OpenStor

Computational Storage
2910

FW version 6.1.7

Product Description

Release date 2024 - 01 – 31

Contents

1. Product Positioning

Challenges to Traditional Data Centers

As the core infrastructure of the digital economy, the data center is deeply integrated into the digital transformation process with the rapid development of new-generation information technologies.

However, traditional data centers face the following problems and challenges:

- High investment cost

  In traditional solutions, computing, storage, and network resources are separated. Therefore, independent storage devices and switches need to be purchased, resulting in high investment costs.

- Slow service rollout

  The traditional service rollout process includes hardware selection, multi-vendor procurement, batch delivery, installation, and joint commissioning. The rollout process is complex and slow.

- Complex O&M

  There are many infrastructure vendors, and configuration management is independent and scattered. Users need to spend a lot of manpower on maintaining storage, computing, and network devices.

OS 2910

The OS 2910 is a new IT infrastructure platform based on the hyper-converged technology. It integrates computing, storage, and network resources into one device and implements visualized management of IT resources through a unified management platform.
The OS 2910 adopts the concepts of "in-depth vertical integration and optimization of software and hardware", "fast deployment", and "unified management" to improve the operation efficiency of core services and reduce the overall investment cost.

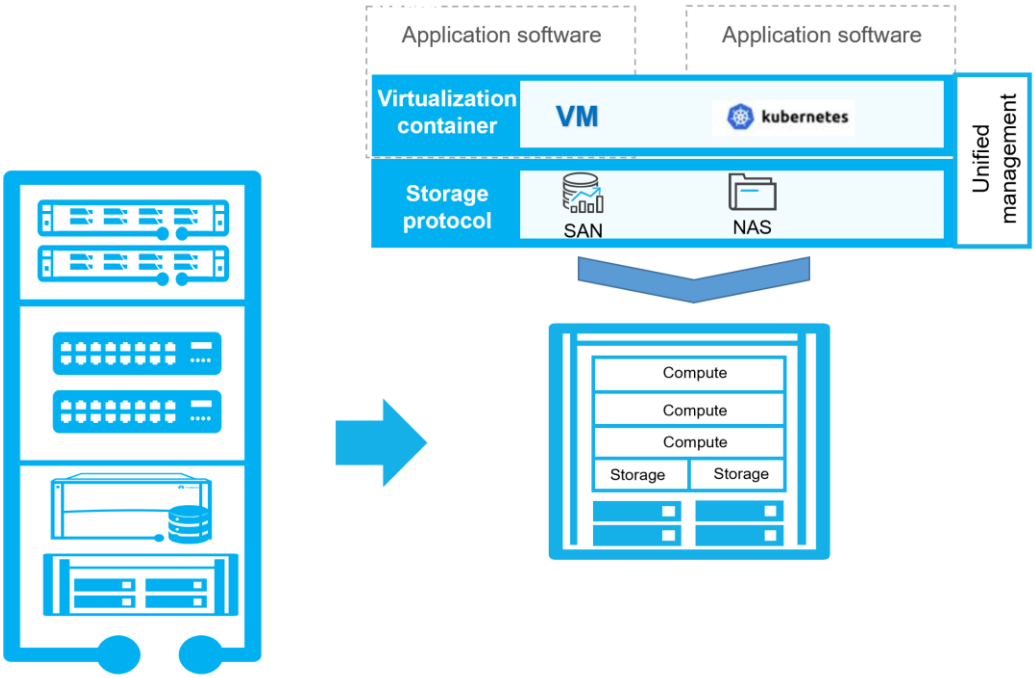

2. Product Features

Simplified management

- Built-in storage, switching, and computing, free of network installation and design.
- System installation is not required, and one-click initialization is quick, achieving quick service rollout.
- Provides unified GUI management and proactive fault diagnosis, simplifying routine O M.
- Hierarchical management and remote application provisioning.

Extremely reliable

- A-A storage architecture, providing a reliable data storage platform.
- Diskless computing design, meaning that system disks and RAID controller cards are not required, and OS faults can be quickly rectified.
- Fully redundant hardware design → Faults of a single fan or expansion board do not affect services
- Fast service support when a single node fails
- Quick automatic recovery

Lightweight and flexible

- High-density design, supporting two storage nodes and three full-width compute nodes in a 4 U space.
- The storage capacity can be expanded online and the disk enclosure are plug-and-play.
- A maximum of 110 disks are supported.
- Compute nodes are plug-and-play, enabling quick capacity expansion.

Solid and Reliable Data Foundation

- Supports various enterprise storage data protection solutions, including high-density snapshot, remote replication, and active-active.
- Supports data reduction technologies to maximize storage efficiency.
- Supports RAID 2.0+, providing up to 92% storage utilization.
- SAN and NAS integration is supported. No additional NAS gateway is required.
- The built-in enterprise storage provides storage services with high concurrency and throughput for database applications. Professional enterprise NAS provides worry-free performance for massive small files.
- Reconstruct 1 TB data within 15 minutes.

3. HW Capabilites

Compute node	Number of nodes	1 to 3
	CPU	C4310, C4314, C4316, C5318Y, C6330, C6326, C6342, C6346, and C6348
	Memory	Each node supports a maximum of 32 memory slots.
	FlexIO card	GE/10GE
	Standard PCIe card	GE/10GE/25GE
	AI card	T4, A10, Atlas300I PRO
	Disk	Compute nodes do not require system disks, and OSs are offloaded to storage nodes.
	Onboard port	4 x 10GE (on management modules P0 and P1)
Storage node	Storage controller	2 nodes
	CPU	28-core 2.6 GHz/controller
	Memory	40 GB/controller
	Onboard port	2 x 10GE (port P2, used as the storage replication port), 4 x 10GE (ports P0 and P1, shared with compute nodes)
	Internal ports	2 x 25GE + 16 x 10GE
	Disk	The system supports a maximum of 110 disks.
	Disks in a controller enclosure	4 SSDs + 6 HDDs
	Disk enclosure for capacity expansion	4 U 3.5-inch SAS disk enclosure x 4, or 2 U 2.5-inch SAS disk enclosures x 4

4. System Logical Architecture

The OS 2910 is a product that provides converged computing, storage, and network services.
Compute nodes and storage nodes share the same subrack to provide virtualization service capabilities and enterprise storage capabilities. A unified management platform is used to manage and maintain resources and systems.

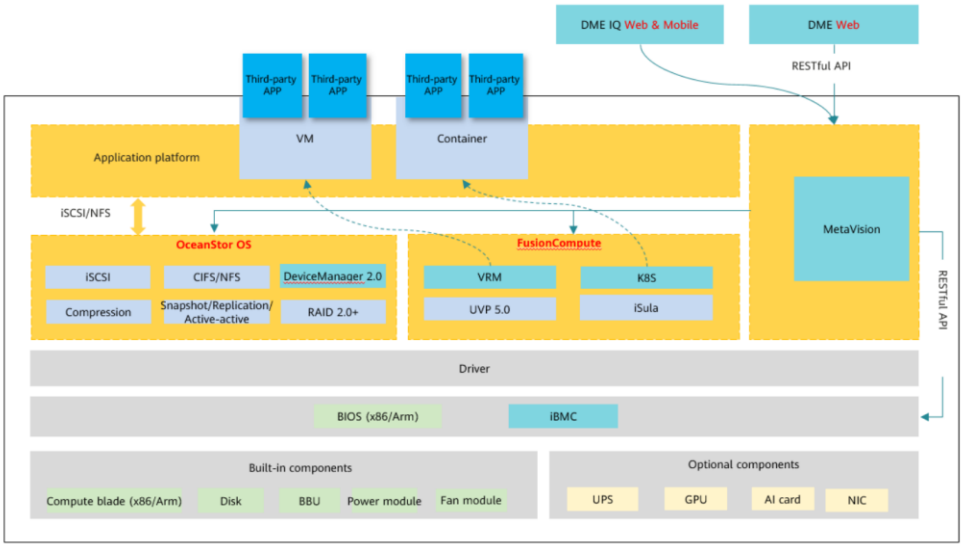

Logical layer		Function Description
Application Service	User Service	Users build service systems based on VMs or Containers.
	System Management	VMs are created to run the management software.
Software Layer	Compute	FusionCompute virtualizes computing and network resources.
	Storage	Storage modules provide storage space for user VMs and management VMs.
MetaVision Management System		The MetaVision unified management platform is used to configure virtualization services, storage services, and manage users, devices, and alarms.
Hardware layer	Management and Computing Converged Node	The OS 2910 consists of three compute nodes and two storage nodes. The management software of the OS 2910 runs on the two compute nodes. A management and computing converged node is an MCNA compute node that provides management and computing functions. This node not only provides user services, but also runs management VMs on it.
	Compute node	Provides computing functions and provides user services.
	Storage node	Provides storage and switching network functions, and provides storage services and network switching services between computing and storage.

5. Product Network Architecture

This section describes the network architecture of the OS 2910. This section helps you understand the network planes, port functions, and internal network principles of the OS 2910, facilitating subsequent network configuration. Figure 5-1 shows the network architecture of the OS 2910. Figure 5-1 uses two MCNAs and one CNA as an example. The network architectures of other devices are similar.

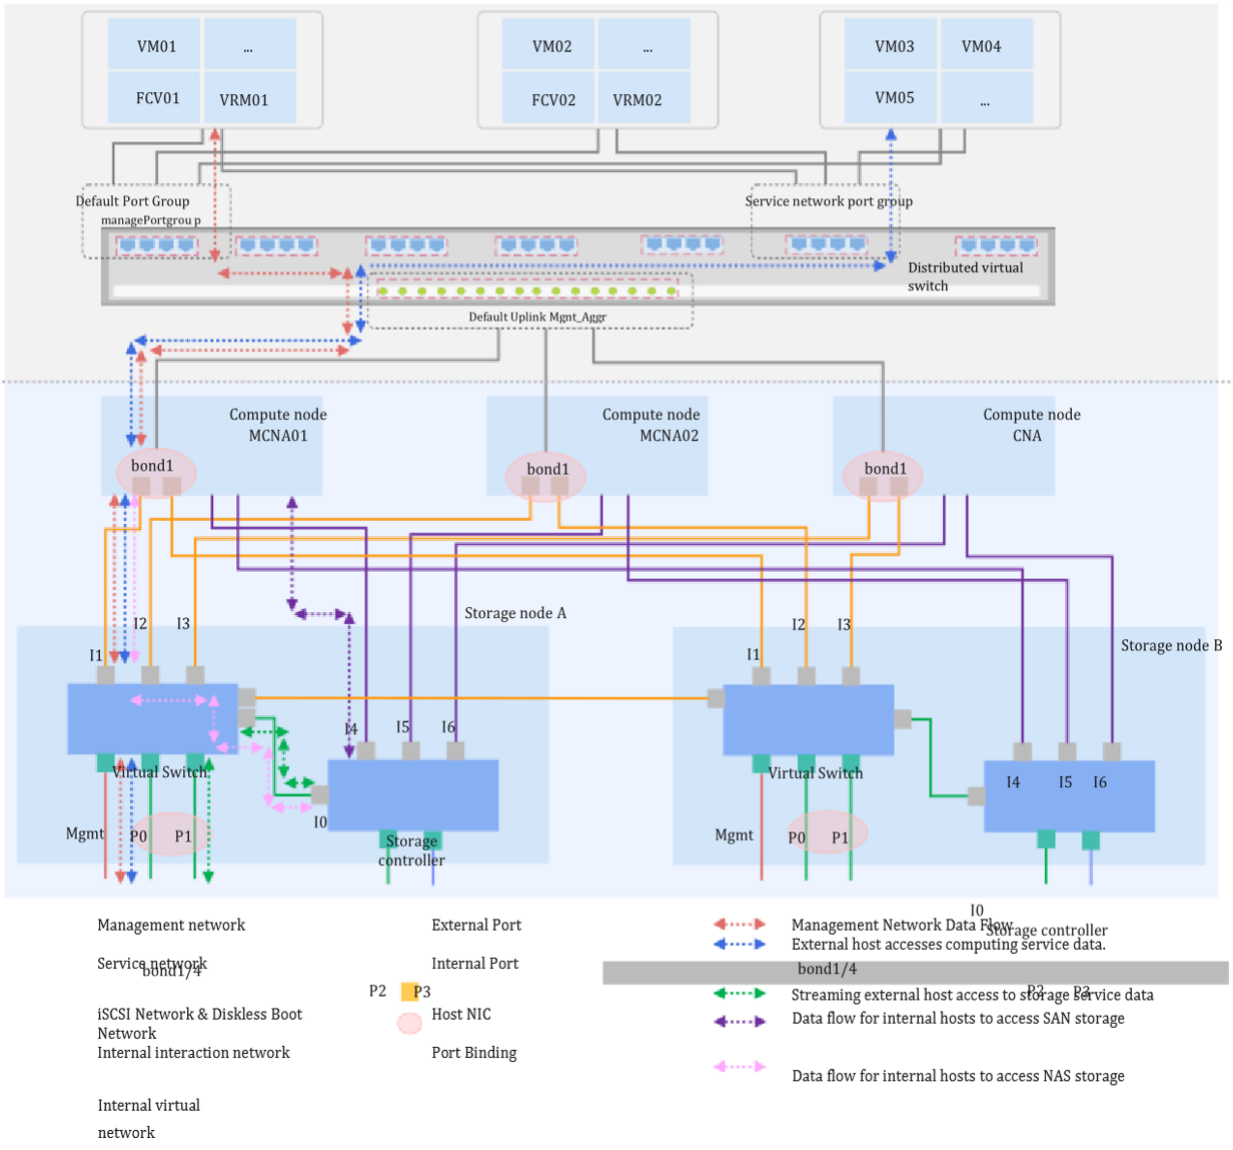

Network element		Function
Physical	Mgmt Port	External port and management port of the device. The external maintenance terminal connects to the system through the Mgmt port. The OS 2910 is connected to the maintenance terminal to manage and maintain the device.
	P0/P1 port	External 10GE port, used for front-end services. External switch/hosts through Ports P0 and P1 are connected to OS 2910 to access storage or computing services.
	P2 Port	External 10GE port used for front-end services. An external switch or host is connected to the OS 2910 through the P2 port to access NAS or SAN storage services. Only storage services can be accessed.
	P3 Port	External port of the device, which is a subtending port. OS 2910 uses P3 ports to cascade disk enclosures to expand storage resources.
	I0 Port	Internal port used for front-end services. When external switches or hosts access NAS or SAN storage services through ports P0 and P1, internal switches or hosts access NAS or SAN storage services through ports I0.
	I1/I2/I3 Port	This port is an internal port used for internal network switching. You can ignore this port.
	I4/I5/I6 Port	Internal port of the appliance, used for internal iSCSI network and diskless boot. Compute nodes are diskless and use LUN space provided by storage nodes as system disks. Ports I4, I5, and I6 provide SAN storage services for compute nodes.
Virtual	Distributed Virtual Switch	Similar to physical switches, distributed virtual machines (DVs) build networks between VMs and provide the capability of communicating with external networks. One end provides a port for connecting to the VM, and the other end provides an uplink for connecting to the physical NIC of the host where the VM is located.
	Port group	A port group is a collection of virtual ports on a DVS. It defines the mode in which VM NICs are connected to physical networks through the DVS. VM NICs connected to the same port group have the same network attributes, and VMs connected to different port groups are isolated from each other.
	Uplink	Uplinks are used by the DVS to connect to the physical NICs of hosts and are used for VM data uplinks.
	VLAN	The management network and service network of the OS 2910 pass through the same port in the system. VLANs are used to isolate the management network and service network. The default VLAN ID of the management plane is 2910.

6. Software Introduction

6.1 Key Virtualization Features

6.1.1 Computing Virtualization

Server virtualization

Physical resources of servers are abstracted into logical resources, and one server becomes several or even hundreds of isolated virtual servers. Hardware such as CPU, memory, disk, and I/O becomes a resource pool that can be dynamically managed. This improves resource utilization and simplifies system management.

Function	Description
CPU virtualization	The OS 2910 virtualizes the CPUs of physical servers into vCPUs for running VMs. When multiple vCPUs are running, the OS 2910 dynamically schedules the capabilities of the physical CPUs among the vCPUs.
Memory virtualization	The OS 2910 supports the hardware-assisted memory virtualization technology, reducing memory virtualization overhead. In addition, the system supports the intelligent memory overcommitment (x86) policy and automatically optimizes and combines various memory overcommitment policies to achieve a high memory overcommitment rate.

VM resource management

Function	Description
VM Lifecycle Management	The OS 2910 provides multiple VM operations, including creating, deleting, starting, stopping, hibernating, cloning, and querying VMs.
VM Template Management	With the VM template function, you can define a standardized template for a VM and use the template to create a VM.
CPU QoS control	VM CPU QoS# is used to ensure the allocation of computing resources, isolate the impact of computing capabilities between VMs caused by different services, meet the requirements of different services on VM computing performance, maximize resource reuse, and reduce costs. When creating a VM, you can specify the CPU QoS based on the CPU performance requirements of the service to be deployed. Different CPU QoSs represent different computing capabilities of VMs. For VMs with specified CPU QoS, the system guarantees the CPU QoS, including the minimum computing capability and resource allocation priority. CPU QoS control is implemented by using the following parameters: •CPU resource share Defines the proportion in which multiple VMs are allocated computing resources when they compete for physical CPU resources. • CPU resource reservation Defines the minimum computing resources allocated when multiple VMs compete for physical CPU resources. •CPU resource limit Controls the upper limit of physical CPU resources occupied by VMs.
VM disk management	By default, OS 2910 disks in thin provisioning mode are attached to VMs. Shared disks can be created for multiple VMs to share at the same time. Supports creation of slave mode or standalone-persistent mode disks.

Dynamic VM resource adjustment

The OS 2910 supports dynamic adjustment of VM resources. Users can dynamically adjust resource usage based on service loads.

Function	Description
Adjusting the number of vCPUs offline or online	Users can increase the number of vCPUs on a VM regardless of whether the VM is offline (stopped) or online. When a VM is offline, you can reduce the number of vCPUs on the VM as required. You can adjust the number of vCPUs on a VM offline or online to flexibly adjust the computing capability when the service load on the VM changes.
Adjusting the memory size offline or online	Users can increase the memory capacity of a VM regardless of whether the VM is offline or online. When a VM is offline, you can reduce the memory capacity of the VM as required. You can adjust the memory size offline or online to meet the requirements for flexible memory adjustment when the service load on a VM changes.
Adding or deleting a NIC offline or online	When a VM is online or offline, you can attach or detach virtual NICs to or from the VM to meet the service requirements on the number of NICs.
Attaching a Virtual Disk Offline or Online	Users can attach virtual disks to VMs regardless of whether the VMs are offline or online. This increases the storage capacity of VMs and enables flexible use of storage resources without interrupting services.

6.1.2 Network Virtualization

Virtual NIC

Each virtual network adapter has its own IP address and MAC address. From the network perspective, the virtual network adapter is the same as the physical network adapter. The OS 2910 improves the I/O performance of virtual NICs by providing functions such as virtual switching, QoS, and uplink aggregation.

Network I/O Control

Network QoS policies provide bandwidth control capabilities and control the amount of data transmitted in and out of a port group to improve network performance. The transmission rate of nonkey services is limited to ensure the transmission rate of key services.

Distributed virtual switch

Similar to physical switches, distributed virtual switches (DVSs) build networks between VMs and provide the capability of communicating with external networks. One end provides a port for connecting to the VM, and the other end provides an uplink for connecting to the physical NIC of the host where the VM is located. By default, a DVS is created for the OS 2910, and you do not need to manually create a DVS.

6.1.3 VM Snapshot

A snapshot records the data and status of a VM disk at a specific time. Users can use snapshots to back up VM disks so that VM data can be restored to the time when the snapshot is created, improving data security and reliability. OS 2910 supports two VM data protection modes: manual snapshot and automatic snapshot.

Function	Description
Manual snapshot	Manually create a snapshot for a specified VM. After the snapshot task is created successfully, the system immediately executes the snapshot creation task.
Automatic snapshot	Periodic snapshot, which is based on the user-defined snapshot policy. (Specify the snapshot start time, frequency, number of retained copies, and so on in the policy) Creates snapshots for VMs periodically. Proper use of automatic snapshot can effectively improve data security and operation fault tolerance.

6.1.4 VM Migration

If the host where a VM is located is faulty, the data store where the VM is located is faulty, or the resource allocation is uneven (for example, overloaded or light), you can migrate the VM from one host to another host or from one data store to another to ensure the normal running of the VM and improve resource utilization.

The OS 2910 supports two migration modes:

- Change Host: Migrate a VM to another host by changing the host. During the migration, disks are not migrated. The position of the disk remains unchanged and the disk is still in the original data store. This operation is supported only by VMs in the Running state.
- Change storage: Migrate the disks of a VM to another datastore by changing the storage. During the migration, the host where the virtual machine is located remains unchanged. This operation is supported only by VMs in the Stopped state.

6.1.5 Dynamic Resource Scheduling

The OS 2910 can detect the usage of global physical resources during VM creation and running, use the intelligent scheduling algorithm to determine the optimal host for VM running, and migrate VMs to the optimal host, improving global service experience. The OS 2910 supports dynamic VM scheduling by configuring dynamic resource scheduling policies, VM scheduling rules, and VM balancing groups.

Dynamic resource scheduling policy

A dynamic resource scheduling policy defines the conditions for triggering VM scheduling. The policy defines the triggering factors for resource scheduling, execution period of scheduling policies, and migration threshold (conservative or aggressive).

- The dynamic resource scheduling function is disabled by default. You need to manually enable and configure the dynamic resource scheduling function before using it. Dynamic resource scheduling supports different triggering factors. You can configure triggering factors based on the CPU usage and memory usage.
  The rate or both determines whether to trigger VM migration.
- The system supports automatic and manual resource scheduling. In automatic scheduling mode, the system
  The system automatically migrates VMs to the most appropriate host. In manual scheduling mode, the system generates scheduling suggestions. You can determine whether to apply the suggestions based on the site requirements.
- Dynamic resource scheduling supports the setting of time-based thresholds to meet scheduling requirements in different time periods.
  Migration brings system overhead. Therefore, you are advised to set the policy to conservative when the service pressure is heavy and to moderate or aggressive when the service pressure is light to avoid affecting service performance.

VM scheduling rules

After the dynamic resource scheduling policy is configured, the system determines the conditions for triggering VM scheduling. The scheduling rules determine how to schedule VMs. Scheduling rules include the affinity and anti-affinity between VMs, and the startup priority of VMs that must/forbidden/preferably/should not run on a host.

VM balancing group

The OS 2910 supports the creation of a VM balancing group and the addition of multiple VMs to the group. After dynamic resource scheduling, the overall load in the cluster is balanced. By default, VMs in the cluster are evenly distributed in load balancing mode. For example, two different services are deployed in a cluster. To ensure stable running of the two services, you can add VMs of the two services to a load balancing group to implement load balancing based on services.

6.1.6 VM Security Group

A security group is a virtual firewall that controls the incoming and outgoing traffic of VMs in a security group, improving VM security. When creating a security group, you need to specify whether the security group type is trustlist or blocklist.

VMs in the security group allow or deny access to communication objects. You can add multiple security rules to a security group. After a VM is added to the security group, the VM is protected by these rules.

Type	Description
Trustlist	If a security group is used as a trustlist, VMs in the security group are denied all access by default. You can add security rules to configure the objects that VMs in the security group are allowed to communicate with.

Blocklist	If the security group is used as a blocklist, all VMs in the security group are allowed to access by default. You can add security rules to prevent VMs in the security group from communicating with each other.

A security group can contain multiple security rules. Lists the attributes that can be specified in a security rule.

Property	Description
Directions	Whether the security rule is applied to the traffic that flows out of the virtual machine or to the traffic that flows into the virtual machine.
Agreement	Indicates the protocols for which the security rule takes effect, such as TCP, UDP, ICMP, or all protocols.
IP address	Specifies the IP addresses for which the security rule takes effect.
Port Range	Ports on which the security rule takes effect.

6.1.7 VM HA

When a host, VM, or data store in a cluster is faulty, the system automatically restarts the faulty VM on a normal host based on the configured HA policy. In this way, the faulty VM can be quickly recovered.

The OS 2910 supports the setting of host, VM, and data store fault and response policies, resource reservation for VM HA in a cluster, and group fault control.

- Host, VM, or data store fault and response policy: Set the fault response policy when the host, VM, or data store is faulty.
- Resource reservation: Resources are reserved for VMs in a cluster to ensure sufficient failover capacity.
- Massive fault control: When a host is faulty in a batch, a large number of VMs are deployed in HA mode.

6.1.8 Safe Power-Off

The graceful power-off function is used to prevent VM software damage, service startup failure, and service data damage when the mains supply is cut off or the device is powered off by pressing the button. To use this function, you need to configure a secure power-off policy. When the entire system is powered off, the system powers off VMs and then devices in sequence based on the power-off policy to prevent VM service loss. To apply the safe power-off policy in the case of a mains failure, you need to purchase the UPS in advance and connect it to the OS 2910.

The safe power-off policy defines the following parameters:

- VM power-off priority: Power off VMs in batches.
- Phase timeout interval: When the power-off duration of a batch of VMs exceeds the threshold, the system automatically enters the next power-off phase.
  Section.
- How to supply power after a mains power failure: whether to immediately power off the device according to the power-off policy after a mains power failure, or Power off the device based on the remaining power supply time of the UPS.

6.1.10 Container Management

The OS 2910 supports the container function. (The single management node mode does not support the container function, and the active/standby mode supports the container function.). After the container function is enabled on MetaVision, you can switch to the FusionCompute page to use the container function. FusionCompute container management provides the enterprise-class K8s platform and provides tenant project management, cluster lifecycle management, container image and application management, and container monitoring and O&M capabilities. FusionCompute container management is the best platform for application modernization.

- Container cluster management: supports Kubernetes cluster life cycle management, including deployment, upgrade, connection, scaling, and configuration.

- Container image management: supports container image repository, allowing users to upload, push, and pull container images and interconnect with the container image repository. Users can create an image repository and synchronize images.
- Containerized application management: Provides visualized management for containerized applications throughout the lifecycle, including application templates and applications.
- Container project management: Container projects can be used to logically isolate tenant resources to meet the requirements of internal enterprises. Container resources of multiple teams are self-serviced and managed.
- Container monitoring and O&M: provides comprehensive K8s event monitoring and performance monitoring capabilities.

6.1.11 Backup and Active-Active

Backup solution

This backup solution uses Huawei eBackup software, FusionCompute snapshot function, and CBT function to back up VMs and disks. eBackup works with FusionCompute to back up specified objects based on specified policies. If VM or disk data is lost or faulty, you can use backup data to restore the VM or disk data.

Active-Active Solution

This active-active solution integrates the HyperMetro feature of OS 2910, the HA capability of FusionCompute, and the DR management function of Huawei UltraVR to ensure high availability of VM services. In the active-active solution, both devices back up each other and are running. The HyperMetro feature enables service data to be written to storage devices at both ends. When a VM, device, or even a fault occurs in a data center, the VM cross-device HA is implemented to ensure quick service recovery and meet customers' requirements for service continuity and reliability.

6.2 Key Storage Features

The OS 2910 has enterprise-class storage capabilities and provides enterprise-class storage features. lists the key storage features supported by the.

Features	Description
Block service	Divide the storage space provided by the storage system into multiple LUNs and map the LUNs to the application server so that the application server can use the storage space provided by the storage system.
File service	The storage space provided by the storage system is divided into file systems and shares are configured. Application servers access the shares to use the storage space.
HyperSnap	Provides the snapshot function for LUNs and file systems. Snapshots do not make complete physical copies of data. They use mapping tables to locate data and quickly provide complete copies of data.
HyperCDP	Second-level data protection. Generates more intensive recovery points for LUNs or file systems on storage devices, providing continuous data protection for users.
HyperClone	Clone LUNs and file systems. Clone replicates the source data locally to generate a data copy for data backup or other applications.
HyperLock	Implements the WORM function of the file system. The write-once-readmultiple technology ensures the integrity, confidentiality, and accessibility of original data in storage media, meeting the increasing requirements for secure data storage.
HyperReplication	Remote replication of LUNs and file systems. Remote replication can generate an available copy in almost real time on storage systems in different regions. This copy can be put into use immediately without data recovery, maximizing the continuity of user applications and data availability.
HyperMetro	Implements the active-active function of LUNs. HyperMetro synchronizes and accesses data between two storage systems in real time, improving resource utilization and enabling seamless service switchover when data access on either storage system fails, ensuring data security and service continuity.
SmartMultiTenant	Implements the multi-tenant function. By creating multiple virtual storage systems in a physical storage system, multiple tenants are provided with shared storage resources that are more flexible, easy to manage, and costeffective to deploy without affecting data security and privacy.
SmartMigration	Implements the data migration function. SmartMigration migrates services from the source LUN to the target LUN without interrupting host services. In this way, the target LUN can replace the source LUN to carry services.
SmartVirtualization	Manages third-party resources. SmartVirtualization enables local storage systems to centrally manage storage resources in third-party storage systems, simplifying storage system management and reducing maintenance costs.
SmartQoS	Implements the service quality function. SmartQoS is used to control LUN performance to ensure the quality of high-priority services.
Hard disk encryption	Encrypts hard disk data. Provides full disk data encryption protection without impacting storage performance.
IP address flapping	Implements IP address flapping. IP address flapping can quickly switch services on a faulty port to another available port of the same type without interrupting host services. After the faulty port recovers, it takes over services again, ensuring service continuity.

7. Hardware Introduction

7.1 4U controller enclosure

This chapter describes the hardware structure, functions and features of each component, front and rear views, and indicators of the controller enclosure.

7.1.1 Overview

A controller enclosure adopts the modular design and consists of the system enclosure, compute nodes, storage nodes, power modules, fan modules, BBU modules, and hard disk drawers.

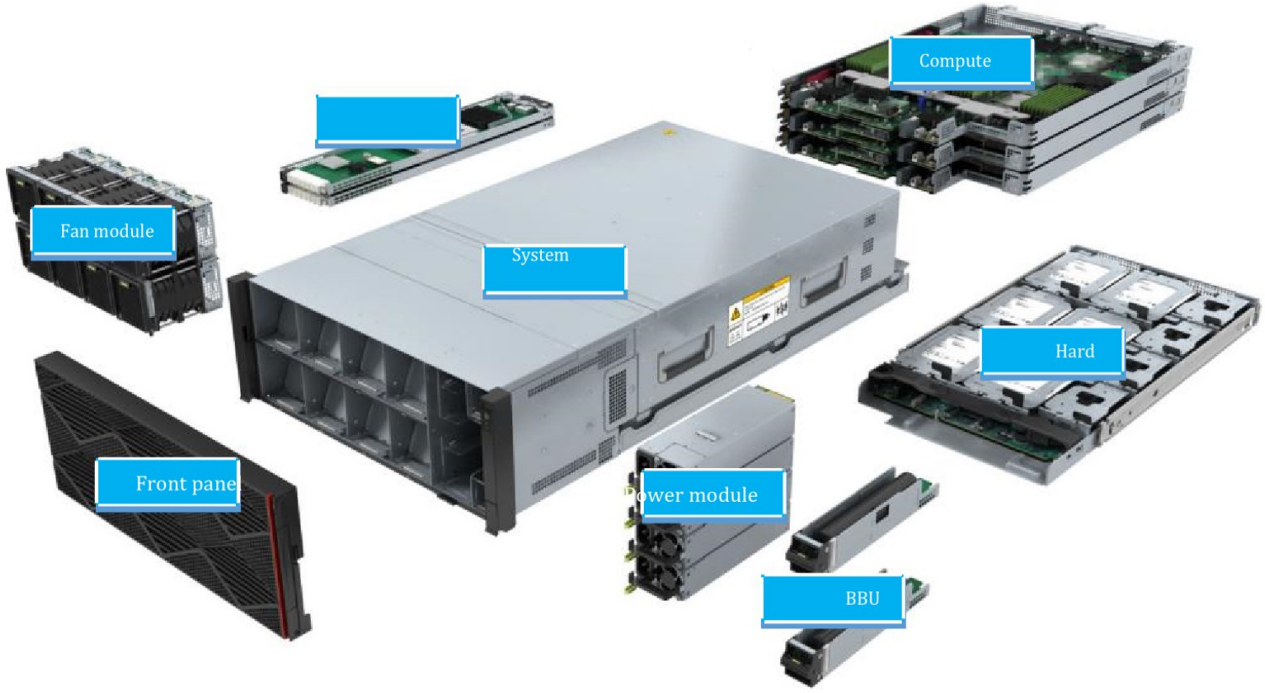

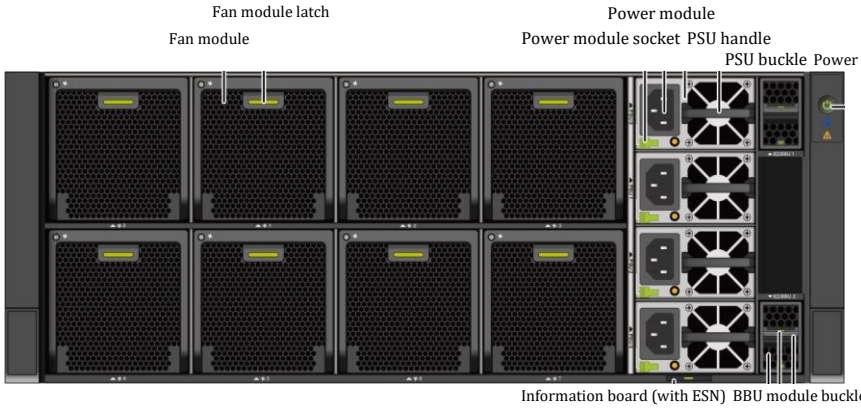

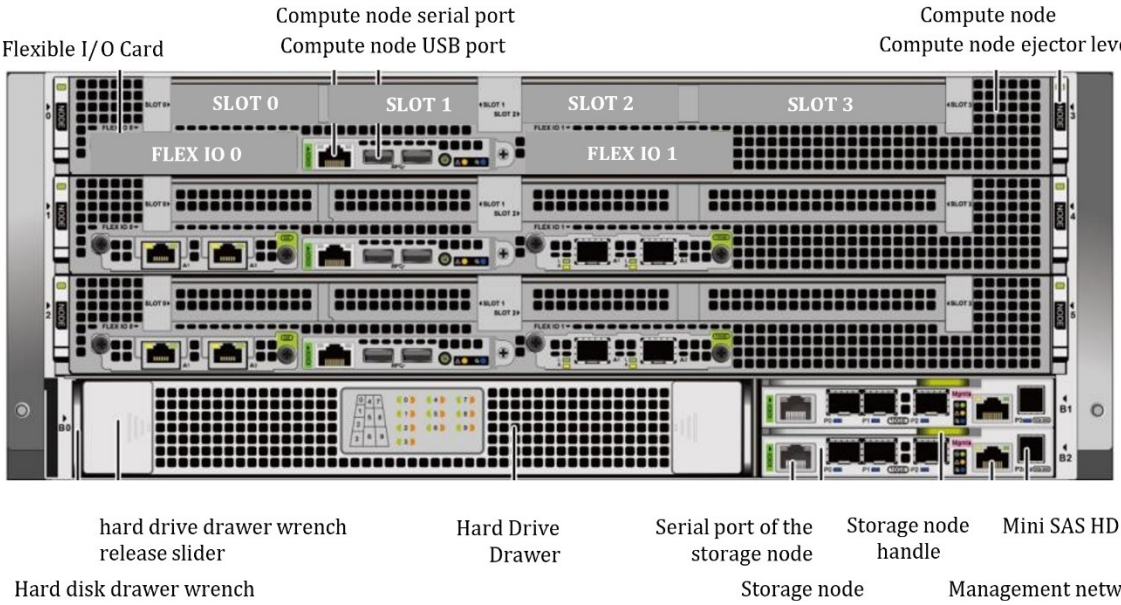

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

- Compute nodes are arranged as compute node 0, compute node 1, and compute node 2 from top to bottom. Storage node A and storage node B are arranged from top to bottom.
- If only one compute node is configured, the compute node must be installed in slot 0. When only two compute nodes are configured, compute nodes must be installed in slots 0 and 1.
- Each compute node and storage node are equipped with a serial port. Users can access the compute node through the serial port.
  Resources and storage resources. This command is used when you cannot log in to the device through the management network port.
- Each compute node supports four standard NICs and two flexible I/O NICs. Slots 0 to 3 are used to install the flexible I/O NICs. Standard cards: Flex IO 0 and Flex IO 1 are used to install flexible I/O cards.
- The hard disk drawer supports two types of hard disk configuration: one which can house a maximum of four 2.5-inch hard disks and six 3.5-inch hard disks, the second can house ten 2.5-inch SSDs.

Hardware Parameters

Name	Parameters
Dimensions (L x W x H)	800mm*447mm*175mm
Weight (excluding the weight of auxiliary materials such as guide rails and cables delivered with the equipment)	76.3 kg (configured with two compute nodes and no hard disk drawer) 81.9 kg (configured with three compute nodes and hard disk drawers)
AC power supply voltage and rated current	2000 W power supply, 200–240 V AC, 10 A, singlephase, 50 Hz/60 Hz
HVDC	Supports 2000 W power supply (240 V DC input), 192–288 V DC, 10 A

7.1.2 Components

This section describes the appearance and ports of the components in the controller enclosure.

7.1.2.1 System Subrack

Through the backplane, the system subrack provides reliable connections for various interface modules, implementing signal and power interconnection between modules.

7.1.2.2 Compute Node

VM services can be deployed on compute nodes. In addition, the compute node can be connected to other devices by using flexible I/O cards.

Notices

For the x86 compute node, only one compute node can be configured in the controller enclosure.

Rear view of the compute node

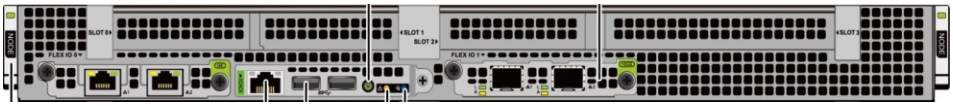

7.1.2.3 Storage Node

Storage nodes provide management ports, including management network ports and serial ports. It also provides service ports, including 10GE optical ports and mini SAS HD expansion ports.

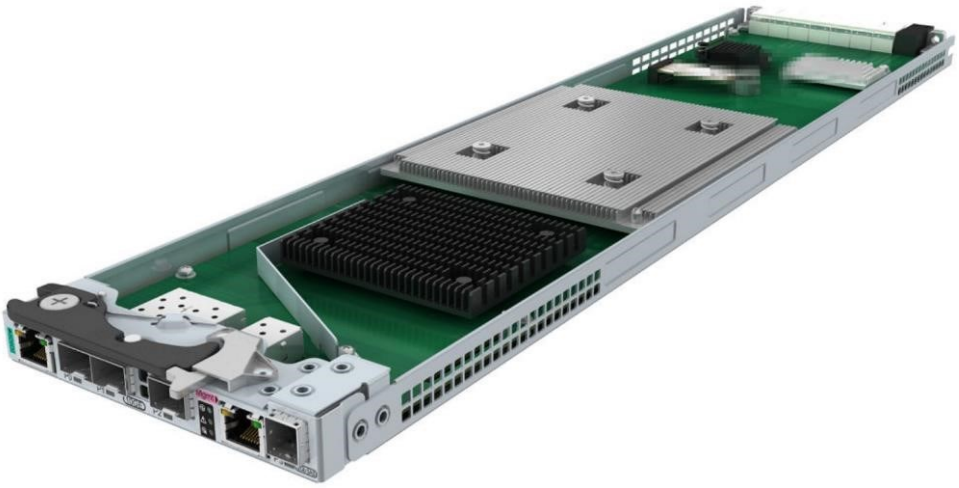

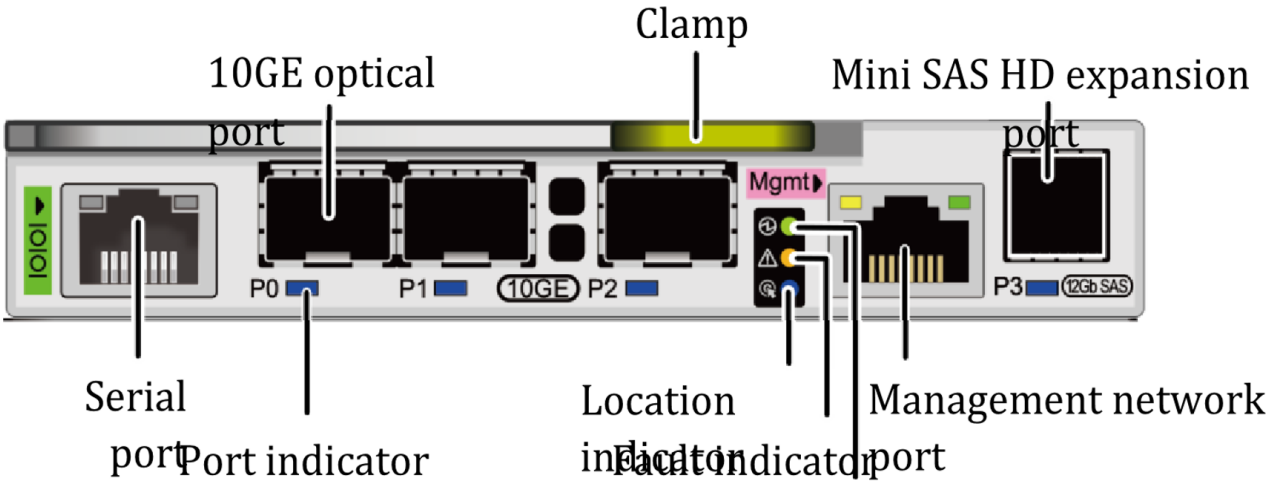

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

7.1.2.4 BBU Module

When the external power supply fails, the BBU provides backup power to ensure the security of service data. When the output of the system power is normal, the BBU works in backup mode. When the external power supply is disconnected, the BBU can continue to supply power to the system. The BBU supports fault isolation. When a BBU is faulty, the system can run properly.

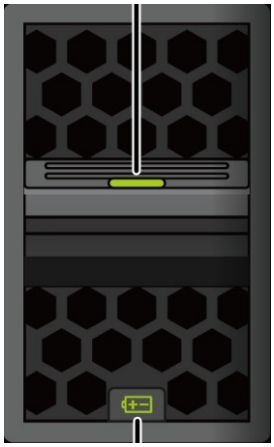

7.1.2.5 Power Supply Module

The PSU is an AC power supply unit (PSU) that supports the normal operation of the controller enclosure in the maximum power consumption mode. Each controller enclosure supports a maximum of four PSUs, and the slots are PSU 0, PSU 1, PSU 2, and PSU 3 from top to bottom.

The number of PSUs installed in the controller enclosure varies according to the number of compute nodes. If only one compute node is configured, the controller enclosure is configured with two PSUs, which are installed in slots PSU 0 and PSU 3. If two or three compute nodes are configured, four PSUs are installed in the controller enclosure.

Notices

To expand the capacity from one compute node to two or three compute nodes, install PSUs in slots PSU 1 and PSU 2.

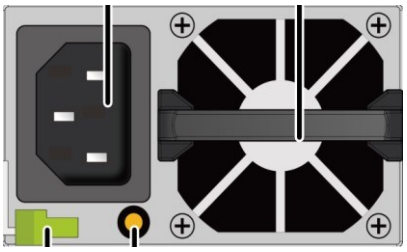

7.1.2.6 Fan Module

Fan modules dissipate heat for the system and ensure that the controller enclosure can run properly in condition of maximum power consumption.

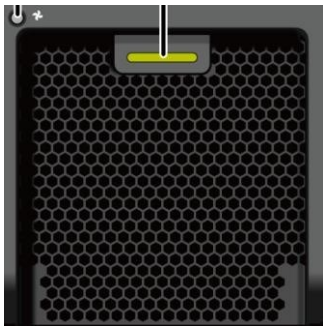

7.1.2.7 Hard Drive Drawer

The disk modules in the controller enclosure are installed in the disk drawer. The disk drawer supports two types of disk configuration: a maximum of four 2.5-inch disks, six 3.5-inch disks, and a maximum of ten 2.5-inch disks. You can view the status of the disk modules in the drawer by observing the disk indicators on the panel.

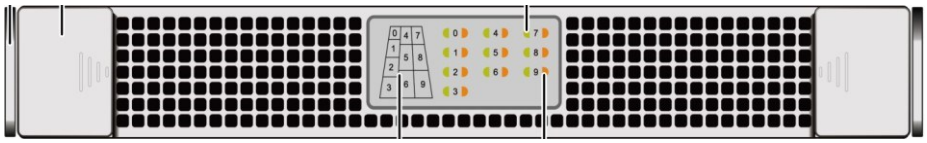

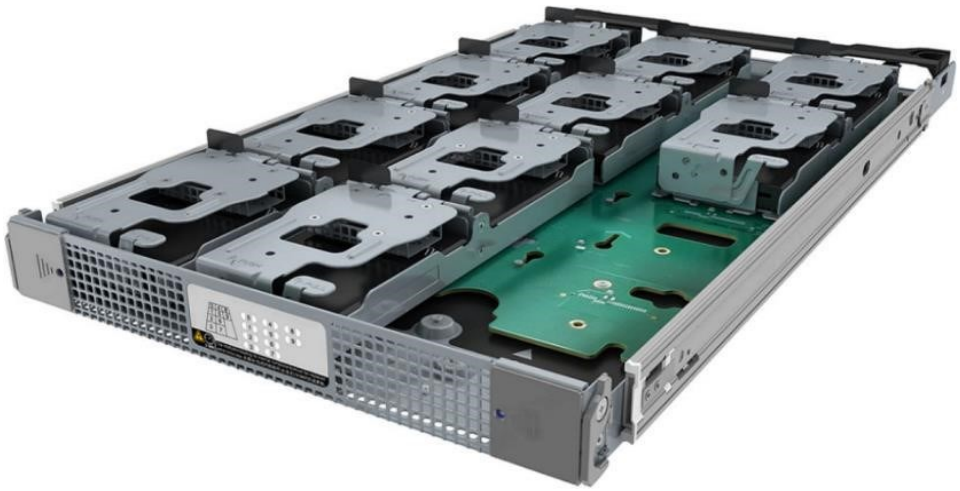

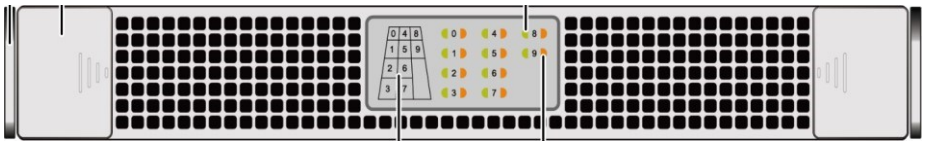

7.1.2.8 Disk Module

The disk modules in the controller enclosure are installed in the disk drawer. Disk modules provide storage capacity for the system. Disk modules are available in 2.5-inch and 3.5-inch specifications.

After the controller enclosure is powered on, you can view the disk indicators on the disk drawer panel to check the disk module status. For details about the indicators on the hard disk drawer, see 7.1.3 Indicators.

7.1.2.9 FlexibleIO Card

This section describes the functions, appearance, and indicator status of a flexibleIO card. For details about the types of FlexIO cards supported by different compute node architectures, see. Specification query tool

Notices

When installing a FlexIO card on a compute node, you must install the FlexIO card in slot 0 and then install the FlexIO card in Slot 1.

7.1.2.9 GE FlexIO Card

The GE flexible I/O card provides two 1 Gbit/s electrical ports. The rate cannot be reduced to 100 Mbit/s or 10 Mbit/s.

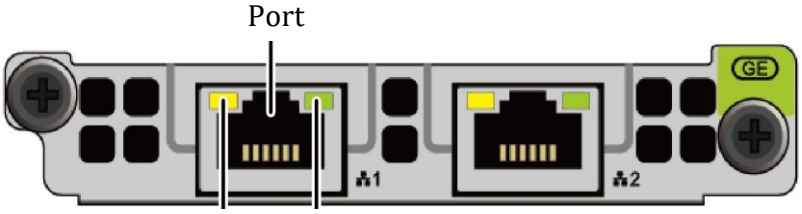

indicator light	Status and Description
Link/Active indicator	● Green: The connection between the and the peer device is normal. Blinking green (2 Hz): Data is being transmitted. ● Off: The connection to the peer device is abnormal.
Port speed indicator	● Steady yellow: The rate is the highest. ● Off: The rate is not the highest.

7.1.2.9.3 10 Gbit/s FlexIO Card

The 10 Gbit/s flexible I/O card provides two 10 Gbit/s optical ports.

Optical modules for controller enclosures cannot be purchased by customers. Use the optical module matching the flexible I/O card in the controller enclosure.

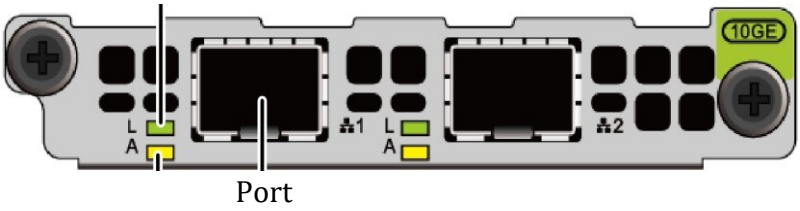

indicator light	Status and Description
Link indicator of the port	● Steady green: The physical link is normal. ● Off: The physical link is not established.
Port active indicator	Steady green: The port is connected but no data is being transmitted. Blinking green (2 Hz): The port is connected and the number of data is displayed. According to the transmission. Off: The port is not connected.

7.1.2.10 Standard Card

Compute nodes support two types of standard cards: half-height AI cards and full-height AI cards.

A half-height AI card can be installed only in slot 0, 1, or 2 of a compute node, and a full-height AI card can be installed only in slot 3 of a compute node.

Half-Height AI Card (Front View)

Full height AI card (front view)

Half-Heigth AI Card (Top View)

7.1.3 Indicators

After the controller enclosure is powered on, you can observe indicators to check the operating status of the controller enclosure.

Module	Indicator Type	Status and Description
BBU Module	Running/Alarm indicator on the BBU	Steady green: The BBU is fully charged. Blinking green (1 Hz): The BBU module is running. Charge. Blinking green at 4 Hz: The BBU module is in the normal state. Discharge. Steady yellow: The BBU module is faulty. Off: The module is not powered on or is hotswappable.
Power module	Power module running/alarm indicator	Green, on: The power supply is normal and no fault occurs. Blinking green (1 Hz): The power input is positive. ?? Blinking green (4 Hz): The power module is in the During the online upgrade. Steady yellow: The power module is faulty. Off: No external power is supplied.
Fan module	Fan module running/alarm indicator	● Steady green: The fan module is running properly. ● Steady yellow: The fan module is faulty. ● Off: The fan module is not powered on.
System subrack	Power indicator/power button on the controller enclosure	Steady green: The controller enclosure is powered on. Blinking green (0.5 Hz): The controller enclosure is just installed. Blinking green (1 Hz): The controller enclosure is being copied. Indicates the machine status. Blinking green (2 Hz): The controller enclosure is in the OS boot phase or is in the power-off process phase. Off: The controller enclosure is not powered on or is in the backup power state.
	Controller enclosure alarm indicator	Yellow, on: The system has a major or higher level. Indicates the alarm. Off: The system is running properly.
	Controller enclosure location indicator	Blinking blue: The controller enclosure is being located. Off: The controller enclosure is not located.

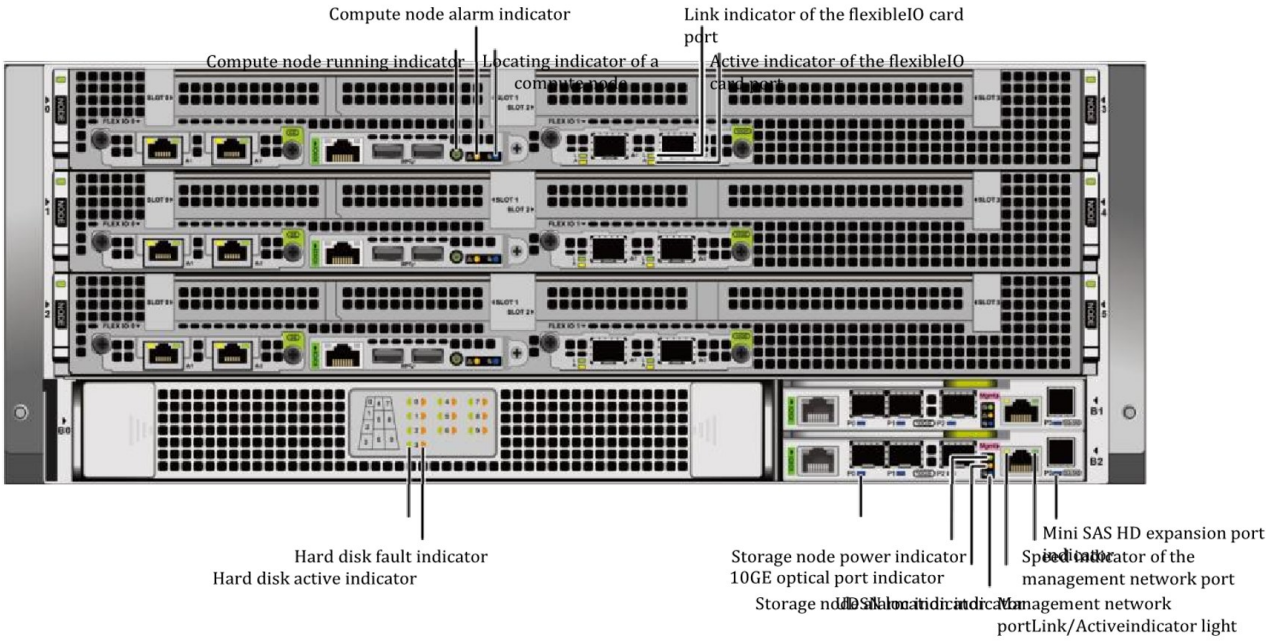

Module	Indicator name	Status and Description
Flexible I/O Card	Link indicator of the flexibleIO card port	Compute nodes support flexible I/O cards at different rates. For details about the indicators on flexible I/O cards, see 7.1.2.9.
	Active indicator of the flexibleIO card port	
Storage node	Storage node running indicator	● Steady green: The storage node is powered on and the system is started. The movement is complete. Blinking green (0.5 Hz): The storage node is justThe is powered on and is in the BIOS boot phase. ● Blinking green (2 Hz): The storage node is in the OS boot phase or is in the power-off process. ● Off: The storage node is not powered on.
	Storage node alarm indicator	Steady yellow: The storage node is faulty. ● Off: The storage node is running properly.
	UDSN location indicator	● Blinking blue (2 Hz): The storage node is being Positioning. ● Off: The storage node is not located.
	Mini SAS HD expansion port indicator	● Blue, on: The port transmission rate is 4 x. 12 Gbit/s. Green on: The port transmission rate is 4 x 6 Gbit/s. or 4 x 3 Gbit/s. ● Steady yellow: The port is faulty. ● Off: The port link is not connected.
	10GE optical port indicator	● Steady blue: The rate is the highest. ● Blinking blue (2 Hz): The speed is the highest. Rate, data is being transferred. ● Steady green: The rate is not the highest. Blinking green (2 Hz): The rate is not the highest. Indicates the rate. Data is being transferred. ● Steady yellow: The optical module or cable on the port is faulty or an unsupported optical module or cable is inserted. Blinking yellow (2 Hz): The phone is being located. ● Off: The port link is not connected.
Compute node	Compute node running indicator	● Steady green: The compute node has been powered on and the system starts. The movement is complete.

		Blinking green (0.5 Hz): The compute node is justThe is powered on and is in the BIOS boot phase. Blinking green (2 Hz): The compute node is in the OS boot phase or is in the power-off process. ● Off: The compute node is not powered on.
	Compute node alarm indicator	● Steady yellow: The compute node is faulty. ● Off: The compute node is running properly.
	Locating indicator of a compute node	● Blinking blue (2 Hz): The compute node is being Positioning. ● Off: The compute node is not located.
Hard Drive Drawer	Hard disk fault indicator	● Steady yellow: The hard disk is faulty or prefailed. ● Off: The hard disk is running properly.
	Hard disk active indicator	● Steady green: The hard disk is installed. Blinking green (4 Hz): The hard disk is in progress. Read and write data. ● Off: The hard disk is not detected or powered off.

7.2 SAS disk enclosure (2 U, 2.5-inch disks)

This chapter describes the hardware structure, functions, front and rear views, and indicators of a disk enclosure.

7.2.1 Overview

A disk enclosure uses a modular design and consists of a system enclosure, expansion modules, power modules, anD disk modules.

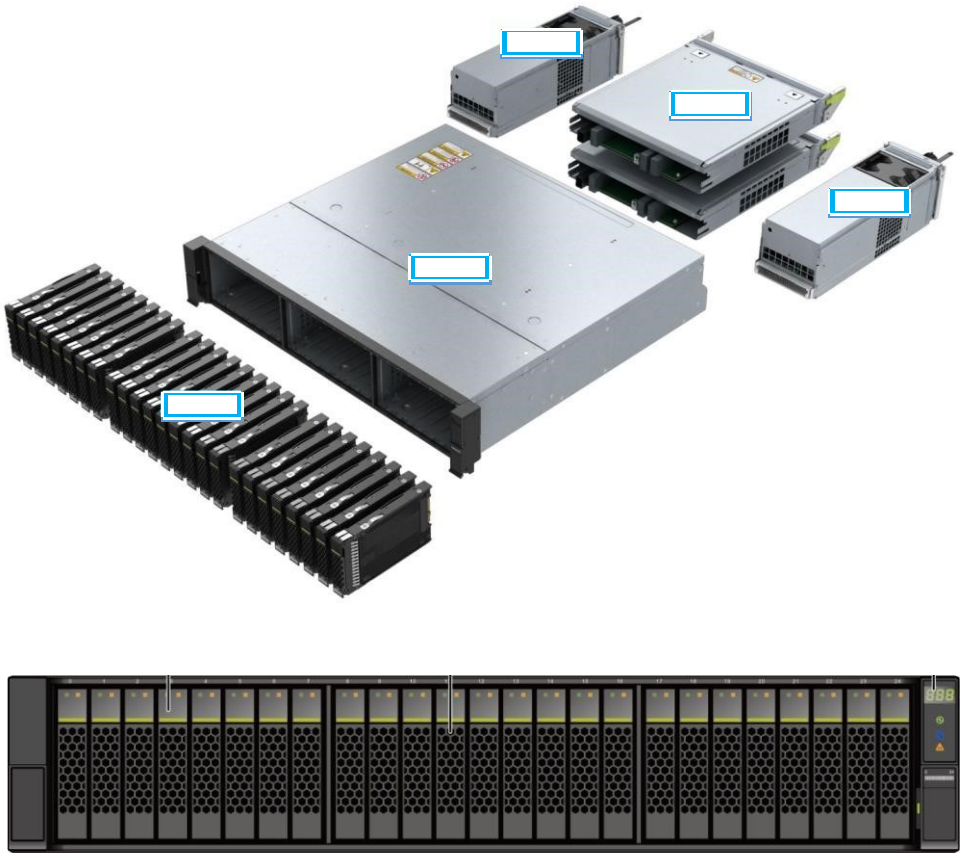

Disk slots are numbered from 0 to 24 from left to right.

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

Hardware Parameters

Name	Parameters
Dimensions (L x W x H)	410mm*447mm*86.1mm
Weight (excluding the weight of auxiliary materials such as guide rails and cables delivered with the device)	● 13.4 kg (excluding the weight of the hard disk unit) ● 19.65 kg (including the weight of the hard disk unit)
AC power supply voltage and rated current	800 W AC power supply (supporting 110 V dual-live-wire input (2 W + PE) (supporting 110 V single-live-wire input) 100–240 V AC ±10%, 10 A, single-phase, 50 Hz/60 Hz
HVDC	800 W power supply (240 V DC input), 192–288 V, 10 A

7.2.2 Components

This section describes the appearance and ports of each component in a disk enclosure.

7.2.2.1 System Subrack

The system subrack provides reliable connections for various interface modules through the backplane, implementing signal and power interconnection between modules.

7.2.2.2 Expansion Module

Each expansion module provides two expansion ports P0 and P1. Expansion modules connect to controller enclosures or disk enclosures through expansion ports and communicate with controller enclosures or disk enclosures. They are the connection points for data transmission between controller enclosures and disk enclosures.

Serial port Mini SAS HD expansion port "P0" Mini SAS HD expansion port "P1"

Alarm indicator of the expansion module

Mini SAS HD expansion port indicator

Power indicator of the expansion

Disk enclosure ID display module

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

Indicator Light	Status and Description
Alarm indicator of the expansion module	● Steady yellow: An alarm is generated for the expansion module. ● Off: The expansion module is running properly.
Power indicator of the expansion module	● Steady green: The expansion module is powered on. ● Off: The expansion module is not powered on.
Mini SAS HD expansion port indicator	● Steady blue: The port transmission rate is 4 x 12 Gbit/s. ● Green on: The port transmission rate is 4 x 6 Gbit/s or 4 x 3 Gbit/s. ● Steady yellow: The port is faulty. ● Off: The port link is not connected.

7.2.2.3 Power Supply Module

Power modules are classified into AC power modules and DC power modules. The power modules support the normal operation of disk enclosures in the maximum power consumption mode. Running.

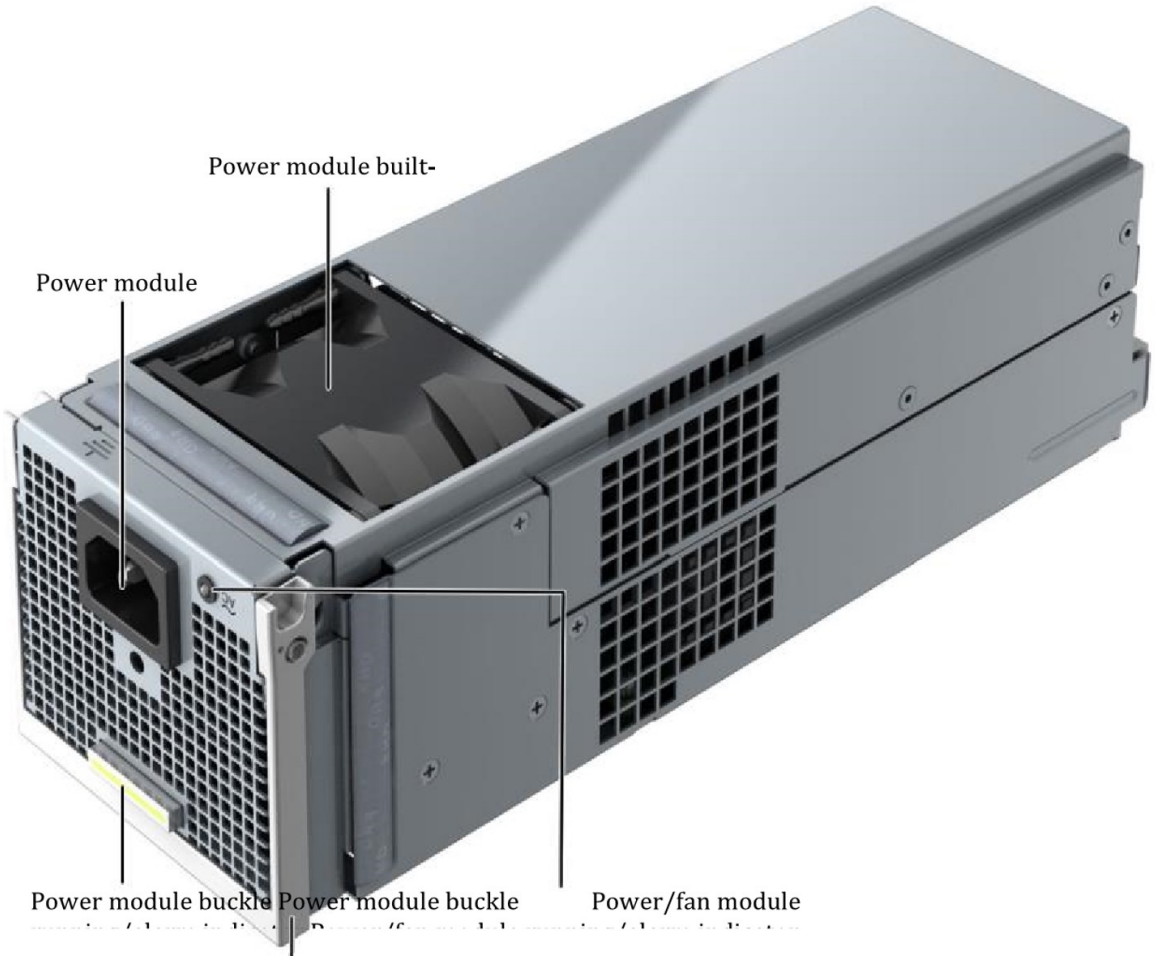

indicator light	Status and Description
Power/Fan module running/alarm indicator	Green, on: The power supply is normal and no fault occurs. Blinking green (1 Hz): The power input is normal and the device is not powered on. Electricity. Blinking green (4 Hz): The power module is being upgraded online. Medium. Steady yellow: The power supply or fan module is faulty. ● Off: No external power input is provided.

7.2.2.4 Disk Module

The disk module provides storage capacity for the system. It can function as a system coffer and store service data, system data, and cached data.

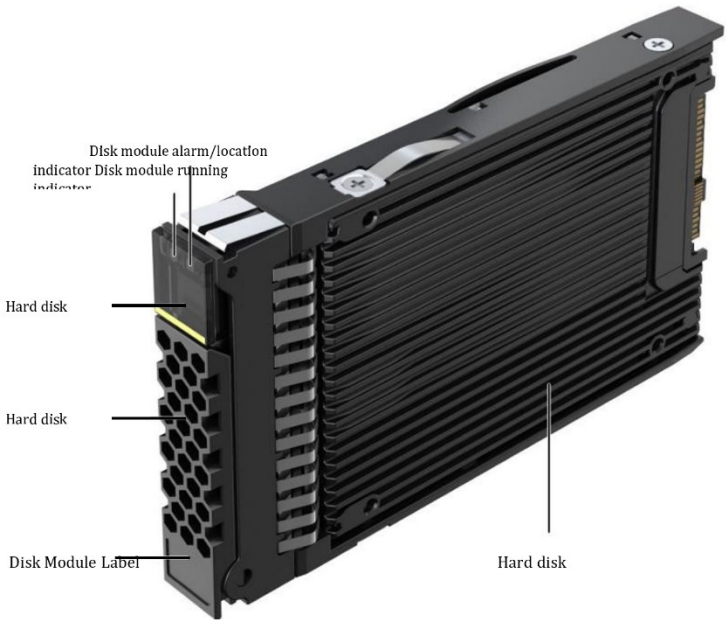

indicator light	Status and Description
Disk module running indicator	Steady green: The disk module is running properly. Blinking green (at 4 Hz or higher): Data is being read and written by the disk module. Off: The disk module is not powered on or fails to be powered on.
Disk module alarm/location indicator	● Steady yellow: The disk module is faulty. ● Blinking yellow (2 Hz): The hard disk is located. ● Off: The disk module is running properly or is hot-swappable.

7.2.3 Indicators

After a disk enclosure is powered on, you can observe indicators to check the current working status of the disk enclosure.

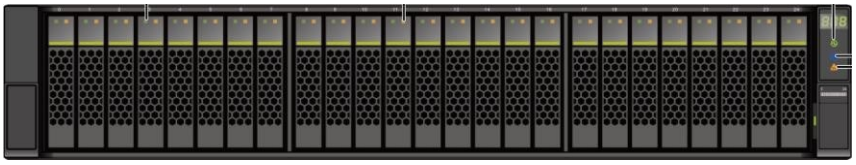

Module	Indicator Type	Status and Description
Disk module	Disk module running indicator	Steady green: The disk module is running properly. Blinking green (at 4 Hz or higher): The module is reading and writing data. Off: The disk module is not powered on or fails to be powered on.
	Disk module alarm/location indicator	Steady yellow: The disk module is faulty. Yellow, blinking at 2 Hz: The location is hard. Off: The disk module is running properly or the disk module is running properly.
System subrack	Disk enclosure location indicator	Blinking blue (2 Hz): The disk enclosure is being Positioning. Off: The disk enclosure is not located.
	Disk enclosure alarm indicator	Steady yellow: An alarm is generated for the disk enclosure. ● Off: The disk enclosure is running properly.
	Power indicator of the disk enclosure	Steady green: The disk enclosure is powered on. Off: The disk enclosure is not powered on.

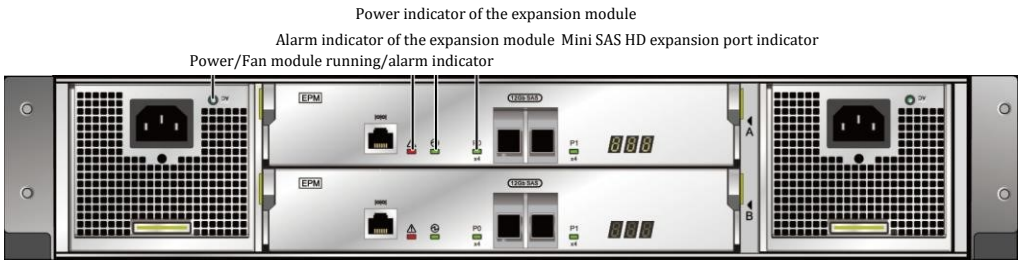

Module	Indicator name	Status and Description
Expansion module	Alarm indicator of the expansion module	● Steady yellow: An alarm is generated for the expansion module. ● Off: The expansion module is running properly.
	Power indicator of the expansion module	● Steady green: The expansion module is powered on. ● Off: The expansion module is not powered on.
	Mini SAS HD expansion port indicator	● Steady blue: The port transmission rate is 4 x 12 Gbit/s. ● Green on: The port transmission rate is 4 x 6 Gbit/s or 4 x. 3 Gbit/s. ● Steady yellow: The port is faulty. ● Off: The port link is not connected.
Power module	Power/Fan module running/alarm indicator	Green, on: The power supply is normal and no fault occurs. Blinking green (1 Hz): The power input is normal and the device is not powered on. Blinking green (4 Hz): The PSU is being upgraded online. Steady yellow: The power supply or fan module is faulty. Off: No external power input is provided.

7.3 SAS disk enclosure (4 U, 3.5-inch disks)

This chapter describes the hardware structure, functions, front and rear views, and indicators of a disk enclosure.

7.3.1 Overview

A disk enclosure uses a modular design and consists of a system enclosure, expansion modules, power modules, fan modules, and disk modules.

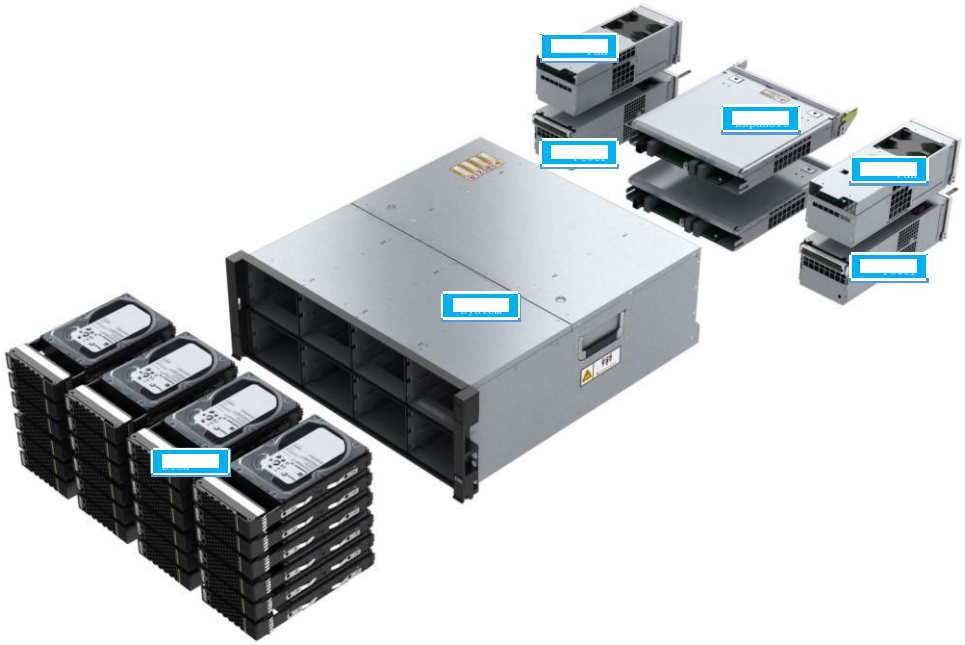

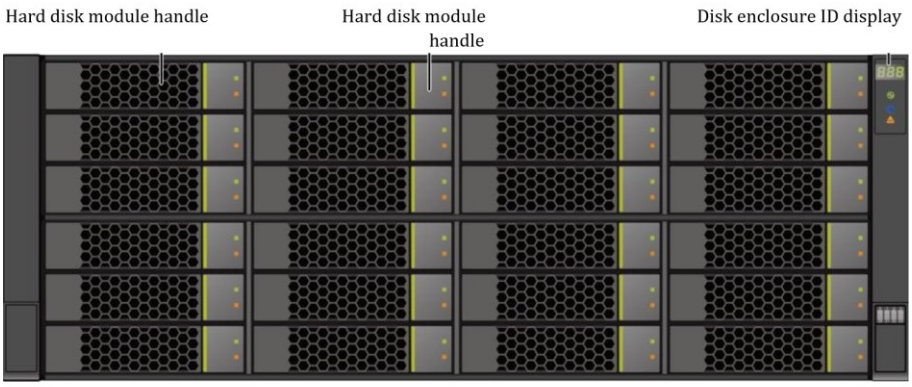

Disk slots are numbered from 0 to 23 from left to right and from top to bottom.

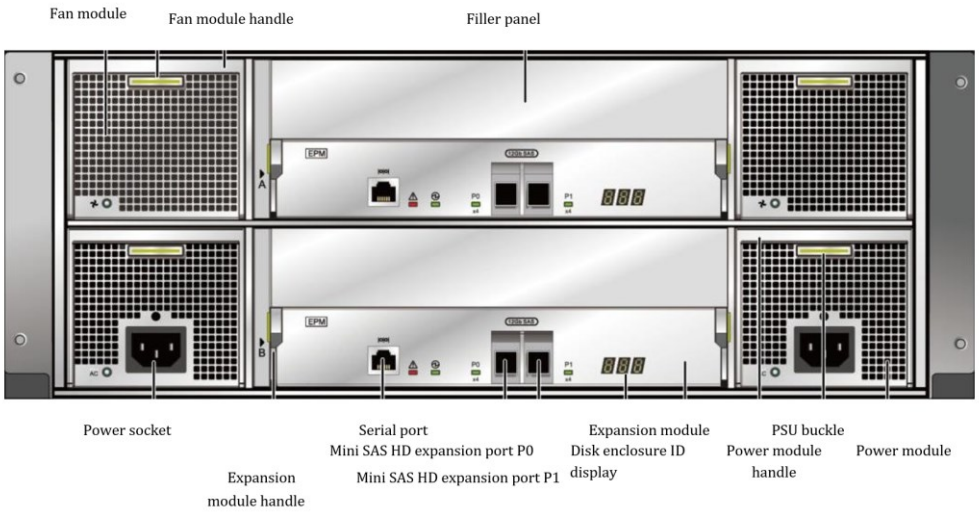

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

Hardware parameters

Name	Parameter
Dimensions (L x W x H)	488mm*447mm*175mm
Weight (excluding the weight of auxiliary materials such as guide rails and cables delivered with the equipment)	● 26.5 kg (excluding the weight of the hard disk unit) ● 43.9 kg (including the weight of the hard disk unit)
AC power supply voltage and rated current	800 W AC power supply (supporting 110 V dual-live-wire input (2 W + PE) (supporting 110 V single-live-wire input) 100–240 V AC ±10%, 10 A, single-phase, 50 Hz/60 Hz
HVDC	800 W power supply (240 V DC input), 192–288 V, 10 A

7.3.2 Component Description

This section describes the appearance and ports of each component in a disk enclosure.

7.3.2.1 System Enclosure

Through the backplane, the system subrack provides reliable connections for various interface modules, implementing signal and power interconnection between modules.

7.3.2.2 Expansion Module

Each expansion module provides two expansion ports P0 and P1. Expansion modules connect to controller enclosures or hard disks through expansion ports.

Enclosure, which enables communication with controller enclosures or disk enclosures and is a connection point for data transmission between controller enclosures and disk enclosures.

Expansion module handle Serial port Mini SAS HD expansion port "P0" Mini SAS HD expansion port "P1"

Alarm indicator of the expansion module

Power indicator of the expansion module

Mini SAS HD expansion port indicator

Disk enclosure ID display

Notices

You can only insert a serial cable into the serial port. Do not insert a network cable into the serial port position.

indicator light	Status and Description
Alarm indicator of the expansion module	● Steady yellow: An alarm is generated for the expansion Module. ● Off: The expansion module is running properly.
Power indicator of the expansion module	● Steady green: The expansion module is powered on. ● Off: The expansion module is not powered on.
Mini SAS HD expansion port indicator	● Steady blue: The port transmission rate is 4 x 12 Gbit/s. ● Green on: The port transmission rate is 4 x 6 Gbit/s or 4 x 3 Gbit/s. ● Steady yellow: The port is faulty. ● Off: The port link is not connected.

7.3.2.3 Power Supply Module

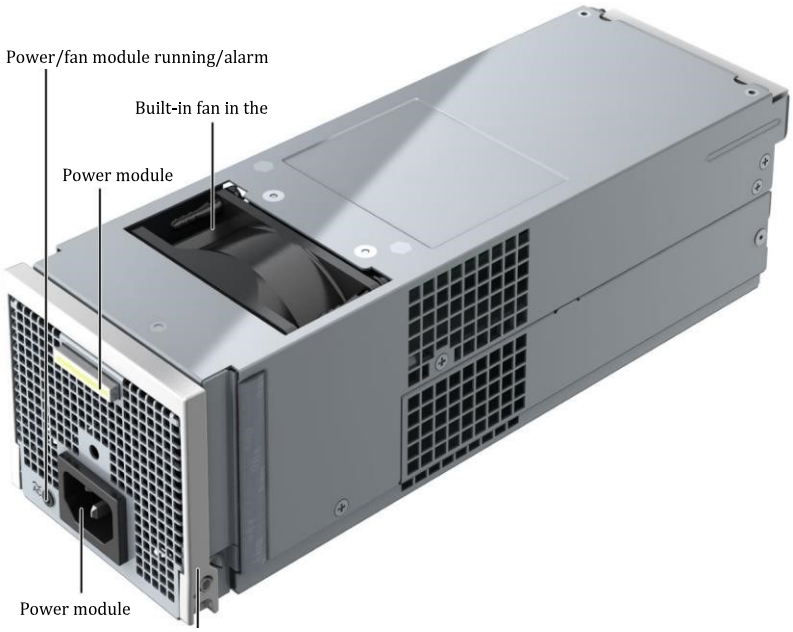

Power module handle Power module handle

Indicator Light	Status and Description
Power module running/alarm indicator	Steady green: The power module supplies power normally and is not faulty. Blinking green (1 Hz): The power input is normal and the device is not powered on. Blinking green at 4 Hz: The PSU is being upgraded online. Steady yellow: The power supply or fan module is faulty. Off: No external power input is provided.

7.3.2.4 Fan Module

Fan modules dissipate heat for the system and ensure that the disk enclosure can run properly at the maximum power consumption.

Fan module running/alarm indicator Fan module running/alarm indicator

indicator light	Status and Description
Fan module running/alarm indicator	● Steady green: The fan module is running properly. ● Steady yellow: The fan module is faulty. ● Off: The fan module is not powered on.

7.3.2.5 Disk Module

Disk modules provide storage capacity for the system to store service data.

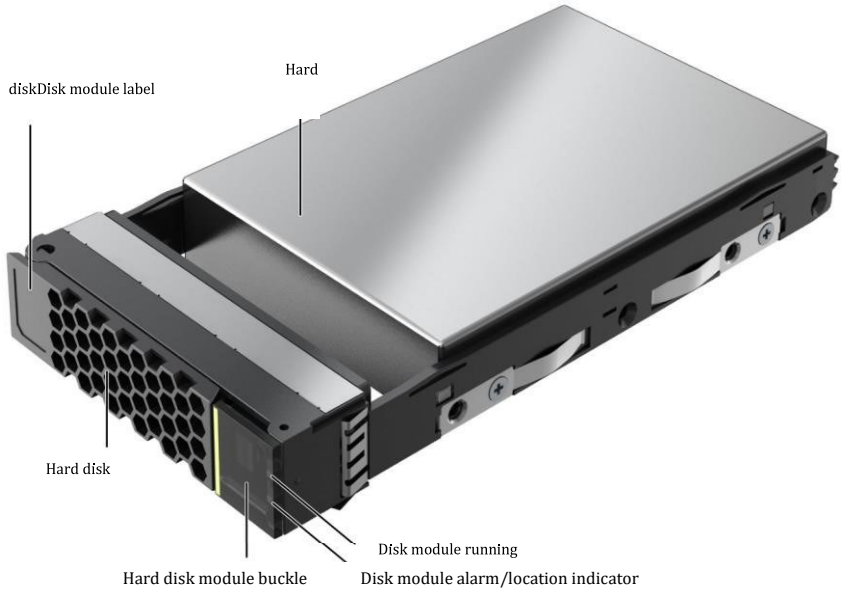

indicator light	Status and Description
Disk module alarm/location indicator	Steady yellow: The disk module is faulty. Blinking yellow (2 Hz): The hard disk is located. Off: The disk module is running properly or the disk module is hot-swappable.
Disk module running indicator	Steady green: The disk module is running properly. Blinking green (at 4 Hz or higher): Data is being read and written by the disk module.

	Off: The disk module is not powered on or fails to be powered on.

7.3.3 Indicators

After a disk enclosure is powered on, you can observe indicators to check the disk enclosure's current operating status.

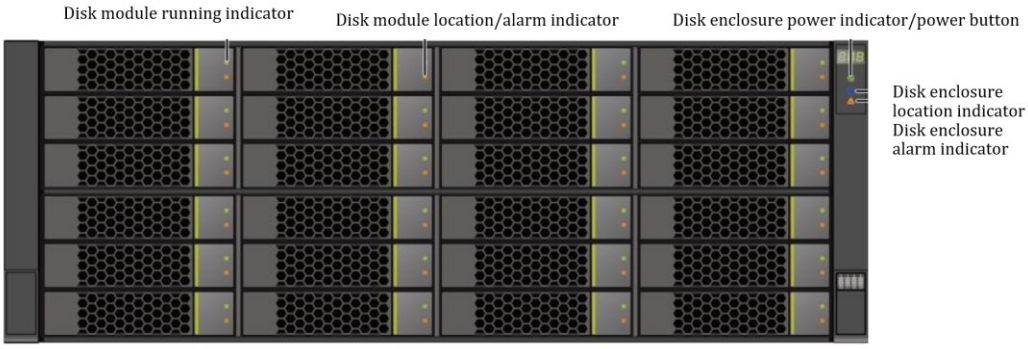

Module	Indicator Type	Status and Description
Disk Module	Disk module running indicator	Steady green: The disk module is running properly. Blinking green (at 4 Hz or higher): The module is reading and writing data. Off: The disk module is not powered on or fails to be powered on.
	Disk module alarm/location indicator	Steady yellow: The disk module is faulty. Yellow, blinking at 2 Hz: The location is hard. Off: The disk module is running properly or the disk module is running properly.
System subrack	Disk enclosure location indicator	Blinking blue (2 Hz): The disk enclosure is being Positioning. Off: The disk enclosure is not located.
	Disk enclosure alarm indicator	Steady yellow: An alarm is generated for the disk enclosure. Off: The disk enclosure is running properly.
	Power indicator of the disk enclosure	Steady green: The disk enclosure is powered on. Off: The disk enclosure is not powered on.

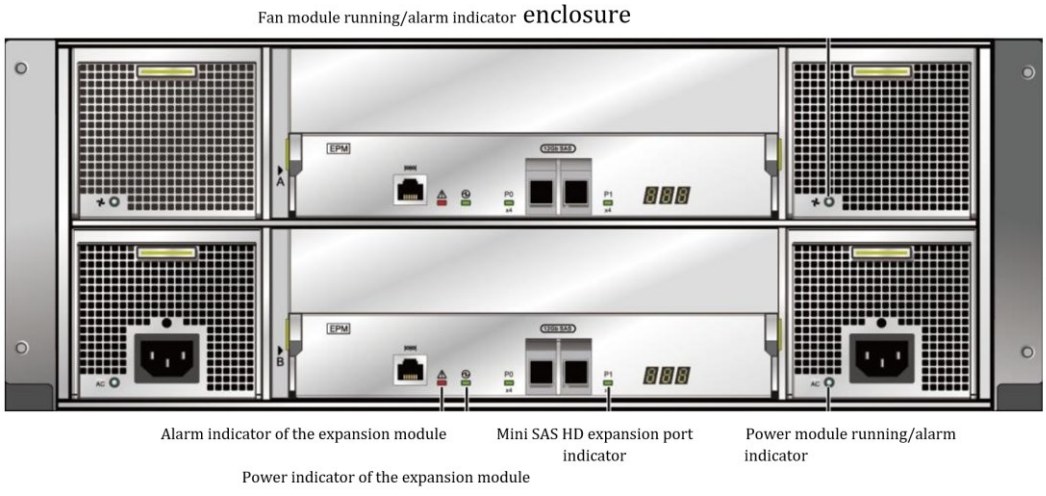

Module	Indicator name	Status and Description
Fan module	Fan module running/alar m indicator	● Steady green: The fan module is running properly. ● Steady yellow: The fan module is faulty. ● Off: The fan module is not powered on.

Power module	Power module running/alar m indicator	Green, on: The power supply is normal and no fault occurs. Blinking green (1 Hz): The power input is normal and the device is not powered on. Blinking green at 4 Hz: The PSU is being upgraded online. Steady yellow: The power module is faulty. Off: There is no external power input.
Expansi on module	Mini SAS HD expansion port indicator	Steady blue: The port transmission rate is 4 x 12 Gbit/s. Steady green: The port transmission rate is 4x6 Gbit/s or 4x3 Gbit/s. Steady yellow: The port is faulty. Off: The port link is not connected.
	Power indicator of the expansion module	Steady green: The expansion module is powered on. Off: The expansion module is not powered on.
	Alarm indicator of the expansion module	Steady yellow: An alarm is generated for the expansion module. Off: The expansion module is running properly.

7.4 Device Cables

The cables used by the OS 2910 include power cables, ground cables, and signal cables. This section describes the cable diagrams and functions, and describes the cable details.

7.4.1 Power cable

A power cable supplies power to the devices in the cabinet. One end of the power cable is connected to the power port of the device, and the other end is connected to the external power supply.

Each AC power module in the controller enclosure and disk enclosure is configured with a PDU power cable.

7.4.2 Grounding cable

Grounding cables are used to ground the equipment and improve the safety of operating the equipment.

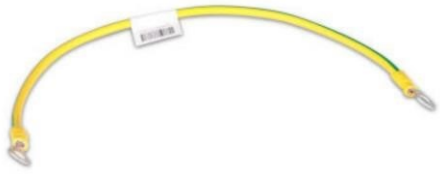

7.4.3 Network Cable

A network cable is required to connect the management network port or other network ports of the device. The device communicates with other devices or servers through the network cable. The device can communicate with the outside through the network cable. One end of a network cable can be connected to the management network port or other network ports of the device, and the other end can be connected to a network switch or an application server.

7.4.4 Serial Cables

After the serial port is connected to the maintenance terminal through a serial cable, you can manage and maintain the shelf through the serial port. This command is used when you cannot log in to the device through the management network port. One end of the serial cable is an RJ-45 connector, which is used to connect to the serial port of the shelf. The other end is a DB-9 port, which is used to connect to the maintenance terminal.

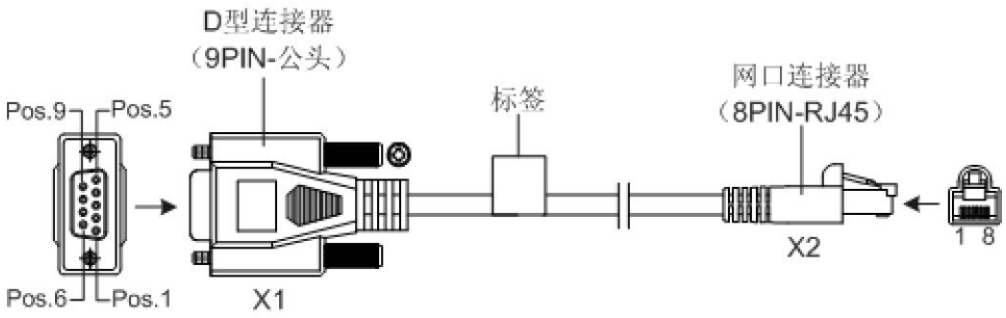

X1 (DB9)	X2 (RJ45)	signalling
8	1	CTS (Clear to Send)
6	2	DSR (Data Set Ready)
2	3	RXD (Receive Data)
5	4	GND
5	5	GND
3	6	TXD (Transmit Data)
4	7	DTR (Data Terminal Ready)
7	8	RTS (Request to Send)

7.4.5 Mini SAS HD Cable

Mini SAS HD cables are used to connect controller enclosures to disk enclosures and disk enclosures to disk enclosures.

7.4.6 Optical Fiber

The device communicates with the FC switch through optical fibers. One end of the optical fiber can be connected to the optical port of the device, and the other end can be connected to the FC switch or application server. Both ends of the optical fiber are LC connectors.

When connecting cables, select cables based on the actual requirements and the labels on the cables.