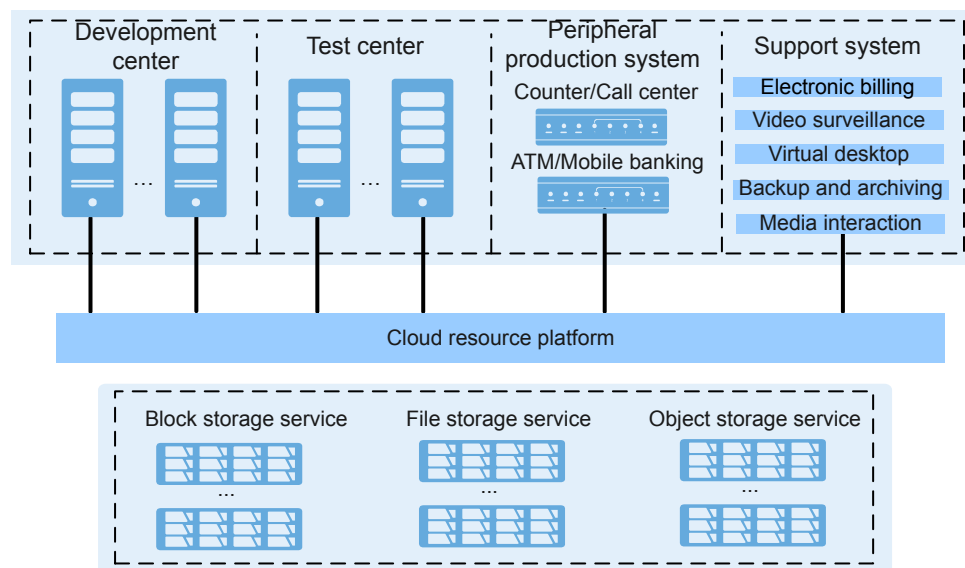
To face challenges presented by new financial organizations and Internet-based financial trends, traditional banks are forced to develop and roll out innovative products as soon as possible, addressing the changing business needs of customers. Traditional banks require storage systems that are capable of:

- Enabling agile resource distribution and reclamation to drive business innovation
- Providing the ability to handle ten times the peak-hour load compared to traditional channels while still ensuring 24x7 availability
- Supporting quick deployment and simplified service distribution procedure for improved efficiency
- Providing high space utilization, concurrent data access and analysis, and the flexible scaling of capacity and performance

To respond to the preceding requirements for storage systems, FusionStorage offers the following advantages in the DevCloud for banking:

- FusionStorage uses x86 servers as its hardware platform to support any type of storage service, including the file storage service, object storage service, and block storage service. These services can be flexibly rolled out or reclaimed when required.
- With excellent performance and solid reliability, FusionStorage is capable of providing stable and secure services even during peak hours.
- FusionStorage can be quickly deployed in the development and test environment, enabling minute-level elastic resource distribution.
- Over 70% of space utilization is ensured. Capacity is flexibly expanded on demand while performance increases linearly.

Figure 7-1 Specific applications in DevCloud for banking



e-Government Cloud

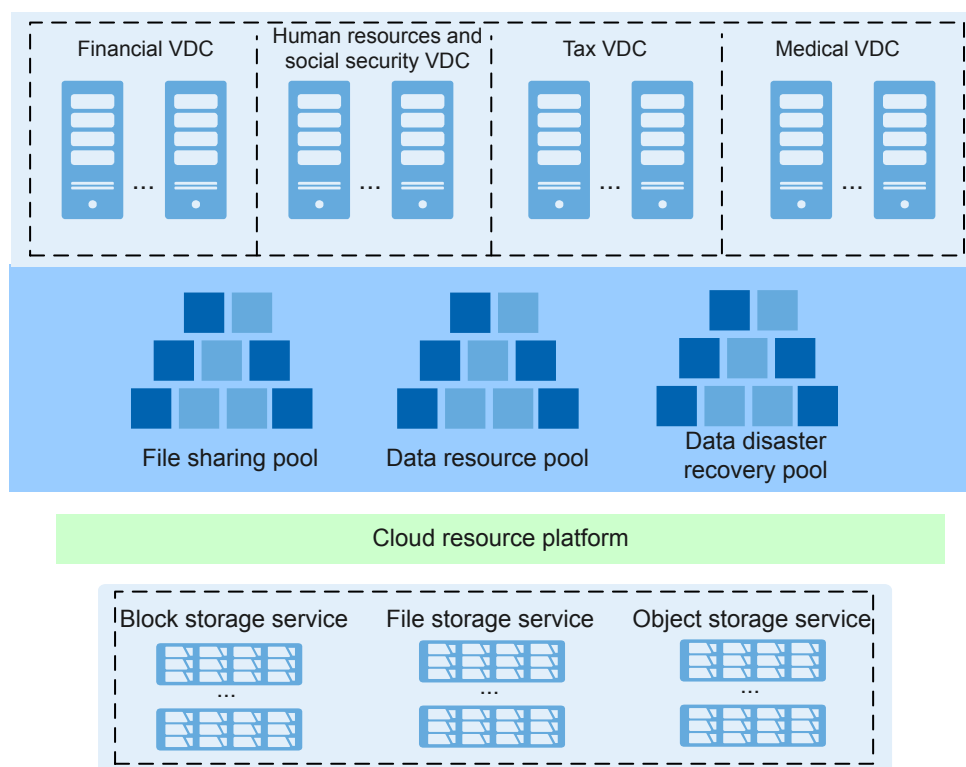
To speed up the information-centric development of governments and help governments provide efficient and convenient services for citizens, developing an e-Government Cloud system is recommended. This cloud system requires storage systems that are capable of:

- Providing one-stop services and allow for data centralization and sharing to maximize investments and facilitate Big Data analysis
- Coping with the complicated interconnection among e-government systems and support fast and flexible capacity expansion as data explosively grows
- Supporting on-demand allocation of storage space and enabling fast service distribution and adjustment

To respond to the preceding requirements for storage systems, FusionStorage offers the following advantages in e-Government Cloud:

- The block storage service perfectly suits four basic core databases and virtual machine applications. Data can be centrally shared to reduce copy operations and capacity can be flexibly scaled when necessary.
- FusionStorage uses x86 servers as its hardware platform to support any type of storage service, including the file storage service, object storage service, and block storage service. These services can be flexibly rolled out or reclaimed when required.
- FusionStorage provides a unified NMS, where the file storage service, object storage service, or block storage service can be distributed at a second level. Furthermore, storage resources are allocated based on specific service levels, addressing the performance and capacity requirements of different services.

Figure 7-2 Specific applications in e-Government Cloud



Virtualization for Carriers

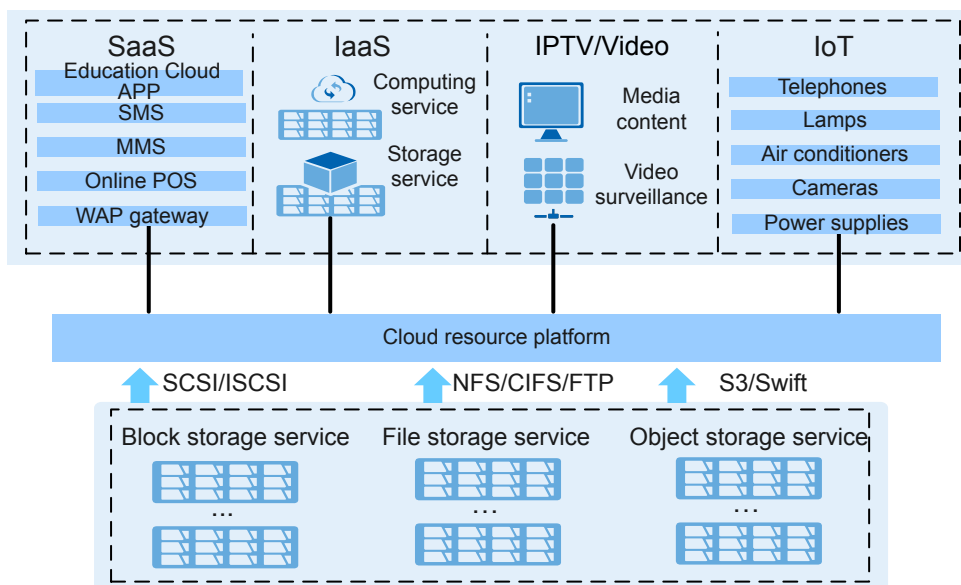
To respond to the challenges posed by emerging IT vendors, traditional telecom carriers are driven to cloudify their networks. Carriers are transforming their services and applications into over the top (OTT) applications, Video On Demand (VOD), Internet of Things (IoT), and public cloud service. Traditional telecom carriers require storage systems that are capable of:

- Providing flexible expansion capabilities to break performance bottlenecks that are generated as storage requirements grow
- Supporting end-to-end platforms for allocating, managing, and maintaining cloud resources, accelerating service distribution and reducing the total cost of ownership (TCO)

To respond to the preceding requirements for storage systems, FusionStorage offers the following advantages in virtualization for carriers:

- A single storage system can build up a massive resource pool whose capacity can be expanded on demand.
- FusionStorage provides a unified NMS, where the file storage service, object storage service, and block storage service can be quickly distributed and rolled out, greatly shortening the time spent on service delivery and capacity expansion.

Figure 7-3 Specific applications in virtualization for carriers



FusionCloud Private Cloud

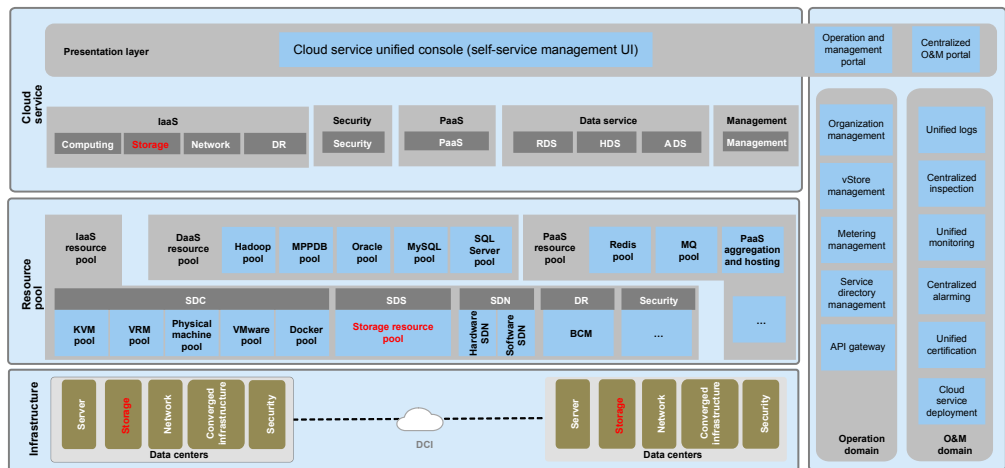
In FusionCloud private cloud scenarios, the object storage service is needed to address challenges in Internet storage posed by massive volumes of data which increases rapidly and unpredictably. This requires storage systems that are capable of:

- Supporting the object storage service while inheriting advantages of block storage and file storage by providing high-speed direct access to disks and distributed shares
- Supporting storage planning and expansion for Internet websites

To respond to the preceding requirements for storage systems, FusionStorage offers the following advantages in FusionCloud private cloud scenarios:

- FusionStorage provides the object storage service and supports storage of unstructured data, such as files, images, and video/audio data.
- FusionStorage supports on-demand linear capacity expansion and provides massive, secure, highly reliable, and cost-effective data storage capabilities.

Figure 7-4 Specific applications in the FusionCloud private cloud



8 Technical Specifications

About This Chapter

Technical specifications of the FusionStorage include reliability specifications and data recovery speed.

8.1 Reliability Specifications

This section describes the reliability specifications of FusionStorage.

8.1 Reliability Specifications

This section describes the reliability specifications of FusionStorage.

Table 8-1 lists the reliability specifications of FusionStorage.

Table 8-1 Reliability specifications

Specifications	Value
Reliability	<p>The redundancy ratio supported by the file storage service can be +1, +2, +3, +4, +2:1, or +3:1. You can set a redundancy ratio for any empty directory.</p> <p>The redundancy ratio supported by the object storage service can be +1, +2, +3, +4, +2:1, or +3:1. You can set the redundancy ratio for any account.</p> <p>The block storage service supports two storage modes: two-copy storage and three-copy storage. The number of copies can be set at your initial configuration and is unchangeable once the system starts providing the block storage service.</p> <p>The redundancy ratio supported by the block storage service can be +1, +2, +3. You can set the redundancy ratio for any storage node or cabinet.</p>
Data recovery speed	Data in the file storage service and object storage service is recovered at a speed of 1 TB/h.

9 Recommended Hardware Configurations

About This Chapter

This chapter describes the recommended hardware configurations of FusionStorage.

9.1 Hardware Components

This section introduces hardware recommended for FusionStorage.

9.2 Storage Nodes

This section describes types of storage nodes providing the FusionStorage file storage service, object storage service, and block storage service.

9.3 Switches

The FusionStorage supports multiple models of switches such as CE6800, CE5800 and SX6018.

9.4 Standard IT Cabinets

It is recommended that the FusionStorage be installed in huawei IT standard cabinets.

9.5 Optional Hardware

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

9.6 Recommended Cabinet Configurations

This section describes typical cabinet configurations of FusionStorage.

9.7 Environmental Specifications

This section describes environmental specifications of FusionStorage.

9.8 Standards Compliance

This chapter lists the protocol standards, safety and electromagnetic compatibility (EMC) standards, and industry standards that FusionStorage complies with. It also lists the certificates obtained by FusionStorage.

9.1 Hardware Components

This section introduces hardware recommended for FusionStorage.

Table 9-1 describes each hardware component of FusionStorage.

Table 9-1 Hardware components

Hardware	Recommended Model	Description
Cabinet	Standard IT cabinet	Providing 42 U (1 U = 44.45 mm) of installation space
Storage node	5288 V3	36-slot node
	RH2288 V3	12- or 25-slot node
	RH2288H V3	12- or 25-slot node
	2288H V5	16-slot node
	Taishan 2280	12-slot node
Switch	Huawei CE6851-48S6Q-HI	10GE switch
	Huawei CE5855-48T4S2Q-EI	GE switch
	SX6018	InfiniBand switch
8-port KVM		Providing eight KVM ports
Modem		Providing SMS-based alarm notification

9.2 Storage Nodes

This section describes types of storage nodes providing the FusionStorage file storage service, object storage service, and block storage service.

9.2.1 Nodes Providing the File Storage Service

The file storage service is provided using RH2288 V3 12-slot nodes or 5288 V3 36-slot nodes.

9.2.1.1 RH2288 V3 12-slot Nodes

Function

Based on requirements of the Internet, Internet Data Center (IDC), cloud computing, enterprise markets, and telecom service applications, Huawei launches the RH2288 V3 12-slot 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, nodes RH2288 V3 12-slot nodes are suitable for storage services including distributed storage, data mining, electronic photo albums, videos, basic applications of enterprises, and telecom service applications.

Exterior

Device

Figure 9-1 shows the exterior of a RH2288 V3 12-slot node overall appearance.

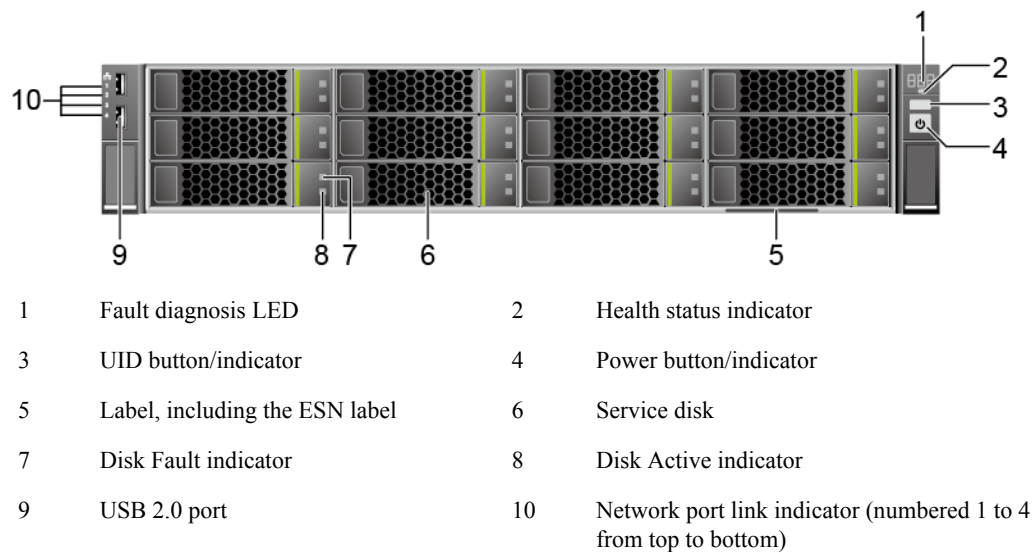
Figure 9-1 Exterior of a RH2288 V3 12-slot node



Front View

Figure 9-2 shows the front panel of a nodeRH2288 V3 12-slot node.


Figure 9-2 Front panel of a RH2288 V3 12-slot node



- | | | | |
|---|--------------------------------|----|--|
| 1 | Fault diagnosis LED | 2 | Health status indicator |
| 3 | UID button/indicator | 4 | Power button/indicator |
| 5 | Label, including the ESN label | 6 | Service disk |
| 7 | Disk Fault indicator | 8 | Disk Active indicator |
| 9 | USB 2.0 port | 10 | Network port link indicator (numbered 1 to 4 from top to bottom) |

The front panel of a RH2288 V3 12-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in [Figure 9-3](#).


Figure 9-3 Disk slots on the front panel of a RH2288 V3 12-slot node



Rear View

[Figure 9-4](#) shows the rear panel of a RH2288 V3 12-slot node.

Figure 9-4 Rear panel of a RH2288 V3 12-slot node



- | | | | |
|----|--|----|--|
| 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 from left to right) | 8 | System disk |
| 9 | Onboard network interface card (NIC) | 10 | UID indicator |
| 11 | USB 3.0 port | 12 | Management network port |
| 13 | VGA port | 14 | Serial port |

Ports

[Table 9-2](#) and [Table 9-3](#) show ports provided by a RH2288 V3 12-slot node.

Table 9-2 Ports on the front panel

Name	Type	Quantity	Description
USB port	USB 2.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>


Table 9-3 Ports on the rear panel

Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
USB port	USB3.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Serial port	DB9	1	<p>A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default baud rate is 115200 bit/s.</p> <p>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.</p>
Network port	-	-	The port types and quantity vary according to the configured NIC.

Physical Structure

Figure 9-5 shows the basic structure of a RH2288 V3 12-slot node.

Figure 9-5 Basic structure of a node



1	System disk	2	PCIe card (on the mainboard)
3	I/O module	4	Internal cable rack
5	Subrack	6	Power module
7	Power backplane	8	Air director
9	Supercapacitor tray	10	Supercapacitor
11	Disk backplane	12	Fan bracket
13	Fan module	14	Service disk
15	Heat radiator	16	DIMM
17	CPU	18	RAID controller card
19	Mainboard	20	NIC

Technical Parameters

Table 9-4 lists the technical parameters of a RH2288 V3 12-slot node.

Table 9-4 Technical parameters of a RH2288 V3 12-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 708 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 30 kg ● Package weight: 5 kg
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 100 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	470 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 64 GB NVDIMM: 8 GB
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● The following two configuration modes are available for service disks: <ul style="list-style-type: none"> - 12 x 3.5-inch service disks (SATA) in slots 0 to 11 - 1 x 3.5-inch service disk (SSD) in slot 0 and 11 x 3.5-inch service disks (SATA) in slots 1 to 11 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.1.2 5288 V3 36-slot Nodes

Function

Based on its x86 hardware structure, FusionStorage pools HDDs, SSDs, and other hardware storage media using distributed techniques to provide an interface complying with the

industry standard for upper-layer applications and clients, as well as fully converged storage capabilities.

Exterior

Device

Figure 9-6 shows the exterior of a 5288 V3 36-slot node.


Figure 9-6 Exterior of a node



Front View

Figure 9-7 shows the front panel of a 5288 V3 36-slot node.


Figure 9-7 Front panel of a 5288 V3 36-slot node



- | | | | |
|---|-----------------------------|----|--------------------------------|
| 1 | Network port link indicator | 2 | USB 2.0 port |
| 3 | Disk Fault indicator | 4 | Disk Active indicator |
| 5 | Service disk | 6 | Fault diagnosis LED |
| 7 | Health status indicator | 8 | UID button/indicator |
| 9 | Power button/indicator | 10 | Label, including the ESN label |

The front panel of a 5288 V3 36-slot node provides 24 service disk slots numbered from 0 to 23 from top to bottom and from left to right, as shown in [Figure 9-8](#).

Figure 9-8 Disk slots on the front panel of a 5288 V3 36-slot node




Rear View

[Figure 9-9](#) shows the rear view of a node5288 V3 36-slot node.

NOTE

This section uses the 10GE NIC as an example.


Figure 9-9 Rear panel of a 5288 V3 36-slot node



1	System disk	2	Disk Fault indicator
3	Disk Active indicator	4	I/O module
5	Power module port	6	Power module indicator
7	Power module	8	Serial port
9	VGA port	10	Management network port
11	USB 3.0 port	12	UID indicator
13	PCIe slot	14	Link indicator
15	Data transfer indicator	16	Service network port
17	Service disk	-	-

The rear panel of a 5288 V3 36-slot node provides 12 service disk slots numbered from 24 to 35 from top to bottom and from left to right, as shown in [Figure 9-10](#).

Figure 9-10 Disk slots on the rear panel of a 5288 V3 36-slot node



Ports

[Table 9-5](#) and [Table 9-6](#) show ports provided by a 5288 V3 36-slot node.

Table 9-5 Ports on the front panel

Name	Type	Quantity	Description
USB port	USB 2.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>


Table 9-6 Ports on the rear panel

Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
USB port	USB3.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Serial port	DB9	1	<p>A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default baud rate is 115200 bit/s.</p> <p>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.</p>
Network port	-	-	The port types and quantity vary according to the configured NIC.

Physical Structure

Figure 9-11 shows the basic structure of a 5288 V3 36-slot node.

Figure 9-11 Basic structure of a 5288 V3 36-slot node



1	System disk	2	I/O module
3	Power module	4	NIC
5	PCIe card	6	RAID controller card
7	DIMM	8	CPU radiator
9	CPU	10	Mainboard
11	Power backplane	12	Service disk
13	Service disk backplane	14	Subrack
15	Air director	16	Fan module
17	Service disk backplane	-	-

Technical Parameters

Table 9-7 lists the technical parameters of a 5288 V3 36-slot node.

Table 9-7 Technical parameters of a 5288 V3 36-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	175 mm x 447 mm x 748 mm (6.89 in. x 17.60 in. x 29.45 in.), 4 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 57 kg (125.69 lb) ● Package weight: 15 kg (33.08 lb)
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 200 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	837 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 208 GB PCIE-SSD: 1.6 TB
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● 1 x 3.5-inch service disk (SSD) in slot 0 ● 35 x 3.5-inch service disks (SATA) in slots 1 to 35 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.2 Nodes Providing the Object Storage Service

The object storage service is provided using RH2288 V3 12-slot nodes or 5288 V3 36-slot nodes.

9.2.2.1 RH2288 V3 12-slot Nodes

Function

Based on requirements of the Internet, Internet Data Center (IDC), cloud computing, enterprise markets, and telecom service applications, Huawei launches the RH2288 V3 12-slot 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, RH2288 V3 12-slot 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 9-12 shows the RH2288 V3 12-slot node overall appearance.


Figure 9-12 RH2288 V3 12-slot node overall appearance



Front view

Figure 9-13 shows the front panel of a RH2288 V3 12-slot node.


Figure 9-13 Front panel



- | | | | |
|---|---|----|--|
| 1 | Fault diagnosis LED | 2 | Health status indicator |
| 3 | UID button/indicator | 4 | Power button/indicator |
| 5 | Label, including the label of equipment serial number (ESN) | 6 | Service disk |
| 7 | Disk Fault indicator | 8 | Disk Active indicator |
| 9 | USB 2.0 port | 10 | Network port link indicator (numbered 1 to 4 from top to bottom) |

The front panel of a RH2288 V3 12-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in [Figure 9-14](#).


Figure 9-14 Disk slots on the front panel of a RH2288 V3 12-slot node



Rear view

[Figure 9-15](#) shows the rear panel of a RH2288 V3 12-slot node.

Figure 9-15 Rear panel of a



- | | | | |
|---|----------------------------------|---|--|
| 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 from left to right)	8	System disk
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 9-8 and **Table 9-9** show ports provided by a RH2288 V3 12-slot node.

Table 9-8 Ports on the front panel

Name	Type	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, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.

Table 9-9 Ports on the rear panel


Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
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 node may operate abnormally if it is connected to an abnormal USB device.
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.

Port	Type	Quantity	Description
Serial port	DB9	1	<p>A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default baud rate is 115200 bit/s.</p> <p>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.</p>
Network port	-	-	The port types and quantity vary according to the configured NIC.

Physical Structure

Figure 9-16 shows the basic structure of a RH2288 V3 12-slot node.

Figure 9-16 Basic structure



5	Subrack	6	Power module
7	Power backplane	8	Air director
9	Supercapacitor tray	10	Supercapacitor
11	Disk backplane	12	Fan bracket
13	Fan module	14	Service disk
15	Heat radiator	16	DIMM
17	CPU	18	RAID controller card
19	Mainboard	20	NIC

Technical Parameters

Table 9-10 lists the technical parameters of a RH2288 V3 12-slot node.

Table 9-10 Technical parameters of a RH2288 V3 12-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 708 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 30 kg ● Package weight: 5 kg
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 100 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	470 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 64 GB NVDIMM: 8 GB

Category	Parameter	Value
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● The following two configuration modes are available for service disks: <ul style="list-style-type: none"> - 12 x 3.5-inch service disks (SATA) in slots 0 to 11 - 1 x 3.5-inch service disk (SSD) in slot 0 and 11 x 3.5-inch service disks (SATA) in slots 1 to 11 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.2.2 5288 V3 36-slot Nodes

Function

Based on its x86 hardware structure, FusionStorage pools HDDs, SSDs, and other hardware storage media using distributed techniques to provide an interface complying with the industry standard for upper-layer applications and clients, as well as fully converged storage capabilities.

Exterior

Entire device

Figure 9-17 shows the entire appearance of a 5288 V3 36-slot node.


Figure 9-17 Appearance



Front view

Figure 9-18 shows the front view of a 5288 V3 36-slot node.

Figure 9-18 Front panel of a 5288 V3 36-slot node




- | | | | |
|---|-----------------------------|---|-----------------------|
| 1 | Network port link indicator | 2 | USB 2.0 port |
| 3 | Disk Fault indicator | 4 | Disk Active indicator |
| 5 | Service disk | 6 | Fault diagnosis LED |
| 7 | Health status indicator | 8 | UID button/indicator |

- 9 Power button/indicator
- 10 Label, including the label of ESN

The front panel of a 5288 V3 36-slot node provides 24 service disk slots numbered from 0 to 23 from top to bottom and from left to right, as shown in **Figure 9-19**.

Figure 9-19 Disk slots on the front panel of a 5288 V3 36-slot node




Rear view

Figure 9-20 shows the rear view of a 5288 V3 36-slot node.

NOTE

This section uses the 10GE NIC as an example.


Figure 9-20 Rear panel of a 5288 V3 36-slot node



- | | | | |
|----|-------------------------|----|-------------------------|
| 1 | System disk | 2 | Disk Fault indicator |
| 3 | Disk Active indicator | 4 | I/O module |
| 5 | Power module port | 6 | Power module indicator |
| 7 | Power module | 8 | Serial port |
| 9 | VGA port | 10 | Management network port |
| 11 | USB 3.0 port | 12 | UID indicator |
| 13 | PCIe slot | 14 | Link indicator |
| 15 | Data transfer indicator | 16 | Service network port |
| 17 | Service disk | - | - |

The rear panel of a 5288 V3 36-slot node provides 12 service disk slots numbered from 24 to 35 from top to bottom and from left to right, as shown in **Figure 9-21**.

Figure 9-21 Disk slots on the rear panel of a 5288 V3 36-slot node



APIs

Table 9-11 and **Table 9-12** show ports provided by a 5288 V3 36-slot node.

Table 9-11 Ports on the front panel

Name	Type	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, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.

Table 9-12 Ports on the rear panel


Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
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 node may operate abnormally if it is connected to an abnormal USB device.

Port	Type	Quantity	Description
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Serial port	DB9	1	A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default 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 9-22](#) shows components of a 5288 V3 36-slot node.

Figure 9-22 Components of a 5288 V3 36-slot node



1	System disk	2	I/O module
3	Power module	4	NIC
5	PCIe card	6	RAID control card
7	Memory	8	CPU radiator
9	CPU	10	Main board
11	Power backplane	12	Service disk
13	Service disk backplane	14	Subrack
15	Air director	16	Fan module
17	Service disk backplane	-	-

Technical Parameters

Table 9-13 lists the technical parameters of a 5288 V3 36-slot node.

Table 9-13 Technical parameters of a 5288 V3 36-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	175 mm x 447 mm x 748 mm (6.89 in. x 17.60 in. x 29.45 in.), 4 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 57 kg (125.69 lb) ● Package weight: 15 kg (33.08 lb)
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 200 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	837 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 80 GB NVDIMM: 8 GB
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● 1 x 3.5-inch service disk (SSD) in slot 0 ● 35 x 3.5-inch service disks (SATA) in slots 1 to 35 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.3 Nodes Providing the Block Storage Service

The block storage service is provided using RH2288 V3 12-slot nodes, RH2288H V3 12-slot nodes, RH2288 V3 25-slot nodes, RH2288H V3 25-slot nodes, or 5288 V3 36-slot nodes.

9.2.3.1 2288H V5 12-slot Nodes

Function

The Huawei 2288H V5 12-slot node is a 2U 2-socket rack server developed for Internet, Internet data center (IDC), cloud computing, enterprise, and telecom service applications. Marked H22H-05 on the nameplate, the 2288H V5 is ideal for IT core services, cloud computing virtualization, high-performance computing, distributed storage, big data processing, enterprise or telecom service applications, and other complex workloads. It combines low power consumption with high scalability and reliability, and easy deployment and management.

Exterior

Entire device

Figure 9-23 shows the 2288H V5 12-slot node overall appearance.

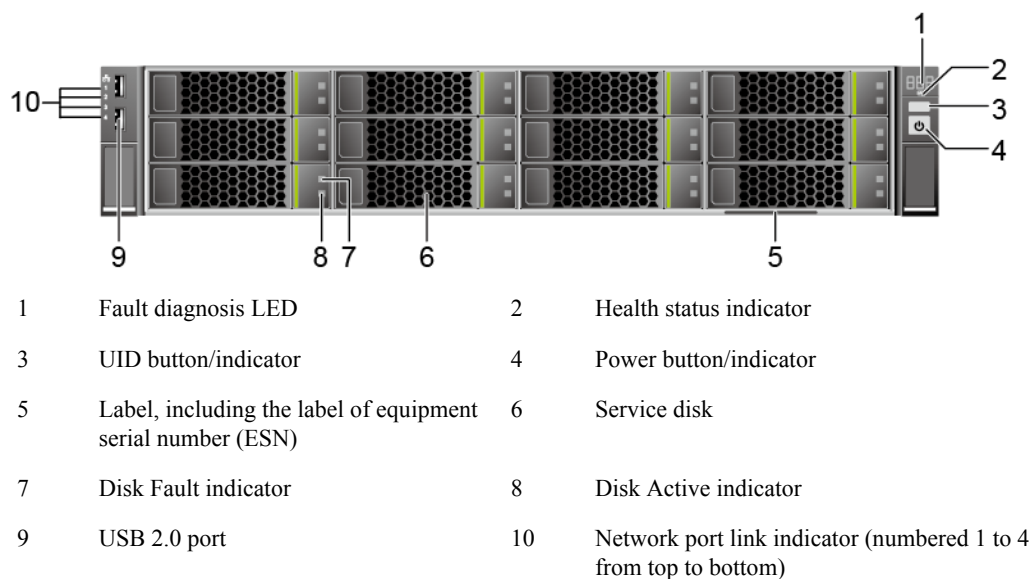
Figure 9-23 2288H V5 12-slot node overall appearance



Front view


Figure 9-24 shows the front panel of a 2288H V5 12-slot node.

Figure 9-24 Front panel of a 2288H V5 12-slot node



The front panel of a 2288H V5 12-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in [Figure 9-25](#).


Figure 9-25 Disk slots on the front panel of a 2288H V5 12-slot node



Rear view

[Figure 9-26](#) shows the rear panel of a 2288H V5 12-slot node.

Figure 9-26 Rear panel of a 2288H V5 12-slot node



1	I/O module 1	2	Data transmission status indicator
3	Connectivity status indicator	4	I/O module 2
5	I/O module 3	6	Power supply unit (PSU) 2
7	PSU sockets	8	PSU indicator
9	PSU 1	10	Flexible NIC (optional)
11	USB 3.0 ports	12	UID indicator
13	Serial port	14	Management network port
15	GE electrical ports	16	VGA port
17	Connection status indicator/Data transmission status indicator	18	10GE electrical ports
19	Transmission rate indicator	20	Transmission rate indicator
21	10GE optical ports	22	Connection status indicator/Data transmission status indicator

APIs

Table 9-14 and **Table 9-15** show ports provided by a 2288H V5 12-slot node.

Table 9-14 Ports on the front panel

Name	Type	Quantity	Description
USB port	USB 2.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE</p> <p>Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>

Table 9-15 Rear Panel


Name	Type	Quantity	Description
PSU socket	—	1 or 2	Used to connect PSUs.
10GE electrical port	10G BASE-T	2	<p>The mainboard provides two 10GE electrical ports or two 10GE optical ports for users to choose.</p> <p>NOTE</p> <ul style="list-style-type: none"> 10GE electrical ports do not support 10 Mbit/s or 100 Mbit/s networks and the rate cannot be forcibly set to 1000 Mbit/s. 10GE optical ports do not support 10 Mbit/s or 100 Mbit/s networks.
10GE optical port	10G SFP+	2	
GE electrical port	GE BASE-T	2	<p>Node service network port</p> <p>NOTE</p> <p>This port does not support forcible rates or 10 Mbit/s and 100 Mbit/s networks.</p>
USB port	USB 3.0	2	<p>The USB ports allow USB devices to be connected to the node.</p> <p>NOTICE</p> <p>Before connecting an external USB device, check that the USB device functions properly. A node may operate abnormally if an abnormal USB device is connected.</p>

Name	Type	Quantity	Description
Management network port	GE BASE-T	1	The 1000 Mbit/s Ethernet port is used for node management.
VGA port	DB15	1	The VGA port is connected to a terminal, such as a monitor or KVM.
Serial port	RJ45	1	The serial port is used as the system serial port by default. You can set it as the iBMC serial port by using the iBMC command. This port is used for debugging.

Physical Structure

Figure 9-27 shows components of a 2288H V5 12-slot node.

Figure 9-27 Components of a 2288H V5 12-slot node



- | | | | |
|---|--------------------|----|---------------------------|
| 1 | I/O module 1 | 2 | I/O module 2 |
| 3 | PSU | 4 | I/O module 3 |
| 5 | Chassis | 6 | Supercapacitor tray |
| 7 | Air duct | 8 | Front hard disk backplane |
| 9 | Fan module bracket | 10 | Fan module |

11	Front hard disk	12	Mainboard
13	Flexible NIC	14	RAID controller card
15	DIMM	16	CPU
17	Heat sink	—	—

Technical Parameters

Table 9-16 lists the technical parameters of a 2288H V5 12-slot node.

Table 9-16 Technical parameters of a 2288H V5 12-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 34.1kg ● Package weight: 5 kg
Environment	Operating temperature	5°C to 35°C
	Operating humidity	8% RH to 90% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● AC: 100 V AC to 240 V AC ● High voltage DC: 192 V DC to 288 V DC
Power consumption	Maximum power consumption	474 W
Example	Processor	Intel Skylake 4109T V5 8-core 2.0 GHz Quantity: 2
	Memory	64 GB
	Disk	<ul style="list-style-type: none"> ● 2 x 3.5-inch system disks (SAS) ● 12 x 3.5-inch service disks (SATA) <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.3.2 2288H V5 16-slot Nodes

Function

The Huawei 2288H V5 16-slot node is a 2U 2-socket rack server developed for Internet, Internet data center (IDC), cloud computing, enterprise, and telecom service applications. Marked H22H-05 on the nameplate, the 2288H V5 is ideal for IT core services, cloud computing virtualization, high-performance computing, distributed storage, big data processing, enterprise or telecom service applications, and other complex workloads. It combines low power consumption with high scalability and reliability, and easy deployment and management.

Exterior

Entire device

Figure 9-28 shows the 2288H V5 16-slot node overall appearance.

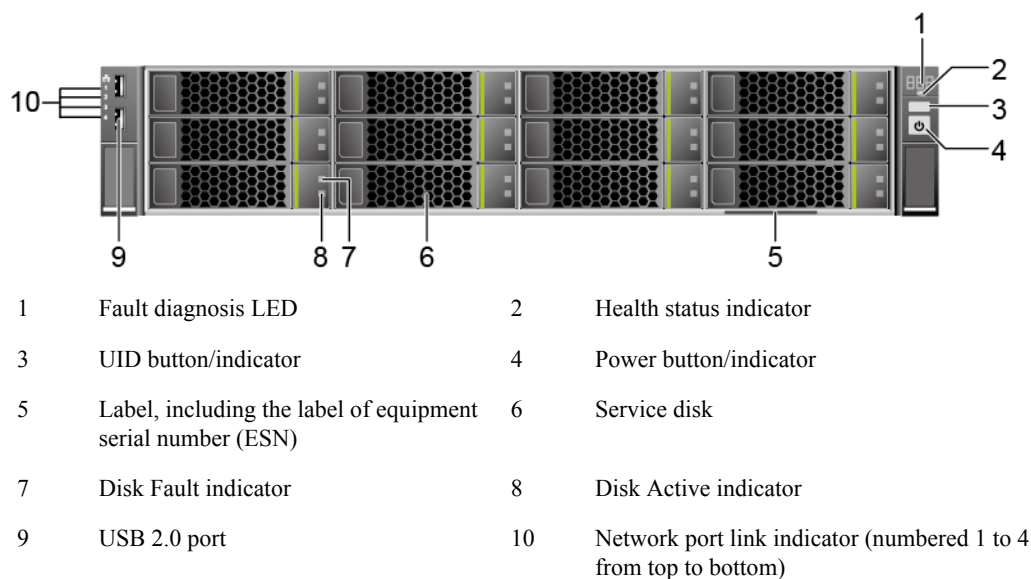
Figure 9-28 2288H V5 16-slot node overall appearance



Front view


Figure 9-29 shows the front panel of a 2288H V5 16-slot node.

Figure 9-29 Front panel of a 2288H V5 16-slot node



The front panel of a 2288H V5 16-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in [Figure 9-30](#). Details about the remaining four built-in disks, see section "Replacing a Built-in Service Disk Module" in *Parts Replacement*.


Figure 9-30 Disk slots on the front panel of a 2288H V5 16-slot node



Rear view

[Figure 9-31](#) shows the rear panel of a 2288H V5 16-slot node.

Figure 9-31 Rear panel of a 2288H V5 16-slot node



1	I/O module 1	2	Data transmission status indicator
3	Connectivity status indicator	4	I/O module 2
5	I/O module 3	6	Power supply unit (PSU) 2
7	PSU sockets	8	PSU indicator
9	PSU 1	10	Flexible NIC (optional)
11	USB 3.0 ports	12	UID indicator
13	Serial port	14	Management network port
15	GE electrical ports	16	VGA port
17	Connection status indicator/Data transmission status indicator	18	10GE electrical ports
19	Transmission rate indicator	20	Transmission rate indicator
21	10GE optical ports	22	Connection status indicator/Data transmission status indicator

APIs

Table 9-17 and **Table 9-18** show ports provided by a 2288H V5 16-slot node.

Table 9-17 Ports on the front panel

Name	Type	Quantity	Description
USB port	USB 2.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE</p> <p>Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>

Table 9-18 Rear Panel


Name	Type	Quantity	Description
PSU socket	—	1 or 2	Used to connect PSUs.
10GE electrical port	10G BASE-T	2	<p>The mainboard provides two 10GE electrical ports or two 10GE optical ports for users to choose.</p> <p>NOTE</p> <ul style="list-style-type: none"> 10GE electrical ports do not support 10 Mbit/s or 100 Mbit/s networks and the rate cannot be forcibly set to 1000 Mbit/s. 10GE optical ports do not support 10 Mbit/s or 100 Mbit/s networks.
10GE optical port	10G SFP+	2	
GE electrical port	GE BASE-T	2	<p>Node service network port</p> <p>NOTE</p> <p>This port does not support forcible rates or 10 Mbit/s and 100 Mbit/s networks.</p>
USB port	USB 3.0	2	<p>The USB ports allow USB devices to be connected to the node.</p> <p>NOTICE</p> <p>Before connecting an external USB device, check that the USB device functions properly. A node may operate abnormally if an abnormal USB device is connected.</p>

Name	Type	Quantity	Description
Management network port	GE BASE-T	1	The 1000 Mbit/s Ethernet port is used for node management.
VGA port	DB15	1	The VGA port is connected to a terminal, such as a monitor or KVM.
Serial port	RJ45	1	The serial port is used as the system serial port by default. You can set it as the iBMC serial port by using the iBMC command. This port is used for debugging.

Physical Structure

Figure 9-32 shows components of a 2288H V5 16-slot node.

Figure 9-32 Components of a 2288H V5 16-slot node



- | | | | |
|---|--------------------|----|---------------------------|
| 1 | I/O module 1 | 2 | I/O module 2 |
| 3 | PSU | 4 | I/O module 3 |
| 5 | Chassis | 6 | Supercapacitor tray |
| 7 | Air duct | 8 | Front hard disk backplane |
| 9 | Fan module bracket | 10 | Fan module |

11	Front hard disk	12	Mainboard
13	Flexible NIC	14	RAID controller card
15	DIMM	16	CPU
17	Heat sink	—	—

Technical Parameters

Table 9-19 lists the technical parameters of a 2288H V5 16-slot node.

Table 9-19 Technical parameters of a 2288H V5 16-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 32.7kg ● Package weight: 5 kg
Environment	Operating temperature	5°C to 35°C
	Operating humidity	8% RH to 90% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● AC: 100 V AC to 240 V AC ● High voltage DC: 192 V DC to 288 V DC
Power consumption	Maximum power consumption	654 W
Example	Processor	Intel Skylake 4114 V5 8-core 2.0 GHz Quantity: 2
	Memory	64 GB
	Disk	<ul style="list-style-type: none"> ● 2 x 3.5-inch system disks (SAS) ● 16 x 3.5-inch service disks (SATA) in slots 0 to 15 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.3.3 Taishan 2280 12-slot Nodes

Function

The Taishan 2280 12-slot node is a 2-socket rack server. It combines high-performance computing with large storage capacity, low power consumption, easy management, and easy deployment, and is ideal for Internet, distributed storage, and cloud computing applications.

Exterior

Entire device

Figure 9-33 shows the Taishan 2280 12-slot node overall appearance.


Figure 9-33 Taishan 2280 12-slot node overall appearance



Front view

Figure 9-34 shows the front panel of a Taishan 2280 12-slot node.


Figure 9-34 Front panel of a Taishan 2280 12-slot node



3	UID button/indicator	4	Power button/indicator
5	Label, including the label of equipment serial number (ESN)	6	Service disk
7	Disk Fault indicator	8	Disk Active indicator
9	USB 2.0 port	10	Network port link indicator (numbered 1 to 4 from top to bottom)

The front panel of a Taishan 2280 12-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in [Figure 9-35](#).


Figure 9-35 Disk slots on the front panel of a Taishan 2280 12-slot node



Rear view

[Figure 9-36](#) shows the rear panel of a Taishan 2280 12-slot node.

Figure 9-36 Rear panel of a Taishan 2280 12-slot node



1	PSU	2	PSU indicator
3	PSU socket	4	I/O module 2
5	Data transmission status indicator for the management port	6	Connectivity status indicator for the management network port
7	Onboard PCIe slot	8	Hard disk activity indicator
9	Hard disk fault indicator	10	I/O module 1
11	Electrical port connectivity status indicator	12	GE electrical port
13	Electrical port data transmission status indicator	14	Optical port connectivity status indicator
15	Optical port data transmission status indicator	16	10GE optical port

17	UID indicator	18	USB 2.0 port
19	Management network port (Mgmt)	20	Video graphics array (VGA) port
21	Serial port	—	—

APIs

Table 9-20 and **Table 9-21** show ports provided by a Taishan 2280 12-slot node.

Table 9-20 Ports on the front panel

Name	Type	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, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.

Table 9-21 Rear Panel


Name	Type	Quantity	Description
PSU socket	—	2	Used to connect PSUs.
GE electrical port	—	2	The mainboard provides GE electrical ports. NOTE When the maximum transmission unit (MTU) of the TaiShan 2280 10GE electrical port is less than 6000 bytes but the MTU of the peer port is greater than 6000 bytes, the TaiShan 2280 cannot receive large packages. In this case, the communication is abnormal.

Name	Type	Quantity	Description
10GE optical port	—	2	<p>The mainboard provides 10GE optical ports.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● 10GE optical ports do not support GE autonegotiation. ● When the MTU of the TaiShan 2280 10GE optical port is less than 6000 bytes but the MTU of the peer port is greater than 6000 bytes, the TaiShan 2280 cannot receive large packages. In this case, the communication is abnormal.
USB port	USB2.0	2	<p>The USB port is connected to a USB device.</p> <p>NOTICE</p> <p>Before connecting to an external USB device, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.</p>
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Video graphics array (VGA) port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
Serial port	DB9	1	The serial port is used as the system serial port by default. You can set it as the iBMC serial port on the iBMC CLI. The port is used for debugging.

Physical Structure

Figure 9-37 shows components of a Taishan 2280 12-slot node.

Figure 9-37 Components of a Taishan 2280 12-slot node



1	I/O module 1	2	PCIe card on the mainboard
3	PCIe card on a riser card	4	I/O module 2
5	Cable management arm (CMA)	6	Chassis
7	PSU	8	PSU backplane
9	Air duct	10	Supercapacitor tray
11	Supercapacitor	12	Front hard disk backplane
13	Fan bracket	14	Fan module
15	Front hard disk	16	Heat sink
17	CPU	18	DIMM
19	SATADOM	20	RAID controller card
21	Mainboard	—	—

Technical Parameters

Table 9-22 lists the technical parameters of a Taishan 2280 12-slot node.

Table 9-22 Technical parameters of a Taishan 2280 12-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm, 2 U (1 U = 44.45 mm)

Category	Parameter	Value
	Weight	<ul style="list-style-type: none"> ● Net weight: 30kg ● Package weight: 5 kg
Environment	Operating temperature	10°C to 35°C
	Operating humidity	10% RH to 90% RH
Power supply	Input voltage	100 V AC to 240 V AC (50Hz/60Hz)
Power consumption	Maximum power consumption	371 W
Example	Processor	Model: Hi1616 32-core 2.4 GHz Quantity: 2
	Memory	DDR4 RDIMM: 80GB PCIE-SSD: 3.2 TB, 1 PCS
	Disk	<ul style="list-style-type: none"> ● 2 x 3.5-inch system disks (SAS) ● 12 x 3.5-inch service disks (SATA) in slots 0 to 11 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.3.4 RH2288 V3 12-slot Nodes

Function

Based on requirements of the Internet, Internet Data Center (IDC), cloud computing, enterprise markets, and telecom service applications, Huawei launches the RH2288 V3 12-slot 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, RH2288 V3 12-slot 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 9-38 shows the RH2288 V3 12-slot node overall appearance.


Figure 9-38 RH2288 V3 12-slot node overall appearance



Front view

Figure 9-39 shows the front panel of a RH2288 V3 12-slot node.


Figure 9-39 Front panel



1	Fault diagnosis LED	2	Health status indicator
3	UID button/indicator	4	Power button/indicator
5	Label, including the label of equipment serial number (ESN)	6	Service disk
7	Disk Fault indicator	8	Disk Active indicator
9	USB 2.0 port	10	Network port link indicator (numbered 1 to 4 from top to bottom)

The front panel of a RH2288 V3 12-slot node provides 12 service disk slots numbered from 0 to 11 from up to down, as shown in **Figure 9-40**.


Figure 9-40 Disk slots on the front panel of a RH2288 V3 12-slot node



Rear view

Figure 9-41 shows the rear panel of a RH2288 V3 12-slot node.

Figure 9-41 Rear panel of a RH2288 V3 12-slot node



- | | | | |
|----|---|----|--|
| 1 | Power module | 2 | Power module indicator |
| 3 | Power socket of the power module | 4 | I/O module (providing slots 6 to 8 from top to bottom) |
| 5 | Link indicator | 6 | Data transfer indicator |
| 7 | Onboard PCIe slot (numbered 4 and 5 from left to right) | 8 | System disk |
| 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 9-23 and **Table 9-24** show ports provided by a RH2288 V3 12-slot node.

Table 9-23 Ports on the front panel

Name	Type	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, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.


Table 9-24 Ports on the rear panel

Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
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 node may operate abnormally if it is connected to an abnormal USB device.
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Serial port	DB9	1	A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default 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 9-42 shows the basic structure of a RH2288 V3 12-slot node.

Figure 9-42 Basic structure



1	System disk	2	PCIe card (on the mainboard)
3	I/O module	4	Internal cable rack
5	Subrack	6	Power module
7	Power backplane	8	Air director
9	Supercapacitor tray	10	Supercapacitor
11	Disk backplane	12	Fan bracket
13	Fan module	14	Service disk
15	Heat radiator	16	DIMM
17	CPU	18	RAID controller card
19	Mainboard	20	NIC

Technical Parameters

Table 9-25 lists the technical parameters of a RH2288 V3 12-slot node.

Table 9-25 Technical parameters of a RH2288 V3 12-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm, 2 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 30kg ● Package weight: 5 kg

Category	Parameter	Value
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 100 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	470 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 128 GB PCIE-SSD: 800GB/1.2 TB/1.6 TB/3.2 TB, 1 to 2 PCS
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● 12 x 3.5-inch service disks (SATA) in slots 0 to 11 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.2.3.5 5288 V3 36-slot Nodes

Function

Based on its x86 hardware structure, FusionStorage pools HDDs, SSDs, and other hardware storage media using distributed techniques to provide an interface complying with the industry standard for upper-layer applications and clients, as well as fully converged storage capabilities.

Exterior

Entire device

Figure 9-43 shows the entire appearance of a 5288 V3 36-slot node.


Figure 9-43 Appearance



Front view

Figure 9-44 shows the front view of a 5288 V3 36-slot node.

Figure 9-44 Front panel of a 5288 V3 36-slot node




- | | | | |
|---|-----------------------------|---|-----------------------|
| 1 | Network port link indicator | 2 | USB 2.0 port |
| 3 | Disk Fault indicator | 4 | Disk Active indicator |
| 5 | Service disk | 6 | Fault diagnosis LED |
| 7 | Health status indicator | 8 | UID button/indicator |

- 9 Power button/indicator
- 10 Label, including the label of ESN

The front panel of a 5288 V3 36-slot node provides 24 service disk slots numbered from 0 to 23 from top to bottom and from left to right, as shown in **Figure 9-45**.

Figure 9-45 Disk slots on the front panel of a




Rear view

Figure 9-46 shows the rear view of a 5288 V3 36-slot node.

NOTE

This section uses the 10GE NIC as an example.


Figure 9-46 Rear panel of a 5288 V3 36-slot node



- | | |
|----------------------------|----------------------------|
| 1 System disk | 2 Disk Fault indicator |
| 3 Disk Active indicator | 4 I/O module |
| 5 Power module port | 6 Power module indicator |
| 7 Power module | 8 Serial port |
| 9 VGA port | 10 Management network port |
| 11 USB 3.0 port | 12 UID indicator |
| 13 PCIe slot | 14 Link indicator |
| 15 Data transfer indicator | 16 Service network port |
| 17 Service disk | - - |

The rear panel of a 5288 V3 36-slot node provides 12 service disk slots numbered from 24 to 35 from top to bottom and from left to right, as shown in **Figure 9-47**.

Figure 9-47 Disk slots on the rear panel of a



APIs

Table 9-26 and **Table 9-27** show ports provided by a 5288 V3 36-slot node.

Table 9-26 Ports on the front panel

Name	Type	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, check that the USB device operates properly. A node may operate abnormally if it is connected to an abnormal USB device.

Table 9-27 Ports on the rear panel


Port	Type	Quantity	Description
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor.
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 node may operate abnormally if it is connected to an abnormal USB device.

Port	Type	Quantity	Description
Management network port (Mgmt)	Ethernet port	1	The 1000 Mbit/s Ethernet port is used to manage the node.
Serial port	DB9	1	A three-pin serial port (that contains only the PIN2 RX, PIN3 TX, and PIN5 GND signals in the DB9 connector, no signals in other pins). The default 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 9-48](#) shows components of a 5288 V3 36-slot node.

Figure 9-48 Components of a 5288 V3 36-slot node



1	System disk	2	I/O module
3	Power module	4	NIC
5	PCIe card	6	RAID control card
7	Memory	8	CPU radiator
9	CPU	10	Main board
11	Power backplane	12	Service disk
13	Service disk backplane	14	Subrack
15	Air director	16	Fan module
17	Service disk backplane	-	-

Technical Parameters

Table 9-28 lists the technical parameters of a 5288 V3 36-slot node.

Table 9-28 Technical parameters of a 5288 V3 36-slot node

Category	Parameter	Value
Size and weight	Dimensions (H x W x D)	175 mm x 447 mm x 748 mm (6.89 in. x 17.60 in. x 29.45 in.), 4 U (1 U = 44.45 mm)
	Weight	<ul style="list-style-type: none"> ● Net weight: 57 kg (125.69 lb) ● Package weight: 15 kg (33.08 lb)
Environment	Operating temperature	5°C to 35°C
	Operating humidity	20% RH to 80% RH
Power supply	Input voltage	<ul style="list-style-type: none"> ● 1200 W high-voltage AC: The power input is 200 V AC to 240 V AC ● 240 V high-voltage DC: The power input is 192 V DC to 288 V DC
Power consumption	Maximum power consumption	837 W
Example	Processor	Model: Intel Xeon 8-core 2.1 GHz Quantity: 2
	Memory	DDR4 RDIMM: 208 GB PCIE-SSD: 1.6 TB
	Disk	<ul style="list-style-type: none"> ● 2 x 2.5-inch system disks (SAS) ● 36 x 3.5-inch service disks (SATA) in slots 0 to 35 <p>System disks must be of the same disk type and capacity. 4KN disk and 512N/512e disks cannot serve as system disks at the same time.</p>

9.3 Switches

The FusionStorage supports multiple models of switches such as CE6800, CE5800 and SX6018.

9.3.1 CE6851-48S6Q-HI

CE6851-48S6Q-HI switches provide 10GE access.

Function


CE6851-48S6Q-HI can provide forty-eight 10GE SFP+ Ethernet optical ports and six 40GE QSFP+ Ethernet optical ports for node interconnection and communication. Each 40GE QSFP+ Ethernet optical port can be divided into four 10GE ports.

Exterior

Front View

Figure 9-49 shows the front view of the switch.

Figure 9-49 Front view of the switch (housing AC power modules)




1	Power module 1	2	Fan module 1
3	Console port	4	ETH management network port
5	USB port	6	Fan module 2
7	Power module 2	-	-

Rear View

Figure 9-50 shows the rear view of the switch.

Figure 9-50 Rear view of the switch



1	48 x 10GE SFP+ Ethernet optical ports	2	6 x 40GE QSFP+ Ethernet optical ports
---	---------------------------------------	---	---------------------------------------

Ports

Table 9-29 lists the ports of the switch.

Table 9-29 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 9-51 shows the numbering of 10GE SFP+ Ethernet optical ports on a CE6851-48S6Q-HI switch.

Figure 9-51 10GE SFP+ Ethernet optical ports on a CE6851-48S6Q-HI switch



Technical Specifications

Table 9-30 lists the technical specifications of the switch.

Table 9-30 Technical specifications of the switch

Item	Value	
Physical specifications	Dimensions (W x D x H)	442.0 mm x 420.0 mm x 43.6 mm (17.4 in. x 16.5 in. x 1.72 in.)
	Fan quantity	2
	Max. power consumption	245 W
	Weight	8.7 kg
System specifications	Processor	1.2 GHz, quad-core.
	DRAM memory	2 GB
	NOR flash	16 MB
	NAND flash	1 GB

Item		Value
Power specifications	Rated input voltage range	200 V AC to 240 V AC; 50 Hz/60 Hz
	Maximum input voltage range	90 V AC to 290 V AC; 45 Hz to 65 Hz
Environmental specifications	Temperature	<ul style="list-style-type: none"> Operating temperature: 0°C to 40°C (0 m to 1800 m). <p>NOTE When the altitude is between 1800 m and 5000 m, the highest operating temperature reduces 1°C every time the altitude increases 220 m.</p> <ul style="list-style-type: none"> Storage temperature: - 40°C to +70°C.
	Relative humidity	5% RH to 95% RH (non-condensing)
	Altitude	< 5000 m

9.3.2 CE6855-48S6Q-HI

CE6855-48S6Q-HI switches provide 10GE access.

Function


CE6855-48S6Q-HI can provide forty-eight 10GE SFP+ Ethernet optical ports and six 40GE QSFP+ Ethernet optical ports for node interconnection and communication. Each 40GE QSFP+ Ethernet optical port can be divided into four 10GE ports.

Exterior

Front View

Figure 9-52 shows the front view of the switch.

Figure 9-52 Front view of the switch (housing AC power modules)




- | | | | |
|---|----------------|---|-----------------------------|
| 1 | Power module 1 | 2 | Fan module 1 |
| 3 | Console port | 4 | ETH management network port |
| 5 | USB port | 6 | Fan module 2 |

7 Power module 2 - -

Rear View

Figure 9-53 shows the rear view of the switch.

Figure 9-53 Rear view of the switch



Ports

Table 9-31 lists the ports of the switch.

Table 9-31 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 9-54 shows the numbering of 10GE SFP+ Ethernet optical ports on a CE6855-48S6Q-HI switch.

Figure 9-54 10GE SFP+ Ethernet optical ports on a CE6855-48S6Q-HI switch



Technical Specifications

Table 9-32 lists the technical specifications of the switch.

Table 9-32 Technical specifications of the switch

Item		Value
Physical specifications	Dimensions (W x D x H)	442.0 mm x 420.0 mm x 43.6 mm (17.4 in. x 16.5 in. x 1.72 in.)
	Fan quantity	2
	Max. power consumption	216 W
	Weight	8.7 kg
System specifications	Processor	1.2 GHz, quad-core.
	DRAM memory	2 GB
	NOR flash	16 MB
	NAND flash	1 GB
Power specifications	Rated input voltage range	100 V AC to 240 V AC; 50 Hz/60 Hz
	Maximum input voltage range	90 V AC to 290 V AC; 45 Hz to 65 Hz
Environmental specifications	Temperature	<ul style="list-style-type: none"> ● Operating temperature: 0°C to 40°C (0 m to 1800 m). <p>NOTE When the altitude is between 1800 m and 5000 m, the highest operating temperature reduces 1°C every time the altitude increases 220 m.</p> <ul style="list-style-type: none"> ● Storage temperature: - 40°C to +70°C.
	Relative humidity	5% RH to 95% RH (non-condensing)
	Altitude	< 5000 m

9.3.3 CE5855-48T4S2Q-EI

CE5855-48T4S2Q-EI switches provide GE access.

Function

CE5855-48T4S2Q-EI (CE5855 for short) is a device that connects servers and enables data communication and active/standby configuration for connected components.

Exterior

Front View

Figure 9-55 shows the front panel of a CE5855 switch.

Figure 9-55 Front panel of a CE5855 switch equipped with AC power modules




1	Power module 1	2	Fan module 1
3	Console port	4	ETH management network port
5	USB port	6	Fan module 2
7	Power module 2	-	-

Rear View

Figure 9-56 shows the rear panel of a CE5855 switch.

Figure 9-56 Rear panel of a CE5855 switch



1	48 x 10/100/1000BASE-T Ethernet electrical ports	2	4 x 10GE SFP+ Ethernet optical ports
3	2 x 40GE QSFP+ Ethernet optical ports	-	-

Ports

Table 9-33 lists the ports of a CE5855 switch.


Table 9-33 Ports of a CE5855 switch

Port Name	Description
10/100/1000 Base-T Ethernet port	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
10GE SFP+ Ethernet optical ports	Used for receiving and sending 10GE services

Port Name	Description
40GE QSFP+ Ethernet optical ports	Used for receiving and sending 40GE services

Figure 9-57 shows the numbering of the forty-eight 10/100/1000 Base-T Ethernet ports on a CE5855 switch.

Figure 9-57 Ethernet ports on a CE5855 switch



Technical Specifications

Table 9-34 lists the technical specifications of a CE5855 switch.

Table 9-34 Technical specifications of a CE5855 switch

Item	Specifications	
Physical specifications	Dimensions (W x D x H)	442.0 mm x 420.0 mm x 43.6 mm (17.4 in. x 16.54 in x 1.72 in.)
	Fan quantity	2
	Max. power	103 W
	Max. weight	8.4 kg
System specifications	Processor	1.2 GHz, dual-core.
	DRAM memory	2 GB
	NOR flash	16 MB
	NAND flash	1 GB
Power specifications	Rated input voltage range	100 V AC to 240 V AC; 50 Hz/60 Hz
	Maximum input voltage range	90 V AC to 290 V AC; 45 Hz to 65 Hz

Item		Specifications
Environmental specifications	Temperature	<ul style="list-style-type: none"> Operating temperature: 0°C to 40°C (0 m to 1800 m). <p>NOTE When the altitude is between 1800 m and 5000 m, the highest operating temperature reduces 1°C every time the altitude increases 220 m.</p> <ul style="list-style-type: none"> Storage temperature: - 40°C to +70°C.
	Relative humidity	5% RH to 95% RH (non-condensing)
	Altitude	< 5000 m

9.4 Standard IT Cabinets

It is recommended that the FusionStorage be installed in huawei IT standard cabinets.


Function

An huawei standard IT 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 the huawei standard IT cabinet is in sand texture black. [Figure 9-58](#) shows exterior of the cabinet.


Figure 9-58 Exterior of the cabinet



Structure

The cabinet consists of the racks, front doors, back doors, side panels, cover, and mounting bars, as shown in [Figure 9-59](#).

Figure 9-59 Cabinet structure



1	Cover	2	Rear door
3	Side panel	4	Mounting bar
5	Rack	6	Front door

Technical Specifications

[Table 9-35](#) lists the technical specifications of the cabinet.

Table 9-35 Technical specifications of the 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
Weight	<ul style="list-style-type: none">● 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 temperature	<ul style="list-style-type: none">● Long term: 0°C to 50°C● Short term: - 5°C to +50°C
Operating humidity	<ul style="list-style-type: none">● 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.

9.5 Optional Hardware

SMS modems and keyboard, video, and mouse (KVM) devices can be configured for the FusionStorage 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 9-60 shows the exterior of an SMS modem.

Figure 9-60 Exterior of an SMS modem



KVM

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

Figure 9-61 shows the exterior of a KVM.

Figure 9-61 Exterior of a KVM



Table 9-36 lists the technical specifications of a KVM.

Table 9-36 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

Parameter	Value
Type	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

9.6 Recommended Cabinet Configurations

This section describes typical cabinet configurations of FusionStorage.

The number of nodes in a fully configured cabinet depends on the power supply capability of the cabinet. Different power supplies support different typical cabinet configurations, as listed in [Table 9-37](#).

Table 9-37 Typical cabinet configurations

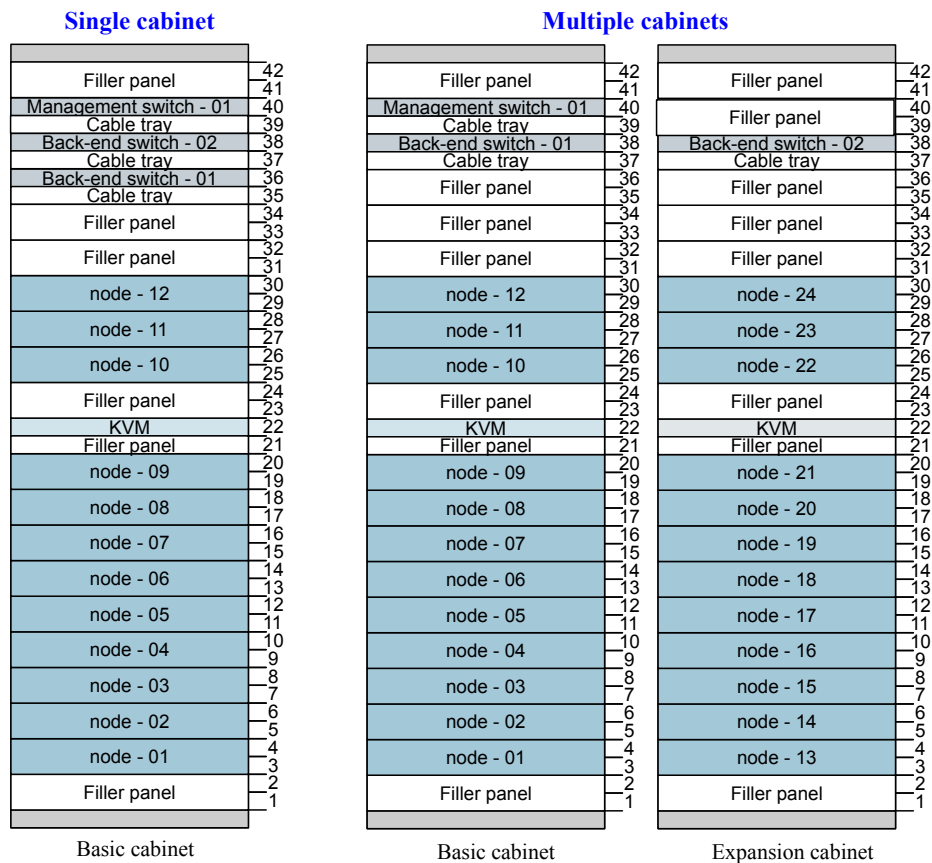
Node Type	220 V/63 A Power Supply and 10GE Networking				220 V/63 A Power Supply and IB Networking				220 V/32 A Power Supply and 10GE Networking				220 V/32 A Power Supply and IB Networking			
	Inter-cabinet		Intra-cabinet		Inter-cabinet		Intra-cabinet	Inter-cabinet		Intra-cabinet		Inter-cabinet		Intra-cabinet		
	First cabinet	Non-first cabinet	First cabinet	Non-first cabinet	First cabinet	Non-first cabinet		First cabinet	Non-first cabinet	First cabinet	Non-first cabinet	First cabinet	Non-first cabinet			
12-slot node (2 U)	12	12	12	12	12	12	12	12	12	12	12	12	12	12	12	
25-slot node (2 U)	12	12	12	12	12	12	12	10	11	10	11	10	11	10	10	

36-slot node (4 U)	8	8	8	8	8	8	8	8	7	7	6	7	7	7	7
--------------------	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

Cabinet Configurations for Scenarios Where the 220 V/63 A Power Supply, 10GE Networking, and 25- or 12-Slot Storage Nodes Are Adopted

Figure 9-62 shows typical single- and multi-cabinet configurations.

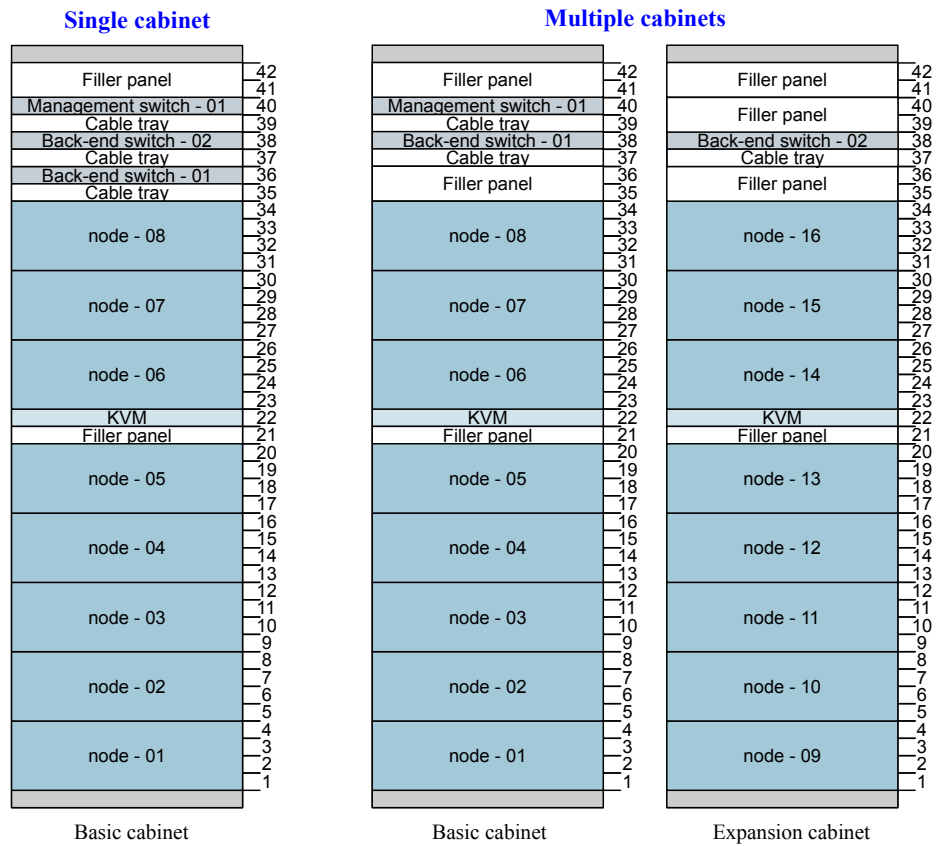
Figure 9-62 Single- and multi-cabinet configurations for 25- and 12-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/63 A Power Supply, 10GE Networking, and 36-Slot Storage Nodes Are Adopted

Figure 9-63 shows typical single- and multi-cabinet configurations.

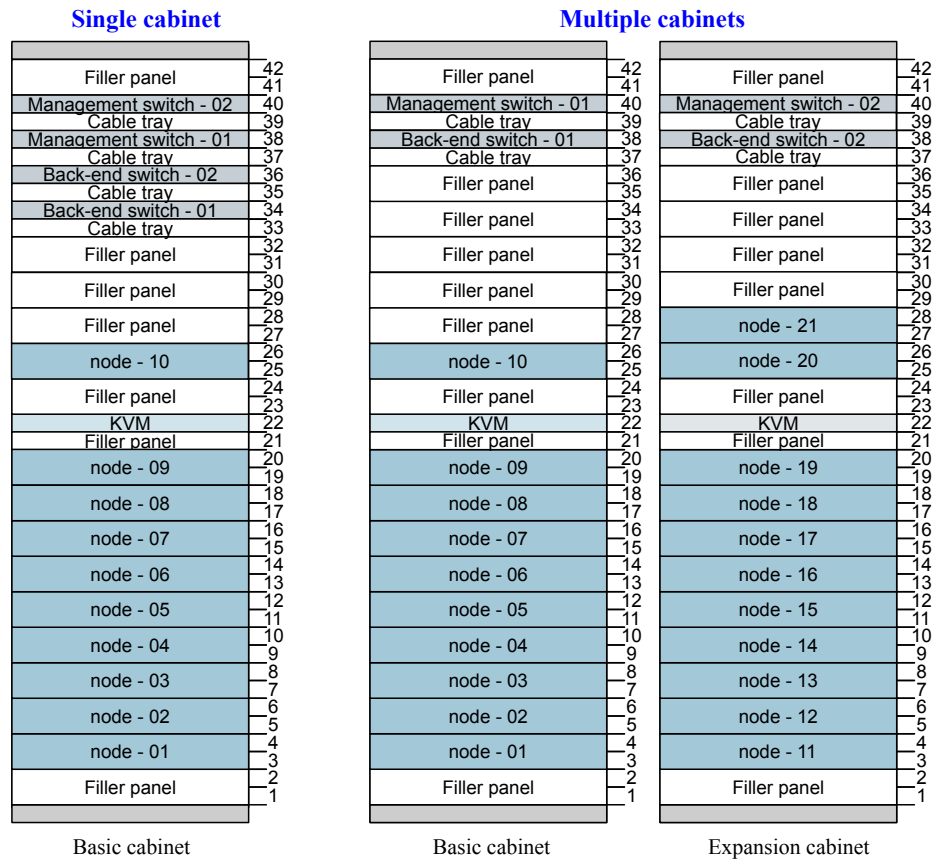
Figure 9-63 Single- and multi-cabinet configurations for 36-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/32 A Power Supply, 10GE Networking, and 25-Slot Storage Nodes Are Adopted

Figure 9-64 shows typical single- and multi-cabinet configurations.

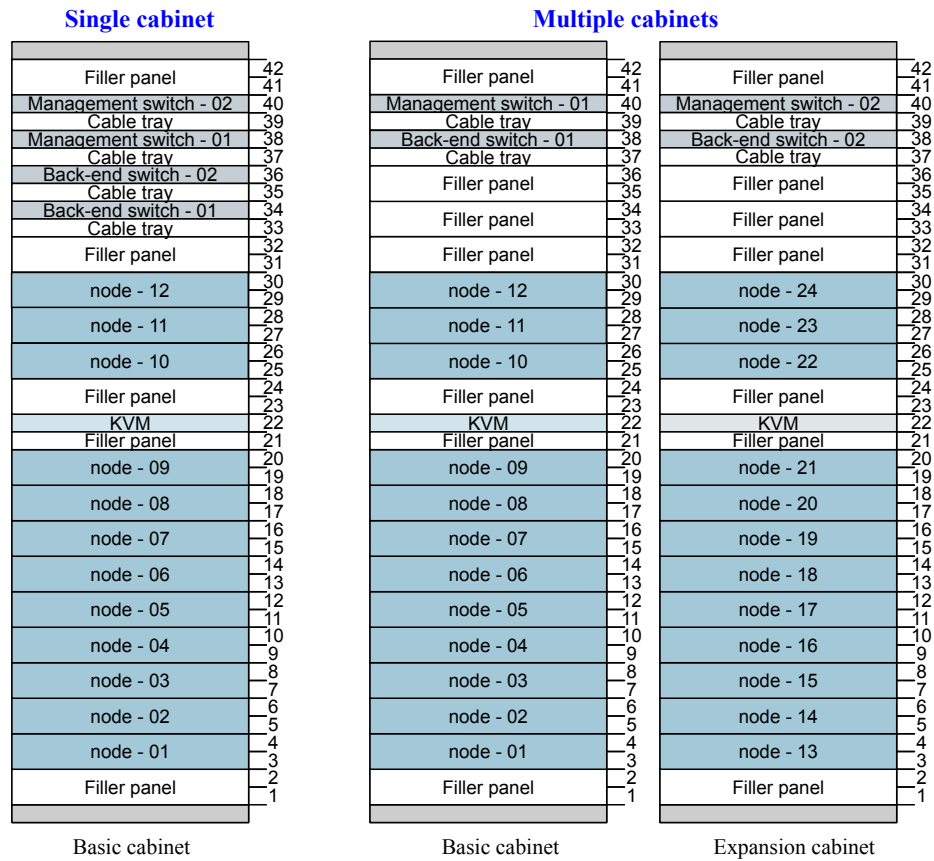
Figure 9-64 Single- and multi-cabinet configurations for 25-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/32 A Power Supply, 10GE Networking, and 12-Slot Storage Nodes Are Adopted

Figure 9-65 shows typical single- and multi-cabinet configurations.

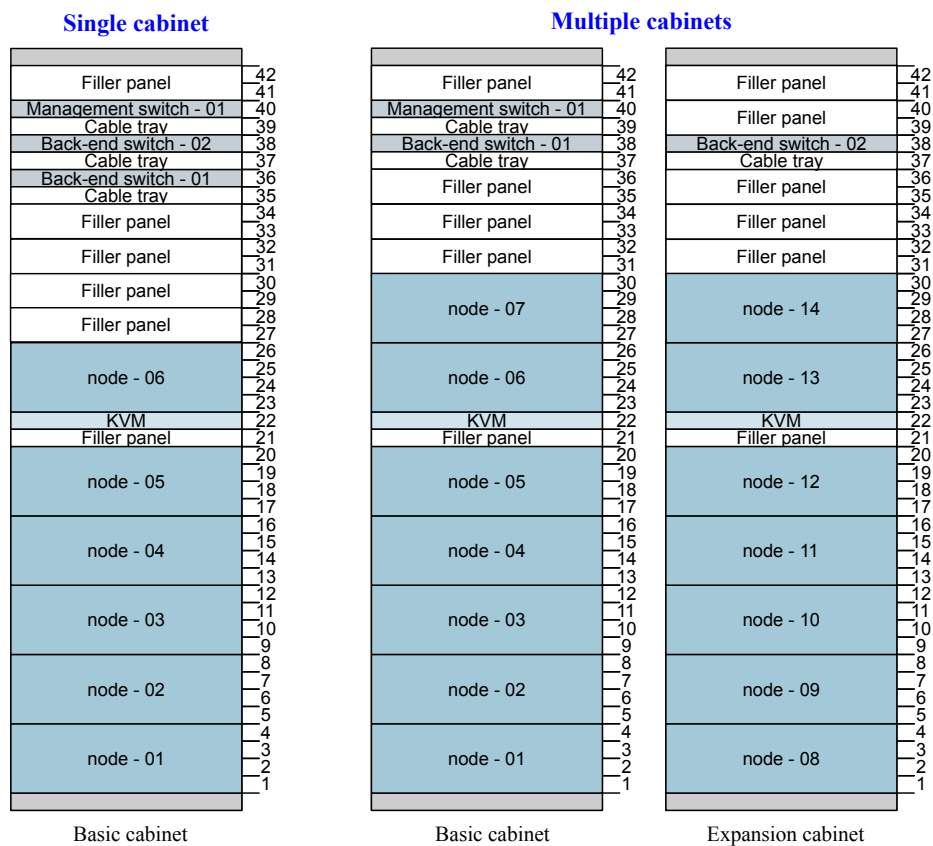
Figure 9-65 Single- and multi-cabinet configurations for 12-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/32 A Power Supply, 10GE Networking, and 36-Slot Storage Nodes Are Adopted

Figure 9-66 shows typical single- and multi-cabinet configurations.

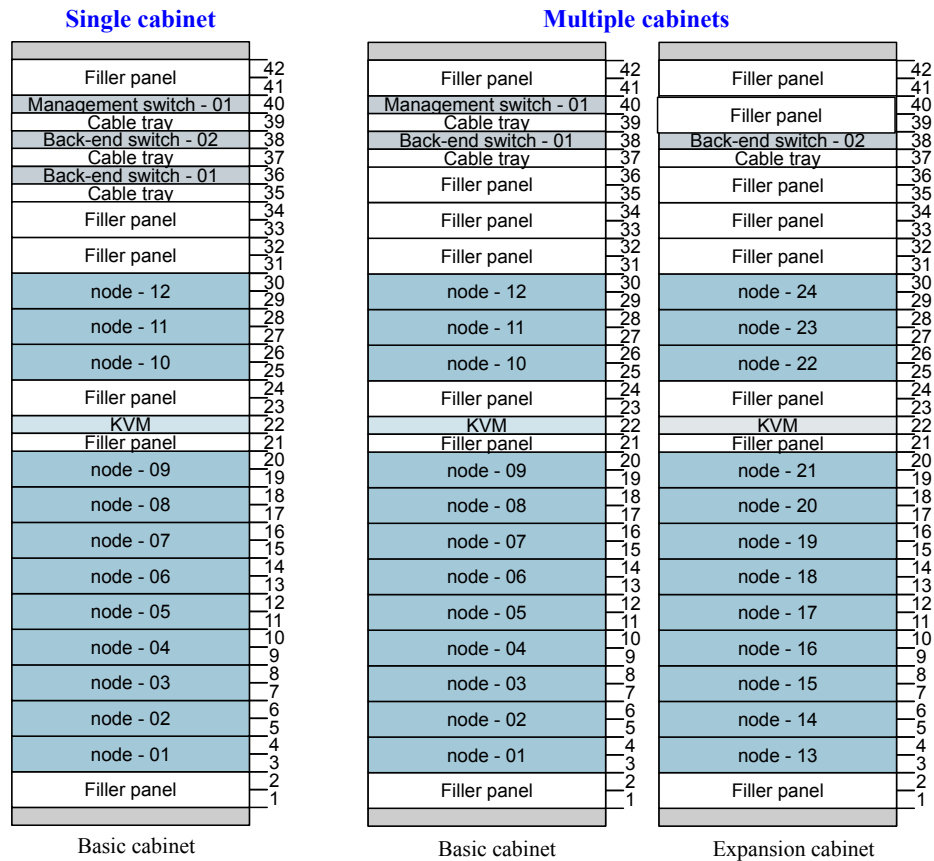
Figure 9-66 Single- and multi-cabinet configurations for 36-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/63 A Power Supply, IB Networking, and 25- or 12-Slot Storage Nodes Are Adopted

Figure 9-67 shows typical single- and multi-cabinet configurations.

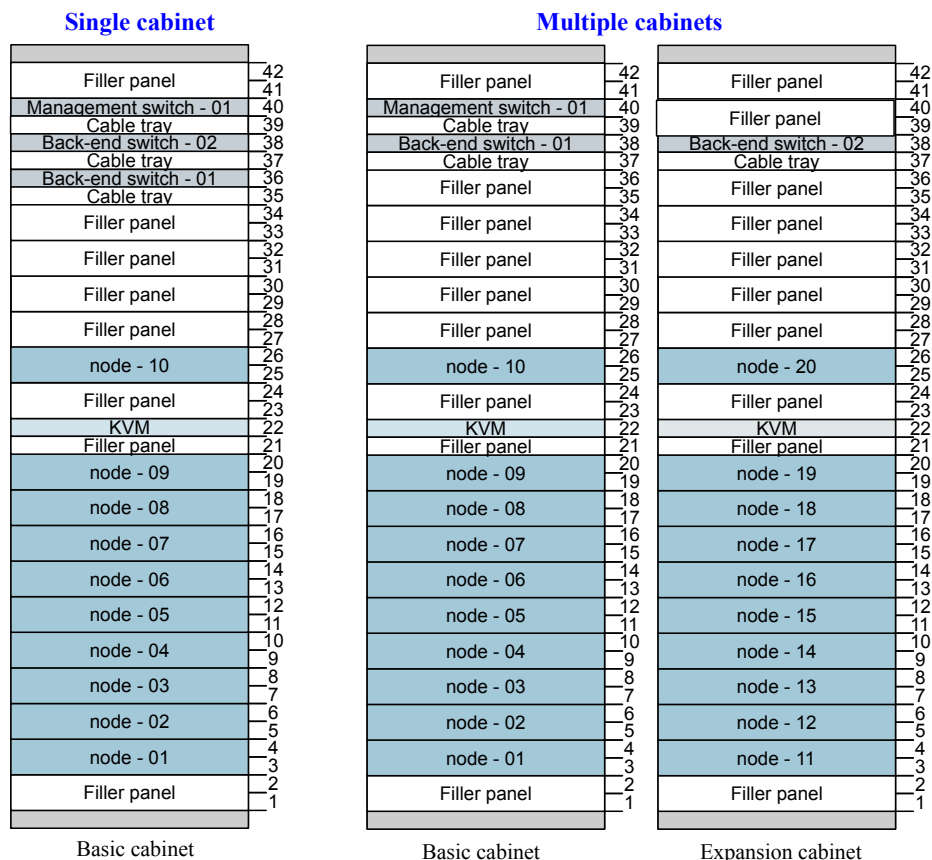
Figure 9-67 Single- and multi-cabinet configurations for 25- and 12-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/32 A Power Supply, IB Networking, and 25-Slot Storage Nodes Are Adopted

Figure 9-68 shows typical single- and multi-cabinet configurations.

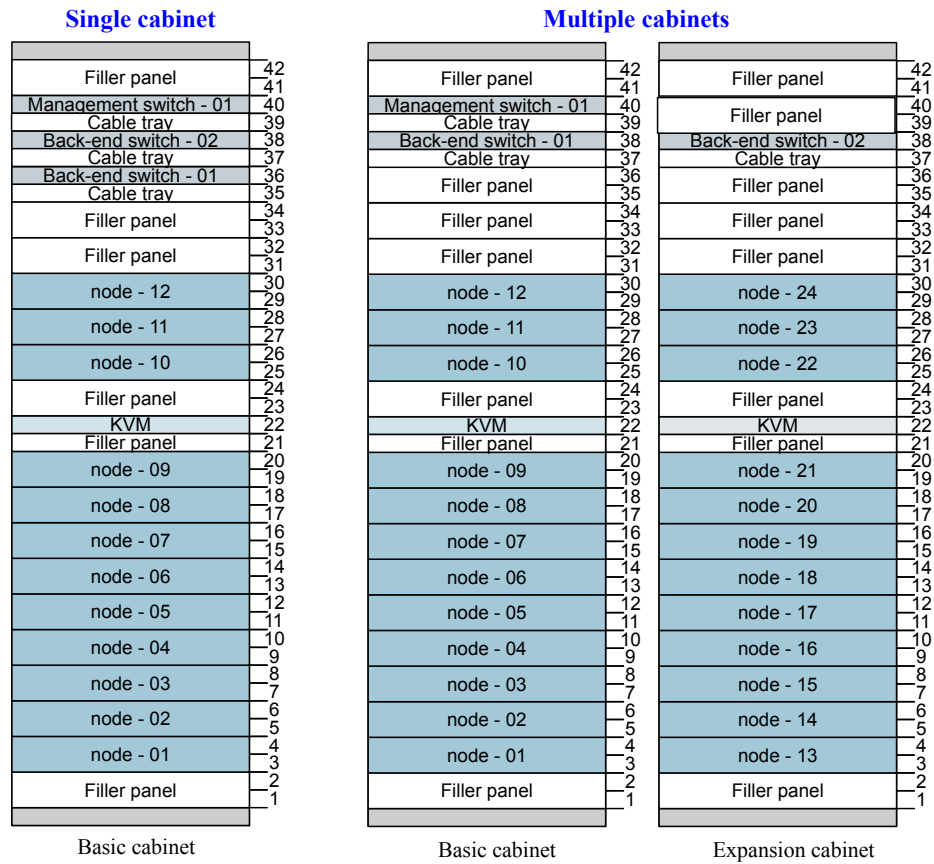
Figure 9-68 Single- and multi-cabinet configurations for 25-slot storage nodes



Cabinet Configurations for Scenarios Where the 220 V/32 A Power Supply, IB Networking, and 12-Slot Storage Nodes Are Adopted

Figure 9-69 shows typical single- and multi-cabinet configurations.

Figure 9-69 Single- and multi-cabinet configurations for 12-slot storage nodes



9.7 Environmental Specifications

This section describes environmental specifications of FusionStorage.

Table 9-38 lists the environmental specifications of FusionStorage.

Table 9-38 Environmental specifications

Specifications	Value
Operating temperature ^a	5°C to 35°C
Storage temperature	-40°C to +65°C -40°C to +70°C
Temperature gradient	< 20°C/hour
Long-time storage temperature	21°C to 27°C
Operating humidity	8% RH to 90% RH (non-condensing) 20% RH to 80% RH (non-condensing)

Specifications	Value
Storage humidity	5% RH to 95% RH (non-condensing)
Humidity gradient	< 20% RH/hour
Long-time storage humidity	30% RH to 69% RH (non-condensing)
Altitude	≤ 3000 m
Noise ^b	< 72dBA
a: When the system runs at an altitude lower than 900 m (2952.72 ft.), the operating temperature ranges from 5°C to 35°C (41°F to 95°F). When the system runs at an altitude higher than 900 m, the operating temperature decreases by 1°C (33.8°F) every time the altitude increases by 300 m (984.24 ft.).	
b: the maximum noise generated by a storage node when the operating temperature is 23°C (73.4°F)	

9.8 Standards Compliance

This chapter lists the protocol standards, safety and electromagnetic compatibility (EMC) standards, and industry standards that FusionStorage complies with. It also lists the certificates obtained by FusionStorage.

Protocol Standards

Table 9-39 lists the protocol standards that FusionStorage complies with.

Table 9-39 Protocols standards

Name	Standard 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
S3	Amazon Simple Storage Service

Safety and EMC Standards

Table 9-40 lists the safety and EMC standards that FusionStorage complies with.

Table 9-40 Safety and EMC standards

Name	Standard No.
Standards of the Information Technology Equipment Safety	GB4943-2001
International Electrotechnical Commission (IEC) standards	IEC 60950-1
North America safety standard	UL 60950-1
US EMC standards	FCC, 47 CFR Part 15, Subpart B
European safety standards	EN 60950-1
European EMC standards	EN 55024: 1998+A1+A2

Industry Standards

Table 9-41 lists the industry standards that FusionStorage complies with.

Table 9-41 Industry standards

Name	Standard No.
Ethernet standards	IEEE 802.3
Fast Ethernet standards	IEEE 802.3u
GE standards	IEEE 802.3z
IEEE standard test port and boundary-scan architecture	IEEE 1149.1-2001
Failure mode and effects analysis (FMEA)	IEC 812
Reliability, maintainability and availability standards	IEC 863
Environmental protection standards	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

Certification

Table 9-42 lists the certificates obtained by FusionStorage.

Table 9-42 Certificates

Certificate Name	Description
CB	The IEC System for Conformity Testing and Certification of Electrical Equipment (IECEE) is based on the use of specific IEC standards for electrical equipment. The Certification Bodies (CB) Scheme is applicable to electrical equipment within the scope of IEC standards for safety, accepted for use in the IECEE. The Scheme becomes operative for such standards as soon as at least one National Certification Body has declared their recognition of CB Test Certificates. The CB scheme is designed for eliminating the international commerce barriers resulting from compliance with certifications and approval guidelines of different countries.
CCC	The China Compulsory Certificate (CCC) mainly involves products related to human health and security, animal and plant life and health, environmental protection, and public security.
CE	Conformité Européenne (CE) refers to the certification required for products to be sold in Europe. 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	The C-Tick is an identification trademark registered to the Australian Communications and Media Authority (ACMA), which signifies compliance with applicable EMC standards and also provides a traceable link between the equipment and the supplier.
FCC	The Federal Communications Commission (FCC) proves that this equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules.
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	Registration, Evaluation, Authorization and Restriction of Chemicals (REACH) is the European Community Regulation on chemicals and their safe use (EC 1907/2006). The law entered into force on 1 June 2007.
RoHS	Restriction of 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 amount of hazardous substances arising from production.
SASO	The Saudi Arabian Standards Organization (SASO), authorized by Saudi Arabian government, organizes and implements International Conformity Certification Program (ICCP) for market access. Products exported to Saudi Arabia must meet SASO's market access requirements and obtain the COC certificate before custom clearance.

Certificate Name	Description
SONCAP	As required by Standards Organization of Nigeria, products exported to Nigeria must be certified by SON Conformity Assessment Programme (SONCAP) before custom clearance.
UL	The Underwriters Laboratories (UL) is a non-profit agency engaged in product safety testing.
VCCI	Voluntary Control Council for Interference (VCCI) is a Japanese organization governing electromagnetic interference.
WEEE	The EU Directive on Waste of Electric and Electronic Equipment. Electrical and electronic products sold in the EU market must comply with this directive and bear the symbol of a crossed out wheelie bin.
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 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 unless the equipment is proven 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 related to energy saving, security, and EMC throughout the whole process in the area of product design, product manufacture, production waste disposal, reclamation, and recycling (from cradle to cradle).