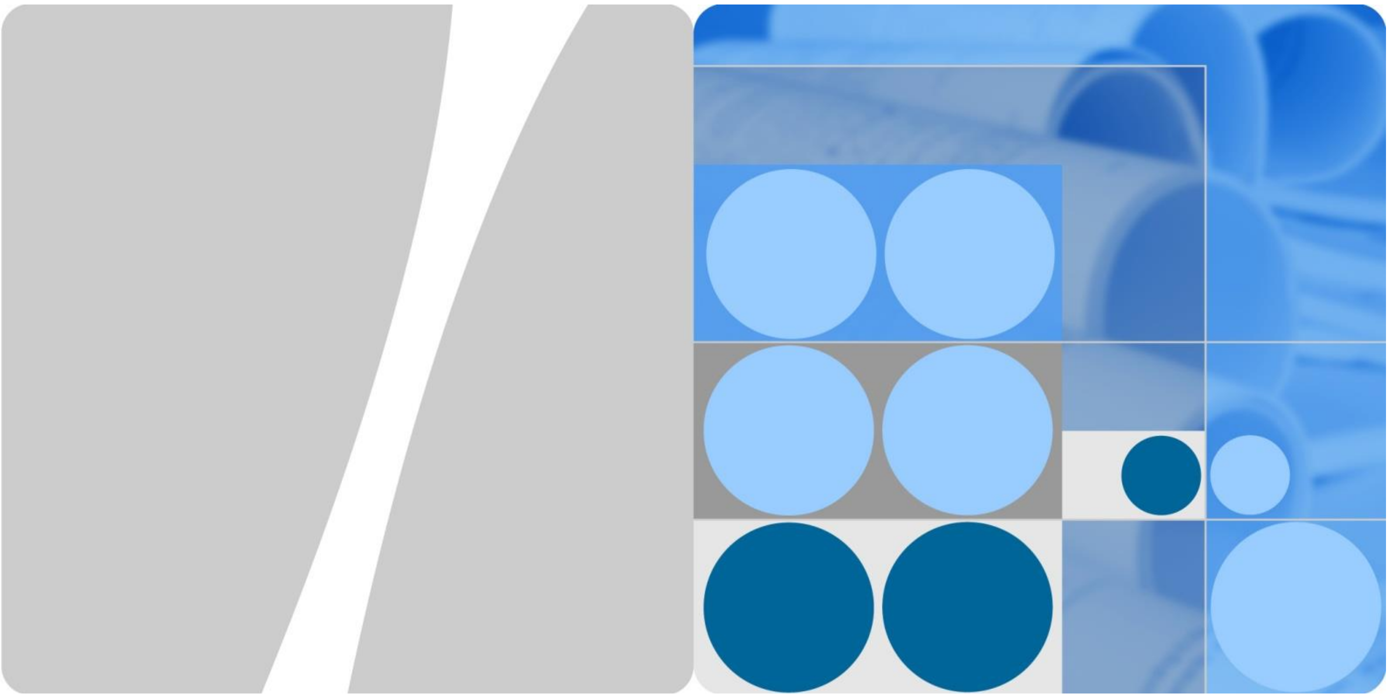

Huawei FusionSphere 6.3.1

Technical White Paper (Server Virtualization)

Issue 1.0

Date 2018-09-30

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

Contents

1 Executive Summary

Driven by agile IT, cloud computing comes into being as a combination of technologies, rather than a new technology. Before cloud computing is introduced, deploying a set of services for an enterprise requires network planning, capacity planning, device modeling, ordering, payment, delivery, shipment, installation, deployment, and commissioning. The entire process takes several weeks or even months for large-scale projects. After cloud computing is introduced, the entire process is shortened to less than one hour.

According to Moore's law, chip performance will be doubled every 18 months. The reverse Moore's law indicates that the vendors that cannot follow Moore's law will be eliminated. The competition in the IT industry is fierce. Cloud computing can help enterprises to improve the provisioning efficiency of the IT infrastructure, whereas those enterprises that do not apply cloud computing are lagged behind in product or service expansion.

In addition, cloud computing can provide a series of other benefits, for example, increasing the reuse ratio, reducing power consumption, and reducing maintenance costs. These benefits become tiny when compared with its efficiency improvement capability to cope with the reverse Moore's law.

Cloud computing virtualizes different infrastructures into standard service components and then automatically combines these service components to meet user requirements. Cloud puts emphasis on virtualization, standardization, and automation.

FusionSphere is a sophisticated infrastructure as a service (IaaS) cloud computing solution. Besides virtualization, standardization, and automation features, the FusionSphere solution, backed by Huawei's twenty years' experience in telecom products, provides you with openness, security, and reliability.

This document provides visibility into the technologies used in the FusionSphere solution. Upon completion of this document, you will be able to understand:

- Virtualization, standardization, and automation of FusionSphere
- Openness, security, and reliability of FusionSphere
- Components and technologies involved by FusionSphere
- Technology selection to meet your service requirements

This document comprises the following chapters:

Chapter 1: provides a brief introduction to cloud computing and the cloud platform and provides guidance on this document.

Chapter 2: provides a brief introduction to each component, and describes how the cloud concepts are reflected in these components.

Chapter 3: describes the standardization capabilities of FusionSphere components.

Chapter 4: describes the automation capabilities of FusionSphere.

Chapter 5: describes the key technologies used for standardization and automation.

Chapter 6: describes Huawei's efforts in providing openness, security, and reliability of FusionSphere.

Chapter 7: provides a conclusion and the path to the FusionSphere Foundation Edition.

2 FusionSphere Overview

2.1 Abstract of This Chapter

Upon completion of this chapter, you will be able to understand:

- Components of FusionSphere
- Technology map from the user's perspective and the value of each technology
- How to construct your services using these technologies

2.2 Components and Functions

FusionSphere virtualizes hardware resources using the virtualization software deployed on physical servers so that one physical server can function as multiple virtual servers. FusionSphere can maximize resource utilization by centralizing existing VM workloads on some servers and thereby releasing more servers to carry new applications and solutions. In addition, FusionSphere can provide virtual data center technologies for enterprises and carriers and support cross-site disaster recovery (DR) capabilities.

Figure 2-1 shows the components of the FusionSphere solution.

Figure 2-1 FusionSphere components

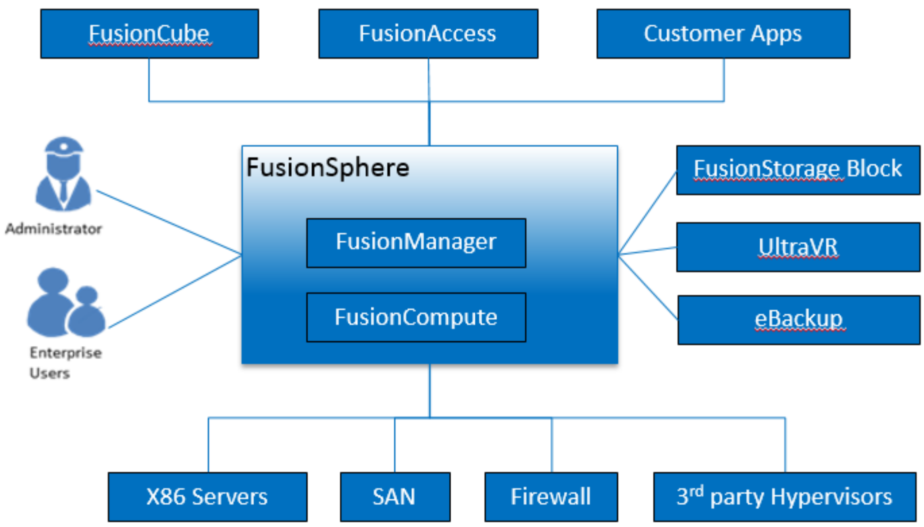

Components and their functions are as follows:

- FusionCompute: virtualizes x86 physical servers and storage area network (SAN) devices and supports software-defined networking.
- FusionManager: exercises FusionCompute capabilities and integrates the automatic management capabilities of firewalls and load balancers to provide a virtual data center management solution for enterprises and carriers.
- FusionStorage Block: provides storage virtualization capabilities based on x86 servers.
- UltraVR: supports cross-site DR.
- eBackup: suppors backup for virtualization machine

2.3 FusionSphere Technology Map

The following table provides a technology map from the standpoint of user experience. After reading this document, you will be familiarized with the technologies involved.

As mentioned before, cloud computing virtualizes different infrastructures into standard service components and then automatically combines these service components to meet user requirements. This technology map organizes all technologies involved in terms of virtualization, standardization, and automation. FusionSphere is an open system. Therefore, this technology map also describes technologies in terms of openness.

Table 2-1 lists the technologies involved in FusionSphere.

Table 2-1 FusionSphere technology map

Component	Virtualization	Openness	Standardization	Automation
FusionCompute	Computing virtualization Storage virtualization Network virtualization Clustering Software load balancer Software-defined networking	General x86 server General SAN storage	VM Virtual storage Virtual network plane Virtual gateway Virtual load balancer	Automatic provisioning of standard components Component service quality assurance

Component	Virtualization	Openness	Standardization	Automation
FusionManager	Firewall virtualization Load balancer virtualization Application virtualization	Heterogeneous firewall F5 load balancer Heterogeneous computing resource virtualization Cloud storage	Virtual firewall Virtual load balancer Virtual data center (VDC) Virtual private cloud (VPC) Virtual application	
FusionStorage Block	Storage virtualization	x86 server	Virtual storage	

These technologies used in the FusionSphere solution will be introduced in the following chapters in detail.

3 Standardization Capabilities of the Cloud Platform

Cloud computing virtualizes different infrastructures into standard service components and then automatically combines these service components to meet user requirements.

Toy blocks are regarded as delicate toys because they are stably shaped, easy-to-replace, freely combinable, and recyclable. Since cloud is born for agile IT, the FusionSphere platform provides a series of standardization capabilities that help users to construct their own systems just like building toy blocks.

This chapter describes the standardization capabilities of the cloud platform. Through this chapter, you can understand what basic services the cloud platform provides for you. Moreover, you can better understand these components and therefore use them to construct your own system.

The following sections describe the FusionSphere standardization capabilities by product component.

3.1 Standardization Capabilities of FusionCompute

3.1.1 VM

The x86 virtualization technology virtualizes x86 servers using virtualization software and provides standard VMs for end users. These VMs possess serialized hardware configurations and apply the same driver, just like a series of products from the same vendor.

FusionCompute is such a virtualization system that can virtualize an x86 server into multiple VMs. Then end users can install software, attach disks, modify configurations, and adjust networks on the VMs.

For end users, VMs can be more rapidly provisioned than physical machines, and their configurations and networking can be more easily modified. For maintenance personnel, the maintenance cost of VMs is significantly reduced because hardware is reused by VMs and the cloud platform supports automatic maintenance. For system administrators, the resource quantity and utilization trend are visible enough.

3.1.2 Virtual Storage

FusionCompute centrally manages the virtual storage resources provided by SAN devices, FusionStorage Block, and local storage of computing nodes and allocates the storage resources to VMs as virtual disks.

End users can use these virtual disks on VMs like using local disks on x86 servers, for example, they can format these virtual disks, read data from or write data into them, install file systems, and install operating systems (OSs). Moreover, virtual storage supports the snapshot function and resizing, which cannot be implemented on physical hard disks.

Administrators only need to manage the SAN device, instead of managing specific disks. Because SAN devices are reliable, the workloads for replacing hard disks are significantly decreased for administrators. In addition, virtual storage supports various features that are not supported by physical disks, such as thin provisioning, QoS, and migration. Therefore, virtual storage has distinct cost advantages over physical disks.

3.1.3 Virtual Network

FusionCompute has the function of distributed virtual switches (DVSs) and provides independent network planes for VMs. Similar to physical switches, DVSs provide different network planes isolated by VLANs. This technology has the following features:

- VMs on the same host but in different VLANs cannot communicate with one another.
- VMs on the same host and in the same VLAN can communicate with one another at layer 2. In this case, VMs use swap memory to exchange data. Therefore, the data exchange between VMs requires less network bandwidth.
- VMs on different hosts but in the same VLAN can communicate with one another through an external switch.

With this technology, you can regard each VLAN as an independent network plane. Therefore, service isolation is implemented by assigning VLANs to different VMs.

3.2 Standardization Capabilities of FusionManager

3.2.1 Virtual Data Center

A virtual data center (VDC) provides users with the capability to manage their virtual resources. All virtual resources belonging to users or organizations are included in one or more VDCs. Different organizations can use resources in different VDCs. An organization can only view its own VDC.

- A VDC manages the following resources and configuration information:
- Organizations and users in organizations: Multiple users in the same organization can be configured to share resources in the same VDC.
- Resource quotas: Resource quotas can be configured to limit the quantity of resources, such as computing capabilities, memory, network bandwidths, and VLANs, available to an organization.
- Organization assets: All assets belong to an organization.
- Networking: The network topology displays relationships of different assets.
- Applications: All applications belong to an organization.

VDCs allow users to effectively manage their own assets.

3.2.2 Virtual Private Cloud

A Virtual Private Cloud (VPC) displays the asset networking relationships for users.

Figure 3-1 shows the general asset networking.

Figure 3-1 General asset networking

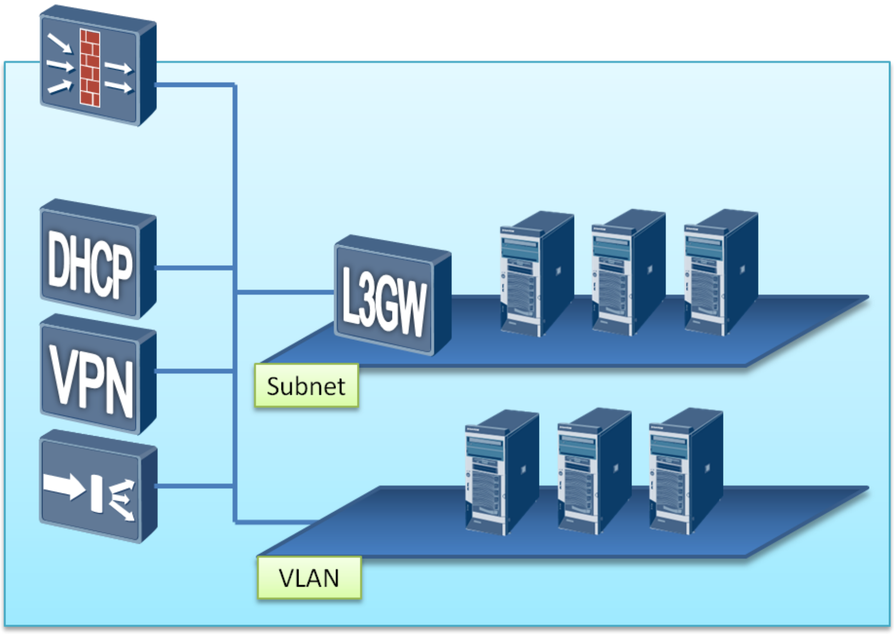

A VPC can provide the following services:

- Virtual firewall (vFW) service: protects a VPC against threats from other networks.
- Virtual gateway service: helps to formulate one or multiple network planes in a VPC.
- Virtual DHCP service: assigns IP addresses to VMs.
- Virtual private network (VPN) service: allows networks of different cloud systems to communicate.
- Elastic IP address service: enables VMs in a VPC to connect to the Internet.
- Virtual load balancers: provides load balancing services.

A VPC enables administrators to control network topologies, facilitating end user management.

The VPC enables agile IT, provides the capacity to quickly deploy networks, and simplifies network resource management. The VPC relieves administrators from recording VLAN and IP address allocation information and avoids application configuration conflicts on a firewall or switch.

3.2.3 Operation Capabilities

Services are resource templates for which users can apply, including VDCs, Elastic Cloud Servers (ECSs), application instances, and Elastic Volume Service (EVS).

The domain service administrator can define service catalogs as required. The following items can be defined:

- Service name, description, and icon
- Parameters that can be specified when users apply for services, for example, VM specification parameters.
- Service parameters that can be specified by administrators during service application approval, for example, a static IP address for a VM.
- Service parameters that are locked and cannot be specified when users apply for services (for example, locking the VM OS as Windows 7)
- Service approval policy, that is, whether a service needs to be approved

3.3 Standardization Capabilities of FusionStorage Block

3.3.1 Virtual Storage

FusionStorage Block is a distributed storage system. It virtualizes the local hard disks of x86 servers to provide virtual storage resources for VMs.

The virtual disks provided by FusionStorage Block for end users are actually distributed on different x86 servers, which significantly improves data read and write efficiency. Furthermore, the distributed system architecture provides easy capacity expansion, and faults can be easily rectified.

FusionStorage Block serves as the storage component of FusionSphere. Besides, FusionSphere can virtualize external SAN devices to provide virtual storage resources.

3.4 Standardization Capabilities of UltraVR

3.4.1 Virtual Host Disaster Recovery

UltraVR provides one remote replication disaster recovery (DR) plans. The plan can:

- Use the remote replication function of the Huawei storage device to copy the VM data stored at the production site to the DR site.
- Use UltraVR software to copy the VM configuration data from the production site to the DR site and to manage the system DR plan.
- Start the DR plan to restore the VMs using the VM configuration data and user data backed up on the DR site if a disaster occurs at the production site and causes VM failures.

The RPO of this plan is the data replication interval between arrays, and the RTO of this plan is the duration for the entire system switchover and VM startup, which is less than 4 hours in the typical scenario where 3000 VMs are deployed.

UltraVR provides the following functions:

- Centralized management of recovery plans
  - Supports recovery plan creation and management.
  - Automatically discovers and displays the VMs that are protected by the storage array.

  - Maps VMs to the related resources (such as resource clusters, storage devices, or networks) at the DR site.
- Automatic switchover
  - Supports one-click startup of the recovery plan.
  - Automatically changes the remote replication standby LUNs to the active LUNs for DR implementation.
  - Monitors site availability and sends an alarm when a site is faulty.
  - Manages and monitors the recovery plan process.
- DR tests without service interruption (DR drilling)
  - Automatically performs recovery tests.
  - Performs recovery tests using storage snapshots to prevent loss of replicated data.
  - Automatically cleans the testing environment after a test.
  - Supports test results query and export.
- Planned migration
  - Automatically performs scheduled migration and stops the VMs to be protected at the production site.
  - Before migration, make sure that VM data is completely replicated to ensure application consistency.
  - Perform data synchronization to forcibly replicate the data on the stopped VMs to the DR site.

3.5 Other Standardization Capabilities

3.5.1 Backup

FusionSphere can integrate the VM backup capability provided by eBackup.

eBackup can periodically back up VMs using the VM snapshot capability of FusionCompute based on the specific backup policies configured for VMs.

The eBackup VM backup solution has the following characteristics:

- No backup agent needs to be installed on the VM to be backed up.
- VM data can be backed up regardless of whether the VM is in the running or stopped state.
- Backup and restoration can be performed for VMs using different storage resources, such as FusionStorage Block or virtualized storage resources.
- VM data can be backed up to various storage devices, including virtual disks attached to the eBackup-installed VM and external storage devices using the NFS or CIFS, such as NAS devices. The virtual disks can be local disks of the VM or be provided by SAN or NAS storage devices.

Figure 3-2 Working principle of eBackup

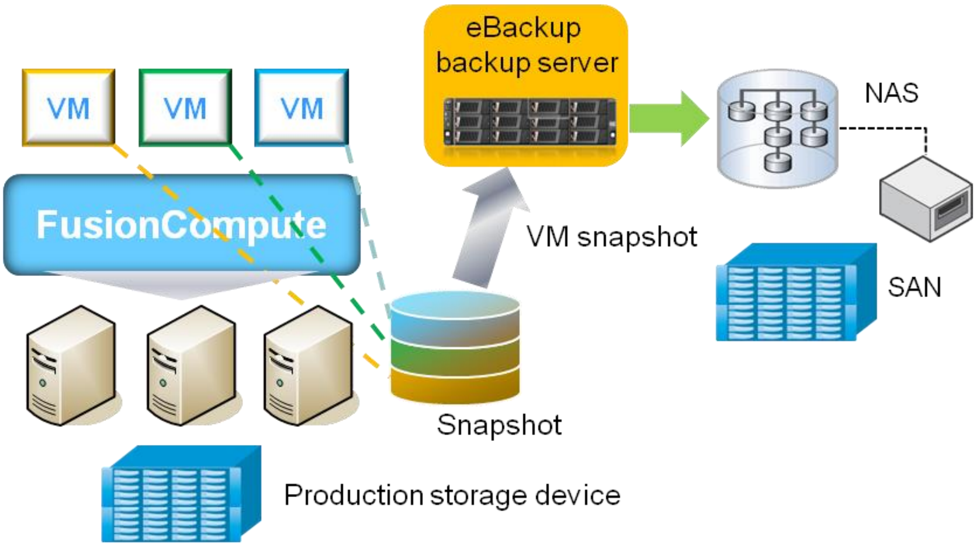

4 Automation Capabilities

4.1 Overview

The core cloud concepts are virtualization, standardization, and automation. Chapter 3 describes the standardization components of FusionSphere, and this chapter describes its automation capabilities.

FusionSphere provides various automation functions to simplify routine management.

According to the enhanced Telecom Operations Map (eTOM), pay attention to the following key points in routine management:

- Policy: How can we provision resources to meet service requirements and how can we adjust resource allocation policies among different services?
- Infrastructure: What infrastructures do we have, what are their usage and states, and how do these infrastructures provide support for products based on specified policies?
- Product/Service: What services do we provide externally and what qualities do these services have? How do the infrastructures provide support for these products and services?

- Provisioning: How is a product/service instance provisioned?
- Assurance: How does the system ensure the quality of a provisioned product or service instance?
- Measurement: What is the amount of the product or service instance to be used, what is the user preference, and how can we adjust resources and policies to better provision the product or service?

For FusionSphere, the products or services to be provisioned are the standardization components described in chapter 3. The automated management capabilities of FusionSphere allow system administrators to easily manage these products or services and end users to conveniently use these products or services.

4.2 Provisioning and Usage of Standardization Components

The system supports the complete life cycle management of all standardization components, including provisioning, creation, application, adjustment, configuration, and reclamation. You can perform these management operations on the portal or using web service interfaces.

4.3 Quality Assurance

FusionSphere provides various active and passive management mechanisms for administrators to ensure service quality.

Active management allows the management system to automatically ensure the service quality based on the policies configured by administrators. Passive management provides related information for administrators to manage components.

The following sections describe the active and passive management capabilities of FusionSphere.

4.3.1 Active Management

The FusionSphere system contains various active management functions, most of which are invisible to administrators, such as VM breakdown detection, hibernation detection, attack detection of the virtual network, and fault detection of the management system. This section describes only common active management functions.

VM HA

FusionSphere supports VM HA to enable affected VMs to be automatically restarted in a specific resource pool in the event of a failure, thereby improving VM availability.

After VM HA is enabled for a cluster, administrators can enable HA for a VM during VM creation.

The FusionCompute system periodically detects VM statuses. If FusionCompute detects that a VM is faulty due to a physical server or software failure, VM HA allows the VM to be automatically restarted on another production physical server with spare capacity.

During the restart, the VM reloads the OS. Data that is not saved to the hard disk before the VM HA is lost.

If this function is not enabled for a VM, the VM stops working after a fault occurs. In this case, administrators need to manually restart the VM.

VM DRS

Dynamic Resource Scheduling (DRS) enables the system to flexibly schedule resources and achieve load balancing using the intelligent algorithm based on the system load, thereby providing enhanced user experience.

DRS policies define scheduling thresholds and periods during which the policies take effect for a cluster. During the set period, if the CPU or memory load on a host exceeds the scheduling threshold, the system migrates some VMs to other light-load hosts to balance the loads.

VM QoS

VM QoS allows users to control the quality of service (QoS) for computing resources (such as CPUs and memory) in a flexible manner.

This feature implements measurement of computing capabilities and limits the computing capabilities of VMs to a specific range. With this feature enabled, VMs requiring different computing capabilities do not affect each other. This feature also optimizes allocation of computing resources and resource reuse, thereby reducing the cost and improving user satisfaction.

VM QoS comprises CPU QoS and memory QoS.

- CPU QoS

  CPU QoS ensures optimal allocation of computing resources for VMs and prevents resource contention between VMs due to different service requirements. Therefore, CPU QoS can effectively increase resource utilization and reduce costs.

  CPU QoS values can be set during VM creation based on the planned VM services. CPU QoS determines VM computing capabilities. The system ensures the VM CPU QoS by setting the minimum computing capability and the computing capability upper limit for VMs.

- 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, thereby increasing memory utilization. In this case, memory QoS is used to reserve the minimum memory for reliable running of VMs and thereby ensures optimal use of memory.

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

Automatic Application Instance Scaling

Application Instance Scaling allows FusionSphere to automatically schedule the number of VMs in a scaling group for an application instance based on the configured policy.

This feature provides the following benefits for users by dynamically scheduling computing resources based on service loads:

- Improves user experience and enhances service availability.
- Maximizes resource utilization.

- Prevents services from being affected during peak hours.

Users can set the following types of resource policies for application instances:

- Intra-group scaling policies

  Intra-group scaling policies apply to a single application instance.

  During application instance operations, the system dynamically adjusts the resources for the application instance based on the current loads.

  When detecting that the resources for an application instance are insufficient, the system adds VMs and installs application software on the VMs to ensure the application instance functions optimally.

  When the system detects that resources for an application instance are underused, it automatically removes VMs from the application instance to release resources and conserve energy.
- Inter-group resource reclaiming policies

  When system resources are insufficient, the system forcibly releases resources used by lower-priority application instances to ensure higher-priority application instances function properly based on the inter-group resource recycling policy.
- Time-based scheduling policies

  Time-based scheduling policies allow resources to be assigned to different application instances during 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.

Automatic VM Backup

VM Backup uses Huawei eBackup backup software, the FusionCompute snapshot function, and the Changed Block Tracking (CBT) function to back up VM data.

When a VM becomes faulty or its data is lost, the VM can be restored using its data backups stored on local virtual disks or the shared NAS storage device connected to eBackup.

The eBackup enables automatic periodic VM backup based on the configured backup policies. Details are as follows:

- Different backup policies can be configured for different VMs or VM groups. The backup system supports up to 200 backup policies.
- Different backup intervals and backup windows can be configured for full backup, incremental backup, and differential backup. For example, the backup system can be configured to perform a full backup weekly and an incremental backup daily. Also, the system can be configured to perform only one full backup and then perform incremental backups at specified intervals.
- The data backup retention duration and the automatic expired data deletion functions can be configured.
- Backup policy priorities can be configured.

4.3.2 Passive Management

Passive management allows administrators to manage components by providing the administrators with reports, statistics, alarms, events, and other related information.

4.4 Summary

The ultimate objective of management is to let users "set it and forget it". To achieve this objective, FusionSphere will incorporate various policy-based active management capabilities in future.

5 Key Standardization and Automation Technologies

5.1 Computing Virtualization

The Unified Virtualization Platform (UVP) is used to enable the computing virtualization of FusionSphere. The UVP uses the bare-metal architecture and builds a software layer between underlying hardware and upper-layer OSs. By abstracting x86 servers and transforming physical resources in the servers, including the CPUs, memory, and I/O, into logical resources that can be uniformly managed, scheduled, and allocated, the UVP constructs VM execution environments that can run simultaneously and are isolated from one another. It also provides various HA features, such as live migration and DRS, to implement dynamic resource allocation for application instances. Therefore, the UVP helps to improve resource utilization, reduce the OPEX, speed up service response, and enhance service agility.

The UVP can virtualize computing, storage, and network resources and support various advanced features, including HA, security, and easy management. It is widely applied to various scenarios, such as server consolidation, virtual desktop, scientific computing, and web application.

The UVP is a key technical platform in Huawei FusionCloud data center virtualization solution. It applies to key application domains for enterprise users. It uses the bare-metal x86 architecture and Kvm technologies and functions as an enhanced server virtualization platform.

5.2 Storage Virtualization

FusionSphere provides host-based storage virtualization, hiding the complexity of storage device types and bypassing performance bottlenecks. Storage virtualization is implemented by virtualizing physical storage resources into logical resources. Resources of different storage devices are integrated and provided in a unified manner. Based on the storage virtualization technology, FusionSphere provides multiple storage services to improve storage utilization, reliability, maintainability, and user experience.

FusionSphere hides peculiarities of different storage devices, including IP SAN, Fiber Channel (FC) SAN, and network-attached storage (NAS) devices, and local disks and offers file-level service operations based on file systems. FusionSphere provides comprehensive storage-related functions, such as thin provisioning, incremental snapshots, storage cold and live migration, linked clones, and VM disk capacity expansion.

In addition, FusionSphere supports data store expansion, which allows one data store to manage multiple physical LUNs, thereby effectively improving the data store expansibility.

5.3 DVS

FusionSphere supports DVSs running on computing nodes. Complete virtual switching is implemented on server CPUs. A virtual port is assigned to a virtual NIC of a VM for virtual switching, and physical NICs of a server function as virtual switching uplink ports.

A DVS has the following features:

- High performance in internal server communication. A DVS can implement layer 2 forwarding of packets among VMs on a server using software.
- Moderate performance in cross-server communication. For a server CPU, the cross-server communication requests must be forwarded by a physical switch. Therefore, the virtual switching performance is inferior to a physical switch.
- Flexible scalability. Unlike physical switches that use layer 3 chips, servers use only software to implement virtual switching, which provides flexible and rapid scalability to better extend cloud computing networks.
- Large size of server memory. The layer 2 switching capability and access control list (ACL) capability of a server are much greater than those of a physical switch.

In addition, FusionSphere supports virtual switching capability provided by single-root I/O virtualization (SR-IOV) NICs.

A DVS provides virtual network security capability, preventing VM IP address spoofing.

5.4 Backup

Depending on different service scenarios, application types, numbers of users, IT infrastructure conditions, and investment budgets, FusionSphere delivers the following backup plans:

- Third-party backup server agent-based backup solution (traditional backup)
- eBackup server agentless backup solution (VM-based backup)
- IP SAN-based backup solution (storage-based backup)

Suggestions on selecting a backup solution are as follows:

- Use the eBackup server agent-less backup solution in most scenarios.
- Use the IP SAN-based backup solution if services must be rapidly restored after the storage at the production site is faulty.
- Select the third-party server agent-based backup solution if users' historical data needs to be restored or archived.
- Use Huawei HDP3500E devices if the third-party server agent-based backup solution is selected. To make a configuration quotation for HDP3500E devices, contact Huawei storage engineers.

5.5 Disaster Recovery

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

Depending on different application scenarios, application types, user scales, IT infrastructure conditions, and investment budgets, FusionSphere delivers the following DR plans:

- Storage replica-based DR plan of UltraVR
- Metropolitan active-active DR
- Application DR based on the continuous data protector (CDP)

Figure 5-1 FusionSphere DR solution

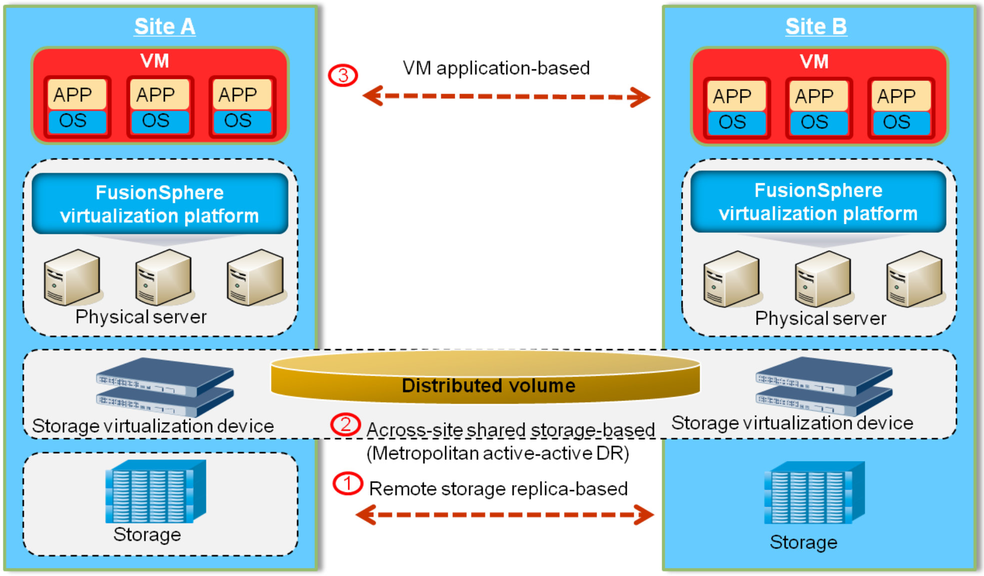

Suggestions on DR plan selection:

- For common scenarios, UltraVR plan is recommended.
- If the RPO is required to be zero and RTO must be short in case any disaster occurs, the metropolitan active-active DR plan is recommended.
- If both VMs and physical servers are deployed at the production and DR sites, the CDP-based application DR plan is recommended.

5.6 Virtual Data Center

The VDC provides the following functions:

- Network device virtualization: The VDC technology virtualizes network devices, such as firewalls, load balancers, layer 2 networks, layer 3 gateways, DHCP servers, and VPN devices, and provides virtual resources for end users. Network administrators can flexibly deploy networks using the virtualized network resources to meet their service requirements.

- Isolation: The VDC technology enables different network resources, including VLANs, IP address segments, and virtual firewalls, to be allocated to different departments to isolate networks of the departments.
- Virtualization decouples networks of different departments, organizations, and services. Therefore, network administrators only need to manage several simple and independent network systems instead of multiple deeply integrated network systems sharing only one set of network devices.
- Tracking of resource allocation: All allocated virtual resources, such as VLANs and IP addresses, will be recorded for further management and query.
- Controllability on key resources: Scarce resources, such as public IP addresses, public network bandwidth, and VPNs, can be properly managed and then allocated based on service requirements of departments.
- Self-service: Self-service portals are provided. Customer network administrators can maintain their networks only after simple training.

In the virtual data center solution, different departments or services can be isolated using VPCs. One VPC functions as a secure network and consists of the following elements:

- A vFW
- Multiple network planes
- The DHCP service, VPN, load balancer, and layer 3 gateway deployed in the network

Figure 5-2 Networking of a VPC

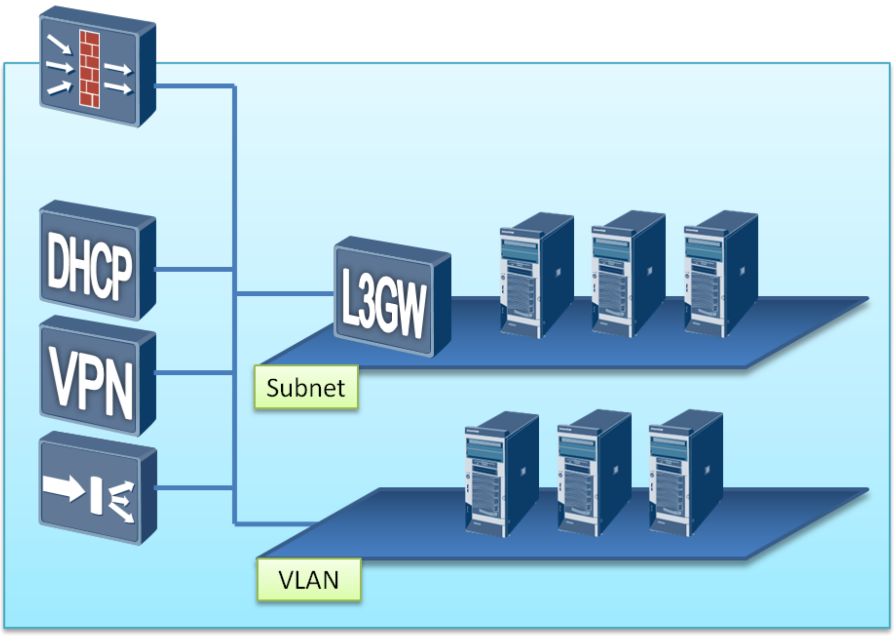

The VDC solution provides security groups, elastic IP addresses, bandwidth control, Super VLANs, floating IP addresses, public-private IP address binding, and the capabilities to access external networks. The VDC solution also provides the metering function and QoS and quota management to meet carriers' requirements.

FusionSphere also supports the multi-DC management feature to consolidate resources of multiple data centers and provide unified computing resources for the upper-layer applications. In addition, the resource SLA feature enables FusionSphere to allocate computing, storage, and network resources from proper resource pools to meet the requirements of different application instances or tenants.

6 Openness, Security, and Reliability

6.1 Openness

In addition to Huawei-owned hardware and software, FusionSphere is compatible with extensive devices from other vendors. This section describes the hardware and software supported by FusionSphere.

6.1.1 Open APIs

Open APIs allow users to connect their own business operation systems to the cloud platform. After a user's business operation system is connected to the cloud platform, the user can provision cloud services and manage end users on their own business operation system. The open APIs provided by FusionSphere allow users to provision services and manage end users using their own operation systems. The users can add and cancel cloud services using the APIs. The Open APIs are developed based on the Simple Object Access Protocol (SOAP) Representational State Transfer (REST) protocols.

6.1.2 Infrastructure

So far, it is formally released in the compatibility test report, which are from mainstream providers in the industry. The report is periodically updated. You can obtain the latest version from the sales personnel.

New devices can be added to the FusionSphere system after driver development and verification.

6.1.3 Virtualization Platform

FusionCompute, developed by Huawei, serves as the cloud platform in the FusionSphere solution. FusionSphere also supports vSphere as the underlying virtualization platform.

6.1.4 Cloud Storage Platform

FusionSphere supports the cloud storage platform that can provide web disks or object-based storage (OBS) services for end users.

- Huawei UDS
- China Mobile Big Cloud

6.2 Security

Huawei provides the virtualization platform security solution to face the threats and challenges posed to the cloud computing system. Figure 6-1 shows the structure of the virtualization platform security solution.

Figure 6-1 Structure of the virtualization platform security solution

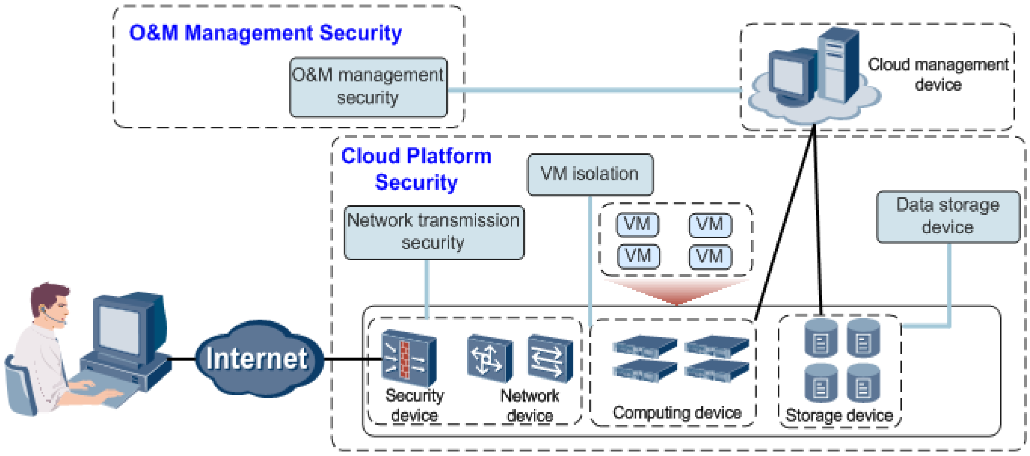

Each layer of the structure is described as follows:

- Cloud Platform Security
  - Data security

    The integrity and confidentiality of user data are ensured using user data isolation, data access control, residual information protection, and data backup.
  - VM isolation

    VMs running on the same physical server are isolated to prevent data theft and malicious attacks. Users can only use VMs to access resources belonging to their own VMs, such as hardware and software resources and data.
  - Network transmission security

    Network plane isolation, firewalls, and encrypted transmission are used to ensure the security of service operation and maintenance.
- O&M Management Security

  Security is ensured from aspects of user accounts, user rights, logs, and data transmission.

In addition, the security of each physical host is ensured by repairing web application loopholes, hardening the operating system (OS) and database, and installing patches and antivirus software.

FusionSphere also supports the VM-Deployed Antivirus feature, which allows the antivirus function to be implemented by a dedicated reinforced secure VM, thereby preventing the antivirus function from consuming resources on other VMs and enhancing the antivirus performance.

6.3 Reliability

FusionSphere is a cloud platform oriented for services. Hardware resources used by FusionSphere are abstracted into virtual resources, and therefore hardware faults are decreased. Virtualization endows the resource pools with the redundancy fault tolerance mechanism, which further improves system reliability. Besides, FusionSphere provides various measures for ensuring system reliability.

6.3.1 Architecture Reliability

Dual network planes: In the FusionSphere solution, all network connections are completed in dual planes, including installing network intelligent cards (NICs), routing cables, and deploying access switches, aggregation switches, and firewalls. This physically ensures reliable communication.

Plane-based communication: The cloud computing system is divided into the management plane, storage plane, and service plane. FusionSphere employs the detached-plane architecture to ensure the reliability and security of various network plane data. Different planes are isolated by VLANs so that the fault of a single plane exerts no impact on other planes.

Active/Standby management nodes: The active and standby management nodes of FusionSphere use the heartbeat detection mechanism of the management plane. The standby node monitors the health status of the active node in real time. Once detecting a fault is detected on the active node, the standby management node takes over services from the active node.

Traffic control: The traffic control mechanism helps the management node provide concurrent services of high availability without system collapse duo to excessive traffic.

Fault detection: The system provides fault detection and alarm reporting functions and the tool for displaying faults on web browsers. When a cluster is running, users can detect cluster management and load balancing using a data visualization tool to detect faults, including load balancing problems, abnormal processes, or hardware performance deterioration trends.

Data consistency check: FusionSphere automatically audits and restores key resource data, periodically audits VMs, and checks volume information to ensure volume data and status consistency. When detecting an exception, FusionSphere automatically generates a log and provides maintenance instructions.

Management data backup and restoration: The system that can connect to a third-party File Transfer Protocol (FTP) server supports periodic local and remote backup of configuration and service data on management nodes. If the management node service becomes abnormal and cannot be automatically restored, it can be rapidly restored using the local data backup. 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 within one hour. With this service, the restoration duration is reduced.

6.3.2 FusionCompute Reliability

VM live migration: This feature ensures high reliability of the customer system. If a fault occurs on a running physical machine, its services can be migrated to other properly running machines before the situation turns worse. With this feature enabled, VM services are not interrupted by the replacement of server hardware.

Storage live and cold migration: FusionSphere offers cold migration and live migration for VM disks. Cold migration is to move VM disks from one data store to another when the VM is stopped. Live migration is to move VM disks from one data store to another without service interruption. With this feature enabled, VM services are not interrupted by the replacement of storage devices.

VM load balancing: If a new VM is started, VMs are live migrated, or computing nodes are remotely restarted due to faults, the system working in load balancing mode dynamically distributes loads based on the load status of each physical computing server in a cluster.

VM HA: If the physical server breaks down or restarts abnormally, the system can migrate the VMs with HA enabled to other computing servers, ensuring rapid restoration of VMs.

VM isolation: The virtualization technology can virtualize one physical server into multiple VMs. VMs are separated from each other. If a virtual machine fails, other virtual machines can still work properly. Users can use VMs the same as using physical machines.

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.

6.3.3 FusionManager OM Reliability

Management node HA deployment: The FusionSphere management system works in active/standby mode. The active node provides services through the floating IP address.

If the active node process is faulty or the OS on the active node or the host breaks down, the standby node takes over service processing.

During the switchover, the floating IP address is configured and the MAC address is updated on the gateway. All processes detected by the original active node start on the standby node and provide services.

Data consistency between the active and standby hosts: FusionSphere uses databases working in active/standby mode. The active database performs data read and write operations. If data in the active database changes, the changes will be synchronized to the standby database. To ensure the performance of the active database, asynchronous synchronization is performed between the active and standby databases. This prevents data loss if an active/standby database switchover occurs.

Real-time backup of management data: Manual backup is performed before an important operation, such as a system upgrade or critical data modification, is performed for FusionSphere. If the important operation fails or the operation has not achieved the expected result, the data backup can be used to restore the FusionSphere system to minimize the impact on services. Therefore, you need to back up the data of the management node in advance.

FusionManager supports the following functions:

- Interworking with third-party FTP servers to back up management data.
- Uploading the management data backups of each component to the third-party FTP servers.
- Instantly backing up the management data of components, such as FusionCompute and FusionManager.
- Querying the backup status.

6.3.4 Network Reliability

The network subsystem takes four measures to enhance system reliability:

- Uses the NIC binding technology to improve the availability of server ports.
- Uses the switch stacking technology to virtualize two switches into one, improving the link utilization and the reliability of access switches.
- Uses the Smart Link technology to connect aggregation switches.
- Uses the Virtual Router Redundancy Protocol (VRRP) to deploy active and standby routers at the core router side, improving the availability of the core network.

Figure 6-2 Networking providing high reliability

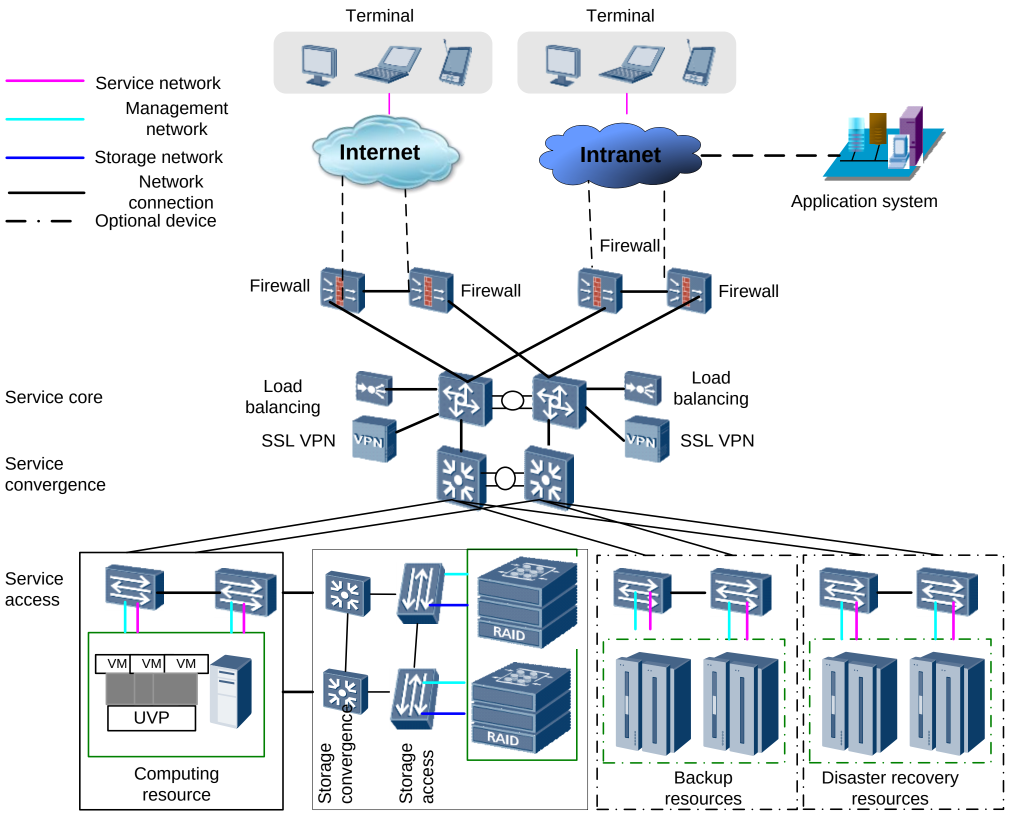

Computing nodes support the redundancy deployment of storage initiator modules, and VMs on these nodes can access the storage system using standard protocols, such as iSCSI. The load balancing, switch stacking, and clustering technologies of multiple NICs provide physical redundant storage paths.

Traffic control over virtualized networks allows users to configure the outbound bandwidth based on the network plane or the virtual NIC.

Multiple NICs provided by the physical server work in Bonding mode for reliability and load balancing.

6.4 Hardware Reliability

Hardware reliability refers to server hardware redundancy and fault detection. The following describes hardware reliability using Huawei RH2285 as an example.

Memory reliability: Self-developed servers, such as the RH2285, employ the Error Checking and Correcting (ECC) technology to detect 2-bit memory errors and correct 1-bit memory errors.

Hard disk hot swap: Self-developed servers, such as the RH2285, support hot swappable hard disks, including the SAS and SATA disks.

Hard disk RAID: Self-developed servers, such as the RH2285, support several RAID modes, such as RAID 0, RAID 1, and RAID 5, and support hot spare disks for RAID groups.

Online scheduled disk fault detection and precaution: Huawei FusionCloud data center virtualization solution uses the industry-leading self-monitoring, analysis and reporting (S.M.A.R.T.) technology to detect and manage the hard disks that are based on the Advanced Technology Attachment (ATA) and Internet Small Computer Systems Interfaces (iSCSIs).

Power supply reliability: Self-developed servers, such as the RH2285, involved in Huawei cloud data center solution are equipped with two power supply units (PSUs) that can generate alarms when a fault occurs. PSUs work in 1+1 redundancy mode and are hot swappable.

System monitoring: Self-developed servers, such as the RH2285, involved in Huawei cloud data center solution can monitor the temperature of key heat components in real time, such as the CPU and memory. They collaborate with intelligent fans to ensure reliable system running.

Onboard software reliability: BMC software supports double images. If one image is damaged in the Flash, the BMC starts from the other image. This prevents the failure in the system startup.

7 Conclusion

Cloud is born for agile IT. Huawei makes every effort to provide you with the FusionSphere solution that can improve the IT provisioning efficiency and help you win the fierce competition.

Cloud computing virtualizes different infrastructures into standard service components and then automatically combines these service components to meet user requirements. This document provides visibility into the standardization and automation capabilities provided by the FusionSphere solution and key technologies used in this solution. Furthermore, it provides some related documents for your reference.

You can click the following link and download a FusionSphere Foundation Edition for trial use:

http://enterprise.huawei.com/en/solutions/IT-solutions/cloud/index.htm

A Acronyms and Abbreviations

Acronym and Abbreviation	Full Name
ACL	access control list
BMC	baseboard management controller
CBT	Changed Block Tracking
CIFS	Common internet file system
CNA	Computing Node Agent
CPU	central processing unit
DHCP	Dynamic Host Configuration Protocol
DRS	dynamic resource scheduler
ECC	Error checking and correcting
eTOM	enhanced Telecom Operations Map
EVS	elastic virtual switch
FC SAN	Fiber Channel storage area network
FTP	File Transfer Protocol
HA	high availability
HBA	Host Bus Adapter
IaaS	Infrastructure as a Service
I/O	Input and Output
IP SAN	IP storage area network
IPsec	Internet Protocol Security
L2TP	Layer Two Tunneling Protocol
LUN	logical unit number

Acronym and Abbreviation	Full Name
MAC	Media Access Control
NAPT	Network Address Port Translation
NAS	network attached storage
NAT	Network Address Translation
NFS	Network File System
NVDIMM	non-volatile dual in-line memory module
PVSCSI	Paravirtual SCSI
QoS	quality of service
RAID	redundant array of independent disks
SAN	storage area network
SCSI	Small Computer Systems Interface
SNAT	Source Network Address Translation
SR-IOV	single-root I/O virtualization
UDS	Universal Distributed Storage
USB	Universal Serial Bus
UVP	Unified Virtualization Platform
VDC	virtual data center
VLAN	virtual local area network
VM	Virtual machine
VPC	virtual private cloud
VPN	virtual private network