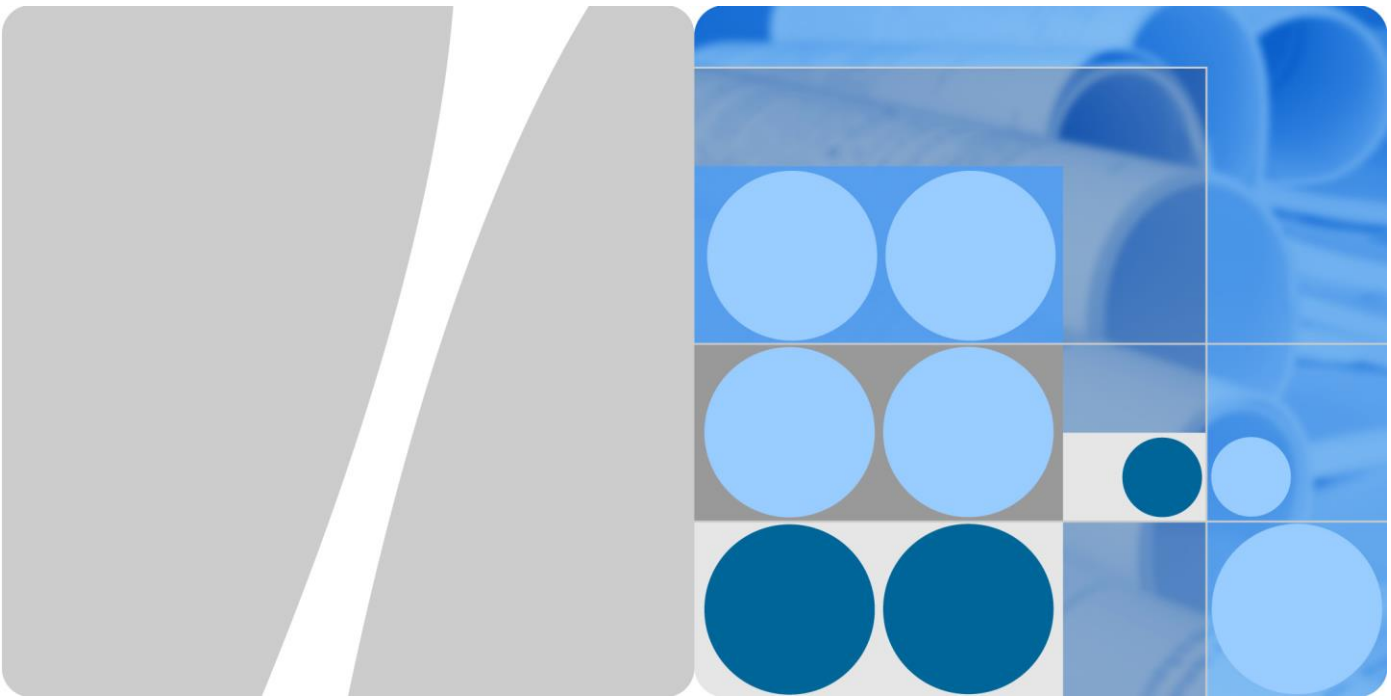

FusionSphere Virtualization Suite 6.3.1

Product Description

Issue 01

Date 2018-10-20

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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://www.huawei.com

Email: support@huawei.com

Contents

1 Introduction to FusionSphere Virtualization Suite

1.1 FusionSphere Virtualization Suite Overview

Huawei FusionSphere virtualization suite is an industry-leading virtualization solution. This solution significantly improves data center infrastructure efficiency and provides the following benefits for customers:

- Improve infrastructure resource utilization data centers.
- Significantly accelerate service rollout.
- Substantially reduce power consumption in data centers.
- Leverage high availability and powerful restoration capabilities of virtualized infrastructure to provide rapid fault recovery for services, thereby cutting data center costs and increasing system runtime.

FusionSphere virtualization suite virtualizes hardware resources using the virtualization software deployed on physical servers, so that one physical server can function as multiple virtual servers. FusionSphere maximizes resource utilization by centralizing existing VMs workloads on some servers and therefore releasing more servers to carry new applications and solutions.

Application Scenarios

- **Single-Hypervisor Scenarios**

  Single-hypervisor applies to scenarios in which an enterprise only uses FusionCompute as a unified operation, maintenance, and management platform to operate and maintain the entire system, including monitoring resources, managing resources, and managing the system.

  FusionCompute virtualizes hardware resources and centrally manages virtual resources, service resources, and user resources. It virtualizes computing, storage, and network resources using the virtual computing, virtual storage, and virtual network technologies. FusionCompute centrally schedules and manages virtual resources using a unified interface, thereby reducing the operating expense (OPEX) and ensuring high system security and reliability.

- **Multi-Hypervisor Scenarios**

  FusionManager used in a multi-hypervisor scenario provides the following functions:

- Unified management and maintenance of resources and services provided by both the Huawei FusionCompute and VMware vCenter hypervisors added to the system.
- Unified alarm monitoring, reporting, and management for hypervisors and physical devices added to the system.

- **Private Cloud Scenarios**

  Private clouds apply to scenarios in which enterprise departments need to independently manage their virtual resources and services. During service provisioning, administrators and tenants separately implement their tasks. Administrators can manage all resources within the system while tenants can manage only resources within virtual data centers (VDCs) they belong to.

  Based on service requirements, private cloud scenarios can be categorized into multi-tenant shared virtual private cloud (VPC) and multi-tenant private VPC scenarios.

  - **Multi-tenant shared VPC scenario**

    In this scenario, administrators allocate virtual resources in the VDCs to tenants. When creating VMs in the tenant view, tenants can use the network created by administrators in a shared VPC. If an administrator is also added to the VDC during VDC creation, the administrator can perform all the tasks in the administrator view and tenant view. This scenario applies to the following conditions that:

    - Enterprise departments need to separately manage their own services and virtual resources (including VMs and disks but excluding networks).
    - The networks to be used by all enterprise departments are centrally planned and maintained by the administrator.

  - **Multi-tenant private VPC**

    In this scenario, administrators allocate virtual resources in the VDCs to tenants. Tenants can create VPC, network, or VMs in tenant view. If the tenant need to share networks, the administrator can create a shared VPC for all tenants to use. If the administrator is also added to the VDC during VDC creation, the administrator can perform all the tasks in administrator view and tenant view. This scenario applies to the following conditions that:

    - Enterprise departments need to separately manage their own services and virtual resources (including networks, VMs, and disks).
    - Enterprise departments need to separately plan, create, and maintain their own networks.

1.2 FusionSphere Virtualization Suite Characteristics and Functions

On-demand Resource Allocation for Applications

FusionSphere allows users to add or reduce VM resources on demand anytime without interrupting applications.

Virtual Resource SLA

FusionSphere allows users to define service level agreement (SLA) policies to control VM resources, thereby allocating physical resources based on application importance.

Centralized VDC Management

FusionSphere provides users with VM management capabilities, including creating, deploying, converting, and migrating VMs. FusionSphere also hides differences between heterogeneous mainstream virtualization platforms to allow users to manage the platforms in a unified manner.

High Compatibility

FusionSphere supports x86-based servers, various storage devices, mainstream Linux and Windows operating systems (OSs), allowing mainstream applications to run on virtualization platforms.

Automatic Scheduling

FusionSphere automatically migrates workloads based on preset policies, thereby optimizing resource allocation, system response efficiency, and user experience.

Comprehensive Rights Management Functions

FusionSphere provides comprehensive rights management functions, allowing authorized users to manage system resources based on their specific roles and assigned permissions.

Intelligent Application Management

FusionSphere provides the sophisticated approval mechanism-based service catalog service and user-defined template service, facilitating custom application deployment.

Sophisticated Metering

FusionSphere collects information about resource usage for each user and reports the statistics to third-party systems to calculate service charges.

Various O&M Functions

FusionSphere provides various operation and maintenance (O&M) functions and tools to improve system O&M efficiency.

Cloud Security

FusionSphere is compliant with local information security laws and regulations and incorporates various security measures to provide end-to-end protection for user access, management and maintenance, data, networks, and virtualization services.

1.3 FusionCompute Product Orientation

The FusionCompute is a cloud operating system (OS). It virtualizes computing, storage, and network resources, and implements centralized management and scheduling of the virtual resources through a unified interface.

The FusionCompute provides high system security and reliability and reduces operational costs. It helps carriers and enterprises build secure, green, and energy-saving data centers.

Figure 1-1 shows the FusionCompute position in the virtualization suite.

Figure 1-1 FusionCompute position in the Huawei cloud computing solution

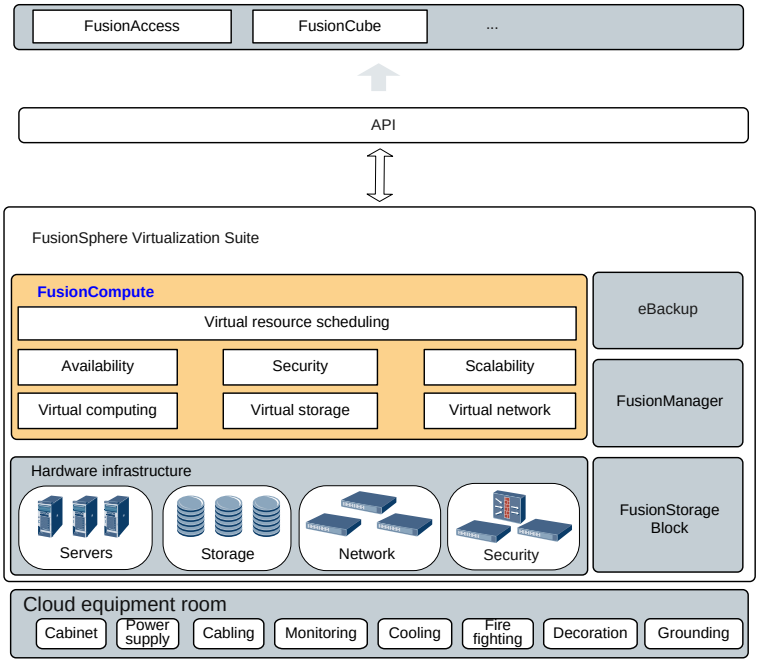

Cloud Facilities

Cloud facilities refer to the auxiliaries and space required by the cloud data center, including the power supply system, fire-fighting system, wiring system, and cooling system.

Huawei has been devoted to continuously enhancing the competitiveness of the data center based on the concept called SAFE, which focuses on smartness, availability, flexibility, and efficiency.

FusionSphere virtualization suite

The FusionSphere virtualization suite virtualizes hardware resources using the virtualization software deployed on physical servers, so that one physical server can be used as multiple virtual servers. The FusionSphere allocates all current workloads to VMs on some servers, so that new applications and solutions can be deployed on the servers whose original workloads are migrated to other VMs.

- Hardware Infrastructure Layer

  Hardware infrastructure consists of servers, storage devices, network devices, and security devices. These resources allow customers to build different scale systems and expand its capacity based on actual needs and to use applications ranging from entry level to enterprise level. Various devices provide customers with multiple and flexible choices.
- FusionManager

The FusionManager monitors and manages hardware and software of cloud computing. It provides automatic resource provisioning and automatic operation and maintenance (O&M) for the infrastructure. Additionally, it provides a web user interface (UI) to administrators to operate and manage the resources in the system.

- FusionStorage Block

  FusionStorage Block is the distributed storage software that integrates storage and computing capabilities. It can be deployed on general-purpose x86 servers to consolidate the local disks on all the servers into a virtual storage resource pool that provides the block storage function.

- eBackup

  The VM backup scheme uses the Huawei eBackup backup software combined with the snapshot backup function and the Changed Block Tracking (CBT) backup function of the FusionCompute to back up VM data.

1.4 FusionCompute Product Features

Unified Virtualization Platform

FusionCompute uses virtualization management software to create high-performance, operable, and manageable virtual machines (VMs) over computing resources. FusionCompute provides the following functions:

- Allocating VM resources on demand.
- Supporting various operating systems (OSs).
- Isolating VMs to ensure quality of service (QoS).

Support for Various Hardware Platforms

FusionCompute supports various x86 servers and is compatible with various storage devices.

Big Cluster

A cluster supports up to 64 hosts and 3000 VMs.

Automatic Resource Scheduling

FusionCompute allows users to define service-level agreement (SLA) policies, fault reporting criteria, and fault rectification policies.

- FusionCompute implements centralized IT resource scheduling, heat management, and energy consumption management, reducing maintenance costs.
- FusionCompute dynamically schedules resources based on the load of servers and services, achieving load balancing across servers and service provisioning systems and ensuring optimal system response and user experience.

Rights Management

FusionCompute provides rights management and allows users to manage system resources based on their roles and rights.

Comprehensive O&M

FusionCompute provides various operation and maintenance (O&M) tools to control and manage services, improving O&M efficiency. FusionCompute provides the following operation and maintenance tools:

- Black box

  The black box enables carriers or enterprises to rapidly locate faults based on logs and program heaps. It reduces the time to locate fault and improves O&M efficiency.
- Automatic health check

  The automatic health check helps FusionCompute automatically detect system faults in a timely manner and generate alarms, ensuring timely O&M of VMs.
- Web interfaces

  FusionCompute provides web interfaces, through which users can monitor and manage all hardware resources, virtual resources, and service provisioning.

Cloud Security

FusionCompute complies with local information security laws and regulations. It adopts various security measures and policies to provide end-to-end protection to user access, management and maintenance, data, networks, and virtualization.

1.5 FusionManager Product Positioning

FusionManager provides the following management functions for the cloud platform: service management, service automation, and service and resource assurance, as shown in Figure 1-2.

Figure 1-2 FusionManager product positioning

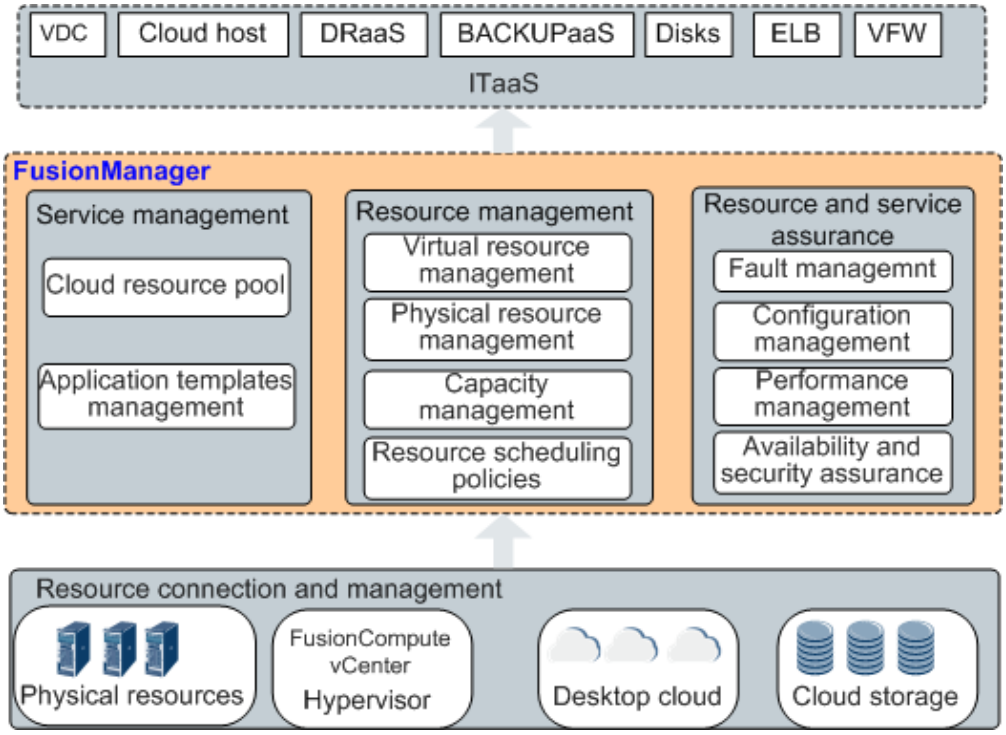

- In regard to underlying resources that are supported by FusionManager, FusionManager can manage physical resources, virtualization software developed by both Huawei and third parties, desktop cloud, cloud storage, and various other cloud services. Physical resources can be computing, storage, and network devices.
- In regard to upper-layer interoperability, FusionManager provides various IT as a service (ITaaS) functions for upper-layer system users.
- In regard to the capabilities of FusionManager, it supports connection of both virtual and physical resources to the system, and also supports automatic management and availability and security assurance for the added resources.

In conclusion, with its featured automatic cloud service management and intelligent resource operation and maintenance capabilities, FusionManager is positioned to provide simple and agile management functions for cloud data centers.

1.6 FusionManager Product Features

FusionManager supports automatic cloud service management and intelligent resource operation and maintenance. It aims to provide simple and agile management functions for cloud data centers, and supports various cloud virtualization solutions, such as Huawei's FusionSphere and ManageOne.

2 System Architecture

2.1 Logical Architecture of FusionSphere Virtualization Suite

Figure 2-1 shows the logical architecture of FusionSphere virtualization suite.

Figure 2-1 Logical architecture of FusionSphere virtualization suite

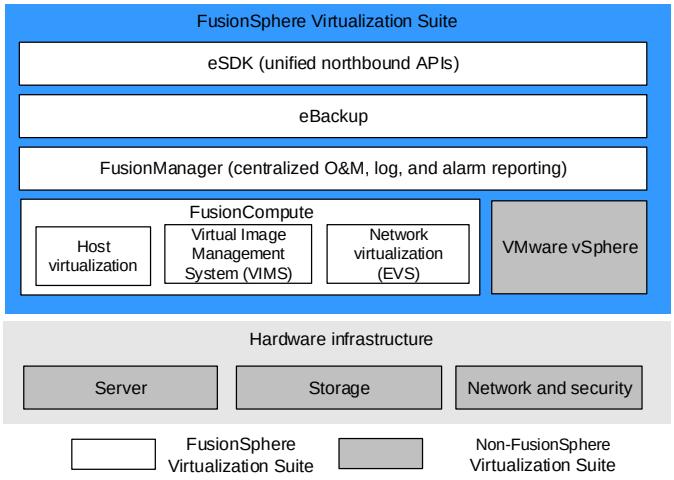

Table 2-1 describes the key components in the FusionSphere virtualization suite.

Table 2-1 Key components in the FusionSphere virtualization suite

Component	Description
FusionCompute	Mandatory FusionCompute is a cloud operating system (OS). It virtualizes hardware resources and centrally manages virtual resources, service resources, and user resources. FusionCompute virtualizes computing, storage, and network resources using the virtual computing, virtual storage, and

Component	Description
	virtual network technologies. It centrally schedules and manages virtual resources over unified interfaces. FusionCompute provides high system security and reliability and reduces operational costs, helping carriers and enterprises build secure, green, and energy-saving data centers.
FusionManager	Optional FusionManager is the cloud management software. It monitors and manages hardware and software on the cloud platform. For example, it manages both homogeneous and heterogeneous (VMware vCenter) cloud resource pools, and reports alarms for both hardware and software devices. FusionManager also offers a web user interface (UI) for administrators to perform O&M operations on the resources.
FusionSphere eSDK	Optional FusionSphere Ecosystem Software Development Kit (eSDK) contains unified northbound interfaces of FusionSphere virtualization suite, allowing third-party network management systems (NMSs) and O&M systems to seamlessly connect to FusionSphere. FusionSphere eSDK makes FusionSphere virtualization suite capabilities open, including VM lifecycle management, advanced virtualization functions, and O&M functions.
eBackup	Optional eBackup is a piece of backup software for virtualization. It leverages FusionCompute snapshot function and the Changed Block Tracking (CBT) function to back up VM data.

2.2 FusionCompute Logical Architecture

Figure 2-2 shows the architecture of FusionCompute.

Figure 2-2 Architecture of FusionCompute

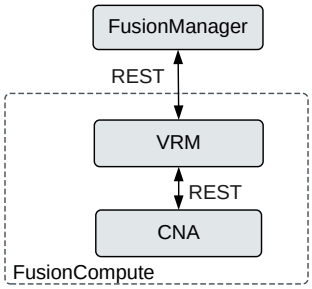

CNA: Computing Node Agent VRM: Virtualization Resource Management

Table 2-2 describes the functions of FusionCompute components.

Table 2-2 Functions of FusionCompute components

Module	Function
CNA	The CNA provides the following functions: • Implementing the virtual computing function. • Managing the VMs running on the CNA. • Managing the computing, storage, and network resources of the CNA.
VRM	The VRM provides the following functions: • Managing block storage resources in the cluster. • Manages network resources, such as IP addresses, virtual local area network (VLAN) numbers, and DHCP severs in the cluster and allocates IP addresses to VMs. • Managing the life cycle of VMs in the cluster and distributing and migrating VMs across CNAs. • Dynamically adjusting resources in the cluster. • Implementing centralized management of virtual resources and user data and providing elastic computing, storage, and IP address services. • Allowing O&M engineers to remotely access FusionCompute through a web interface to perform resource monitoring and management and view resource statistics reports.

2.3 FusionManager Logical Architecture

Figure 2-3 shows the logical architecture of FusionManager.

Figure 2-3 FusionManager logical architecture

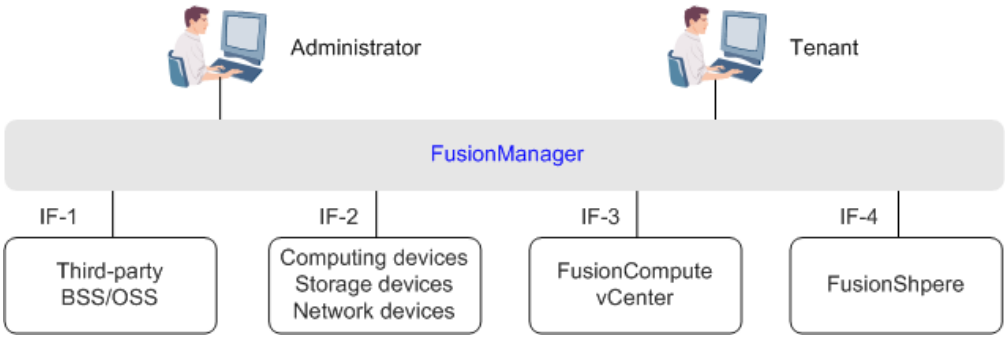

2.4 Interfaces and Protocols

The FusionCompute provides standard and open protocol interfaces to interwork with various devices. Figure 2-4 shows the interfaces and protocols supported by the FusionCompute.

Figure 2-4 Interfaces and protocols supported by the FusionCompute

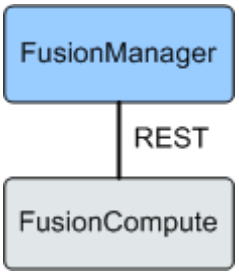

Table 2-3 describes the FusionCompute interfaces.

Table 2-3 FusionCompute interfaces

Interface Name	Interface Function	Interworking Entity	Protocol
External interfaces provided by the Virtualization Resource Management (VRM)	Allow the FusionManager to obtain configuration and alarm information about virtualized resources using FusionCompute. Allow the FusionCompute to obtain instructions from the FusionManager to manage VMs.	FusionManager	REST

Table 2-4 lists the interfaces and protocols used by FusionManager.

Table 2-4 Interfaces and protocols used by FusionManager

Interface Number	Interface Type	Applicable Subsystem	Function
IF1	REST	FusionManager<->upper-layer network management system (NMS)	FusionManager uses REST interfaces to communicate with the upper-layer NMS system, for example, to report alarms and performance statistics to the NMS system. The upper-layer NMS system identifies and authenticates FusionManager using Hypertext Transfer Protocol (HTTP) or Hypertext Transfer Protocol Secure (HTTPS).
IF-2	SNMP, IPMI, SSH, HTTP, HTTPS, TLV, or SMI-S	FusionManager<->computing, storage, and network devices	Computing, storage, and network devices can use Simple Network Management Protocol (SNMP), Intelligent Platform Management Interface (IPMI), Secure Shell (SSH), Hypertext Transfer Protocol (HTTP), Hypertext Transfer Protocol Secure (HTTPS), Type-Length-Value (TLV), or Storage Management Initiative—Specification (SMI-S) protocols to connect to FusionManager.
IF-3	REST, SOAP	FusionManager<->third-party operating or operation and maintenance (O&M) system	FusionManager uses northbound interfaces to provide third-party operating or O&M systems with various functions, including resource management, VM management, backup management, disk management, and network management.
IF-4	REST	Local FusionManager<->ServiceCenter	Local FusionManager systems use REST interfaces to communicate with the ServiceCenter system.

3 Deployment Plan

3.1 Small-Scale Deployment Plan

In a small-scale deployment scenario, the number of hosts in the system ranges from 3 to 50, and the number of VMs ranges from 1 to 1000.

Table 3-1 Small-scale deployment plan

Component	Deployment Mode	Specifications
Virtualization Resource Management (VRM), the FusionCompute management component **NOTE** A Virtualization Resource Management (VRM) node is a FusionCompute management node in the FusionSphere system and manages resources in host clusters.	Deployed in 1+1 active/standby mode on physical servers or VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
FusionManager	Deployed in 1+1 active/standby mode on VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
eBackup	Deployed on a VM	It is advisable to deploy one eBackup system for every 200 VMs. For details about configuration of other resources, see **User Guide** > **Virtual Backup** > **Overview** > **Technical Specifications** in the *OceanStor BCManager eBackup Product*

Component	Deployment Mode	Specifications
		Documentation.

3.2 Medium-Scale Deployment Plan

In a medium-scale deployment scenario, the number of hosts in the system ranges from 51 to 200, and the number of VMs ranges from 1001 to 5000.

Table 3-2 Medium-scale deployment plan

Component	Deployment Mode	Specifications
Virtualization Resource Management (VRM), the FusionCompute management component	Deployed in 1+1 active/standby mode on physical servers or VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
FusionManager	Deployed in 1+1 active/standby mode on VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
eBackup	Deployed on a physical server or VM	It is advisable to deploy one eBackup system for every 200 VMs. For details about configuration of other resources, see **User Guide** > **Virtual Backup** > **Overview** > **Technical Specifications** in the *OceanStor BCManager eBackup Product Documentation*.

3.3 Large-Scale Deployment Plan

In a large-scale deployment scenario, the number of hosts in the system ranges from 201 to 1000, and the number of VMs ranges from 5001 to 10,000.

Table 3-3 Large-scale deployment plan

Component	Deployment Mode	Specifications
Virtualization Resource Management (VRM), the FusionCompute management component	Deployed in 1+1 active/standby mode on physical servers or VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
FusionManager	Deployed in 1+1 active/standby mode on VMs	Network interface card (NIC): 2 x 10 Gbit/s (recommended) For details about configuration of other resources, see 9.1 Deployment Rules
eBackup	Deployed on a physical server	It is advisable to deploy one eBackup system for every 200 VMs. For details about configuration of other resources, see **User Guide** > **Virtual Backup** > **Overview** > **Technical Specifications** in the *OceanStor BCManager eBackup Product Documentation*.

3.4 FusionCompute Host Requirements

A host is a physical server where the virtualization software is deployed. Table 3-4 lists the host requirements.

NOTE
If the hosts have been used before, restore the hosts to factory settings before configuring the basic input/output system (BIOS).

Table 3-4 Host requirements

Item	Requirement
CPU	• Intel 64-bit CPU • The CPU supports hardware virtualization technology, such as Intel VT-x, and the BIOS system must have the CPU virtualization function enabled. • The CPUs in one cluster must have the same model. Otherwise,

Item	Requirement
	the VM migration between hosts may fail. To ensure CPU consistency, you can use the servers of the same model in one cluster. • If the host is connected to FusionStorage Block, configure the CPUs based on the CPU configuration section provided in **Installation and Commissioning** > **Installation Process (FusionCompute)** > **Introduction** > **System Requirements** in the FusionStorage Block product documentation of the required version. • If eBackup is connected, two more vCPUs need to be reserved for the host management domain. **NOTICE** If the CPU virtualization function is disabled on a host, VMs cannot be created on the host.
Memory	Minimum memory size: 8 GB If the host is used to deploy a management VM, the host memory must be greater than or equal to the total of the management VM memory and the host management domain memory. If you need to configure the user-mode switching specification for a host, reserve another 5 GB to the original host management domain memory size. If the host connects to FusionStorage Block, the reserved memory of the management domain must be greater than the total of the reserved memory of the management domain before the host is connected to FusionStorage Block and the FusionStorage Block memory. For details about the FusionStorage Block memory requirements, see **Installation and Commissioning** > **Installation Process (FusionCompute)** > **Introduction** > **System Requirements** in the FusionStorage Block product documentation of the required version. Recommended memory size: ≥ 48 GB When Huawei servers are installed, the memory needs to be set based on the recommended configuration. Otherwise, the system cannot achieve optical performance. For details about the recommended configuration, see http://support.huawei.com/onlinetoolsweb/smca/. **NOTE** V5 servers must use recommended configurations. Otherwise, the system performance deteriorates obviously.
Hard disk	Minimum memory size: 70 GB For computing hosts, it is recommended that two more than 300 GB SAS disks with the available space of 70 GB form RAID 1 used as the system disk. If memory overcommitment is to be used, two independent SAS disks are required to form RAID 1 used for memory overcommitment. Memory consumed for memory overcommitment = Memory of a server x 60% If service VMs use local storage, you need to plan independent local storage for them. It is recommended that local disks form RAID 1 to provide storage space.

Item	Requirement
Network port	• Number of network ports: ≥ 1 • Recommended number of network interface cards (NICs): 6 • Recommended network port rate: > 1000 Mbit/s • Use the GE or 10GE network depending on the amount of the estimated network traffic. The network load should be lower than 60% of the bandwidth.
Redundant Array of Independent Disks (RAID)	• For the server accommodating the VRM node: If the number of disks on the server is greater than 2, form RAID 10. If the number of disks on the server is 2, form RAID 1. • The configurations of compute nodes are as follows: – Configure hard disks 1 and 2 as RAID 1 for installing the host OS to improve storage reliability. When setting the boot device in the host BIOS, set the first boot device to a RAID 1 disk. – If the host has multiple disks, set up RAID 5 with all the disks except disks 1 and 2. – If the customers have special requirements on RAID, adjust the configuration to suit their requirements. **NOTE** The RAID cards on certain servers require that server disks must form RAID arrays. Otherwise, the host OS cannot be installed. For details about the RAID card requirements, see the product documentation delivered with the servers.

3.5 Networking Security and Openness

This section describes the networking security protection policies of the FusionManager system and provides suggestions on secure usage.

Logical Deployment

Figure 3-1 shows the logical deployment of FusionManager.

Figure 3-1 Logical deployment

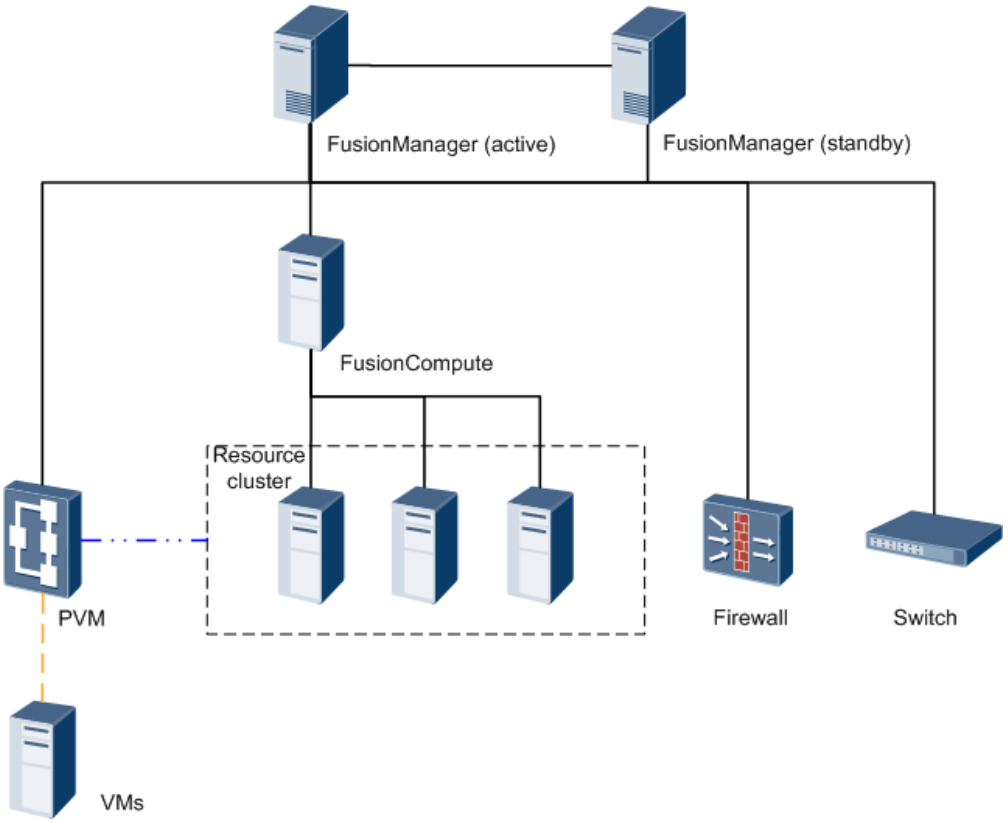

FusionManager nodes are deployed in active/standby mode. The FusionCompute node interworks with FusionCompute to manage virtual resources, switches, and firewalls.

The Proxy VM (PVM) is automatically created during service deployment and provides software installation and application instance monitoring services for user VMs.

Logical Networking

Figure 3-2 shows the logical networking of FusionManager.

Figure 3-2 Logical networking

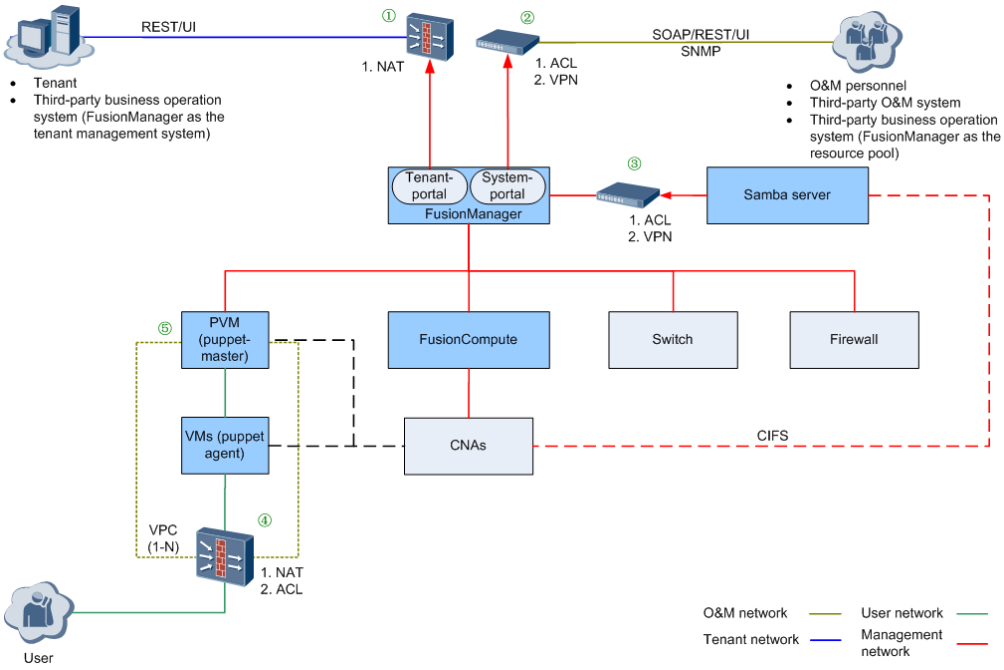

Observe the following security protection rules in the logical networking:

1. Enable the Network Address Translation (NAT) function of the firewall to map the tenant application programming interfaces (APIs) and user interfaces (UIs), sort ports by the criteria of the External+Tenant side+Administrator side/Tenant side type in *FusionManager Communication Matrix (Server Virtualization)* of the FusionManager system to an untrusted network, thereby opening these interfaces. In this way, tenants can manage tenant resources using the tenant UIs, and third-party business operation systems can manage tenant resources using tenant APIs.
2. Operation and maintenance (O&M) personnel can manage system resources using the management UIs, and third-party O&M systems and business operation systems can manage system resources using management APIs. The third-party O&M systems and business operation systems reside on the same trusted network (trust zone) as the FusionManager system and its managed devices, and they communicate with one another over virtual private networks (VPNs). In this case, the device network and system internal ports are visible to O&M systems and business operation systems.

   If the customer poses high requirements for system security, for example, requiring isolation from the device plane, configure access control lists (ACLs) (sort ports by the criteria of the External+Management plane+Management plane/Tenant plane type in *FusionManager Communication Matrix (Server Virtualization)* and enable these ports) on the switches. This method is recommended to improve system security.
3. Connect the Samba server provided by the customer to the FusionManager system, enabling users to upload and download ISO images and software. Typically, the Samba server resides on the same trusted network (trust zone) as the FusionManager system and communicates with it over the VPN.

   The customer shall assess and ensure the security hardening and defense of the Samba server. To ensure network system security, Huawei recommends that:

   – If the Samba server resides on different trusted networks from the FusionManager system, configure ACL rules on the gateway switches to ensure security.

- If the Samba server resides on the same trusted network as the FusionManager system, configure VPNs to enable their communication.

4. To protect access to user VMs, configure ACL rules or NAT policies for packet filtering.
5. The software installation server (PVM) connects to both the management network and the user network. To prevent malicious users from intruding into the management network, enable the PVM server to use IPTABLE for isolating the user plane.

Access to Tenant APIs and Tenant UI

As shown in **1** in Figure 3-2, the firewall NAT function enables the system to map the FusionManager tenant API and tenant UI ports to an untrusted network. This function implements secure access to tenant APIs and tenant UI.

For details about the supported ports, view the *FusionManager Communication Matrix (Server Virtualization)* document and filter out the **External+Tenant plane+Management plane/Tenant plane** ports.

The following section provides an example of configuring FusionManager tenant API and tenant UI accessible from external users.

Configuring FusionManager Tenant API and Tenant UI Accessible from External Users

This example provides you the typical networking that allows FusionManager tenant API and tenant UI to be accessible from external users and also describes the configuration procedure. Figure 3-3 shows the networking diagram.

NOTE

If the FusionManager management IP address has been changed after the configuration, you need to configure tenant API and tenant UI accessibility again.

Figure 3-3 Typical networking

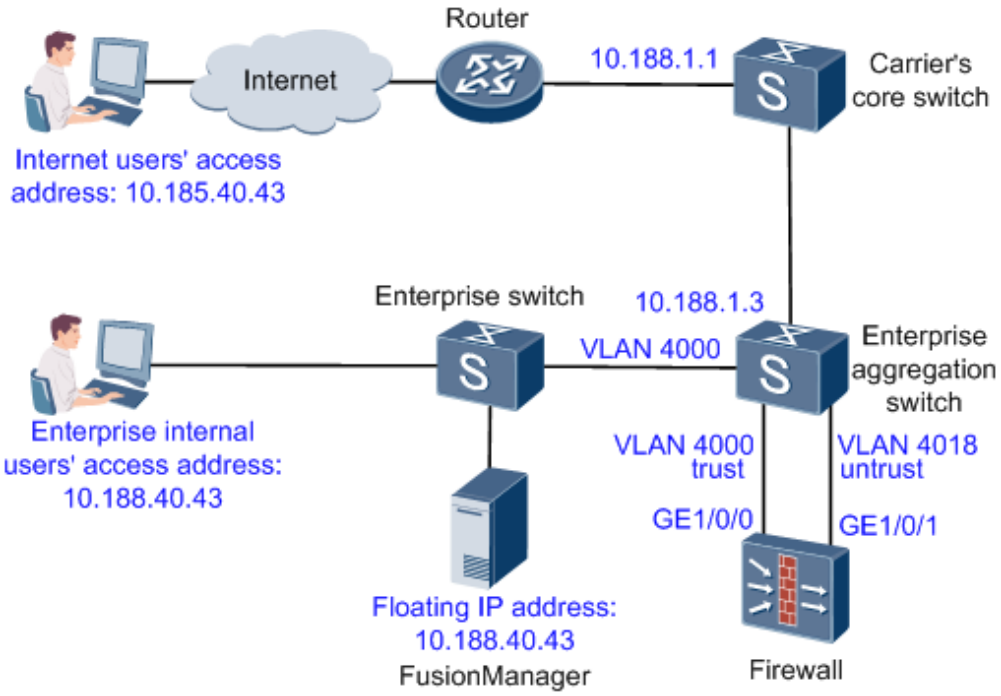

Description

- The carrier network is already connected to the Internet, and Internet users' access to the FusionManager public IP address **10.185.40.43** can be correctly routed to the carrier's core switch **10.188.1.1**.
- The FusionManager floating IP address is **10.188.40.43**.
- The following FusionManager tenant API and tenant UI ports are required to be open for external users' access.
  - TCP ports: 21, 80, 443, 543, 6081, 30000 to 30100
  - For details about the supported ports, view the *FusionManager Communication Matrix (Server Virtualization)* document and filter out the **External+Tenant plane+Management plane/Tenant plane** ports.
- The physical gateway location of FusionManager is Firewall, and the firewall networking mode is static-route.
- The GigabitEthernet 1/0/0 port on the firewall is used as the trust port, and the GigabitEthernet 1/0/1 port is an untrust port.
- The following key network devices are used:
  - Firewall: Huawei Eudemon8000E-X3&8000E-X8&8000E-X16 V200R001C01SPC900
  - Enterprise aggregation switch: Huawei Quidway S5300 V200R001C00SPC300

Procedure

1. Configure the carrier's core switch.

   #The following command enables the core switch to receive Internet users' FusionManager access requests and route the requests to the enterprise aggregation switch:

ip route-static 10.185.40.43 255.255.255.255 10.188.1.3

2. Configure the enterprise aggregation switch.

   #The following commands are used to configure the VLAN for the system to connect to the firewall:

   Vlan 4018

   interface Vlanif4018

   ip address 192.168.185.1 255.255.255.252

   ip route-static 0.0.0.0 0.0.0.0 10.188.1.1 (If this route already exists, you do not need to configure it again.)

   #The following command enables the system to route the received FusionManager network segment access requests from users to the firewall:

   ip route-static 10.0.0.0 8 192.168.185.2

3. Configure the firewall.

   #The following commands are used to configure the virtual firewall vfwfmnat used by the FusionManager management network:

   ip vpn-instance vfwfmnat

   route-distinguisher 121:1

   #The following commands are used to configure the untrust zone subinterface of the virtual firewall vfwfmnat:

   interface GigabitEthernet 1/0/1.4018

   vlan-type dot1q 4018

   ip binding vpn-instance vfwfmnat

   ip address 192.168.185.2 255.255.255.252

   firewall zone vpn-instance vfwfmnat untrust

   add interface GigabitEthernet 1/0/1.4018

   #The following commands are used to configure the trust zone subinterface of the virtual firewall vfwfmnat:

   interface GigabitEthernet 1/0/0.4000

   vlan-type dot1q 4000

   ip binding vpn-instance vfwfmnat

   ip address 10.188.40.1 255.255.255.0

   firewall zone vpn-instance vfwfmnat trust

   add interface GigabitEthernet 1/0/0.4000

   #The following command is used to configure the default route on the virtual firewall vfwfmnat:

   ip route-static vpn-instance vfwfmnat 0.0.0.0 0.0.0.0 192.168.185.1

   #The following commands are used to configure the default packet filtering rule between security zones on the virtual firewall vfwfmnat:

   firewall packet-filter default permit interzone vpn-instance vfwfmnat local trust direction inbound

   firewall packet-filter default permit interzone vpn-instance vfwfmnat local trust direction outbound

   firewall packet-filter default permit interzone vpn-instance vfwfmnat local untrust direction inbound

firewall packet-filter default permit interzone vpn-instance vfwfmnat local untrust direction outbound

firewall packet-filter default permit interzone vpn-instance vfwfmnat trust untrust direction inbound

firewall packet-filter default permit interzone vpn-instance vfwfmnat trust untrust direction outbound

#The following commands are used to configure NAT translation rules allowing public IP address and port translation to private IP addresses and ports for Internet users to gain access to FusionManager:

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 21 inside 10.188.40.43 21 no-reverse

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 80 inside 10.188.40.43 80 no-reverse

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 443 inside 10.188.40.43 443 no-reverse

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 543 inside 10.188.40.43 543 no-reverse

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 6081 inside 10.188.40.43 6081 no-reverse

#Use commands similar to the following to configure NAT rules for ports 30000 to 30100:

nat server vpn-instance vfwfmnat zone untrust protocol tcp global 10.185.40.43 30000 inside 10.188.40.43 30000 no-reverse

4. Configure the trust domain settings.

   To ensure secure access to FusionManager from a public network, you need to add the public IP address to the trust domain of FusionManager.

   a. Use **PuTTY** to log in to the active FusionManager node.

      Ensure that the active FusionManager node management IP address and username **galaxmanager** are used to establish the connection. The default password of user **galaxmanager** is **Huawei@CLOUD8**.In FusionManager standalone deployment mode, use the FusionManager management IP address to log in to the system.

   b. Run the following command to disable logout on timeout:

      TMOUT=0

   c. Run the following command to add the public IP address or domain name to the trust domain:

      white-ls-cfg add *IP address or domain name*

      For example, run the following command to add the user's public IP address for connecting to FusionManager to the trust domain:

      white-ls-cfg add 10.185.40.43

   NOTE

   To view trust domain configuration information or delete a trust configuration item, use the **white-ls-cfg --help** command.

   d. If FusionManager is deployed in active/standby mode, log in to the standby FusionManager node and repeat the preceding steps to add the required IP address or domain name to the trust domain.

4 Product Functions

4.1 FusionCompute Product Functions

4.1.1 Virtual Computing

Server Virtualization

Server virtualization enables physical server resources to be converted to logical resources. With virtualization technology, a server can be divided into multiple virtual computing resources that are isolated with each other. CPU, memory, disks, and I/O resources become pooled resources that are dynamically managed. Server virtualization increases resource utilization rate, simplifies system management, and implements server integration. In addition, the hardware-assisted virtualization technology increases virtualization efficiency and enhances VM security.

Server virtualization involves the following features:

- Bare metal architecture

  FusionCompute hypervisor adopts the bare-metal architecture and can directly run on servers to virtualize hardware resources. The bare metal architecture ensures that VMs have near-native performance, high reliability, and scalability.

- CPU virtualization

  FusionCompute converts physical CPUs to virtual CPUs (vCPU) for VMs. When multiple vCPUs are running, FusionCompute dynamically allocates CPU capabilities among the vCPUs.

- Memory virtualization

  FusionCompute adopts the hardware-assisted virtualization technology to reduce memory virtualization overhead. It also adopts memory overcommitment technologies to maximize the memory utilization. FusionCompute supports the following memory overcommitment technologies:

  - Memory ballooning: The system dynamically reclaims the free memory from a VM and allocates it to other VMs. Applications on the VMs are not aware of memory reclamation and allocation. The total amount of memory used by all VMs on a physical server cannot exceed the amount of the physical memory.
  - Memory swapping: The system swaps out data on the reserved VM memory to an external storage file to free the reserved memory and get back the data when required.

  - Memory sharing: Multiple VMs share the memory page on which the data content is the same.
- GPU passthrough

  In FusionCompute, a Graphic Processing Unit (GPU) on a physical server can be directly attached to a specified VM to improve VM graphic processing capabilities.
- USB passthrough

  In FusionCompute, a USB device on a physical server can be directly attached to a specified VM. This feature allows users to use USB devices in virtualization scenarios.

VM Resource Management

VM Resource Management allows administrators to create VMs using a VM template or in a custom manner, and manage cluster resources. This feature provides the following functions: automatic resource scheduling (including load balancing mode and dynamic energy-saving mode), VM life cycle management (including creating, deleting, starting, restarting, hibernating, and waking up VMs), storage resource management (including managing common disks and shared disks), VM security management (including using custom VLANs), and online VM QoS adjustment (including setting CPU QoS and memory QoS).

- VM life cycle management

  The VM life cycle management function allows users to adjust the VM status based on service load. User can perform the following operations to a VM:

  - Create, delete, start, stop, restart, or query a VM.

    After receiving a VM creation request, FusionCompute selects proper physical resources to create the VM based on the specified requirements. The requirements include the VM specifications (such as the number of vCPUs, memory size, and system disk size), image specifications, and network specifications.

    After the VM is created, FusionCompute monitors the VM running status and its attributes.

    Users can stop, restart, and delete VMs as required.
  - Hibernate or wake up a VM.

    Users can hibernate idle VMs when the service load is light and wake up the hibernated VMs when the service load is heavy. This improves system resource utilization.
- VM template management

  VM templates can be customized for creating VMs.
- CPU Quality of service (QoS)

  The CPU QoS ensures optimal allocation of computing resources for VMs and prevents resource contention between VMs due to different service requirements. It effectively increases resource utilization and reduces costs.

  During creation of VMs, the CPU QoS is specified based on the service to be deployed. After the VMs are created, the system dynamically binds the vCPUs to physical CPUs based on the CPU QoS. In this way, a pool of physical CPUs that are bound to different vCPUs with the same CPU QoS is created on a server.

  The CPU QoS determines the VM computing power. The system ensures the VM CPU QoS by setting the minimum computing capability and the computing capability upper limit for VMs.

  CPU QoS contain the following parameters:

  - **CPU quota**

CPU quota defines the proportion based on which CPU resources to be allocated to each VM when multiple VMs compete for physical CPU resources.

This section uses a host (physical server) that uses a single-core, 2.8 GHz CPU as an example to describe how CPU quota works. Three VMs (A, B, and C) run on the host, and their quotas are set to 1000, 2000, and 4000, respectively. When the CPU workloads of the VMs are heavy, the system allocates CPU resources to the VMs based on the CPU quotas. VM A with 1000 CPU quota can obtain a computing capability of 400 MHz. VM B with 2000 CPU quota can obtain a computing capability of 800 MHz. VM C with 4000 CPU quota can obtain a computing capability of 1600 MHz. The computing capability calculation is more complex in actual use.

The CPU quota takes effect only when resource contention occurs among VMs. If the CPU resources are sufficient, a VM can exclusively use physical CPU resources on the host if required. For example, if VMs B and C are idle, VM A can obtain all of the 2.8 GHz computing capability.

- **CPU reservation**

  CPU reservation defines the minimum CPU resources to be allocated to each VM when multiple VMs compete for physical CPU resources.

  If the computing capability calculated based on the CPU quota of a VM is less than the CPU reservation value, the system allocates the computing capability to the VM according to the CPU reservation value. The offset between the computing capability calculated based on the CPU quota and the CPU reservation value is deducted from computing capability of other VMs based on their CPU quotas and is added to the VM.

  If the computing capability calculated based on the CPU quota of a VM is greater than the CPU reservation value, the system allocates the capability to the VM according to the CPU quota.

  For example, three VMs (A, B, and C) run on the host that uses a single-core physical CPU, their quotas are set to 1000, 2000, and 4000, respectively, and their CPU reservation values are set to 700 MHz, 0 MHz, and 0 MHz, respectively. When the CPU workloads of the three VMs are heavy:

  - According to the VM A CPU quota, VM A should have obtained a computing capability of 400 MHz. However, its CPU reservation value is greater than 400 MHz. Therefore, VM A obtains a computing capability of 700 MHz according to its CPU reservation value.
  - The system deducts the offset (700 MHz minus 400 MHz) from VMs B and C based on their CPU quota.
  - VM B obtains a computing capability of 700 (800 minus 100) MHz, and VM C obtains a computing capability of 1400 (1600 minus 200) MHz.

  The CPU reservation takes effect only when resource contention occurs among VMs. If the CPU resources are sufficient, a VM can exclusively use physical CPU resources on the host if required. For example, if VMs B and C are idle, VM A can obtain all of the 2.8 GHz computing capability.

- **CPU limit**

  CPU limit defines the upper limit of physical resources that can be used by a VM. For example, if a VM with two virtual CPUs has a CPU limit of 3 GHz, each virtual CPU of the VM can obtain a maximum of 1.5 GHz computing resources.

- Memory QoS

  Memory QoS allows VM memory to be intelligently allocated based on the preset percentage of reserved memory. Memory overcommitment technologies, such as

memory ballooning, are used to provide more virtual memory resources. Users are unaware of the memory overcommitment.

The user can set the reserved memory percentage based on service requirements. The main principle of memory overcommitment is to first use the physical memory.

Memory QoS contains the following parameters:

- **Memory quota**

  Memory quota defines the proportion based on which memory resources to be allocated to each VM when multiple VMs compete for physical memory resources.

  The system allocates memory resources to VMs based on the proportion when VMs apply for memory resources or hosts release free memory resources (such as when VMs are migrated or stopped).

  CPU resources can be scheduled in real time. Memory resources are scheduled subtly and continuously when VMs are running until the configured VM memory resources are allocated to VMs.

  For example, three 4 GB memory VMs run on a 6 GB memory host, and their memory quotas are set to 20480, 20480, and 40960, respectively. In this case, the memory allocation ratio is 1:1:2. When the memory workloads on the three VMs gradually increase, the system subtly adjusts memory resources on the VMs based on the memory quotas until the VMs obtain 1.5 GB, 1.5 GB, and 3 GB memory, respectively.

  The memory quota takes effect only when resource contention occurs among VMs. If the memory resources are sufficient, a VM can exclusively use physical memory resources on the host if required. For example, if the memory resources required by VMs B and C are less than the reserved memory values, and VM A has more memory workloads to handle, VM A can use memory resources from free memory resources, and memory resources on VMs B and C, until the memory resources obtained by VM A reach the upper limit, or the free memory resources are used up and memory resources on VMs B and C drop to the reserved values. For example, if VM C is not under memory pressure and has 1 GB memory reserved, VMs A and B theoretically can obtain a maximum of 2.5 GB memory resources each.

- **Memory reservation**

  Memory reservation defines the minimum memory resources to be allocated to each VM when multiple VMs compete for memory resources.

  A VM exclusively uses its reserved memory, that is, if certain memory resources are reserved for a VM, other VMs cannot use the memory resources even if the memory resources are idle.

- **Memory limit**

  Memory limit defines the upper limit of physical resources that can be used by a VM. For example, if a VM with two virtual CPUs has a CPU limit of 3 GHz, each virtual CPU of the VM can obtain a maximum of 1.5 GHz computing resources. When multiple VMs are started, they will complete for memory resources. To improve memory utilization and reduce idle memory, users can set the memory limit parameter in the configuration file when creating a VM so that the memory allocated to the VM does not exceed the upper limit.

- Cluster resource management, including dynamic resource scheduling for load balancing or energy saving.

  When detecting an idle VM, the system releases some memory and CPU resources of the VM to the virtual resource pool based on the preset conditions. Users can monitor the dynamic resources on FusionCompute.

NOTE

The lower CPU overcommitment rate for a host indicates better host and VM performance. You are advised to set the CPU overcommitment rate to at most 4:1 when provisioning services and performing maintenance operations.

- VM statistics

  The system collects information about the resource usage for user VMs and disks.

Dynamic VM Resource Adjustment

FusionCompute dynamic VM resource adjustment allows users to perform the following operations:

- Adjusting the number of vCPUs for VMs in the running or stopped state

  Users can add vCPUs for VMs based on the service load, irrespective of whether the VMs are running or stopped and delete vCPUs for VMs when VMs are stopped. This allows computing resources to be adjusted in a timely manner.

- Adjusting the memory size for VMs in the running or stopped state

  Users can increase the memory size for VMs based on the service load, irrespective of whether the VMs are running or stopped and decrease the memory size for VMs when are stopped. This allows memory resources to be adjusted in a timely manner.

- Adding or deleting NICs for VMs in the running or stopped state

  Users can add or delete virtual NICs for running or stopped VMs to meet service requirement for NICs.

- Attaching virtual disks for VMs in the running or stopped state

  Users can expand the storage capacity for VMs, irrespective of whether the VMs are running or stopped.

NOTE

When a VM uses virtualized storage and is in the running or stopped state, users can expand the VM storage capacity by enlarging the capacity of existing disks on the VM.

Distributed Resource Scheduling and Power Management

FusionCompute provides various pooled virtual resources, such as computing resources, storage resources, and virtual network resources. The FusionCompute intelligently schedules the virtual resources based on the system load to ensure optimal resource allocation as well as high reliability and availability.

FusionCompute supports the following type of resource scheduling:

- Resource scheduling for load balancing

  When detecting that the load on servers in a cluster varies significantly and exceeds the preset thresholds, the system automatically migrates VMs based on the predefined load balancing policy to achieve load balancing.

- Resource scheduling for energy saving

  Resource scheduling for energy saving can be enabled only after resource schedule for load balancing is enabled. When detecting that the load on a cluster is light, the system automatically migrates VMs to some servers based on the predefined energy saving policy and powers off the idle servers.

VM Live Migration

FusionCompute supports VM migration between the hosts that share the same data stores without interrupting services. This reduces the service interruption time caused by server maintenance and saves energy for data centers.

4.1.2 Virtual Network

Virtual NIC

Each virtual network interface card (NIC) has an IP address and a MAC address. It has the same functions as a physical NIC on a network. FusionCompute supports intelligent NICs (iNICs), which use multiple queues, virtual swapping, quality of service (QoS), and uplink aggregation to improve NIC I/O performance.

Network I/O Control

The network QoS policy enables bandwidth configuration control.

- Bandwidth control based on the send direction and receiving direction of a port group member port

  Traffic shaping and bandwidth priority are configured for each port in a port group to ensure network QoS.

DVS

Each host connects to a distributed virtual switch (DVS), which functions as a physical switch. The DVS connects to VMs through a virtual port in the downstream direction and to the host NIC in the upstream direction. The DVS implements network communication between hosts and VMs.

In addition, the DVS ensures unchanged network configuration for VMs when the VMs are migrated across hosts.

4.1.3 Virtual Storage

Virtual Storage Management

Storage virtualization abstracts resources on storage devices into data stores. Data stores are logical containers that hide the peculiarities of physical storage devices from VMs and provide a unified model for storing VM files. Storage virtualization helps the system better manage virtual infrastructure storage resources with improved resource utilization and flexibility and increased application uptime.

The following storage units can be encapsulated as data stores:

- Logical unit numbers (LUNs) on storage area network (SAN) storage, including Internet Small Computer Systems Interface (iSCSI) and fibre channel (FC) SAN storage
- File systems on network attached storage (NAS) devices
- FusionStorage Block storage resource pools
- Local hard disks

Data stores support the following file system formats:

- Virtual Image Management System (VIMS)

  The VIMS is a high-performance file system that is designed for storing VM files. The VIMS data can be stored on any Small Computer System Interface (SCSI)-based local or shared storage device, such as FC, Fibre Channel over Ethernet (FCoE), and iSCSI SAN devices.

- Network File System (NFS)

  The NFS runs on NAS devices. FusionSphere supports the NFS V3 protocol. It can connect to the custom NFS disks on the NFS server and have these disks attached to meet storage requirements.

- EXT4

  FusionSphere supports virtualization of server local disks.

Thin-Provisioning Virtual Storage

Thin-provisioning virtual storage enables flexible, on-demand allocation of storage space, which improves storage utilization. This function allows more virtual memory space to be allocated than the physical memory available. The physical memory space is allocated only for the virtual memory space where data is recorded. The virtual memory space where no data is recorded does not occupy the physical memory space.

This function is configured for virtual disks. Administrators can set a disk to **Common** or **Thin-provisioning**.

- Storage device independent

  This function is not dependent on the operating system (OS) type or hardware of the storage device. It can be configured for any storage device that runs a virtual image management system.

- Capacity monitoring

  This function enables alarms over data store usage. If the data usage exceeds the preset threshold, an alarm will be displayed.

VM Snapshot

Users can save the static data of a VM at a specific moment as a snapshot. The data includes information about all disks attached to the VM at the snapshot-taking moment. The snapshot can be used to restore the VM to the state when the snapshot was taken. This function applies to data backup and disaster recovery systems (for example, eBackup) to improve system security and availability.

Storage Live Migration

This function enables disks on VMs to be migrated to other storage units when the VMs are running. The disks can be migrated between different storage devices or storage units on one storage device under virtual storage management. With this function enabled, storage resources of VMs can be dynamically migrated, thereby facilitating device maintenance.

4.1.4 Availability

VM Live Migration

In FusionCompute, this feature enables VMs to be migrated from one host to any host across computing clusters. During the migration, services are not interrupted. If the migration fails,

the VM on the destination server will be destroyed. The user can still use the VM on the source server. This reduces the service interruption time by migrating VMs from the physical server to be maintained to another physical server and saving energies for the data center.

VM Fault-based Migration

If a VM becomes faulty, FusionCompute automatically restarts the VM. In the process of configuring clusters, the user can enable or disable the high availability (HA) function. The system periodically checks the VM status. When the system detects that the physical server on which a VM runs is faulty, the system will restart the VM on the original physical server or another physical server based on the host fault processing policies so that the VM can be restored in a timely manner. Because the restarted VM will be recreated and loaded with the OS like a physical server, the unsaved data is lost when the VM encountered the error.

The system can detect errors on the hardware and system software that cause VM failures.

4.1.5 Security

Virtual Network Access Control

The network range of the VMs can be divided by configured virtual local area network (VLAN) IDs of network interface cards (NICs) on VMs.

- The port group to which a VM NIC belongs to can be dynamically modified, and thereby the NIC VLAN ID can also be dynamically changed.
- When the NIC VLAN ID is dynamically changed, the NIC VLAN can also be changed by binding a new VLAN to the NIC without adding a NIC.

4.2 FusionManager Product Functions

4.2.1 Role-based Access Control

Role-based Access Control (RBAC) assigns an administrator the rights to perform specified operations to specific roles in different domains and virtual data centers (VDCs). The administrator can manage users, roles, domains, and VDCs to separate operation rights and service data of different users.

User Management

- Users are classified into system administrators and tenants.
  - A system administrator has permission to manage all resources in the system.
  - A tenant can manage only the resources in the VDC to which the tenant belongs and perform the allowed operations, for example, creating, starting, or stopping a VM or creating an application.
- Password Management

  Administrators can configure a password policy to ensure that all the passwords entered in the system meet data security requirements. The policy can be specified as required, for example, specifying the password length, password reuse frequency, and restrictions on the use of a username and special characters in a password.

Role Management

FusionManager provides the following default roles: system super administrator, system viewer, VDC administrator, VDC user, and system operator.

Domain Management

- Domain-based resource management

  Administrators can create different domains and associate resource clusters with the domains to implement domain-based resource management.

- Domain-based user management

  Administrators can grant users different rights based on the domains to which the users belong. Domains are defined for the system administrator. A service administrator can manage multiple domains, and a domain can be managed by multiple administrators.

- Rights- and domain-based management

  Users belonging to different domains or resource clusters are assigned different permissions by the implementation of rights- and domain-based management.

VDC Management

A VDC is the basic unit of virtual resources in FusionManager. A VDC is managed by its administrator. The system administrator can configure the range and quotas for resources in a VDC.

The administrator who creates a VDC can create an administrator or set an existing administrator as the VDC administrator.

Single Sign-On

The single sign-on function allows an administrator who has logged in to FusionManager to directly log in to FusionCompute without entering the FusionCompute login username and password.

4.2.2 Computing Virtualization

Computing virtualization enables physical server resources to be converted to logical resources so that one server can function as multiple or hundreds of standalone virtual servers. Hardware resources, such as the CPUs, memory, disks, and I/O resources become dynamically managed resource pools. This enhances the resource utilization rate and simplifies system management. In addition, hardware-assisted virtualization technology increases virtualization efficiency and enhances VM security.

VM Management

You can create VMs by importing templates or creating application instances. VM management includes managing the lifecycle of VMs and scheduling VM resources. For details, see 4.2.7 VM Management.

Resource Cluster Management

A resource cluster consists of virtual computing, storage, and network resources. Resource clusters are isolated from one another. For details about the resource cluster, see 4.2.5 Resource Cluster Management.

Large VM Memory

The Large VM Memory feature allows VMs that have up to 1 TB virtual memory to run properly in the system and keeps up with the requirements of memory-consuming services.

Memory Overcommitment

After memory overcommitment is enabled for a host, VMs running on the host share the memory resources, increasing VM density on the host.

Cross-CPU VM Live Migration

Hardware devices that compose a virtualization cluster may be purchased at different times and use CPUs of different generations. With the Cross-CPU VM Live Migration feature, VMs are allowed to be live migrated between hosts that use CPUs of different generations. This feature masks the unique advanced features of the host CPUs and provides the same CPU features to all hosts in a cluster. Therefore, VMs can be migrated to any host in the cluster, but they cannot use the masked advanced CPU features of the host.

Guest NUMA

Guest NUMA presents a topology view of memory and CPU resources on each host to VMs, the VM user can configure VM CPUs and memory using the third-party software (such as Eclipse) based on this topology, so that VMs can obtain the most easy-to-access memory based on the topology, thereby reducing access latency and improving VM performance.

4.2.3 Network Virtualization

FusionManager can allocate network resources, such as subnets and VLANs to VMs using distributed virtual switches (DVSs) to implement network virtualization.

DVS

A DVS works similarly to a physical switch. The DVS connects to network interface cards (NICs) on hosts that provide the VM resources in the uplink and connects to VMs through port groups in the downlink. VMs communicate with external networks over the uplink of the DVS.

Subnet Management

Subnet Management isolates VMs in each subnet at layer 2 in the FusionManager system.

VLAN Pool

VLAN Pool provides VLANs for port groups on DVSs.

FusionManager also supports the Virtual eXtensible Local Area Network (VXLAN) function, which breaks through VLAN restrictions on the number of networks and applies to the large-scale network deployment in multi-tenant environments.

IPv6

FusionManager supports IPv6 for service VMs. Both VLAN and VXLAN VMs can use IPv6 addresses to communicate. Each NIC of the VM can have an independent IPv4 or IPv6 address or have both addresses.

NOTE
Virtualization suite 6.3.0 and 6.3.1 do not support IPv6.

VPC Management

A virtual private cloud (VPC) provides an exclusive and isolated network container for application instance provisioning. Administrators can add virtual firewalls and various networks to the VPC.

FusionManager supports three types of networks in a virtual private cloud (VPC): direct network, internal network, and routed network.

- Direct network: A direct network itself contains no network resources, but it has direct access to external networks. Therefore, VMs created on a direct network actually uses the IP addresses provided by the connected external network, which can be the existing network of the customer or the Internet.

NOTE
When creating VMs in a direct network, you are advised to configure security groups to isolate the VMs and implement VM access security control.

- Internal network: An internal network is isolated from other networks and uses exclusive network resources. Therefore, you can deploy services that have high security requirements, for example, database server services, on an internal network to ensure high data security.
- Routed network: A routed network provides flexible communication functions and supports various services. Routed networks in a VPC can communicate with each other or the Internet. A routed network also provides elastic IP address, access control list (ACL), destination network address translation (DNAT), and virtual private network (VPN) services. Before creating a routed network in a VPC, ensure that you have applied for a virtual router for the VPC.

Figure 4-1 Networks

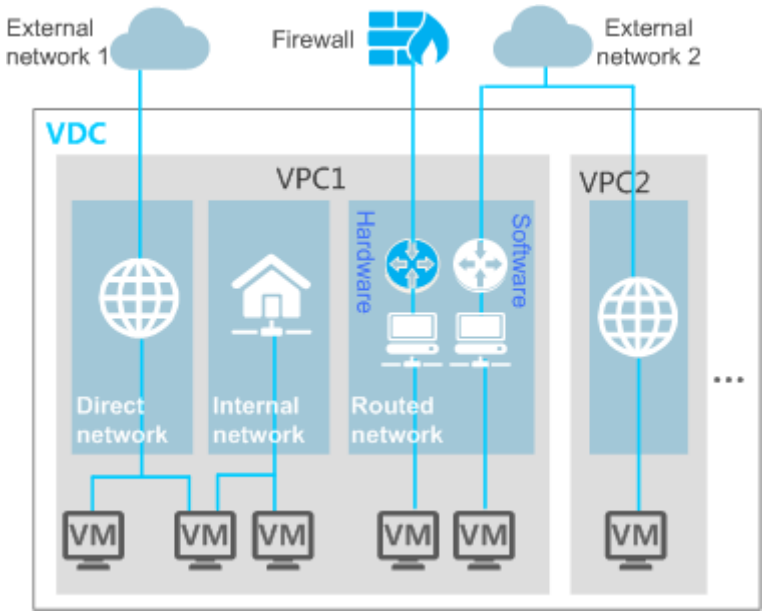

Advanced VPC Network Attributes

- Router

  You can apply for a virtual router in a virtual private cloud (VPC) to create routed networks. Routed networks in a VPC can communicate with each other. VMs in a routed networks can use advanced network functions including elastic IP address, destination network address translation (DNAT), source network address translation (SNAT), access control list (ACL), and virtual private network (VPN). FusionManager supports the following types of routers:

  - Hardware router: a firewall instance created on a physical firewall to provide router functions
  - Software router: a system service VM, which is also referred to as a router VM, created on FusionManager to provide router functions

  Table 4-1 lists the advanced network functions supported by hardware routers and software routers.

Table 4-1 Advanced network functions supported by hardware routers and software routers

Router Type	Elastic IP Address	NAT	ACL	VPN	VXLAN
Software router	Supported	Supported	Supported	Supported (by IPSec VPN)	Supported
Hardware router	Supported	Supported	Supported	Supported	Not supported

- Elastic IP address

By binding a public IP address to the VMs or private IP addresses associated with a routed network, the elastic IP address service allows services in a virtual private cloud (VPC) accessible to external services using a fixed public IP address.

- NAT

  Network address translation (NAT) is used to translate between service private IP addresses and public IP addresses. Compared with the elastic IP address service, the NAT service uses less public IP addresses when interconnecting private networks with public networks. Therefore, the NAT service can be used in scenarios in which public IP address resources are limited.

  Network Address Translation (NAT) can be categorized as Destination Network Address Translation (DNAT) and Source Network Address Translation (SNAT). DNAT is used to establish a mapping between private IP addresses and public IP addresses, allowing public network users to access internal services in a private network. SNAT is used to translate the source IP address contained in an IP packet into another public IP address, allowing internal users to access external networks using a shared public IP address.

- ACLs

  Access Control List (ACL) is an access security control method that allows permitted users to access only the specified network resources by controlling the data packets sent to or from a port.

- VPN

  Virtual private network (VPN) connections are used to connect routed networks in a virtual private cloud (VPC) to user networks using VPN tunnels, which allow users to access service resources in the VPC.

  FusionManager supports two types of VPNs: IP Security Protocol (IPSec) and Layer 2 Tunneling Protocol (L2TP).

  - IPSec VPN: commonly used for communication between enterprise headquarters and branches by setting up an encrypted VPN tunnel between the VPN gateway and the remote user network gateway.
  - L2TP VPN: commonly used for enterprise employees on business trips to connect to the enterprise network over the Internet. Such connections are set up between the VPN clients and the VPN gateway.

  Hardware routers support both IPSec and L2TP VPNs. Software routers only supports IPSec VPNs.

NOTE

IPSec VPN connections in different VPCs cannot be configured with the same peer IP address. Therefore, you need to configure multiple IP addresses for the peer device to communicate with IPSec VPN connections in different VPCs.

VPC Security Management

FusionManager provides ACLs, security groups to ensure the security of application systems and VMs in VPCs.

- ACLs

  Access Control List (ACL) is an access security control method that allows permitted users to access only the specified network resources by controlling the data packets sent to or from a port.

- Security group

  Security group is a VM access security control method that specifies the communication scope of VMs in the same or different security groups. You can define different access

rules for a security group, and these rules take effect for all VMs added to this security group.

4.2.4 Storage Virtualization

Storage virtualization technology abstracts resources on storage devices into data stores. Data stores, working similarly to file systems, are logical containers that hide the specifics of physical storage devices from VMs and provide a unified model for storing VM files. Storage virtualization helps the system better manage virtual infrastructure storage resources with improved resource utilization and flexibility and increased application uptime.

The FusionManager system can manage storage resources on Internet Protocol storage area network (IP SAN), Fibre Channel (FC) SAN, network attached storage (NAS), and FusionStorage devices. It can also allocate these resources to resource clusters as data stores, providing storage resources for VMs in the resource clusters.

Thin Provisioning

Thin provisioning of virtual storage resources enables flexible, on-demand allocation of storage space, which improves storage utilization. Different from traditional thick provisioning, thin provisioning provides more storage space than the physical host has available. The system allocates physical storage space only when data is written into virtual storage, improving storage utilization.

This feature is configured for virtual disks of specified levels. Administrators can set a disk to common or thin-provisioning mode.

VM Snapshot

A VM snapshot records the VM status at a specific time point and can be used to restore VM settings to the point at which the snapshot was taken. A VM snapshot captures the entire status of the VM, including information about all VM disks. VM snapshots can be used in data backup and disaster recovery (DR) scenarios to improve system security and reliability.

Virtual Storage QoS

Virtual Storage QoS allows users to set an input/output (I/O) upper limit for each VM disk. The I/O upper limit controls the ability of a VM to obtain storage resources. Therefore, among the VMs attached to the same storage device, I/O-extensive VMs or VMs with abnormal storage I/O will not affect other VMs' access to the storage device.

Raw Device Mapping

Raw device mapping (RDM) allows VMs to directly access logical unit numbers (LUNs) on physical storage devices over FC or iSCSI.

Virtual Storage Live Expansion

Virtual Storage Live Expansion allows the VM disk capacity to be expanded when the VM is running. Whether the live expansion takes effect in service depends on the virtualization software.

Virtual Storage Migration

Virtual Storage Migration allows administrators to migrate VM disks to other data stores on the same storage device or to different storage devices. This feature implements dynamic storage resource scheduling and helps ensure service continuity during device maintenance or resource scheduling.

If VMs are stopped, VM disks can also be migrated between LUNs and virtualized storage.

Storage Life Cycle Management

Storage Life Cycle Management allows administrators to create or delete virtual disks on any data stores and also attach a disk to a VM or detach the disk from a VM.

4.2.5 Resource Cluster Management

A resource cluster is a group of computing, storage, and network resources in the system. Different resource clusters are isolated from one another. A hypervisor may contain multiple resource clusters.

In the FusionSphere solution, administrators operate resource clusters, for example, create or delete a resource cluster, on the in-use hypervisor, such as FusionCompute. Administrators query resource cluster information, monitor resource cluster performance, and schedule resource clusters on FusionManager.

Querying a Resource Cluster

On FusionManager, administrators can query basic information about a resource cluster, such as the cluster name, domain, hypervisor, VMs, host resources, and storage resources of the resource cluster.

Monitoring Resource Cluster Performance

On FusionManager, administrators can query the performance monitoring data of a resource cluster by week, month, year, or a custom period. Performance monitoring indicators include the average CPU usage, average memory usage, average network rate, and hosts with the highest CPU usages.

Configuring Resource Scheduling Policies

On FusionManager, administrators can configure and view resource scheduling policies. With resource scheduling policies configured, the system keeps monitoring the service loads of hosts in a resource cluster. To ensure load balance in the resource cluster, when detecting that the service loads of some hosts are heavy, the system automatically migrates the VMs across hosts or provides manual VM migration suggestions for users to perform.

4.2.6 Resource SLA Management and Scheduling

The requirements for resources may vary depending on different applications or tenants, for example, some important applications demand high-performance and high-security resources. To satisfy diverse resource requirements, FusionManager provides the resource service level agreement (SLA) function to implement on-demand resource allocation.

FusionManager allows administrators to define different computing SLA levels for availability zones (AZs) and resource clusters by their computing performance, reliability,

hardware models, and other factors, and for data stores by their storage media, reliability, performance, RAID levels, and other factors.

When a user applies for resources, FusionManager allocates matched cluster and data store resources based on the requested computing and storage SLA levels.

4.2.7 VM Management

VM management mainly includes the VM life cycle management and VM resource management functions of FusionManager.

VM Life Cycle Management

- Creating a VM

  On FusionManager, administrator and tenants can create a VM by:

  - Creating an application instance
  - Using a VM template
  - Customizing VM specifications
  - Cloning a VM
- Destroying a VM

  Administrators can destroy an unnecessary VM to free system resources by deleting the application instance to which the VM belongs.
- Operating a VM

  Administrators or tenants can perform the following operations for one or more VMs: starting, safely restarting, forcibly restarting, hibernating, safely stopping, and forcible stopping the VMs.
- Migrating a VM

  Administrators can migrate a VM from one host to another.
- Restoring a VM

  If the operating system of a VM is not working properly, the administrator can restore the VM without user data on the VM being affected.
- Creating a VM snapshot

  A VM snapshot is a reproduction of VM data and status created at a certain point of time for a VM. If the VM becomes faulty, the tenant can revert the VM snapshot to restore the VM to its status that remained when the snapshot was taken.
- Monitoring VM performance

  On FusionManager, administrators can query the performance monitoring results of a VM, such as its CPU usage, memory usage, network flow rates, and disk input/output (I/O) rates in a real-time manner or by a specified period (week, month, year, or a custom time period).

VM Resource Adjustment

Based on service load changes in the system, administrators can adjust the resources of a VM by performing the following operations:

- Adjusting the VM quality of service (QoS) settings

  VM QoS settings include the VM CPU and memory QoS settings. If a VM has CPU and memory QoS attributes configured, the system allocates virtual CPU and memory

resources to the VM based on the QoS requirements specified for the VM to ensure VM performance.

- Adjusting the number of VM CPUs

  Administrators can add or delete virtual CPUs for a VM to meet its computing requirements that vary based on service load changes.

- Adjusting the memory size

  Administrators can expand or reduce the memory size of a VM based on the changes to the VM service load.

- Adding or modifying a VM disk

  Administrators can add a disk or modify the disk capacity for a VM based on the changes to the VM service load to flexibly schedule system storage resources.

- Deleting a virtual disk

  Administrators can delete a virtual disk of a VM based on the changes to the VM service load to release storage resources.

- Adding a virtual network interface card (NIC) or modifying NIC attributes

  Administrators can add a virtual NIC or modify NIC attributes for a VM to implement flexible use of network resources.

- Deleting a NIC

  Administrators can delete a virtual NIC of a VM to release network resources and improve resource utilization efficiency.

- Adding a universal serial bus (USB) controller

  Administrators can add a USB controller for VMs created in the VMware hypervisor and then bind USB devices to the VMs.

- Deleting a USB controller

  Administrators can delete a USB controller from the VMware hypervisor to release USB controller resources used by VMs in the hypervisor.

- Binding or unbinding a USB device to or from a VM

  Administrators can bind a host USB device to a VM so that the VM is allowed to access the USB device. Administrators can also unbind the USB device from a VM to release USB device resources in the system.

4.2.8 Third-Party Resource Integration

FusionManager supports management of physical devices and virtualization software from Huawei and third-party vendors in a centralized manner.

Physical Device Adding and Management

- Adding physical devices to FusionManager

  FusionManager can monitor and manage servers, storage devices, physical firewall devices, load balancers, and switches after these devices have been added to it. Administrators can add physical devices to the FusionManager system individually or in batches.

- Operating physical devices

  On FusionManager, administrators can power on, power off, or restart the added physical devices of various models or from different vendors.

- Monitoring physical devices

On FusionManager, administrators can view the real-time or historical performance monitoring statistics for the added physical devices.

- Viewing and handling device alarms

  On FusionManager, administrators can view and handle alarms generated for the added physical devices.

Virtualization Software Adding and Management

- Adding virtualization software to FusionManager

  A hypervisor is a software system that virtualizes computing, network, and storage resources in a cloud environment. For example, FusionCompute is a Huawei hypervisor. FusionManager supports FusionCompute and VMware vCenter hypervisors.

- Managing hypervisors

  Administrators can view attribute information about a hypervisor, such as its name and domain.

  After adding a hypervisor to the FusionManager system, administrators can associate the resource clusters in the hypervisor with FusionManager so that FusionManager can manage the virtual resources in the hypervisor resource clusters.

4.2.9 Physical Resource Management

FusionManager supports management of physical devices, including subracks, servers, storage devices, and switches. In the FusionSphere solution, physical devices must be manually added to FusionManager.

Physical Device Connection Protocols

Physical servers, switches, and storage devices can be connected to the FusionManager system using protocols described in Table 4-2.

Table 4-2 Physical device connection protocols

Category	Device Type	Connection Protocol	Connection Protocol Description	Alarm Reporting Protocol
Server	Blade server	Simple Network Management Protocol (SNMP) and Intelligent Platform Management Interface (IPMI)	• SNMP: used for connecting devices to the system, setting the alarm trap IP address, obtaining blade monitoring information, such as the fan or power source status, and adding alarms generated for servers to FusionManager	Simple Network Management Protocol (SNMP) **NOTE** If a redundant array of independent disks (RAID) on a server is degraded, the server reports an alarm to the system. FusionManager reports alarms for the degrade of only RAID models 1064, 1068, 1078, 2308, and 2208.

Category	Device Type	Connection Protocol	Connection Protocol Description	Alarm Reporting Protocol
			. • IPMI: used for monitoring blade information, such as the CPU and memory usage, and powering on or off a blade.	
	Rack server	IPMI	Used for monitoring server information, such as the CPU and memory usage, and powering on or off a server.	
Storage devices	SAN device and FusionStorage	Type-Length-Value (TLV), Storage Management Initiative—Specification (SMI-S), Hypertext Transfer Protocol (HTTP), and Hypertext Transfer Protocol Secure (HTTPS)	TLV, SMI-S, HTTP, and HTTPS are used for connecting different types of storage devices to FusionManager.	
Switches, firewalls, and load balancers	N/A	SNMP and Secure Shell (SSH)	• SSH: used for connecting switch devices to the system, obtaining the versions of the connected devices, and setting the alarm trap IP address. • SNMP: used for obtaining switch device monitoring information, such as the port status and rate. • Firewalls are connected to	

Category	Device Type	Connection Protocol	Connection Protocol Description	Alarm Reporting Protocol
			FusionManager using SSH and SNMP. • HTTPS: used for connecting load balancers to FusionManager. **NOTE** FusionManager can use the SNMP protocol to obtain detailed switch monitoring information only when the switches are connected to the system using SNMP. If the switches are connected to the system using the SSH protocol, FusionManager can only view the switch online information (regardless of the switch's status of connection to the FusionManager system).	

NOTE

As a client to connect to third-party devices, FusionManager supports a wide range of communications protocols, including SNMP, IPMI, TLV, SMI-S, SSH, HTTP, and HTTPS. In practice, use the protocol that the connected third-party device supports. The guidelines for selecting a recommended protocol are as follows:

- SNMPv1, SNMPv2c, and HTTP are known insecure protocols. You are advised to use the more secure SNMPv3 or HTTPS protocol.
- SNMPv3 supports HMACMD5 and HMACSHA algorithms, which, however, are insecure algorithms. You are advised to use the more secure HMACSHA algorithm.
- FusionManager supports AES (AES128) and DES data encryption algorithms. You are advised to use the more secure AES (AES128) algorithm.

If the connected third-party device only supports an insecure protocol or algorithm, the configuration guideline is as follows:

- If the device does not have high requirements for security, add the device as normal.
- If the device has high requirements for security, for example, it requires to use SNMPv3, contact technical support.

Server Management

- Querying server information

  On FusionManager, you can query the attribute information about a server, such as the server name, management IP address, basic input/output system (BIOS), number of CPUs, memory size, logical hard disk capacity, number of network interface cards (NICs), and number of network ports.

- Operating a server

  On FusionManager, you can power on, forcibly power off, and safely restart a server.

- Monitoring a server

  On FusionManager, you can view the following performance indicators of a server: CPU usage, memory usage, network outbound rate, network inbound rate, disk write rate, and disk read rate. In addition, you can also query the server performance by week, by month, by year, by a custom interval, or export the query result.

Network Device Management

- Querying switch information

  On FusionManager, you can query the attribute information about a switch, such as the switch name, management IP address, model, type, and status.

- Querying the switch port connection status

  On FusionManager, you can query the information about each switch port, such as the status, transmit rate, receive rate, transmitted packet loss rate, received packet loss rate, packet transmit error rate, and packet receive error rate.

Storage Device Management

Querying storage device configuration information: On FusionManager, you can query the attribute information about a storage device, such as the device name, management IP address, model, and status.

4.2.10 Load Balancing Management

FusionManager supports hardware load balancers (LBs). Based on customized load balancing policies, FusionManager allows load balancers to evenly send service requests to associated hosts to balance workloads on these hosts, improving service stability and reliability.

LB devices connected to FusionManager provide hardware LB functions.

Hardware LB

Huawei FusionManager can manage F5 LBs. Administrators can create, query, delete, and modify F5 LB services on FusionManager.

FusionManager supports **F5 BIG-IP LTM 10.2** and **F5 BIG-IP LTM 11.6** load balancers.

4.2.11 Template Management

FusionManager provides the functions for administrators to manage VM templates, VM logical templates, software packages, script files, VM specifications, and application templates.

VM template management

A VM template is a duplicate of a VM, which includes an OS, software packages, and a set of VM specifications.

Administrators can create a VM template to:

- Provision VMs. Creating VMs using a VM template can greatly reduce the time for creating the VMs and configuring operating systems (OSs) for them.
- Provision application instances: A VM template can be used to create multiple VMs for an application instance. An application instance can be used to provision application instances.

VM Logical Template Management

A VM logical template is a set of VM templates with the same specifications and OS settings. When FusionManager functions as a service provider, it provides VM templates to other management platforms using VM logical templates. When the VM templates with the same specifications in different hypervisors are provided to upper-layer management platforms as one VM logical template, the differences of the VM templates in the hypervisors are hidden from the management platforms.

Software Package Management

Software packages are installed on application VMs in an application template to deploy application instances. A software package contains software installation files, commands, and scripts.

Software packages stored on the local PC can be uploaded to FusionManager to:

- Automatic software installation: When you provision application instances using an application template in which the VM template has specific software configured, the system automatically installs the software on the VM that is to run the application instance.
- Software distribution: After you select a software package and specify the target VMs to install the package, the system automatically installs the software on the specified VMs.

VM Specification Management

VM Specifications define the hardware settings, including the number of CPUs, memory space, and system disk space.

You can define specifications for provisioning VMs with the defined specifications. VMs can have different specifications specified to meet different performance requirements.

Application Template Management

An application template is a service model that combines logical components, including VMs, software packages, scripts, and network resources, based on the configured relationship.

You can use an application template to create VMs in batches and rapidly deploy application instances.

4.2.12 Automatic Operation and Maintenance

FusionManager provides a tool that enables visibility for designing application templates. Administrators can quickly deploy applications using the created application templates. Administrators can also configure different resource scheduling policies for applications, so that the system can automatically schedule resources based on the configured scheduling policies and the application service loads.

Application Template Management

FusionManager provides an application template design tool, with which administrators can select the required application elements and drop them onto a visible central panel to rapidly create an application template and define resource relationships in the application.

Automatic application instance deployment

After creating an application template, administrators can rapidly publish application instances using the template.

Automatic Application Scaling

Administrators can configure application resource scaling policies for the system to automatically scale resources of the application instance, for example, add or reduce VMs in the application instance. These operations are triggered based on the scaling policies and the service load of the application instance, thereby achieving optimal resource utilization in the system. Automatic scaling policies include intra-group scaling policies and scheduled task policies.

- Intra-group scaling policies

  Intra-group scaling policies apply to only one application. When detecting that an application instance has a high CPU usage, memory usage, disk input/output (I/O) rate, or network inbound or outbound flow rate, the system automatically adds VMs for this application instance and installs the required software for the VMs. This reduces the overall resource load of the application instance and ensures that the application instance operates properly. Similarly, when detecting that the resources of an application instance is underused, the system automatically removes VMs from the application instance to release resources.

- Scheduled task policies

  Applying to a scaling group, time-based scheduling task policies allow resources to be assigned to different application instances at different time periods. For example, the system assigns resources to the VMs of office users in the daytime and to some public VMs in the evening.

4.2.13 Self-Service Management

FusionManager provides service catalogs for self-service management. Tenants can create application instances using an application template and manage the created application instances.

Service Request

FusionManager provides a self-service portal for administrators to quickly create applications and view the application deployment progress and reports.

Application Management

- Managing the application instance life cycle

  On FusionManager, administrators can start, stop, modify, and delete an application instance as needed.

- Monitoring application instance logs

  On FusionManager, administrators can monitor the running status and changes of application instances to identify and locate exceptions in a timely manner.

4.2.14 Monitoring Management

FusionManager provides the following monitoring management functions: performance monitoring, alarm management, task center management, and operation log management. Administrators can learn the system hardware and software running conditions by viewing the monitoring information, and therefore can handle the exceptions in a timely manner.

Monitoring Management

- Status monitoring

  Administrators can view the summary information about system hardware and software resources from dashboards, view the current status of physical and virtual computing, storage, and network resources, and export the historical monitoring data.

- Capacity monitoring

  Administrators can view the total capacity and usage of physical and virtual computing, storage, and network resources, and find out capacity expansion or reduction requirements in advance, preventing system overload or waste of resources.

- Application monitoring

  After an application is created, tenants need to perform routine maintenance for the application to ensure that the application provides services stably. FusionManager supports application monitoring, which enables tenants to learn the application service running information in real time, so they can handle the exceptions in a timely manner to ensure that the services run stably.

Alarm management

FusionManager monitors hardware and software alarms in the system in a centralized manner. All alarms, including those generated for physical and virtual computing and storage resources, are displayed in a list. FusionManager supports the following alarm severities:

- Critical: indicates the occurrence of a fault that will severely affect system services. Alarms of this severity must be handled immediately even if the fault occurs in non-working time.
- Major: indicates the occurrence of a fault that may affect system service quality. Alarms of this severity must also be handled immediately.
- Minor: indicates the occurrence of a fault does not affect service quality but may cause further exceptions. To prevent more serious faults, alarms of this severity need to be observed or handled if necessary.
- Warning: indicates the occurrence of a fault that may affect service quality. Alarms of this severity need to be rectified based on the error type.

To improve alarm handling efficiency, you can set the following alarm functions:

- Alarm notification by email: allows the system to send the alarms of the specified severity to the configured administrator's email box, so the administrator can obtain alarm information and handle the alarms in a timely manner.
- Alarm threshold setting: allows the system to generate alarms when the resource usage, such as the CPU usage, memory usage, disk input/output (I/O), and network bandwidth, reaches the specified alarm thresholds.

Operation Log Management

Administrators can query and export user operation logs from FusionManager.

4.2.15 Open APIs and SDKs

FusionManager provides software development kits (SDKs) and flexible application programming interfaces (APIs) for users to conduct secondary development.

Open SDKs

- FusionManager provides SDKs that support multiple programming languages, including JAVA and Python, facilitating secondary development.
- FusionManager also provides secondary development guidance, including SDK interface documentation, secondary development guides, and development samples.

Open APIs

FusionManager provides open APIs for the network management system to obtain FusionManager information for management and maintenance.

API Security

The network management system communicates with the open APIs over HTTP or HTTPS. The username and password are used to authenticate the interconnection between open APIs and the network management system. To ensure password security, the FusionManager password is encrypted using the RSA2048 algorithm.

O&M Data Obtaining

- Administrators of the network management system can easily obtain FusionManager operation and maintenance (O&M) data, such as real-time and historical alarm data and resource monitoring data, using open APIs.
- Administrators of the network management system can easily obtain information about FusionManager resources, such as clusters, servers, switches, and VMs, using open APIs.

Resources O&M

Administrators of the network management system can perform operations for FusionManager resources using open APIs, for example, starting, stopping, or restarting VMs, and powering on, powering off, or restarting hosts.

4.2.16 Security Management

FusionManager has a robust security system that ensures security of the application layer, system layer, network layer, and management layer.

- Application-layer security:

Ensures that the FusionManager service system can properly run and provide services for end users.

Application-layer security consists of account security, log security, web security, and rights- and domain-based management.

- System-layer security:

  Ensures that operating systems (DBs) and databases can properly run to support the running of various application software at the application layer.

  System-layer security consists of OS security hardening and GaussDB data security.

- Network-layer security:

  Ensures the effective virtual network isolation and proper running of network devices, such as switches, routers, and firewalls, thereby effectively implementing security policies at the network layer.

  Network-layer security consists of the defense against malformed packet attacks, isolation between administrator and tenant networks, VLAN and subnet, VPC, VPN, ACL, and security group.

- Management-layer security:

  Strengthens management and prevents hidden risks. Management-layer maintenance is penetrating in all the preceding layers.

  The FusionManager product document provides security management suggestions for you, including periodically changing your passwords and updating your certificates.

Application-Layer Security

Table 4-3 describes the policies applied for ensuring application-layer security.

Table 4-3 Application-layer security policies

Security Policy	Description
Account security	• On FusionManager, administrators can change user passwords periodically to ensure password security. • The password must meet the password strength requirements to prevent brute force cracking. • You can configure login using the specified IP address at the specified time period to improve system security.
Log security	• Administrators can view logs to ascertain system running status and operation records, thereby managing user behaviors and locating problems. • An operation log records the operations performed by a user on the system, for example, logging in to the system, logging out of the system, or creating a VM, as well as the result of the operation. Operation logs can help administrators check whether the system is under attacks or if malicious operations have been performed.
Web security	• A verification code is required on the login page. On the web-based login page, the system generates a random verification code. A user can log in to the system only when the username, password, and verification code have been entered correctly.

Security Policy	Description
	• The Web management system will automatically exit if it is not used for a long period of time. • The system supports defense against web application attacks, such as SQL injection and cross-site scripting.
Rights- and domain-based management	FusionManager provides comprehensive rights management functions, enabling users to perform operations in different virtual data centers (VDCs) and domains. This helps isolate user data and ensure the security of system resources. For details, see 4.2.1 Role-based Access Control.

System-Layer Security

Table 4-4 describes the policies applied for ensuring system-layer security.

Table 4-4 System-layer security policies

Security Policy	Description
OS security	FusionManager uses a SUSE Linux OS. The following are the basic security settings that are implemented on the SUSE Linux OS to ensure its secure operating: • Disables unnecessary services, such as Telnet services. • Hardens the secrete shell (SSH) service. • Controls access permission for files and directories. • Records operation logs.
Data security	Basic security settings are implemented to ensure the secure operating of databases. The following security-related measures are taken on a GaussDB database: • Logs the operations performed on the GaussDB database. • Prevents remote access to the database. • Backs up data to restore the database in the event of a database failure.
Whitelist mechanism	The whitelist mechanism is added on the basis of OS hardening to open only necessary commands and files, which helps better protect application systems. Users can upload software packages or other files only after adding signatures on them.
Security patches	Software design defects cause many system vulnerabilities. Regular installation of system security patches can eliminate system vulnerabilities and prevent viruses, worms, and hackers from using these vulnerabilities to attack the system. FusionManager provides the following security patch scheme: • OS security patch: Compatibility tests are periodically performed on

Security Policy	Description
	OS patches and compatible OS patches are released for users to install as required. • FusionManager system security patches: Security patch packages and patch installation guides are released based on security requirements. Users can install patches to fix loopholes.

Network-Layer Security

Table 4-5 describes the policies applied for ensuring network-layer security.

Table 4-5 Network-layer security policies

Security Policy	Description
Defense against malformed packet attacks	Because FusionManager interacts with end users on untrusted networks, it may be vulnerable to malformed packet attacks. FusionManager has been fully tested using tools, such as Codenomicon and xDefend, on its ability to defend against malformed packet floods, ensuring the security of the FusionManager system during interaction with end users.
Isolation between administrator and tenant networks	The administrator network and tenant network are isolated. When logging in to the web client, you can select to enter the administrator view or tenant view. This effectively demarcates the respective resources managed by the administrator and tenant and ensures data security.
VLAN and subnet	VLANs and subnets are used to isolate user VMs, ensuring VM data security.
VPC	A VPC provides VDCs with isolated, secure networks, thereby ensuring network security of each VDC.
VPN	A VPN provides an encrypted communication tunnel between remote users and their VPCs so that the users can directly obtain VPC service resources through the VPN. FusionManager supports two VPN protocols, Internet Protocol Security (IPSec) and Layer Two Tunneling Protocol (L2TP). • The IPSec VPN sets up VPN connections over IPSec. • The L2TP VPN encapsulates packets using L2TP first and then using IPSec.
ACL	The ACL is an access security control method that allows authorized users to access only the specified network resources by controlling data packets sent to or from a port. ACLs can function in both the interzone and intrazone. Intrazone ACL rules control the communication within a VPC, and interzone ACL rules control the communication between a VPC and the public network.
Security group	A security group is a logical group consisting of VMs that pose the same

Security Policy	Description
	security requirements and trust one another in the same region. The security group, functioning similarly to a firewall, is an important security isolation measure that is used to control the network access of one or more cloud servers. Each VM belongs to at least one security group. You must specify its security group when creating a VM. VMs in different security groups cannot communicate with each other by default. However, you can configure settings to allow mutual access between two security groups.

Management-Layer Security

Table 4-6 describes the policies applied for ensuring management-layer security.

Table 4-6 Management-layer security policies

Security Policy	Description
Periodically changing passwords	Huawei recommends that you periodically change your passwords, and the FusionManager product document provides the methods for changing the passwords of all system-wide accounts.
Periodically updating certificates	Huawei recommends that you periodically update your certificates, and the FusionManager product document provides the methods for updating all system-wide certificates.

5 Key Features

5.1 Cross-Host VM Live Migration

This feature allows VMs to be live migrated from one server to another without interrupting user services. The feature applies to planned server maintenance requiring no stop of any user services.

Definition

The live migration feature allows users to migrate VMs from one physical server to another physical server without interrupting services. The VM manager provides quick recovery of memory data and memory sharing technologies to ensure that the VM data remains unchanged before and after the live migration. The VM live migration applies to the following scenarios:

- Before performing operation and maintenance (O&M) operations on a physical server, system maintenance engineers can relocate VMs from this physical server to another physical server. This minimizes risk of service interruption during the O&M process.
- Before upgrading a physical server, system maintenance engineers can relocate VMs from this physical server to other physical servers. This minimizes risk of service interruption during the upgrade process. After the upgrade is complete, system maintenance engineers can relocate the VMs to the original physical server.
- System maintenance engineers can relocate VMs from a light-loaded server to other servers and then power off the server. This helps reduce service operation costs.

Table 5-1 describes the types of VM live migration.

Table 5-1 Types of VM live migration

Migration Type	Subclass	Description
Manual migration	By destination	On the FusionCompute web client, system maintenance engineers manually relocate one VM to another server.
Automatic migration	VM resource scheduling	The system automatically relocates VMs to other servers in the cluster based on the preset VM scheduling

Migration Type	Subclass	Description
		policies.

Benefits

For...	Benefits
Customers	The feature applies to planned server maintenance requiring no stop of any user services.

Dependency

None

5.2 Smart Memory Overcommitment

This feature allows a server to provide virtual memory that can be larger than the server's physical memory size, using various memory technologies, such as memory ballooning, memory sharing, and memory swapping.

Definition

Memory overcommitment allows a VM to use more memory space than the physical host has. Commonly used technologies include memory ballooning, memory sharing, and memory swapping. This feature allows more VMs to be supported by a server since it offers more memory resources than the physical server has.

This feature increases memory utilization, reduces the investment on storage devices, and prolongs the memory service time for servers.

FusionCompute supports the following memory overcommitment technologies:

- Memory ballooning: The system automatically reclaims the unused memory from a VM and allocates it to other VMs to use. Applications on the VMs are not aware of memory reclamation and allocation. The total amount of the memory used by all VMs on a physical server cannot exceed the physical memory of the server.
- Memory swapping: The system swaps out the content on the reserved VM memory to an external storage file to free the reserved memory and gets back the content when required.
- Memory sharing: The VMs that have the same memory content share one memory page.

The memory overcommitment degree is inversely proportional to the actual VM memory usage. Therefore, you must specify the QoS of the memory overcommitment for computing nodes.

After memory overcommitment is enabled, the memory overcommitment policy is used to allocate memory. VMs can use all physical memory when the memory is sufficient. If the memory is insufficient, the system schedules memory resources based on the memory

overcommitment policies by using memory overcommitment technologies to release free memory.

Benefits

For...	Benefits
Customers	This feature helps reduce operating costs. • The feature helps increase VM density when the memory size of computing nodes is fixed. • The feature eliminates the need for storage devices when the VM density of computing nodes is fixed.

Dependency

The total memory size reserved for all VMs running on each computing node cannot exceed the total memory size of virtualization domains on the computing node.

Constraints

- Sufficient space for memory swapping must be configured for hosts to ensure stable running of the memory overcommitment function. The maximum memory overcommitment ratio depends on the swap partition size. The specific calculation formula is as follows:

  Maximum memory overcommitment ratio supported by a host = 1 + (Size of the swap partition — Physical memory size of the virtualization domain x 0.1)/Physical memory size of the virtualization domain
- The memory swap partition and the host OS are configured on the same disk by default (Default size = 30GB). A maximum of 150% overcommitment ratio is supported. If you manually configure the disk size, a minimum of 30 GB is required.

  NOTE

  The default system memory swap partition does not have its own data store. The default memory swap partition is **/dev/xxx**.
- Memory overcommitment is exclusive of SR-IOV passthrough, GPU passthrough, and NVME SSD disk passthrough. Passthrough VMs must exclusively occupy the memory. The memory exclusively used by VMs cannot be exchanged to the space for memory swapping. Memory overcommitment-enabled VMs with memory reservation less than 100% cannot be bound with physical devices.

Precautions

- When the memory usage of a host reaches more than 70%, the VM services on the host are memory-consuming. You are not advised to enable memory overcommitment. If memory overcommitment is enabled, the memory is probably insufficient and the memory swap policy is used to release free memory. As a result, the performance of the VM for without full memory reserved deteriorates.
- After memory overcommitment is enabled, if the host memory is insufficient, the memory performance deteriorates for VMs with low reserved memory. The host memory overcommitment is disabled by default.
- After memory overcommitment is enabled, adjust the alarm threshold of **ALM-15.1000033 Host Memory Usage Exceeds the Threshold** to 80%. If the alarm is

generated and the alarm object is a host in a cluster with memory overcommitment enabled, migrate VMs on the host or stop some VMs on the host based on *Alarm Handling* to prevent large-scale memory swap.

- To balance I/O pressure of the memory swapping on local disks and maximize the memory overcommitment function, you are advised to configure the memory swap partitions on different local virtualized data stores of the same host and deploy the host OS and memory swap partitions on different disks. In addition, you are advised to use high-performance local SSDs as swap partitions. If no free disk is used for memory swapping, check **ALM-15.1000033 Host Memory Usage Exceeds the Threshold** every day (recommended). If the alarm object is a host in a cluster with memory overcommitment enabled, handle the alarm in a timely manner.
- The recommended memory overcommitment ratio is less than 150%. A major alarm is generated if the memory overcommitment ratio exceeds 120%. A critical alarm is generated if the actually used memory size of a VM is greater than 90% of the total memory size of the virtualization domain and swap partition.
- When memory overcommitment is enabled, perform live migration operations. If some VM memory is swapped to the memory swap disk, the migration time becomes long.

5.3 VM HA

This feature allows VMs on a server to automatically start on another properly-running server within a few minutes if their original server becomes faulty. The services running on the VMs can also automatically recover on the new server after the VMs are moved to it.

Definition

The VM high availability (HA) feature ensures quick restoration of a VM. If a VM is faulty, the system automatically recreates the VM on another normal computing node.

When the system detects that a VM is faulty, the system selects a normal host and recreates the VM on the host. The VM HA feature protects VMs against the following failures:

- Restart or restoration of a computing node from a power failure

  When a computing node restarts or restores from a power failure, the system recreates the HA VMs on another computing node.
- Blue screen of death (BSOD) of a VM

  When the system detects that BSOD occurs on a VM and the handling policy configured for this error is HA, the system recreates the VM on another normal computing node.

Benefits

For...	Benefits
Customers	The VM high availability (HA) feature ensures quick restoration of a VM.

Dependency

Ensure that sufficient resources are available for VM HA.

5.4 VM Storage Live Migration

This feature allows VM disks to migrate to another storage source without interrupting services running on the VM. The feature applies to planned storage resource maintenance or migration requiring no stop of any user services.

Definition

This function enables disks on VMs to be migrated to other storage units when the VMs are running. The disks can be migrated between different storage devices or storage units on one storage device under virtual storage management. With this function enabled, storage resources of VMs can be dynamically migrated, thereby facilitating device maintenance.

Benefits

For...	Benefits
Customers	The feature applies to planned storage resource maintenance or migration requiring no stop of any user services.

Dependency

- A **Sharing** disk that has been attached to a VM and disks on a linked clone cannot be migrated.
- A VM in the **Running** state does not allow non-persistent disks to be migrated. Migrate disks after stopping VMs if permitted.
- Migration across FusionStorage Block storage resources is not supported.

5.5 Thin Provisioning

This feature enables dynamic allocation of shared storage resources, allowing users to apply for storage space based on their demands rather than the IT department's plan, and thereby helping customers reduce costs on storage devices by over 50%.

Definition

Thin Provisioning provides users with larger virtual storage space than the actual memory space. The system allocates physical storage space only when data is written into the virtual storage.

Thin Provisioning of the FusionSphere virtualization suite does not depend on the storage device.

Thin Provisioning applies to the user data volumes of VMs. When the declared storage capacity is far beyond user requirements, Thin Provisioning can be used to reduce initial investments for carriers.

Benefits

For...	Benefits

For...	Benefits
Customers	Thin Provisioning increases storage resource utilization and helps customers reduce initial investment on storage.

Dependency

None.

5.6 Anti-Virus Virtualization

This feature allows IT administrators to specify a dedicated VM on the anti-virus virtualization management interface to implement anti-virus measures for all VMs in the system in batches. In this manner, administrators do not need to install an independent anti-virus agent on each VM, preventing anti-virus storms from occurring in the system. The IT administrators can also manage anti-virus software and virus libraries on the anti-virus virtualization management interface in a unified manner, ensuring high security of the entire system.

Description

To protect VMs on hosts against virus attacks, the antivirus function is required. However, if traditional antivirus products are used, the products must be installed on each VM. In this case, the products occupy VM resources and even may cause an antivirus storm when users perform a global scan or antivirus update. To address this problem, FusionSphere provides dedicated antivirus application programming interfaces (APIs), which support secondary development by antivirus product vendors, to offer a VM antivirus solution. In this solution, the antivirus engine is installed on a dedicated secure VM on a host, and a lightweight antivirus driver is installed on the other VMs on the host. All the other user VMs can scan for and remove viruses using the services provided by the secure VM, consuming only a few VM resources.

Benefits

For...	Benefits
Customers	• This feature implements central management on antivirus services. Customers do not need to install and update antivirus databases on each VM. • This feature protects VMs against antivirus storms.

Dependency

The FusionCompute antivirus virtualization function has the following restrictions on the system:

- Only one secure service VM can be deployed on one host.

- For details about how to obtain the OSs supported by guest virtual machines (GVMs), visit Compatibility check assistant.

The FusionCompute antivirus virtualization function has the following restrictions on secure service VMs:

- Memory snapshots cannot be taken for secure service VMs.
- Secure service VMs cannot be stopped or hibernated, because they must provide real-time services.
- A secure service VM cannot be automatically migrated from a host due to dynamic resource scheduling or a host failure, because it is bound to the host.

The FusionCompute antivirus virtualization function has the following restrictions on GVMs:

- Memory snapshots cannot be taken for secure user VMs.
- A secure user VM must be migrated to a host that has a secure service VM deployed due to dynamic resource scheduling or a host failure, because secure user VMs use antivirus functions provided by a secure service VM.

  Therefore, each host in a cluster must have the antivirus virtualization function enabled and a secure service VM deployed.

5.7 Disaster Recovery and Backup

This feature enables enterprises to set up a production-site database and a backup database and to use the UltraVR component of FusionSphere to back up production site data to the backup database. If the production site encounters a disaster, such as a global power failure, earthquake, or fire, the data backed up on the backup database can be used to rapidly restore services, preventing services from being interrupted for a long time in the event of a disaster on the production site.

Definition

Disaster recovery (DR) is the ability to provide continuous services after unexpected problems, such as fire or earthquake. This is achieved by setting up two or more IT systems of the same function in remotely dispersed areas. When one system stops, another system will take over the services from the faulty system.

Backup is the process of copying data to a dump device. A dump device is a tape or disk used to store data copies. When a system is faulty or data loss occurs, the data copy stored in dump devices can be used to restore the system or data.

The FusionSpherevirtualization suite offers the following disaster recovery (DR) plans: metropolitan active-active DR, array-based replication DR, and two-site and three-center DR. It also offers the VM backup and user data backup plans. Customers can choose these plans based on service requirements.

- The metropolitan active-active DR plan allows two sites far from each other to use the HyperMetro feature of the Huawei V3, V5, or Dorado series storage and the high availability (HA) and dynamic resource scheduler (DRS) functions of FusionCompute to implement DR. The two sites (production sites) provide services and serve as DR sites for each other.
- The array copy disaster recovery (DR) scheme is implemented by creating two sites, a production site and a DR site, at two places, using the remote replication function of storage devices to copy the VM data from the production site to the DR site, and using

the DR management software UltraVR to register the VMs on the DR storage at the DR site to a hypervisor and automatically start the VMs.

- The VM backup plan uses the Huawei eBackup software combined with the FusionCompute snapshot backup function and the Changed Block Tracking (CBT) backup function to back up VM data. Working together with FusionCompute, eBackup can back up a specific object based on the configured policy. When a VM is faulty or its data is lost, the VM can be restored using the backup data. The backup data is stored on the shared storage devices connected to the eBackup. The snapshot-based backup and CBT-based backup support full backup and incremental backup.

Benefits

For...	Benefits
Enterprises and carriers	• This feature shortens the system downtime after a disaster occurs. • This feature allows important data to be rapidly restored, minimizing the impact of data loss. • This feature enhances service reliability.

Dependency

For details see DR and Backup User Guide.

5.8 Dynamic Resource Scheduling

FusionSphere monitors global resource usages when creating VMs and when VMs are running. In the monitoring process, FusionSphere uses a dynamic resource scheduling algorithm to determine the optimal host on which the VMs can run, and it also moves the VMs to this optimal host by means such as live migration, ensuring VM stability and improving user experience.

Definition

The distributed resource scheduler (DRS) feature uses intelligent scheduling algorithms to periodically monitor the work load on hosts in a cluster and migrates VMs between the hosts based on the work load to implement load balance. This feature collaborates with the dynamic power management (DPM) to increase resource utilization and reduce power consumption.

- When the system is lightly loaded, the system migrates some VMs to one or more physical hosts, and powers off the idle hosts.
- When the system is heavily loaded, the system starts some VMs and physical hosts and allocates VMs evenly on hosts to ensure resource supply.
- Scheduled tasks can be set to enable different resource scheduling policies at different times based on the system running status to meet user requirements in different scenarios.

Benefits

For...	Benefits

For...	Benefits
Enterprises and carriers	This feature optimizes resource allocation in different scenarios, reduces power consumption, and improves resource utilization.

Dependency

Dynamic resource scheduling does not take effect on a VM that has been bound to a host, a USB flash drive, a passthrough device, or GPU resource group.

5.9 VM Resource QoS Control

This feature allows IT administrators to set an upper limit of resources, such as CPUs, memory, networks, and disk input/output per second (IOPS), available to a VM, preventing non-critical applications or malicious users from preempting shared resources.

Definition

- CPU Quality of service (QoS)

  The CPU QoS ensures optimal allocation of computing resources for VMs and prevents resource contention between VMs due to different service requirements. It effectively increases resource utilization and reduces costs.

  During creation of VMs, the CPU QoS is specified based on the service to be deployed. After the VMs are created, the system dynamically binds the vCPUs to physical CPUs based on the CPU QoS. In this way, a pool of physical CPUs that are bound to different vCPUs with the same CPU QoS is created on a server.

  The CPU QoS determines the VM computing power. The system ensures the VM CPU QoS by setting the minimum computing capability and the computing capability upper limit for VMs.

  CPU QoS contains the following parameters:

  - **CPU quota**

    CPU quota defines the proportion based on which CPU resources to be allocated to each VM when multiple VMs compete for physical CPU resources.

    This section uses a host (physical server) that uses a single-core, 2.8 GHz CPU as an example to describe how CPU quota works. Three VMs (A, B, and C) run on the host, and their quotas are set to 1000, 2000, and 4000, respectively. When the CPU workloads of the VMs are heavy, the system allocates CPU resources to the VMs based on the CPU quotas. VM A with 1000 CPU quota can obtain a computing capability of 400 MHz. VM B with 2000 CPU quota can obtain a computing capability of 800 MHz. VM C with 4000 CPU quota can obtain a computing capability of 1600 MHz. The computing capability calculation is more complex in actual use.

    The CPU quota takes effect only when resource contention occurs among VMs. If the CPU resources are sufficient, a VM can exclusively use physical CPU resources on the host if required. For example, if VMs B and C are idle, VM A can obtain all of the 2.8 GHz computing capability.

  - **CPU reservation**

CPU reservation defines the minimum CPU resources to be allocated to each VM when multiple VMs compete for physical CPU resources.

If the computing capability calculated based on the CPU quota of a VM is less than the CPU reservation value, the system allocates the computing capability to the VM according to the CPU reservation value. The offset between the computing capability calculated based on the CPU quota and the CPU reservation value is deducted from computing capability of other VMs based on their CPU quotas and is added to the VM.

If the computing capability calculated based on the CPU quota of a VM is greater than the CPU reservation value, the system allocates the capability to the VM according to the CPU quota.

For example, three VMs (A, B, and C) run on the host that uses a single-core physical CPU, their quotas are set to 1000, 2000, and 4000, respectively, and their CPU reservation values are set to 700 MHz, 0 MHz, and 0 MHz, respectively. When the CPU workloads of the three VMs are heavy:

- According to the VM A CPU quota, VM A should have obtained a computing capability of 400 MHz. However, its CPU reservation value is greater than 400 MHz. Therefore, VM A obtains a computing capability of 700 MHz according to its CPU reservation value.
- The system deducts the offset (700 MHz minus 400 MHz) from VMs B and C based on their CPU quota.
- VM B obtains a computing capability of 700 (800 minus 100) MHz, and VM C obtains a computing capability of 1400 (1600 minus 200) MHz.

The CPU reservation takes effect only when resource contention occurs among VMs. If the CPU resources are sufficient, a VM can exclusively use physical CPU resources on the host if required. For example, if VMs B and C are idle, VM A can obtain all of the 2.8 GHz computing capability.

- **CPU limit**

  CPU limit defines the upper limit of physical resources that can be used by a VM. For example, if a VM with two virtual CPUs has a CPU limit of 3 GHz, each virtual CPU of the VM can obtain a maximum of 1.5 GHz computing resources.

- Memory QoS

  Memory QoS allows VM memory to be intelligently allocated based on the preset percentage of reserved memory. Memory overcommitment technologies, such as memory ballooning, are used to provide more virtual memory resources. Users are unaware of the memory overcommitment.

  The user can set the reserved memory percentage based on service requirements. The main principle of memory overcommitment is to first use the physical memory.

  Memory QoS contain the following parameters:

  - **Memory quota**

    Memory quota defines the proportion based on which memory resources to be allocated to each VM when multiple VMs compete for physical memory resources.

    The system allocates memory resources to VMs based on the proportion when VMs apply for memory resources or hosts release free memory resources (such as when VMs are migrated or stopped).

    CPU resources can be scheduled in real time. Memory resources are scheduled subtly and continuously when VMs are running until the configured VM memory resources are allocated to VMs.

For example, three 4 GB memory VMs run on a 6 GB memory host, and their memory quotas are set to 20480, 20480, and 40960, respectively. In this case, the memory allocation ratio is 1:1:2. When the memory workloads on the three VMs gradually increase, the system subtly adjusts memory resources on the VMs based on the memory quotas until the VMs obtain 1.5 GB, 1.5 GB, and 3 GB memory, respectively.

The memory quota takes effect only when resource contention occurs among VMs. If the memory resources are sufficient, a VM can exclusively use physical memory resources on the host if required. For example, if the memory resources required by VMs B and C are less than the reserved memory values, and VM A has more memory workloads to handle, VM A can use memory resources from free memory resources, and memory resources on VMs B and C, until the memory resources obtained by VM A reach the upper limit, or the free memory resources are used up and memory resources on VMs B and C drop to the reserved values. For example, if VM C is not under memory pressure and has 1 GB memory reserved, VMs A and B theoretically can obtain a maximum of 2.5 GB memory resources each.

- **Memory reservation**

  Memory reservation defines the minimum memory resources to be allocated to each VM when multiple VMs compete for memory resources.

  A VM exclusively uses its reserved memory, that is, if certain memory resources are reserved for a VM, other VMs cannot use the memory resources even if the memory resources are idle.

- **Memory limit**

  Memory limit defines the upper limit of physical resources that can be used by a VM. For example, if a VM with two virtual CPUs has a CPU limit of 3 GHz, each virtual CPU of the VM can obtain a maximum of 1.5 GHz computing resources. When multiple VMs are started, they will complete for memory resources. To improve memory utilization and reduce idle memory, users can set the memory limit parameter in the configuration file when creating a VM so that the memory allocated to the VM does not exceed the upper limit.

- The network QoS policy enables bandwidth configuration control. The QoS function does not support traffic control among VMs on the same host.
  - Transmit and receive bandwidth control based on a port group member port

    Traffic shaping and bandwidth priority are configured for each port in a port group to ensure network QoS.

Benefits

For...	Benefits
Customers	This feature allows IT administrators to set an upper limit of resources, available to a VM, preventing non-critical applications or malicious users from preempting shared resources.

Dependency

None.

5.10 User-Mode Switching Mode

User-mode switching mode provides high-performance virtual network processing capability to improve network performance of user VMs.

Definition

Data Plane Development Kit (DPDK) is a set of libraries and drivers and used for fast data packet processing on the x86 platform. It uses multiple technologies, including the bypass kernel protocol stack at the abstraction layer, uninterrupted packet sending and receiving in polling mode, memory, buffer area, and queue management optimization, and load balancing among multiple NIC queues and data flows, achieving high-performance packet forwarding in the x86 processor framework and improving VM network performance.

Benefits

For...	Benefit
Customers	Provide high-performance network processing capability.

Dependency

- User-mode switching mode supports Intel 82599ES and XL710 NICs.
- To use user-mode switching mode, perform operations based on Configuring the Huge Page Memory for a Host first and then Configuring User-Mode Switching Specifications for a Host.

5.11 Distributed Virtual Switch

FusionSphere uses distributed virtual switches (DVSs) to offer advanced network functions and manage the virtual network of the entire system in a unified manner, thereby simplifying and enhancing the VM network connections in the system.

Definition

Distributed virtual switch (DVS) management allows the system administrator to configure and maintain physical and virtual ports of DVSs on one or multiple Computing Node Agent (CNA) servers.

Figure 5-1 shows the DVS model.

Figure 5-1 DVS model

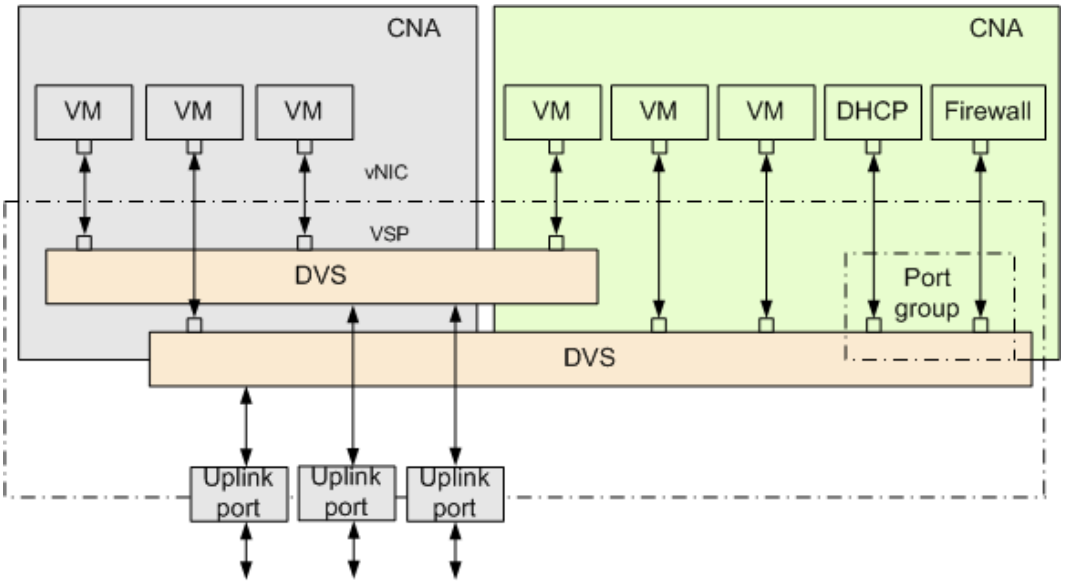

Firewall	Uplink Port: cascading port	vNIC: virtual network interface card
VSP: virtual switch port	N/A	N/A

The DVS model has the following characteristics:

- Multiple DVSs can be configured, and each DVS can serve multiple CNA nodes in a cluster.
- A DVS provides several virtual switch ports (VSP) with their own attributes, such as the rate, statistics, and access control lists (ACL). The ports with the same attributes are assigned to a port group for management. The port groups with the same attributes are allocated to the same virtual local area network (VLAN).
- An uplink port group can be configured for each DVS to enable external communication for VMs. An uplink port group comprises multiple physical NICs working based on load-balancing policies.
- Each VM provides multiple virtual network interface card (vNIC) ports, which connect to VSPs of the switch in one-to-one mapping.

Benefits

For...	Benefits
Customers	• The cloud computing management system, which integrates DVS management, implements centralized management of the virtual networks of all CNA nodes. This significantly reduces the management workload and minimizes misoperations. • Visualized network management

For...	Benefits
	provides customers with a clear virtual network topology and traffic information, helping customers easily maintain the network.

5.12 SR-IOV

The single-root I/O virtualization (SR-IOV) feature enables FusionSphere to use all advantages brought by physical intelligent network interface card performance acceleration.

Definition

SR-IOV enables a physical Peripheral Component Interconnect Express (PCIe) to be shared on a virtual environment using multiple virtual interfaces, offering different virtual functions to different virtual components on a physical server machine. SR-IOV directly allocates I/O data to VMs, which allows the I/O data to bypass the software emulation layer, thereby reducing the I/O overhead at the software emulation layer.

Benefits

For...	Benefits
Customers	Use hardware device functions on networks, which reduces the I/O overhead at the software emulation layer.

Dependency

- SR-IOV can be enabled for Intel 82599ES. If SR-IOV is enabled, a network port of the Intel 82599ES NIC supports a maximum of 63 virtual network ports.
- To create port groups on the SR-IOV-enabled DVS, the port type can only be set to **Access**.
- If a VM uses the SR-IOV-enabled NICs, the following functions become unavailable to the VM: VM hibernation, VM waking up, VM live migration, VM migration as a whole, memory snapshot, memory hot add, NIC hot add and delete, IP-MAC address binding, and security group.
- If the SR-IOV-enabled NICs are used, the VMs can run one of the operating systems (OSs) provided in the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*.

NOTE

To obtain the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*, visit the following websites:

- For enterprises: Visit http://support.huawei.com/enterprise, search for the document by name, and download the document of the required version.
- For carriers: Visit http://support.huawei.com, search for the document by name, and download the document of the required version.

5.13 GPU Passthrough

This feature allows FusionSphere to use the graphic processing capabilities of powerful graphics processing units (GPUs).

Definition

A GPU on a physical server can be directly attached to a specified VM to improve VM graphic processing capabilities.

Benefits

For...	Benefits
Customers	This feature allows customers to use the graphic processing capabilities of powerful GPUs.

Dependency

Table 5-2 and Table 5-3 list the GPUs in passthrough mode and OSs that are compatible with each other.

Table 5-2 GPU passthrough device models and guest OSs (Tesla driver: the CUDA version in the brackets indicates that corresponding to the driver version)

GPU	Vendor ID	Device ID	Supported Latest Driver	Guest OS
NVIDIA Tesla P100	0x10de	0x15f8	396.37 (CUDA 9.2)	Ubuntu Server 16.04 64-bit Debian GNU/Linux 8.0.0 64-bit Windows Server 2012 R2 64-bit EulerOS 2.2 64-bit CentOS 7.3 64-bit
NVIDIA Tesla P4	0x10de	0x1bb3	396.37 (CUDA 9.2)	Ubuntu Server 14.04 64-bit Ubuntu Server 16.04 64-bit Redhat 6.5 64-bit CentOS 7.2 64-bit CentOS 7.4 64-bit Windows Server 2012 R2 64-bit EulerOS 2.3 64-bit YITU (Ubuntu 14.04) 64-bit
NVIDIA	0x10de	0x1b38	396.37	Ubuntu Server 14.04.4

GPU	Vendor ID	Device ID	Supported Latest Driver	Guest OS
Tesla P40			(CUDA 9.2)	64-bit Red Hat 6.5 64-bit CentOS 7.2 64-bit Windows Server 2012 R2 64-bit YITU (Ubuntu 14.04) 64-bit Windows 10 (only support for TH2 and above builds)

Table 5-3 GPUs in passthrough mode and guest OSs (GRID 6.2 driver)

GuestOS\GPU	M60	P40	P4
Windows Server 2008 R2 64 bits (Standard/Data Center/Enterprise)	√	×	×
Windows Server 2008 R2 SP1 64 bits (Standard/Data Center/Enterprise)	√	×	×
Windows Server 2012 R2 64 bits (Standard/Data Center)	√	√	×
Windows Server 2016 64 bits (Essentials/Standard /Data Center)	√	√	√
Windows 7 64 bits (Professional/Ultimate)	√	×	×
Windows 10 64 bits (Professional/Enterprise/Ent_2016_LTSB/Pro_Edu)	√ Not supported for 1511 and 1703	√ Not supported for 1507, 1511, and 1703	√ Not supported for 1507, 1511, and 1703

When a GPU is set to the passthrough mode, for details about the compatible OSs, see the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*.

NOTE

To obtain the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*, visit the following websites:

- For enterprises: Visit http://support.huawei.com/enterprise, search for the document by name, and download the document of the required version.
- For carriers: Visit http://support.huawei.com, search for the document by name, and download the document of the required version.

After a GPU is attached to a VM, the following functions are affected:

- VM HA: Only the automatically distributed VM with a GPU bound support HA. Other VMs with GPUs bound do not support HA.
- VM live migration: A VM with a GPU bound does not support live migration.
- VM hibernation: A VM with a GPU bound does not support hibernation.
- VM snapshot: A VM with a GPU bound does not support memory snapshot creation.
- VM online cloning: A VM with a GPU bound does not support online cloning.
- VM export: A VM with a GPU bound cannot be exported.
- Cluster scheduling policy: This policy does not take effect for a VM with a GPU bound.
- VNC login: VNC login is not supported for the VM with a GPU in passthrough or graphics passthrough mode. Users need to remotely log in to the VM.

5.14 GPU Virtualization

In the FusionSphere virtualization solution, GPU virtualization allows GPUs to be shared and users to remotely process the VDI (Virtual Desktop Infrastructure).

Definition

GPUs on a physical server can be virtualized into multiple vGPUs used by multiple VMs based on the hardware technology. GPUs provide the VMs with 2D graphics the processing and 3D graphics rendering acceleration services and feature high performance and low costs. In this way, the GPUs are shared and user costs are lowered.

Benefits

For...	Benefits
Customers	Shares GPUs and remotely processes the VDI.

Dependency

Table 5-4 lists the GPU virtualization device models and OSs that are compatible with each other.

Table 5-4 Guest OSs supported by GPU virtualization devices (GRID 6.2 driver)

GuestOS\GPU	M60	P40	P4
Windows Server	√	√	√

GuestOS\GPU	M60	P40	P4
2008 R2 64 bits (Standard/Data Center/Enterprise)			
Windows Server 2008 R2 SP1 64 bits (Standard/Data Center/Enterprise)	√	√	Not verified
Windows Server 2012 R2 64 bits (Standard/Data Center)	√	√	√
Windows Server 2016 64 bits (Essentials/Standard /Data Center)	√	√	√
Windows 7 32/64 bits (Professional/Ultimate)	√	√	√
Windows 10 32/64 bits (Professional/Enterprise/Ent_2016_LTSB/Pro_Edu)	√ Not supported for 1511 and 1703	√ Not supported for 1507, 1511, and 1703	√ Not supported for 1507, 1511, and 1703

When a GPU is set to virtualization, for details about the compatible OSs, see the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*.

NOTE

To obtain the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*, visit the following websites:

- For enterprises: Visit http://support.huawei.com/enterprise, search for the document by name, and download the document of the required version.
- For carriers: Visit http://support.huawei.com, search for the document by name, and download the document of the required version.

After a GPU is attached to a VM, the following functions are affected:

- VM live migration: A VM with a GPU bound does not support live migration.
- VM hibernation: A VM with a GPU bound does not support hibernation.
- VM snapshot: A VM with a GPU bound does not support memory snapshot creation.
- VNC login: VNC login is not supported for the VM with a GPU in passthrough or graphics passthrough mode. Users need to remotely log in to the VM.

6 System Principle

6.1 Communication Principles

Communication Planes

The FusionCompute system consists of the following communication planes:

- Management plane: provides a communication plane to implement system monitoring, operation and maintenance (including system configuration, system loading, and alarm reporting), and VM management (such as creating, deleting, and scheduling VMs).
- Storage plane: provides a communication plane for the storage system and storage resources for VMs. This plane is used for storing and accessing VM data (including data in the system disk and user disk of VMs).
- Service plane: provides a plane for virtual network interface cards (VNICs) of VMs to communicate with external devices.

Figure 6-1, Figure 6-2, and Figure 6-3 show the communication between planes in the FusionCompute system.

NOTE

This section uses IP storage area network (SAN) devices as an example to describe how storage devices in the FusionCompute system communicate with other planes.

Figure 6-1 Communication between planes (two NICs)

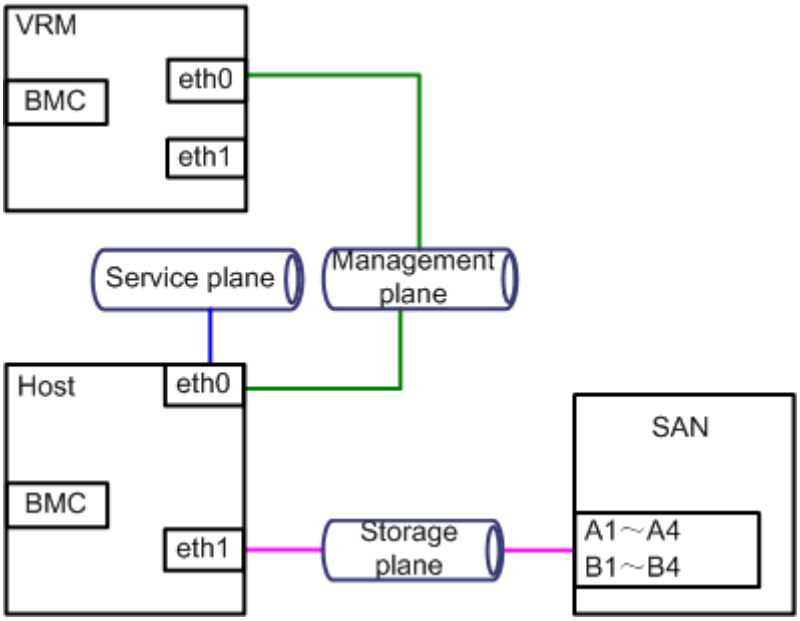

Figure 6-2 Communication between planes (four NICs)

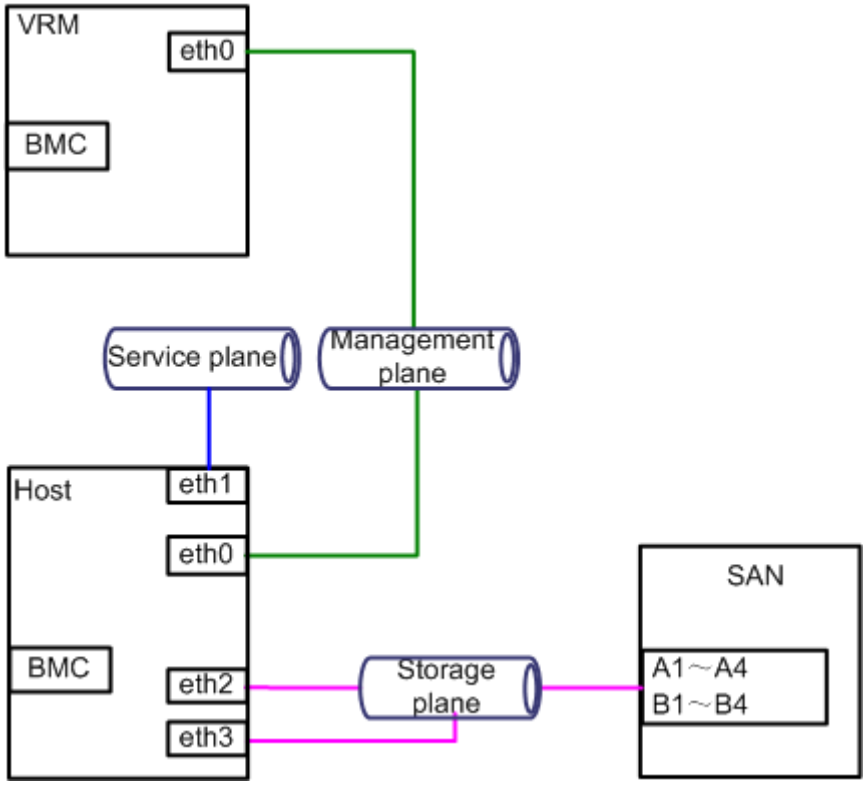

Figure 6-3 Communication between planes (six NICs)

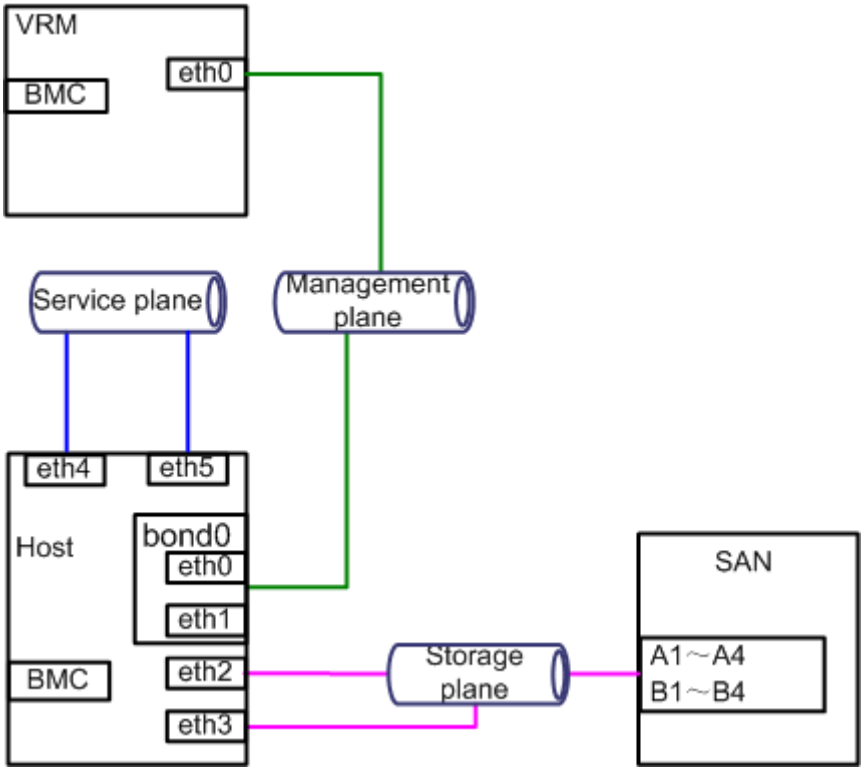

NOTE

VLAN Planning Rules

Table 6-1, Table 6-2 and Table 6-3 list the virtual local area network (VLAN) plan for the planes.

Table 6-1 VLAN plan (two NICs)

Plane	Network Ports	VLAN Planning Rules
Management plane	Network port eth0 on the host	The network port eth0 on each node is assigned to the management plane VLAN, and the VLAN which the port eth0 on the node belongs to becomes the default VLAN of the management plane.
	Network port eth0 on the active and standby Virtualization Resource Management (VRM) nodes	
	BMC network ports on the VRM and host	The switch port connected to the BMC network port on each node is assigned to the BMC plane VLAN, which then becomes the default VLAN for the BMC plane. **NOTE** The BMC network ports can be assigned to an independent BMC plane or to the same VLAN as management network ports. The

Plane	Network Ports	VLAN Planning Rules
		specific assignment depends on the actual network plan.
Storage plane	Storage network ports A1, A2, A3, A4, B1, B2, B3, and B4 on the SAN storage devices	The storage plane is divided into four VLANs: • A1 and B1 are assigned to VLAN 1. • A2 and B2 are assigned to VLAN 2. • A3 and B3 are assigned to VLAN 3. • A4 and B4 are assigned to VLAN 4. • eth1 is assigned to VLAN 1, VLAN 2, VLAN 3 and VLAN 4.
	Storage network port eth1 on the host	
Service plane	Service network port eth0 on the host	The service plane is divided into multiple VLANs to isolate VMs. All data packets from different VLANs are forwarded over the service network ports on the CNA. The data packets are marked with VLAN tags and sent to the service network port of the switch on the access layer.

Table 6-2 VLAN plan (four NICs)

Plane	Network Ports	VLAN Planning Rules
Management plane	Network port eth0 on the host	The network port eth0 on each node is assigned to the management plane VLAN, and the VLAN which the port eth0 on the node belongs to becomes the default VLAN of the management plane.
	Network port eth0 on the active and standby VRM nodes	
	BMC network ports on the VRM and host	The switch port connected to the BMC network port on each node is assigned to the BMC plane VLAN, and the VLAN which the BMC network port on the node belongs to becomes the default VLAN for the BMC plane. **NOTE** The BMC network ports can be assigned to an independent BMC plane or to the same VLAN as management network ports. The specific assignment depends on the actual network plan.
Storage plane	Storage network ports A1, A2, A3, A4, B1, B2, B3, and B4 on the SAN storage devices	The storage plane is divided into four VLANs: • A1 and B1 are assigned to VLAN 1. • A2 and B2 are assigned to VLAN 2. • A3 and B3 are assigned to VLAN 3. • A4 and B4 are assigned to VLAN 4. • eth2 is assigned to VLAN 1 and VLAN 2. • eth3 is assigned to VLAN 3 and VLAN 4. The network port eth2 can communicate with ports A1, A2, B1, and B2 over the layer 2 network. The network port eth3 can communicate with ports A3, A4, B3, and B4 over the layer 2 network. This allows computing resources to
	Storage network ports eth2 and eth3 on the host	

Plane	Network Ports	VLAN Planning Rules
		access storage resources through multiple paths. (Each computing server has eight Internet Small Computer Systems Interface (iSCSI) links to connect to the same storage device.) Therefore, the storage network reliability is ensured.
Service plane	Service network port eth1 on the host	The service plane is divided into multiple VLANs to isolate VMs. All data packets from different VLANs are forwarded over the service network ports on the CNA. The data packets are marked with VLAN tags and sent to the service network port of the switch on the access layer.

Table 6-3 VLAN plan (six NICs)

Plane	Network Ports	Virtual Network Port	VLAN Planning Rules
Management plane	Network ports eth0 and eth1 on the host	bond0	The network ports eth0 and eth1 on each node are assigned to the management plane VLAN, and the VLAN which the network ports eth0 and eth1 on each node belong to become the default VLAN of the management plane.
	Network port eth0 on the active and standby VRM nodes	-	
	BMC network ports on the VRM and host	-	The switch port connected to the BMC network port on each node is assigned to the BMC plane VLAN, and the VLAN which the BMC network port on the node belongs to becomes the default VLAN for the BMC plane. **NOTE** The BMC network ports can be assigned to an independent BMC plane or to the same VLAN as management network ports. The specific assignment depends on network planning.
Storage plane	Storage network ports A1, A2, A3, A4, B1, B2, B3, and B4 on the SAN storage devices	-	The storage plane is divided into four VLANs: • A1 and B1 are assigned to VLAN 1. • A2 and B2 are assigned to VLAN 2. • A3 and B3 are assigned to VLAN 3. • A4 and B4 are assigned to VLAN 4. • eth2 is assigned to VLAN 1 and VLAN 2. • eth3 is assigned to VLAN 3 and VLAN 4. The network port eth2 can communicate with ports A1, A2, B1, and B2 over the layer 2 network. The network port eth3 can communicate with ports A3, A4, B3, and B4
	Storage network ports eth2 and eth3 on the host	-	

Plane	Network Ports	Virtual Network Port	VLAN Planning Rules
			over the layer 2 network. This allows computing resources to access storage resources through multiple paths. (Each computing server has eight iSCSI links to connect to the same storage device.) Therefore, the storage network reliability is ensured.
Service plane	Service network ports eth4 and eth5 on the host	-	The service plane is divided into multiple VLANs to isolate VMs. All data packets from different VLANs are forwarded over the service network ports on the CNA. The data packets are marked with VLAN tags and sent to the service network port of the switch on the access layer.

6.2 Time Synchronization Mechanism

Figure 6-4 shows the time synchronization mechanism of the FusionSphere virtualization suite. Table 6-4 lists the components involved in FusionSphere virtualization suite time synchronization.

Figure 6-4 Time synchronization mechanism of the FusionSphere virtualization suite

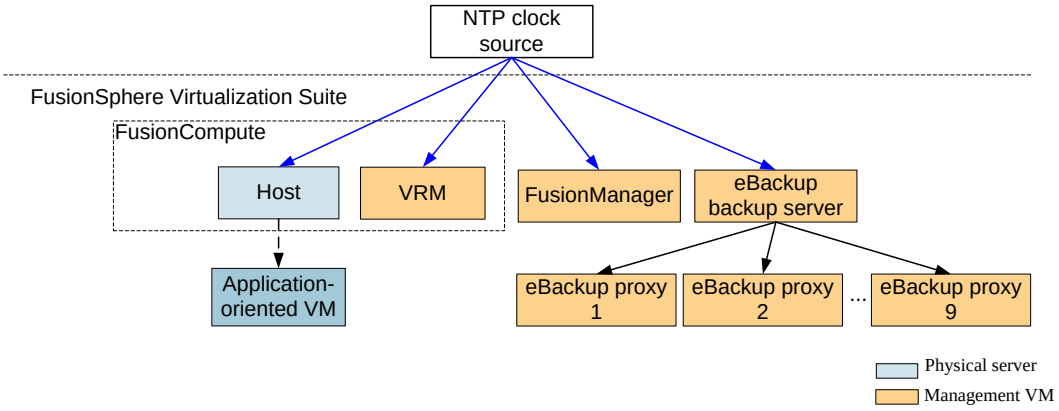

Table 6-4 Components involved in FusionSphere virtualization suite time synchronization

Component	Description
FusionCompute	To ensure precise system time in the FusionSphere virtualization suite, you are advised to configure an external Network Time Protocol (NTP) clock source to serve the FusionCompute and the FusionManager.
FusionManager	
eBackup backup	• Backup server: To ensure time accuracy, an external NTP clock source is required. The backup server synchronizes time with

Component	Description
server	the clock source. • Backup proxy: After the backup server is configured, the backup proxy synchronizes time with the backup server.
Application-oriented VM	Users can select a time synchronization policy based on their specific VM time precision requirements. • (Recommended) Free clock policy: Users customize a VM time synchronization configuration, with which VM time is not affected by the FusionSphere virtualization suite system time. To use this policy, set the clock policy to not synchronizing time with the host clock when you create a VM template. • Synchronizing time with the host on which the VMs are running: VM time is determined by the host time. To use this policy, set the clock policy to synchronizing time with the host clock when you create a VM template. If a user does not expect application-oriented VM time to be determined by the FusionCompute system time or the user is using a reliable clock source, the free clock policy is recommended.

7 Reliability

7.1 FusionSphere System Reliability

Management Nodes in Active/Standby Mode

FusionSphere management nodes are deployed in active/standby mode. If any fault occurs, such as a hardware failure, suspended key process, operating system (OS) panic, or network interruption, the standby mode automatically takes over for the active node in 1 or 2 minutes, continuing to provide virtualization management services.

OS Fault Locating Tool: Black Box

Each FusionCompute host OS has a black box embedded. This black box stores fault information about the virtualization OS kernel, facilitating kernel fault locating and restoration.

Management Data Backup and Restoration

FusionSphere periodically backs up configuration and service data on local and remote devices. If any management node service fails and cannot be automatically restored, FusionSphere allows users to rapidly restore the service using the local data backup. If a catastrophic failure occurs, both the active and standby management nodes are faulty and cannot be restored by restarting. FusionSphere allows users to restore these nodes using the remote data backup, thereby shortening the service restoration duration.

7.2 FusionCompute Software Reliability

VM HA

The VM high availability (HA) feature ensures quick restoration of a VM. If a VM is faulty, the system automatically recreates the VM on another normal computing node.

When the system detects that a VM is faulty, the system selects a normal host and recreates the VM on the host. The VM HA feature protects VMs against the following failures:

- Restart or restoration of a computing node from a power failure

When a computing node restarts or restores from a power failure, the system recreates the HA VMs on another computing node.

- Blue screen of death (BSOD) of a VM

  When the system detects that BSOD occurs on a VM and the handling policy configured for this error is HA, the system recreates the VM on another normal computing node.

VM Live Migration

The live migration feature allows users to migrate VMs from one physical server to another physical server without interrupting services. The VM manager provides quick recovery of memory data and memory sharing technologies to ensure that the VM data remains unchanged before and after the live migration. The VM live migration applies to the following scenarios:

- Before performing operation and maintenance (O&M) operations on a physical server, system maintenance engineers can relocate VMs from this physical server to another physical server. This minimizes the risk of service interruption during the O&M process.
- Before upgrading a physical server, system maintenance engineers can relocate VMs from this physical server to other physical servers. This minimizes the risk of service interruption during the upgrade process. After the upgrade is complete, system maintenance engineers can relocate the VMs to the original physical server.
- System maintenance engineers can relocate VMs from a light-loaded server to other servers and then power off the server. This helps reduce service operation costs.

Table 7-1 describes the types of VM live migration.

Table 7-1 Types of VM live migration

Migration Type	Subclass	Description
Manual migration	By destination	On the FusionCompute web client, system maintenance engineers manually relocate one VM to another server.
Automatic migration	VM resource scheduling	The system automatically relocates VMs to other servers in the cluster based on the preset VM scheduling policies.

VM Load Balancing

In the load balancing mode, the system dynamically allocates the load based on the current load status of each physical server node to implement load balance in a cluster.

Snapshot

The snapshot feature enables the FusionCompute to restore a damaged VM using its snapshots.

A snapshot is a set of system files and directories of a VM kept in storage as they were at some time in the past.

- When a VM is faulty, a user can quickly create a VM based on the backed-up VM snapshot.
- A user can also restore a VM to the time the snapshot is created.

VM Isolation

The isolation feature ensures that all VMs running on the same physical server are independent. Therefore, the faulty VM does not affect other VMs.

VM isolation is implemented based on virtualization software. Each VM has independent memory space, network address space, CPU stack register, and disk storage space.

VM OS Fault Detection

If a VM becomes faulty, the system automatically restarts the faulty VM from the physical server where the VM is located or from another physical server, depending on the preset policy. Users can also configure the system to neglect the faults. The system can detect and address internal errors of VM OSs, such as the blue screen of death (BSOD) on Windows VMs and the panic status of Linux VMs.

Black Box

The black box embedded in the FusionCompute collects information about the system. If a fault occurs, the black box collects and stores the last information about the system. This facilitates fault location.

The black box stores the following information:

- Storage kernel logs
- System snapshots
- Screen output information before the system exits
- Diagnosis information from the diagnosis tool

7.3 FusionCompute Architecture Reliability

Management Node HA

Management nodes work in active/standby mode to ensure high availability (HA). If the active node is faulty, the standby node takes over services from the active node, ensuring uninterrupted service processing of management nodes.

The active and standby nodes check the status of each other using the heartbeat messages sent over the management plane. The active node is automatically determined based on the heartbeat messages.

- Only the active node provides services. The standby node only provides basic functions and periodically synchronizes data with the active node.
- If the active node is faulty, the standby node takes over services from the active node and changes to the active state. The original active node changes to the idle state.

  The active node faults include the network interruption, abnormal state, or faulty service process of the active node.

Management Data Backup and Restoration

The system backs up the configuration data and service data periodically on local and remote devices. If the management node service becomes abnormal and cannot be automatically restored, it can be restored using the local data backup rapidly. If a devastating fault occurs and both the active and standby management nodes are faulty at the same time, and they cannot be restored by restarting, they can be restored using the remote data backup rapidly (within 1 hour). With this service, the time for restoration is reduced.

Traffic Control

The traffic control mechanism helps the management node provide concurrent services of high availability without system collapse due to excessive traffic. The traffic control is enabled for the access point, so that excessive load on the front end can be prevented to enhance system stability. To prevent service failures due excessive traffic, this function is also enabled for each key internal process in the system, such as traffic control on image downloading, authentication, VM services (including VM migration, VM high availability, VM creation, hibernation, waking up, and stopping), and operation and maintenance (O&M).

Fault Detection

The system provides the fault detection and alarm functions, and the tool for displaying fault on web browsers. When a cluster is running, users can monitor cluster management and load balancing using a data visualization tool to detect faults, including load balancing problems, abnormal processes, or hardware performance deterioration trend. Users can view historical record to obtain the information about daily, weekly, and even annual hardware resource consumption.

Data Consistency Check

The FusionCompute periodically performs data consistency checks to ensure data consistency.

The FusionCompute periodically checks consistency of all VM data and disk file data on management nodes. If data inconsistency is found, the system generates an audit log. The maintenance engineers can clear the data inconsistency based on the log.

7.4 High Availability

FusionManager is deployed in active/standby mode. The active node provides services, and the standby node synchronizes data with the active node. When the active node fails, services are automatically switched over to the standby node. In the mean time, an alarm is generated.

High Availability with the Active/Standby Mode

FusionManager uses an active/standby dual-node configuration. The active node uses floating IP addresses to provide services. A heartbeat link is configured between the active and standby nodes. When the active node experiences a process or OS failure or the host on which the active node resides experiences an OS failure, an active/standby switchover occurs. The standby node becomes the active node, and all the processes that formerly ran on the active node start on the standby node to provide services, ensuring service continuity.

Data Consistency with the Active/Standby Mode

The FusionManager database uses an active/standby dual-node configuration. The active database provides data access when both databases are working properly. When the data changes, the system synchronizes the changed data over the network between active and standby database processes to the standby database in real time. This prevents data loss when an active/standby switchover occurs.

Alarm

- When an active/standby switchover occurs, an alarm is generated to alert the system administrator.
- In the following scenarios, FusionManager generates a link interruption alarm:
  - The heartbeat link between the active and standby nodes is interrupted.
  - The OS of the standby node is faulty.
  - The OS of the host on which the standby node resides is faulty.

  With the alarm information, the system administrator can restore the network or the backup node, maintaining high availability for the system.
- When a resource on the active or standby node fails, an alarm is generated. With the alarm information, the system administrator can audit the system running status and restore the system.

8 System Specifications

8.1 FusionCompute Technical Specifications

Management Capability

Table 8-1 Management capability

Item	Specifications
Maximum hosts per Virtualization Resource Management (VRM) node	1000
Maximum clusters per VRM node	32
Maximum VMs per VRM node	• Number of started VMs: 10,000 • Number of registered VMs: 30,000
Maximum hosts per cluster	Maximum number of hosts using virtualized data stores: 64

Host Specifications

Table 8-2 Host specifications

Item	Specifications
Maximum logical CPUs per host	768
Maximum memory size per host	16 TB
Maximum VMs per host	1024
Maximum number of logical unit numbers (LUNs) supported by a host	1024
Maximum number of vCPUs supported by a physical server	4096
Maximum number of virtual NICs	300

Item	Specifications
supported by a physical server	
Maximum number of virtual disks supported by a physical server	2048
Maximum number of non-uniform memory access (NUMA) nodes supported by a host	16
Maximum concurrent live migrated VMs per host	8
Maximum number of DVSs supported by a physical server	4

VM Capacity

Table 8-3 VM capacity

Item	Specifications
Maximum VMs per cluster	8000

Snapshot Capacity

Table 8-4 Snapshot capacity

Item	Specifications
Maximum snapshots per VM	32

Network Capacity

Table 8-5 Network capacity

Item	Specifications
Maximum distributed virtual switches (DVSs) per system	50
Maximum hosts per DVS	1000
Maximum virtual switch ports per DVS	10000

VM Specifications

Table 8-6 VM specifications

Item	Specifications
Maximum vCPUs per VM	255
Maximum virtual NICs per VM	16
Maximum disks per VM	60
Maximum memory size per VM	4 TB
Maximum capacity of a single disk per VM (using virtualized storage)	64 TB

8.2 FusionManager Technical Specifications

Technical Specifications for FusionManager

Table 8-7 shows the technical specifications of FusionManager.

Table 8-7 Technical specifications for FusionManager

Item	Specifications
Maximum number of physical servers in the system	4096
Maximum number of VMs in the system	80,000
Maximum number of templates in the system	1600
Maximum number of clusters in the system	256
Maximum number of storage pools in the system	2000
Maximum number of resource zones in the system	256
Maximum number of availability zones (AZs) in the system	256
Maximum number of hypervisors in the system	256
Maximum number of external networks in a resource zone	4096
Maximum number of service VLAN pools in a resource zone	4000
Maximum number of virtual private clouds (VPCs) in a resource zone	2000

Item	Specifications
Maximum number of elastic IP addresses in a resource zone	8000
Maximum number of Virtual eXtensible Local Area Networks (VXLANs) in a resource zone	4000
Maximum number of virtual data centers (VDCs) in the system	10,000
Maximum number of VPCs in a VDC	400
Maximum number of templates in a VDC	100
Maximum number of direct networks in a VPC	200
Maximum number of routed networks in a VPC	200
Maximum number of internal networks in a VPC	200
Maximum number of security groups in the system	80,000
Maximum number of application templates in the system	10,000
Maximum number of application instances in the system	10,000
Maximum number of elastic scaling groups in an application instance	10
Maximum size of a software package	4 GB
Maximum number of VMs that can be monitored by application instance monitoring	1000
Total number of saved alarms	200,000
Number of users	10,500
Maximum number of operation logs	100,000

8.3 Compatibility

8.3.1 Hardware Compatibility

You can query the servers, network devices, storage devices, and firewalls supported by FusionManager using Compatibility check assistant.

8.3.2 Virtualization Compatibility

8.3.2.1 Compatibility with Huawei Virtualization Systems

FusionManager supports Huawei FusionCompute virtualization systems. It can centrally manage dispersed virtualization systems.

FusionManager supports the following FusionCompute versions:

- FusionCompute V100R005C00
- FusionCompute V100R005C10
- FusionCompute V100R005C10U1
- FusionCompute V100R006C00
- FusionCompute V100R006C00U1
- FusionCompute V100R006C10
- FusionCompute V100R006C10SPC100
- FusionCompute V100R006C10SPC101

8.3.2.2 Compatibility with Third-Party Virtualization Systems

FusionManager can manage the VMware virtualization system.

Table 8-8 lists the virtualization system versions supported by FusionManager.

Table 8-8 Compatible third-party virtualization systems

Vendor	System Version
VMware	vCenter Server 5.1.0
VMware	vCenter Server 5.0.0
VMware	vCenter Server 5.5.0
VMware	vCenter Server 6.0.0

8.3.3 Supported OSs

You can query the VM operating systems (OSs) supported by FusionSphere using Compatibility check assistant.

For details about the OSs that support the FusionCompute RDM function, see the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*.

NOTE

To obtain the *FusionSphere SIA Huawei Guest OS Compatibility Guide (KVM Enterprise Virtualization)*, visit the following websites:

- For enterprises: Visit http://support.huawei.com/enterprise, search for the document by name, and download the document of the required version.
- For carriers: Visit http://support.huawei.com, search for the document by name, and download the document of the required version.

9 Appendix

9.1 Deployment Rules

The specifications of Virtual Resource Management (VRM) nodes vary depending on the number of VMs managed by the VRM nodes.

Specifications Requirements for VRM Nodes

Table 9-1 lists the specifications requirements for VRM nodes.

Table 9-1 Specifications requirements for VRM nodes

System Capacity	VRM VM Specifications
1000 VMs or 50 physical servers	• Number of vCPUs ≥ 4 • Memory size ≥ 5 GB • Disk size ≥ 140 GB
3000 VMs or 100 physical servers **NOTE** It is recommended that VRM nodes are deployed on physical servers to support such specifications.	• Number of vCPUs ≥ 8 • Memory size ≥ 8 GB • Disk size ≥ 140 GB
5000 VMs or 200 physical servers **NOTE** It is recommended that VRM nodes are deployed on physical servers to support such specifications.	• Number of vCPUs ≥ 12 • Memory size ≥ 16 GB • Disk size ≥ 140 GB
10,000 VMs or 1000 physical servers **NOTE** The capacity is supported only when the VRM nodes are deployed on physical servers to support such specifications.	• Number of vCPUs ≥ 30 • Memory size ≥ 40 GB • Disk size ≥ 140 GB

NOTE
In physical deployment scenarios, vCPUs in the table refer to the number of hyper threads.

The required RAID configuration varies depending on the number of hosts and VMs in the system. However, the actual number of hosts and VMs may vary from the following typical configurations. You need to configure RAID to match the higher configuration.

- 50 hosts or 1000 VMs: RAID 10 consisting of six SAS disks or RAID 1 consisting of two SSDs.
- 100 hosts or 3000 VMs: RAID 10 consisting of four SAS disks or RAID 1 consisting of two SSDs.
- 200 hosts or 5000 VMs: RAID 10 consisting of four SAS disks or RAID 1 consisting of two SSDs.
- 1000 hosts or 10,000 VMs: RAID 10 consisting of ten SAS disks or RAID 1 consisting of two SSDs.

All the above mentioned disks must be 15000-RPM Serial Attached SCSI (SAS) disks.

Configuration Requirements for the FusionManager VM

Table 9-2 lists the configuration requirements for the FusionManager VM deployed in All-in-One mode.

Table 9-2 FusionManager VM requirements

Item	Service VM Scale (< 200)	Service VM Scale (200 to 1000)	Service VM Scale (1000 to 10,000)	Service VM Scale (10,000 to 80,000)
VM host	Select the planned host from FusionCompute and select **Bind to the selected host** and **Always reserve resources for the VM**. To create the FusionManager VM on the VMware hypervisor, you do not need to bind the VM to a specific host. • If FusionCompute of the earlier version has been deployed in the running environment, FusionManager VMs are preferentially created on FusionCompute of the earlier version. • If FusionCompute of the earlier version has not been deployed in the running environment, FusionManager VMs are created on FusionCompute 6.3.0 or later versions.			
VM name	Planned FusionManager VM name			
OS	Linux			
OS version	Novell SUSE Linux Enterprise Server11 SP3 64-bit			
Number of vCPUs	4	6	12	16
Memory size (GB)	6	18	24	30
Disk	One 80 GB disk	One 120 GB disk	One 120 GB disk	One 150 GB disk
	When the in-use hypervisor is FusionCompute, non-virtualized local hard disks and SAN storage are recommended. Between these two types of storage, preferentially use non-virtualized local hard disks.set			

Item	Service VM Scale (< 200)	Service VM Scale (200 to 1000)	Service VM Scale (1000 to 10,000)	Service VM Scale (10,000 to 80,000)
	the disk configuration mode to **Common** or **Thick provisioning lazy zeroed** and select **Independent & Persistent**. When the in-use hypervisor is VMware vCenter, VMFS disks are recommended. Set the disk configuration mode to **Thick Provision Lazy Zeroed** or **Thick Provision Eager Zeroed**.			
Number of NICs	1			
Quality of service (QoS) settings	Configure QoS settings for the FusionCompute system as follows: • In the **Hardware** area of the **Configure VM** page, set **Reserved (MHz)** of CPU to a value that equals to the number of vCPUs multiplied by the dominant frequency of the VM host. • In the **Hardware** area of the **Configure VM** page, set **Reserved (MB)** of Memory to the planned memory reservation size. If the in-use hypervisor does not support reserved CPU and memory settings, skip the QoS settings.			
High availability (HA)	When creating a FusionManager VM in the FusionCompute hypervisor, enable the VM HA. When creating a FusionManager VM in other hypervisors, enable the HA function in the in-use hypervisor. Take note of the FusionManager VM name to facilitate subsequent VM maintenance.			
Blue screen of death (BSOD) policy	Select **No processing**.			
Clock synchronization policy	Configure a precise clock source for the FusionManager system. Therefore, to create a FusionManager VM in the FusionCompute hypervisor, leave **Sync time with host** unchecked. To create a FusionManager VM in other hypervisors, configure a precise external clock source for the FusionManager VM.			
Network settings	Specify the distributed virtual switch (DVS) and the port group connected to the management plane.			
Memory swap partition	Select **Enable memory swapping for the VM** (this operation is required only when FusionManager is deployed in FusionCompute of the earlier version).			
Other parameters	Accept the default values of other parameters for the FusionManager VM to be created in the FusionCompute hypervisor. To create a FusionManager VM in other hypervisors, set the required parameters based on the configuration documentation of the in-use hypervisor.			

9.2 Technical Support

Huawei Technologies Co., Ltd. provides customers with comprehensive technical support and service. For any assistance during the product usage or maintenance, visit the following websites to get access to self-service:

- Cloud computing and Big Data information service platform: http://support-it.huawei.com/cloud
- Intelligent robot: http://support.huawei.com/eirobot/
- Interactive community: http://forum.huawei.com/enterprise/zh/forum.html

If the issue cannot be solved using the preceding methods, contact our local office or company headquarters.

- Call the service hotline of your local Huawei office.
- Give your feedback using the information provided on the **Contact Us** page at the technical support websites:
  - For enterprise users: http://support.huawei.com/enterprise
  - For telecom carrier users: http://support.huawei.com