Recommendations

- Huawei Learning Website
  - http://learning.huawei.com/en
- Huawei e-Learning
  - https://ilearningx.huawei.com/portal/#/portal/ebg/51
- Huawei Certification
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_31&lang=en
- Find Training
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_trainingsearch&lang=en

More Information

- Huawei learning APP

Overview of Converged Resource Pools in Cloud DCs

www.huawei.com

Foreword

- This course introduces the definition of data centers (DCs), the architecture of converged resource pools in Huawei cloud DCs, and related software components and functions.

Objectives

- On completion of this course, you will be able to know:
  - Architecture of converged resource pools in cloud DCs
  - Software components and functions of converged resource pools in cloud DCs

Contents

1. **Overview and Development of DCs**
2. Architecture of Converged Resource Pools in Cloud DCs
3. Software of Converged Resource Pools in Cloud DCs

Definition of DCs

What Is a DC?

- Core of the enterprise IT system
- Center of massive data computing, exchange, and storage
- Computing environment for critical information applications
- Environment for centralized management and control of various data, applications, and physical or virtual devices

Four focuses of a DC: reliability, flexibility, energy conservation, and resource utilization

HUAWEI

- Internet Data Center (IDC) is a platform for data storage and circulation. It involves the most intensive data exchange on the Internet.

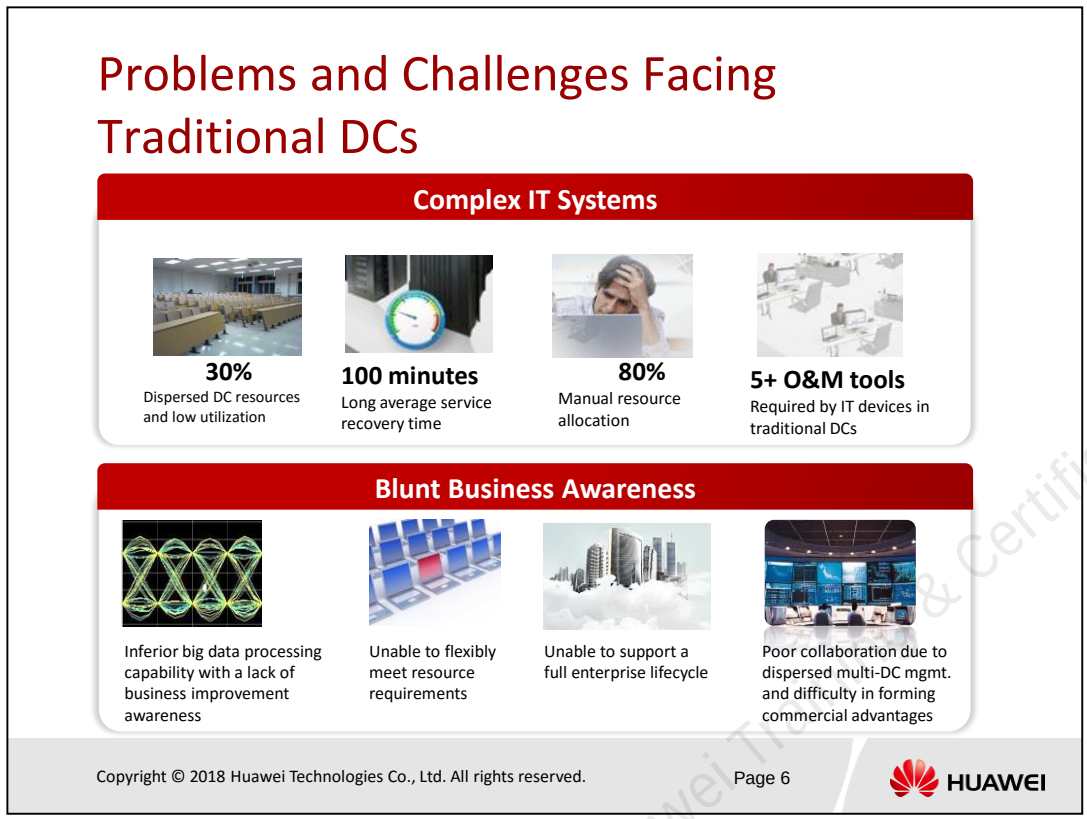

- Inferior server utilization and soaring cost
  - The virtualization technology is one solution for resolving the issue.
  - Frequent problem occurrence with a time-consuming handling process
  - IDC data shows that 20% of faults takes over one day to resolve, although the fault causes usually turn out to be quite simple.
- Slow service rollout
  - Currently, all industries are experiencing faster changes. Limited by complicated IT systems, the speed of rolling out service support systems cannot keep up with the changes and development of industries. In this case, the following questions need to be answered: how can services be effectively managed and revenues be increased? Can new services be fast developed? Can users subscribe services as easily as online shopping or online banking? Can users deliver services fast after purchasing them? Can reselling be supported and new revenue pipelines be explored?
  - Complex management and high management costs
  - According to IDC statistics, annual maintenance costs account for 2/3 of the new equipment purchase costs.
  - According to statistics, 2/3 of DCs adopt more than five management tools, which easily leads to management silos.
  - IDC statistics show that IT O&M costs account for almost half of the TCO. However, administrative personnel are still busy and short-handed.

- Cloud computing is not only a new technology but also a new business model. Cloud is a way of social resource aggregation. Social resources exist in the form of "water molecules" and can be aggregated into cloud when necessary. When the cloud is not used any more, it will dissipate. Cloud-based delivery has widely applied in our daily life. For example, each transaction in the e-commerce is a cloud-delivery process. Cloud computing is the outcome of the Internet development.
- A cloud DC is where cloud services are deployed and are running.

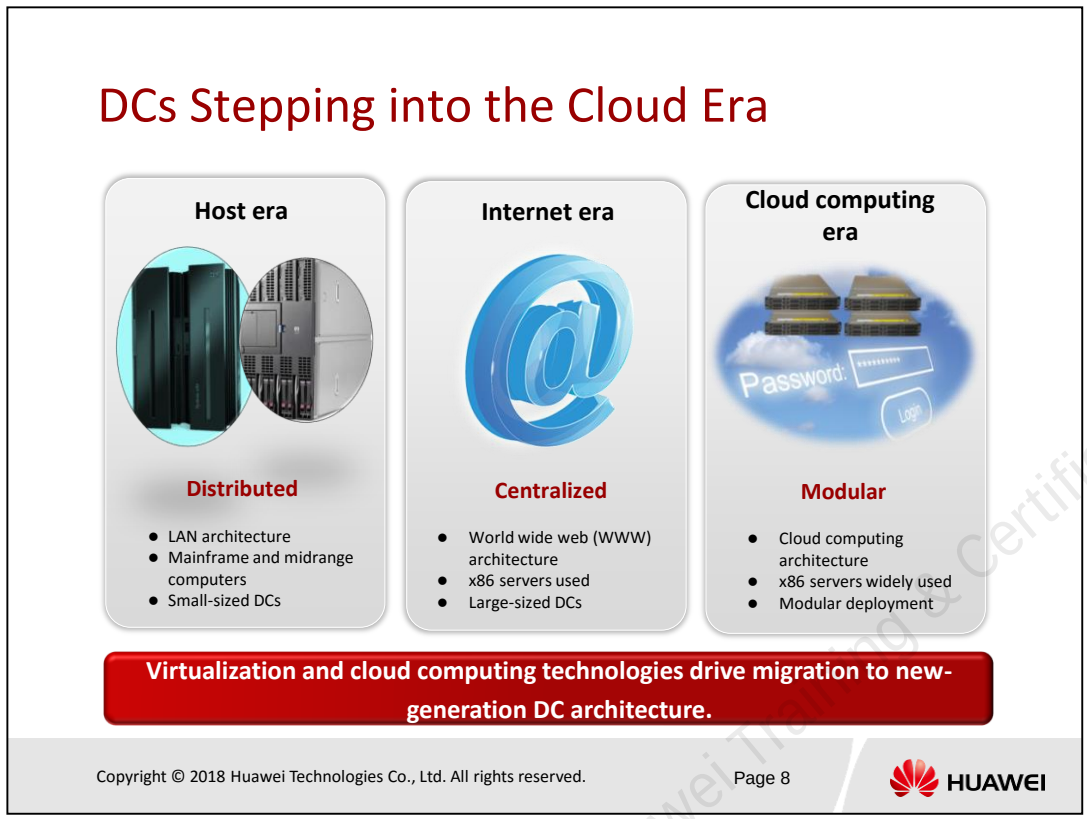

- Host era:
  - LAN architecture for small-sized DCs where mainframe and midrange computers are deployed
  - Physical facilities often encountering single points of failure
  - Dual-system hot backup or local backup
- Internet era:
  - WWW network architecture for large-sized DCs where x86 servers and other models of servers are deployed
  - Centralized management and maintenance
  - Remote backup and DR for DCs
- Cloud computing era:
  - Cloud computing architecture for DCs where x86 servers are deployed
  - Modular deployment and delicacy management
  - VM-based DR and backup

Contents

1. Overview and Development of DCs
2. **Architecture of Converged Resource Pools in Cloud DCs**
3. Software of Converged Resource Pools in Cloud DCs

Overall Architecture of Next-Generation DCs

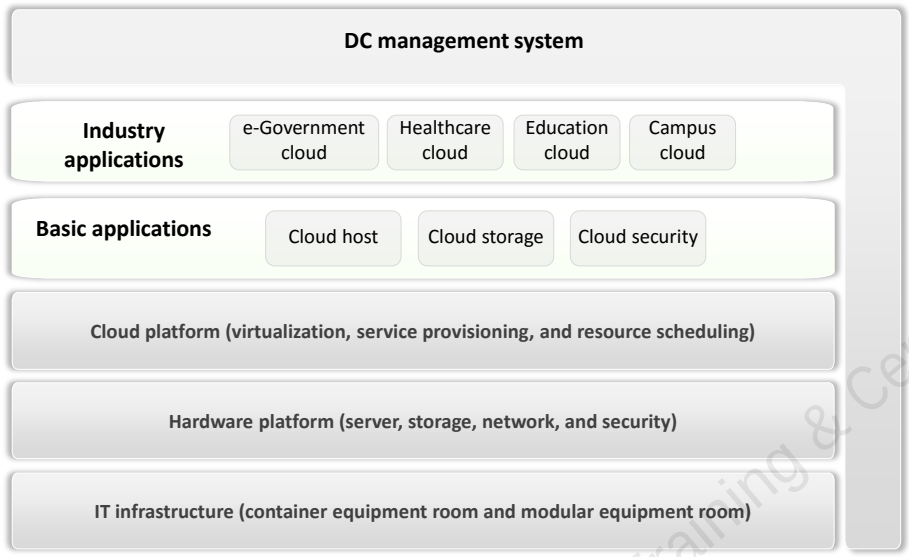

Logical Architecture of Huawei FusionSphere Solution

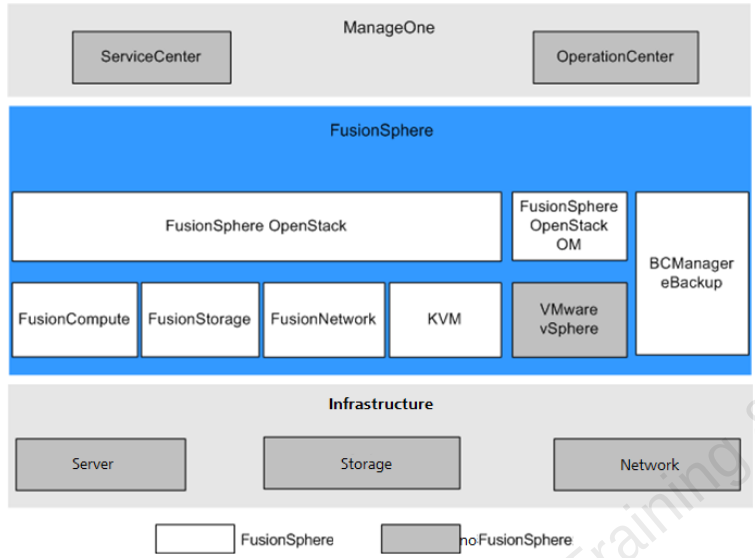

Software Components of the FusionSphere Solution (1)

Component	Function
FusionCompute	Optional FusionCompute is a cloud operating system that virtualizes hardware resources and centrally manages virtual resources, service resources, and user resources. FusionCompute virtualizes computing, storage, and network resources using virtual computing, virtual storage, and virtual networking technologies.
FusionStorage	Optional FusionStorage is distributed storage software that has both storage and computing capabilities. It can be deployed on general x86 servers to provide block storage by combining local disks on all the servers into a virtual storage resource pool.
FusionNetwork	Mandatory FusionNetwork is solution designed for evolving traditional virtual switching to forward-thinking software-defined networking (SDN). FusionNetwork leverages Virtual Extensible LAN (VXLAN) L2 tunnel encapsulation technologies and Huawei SDN controllers to automatically deploy and configure SDN networks, control service level agreement (SLA), and isolate and layer multiple tenants. Note: Functions of FusionNetwork are provided and fulfilled by the Neutron service in the FusionSphere OpenStack system.

Software Components of the FusionSphere Solution (2)

Component	Function
FusionSphere OpenStack	Mandatory FusionSphere OpenStack is enhanced and reinforced by Huawei for enterprises based on Mitaka of the OpenStack community. FusionSphere OpenStack centrally schedules and manages virtual computing, storage, and network resources over unified RESTful APIs. It reduces the operating expense (OPEX) and provides high system security and reliability, helping enterprises build secure and energy-saving private cloud capabilities.
KVM	Optional The server operating system delivered by the FusionSphere OpenStack software package supports KVM virtualization, which enables compute nodes with KVM virtualization installed to be deployed in FusionSphere.
FusionSphere OpenStack OM	Mandatory FusionSphere OpenStack OM, the O&M component of FusionSphere OpenStack, monitors and manages FusionSphere hardware and software and supports infrastructure O&M. It provides a web user interface (UI) for internal O&M personnel to perform operations and management.
BCManager eBackup	Optional BCManager eBackup software is abbreviated to eBackup. It provides E2E backup solutions for all virtualization applications.

Contents

1. Overview and Development of DCs
2. Architecture of Converged Resource Pools in Cloud DCs
3. **Software of Converged Resource Pools in Cloud DCs**

Logical Architecture of FusionCompute

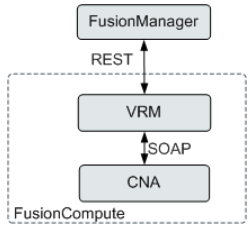

Module	Function
CNA	CNA provides the following functions: • Enables virtual computing. • Manages VMs on compute nodes. • Manages the computing, storage, and network resources on compute nodes.
VRM	VRM provides the following functions: • Manages block storage resources in clusters. • Assigns private IP addresses to VMs based on Dynamic Host Configuration Protocol (DHCP). • Manages compute nodes in clusters, and maps physical computing resources to virtual computing resources. • Manages network resources in clusters, such as IP addresses, VLANs, and DHCPs, and assigns IP addresses to VMs. • Specifies the lifecycle of VMs in clusters, and allocates and migrates VMs on compute nodes. • Dynamically adjusts resources in clusters. • Externally provides elastic computing, storage, and IP address services through unified management of virtual resources and user data. • Provides unified O&M APIs. O&M personnel use WebUI to remotely access FusionCompute for O&M of the entire system, including resource management, monitoring, and reports.

Technical Characteristics of FusionStorage

- Linear and flexible scalability
- High performance
- High reliability
  - Multiple data redundancy protection mechanisms are available.
- Rich advanced storage functions
  - Thin provisioning, synchronous replication DR...
  - Snapshot, snapshot backup, linked clone...

FusionStorage Architecture

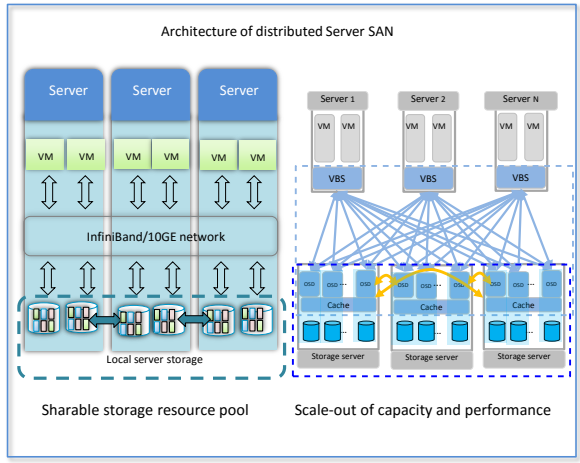

Definition: Server SAN is a storage resource pool which consists of storage of multiple independent servers and incorporates computing and storage resources.

Dedicated devices converted to universal devices

Linear scaling of computing and storage

Simple management and low TCO

FusionSphere

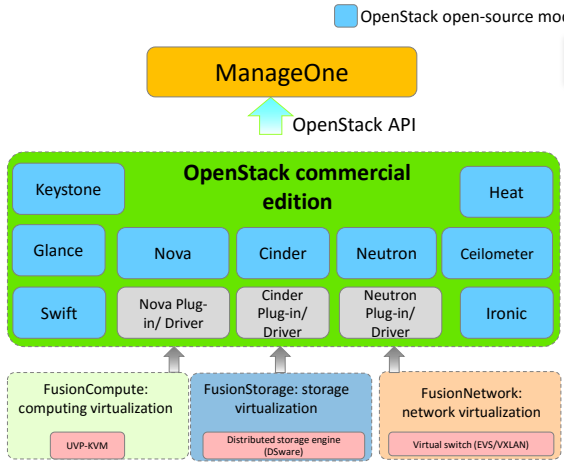

Gold member and gold director of the OpenStack community

- Standard OpenStack APIs
  - Originates from the native APIs of the OpenStack community.
  - Keeps pace with latest OpenStack versions in a rapid way.
- Broad support for third-party vendors
  - Compatible with the OpenStack community ecosystem chain.
  - Supports heterogeneous hypervisors and hardware devices.
- SOA-based loosely coupled architecture
  - Decouples computing, storage, and network.
  - Supports hybrid networking of multi-vendor resource pools.

ManageOne

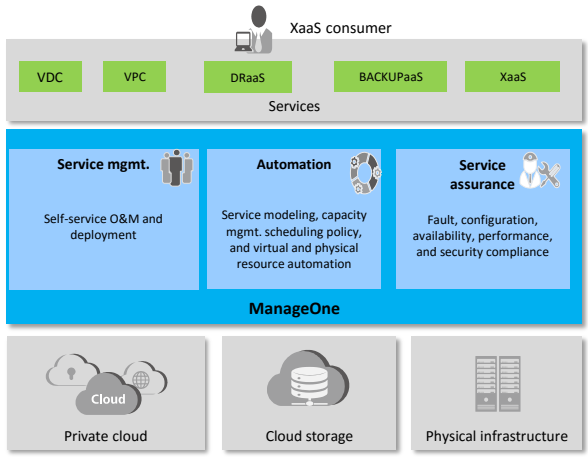

Efficient and intelligent management system

ManageOne provides cloud DC management solutions for customers.

- OperationCenter (OC)

  An O&M center where users perform scenario-specific, proactive, and predictable O&M operations

- ServiceCenter (SC)

  A service center that provides unified service orchestration and automatic management

Agile Controller

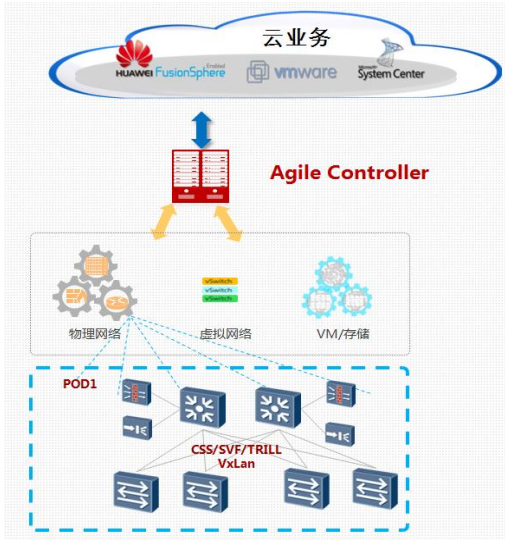

Efficient, simple, and open

- Efficient provisioning of network resources

Activates tenant network services on demand, simplifies the process of applying for network resources, and implements automatic network resource provisioning.

- Simple O&M for virtual and physical networks

Centrally displays virtual and physical networks, enabling visualized management and control of cloud networks.

- Openness

Provides diversified open APIs that quickly interconnect with mainstream cloud platforms.

Huawei Training & Certification

BCManager eBackup

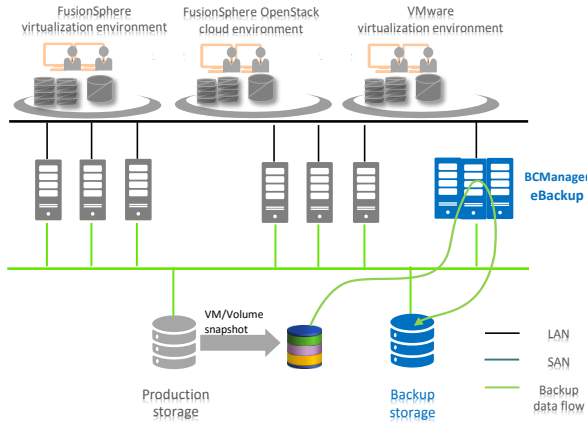

- eBackup is a backup software product designed for FusionSphere and VMware vSphere virtualization and cloud platforms. With VM snapshot and change block tracing (CBT) technologies, eBackup can comprehensively protect VM data. With highlights such as ease-of-use and cost-effectiveness, it meets backup requirements of a large number of VMs.

Quiz

1. Which software components does the FusionSphere cloud DC solution consist of?

- Reference answer: FusionSphere OpenStack, FusionSphere OpenStack OM, ManageOne SC&OC, FusionCompute, BCManager eBackup...

Recommendations

- Huawei Learning Website:
  - http://support.huawei.com/learning/en/newindex.html
- Huawei Support Case Library:
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Cloud Data Center OpenStack Architecture and Principles (Volume 1)

www.huawei.com

Foreword

- This section describes OpenStack and its components for you to have a basic understanding of the OpenStack architecture and principles.

Objectives

- After finishing this course, you will be able to:
  - Know what OpenStack is.
  - Understand the OpenStack project structure.
  - Understand OpenStack component principles.

Contents

1. **OpenStack Overview**
2. OpenStack Component – Nova
3. OpenStack Component – Cinder
4. OpenStack Component – Neutron
5. OpenStack Component – Swift
6. OpenStack Component – Glance
7. OpenStack Component – Keystone

Origin and History of OpenStack

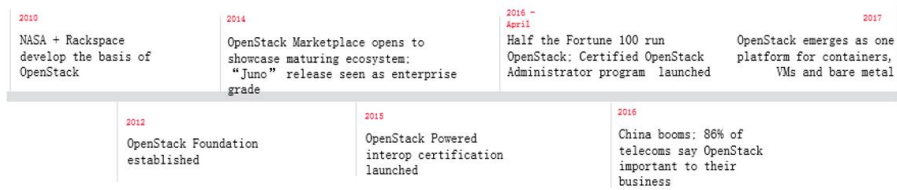

- An OpenStack version is released every 6 months.

What is OpenStack?

- Most popular open-source cloud operating system (OS) kernel
- OS
  - Resource abstraction
    - OpenStack abstracts hardware resources into resource pools based on virtualization and software definition.
  - Resource allocation and load scheduling
    - Based on the requirements of administrators and users, OpenStack allocates resources in the resource pool to different users to carry different applications.
  - Application lifecycle management
    - OpenStack can provide basic application deployment and withdrawal, and automatic adjustment on the application scale.
  - System O&M
    - OpenStack can provide certain system monitoring capabilities.
  - Man-machine interaction
    - OpenStack provides man-machine interfaces to allow users to use APIs, command lines, or a GUI to interact with OpenStack.

OpenStack Is Not Virtualization

- OpenStack architecture positioning and technical scope
  - OpenStack is only the control plane of the system.
  - OpenStack does not contain data plane components, such as hypervisor, storage, and network devices.
- Cloud and virtualization have key differences.

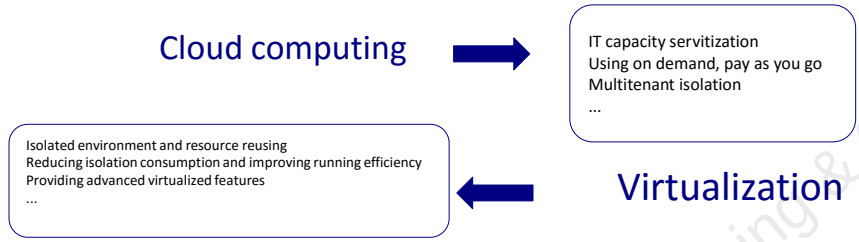

Virtualization is one of the technologies to implement cloud computing, but not a core focus of cloud computing.

OpenStack Is Not a Cloud

- OpenStack is only a key component for cloud construction.
  - Kernel, backbone, framework, bus
- To build a cloud, we need many things.

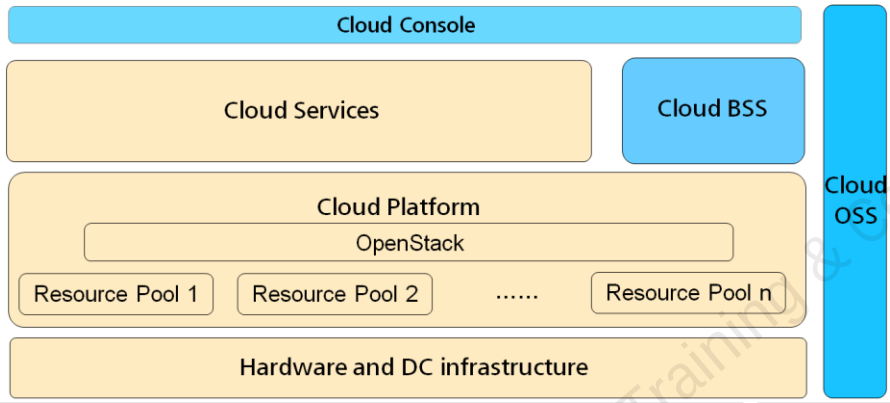

Introduction to OpenStack Services (1)

Service	Project Name	Description
Dashboard	Horizon	Provides a web-based self-service portal to interact with underlying OpenStack services, such as starting an instance, assigning IP addresses, and configuring access control.
Compute	Nova	Manages the computing instance lifecycle in the OpenStack environment. Responds to operations such as generating, scheduling, and recycling VMs on demand.
Networking	Neutron	Ensures that other OpenStack services provide the network as a service (NaaS), such as OpenStack computing. Provides APIs for users to use networks. Its plugin architecture supports many network providers and technologies.
Object Storage	Swift	Swift uses Proxy Server to externally provide HTTP-based REST APIs to perform CRUD operations on accounts, containers, and objects. It has a high error tolerance mechanism based on data replication and a scalable architecture. It is implemented in the same way as a directory needs to be mounted to a file server. In this way, Swift writes objects and files to multiple hard disks to ensure that data is in multiple copies across servers in a cluster.

Introduction to OpenStack Services (2)

Service	Project Name	Description
Block Storage	Cinder	Provides persistent block storage for running instances. Its pluggable driver architecture helps to create and manage block storage devices.
Identity Service	Keystone	Provides authentication and authorization services for other OpenStack services and provides an endpoint directory for all OpenStack services.
Image Service	Glance	Stores and searches for VM disk images. Nova uses Glance during instance deployment.
Telemetry Service	Ceilometer	Monitors and measures the accounting, benchmark, scalability, and statistics of the OpenStack cloud.
Orchestration Service	Heat	Heat can use the local template format or the AWS CloudFormation template format to orchestrate multiple integrated cloud applications through the OpenStack local REST APIs or CloudFormation compatible queue APIs.

Huawei Training & Certification

OpenStack Project Layering

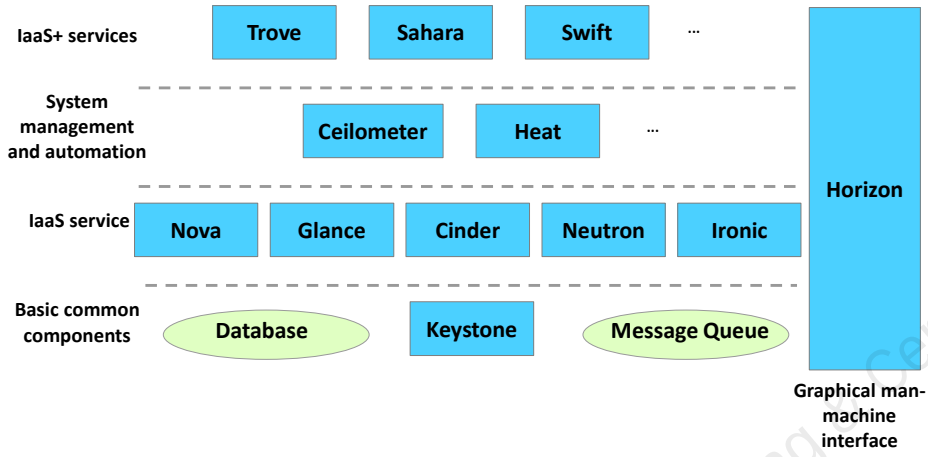

- By now, the OpenStack project covers common service types at the IaaS layer, part of system management and automation services, and some important IaaS+ services.

Relationships Between OpenStack Projects

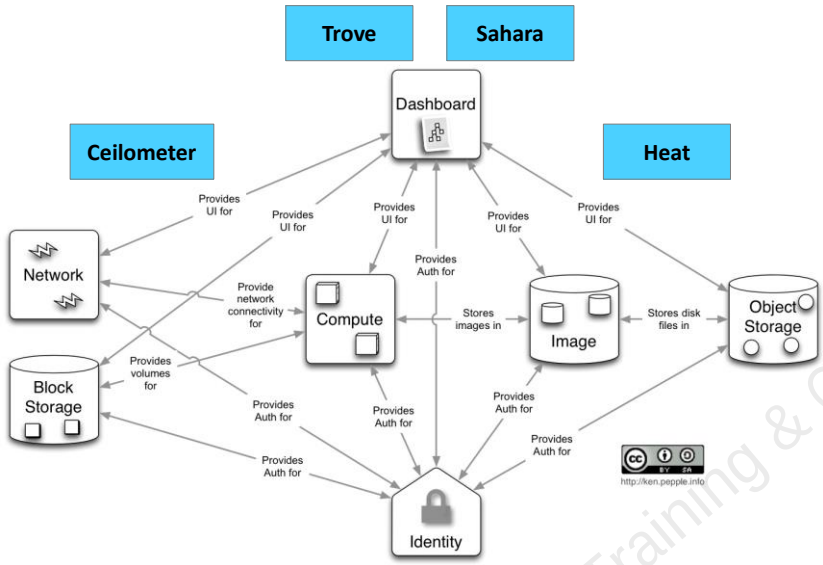

Physical Deployment of OpenStack

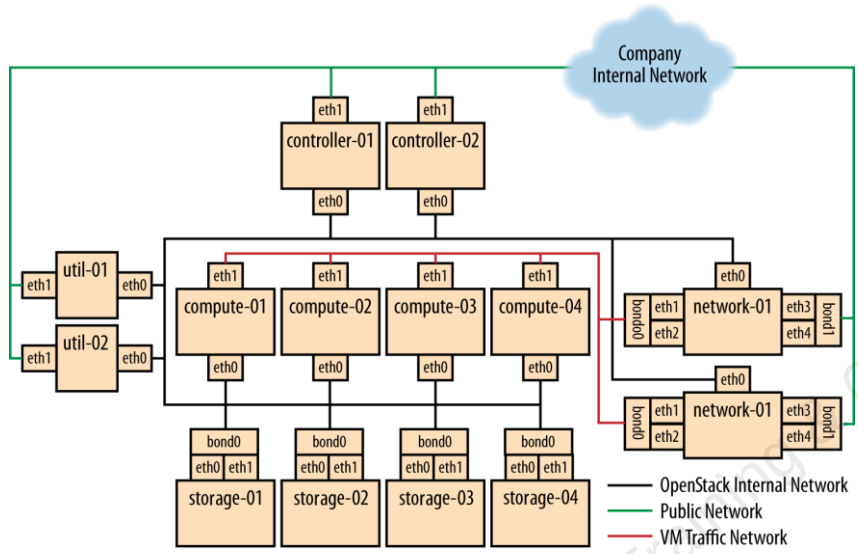

Contents

1. OpenStack Overview
2. **OpenStack Component – Nova**
3. OpenStack Component – Cinder
4. OpenStack Component – Neutron
5. OpenStack Component – Swift
6. OpenStack Component – Glance
7. OpenStack Component – Keystone

Nova Overview

- Nova:
  - Provides computing resource services.
- Nova is responsible for:
  - VM lifecycle management
  - Lifecycle management of other computing resources
- Nova is not responsible for:
  - Management of physical hosts accommodating VMs
  - Comprehensive system status monitoring

Nova Logical Architecture – KVM Scenario

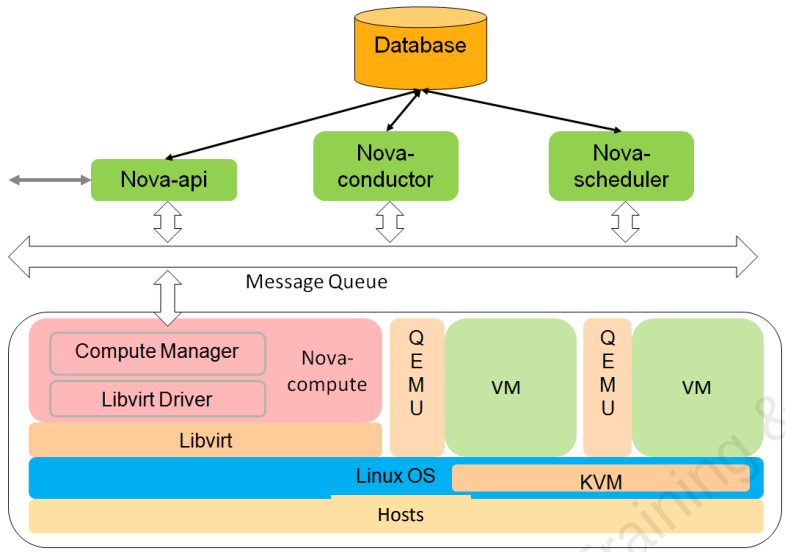

System Architecture of Nova

Module	Function	Common Deployment Position
nova-api	Receives the rest message.	Control node
nova-scheduler	Selects a proper host.	Control node
nova-conductor	Perform database operations and complex flow control.	Control node
nova-compute	Manages the VM lifecycle and resources.	Compute node
nova-novncproxy	Allows users to use VNC to access VMs.	Control node
nova-consoleauth	Authenticates users to use VNC to access VMs.	Control node

Nova Deployment Example

- No central structure.
- Each component has no local persistent status.
- Supports horizontal expansion.
- Generally, nova-api, nova-scheduler, and nova-conductor are deployed on the control node.
- Multiple control nodes are deployed to implement high availability (HA) and load balancing.
- Add control nodes and compute nodes to facilitate system expansion.

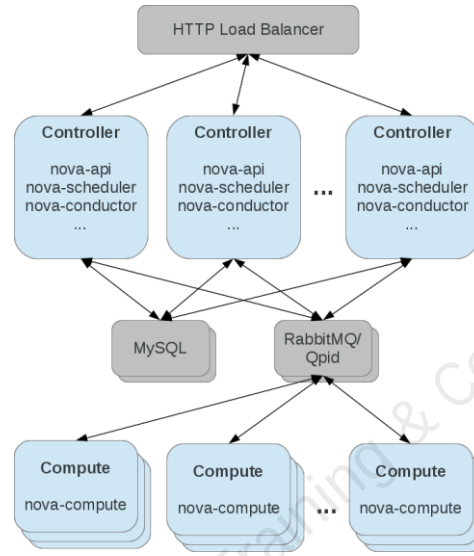

Computing Resource Models of OpenStack Nova

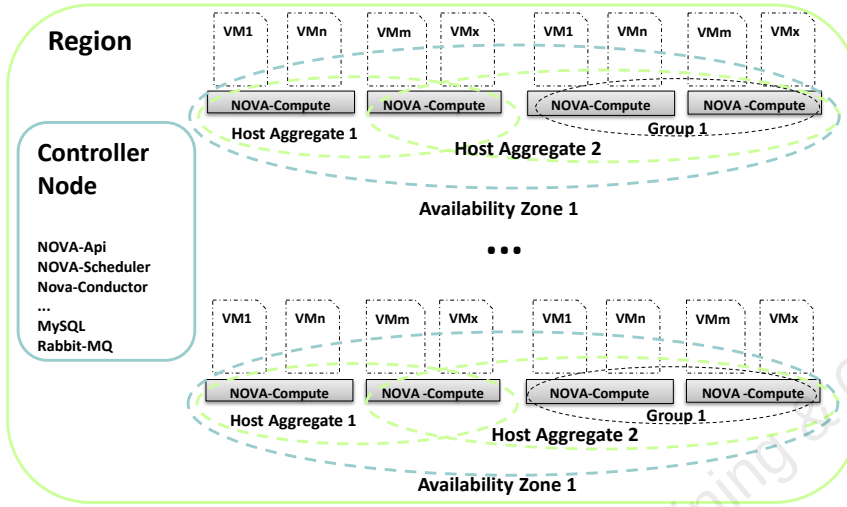

Resource grouping/organization under OpenStack

Typical Operations of Nova

Group	Description
VM lifecycle management	VM creation, deletion, startup, shutdown, restart, reestablishment, specification change, suspension, resume, pending, migration, online migration, lock, unlock, evacuation, rescue, unrescue, holding, holding deletion, holding restoration, backup, image export from the VM, list, detailed information, information query, change, and password change
Volume and snapshot management operation	Essentially, it encapsulates the Cinder API and supports volume creation, deletion, list, information query. Snapshot management includes snapshot creation, deletion, list, and information query.
VM volume operation	Supports volume attaching and detaching, VM volume list, and VM volume information query.
Virtual network operation	Essentially, it encapsulates the Neutron API and supports VM network creation, deletion, list, and information query.
vNIC operation	Supports NIC attaching and detaching, and vNIC list.
VM image operation	Essentially, it encapsulates Glance API and supports the image creation, deletion, list, and detailed information query.
VM HA	Reliability operations of Huawei-developed VMs
Other resources and operations	Flavor, host group, keypair, and quota

- Evacuation: Move an instance from a faulty or stopped compute node to another host in the same environment.
- Rescue: The VM is normal, but the system is faulty (such as password loss and damaged system files). Start the system in a new image and mount the original system to a new partition for recovery. Then restart the system in the original system disk.
- Holding: Release the VM that is not used for a long time from the bottom to save server resources. Users can recover the VM as needed.
- Backup: Back up VMs that are started using an image, upload the backup files to Glance, and retain the number of backup files specified by the user.
- Image export: Local volumes can be exported. Snapshots are created for all volumes of the VM started using the volume. For VMs that are started using a disk, snapshots are created for only the local system disk.

Contents

1. OpenStack Overview
2. OpenStack Component – Nova
3. **OpenStack Component – Cinder**
4. OpenStack Component – Neutron
5. OpenStack Component – Swift
6. OpenStack Component – Glance
7. OpenStack Component – Keystone

What Is Cinder?

- Cinder is an OpenStack component, which is separated from Nova-Volume of OpenStack Folsom.
- The core function of Cinder is to manage volumes and process volumes, volume types, volume snapshots, and volume backup.
- Cinder provides persistent block storage services that provision resources on demand through unified interfaces.
- Cinder uses the driver to access different types of backend storage devices (local storage, network storage, FC SAN, IP SAN, and FusionStorage).

Logical Architecture of Cinder

- Cinder-api
  - The only entry for users to access Cinder and the endpoint of Cinder that receives and processes REST requests
- Cinder-scheduler
  - Selects a proper backend to process tasks according to the preset scheduling filtering policy and weight calculation policy.
- Cinder-volume
  - Interconnects with backend storage devices and converts OpenStack operations into storage operations through the drivers provided by vendors.

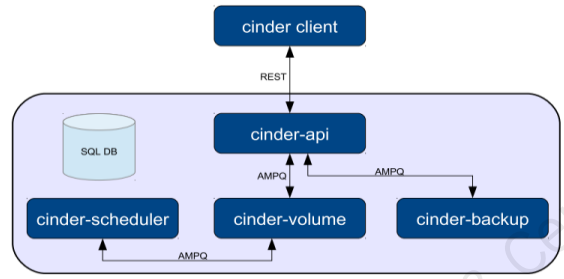

- Cinder Client encapsulates the REST APIs provided by Cinder in Command Line Interface (CLI) mode.
- Cinder API provides the REST APIs externally, parses the operation requirements, and seeks corresponding handling methods using APIs including adding, deleting, modifying, and querying volumes (volume creation from the source volume, image, and snapshot), adding, deleting, modifying, and querying snapshots, backup, volume type management, attaching and detaching (Nova invocation).
- Cinder-scheduler collects capacity and capability information reported by the backend and allocates the volume to specified Cinder-volume according to the set algorithm.
- Cinder-volume is deployed on multiple nodes. Storage vendors insert code to Cinder-volume, and Cinder-volume uses different configuration files and connects to different backend devices. Each storage vendor inserts driver code to Cinder-volume, and Cinder-volume interacts with the devices to collect device capacity and capability information and perform volume operations.
- Cinder-backup backs up volume data to other storage medium (current SWIFT/Ceph/TSM provides the drive).
- SQL DB provides storage volume, snapshot, backup, and service data, supports SQL databases such as MySQL, PG, and MSSQL.

Physical Deployment of Cinder

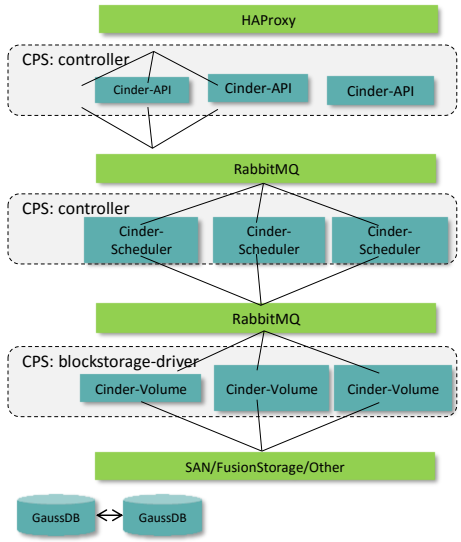

- Cinder-api, Cinder-scheduler, and Cinder-volume can be deployed on one node or separately.
- The API uses the active/active mode, and HAProxy functions as an LB to distribute requests to multiple Cinder APIs.
- Scheduler uses active/active mode. RabbitMQ distributes tasks to three nodes in load balancing mode and receives the capability information reported by Cinder-volume from RabbitMQ. During scheduling, scheduler reserves resources in the database to ensure data consistency.
- Cinder-volume also uses the active/active mode, reports the capacity and capability information of the same backend, and receives requests for processing.
- RabbitMQ supports the active/standby mode or cluster mode.
- GaussDB supports the active/standby mode.

Cinder Logic

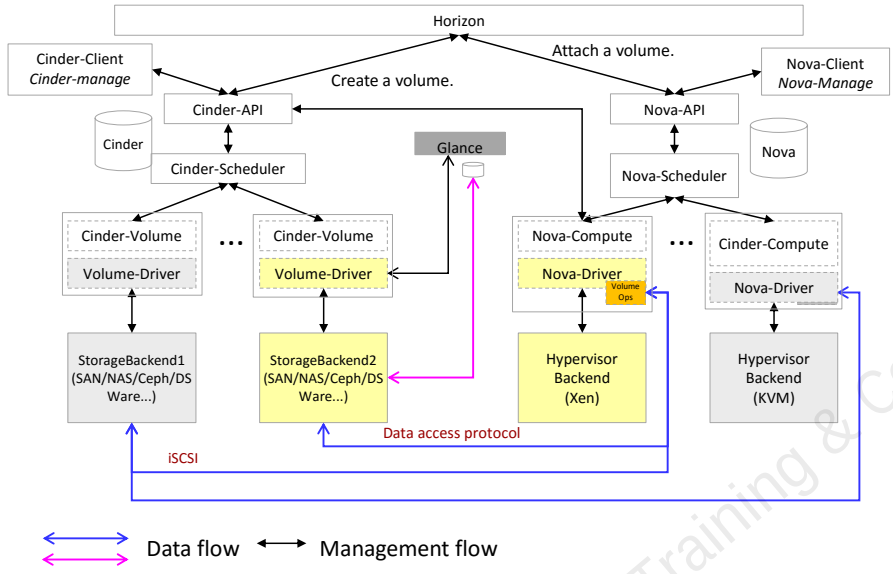

HUAWEI

- Cinder is responsible for the full lifecycle management of volumes (volume creation, deletion, and status maintenance). Nova initiates volume use (attaching and detaching).
- Lines in pink and blue indicate data channels. Other lines indicate management channels.
- Components in each full line box can be deployed across servers.

Contents

1. OpenStack Overview
2. OpenStack Component – Nova
3. OpenStack Component – Cinder
4. **OpenStack Component – Neutron**
5. OpenStack Component – Swift
6. OpenStack Component – Glance
7. OpenStack Component – Keystone

Position of Neutron in OpenStack

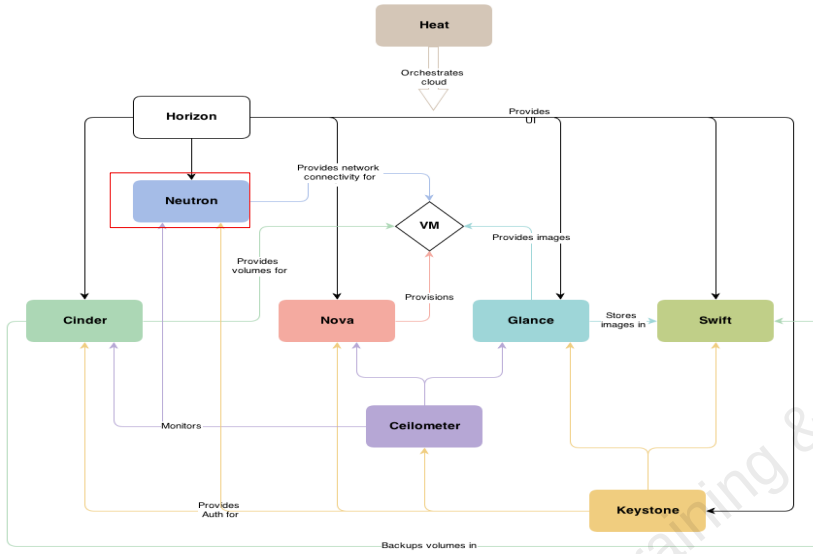

Neutron Overview

- Network connectivity as a service

Neutron management object

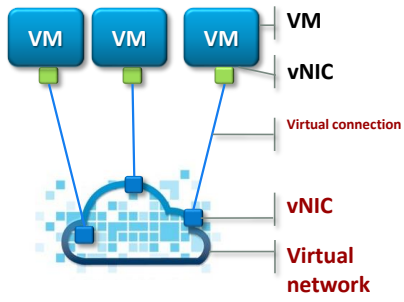

Logical network control for tenants

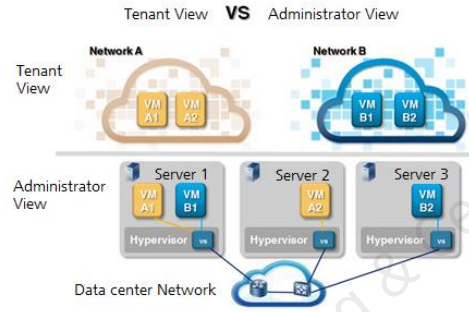

- With the development of the Neutron project, a variety of L2 to L4 network services, such as LBaaS, VPNaaS, FWaaS, have been extended based on basic services.

Network Module - Neutron Structure

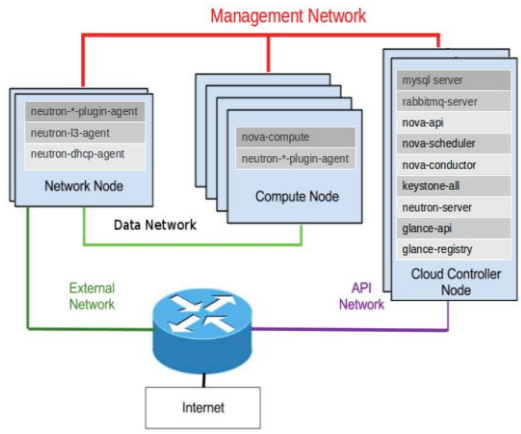

- Network Module:
  - Management Network
  - Data Network
  - External Network
  - API Network

- Management Network
  - Provides internal communication between OpenStack Components. IP addresses on this network should be reachable only within the data center.
- Data Network
  - Provides VM data communication within the cloud deployment. The IP addressing requirements of this network depend on the Networking plug-in that is used.
- External Network
  - Provides VMs with Internet access in some deployment scenarios. Anyone on the Internet can reach IP addresses on this network.
- API Network
  - Exposes all OpenStack APIs, including the Networking API, to tenants. IP addresses on this network should be reachable by anyone on the Internet. The API network might be the same as the external network, because it is possible to create an external-network subnet that is allocated IP ranges that use less than the full range of IP addresses in an IP block.

Neutron Architecture and Components (1)

- Neutron architecture
  - Vendor extends the Neutron API.
  - Vendor can use custom or open source plugins.

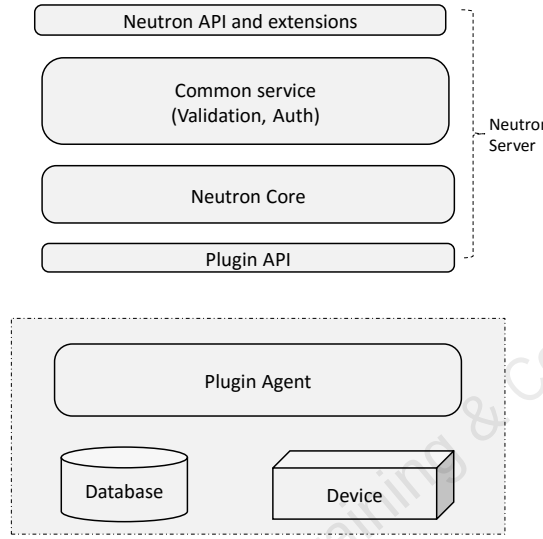

Neutron Architecture and Components (2)

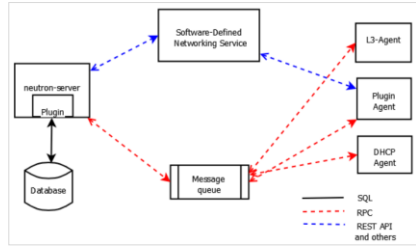

Components	Function	deployment location
Neutron-server & plugin	Receives REST requests, authenticates Keystone, interacts with the database, and provides API functionality.	Controller node
Neutron database	Manages data.	Controller node
L2 agent (OVS-agent)	Connects to the network port to process the packet.	Compute node
DHCP-agent	Provides the DHCP service.	Network node
L3-agent	Provides Layer 3 forwarding services for clients to access external networks.	Network node
RPC	Message queue for communication between plugins and agents	All nodes

Neutron Server Overview (1)

- Neutron Server includes Neutron-common and plugins.

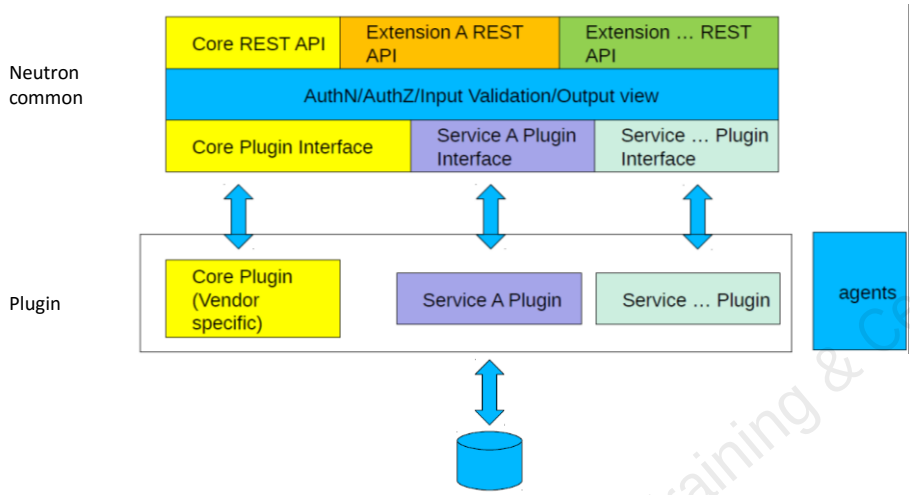

Neutron Server Overview (2)

- API – Core
  - A minimal set of plugin functions including network, subnet, and port query, addition, and deletion, and update operations.
- API – API Extensions
  - They are typically implemented for specific plugins so that tenants can use unique capabilities of these plugins, such as access control (ACL) and QoS.
- Plugin
  - Store the configuration information of the current logical network, determine and store the corresponding relationship between the logical network and the physical network (for example, select a VLAN for a logical network), and communicate with one or more switches to implement the communication.

Neutron Architecture and Components

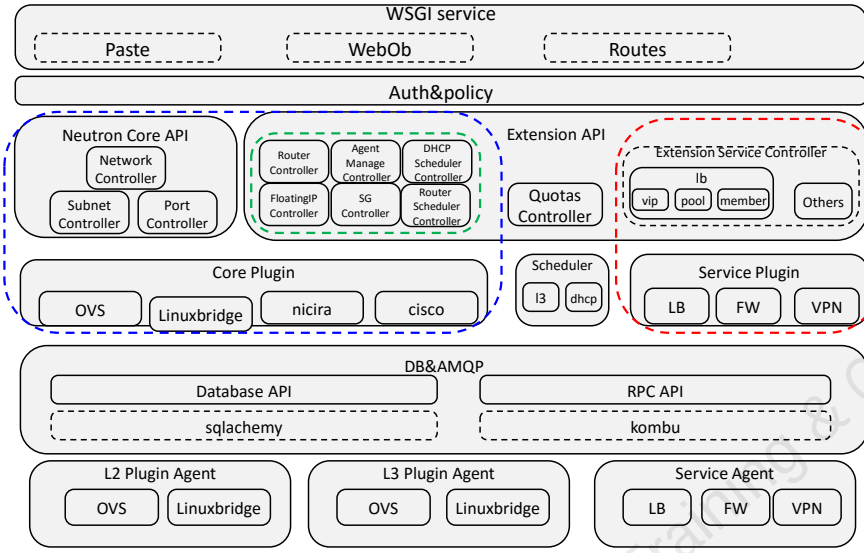

Contents

1. OpenStack Overview
2. OpenStack Component – Nova
3. OpenStack Component – Cinder
4. OpenStack Component – Neutron
5. **OpenStack Component – Swift**
6. OpenStack Component – Glance
7. OpenStack Component – Keystone

Position of Swift in OpenStack

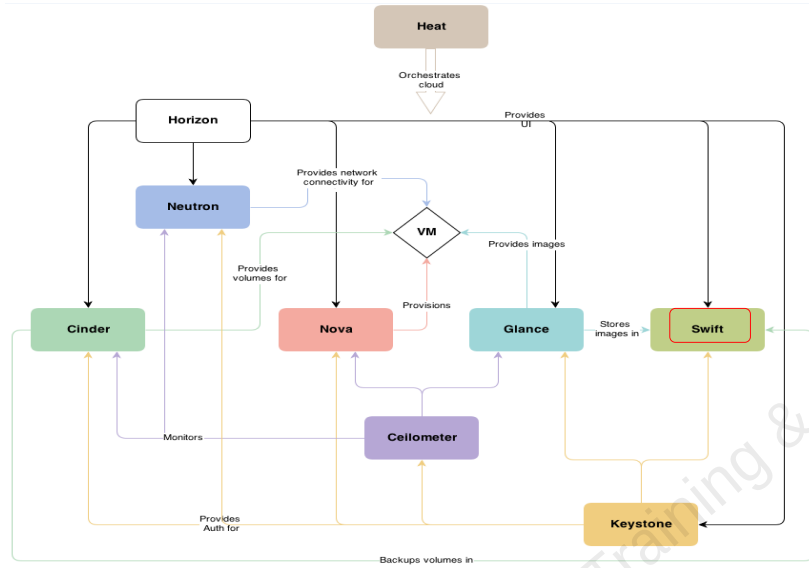

What Is Swift?

- HA and distributed object storage service
- Final consistency model
- Suitable for solving the problem of unstructured data storage in the Internet application scenario.
- Constructed based on a cheap standard hardware storage infrastructure.

- Swift was originally a high-availability (HA) distributed object storage service developed by Rackspace. In 2010, Swift was contributed to the OpenStack open-source community as one of the first core sub-projects, providing the VM image storage service for the Nova sub-project. Swift is constructed on the standard hardware storage infrastructure that is cheap without redundant array of independent disks (RAID). Swift achieves HA and scalability by using consistency hash and data redundancy at the software layer and sacrificing a certain degree of data consistency. It supports the multi-tenant mode, container, and object read/write operation, which is suitable for resolving unstructured data storage problems in Internet application scenarios.

Swift Storage Virtualization – Object Storage

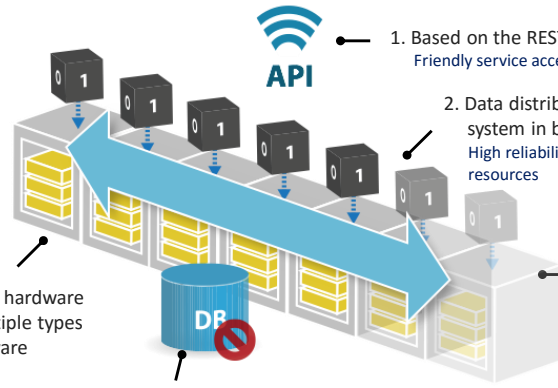

1. Based on the REST APIs
   Friendly service access mode
2. Data distributed in the entire system in balanced manner
   High reliability and efficient use of resources
3. Independent from hardware and supports multiple types of standard hardware
   No customization requirements for dedicated hardware devices
4. Easy to expand
5. No central database.
   No single-point performance bottleneck or single point of failure (SPOF) risks
6. Account/Container/Object
   The three-level storage structure does not need the file system and has N ($N \geq 3$) copies.
   High data reliability

Swift Principles and Architecture

- Consistency hash
- Data consistency model
- Data model
- Ring

Consistency Hash

- By calculation, objects can be evenly distributed to virtual nodes in the virtual space.
- When a node is added or deleted, the quantity of data to be moved is greatly reduced.
- In actual deployment, a proper quantity of virtual nodes needs to be calculated carefully to balance the storage space and workloads.

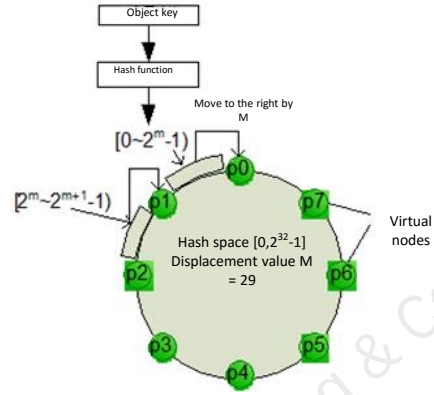

Data Consistency Model

- N: Number of data copies; W: Number of copies that are accepted in the write operation; R: Number of copies in the read operation
- Strong consistency: R + W > N
- Weak consistency: R + W ≤ N
- By default, the Swift configuration is: N = 3, W = 2 > N/2, R = 1 or 2. That is, each object has three copies, which are stored on nodes in different zones. W = 2 indicates that write is successful only when at least two copies are updated. When R = 1, a message indicating a successful operation is returned immediately. In this case, the old version (weak consistency model) may be read. When R = 2, the x-newest=true parameter needs to be added in the read operation request header to read metadata information of two copies at the same time. Then the timestamps are compared to determine which version is the latest (strong consistency model).

Data Model

- The hierarchical data model consists of three layers of logical structures: Account/container/object
- The number of nodes at each layer is not limited and can be extended randomly.

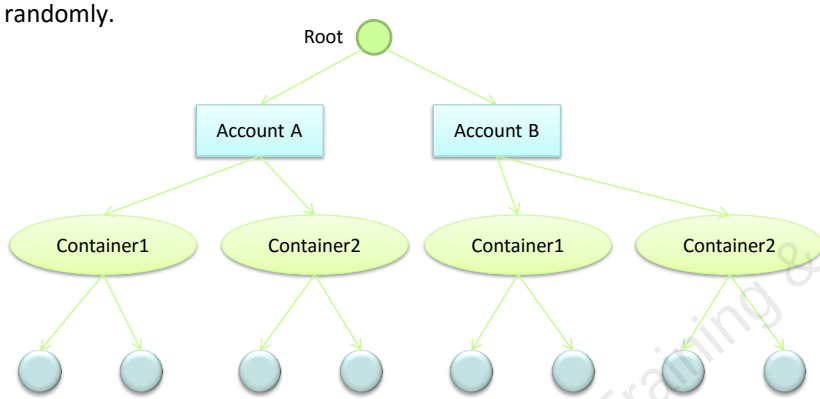

Ring

- Maps a virtual node (partition) to a group of physical storage devices.
- Account ring, container ring, object ring
- Consists of:
  - Storage device list, covering the unique identifier (ID), zone number, weight, IP address, port, device name, and metadata
  - Mapping from a partition to a device
  - Calculation of the displacement of the partition number

- As the key, the hierarchical structure account/container/object uses the MD5 hash algorithm to obtain a hash value. Perform an operation of shifting the first four bytes of the hash value to the right to obtain the partition index number. The shift value is specified by part_shift. Search all device numbers corresponding to the partition where the object is located in the mapping table (replica2part2dev_id) according to the partition index. These devices are deployed in different zones. A zone is only an abstract concept. It can be a machine, rack, even a cluster in a building. To provide the highest level of redundancy, deploy at least five zones. The weight parameter is a relative value and can be adjusted according to the disk size. A larger weight parameter indicates that more space can be allocated and more partitions can be deployed.
- Swift defines the ring for the account, container, and object, respectively. Searching for account and container follows the same process.

Ring

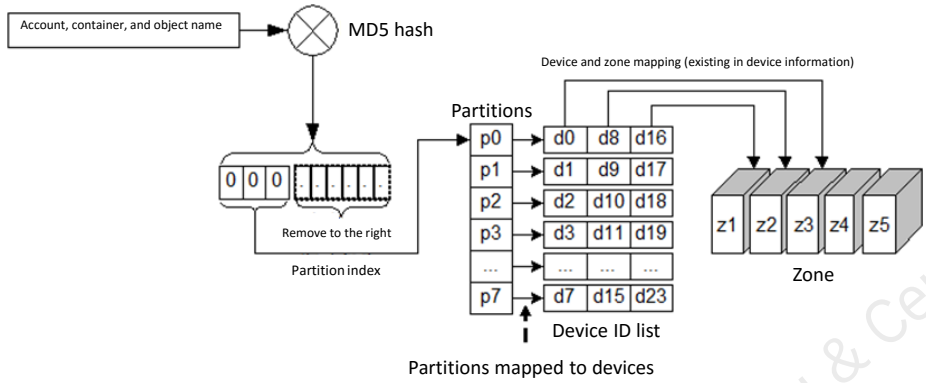

System Architecture

- Symmetric and resource-oriented distributed system architecture design

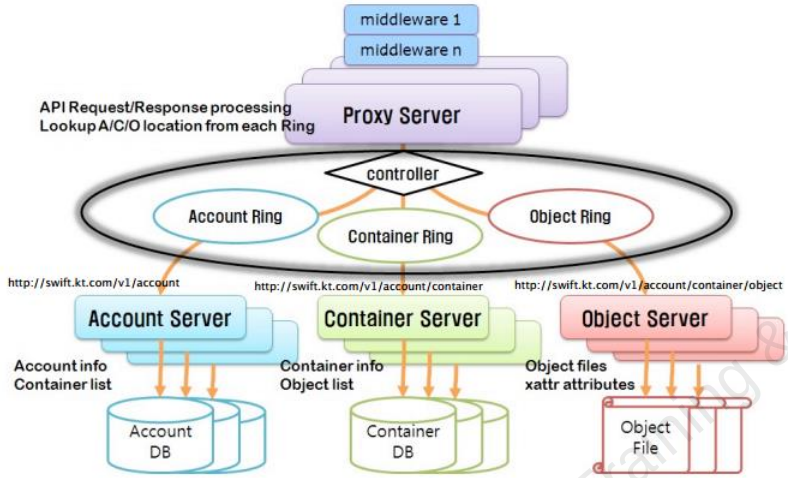

Swift Components (1)

- Proxy Server
  - Provides the object service APIs externally. Due to the stateless REST request protocol, horizontal expansion can be performed to balance the load.
- Account Server
  - Provides account metadata and statistics and maintains the service of the container list. Information of each account is stored in a SQLite database.
- Container Server
  - Provides container metadata and statistical information and maintains the service of the object list. Information of each container is also stored in a SQLite database.
- Object Server
  - Provides object metadata and the content service. The content of each object is stored as a file in a file system. Metadata is stored as a file attribute. It is recommended that you use the XFS file system whose attributes can be extended.

Swift Components (2)

- Replicator
  - Checks whether a local partition copy is consistent with a remote copy. If they are inconsistent, push is used to update the remote copy and ensure that the object to be deleted is removed from the file system.
- Updater
  - If an object cannot be updated immediately due to high load, the task is sequenced in the local file system so that the task can be updated asynchronously after the service is restored.
- Auditor
  - Checks the integrity of the object, container, and account. If a bit-level error is found, the file will be isolated, and other copies will be copied to overwrite the locally damaged copies. Other types of errors are recorded in logs.
- Account Reaper
  - Removes the account to be deleted, and deletes all containers and objects contained in this account.

- Updater: If an object cannot be updated immediately due to high load, the task is sequenced in the local file system so that the task can be updated asynchronously after the service is restored. For example, after an object is created, the container server does not update the object list. In this case, the container update operation is queued in, and Updater scans the queue and updates the object list after the system is restored.

Logical Architecture of Swift

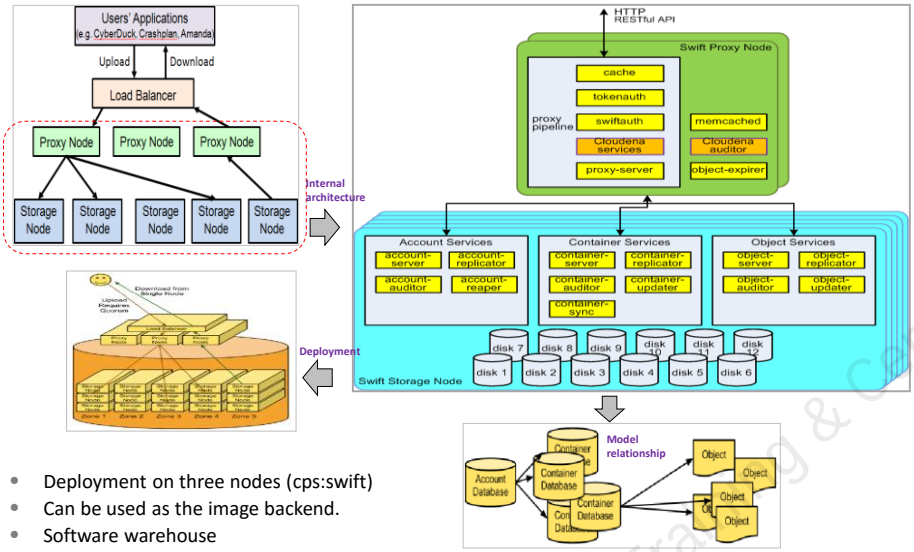

- Deployment on three nodes (cps:swift)
- Can be used as the image backend.
- Software warehouse

Contents

1. OpenStack Overview
2. OpenStack Component – Nova
3. OpenStack Component – Cinder
4. OpenStack Component – Neutron
5. OpenStack Component – Swift
6. **OpenStack Component – Glance**
7. OpenStack Component – Keystone

Huawei Training & Certification

Position of Glance in OpenStack

- Glance provides the common image service capability and image/VM snapshot management function for instance creation.
- Glance is one of the eight core components of OpenStack and is one of the core projects of OpenStack.

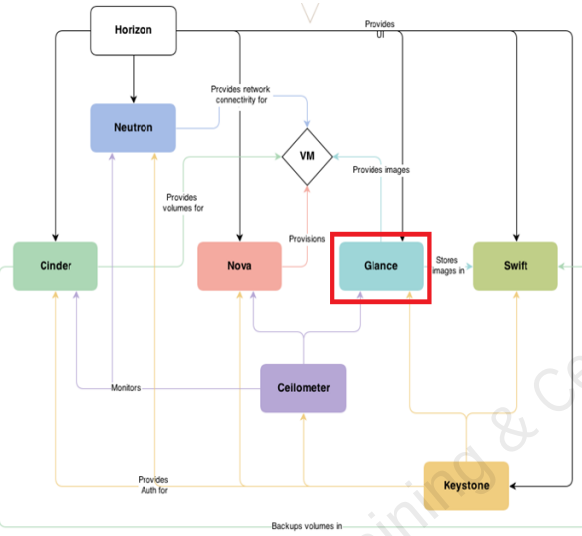

Glance – Image Module

- Glance is used as an independent large-scale image search service. When being used with Nova and Swift, Glance provides the VM image search service for OpenStack. Similar to all OpenStack projects, Glance complies with the following principles:
  - Component-based architecture: Adds new features conveniently.
  - HA: Supports large load.
  - Error tolerance: Independent process address space avoids serial errors.
  - Open standards: Provides reference for APIs driven by the community.

Glance – Image Management

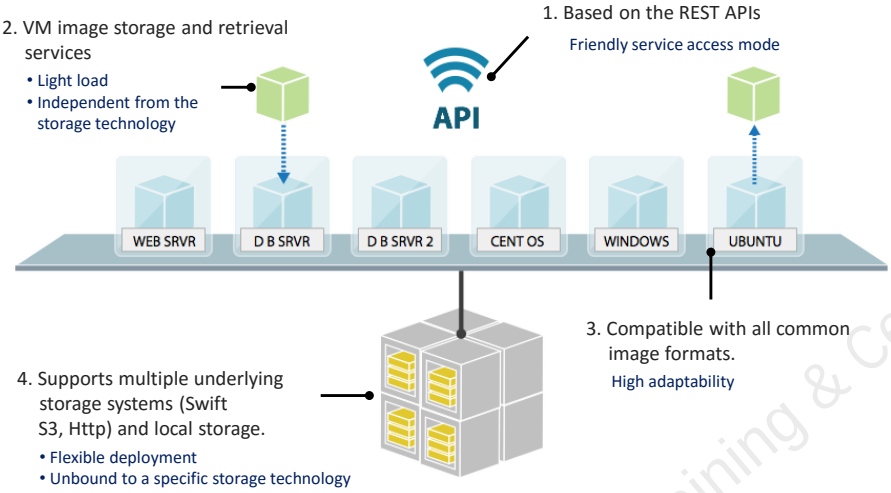

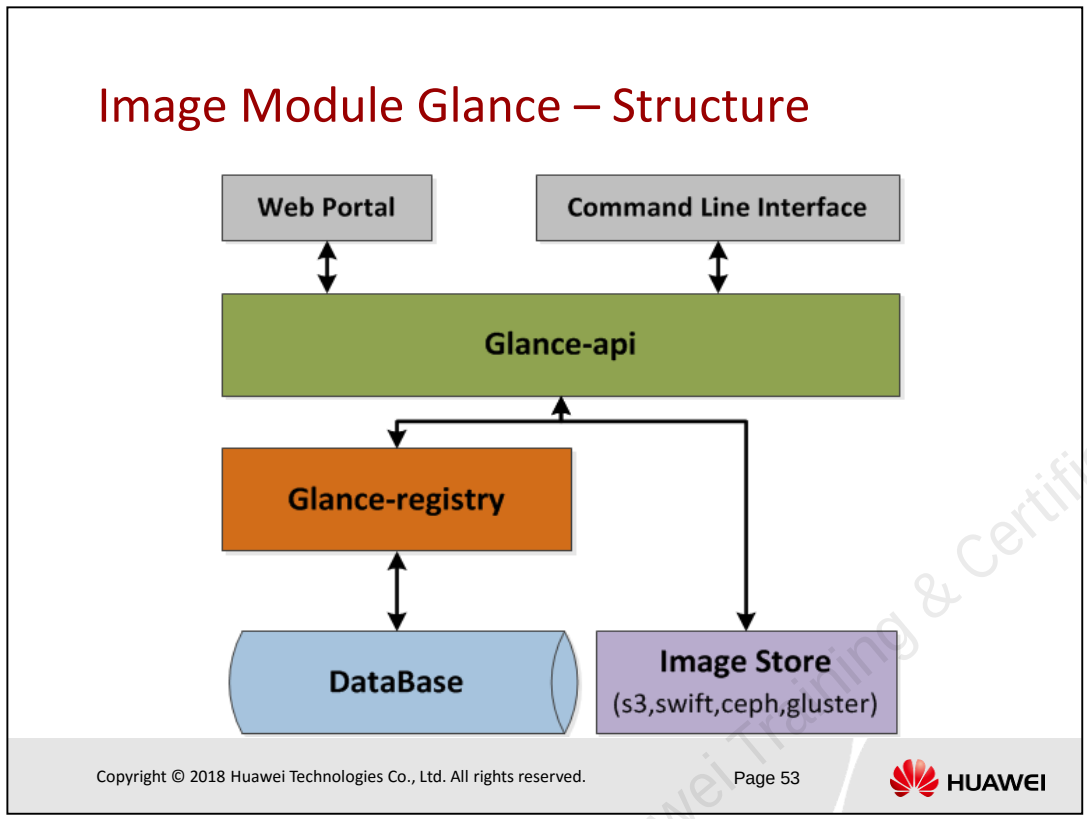

- Glance architecture
  - Glance-api: Similar to nova-api, it receives the REST API request, and uses other modules (glance-registry and image store) to complete operations such as querying, obtaining, uploading, and deleting an image.
  - Glance-registry: Interacts with the MySQL database to store and obtain the image metadata.
  - Image store: An interface layer of storage. With this interface, Glance can obtain images. Image store supports Amazon S3, Swift of OpenStack, and distributed storage such as ceph, sheepdog, and GlusterFS.
- Cache mechanism: Cache the required image file to the target node before the VM creation. During the VM creation, the system searches the image from the cache. If the required image can be found in the cache, the system will not download the image file from Glance. The system uses the image in the local PC to create the disk file of the VM, greatly reducing the VM deployment time.

Glance – Image Management

Image service and storage: Logical architecture

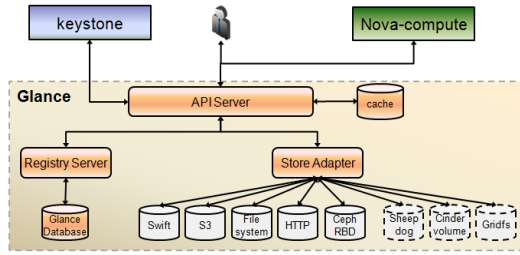

- Image service and storage: Local cache

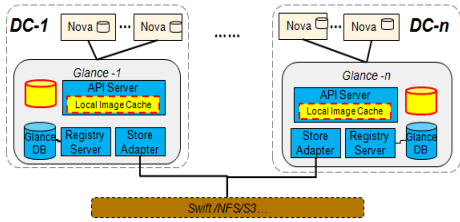

- Image service and storage: Multi-service (multi-data center) replication

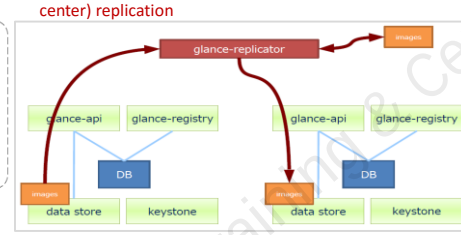

- Unified service interface and private and open images
- Supports multiple storage backends through the driver and supports isolation of image storage space between tenants (only Swift).
- Supports comprehensive image formats.
- New features such as local cache of images, multiple data center (MDC) replication, multiple locations, incremental images, and Cinder-Volume backend simplify image management during large-scale MDC deployment, improve the storage space usage and VM provisioning efficiency, and support image backup.

Introduction to Multiple Image Formats (1)

- Glance supports the following image formats:
  - raw: Unstructured image format
  - vhd: Universal VM disk format, which can be used for VMware, Xen, Microsoft Virtual PC/Virtual Server/Hyper-V, and VirtualBox
  - vmdk: VMware VM disk format that multiple hypervisors support
  - vdi: VM disk format that VirtualBox and QEMU support
  - iso: Archive format of the CD-ROM archive format
  - qcow2: Disk format that can be supported by QEMU and can be dynamically expanded
  - aki: Amazon Kernel image
  - ari: Amazon Ramdisk image
  - ami: Amazon VM image

Introduction to Multiple Image Formats (2)

- The image container format (or Package VM disk format) is similar to a folder. In Swift, it is called Container. Glance supports multiple packed Container formats.
  - ovf: Open VM disk format, which is initiated by VMware and has been supported by multiple virtualization devices.
  - bare: indicates that an image does not contain a container or metadata.
  - aki: Amazon Kernel image
  - ari: Amazon Ramdisk image
  - ami: Amazon VM image

Contents

1. OpenStack Overview
2. OpenStack Component – Nova
3. OpenStack Component – Cinder
4. OpenStack Component – Neutron
5. OpenStack Component – Swift
6. OpenStack Component – Glance
7. **OpenStack Component – Keystone**

Position of Keystone in OpenStack

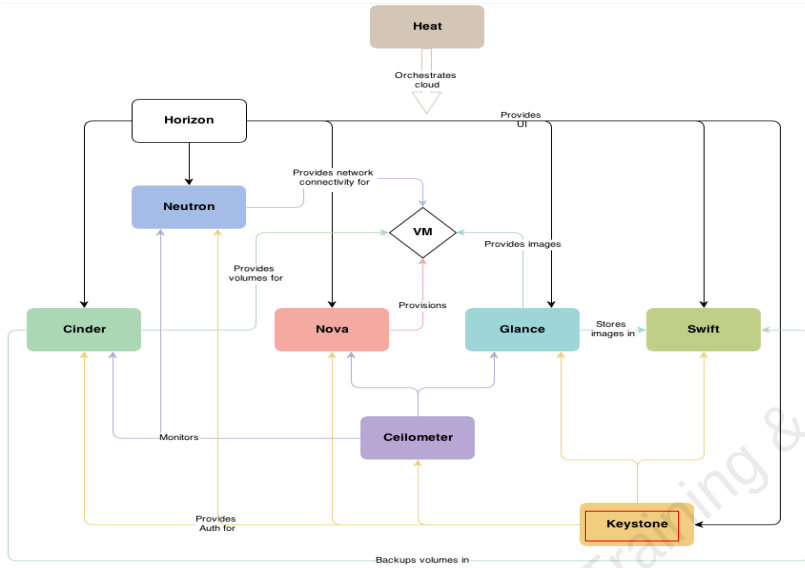

Keystone Overview

- Provides the centralized authentication service, including user authentication and OpenStack service authentication.
- Common terms:
  - User: OpenStack user
  - Role: A collection of rights, which are associated with users to implement access control
  - Group: A collection of users with the same rights. If you set the role for a group, you set the same role for all users in the group.
  - Project or tenant The term tenant is used in previous versions. In the Neutron version, it is called project, which refers to a collection of resources allocated to users.
  - Domain: Defines the management boundary.
  - Service endpoint: URL of a service, which is used to search for services by users

Keystone Object Model (1)

- Region: represents a data center in Keystone.
- Policy: Access control policy, which defines API access control rules
- Assignment: refers to a triplet (actor, target, and role) in which the actor contains a user and a group, and the target contains a domain and a project. Each assignment represents a permission assignment operation.
- Token: Service access credentials of users. It represents user account information and contains user, scope information (project, domain, and trust), and role information.

Keystone Object Model (2)

- User, project, domain, role, and assignment:

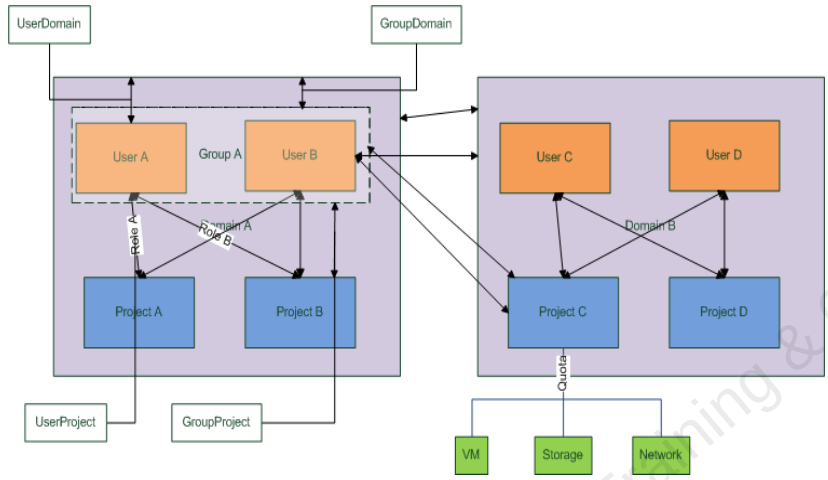

Keystone Object Model (3)

- Region, service, and endpoint:

Introduction to Token

- Token
  - After a user provides a group of valid user information to Keystone, Keystone returns a token to the user.
  - A token contains the user information, user role, token scope, and token validity period.
  - The user can use this token to access other OpenStack services.
- Why token is used?
  - It is easier to use the token to access the API service than using the username and password.
  - The token has a validity period. It is safer to cache the token on the client than caching the username and password.

Application of Token in OpenStack

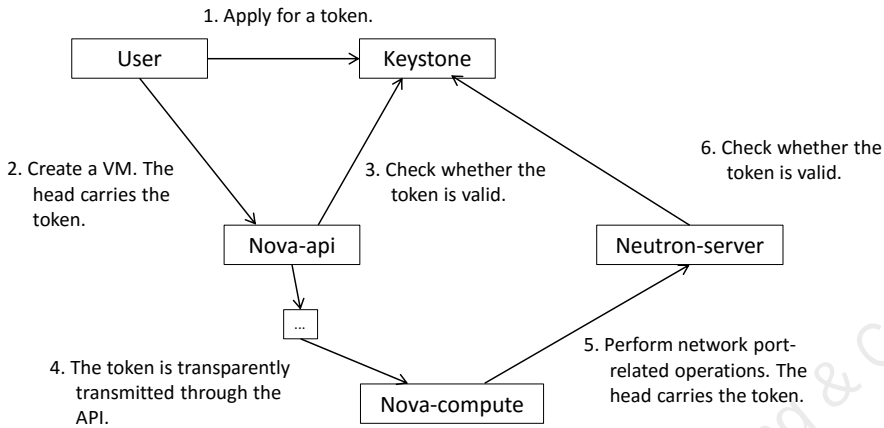

Recommendations

- Huawei Learning website
  - http://support.huawei.com/learning/en/newindex.html
- Huawei Support case library
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Cloud Data Center OpenStack Architecture and Principles (Volume 2)

www.huawei.com

Foreword

- To begin with, this course describes the concept of FusionSphere OpenStack.
- Then, this course gives an introduction to mandatory and optional modules of FusionSphere and describes each module. You will have a basic understanding of the overall FusionSphere framework.
- Finally, this course introduces and describes data protection and service continuity disaster recovery (DR) products and solutions for FusionSphere.

Objectives

- After finishing this course, you will be able to:
  - Understand FusionSphere.
  - Have a good command of FusionSphere components and the overview of each component.
  - Obtain FusionSphere compatibility and specifications information.
  - Master FusionSphere DR products and solutions.

Contents

1. **FusionSphere Concept**
2. FusionSphere Versions
3. FusionSphere Components and Description
4. FusionSphere Compatibility and Specifications
5. FusionSphere Cloud Data Center Deployment
6. FusionSphere DR Solution

Introduction to FusionSphere

- Concept
- A Huawei-developed cloud operating system (OS) designed for multiple industries
- Dedicated for cloud design and optimization
  - Provides powerful virtualization functions and resource pool management.
  - Abundant cloud basic service components and tools
  - Open APIs
- Horizontally integrates physical and virtual resources in data centers (DC) and vertically optimizes the service platform, simplifying cloud computing deployment and usage for enterprises.

Compared with VMware

	FusionSphere	VMware
Virtual infrastructure layer	FusionCompute CNA	VMware ESXi
Virtual resource management layer	FusionCompute VRM	VMware vCenter Server
Cloud resource management layer	FusionManager	VMware vCloud Director

Contents

1. FusionSphere Concept
2. **FusionSphere Versions**
3. FusionSphere Components and Description
4. FusionSphere Compatibility and Specifications
5. FusionSphere Cloud Data Center Deployment
6. FusionSphere DR Solution

Existing FusionSphere Version Information

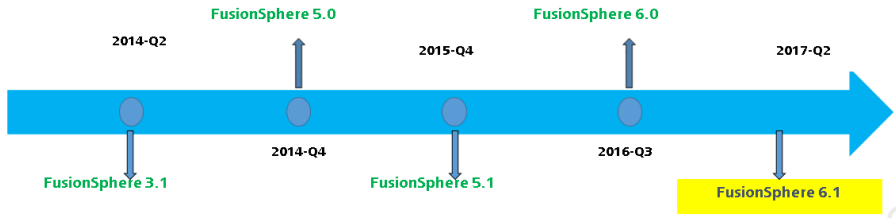

Contents

1. FusionSphere Concept
2. FusionSphere Versions
3. **FusionSphere Components and Description**
4. FusionSphere Compatibility and Specifications
5. FusionSphere Cloud Data Center Deployment
6. FusionSphere DR Solution

FusionSphere Components

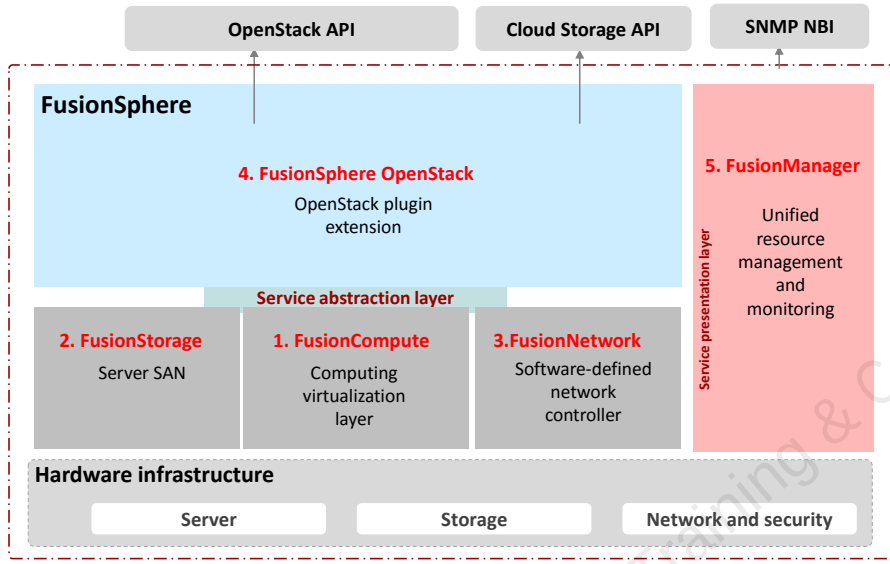

OpenStack Modules Realizes a Unified Model in the Virtualization Environment.

- OpenStack is introduced to FusionSphere to implement unified resource abstraction, management, and allocation in the heterogeneous virtualization environment.

OpenStack plugin extension

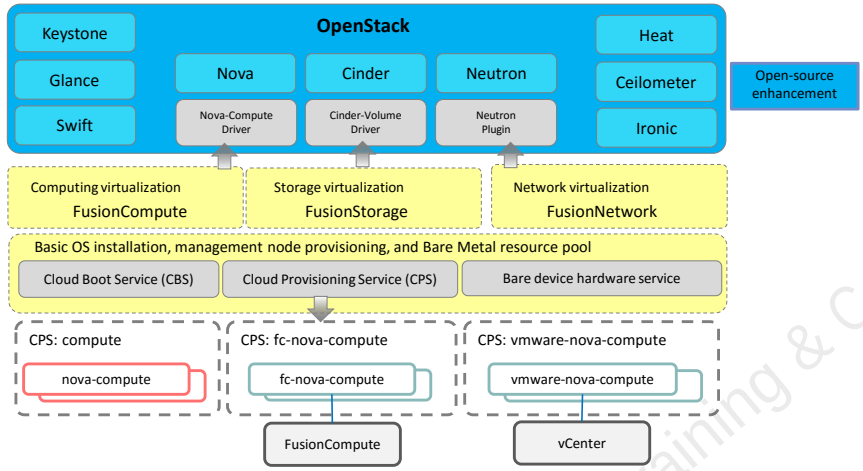

FusionCompute Compute Device Virtualization

- FusionCompute is a mandatory function module used for establishing a virtual environment.

  Relationships between FusionCompute modules

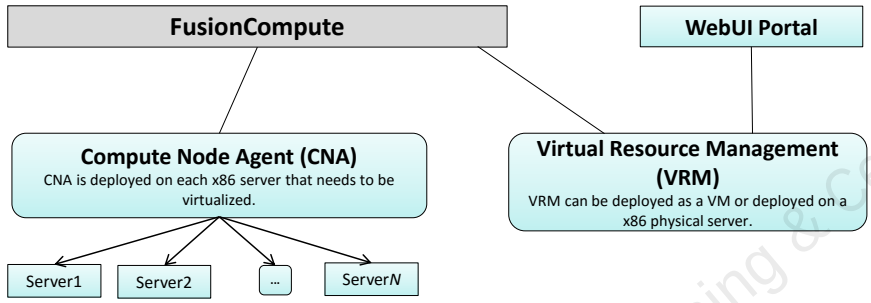

- CNA is deployed on a server that needs to be virtualized.
- VRM can be deployed as a VM or deployed on a physical server. VRM externally provides the web operation page for management and maintenance personnel.

FusionStorage Server SAN (Optional Module)

- Huawei FusionStorage is software-defined distributed block storage software.
- FusionStorage provides high-scalability, high-performance, and high-reliability block storage for FusionSphere, VMware, and physical database environments. It can be purchased and used independently. It is optional for building a FusionSphere environment.

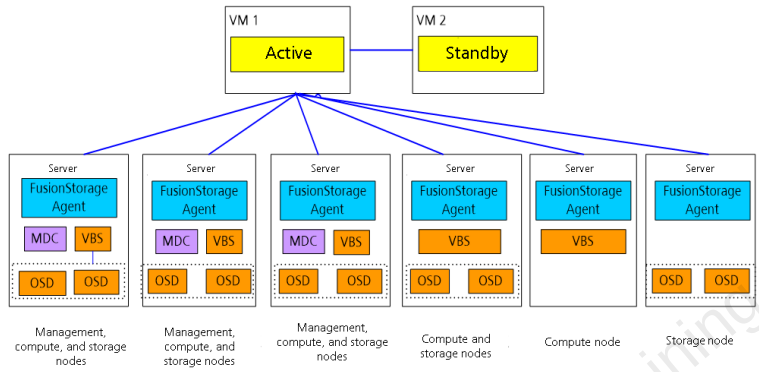

FusionNetwork SDN Controller (Optional Module)

- Huawei software-defined network function module
- FusionNetwork is a component used to set up or use advanced network functions and flexibly configure and manage network functions. It is optional for building a FusionSphere environment.

Relationships between FusionNetwork modules

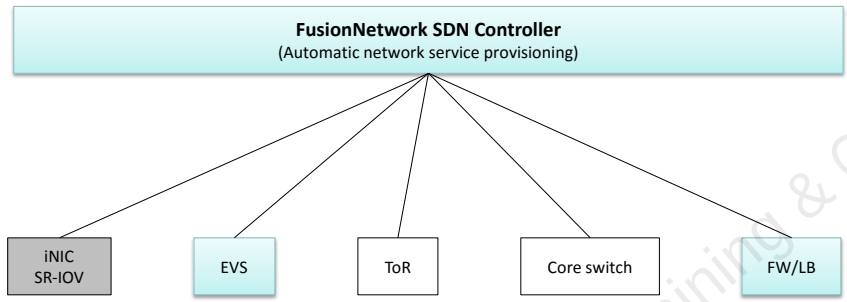

- FusionNetwork will be implemented using Neutron of FusionSphere OpenStack to access Agile Controller of Huawei SDN products.

Unified Management of FusionSphere OpenStack OM Resources

- Huawei FusionSphere OpenStack OM centrally manages and monitors resources in the cloud computing environment.

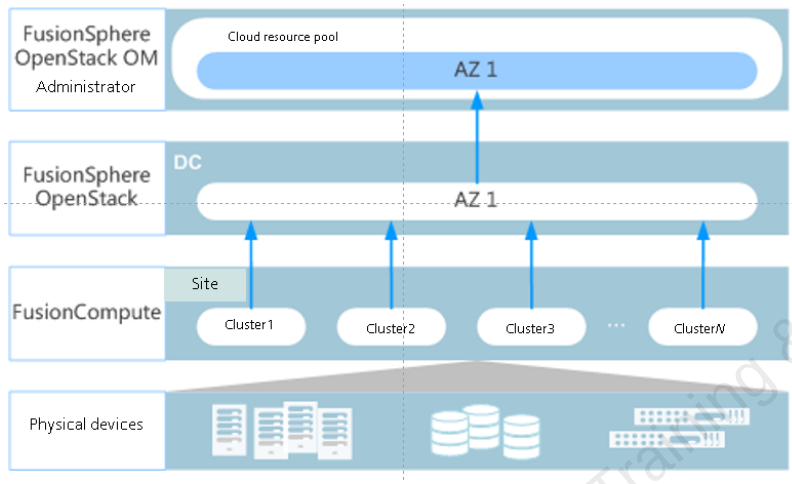

Contents

1. FusionSphere Concept
2. FusionSphere Versions
3. FusionSphere Components and Description
4. **FusionSphere Compatibility and Specifications**
5. FusionSphere Cloud Data Center Deployment
6. FusionSphere DR Solution

FusionSphere Compatibility – Hardware&Software

- Hardware compatibility includes:
  - Server
  - Storage device
  - Network devices
  - I/O devices, such as NICs, RAID cards, HBA cards, and GPU
  - Guest OS, LinuxOS Family, WindowsOS Family
- Software compatibility includes:
  - FusionSphere application compatibility list
    - Database software, middleware, web server software, email, HA, DR, security software, and enterprise applications
  - FusionSphere OpenStack application compatibility list
    - Database software, web server software, file server software, DNS software, and ICT applications

- For details about hardware compatibility, visit: http://support.huawei.com/onlinetool/datums/fusioncloud/comptool/index.en.jsp
- No link is available for querying software compatibility. If you want to query it, contact Huawei technical support or visit:
- http://support-open.huawei.com/ready/index.jsf
- For details about more Huawei product compatibility, visit:
- http://support-open.huawei.com/ready/index.jsf

FusionSphere Specifications (1)

Management Capability	Specifications	Specifications of Physical Hosts	Specifications
Maximum number of hosts supported by a single VRM	1024	Maximum logical CPU cores of each physical host	512
Maximum sites supported by a VRM domain	16	Maximum physical memory of each physical host	16 TB
Number of host clusters supported by a single VRM	32	Maximum number of VMs on each physical host	1024
Maximum number of VMs supported by a single VRM	10,000 running VMs or 30,000 registered VMs	Maximum number of LUNs mounted to each physical host	512
Number of vSwitches supported by each VRM	128	Maximum volume supported by each physical host	2048
Number of hosts supported by a single logical cluster	128 (LUN storage)/ 64 (virtual storage)	Number of NUMA nodes on each physical host	16
Number of VMs supported by a single logical cluster	3000	Maximum concurrently live migrated VMs per host	8
Maximum number of physical machines	4096		
Maximum number of VMs	80,000		

FusionSphere Specifications (2)

VM Specifications	Specifications	Storage Capacity Specifications	Specifications
Number of vCPUs supported by a single VM	128	Maximum number of volumes supported by a VRM	20000
Number of NICs supported by a single VM	12	Maximum number of hard disks supported by distributed storage	2000
Number of virtual disks on a VM	60	Maximum number of snapshots supported by the system	48000
Memory supported by a single VM	6 TB	**Network capacity specifications**	**Specifications**
Maximum virtual disk capacity of a VM	64 TB	Maximum number of distributed switches supported by the system	128
Maximum number of snapshots for a VM	32	Maximum number of hosts managed by a distributed switch	1024
		Maximum number of port groups supported by a single distributed switch	4096

Contents

1. FusionSphere Concept
2. FusionSphere Versions
3. FusionSphere Components and Description
4. FusionSphere Compatibility and Specifications
5. **FusionSphere Cloud Data Center Deployment**
6. FusionSphere DR Solution

FusionSphere in Cloud Data Center Scenarios

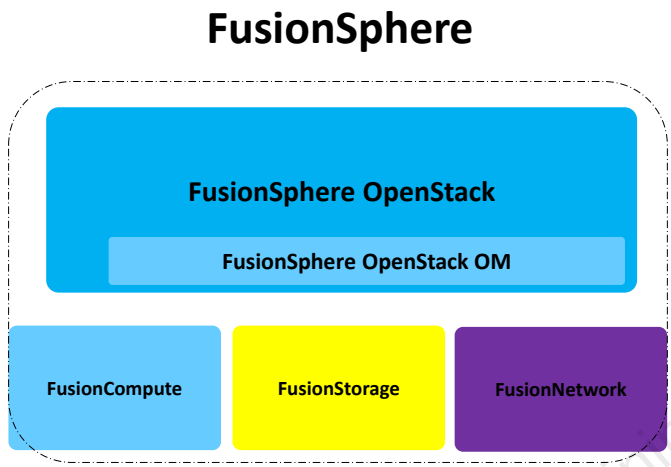

- FusionSphere OpenStack OM, FusionCompute, and OpenStack are delivered to customers as independent products.

Typical Deployment of FusionSphere in Cloud DC Scenarios (1)

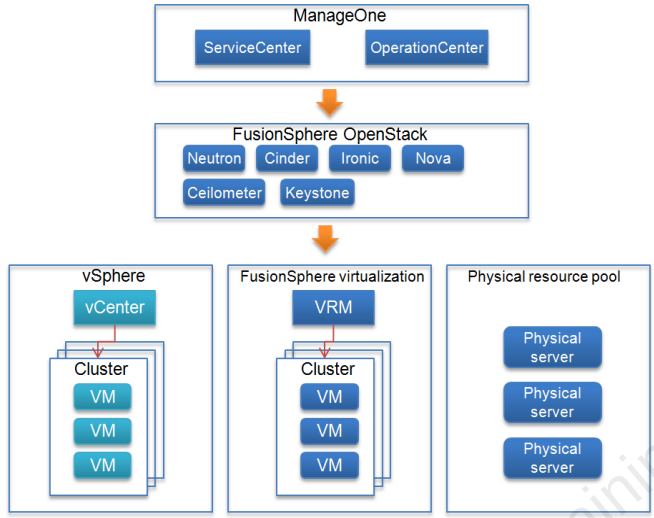

- Huawei private cloud and hybrid cloud solutions use this architecture.

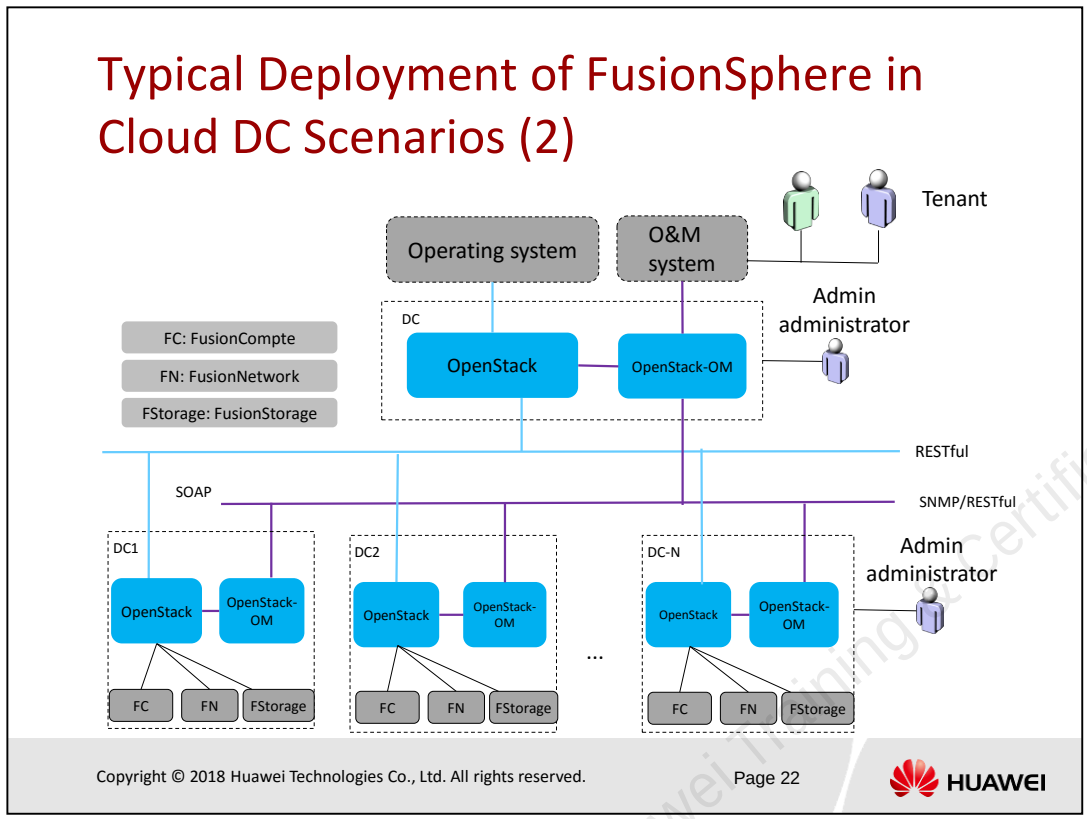

- In DC, OpenStack-OM supports local O&M and performance, capacity, monitoring, and alarm management.
- In cascading OpenStack, OpenStack-OM implements centralized O&M, supports global resource configuration, and centrally manages global performance, capacity, monitoring, and alarm.
- HUAWEI CLOUD uses this architecture.

Contents

1. FusionSphere Concept
2. FusionSphere Versions
3. FusionSphere Components and Description
4. FusionSphere Compatibility and Specifications
5. FusionSphere Cloud Data Center Deployment
6. **FusionSphere DR Solution**

eBackup Products: Huawei Data Protection

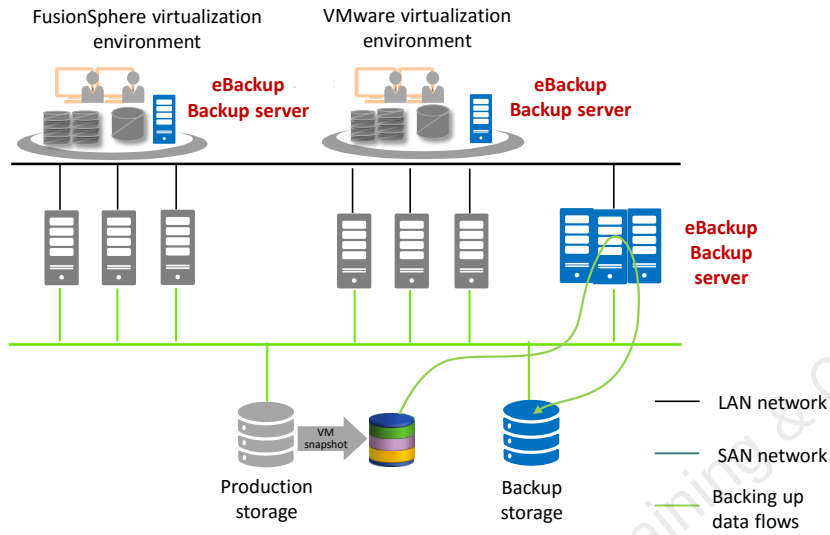

BCManager Product: Huawei Service Continuity DR

OceanStor BCManager supports unified management of multiple DR scenarios. It provides one-click DR test and recovery for Oracle, IBM DB2, Microsoft SQL Server database, Microsoft Exchange Server applications, NTFS file system, NAS file system, VMware, and FusionSphere virtualization environment.

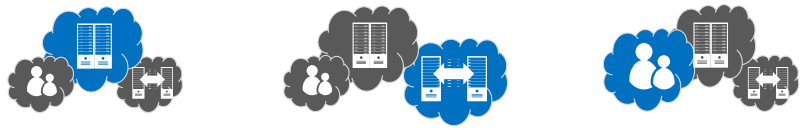

Service integration		Management automation		DR visualization	
Automatic application identification	Consistency of application data	One-click management	Customized process	Visualized topology	DR process monitoring
Virtualized environment protection	Policy-based protection	Automatic test	Automatic recovery	Statistics report	RPO/RTO monitoring

BCManager supports the following functions: local protection, local HA, active/standby DR, active-active DC, and geo-redundant DR solution.

HUAWEI

- OceanStor BCManager is management software developed by Huawei for protection and DR of host applications and VMs.
- Application DR is based on value-added functions of array storage devices and is compatible with multiple platforms including Windows, Linux, AIX, and HP-UX. Application DR provides data consistency protection, application protection, and application DR for hosts and application software on the hosts on these platforms.
- Virtualization DR protects VM in virtualization environments, such as FusionSphere (in storage-based replication or host-based replication mode) and vSphere hypervisors (in host-based replication mode).

Quiz

1. Which of the following modules are included in the overall architecture of Huawei FusionSphere? ( )
   A. FusionComopute
   B. FusionStorage
   C. FusionNetwork
   D. FusionCloud

- 1: BC

Recommendations

- Huawei Learning website
  - http://support.huawei.com/learning/en/newindex.html
- Huawei Support Case Library
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Cloud Data Center OpenStack Deployment (Including OM)

www.huawei.com

Huawei Training & Certification

Foreword

- This mainly describes Cloud Data Center FusionSphere architecture, and installation and deployment of FusionSphere OpenStack.

Page 2

Objectives

- Upon completion of this course, you will be able to:
  - Learn how to install FusionSphere OpenStack
  - Understand FusionSphere OpenStack deployment process
  - Learn how to configure FusionSphere OpenStack
  - Learn about FusionSphere OpenStack OM installation and configuration

Page 3

Contents

1. **FusionSphere Introduction**
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. Configuring FusionSphere OpenStack
5. Installing and Configuring FusionSphere OpenStack OM

Huawei Training & Certification

FusionSphere logical Architecture

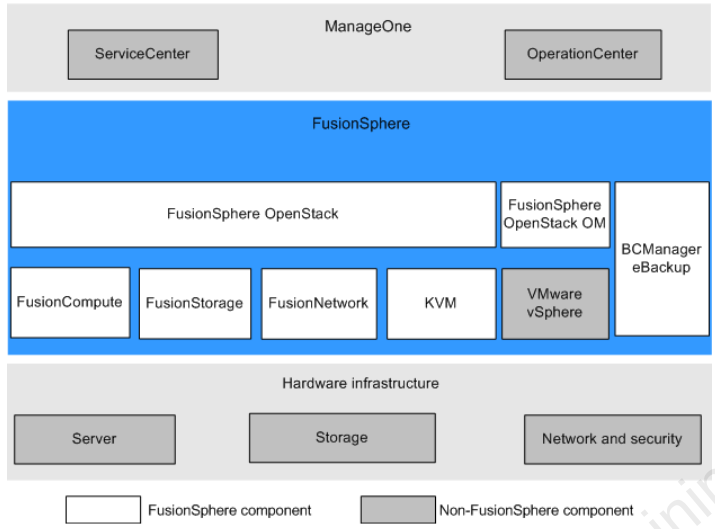

Key Components in the FusionSphere Solution

Component	Description
FusionCompute	Optional. 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 virtual network technologies.
FusionStorage	Optional. FusionStorage is a piece of distributed storage software that has both storage and computing capabilities. It can be deployed on multiple general x86 servers to combine the local disks on all the servers into a virtual storage resource pool to provide the block storage function.
FusionNetwork	Mandatory. FusionNetwork is solution designed for traditional virtual switching to evolve to forward-thinking software-defined networking (SDN). FusionNetwork leverages Virtual Extensible LAN (VXLAN) layer 2 tunnel encapsulation technologies and Huawei SDN controllers to automatically deploy and configure SDN networks, control service level agreement (SLA), and isolate multiple tenants. NOTE: FusionNetwork provides network functions based on the Neutron service in the FusionSphere OpenStack system.

Key Components in the FusionSphere Solution

Component	Description
FusionSphere OpenStack	Mandatory. Huawei FusionSphere OpenStack is an enhanced enterprise edition developed based on OpenStack Mitaka. FusionSphere OpenStack centrally schedule and manages virtual computing, storage, and network resources over unified RESTful interfaces. It also reduces operating expense (OPEX) and provides high system security and reliability, helping enterprises build secure, energy-saving data centers.
KVM	Optional. The server OS delivered with the FusionSphere OpenStack software packages provides the KVM virtualization function used for deploying KVM compute nodes in the FusionSphere system.
FusionSphere OpenStack OM	Mandatory. FusionSphere OpenStack OM is a piece of operation and maintenance (O&M) software for FusionSphere OpenStack. FusionSphere OpenStack OM monitors and manages hardware and software deployed on the cloud platform, supports system O&M, and provides a web user interface (UI) for administrators to perform operations and manage resources.
BCManager eBackup	Optional. BCManager eBackup (eBackup for short in follow-up sections) is a piece of backup software that provides end-to-end backup solutions for all virtualized applications in the FusionSphere system.

Page 7

Contents

1. FusionSphere Introduction
2. **Installing FusionSphere OpenStack**
   - Deployment Plan and Installation Process
     - Preparing for Installation
     - Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. Configuring FusionSphere OpenStack
5. Installing and Configuring FusionSphere OpenStack OM

Deployment Plan

- Different FusionSphere management system deployment plans are provided, varying depending on the FusionSphere management scale.

3 to 20 managed hosts and 1 to 100 managed VMs

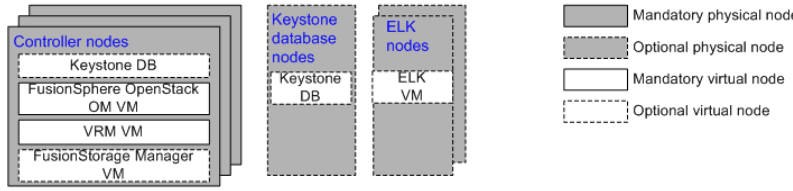

20 to 50 managed hosts and 1 to 1000 managed VMs

Physical Node Type	Quantity (Unit)	Component	Specifications
Controller nodes	3	•OpenStack management component •FusionSphere OpenStack OM VM •VRM VM •FusionStorage Block Manager VM •fc-nova-compute connected to the FusionCompute cluster •vmware-nova-compute connected to the vCenter clusters	•CPU: • Recommended: 20 cores • Minimum: 12 cores •Memory: • Recommended: 192 GB • Minimum: 106 GB •Network interface card (NIC): 2 x 10 Gbit/s (recommended)

Software Architecture

- A host refers to a physical server with the FusionSphere OpenStack system deployed. A host can be assigned different roles to serve as a controller node or compute node.
- A storage resource can be a storage unit provided by an OBS device, or the host local storage.

Node	Description
Controller node	A controller node has the following roles assigned: sys-server, controller, and measure.
Compute node	A compute node has the compute role assigned, and it is not a controller node.

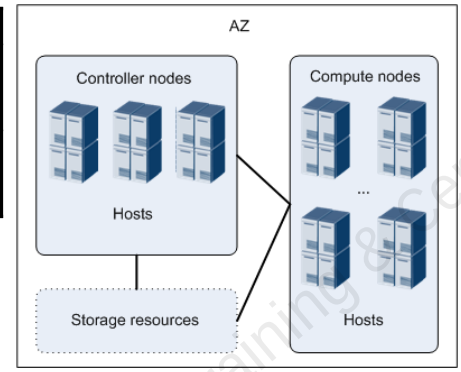

Objects

- Component
  - Component is the minimum object for installation and deployment. A component can be independently started, restarted, and configured, and its status can be monitored.
- Service
  - Service is a piece of distributed software that can provide independent functions, be upgraded, and be patched. A service is a collection of a group of components.
- Role
  - Role is a concept introduced by Huawei for service installation. A role defines a group of hosts that have the same property or function in a template. For example, the compute role is a collection of templates that provide the management service for compute resources. You can configure the range of the hosts that can be assigned a role, and then the CPS automatically selects all the qualified hosts to deploy all components the role supports. A host can be assigned multiple roles.

Roles (1/2)

- sys-server: provides the application programming interface (API) proxy and load balancing services.
- auth: provides the authentication function (shared within the system).
- image: provides the image management service (shared with DC).
- controller: provides the FusionSphere OpenStack management service.
- router: provides the layer 3 network, dynamic host configuration protocol (DHCP), and metadata services.
- compute: provides the computing virtualization function of the host in FusionSphere OpenStack.
- blockstorage-driver: provides the storage virtualization management service.
- blockstorage: provides the FusionStorage distributed storage service.
- swift: provides the object storage management service.

Roles (2/2)

- database: provides the FusionSphere OpenStack management data storage service.
- mongoDB: provides the FusionSphere OpenStack sampling data storage service.
- zooKeeper: provides the distributed cluster service.
- rabbitMQ: provides the distributed communication service.
- baremetal: provides the bare metal server management service.
- loadBalancer: provides load balancing capabilities for networks.
- sys-client: provides configuration proxy and component status monitoring capabilities.

Deployment Model - Cluster Mode

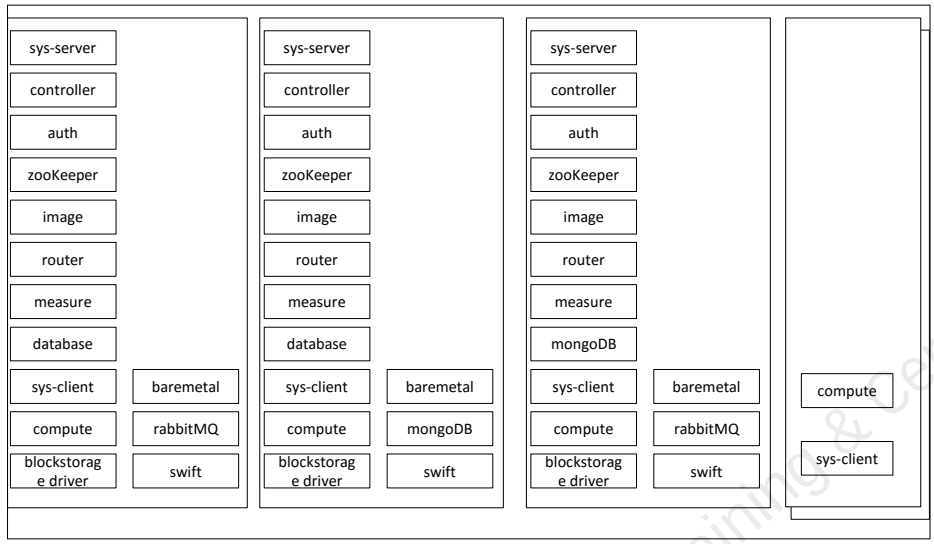

Deployment Model - Thin Mode

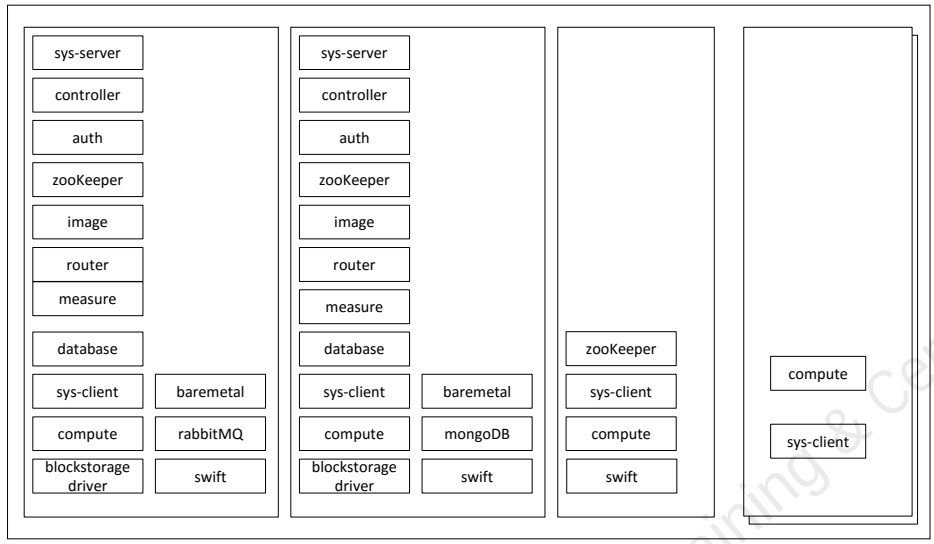

Installation Process

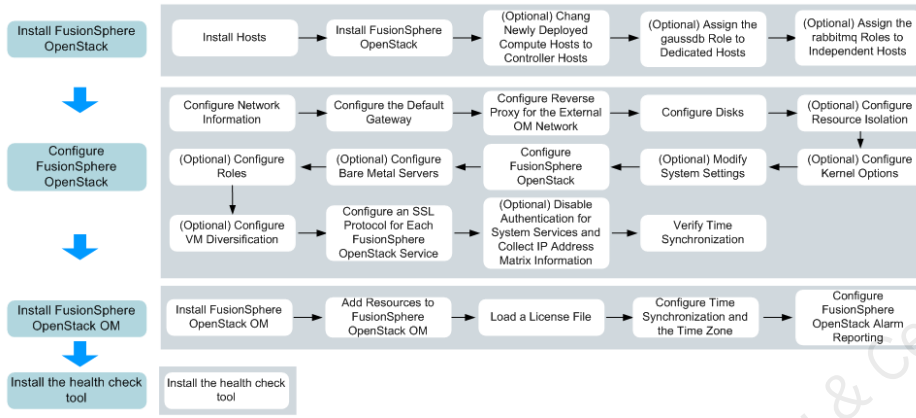

Page 16

Contents

1. FusionSphere Introduction
2. **Installing FusionSphere OpenStack**
   - Deployment Plan and Installation Process
   - Preparing for Installation
   - Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. Configuring FusionSphere OpenStack
5. Installing and Configuring FusionSphere OpenStack OM

Page 17

Windows PC Requirements

Item	Description
Login account	Only administrators can log in to the local PC. If logging in to the local PC as a guest, ensure that the account has administrator permission.
OS	Supported OSs: Windows 7 32-/64-bit, Windows Server 2008 32-/64-bit
Firewall	The Windows firewall must be disabled.
Network	If hosts are to be PXE installed, network requirements are as follows: •The local PC network ports are connected to the same switch as the management plane network ports of the hosts to be installed. •The local PC network ports must be assigned IP addresses in the temporary network segment planned for installation. •On the switch, the VLAN settings for the local PC network ports and the PXE ports must be consistent. The VLAN must be untagged. Set the PVIDs for the PXE ports to this VLAN ID. •If the local PC needs to access the host baseboard management controller (BMC) to enable you to monitor, manage, and perform remote operations on the host, ensure that the local PC can communicate with host BMC ports.
Hard disk	The partition in which the OS is to be installed has more than 20 GB of free space.
Tool	•A decompression tool has been installed on the PC. •PuTTY has been installed on the local PC for logging in to hosts. •The local PC has JAVA 1.7.0 and later installed.

Software

- FusionSphere OpenStack + FusionSphere OpenStack OM + FusionCompute software package list

Software Package	Description	How to Obtain
FusionSphere OpenStack V100R006C10SPC300.tar.gz	FusionSphere OpenStack software package	For carrier users, visit http://support.huawei.com/carrier/ and choose Software > Carrier IT > Cloud Computing > FusionCloud > FusionSphere > FusionSphere OpenStack > V100R006C10SPC100. For enterprise users, visit http://support.huawei.com/enterprise and choose Support > Software Download > Cloud Computing > FusionCloud > FusionSphere OpenStack > V100R006C10SPC100.
FusionSphere OpenStack V100R006C10SPC300_InstallTool.tar.gz	FusionSphere OpenStack installation tool for Windows OSs	
FusionSphere OpenStack V100R006C10SPC300.iso	FusionSphere OpenStack installation image	
FusionSphere OpenStack V100R006C10SPC300_OM.zip	FusionSphere OpenStack OM software package	
FusionCompute V100R006C10SPC101_Installer.zip	FusionSphere installation tool	For carrier users, visit http://support.huawei.com/carrier/ and choose Software > Carrier IT > Cloud Computing > FusionCloud > FusionSphere > FusionCompute > V100R006C10SPC101. For enterprise users, visit http://support.huawei.com/enterprise and choose Support > Software Download > Cloud Computing > FusionCloud > FusionCompute > V100R006C10SPC101.
FusionCompute V100R006C10SPC101_CNA.iso	Operating system (OS) for FusionCompute hosts	
FusionCompute V100R006C10SPC101_DC.zip	VRM VM template, used for installing the VRM node on a FusionSphere OpenStack VM	

- In order to ensure that the obtained software package can be used normally, the software package needs to be checked for completeness and correctness after obtaining

FusionSphere OpenStack Data

- Example of planned network data

Type	VLAN ID	Remarks
external_om	VLAN 4005	
external_api	VLAN 4004	
external_base	VLAN 4006	If the RabbitMQ service and the FNM service are connected over the External Base plane and FusionCompute, plan at least two IP addresses used for RabbitMQ and FNM.
internal_base	VLAN Untag(VLAN 4007)	默认为“172.28.0.0/20”

FusionSphere OpenStack OM Data

- lists the data required for installing FusionSphere OpenStack OM.

Parameter	Description	Example Value
Application Scenario	Specifies the delivery scenarios for FusionSphere OpenStack OM.	FusionSphere
Type	Specifies the installation mode. Specify the installation object (FusionSphere OpenStack OM or ServiceCenter) and the FusionSphere OpenStack OM type.	AllinOne OM
Install Service OM	Select yes.	yes/no
HA Mode	Specifies the deployment mode..	Active/Standby
System Scale	Specifies the scale of the VMs managed by FusionSphere OpenStack OM.	0~512VMs
CPU Binding	Specifies whether the FusionSphere OpenStack OM VM CPU is to be bound to the FusionSphere OpenStack OM server CPU.	Disabled
Cascading Mode	Specifies the FusionSphere OpenStack OM cascading type.	Standalone
API Float IP	Specifies the floating IP address of the External API network hosting the FusionSphere OpenStack OM VM.	192.168.100.90
OM Float IP	Specifies the floating IP address of the External OM network hosting the FusionSphere OpenStack OM VM.	

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
   - Deployment Plan and Installation Process
   - Preparing for Installation
   - Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. Configuring FusionSphere OpenStack
5. Installing and Configuring FusionSphere OpenStack OM

Page 22

Installation Guide

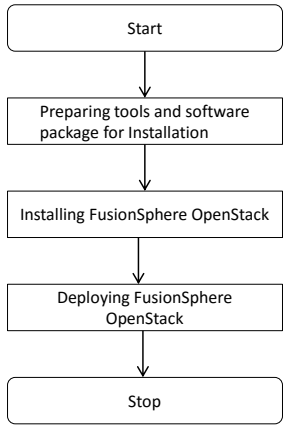

- **Preparing tools and software package for Installation**
  - PC
  - Install first node
- **Installing FusionSphere OpenStack**
  - Install using PXE: Use the Windows installation tool to PXE install the operating system for the hosts in FusionSphere OpenStack. In this way, you need to deploy the installation tool on the Windows operating system, the first node after the success of the installation, it will automatically install other hosts
- **Deploying FusionSphere OpenStack**
  - Use FusionSphere OpenStack to install and deploy the FusionSphere OpenStack on the host where the operating system is installed.

Host OS Installation Guidelines

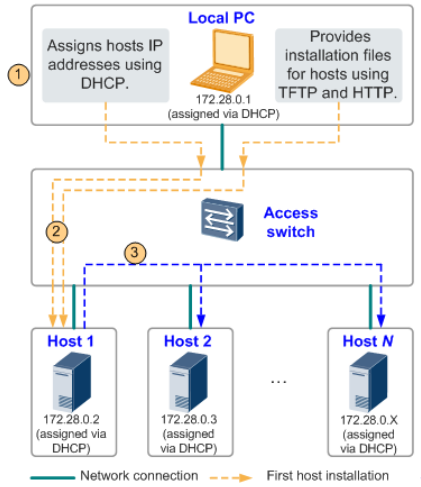

1. Get the installation tool and software packages ready.
   Deploy the installation tool on the local PC for it to work as the DHCP, TFTP, and HTTP server. Host 1 works as the client of the three services.
   Save the software packages to be installed on the local PC.
2. Install Host 1 using the local PC.
   After the Host 1 server is powered on, the local PC assigns Host 1 an IP address using DHCP so Host 1 can communicate with the local PC. Then you can upload the software packages to Host 1 using TFTP and HTTP to install OSs and basic services for it.
3. Install other hosts using Host 1.
   After the servers of other hosts are powered on, the hosts automatically connect to Host 1. Then Host 1 assigns these hosts IP addresses and provides them with the software packages for installing OSs and basic services on them.

Start the Host Installation Tool

- Decompress the installation tool package **FusionSphere OpenStack V100R006C10SPC300_InstallTool.tar.gz** to a root directory, for example, **D:**. If you decompress the package in a non-root directory, the installation may fail due to long absolute paths of some files in the tool package.
- In the windows folder, right-click **InstallFSP.exe** and choose **Run as administrator** to start the host installation tool. If **Run as administrator** is unavailable on your local PC, double-click **InstallFSP.exe** in the **windows** folder to start the host installation tool.

Select the Installation Tool language and Software Packages Used for the Installation

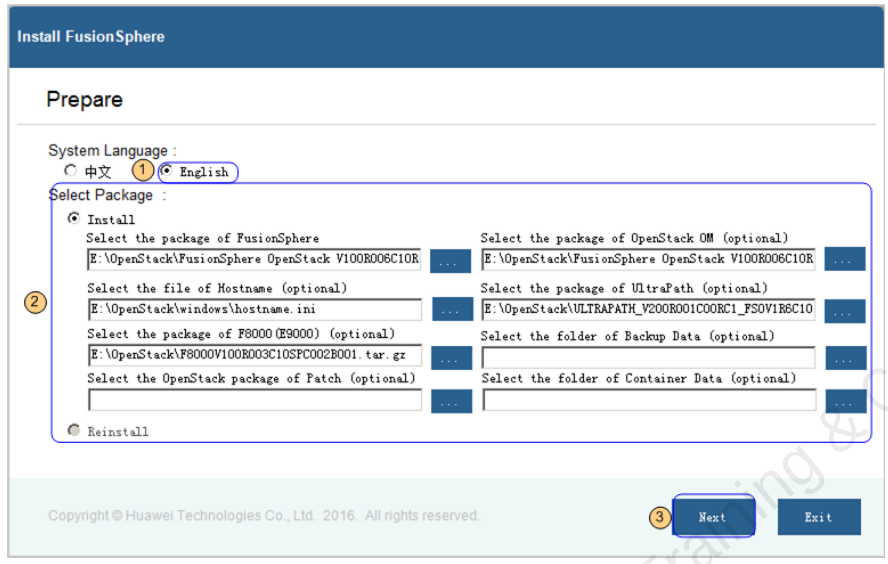

- Select the installation tool language.
- Select Install and select the required installation packages.
- Click Next to decompress the software package. After about 3 minutes, the system automatically switches to the System Config page when the decompress is complete.

Configure System Settings

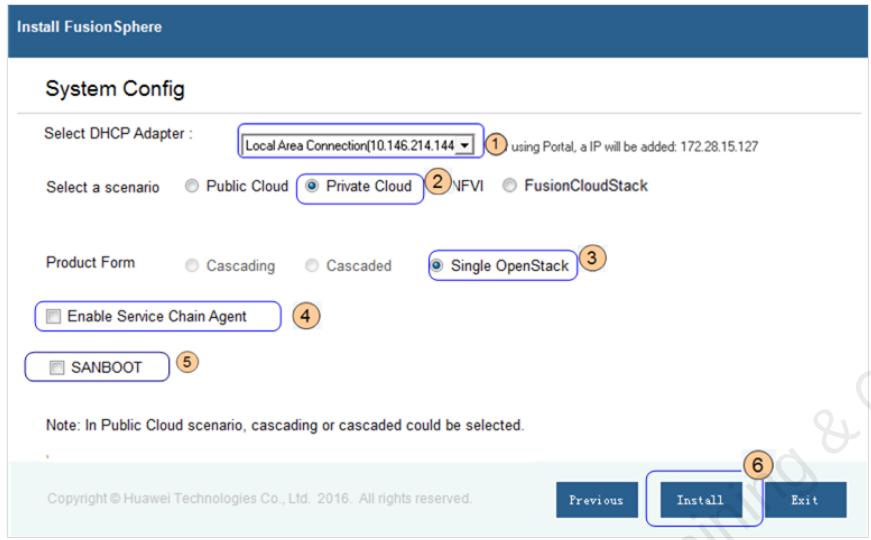

HUAWEI

- Select an IP address in network segment 172.28.X.X.
- Select a scenario: Private Cloud.
- Select the product form: Single OpenStack.
- Do not select Enable ServiceChain Agent.
- Determine whether to select SANBOOT.
- Click Install and perform 3 within 10 minutes. Otherwise, the installation may fail.

Install the First Node Using the local PC

- Configure the first host to boot from PXE, power on the host, and bootstrap its OS installation from your local PC, enabling it to work as the first node in the system.
- Log in to the iBMC interface of the RH server

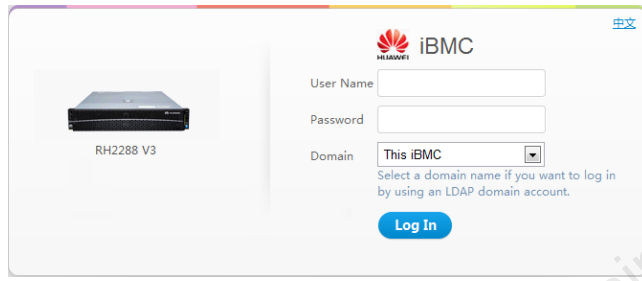

Install the First Node Using the local PC

- Login, select "Remote Control -> Remote Virtual Console (Shared Mode)", enter the virtual KVM, click , select restart the server or power on the server, press the F11 reboot into the PXE boot.

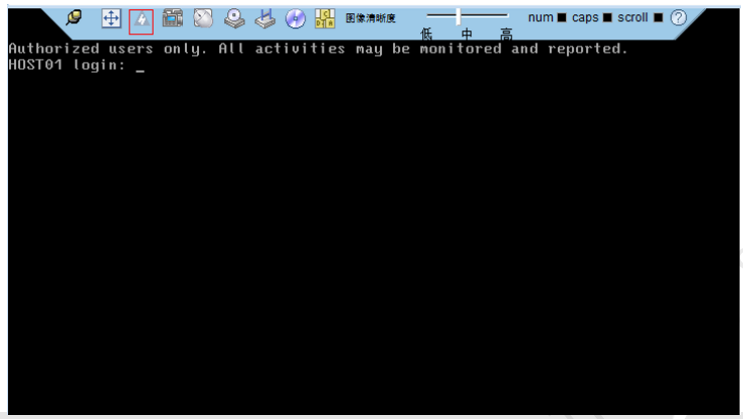

Install the First Node Using the local PC

- The installation takes about 10 to 20 minutes. After the first host is installed, the system displays the following page. Click Exit on this page.

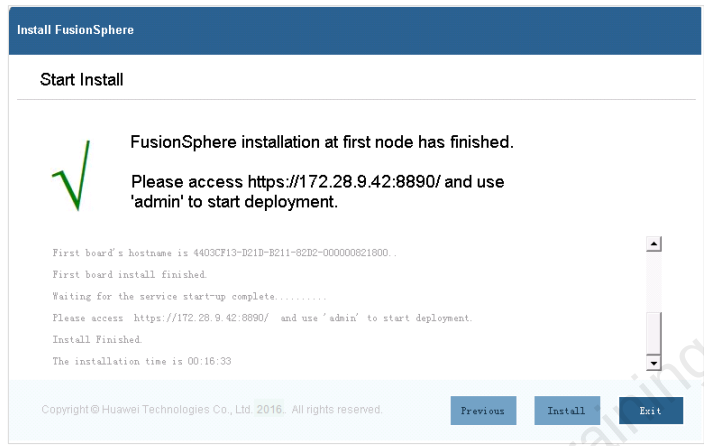

HUAWEI

- The installation takes about 10 to 20 minutes. After the first host is installed, the system displays the following page. Click Exit on this page.
- Note
  - To ensure system security, you are not allowed to directly log in to the installed host using an SSH connection as user root. You can use username fsp and its default password Huawei@CLOUD8 to log in to the host through an SSH connection. If you need to have root user permission for the host after logging in to it as user fsp, run su to switch to user root. The default password is Huawei@CLOUD8!.
  - To improve system security, it is recommended that you change the password of each user in a timely manner after installing the hosts. For details about password change operations, see Security Management.
  - After the first host is installed, the system automatically assigns it an internal management IP address in the 172.28.0.2 network segment. In subsequent installation and configuration phases, any use of "the first host" or "the first node" refers to this host.

Bootstrap OS Installation of Other Hosts From the First Node

- In the address box of the local PC, enter **https://172.28.9.42:8890** to log in to the FusionSphere OpenStack web client and navigate to the **Discover Host** page.
- The login username is admin, and its default password is **FusionSphere123**. Change the password upon your first login.

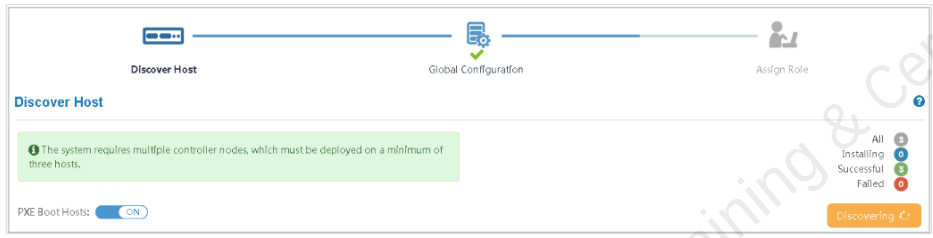

- Configure the other hosts to boot from PXE and power them on.

Bootstrap OS Installation of Other Hosts From the First Node

- Confirm that all hosts have been installed.

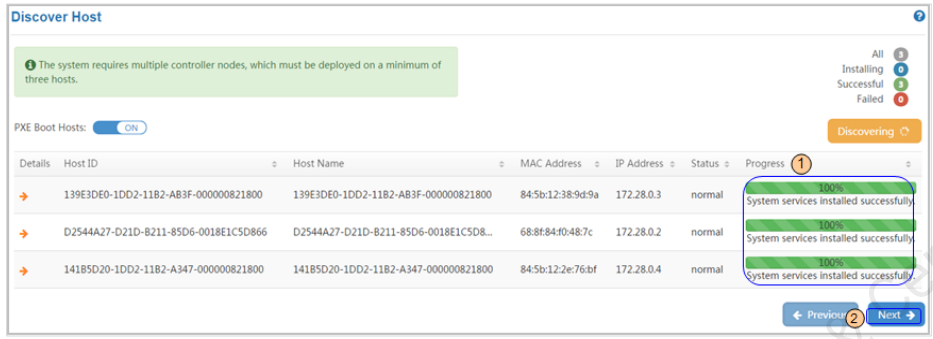

- If Process shows all target hosts and their installation progress at 100%, all the hosts have been installed.
- Confirm that all hosts have been installed and PXE Boot Hosts is set to OFF.
- To better identify hosts, click the names in the Host Name column and modify them in the format of chassis name-Host+serial number, for example, S1-RN01-SRN01-Host01.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. **Deploying FusionSphere OpenStack Service**
4. Configuring FusionSphere OpenStack
5. Installing and Configuring FusionSphere OpenStack OM

Page 33

Deploying FusionSphere OpenStack Services

- Configure basic settings.
- Configure AZ proxies.
- Configure DNS settings.
- Configure the NTP server.
- Configure advanced settings.

Configure Basic Settings

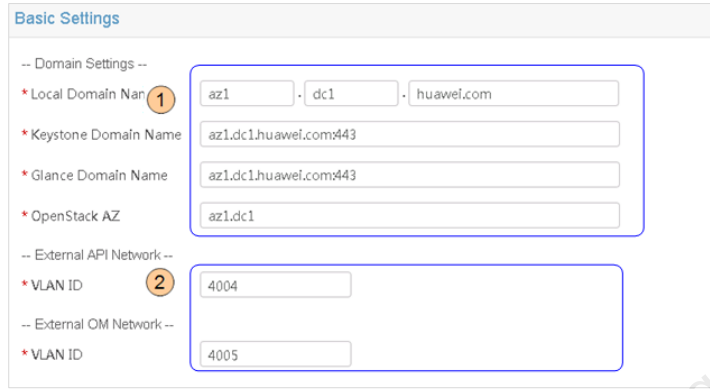

HUAWEI

- Local Domain Name
  - Domain to which the current FusionSphere OpenStack system belongs
  - Format: AZ name.DC name.Domain name suffix
- Keystone Domain Name
  - Domain name of the AZ where the Keystone service is deployed. The Keystone service can be shared by several AZs. Therefore, the current AZ may use the Keystone service deployed in another AZ.
  - Format: AZ name.DC name.Domain name suffix:Port number
- Glance Domain Name
  - Domain name of the AZ where the Glance service is deployed. The Glance service can be shared by several AZs. Therefore, the current AZ may use the Glance service deployed in another AZ.
  - Format: AZ name.DC name.Domain name suffix:Port number
- OpenStack AZ
  - OpenStack AZ is not used in any current Private Cloud scenarios. Keep this parameter the same as the Local Domain Name value. The system automatically configures the AZ and DC information during installation by default. If multiple FusionSphere OpenStack systems are deployed in one AZ, set the OpenStack AZ value for each system and do not use the automatically configured AZ and DC information.

AZ Proxies (1)

- Forward proxy ip: Network that allows internal services to communicate with external AZ services through the External API network.
- Reverse proxy ip: Allows external services to communicate with internal AZ services through the External API network.

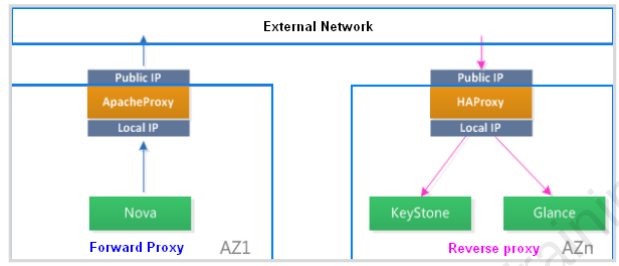

- AZ proxies are used to enable communication between services inside and outside of an AZ. external AZ services include the services of other AZs and third-party services. Therefore, AZ proxies must be configured even if the system uses only one AZ.

AZ Proxies (2)

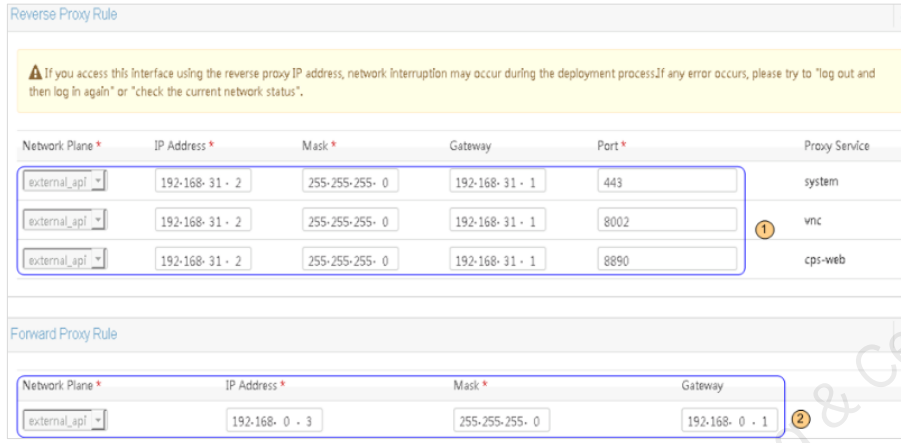

Configure the DNS Server

- DNS specifies the IP address used by the AZ to communicate with external networks.
- There are two types of domain name settings :
  - No external DNS server is available
  - An external DNS server is available.

- If the DC uses only one AZ, set it to OFF. If the DC uses multiple AZs, which means that multiple FusionSphere OpenStack systems share the Keystone and Glance services, set it to ON.

Configure the NTP Server

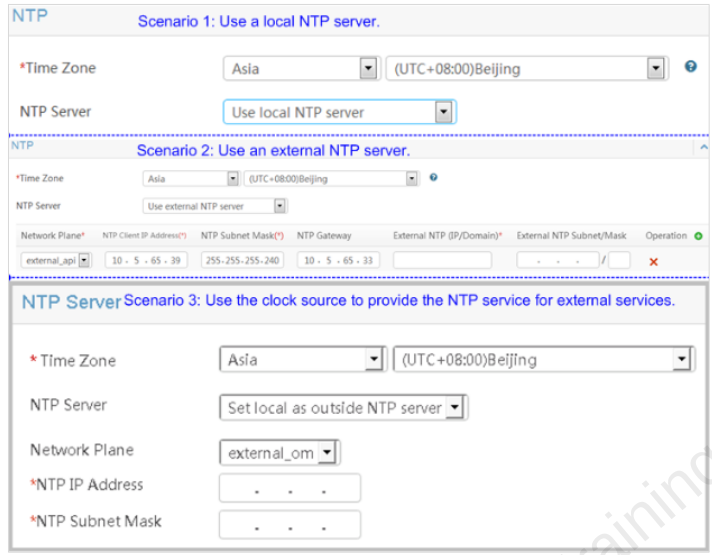

Configure Advanced Settings

- Deploy Database for Keystone: After this function is enabled, Keystone databases will be automatically deployed on two dedicated hosts. One is the controller host that has no default databases deployed, the other is a random compute host.
- MAC Address Pool: This function is used to configure a reserved MAC address segment. The format is fa:16:3e:00:00:00. The first three or four segments can be reserved. The last two segments of the MAC address must be **00:00**.

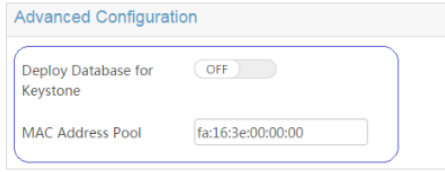

Assign Roles

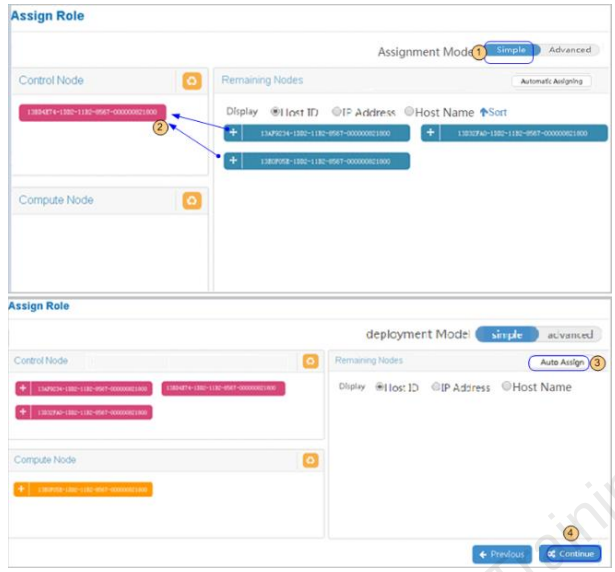

- Select the simple or advanced deployment model for assigning roles based on the number of the deployed controller nodes.
  - If three controller nodes are deployed, assign roles using the simple deployment model.
  - If more than three controller nodes are deployed, assign roles using the advanced deployment model.
- Set the deployment model to Simple.
- Manually drag the other two planned controller nodes to the Controller Node area.
- Click Auto Assign.
- Click Continue. The other hosts will be automatically added to the Compute Node area.

View the Deployment Progress

- View the deployment progress. FusionSphere OpenStack is installed and configured after the FusionSphere OpenStack web client displays all the target hosts and their progress at 100%.

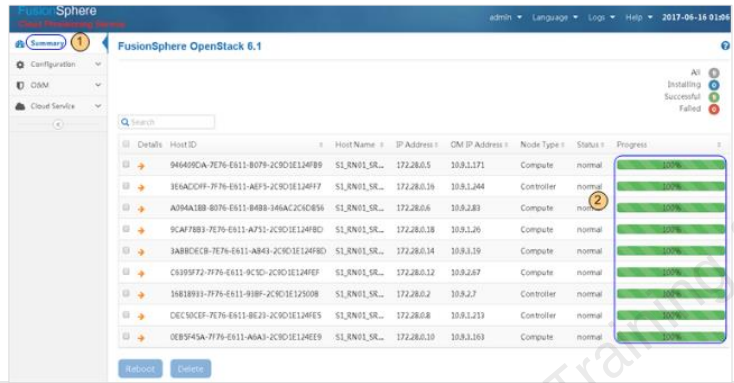

HUAWEI

- After the deployment task starts, the Summary page is displayed, showing the deployment progress for each host. After the progress of each host reaches 100%, the deployment is complete.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Page43

Configuring Networks

- **Physical Networks**: You must configure it before creating a virtual machine
- **NIC mapping**: A newly created physical network can provide network services only after it is mapped to physical network ports on hosts
- **Network planes**: used for FusionSphere cloud platform network platform

- The physical networks to be used by VMs must be configured before you create the VMs. Therefore, you are advised to create all the networks to be used by VMs and other services in this step.
- A newly created physical network can provide network services only after it is mapped to physical network ports on hosts.
- Hosts can be grouped automatically or manually.

Configure physical networks

- Choose **Configuration > Network**.
  - Create service-related physical networks.

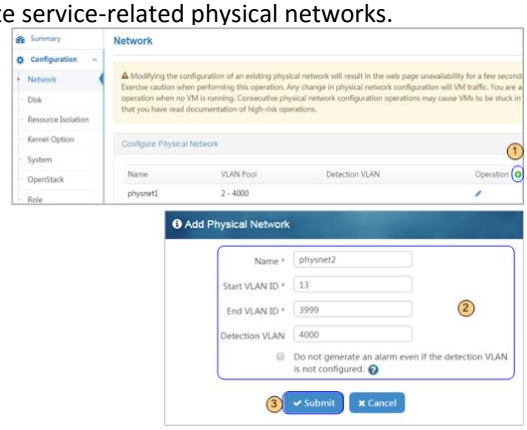

By default, the system creates a physical network and the system plane is carried on the physical network. You can add physical networks as required.

- Start VLAN ID
  - Specifies the start VLAN. Set this parameter based on your network plan.
- End VLAN ID
  - Specifies the end VLAN. Set this parameter based on your network plan.
- Detection VLAN
  - In KVM scenarios, set this parameter to prevent VM brain splits when VM HA is required.
  - Before services are provisioned, configure a detection VLAN for each physical network to prevent VM brain splits for VM HA. Ensure that this VLAN has been configured on the switch. If you configure detection VLANs after services are provisioned, network configuration changes will exert adverse impacts on services.
  - Ensure that this VLAN is not included in the VLAN range configured for the physical network.

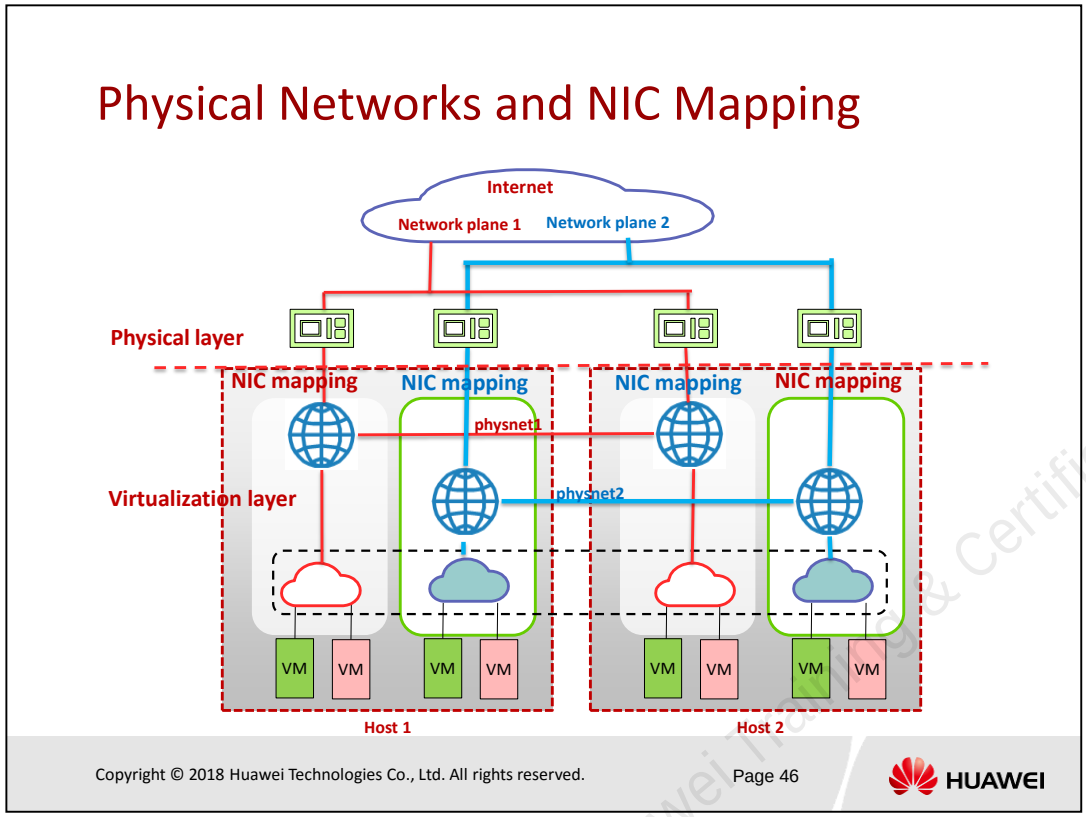

- After the FusionSphere OpenStack system is installed, a physical network named physnet1 is automatically created. Default VLANs for this network range from 1 to 4000. The Neutron service uses VLANs in this range to provide networks for tenants and uses network nodes to communicate with external networks. After a tenant creates a network, the Neutron service automatically allocates a VLAN from this range to the network. VLANs whose IDs are greater than 4000 are used by system ports. The two VLAN ranges isolate tenant networks from the system plane. If you have planned VLANs for the virtual networks of tenants and the system plane and these planned VLANs can communicate with external networks, you can set VLANs to static mode, regardless of the VLAN range of the physnet1 network.

NIC Mapping Mode

- **Kernel-Mode OVS Network Bridge**: specifies a kernel-mode OVS network bridge.
- **User-Mode EVS Network Bridge**: specifies a user-mode EVS network bridge, which is selected if a high-performance network softswitch is required.
- **PCI Passthrough**: specifies a PCI passthrough NIC plane. This option is available only when hosts use iNICs.

- Kernel-Mode OVS Network Bridge: specifies a kernel-mode OVS network bridge.
- User-Mode EVS Network Bridge: specifies a user-mode EVS network bridge, which is selected if a high-performance network softswitch is required. Before configuring an EVS network bridge, ensure that the following conditions are met: NICs support the DPDK I/O capability.
- The NIC type is either 82599EB or MT27520. A mixed use of the two types of NICs is not allowed.
- A maximum of four NICs are configured for a host.
- NICs are bound either in active/standby or LACP mode.
- PCI Passthrough: specifies a PCI passthrough NIC plane. This option is available only when hosts use iNICs. If you select this option for the first time, a dialog box will be displayed prompting you to configure the number of virtual network ports for the NIC.

NIC Mapping Bond Mode

- A kernel-mode OVS network bridge or a user-mode EVS network bridge allows NICs to be bound as a NIC bond and allows users to specify the binding mode. A PCI passthrough-enabled NIC does not support NIC binding.
- Bond mode:
  - Active/Standby mode
  - Source MAC Address-Based Load Balancing
  - LACP(Layer 2 Load Balancing or Layer 3 Load Balancing)

NIC Mapping Configuration

- Drag the target NICs to the area of the newly created physical network in the lower part
- After selecting the network card to set the plane type of network mapping. Currently, you can set plane types such as "Kernel OVS Bridge" and "Straight Through". In a private cloud scenario, when FusionSphere OpenStack is connected to a FusionCompute, only using the OVS type

- Drag the target NICs to the area of the newly created physical network in the lower part
- After you drag NICs to a physical network, set the network mapping mode.
- A kernel-mode OVS network bridge or a user-mode EVS network bridge allows NICs to be bound as a NIC bond and allows users to specify the binding mode. A PCI passthrough-enabled NIC does not support NIC binding.

Network Planes

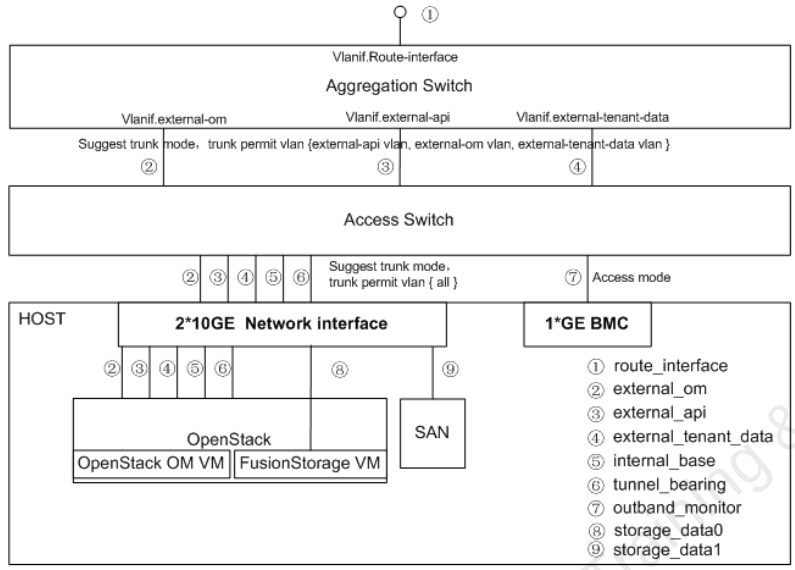

- Route Interface (connecting the aggregation switch and the core router):Connects the internal network and external networks.
- External OM (O&M management plane):connects to external networks and provides network support for the following functions: An internal user can log in to a host using Secure Shell (SSH) and perform operations on the host.The system monitors the host status and performance and uploads the data.Allow the system to connect to FTP or SFTP devices.
- External API (external API management plane):Connects to external networks and provides network support for the following functions:
  - Users access the web client from the Internet.
  - Users log in to VMs using Virtual Network Computing (VNC).
  - The system interconnects with Universal Distributed Storage (UDS) devices.

Network Plane Configuration

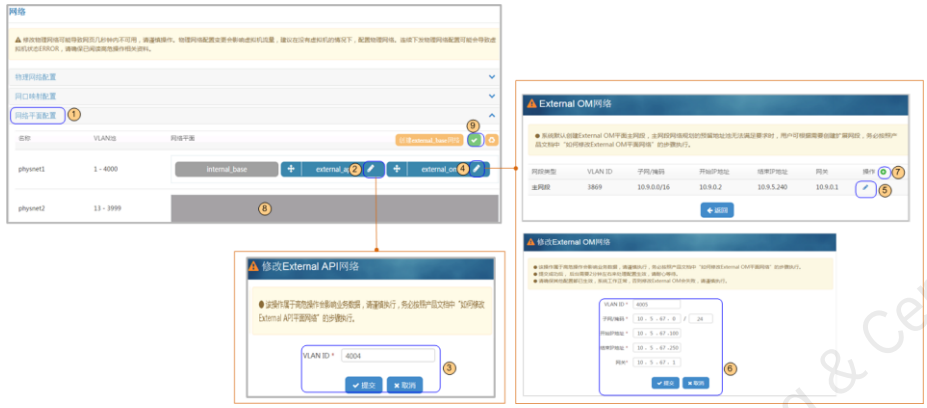

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Configure the Default Route

- On the FusionSphere OpenStack web client, choose **Configuration > System**.
- In **Default Gateway of Host Network Configuration**, set the gateway for the network segment to which the External OM network plane belongs.
  - The routes that you manually configured in the host OS will be automatically overwritten by the default routes of the gateway.

Modify the Reverse Proxy Information for the External OM Network

- On the FusionSphere OpenStack web client, choose **Configuration > System**.
- In the **Reverse Proxy** area, click . Configure the reverse proxy information for the external OM network.

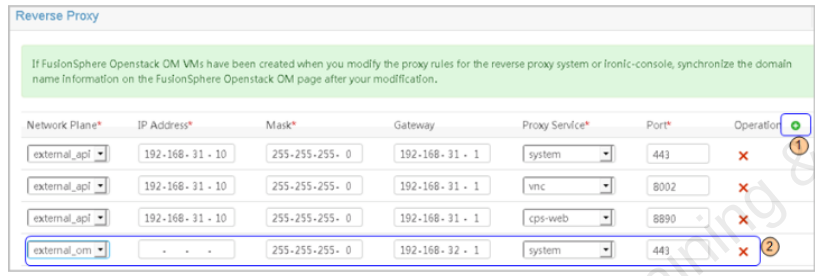

- Configure the reverse proxy information for the external OM network.
  - Network Plane: Select external_om.
  - IP Address: Set this parameter to an idle IP address that is not contained in the DHCP IP address segments on the external OM network.
  - Mask: Set the subnet mask for the external OM network.
  - Gateway: Set the gateway for the external OM network.
  - Proxy Service: Select system.
  - Port: Enter the port number of the system service. The default value is 443.
- After configuring the reverse proxy information, click Submit in the lower part of the page.
- Switch to the Summary page to establish automatic Keystone connection. If no error message is displayed, the connection is successful.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Configuring Disks (1)

- After installing and configuring FusionSphere OpenStack on the web client, configure the disk sizes based on some roles assigned to the management and compute groups. Then FusionSphere OpenStack automatically creates logical disks based on the configured disk sizes.

Role	Logical Disk	Description
mongodb	ceilometer-data	
image、compute	image image-cache	image: specifies the temporary disk space required by Glance. Controller node image partition: specifies the space for storing management VM images and management data and for providing disk space for management VMs. Controller node image-cache partition: specifies the space for caching management VM images or for backend storage connected to Cinder. Compute node image partition: specifies the VM disk space using local disks. Compute node image-cache partition: specifies the cache image space for service VMs that use local disks or for backend storage connected to Cinder. If connecting to multiple storage backends, Cinder needs to use the image-cache space provided by the host with the blockstorage-driver-kvmXXX role assigned.
rabbitmq	rabbitmq	rabbitmq: specifies the size of the disk partition attached to Rabbitmq. The specifications are related to the VM scale.
swift	swift	swift: specifies the host space fused for storing software installation packages and images.

Configuring Disks (2)

- Configure local disks for each management group.

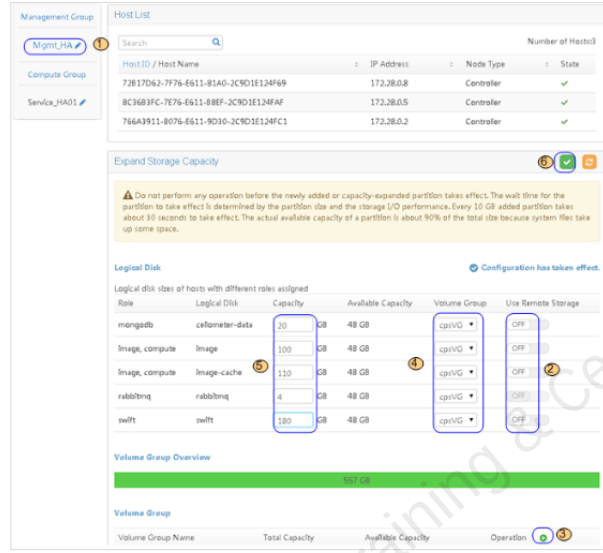

- The disk space can only be expanded.
- Modifications to disk space do not take effect immediately. Capacity expansion, newly added disk partitions, and I/O performance determine how long the changes will take. It takes about 30 seconds to expand or add 10 GB of space.

Configuring Disks (3)

- Configure disks for each compute group.

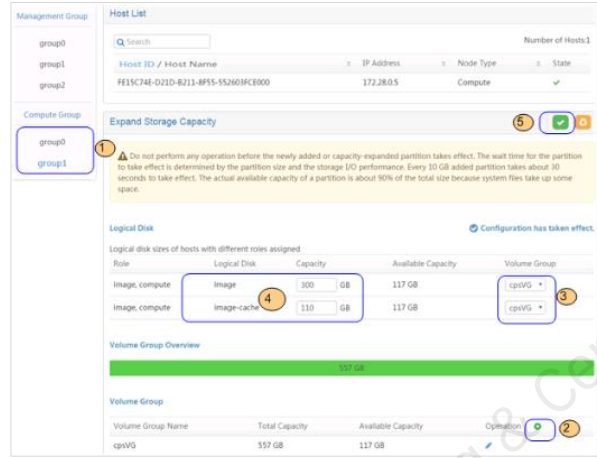

Configuring Disks (4)

- Configure disks on remote storage for each management group.

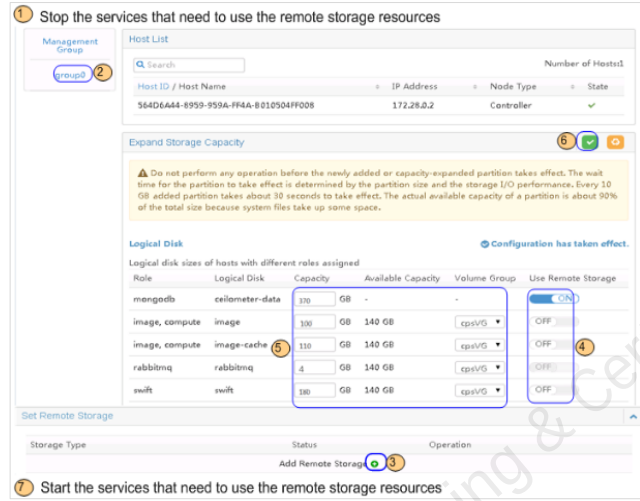

Set Remote Storage

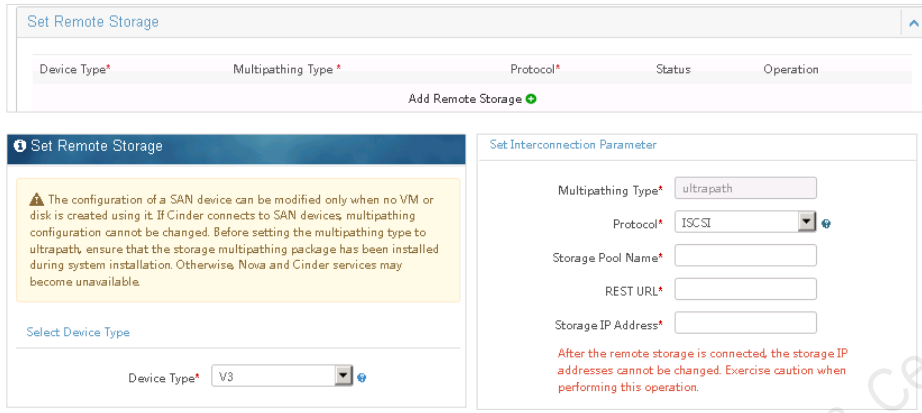

HUAWEI

- Set Remote Storage: The system service partitions are deployed in the cpsVG by default. You can create a dedicated volume group for a service

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Page 61

Resource Isolation Configuration

- Resource isolation limits the memory and vCPUs used by IaaS processes, preventing resource competition between IaaS processes and service VMs.
- By default, resources are not isolated in the management group. 2 vCPU and 2G RAM are isolated in the computer group.

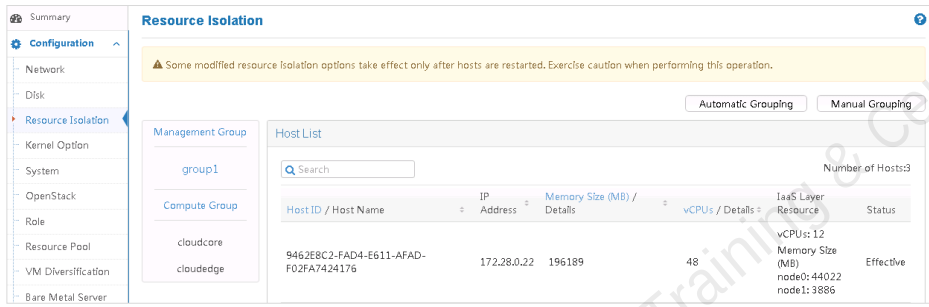

HUAWEI

- Resource isolation must be configured if you wan to deploy service VMs on management nodes according to the data plan. If FusionStorage Block is connected to FusionSphere OpenStack, resource isolation must be configured.
  - Automatic grouping: The system automatically groups hosts into one management group and one compute group based on the roles assigned to the hosts
  - Manual grouping: You can group hosts based on service requirements. If multiple types of upper-layer NEs, are to be deployed together, different resources need to be isolated. In this case, you can create multiple management groups or compute groups and add specified hosts to them to isolate resources for the hosts.When manually grouping hosts, add all hosts with management roles assigned to the management group.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Configuring FusionSphere OpenStack

- After installing FusionSphere OpenStack, configure its services.
  - Configure Glance
  - Configure Cinder
  - Configure Neutron
  - Configure Nova
  - Configure Ceilometer
  - Configure Rabbitmq
  - Configure Health

Configure Glance

- Choose **Configuration > OpenStack** on the FusionSphere OpenStack web client.
- Configure Glance parameters.
  - **Swift**: The Swift service in FusionSphere OpenStack is used. The Swift service is a distributed storage service. When it is used, data will be concurrently stored on three controller nodes, improving data reliability.
  - **Local Storage**: Local storage is used.
  - **OBS Service**: The OBS service is used. Set the following parameters if the backend storage type is set to OBS Service:

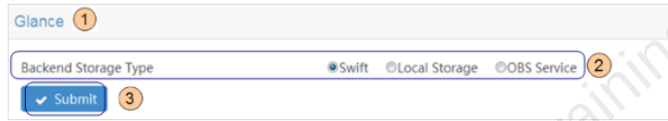

Configure Cinder

- Configure the storage network plane and storage multipathing in cinder page. After the configuration is complete, **in Configuration > Resource Pool** Management, complete the configuration of the storage resource pool
- **Configure Network Plane**
  - Storage Data 0
  - Storage Data 1
- **Configure Storage Multipathing**
  - Ultrapath
  - Multipath
  - Disabled

- In cloud data center scenarios, note the following requirements:
  - If FusionSphere OpenStack connects to FusionCompute or VMware vCenter, configure the backend storage after adding FusionCompute or VMware vCenter to the FusionSphere OpenStack system
  - If KVM hosts require storage resources, configure backend storage here.

Configure Cinder Network Plane

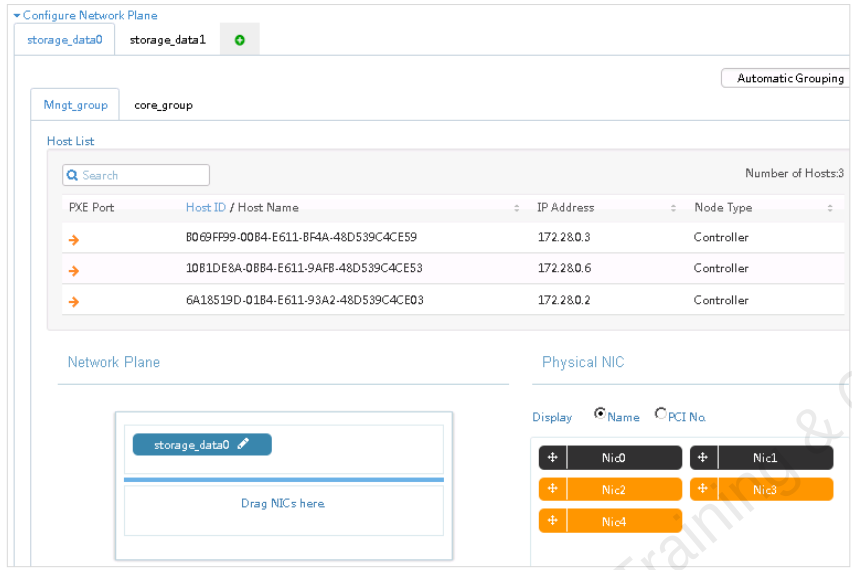

HUAWEI

- Set Use Storage Multipathing.
  - If Huawei SAN storage is in use, set Use Storage Multipathing to ultrapath and ensure that you have obtained the multipathing driver package and installed the driver.
  - If SAN storage devices of other vendors are in use and the native OpenStack multipathing software is to be installed, set Use Storage Multipathing to multipath. The universal multipathing function supports only FC SAN devices.
  - If storage multipathing is not required, set Use Storage Multipathing to Disabled.

Configure Neutron

- Configure Neutron parameters.

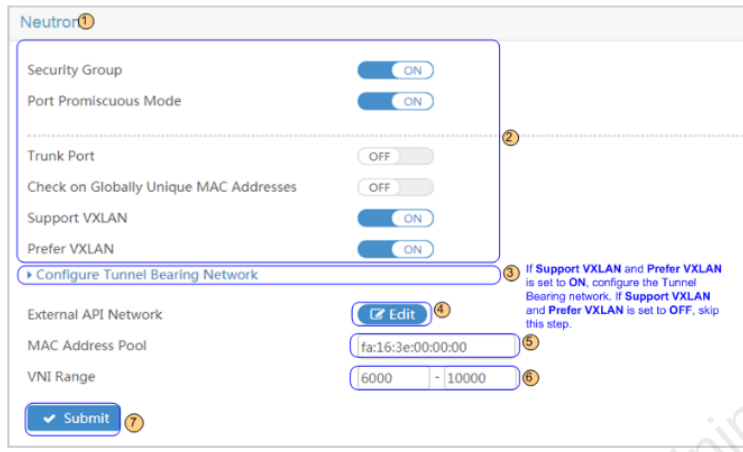

- Support VXLAN
  - If this function is enabled, the system supports the VXLAN network type. VXLAN is a virtual Ethernet established based on the physical network. It uses layer-3 protocols to encapsulate layer-2 packets and enables the layer-2 network to extend to layer 3. The VLAN feature allows users to create only 4094 isolated virtual networks, whereas the VXLAN feature breaks through this limitation and allows users to create about 16 isolated virtual networks. VXLANs support large-scale network deployment in multi-tenant environments.
  - In VLAN-based networking, if the Tunnel Bearing network plane is deleted, disable this function.
  - This function takes effect globally. Once it is disabled, you cannot create VMs that use VXLAN in clusters. Therefore, do not disable this function if VMs using VXLAN are running in the system.
- Prefer VXLAN
  - Prefer VXLAN: After this function is enabled, the Virtual Extensible Local Area Network (VXLAN) type is used by default. This parameter is valid when Support VXLAN is set to ON.
- External API Network
  - If Configure subnet is set to ON, the start and end IP addresses must be idle on the network. The gateway must be consistent with that of the reverse proxy on the External API network.

Configure Nova

- Configure Nova parameters.
  - Record BIOS logs and Serial logs
  - Memory Preemption
  - Bind NIC Interrupt
  - VM Watchdog
  - Management host group
  - Cold Migration

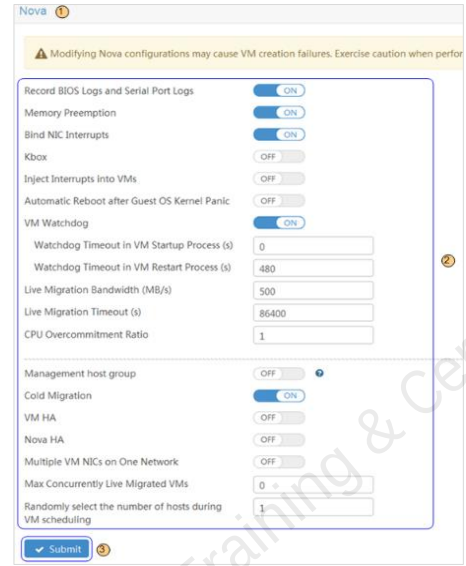

- Management host group
  - Before enabling this function, ensure that FusionSphere OpenStack OM VMs have been created. Otherwise, specifications modification or migration operations on FusionSphere OpenStack OM VMs may fail.
  - Management host group: After this function is enabled, the system creates a management host group named manage-aggr in the AZ manage-az, and adds all controller hosts to the group. Service VMs cannot be created on controller hosts.
  - After this function is enabled, you can add hosts to or delete hosts from the group. In some large-scale scenarios, if some roles such as gaussdb and rabbitmq are assigned to dedicated hosts, FusionSphere OpenStack OM VMs are deployed on dedicated hosts, or the role fc-nova-compute or vmware-nova-compute is to be assigned to dedicated hosts, select these hosts and other non-controller hosts from the host list and add them to the management host group.
- Multiple VM NICs on One Network
  - Multiple VM NICs in One Network: After this function is enabled, multiple VM NICs can belong to the virtual network that is created in the same Neutron. This function is typically used for SR-IOV NICs when multiple vNICs on the same network pass through to one VM. In this case, ensure that the network ports on the VM automatically form bonds to increase the VM NIC bandwidth and improve data reliability. Enable this

function in private cloud scenarios.

Configure Ceilometer

- Configure Ceilometer parameters.
  - Collection Interval
  - Monitoring Indicator Storage Duration

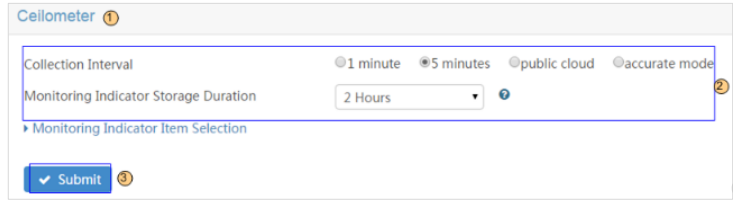

- Collection Interval
  - The Ceilometer service monitors and measures the FusionSphere OpenStack cloud resources for billing, baseline management, scaling, and data collection.
- Monitoring Indicator Storage Duration
  - The system saves monitoring data based on the selected duration. For example, if you set the parameter to two hours, the system saves monitoring data generated within the latest two hours. Data generated earlier will be deleted.
  - The system offers multiple duration options for you to select based on service requirements. If you need to switch the duration after the system has been running for a period of time, but there is a big gap between the two duration options, for example, switching from seven days into two hours, the system will delete a large amount of historical data, which will cause high I/O and adversely affect the MongoDB functions. In this case, you can change the duration into a moderate value first, for example, changing it into three days first and then to two hours.

Configure Rabbitmq

- Configure RabbitMQ parameters. You can set
  - Rabbitmq SSL protocol version
  - Configure Network

- Rabbitmq SSL protocol version
  - Determines whether to use an encrypted access to RabbitMQ.

Configure Health Check (1)

- Configure health check parameters.
  - The health check function allows the system to detect whether the network switch settings map the host settings, whether the network plane settings are correct, and whether the network planes are in the healthy state. If detecting a fault, the system generates an alarm prompting the maintenance personnel to handle.
  - In cloud data center scenarios, if KVM, FusionCompute, and VMware resource pools are available, this function takes effect only for KVM hosts after it is configured.

Configure Health Check (2)

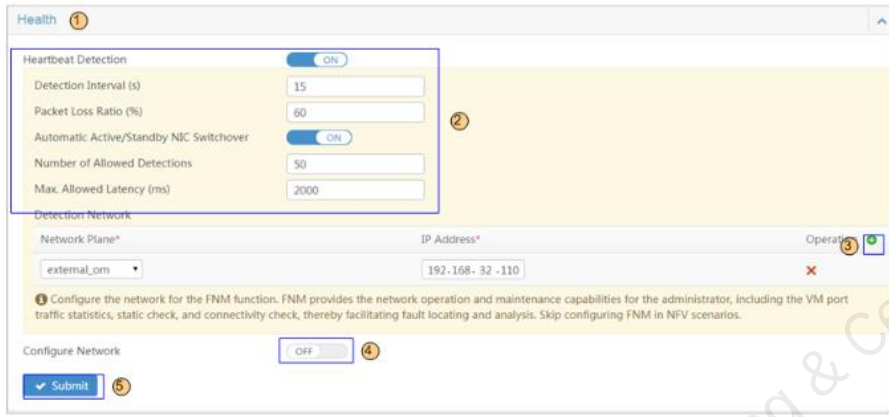

- Network Plane
  - Specifies the network plane to be detected.
  - When management network planes are deployed on the same physical network, select one of them to enable detection. You are recommended to select the External OM network plane with an external gateway.
- IP Address
- Configure Network
  - Configure the IP address used by the FNM function. FNM provides the network operation and maintenance capabilities for the administrator, including the VM port traffic statistics, static check, and connectivity check, thereby facilitating fault locating and analysis. Skip configuring FNM in NFV scenarios because it is not required in the scenarios.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Page 74

Configuring Roles

- During the Deploying FusionSphere OpenStack Services phase, hosts are assigned different roles to function as controller or compute nodes. Some special roles, such as **router**, **dhcp** and **blockstorage-driver/blockstorage-driver-kvmxxx**, must be deployed based on service requirements. Therefore, after the Deploying FusionSphere OpenStack Services phase, adjust roles based on your service requirements.
- Adjust roles:
  - Add a role
  - Delete a role.
  - Migrate a role

Add a Role

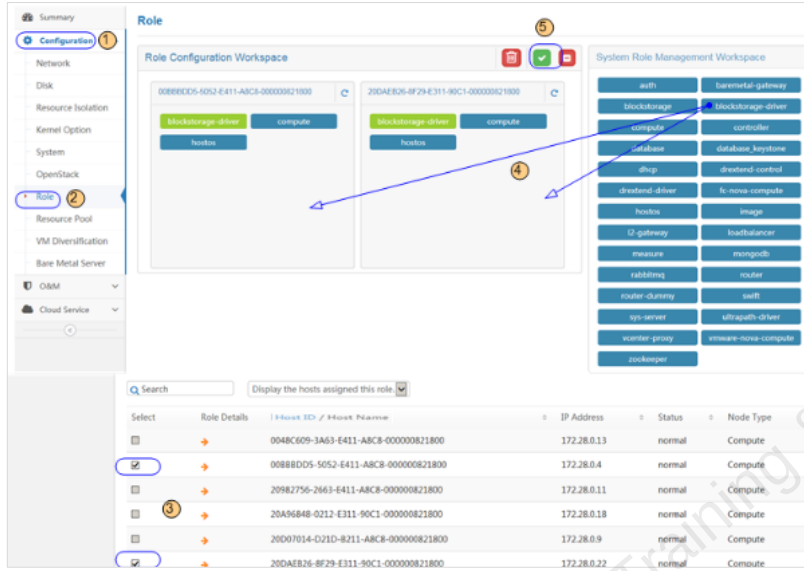

HUAWEI

- Choose Configuration > Role.
- Select two compute nodes.
- Drag the blockstorage-driver/blockstorage-driver-kvmxxx from System Role Management Workspace to the two hosts in Role Configuration Workspace.
- Submit.

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. **Configuring FusionSphere OpenStack**
   - Configuring Networks
   - Configuring a Reverse Proxy for the External OM Network and the Default Route
   - Configuring Disks
   - Configuring Resource Isolation for FusionSphere OpenStack
   - Configuring FusionSphere OpenStack
   - Configuring Roles
   - Configuring the Maximum Connections of a Database

Configuring the Maximum Connections of a Database (1)

- Use PuTTY to log in to the first host. Ensure that the reverse proxy IP address on the External API network and username **fsp** are used to establish the connection. The default password of user **fsp** is **Huawei@CLOUD8**.
- Run the following command and enter the password of user root to switch to user **root: su – root**.
- Run the following command to disable logout on timeout: **TMOUT=0**
- Importing Environment Variables.
- Run the following commands to change the number of GaussDB connections:
  - **cps template-params-update --service gaussdb gaussdb --parameter database_max_connections=3000**
  - **cps commit**

- To improve the processing efficiency of large-scale concurrent service flows, for example, concurrently creating 100 VMs, adjust the maximum number of connections of a database.
- You are advised to change the maximum number of database connections for private cloud scenarios.

Configuring the Maximum Connections of a Database (2)

- If independent databases are deployed for the cascading services, run the required command to change the number of the connections of each database based on site requirements.
  - **cps template-params-update --service gaussdb_keystone gaussdb_keystone --parameter database_max_connections=3000**
  - **cps template-params-update --service gaussdb_nova gaussdb_nova --parameter database_max_connections=3000**
  - **cps template-params-update --service gaussdb_neutron gaussdb_neutron --parameter database_max_connections=3000**
  - **cps template-params-update --service gaussdb_cinder gaussdb_cinder --parameter database_max_connections=3000**
  - **cps commit**

Contents

1. FusionSphere Introduction
2. Installing FusionSphere OpenStack
3. Deploying FusionSphere OpenStack Service
4. Configuring FusionSphere OpenStack
5. **Installing and Configuring FusionSphere OpenStack OM**

Page 80

FusionSphere OpenStack OM VMs

- The FusionSphere OpenStack OM VM must use the local hard disks.
- Operating System (OS): Huawei-developed OS (based on Novell SUSE Linux Enterprise Server 11 SP3 64-bit)
- FusionSphere OpenStack OM in all-in-one mode

Flavor	vcpus	Memory(G)	Disk(G)
0-512 VM	4	6	80
512-5120 VM	6	18	80
5120-10000 VM	12	24	120
10000-80000 VM	16	30	150

Install FusionSphere OpenStack OM (1)

- Upload the software package

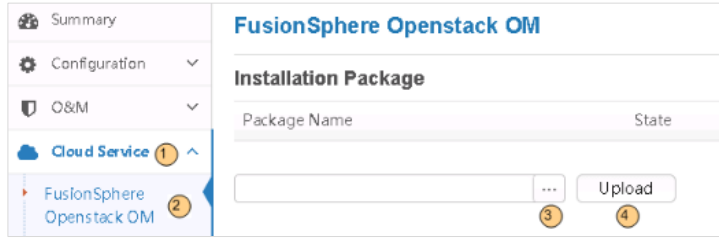

Install FusionSphere OpenStack OM (2)

- Create FusionSphere OpenStack OM VMs.

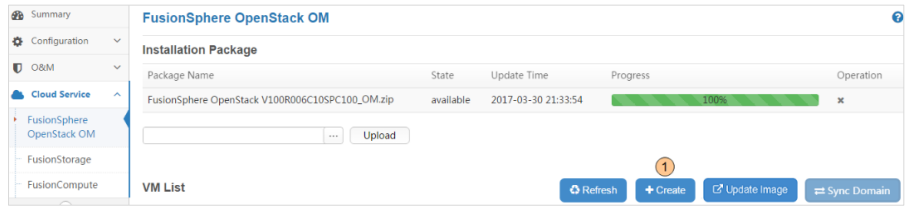

- Configure basic settings:
  - Application Scenario: Select FusionSphere.
  - Type: Select AllinOne OM.
  - Install Service OM: Select yes.
  - HA Mode: Select a VM high availability (HA) mode.
  - System Scale: Select the desired VM management scale. Greater than 512 VMs must be selected when the FusionSphere Service OM is deployed.
  - CPU Binding: Determine whether to bind the VM CPU to the server CPU.
  - Cascading Mode: Select Standalone.
- Configure API network settings:
  - Floating IP Address: Enter the API network floating IP address.
- Configure OM network settings:
  - Floating IP Address: Enter the OM network floating IP address.

Install FusionSphere OpenStack OM (3)

- Configure FusionSphere OpenStack OM VM parameters

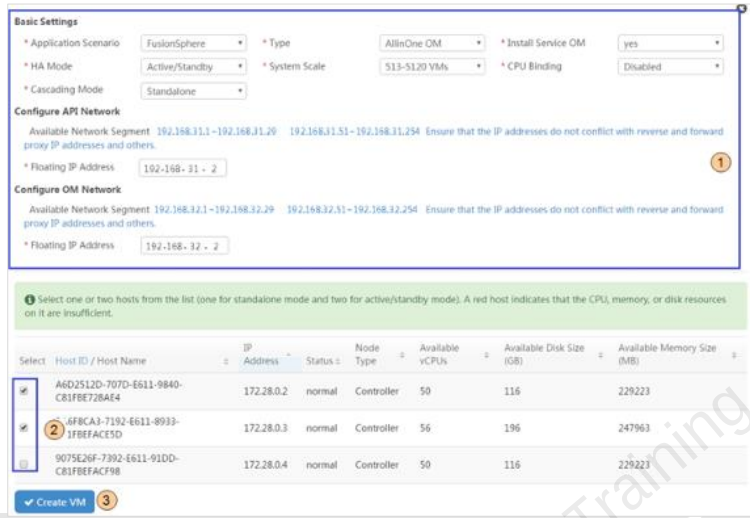

Initial Configuration Process

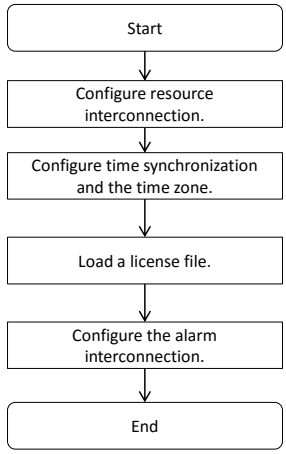

- Configure resource interconnection.
  - You can click ⇄ Sync Domain in the FusionSphere OpenStack web client to interconnect with Keystone and the FusionSphere OpenStack instance.
- Configure time synchronization and the time zone.
  - You can log in to the FusionSphere OpenStack OM web client and choose **System** > **System Configuration** > **Time Management**.
- Load a license file.
  - Load the license file to authorize the system to provide the licensed services in the permitted period. You can log in to the FusionSphere OpenStack OM web client and choose **System** > **System Configuration** > **License Management** to import the new license file.
- Configure the alarm interconnection.
  - The domain name resolution information needs to be configured when FusionSphere OpenStack alarms need to be reported to FusionSphere OpenStack. You can log in to the OpenStack web client and choose **System** > **DNS** to add DNS rules. The domain name is local-oam.domainname.com, and the IP address is the FusionSphere OpenStack floating IP address.
  - You can log in to the FusionSphere OpenStack web client to configure alarm service interconnection information.

Adding Resources to FusionSphere OpenStack OM

- Locate the FusionSphere OpenStack OM VM, click next to **Connect Keystone**, and configure parameters to interconnect FusionSphere OpenStack OM and FusionSphere OpenStack.
- Configure DNS settings.
- Choose **Configuration > System**. In the **DNS** area in the middle of the page, locate the **Operation** column, and click to add the DNS server for the FusionSphere OpenStack OM service.

Loading a License File

- On FusionSphere OpenStack Management Console, choose **System > System Configuration > License Management**. The License Management page is displayed.
- Apply for a license file using the equipment serial number (ESN).

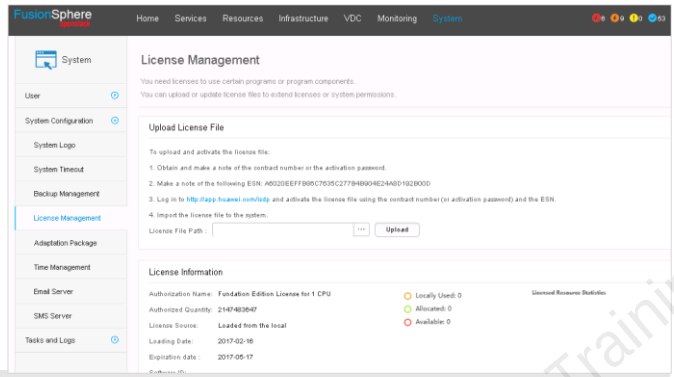

- Get the ESN number
- At support.huawei.com, apply for a license file
- Start to upload the license file, click "OK"

Configuring Time Synchronization and the Time Zone

- In the FusionSphere OpenStack Management Console, choose **System > System Configuration > Time Management**. The Time Management page is displayed.

Configuring FusionSphere OpenStack Alarm Reporting

- Configure the floating IP address of the alarm component.
  - Enter **https://Reverse proxy IP address:8890** in the address box of your browser to log in to the FusionSphere OpenStack web client as user **admin**.
  - On the FusionSphere OpenStack web client, choose **Configuration > System**.
  - Click **Alarm Configuration**. Configure the floating IP address of the alarm component. After the configuration, click **Submit**.
- Enable FusionSphere OpenStack alarms to be reported to FusionSphere OpenStack OM.
  - On the FusionSphere OpenStack web client, choose **Cloud Service > FusionSphere OpenStack OM**.
  - Click  o the right of **Connect Alarm Service**.

More Information

- Huawei Learning website
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei Support case library
  - http://support.huawei.com/enterprise/servicecenter?lang=zh

Thank You

www.huawei.com

FusionCompute Principles and Deployment in Cloud Data Center Scenarios

www.huawei.com

Foreword

- This course introduces FusionCompute and the FusionCompute installation process. After understanding the process, trainees will perform installation and deployment tests and learn installation and deployment guides to independently create FusionCompute hypervisors and log in to the FusionComopute web client. In addition, this slide introduces the basic operations to achieve interconnection of FusionCompute and FusionSphere.

Objectives

- Upon completion of this course, you will be able to:
  - Learn the FusionCompute architecture.
  - Be familiar with FusionCompute key functions.
  - Be familiar with the FusionCompute installation and configuration process.
  - Achieve interconnection between FusionSphere and FusionCompute.

Contents

1. **FusionCompute Architecture**
2. FusionCompute Features
3. FusionCompute Deployment
4. FusionSphere OpenStack Interconnection Configuration
5. Initial Configuration of FusionCompute
6. Interconnection of FusionCompute and FusionSphere OpenStack
7. FusionCompute O&M

Introduction to FusionCompute

- FusionCompute is a virtualization platform that is deployed on physical servers to virtualize the server hardware resources, including CPUs, memory, and network interface cards (NICs).
- FusionCompute includes the CNA and VRM.

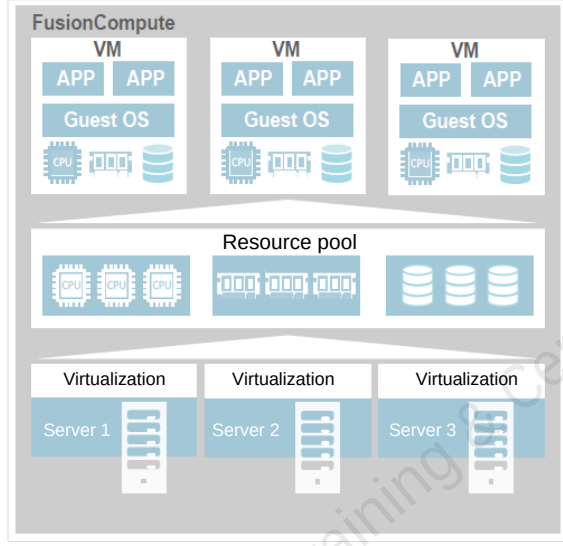

FusionCompute CNA Computing Device Virtualization

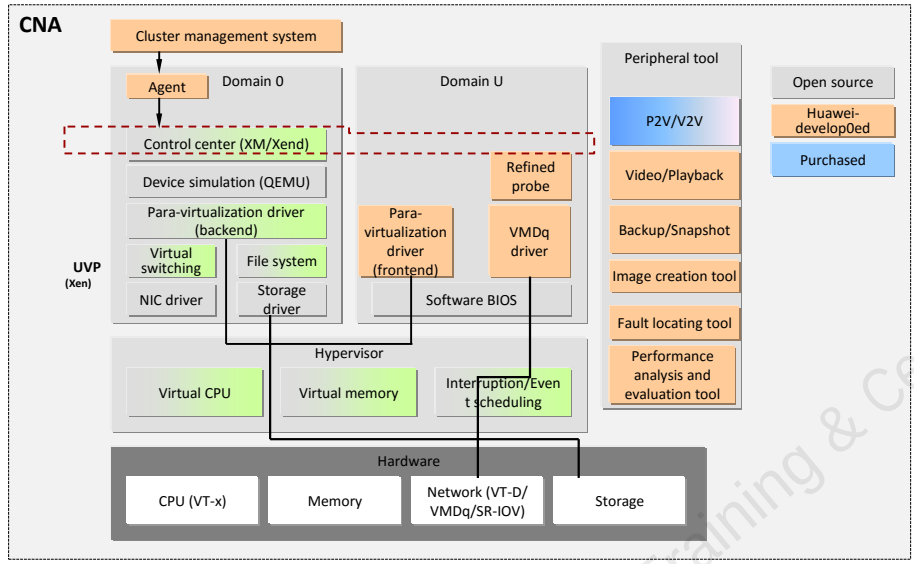

HUAWEI

- This figure shows the logical architecture of a FusionCompute CNA node.
- In the figure, except the cluster management system and agent, other modules are called the UVP.
- VRM is the cluster management system. Agent functions as the bridge between VRM and CNA to form the FusionCompute hypervisor.

FusionCompute VRM Layer

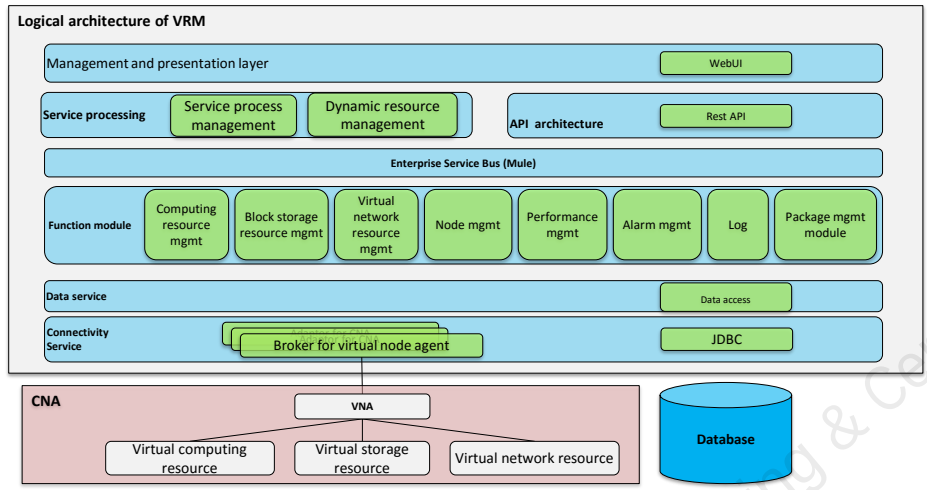

VRM is the FusionCompute management center, which is responsible for resource allocation and scheduling, and provides unified operation and maintenance, resource monitoring, and resource management.
VNA is deployed on the CNA and used to configure and manage computing, storage, and network virtualization.

- This figure shows the logical architecture of a FusionCompute VRM node.
- The service package includes multiple modules. Management nodes manage all CNA nodes. Computing, block, and network resource management focus on the resource pooling of CNA nodes. VRM nodes manage all CNA nodes and provide services, such as performance monitoring and log management.
- In this figure, Broker Virtual Node Agent and VNA connect the VRM nodes to the CNA nodes. Please understand the FusionCompute VRM nodes by combining FusionCompute CNA computing device virtualization described in the previous slide.

Contents

1. FusionCompute Architecture
2. **FusionCompute Features**
3. FusionCompute Deployment
4. FusionSphere OpenStack Interconnection Configuration
5. Initial Configuration of FusionCompute
6. Interconnection of FusionCompute and FusionSphere OpenStack
7. FusionCompute O&M

VIMS Virtual Cluster File System

- VIMS: the basis of advanced features, such as the thin provisioning, snapshot, and storage live migration.

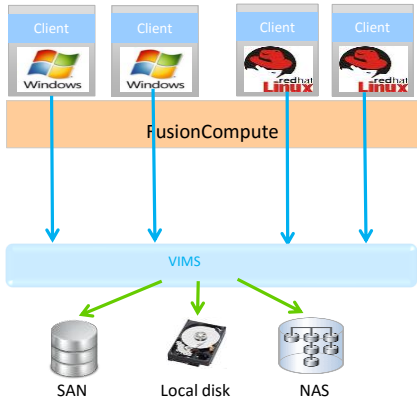

- **Technical features**
  - Be compatible with FusionStorage, SAN, IP SAN, NAS, and local disks.
  - Support SSDs, dynamic disks, and differential disks.
- **Application scenarios**
  - Used for VMs that require advanced storage features, such as the storage live migration, snapshot, and linked clone.
- **Restrictions**
  - If the VIMS is used, a host cluster can contain a maximum of 32 hosts.

Provides better support for deploying cluster software on the virtualization platform.

- Virtual Image Management System (VIMS) is the core of the shared storage technology in FusionCompute computing clusters and also a basic technology for VM advanced features.
- VIMS does not require manual configuration. After you add shared storage to a computing cluster, VIMS will be automatically created on the shared storage when the storage is formatted.

VM Live Migration (vMotion)

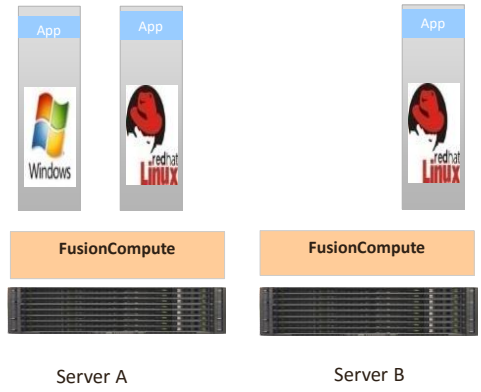

- Technical features
  - Based on the memory compression technology, VM live migration efficiency is improved by one time.
  - VM disks can be live migrated from a shared storage device to local disks.
- Application scenarios
  - Short-time service interruption is tolerable but services must be quickly restored, such as lightweight database and desktop cloud services.

VM live migration does not interrupt VM services and insensible to users.

HUAWEI

- FusionCompute computing clusters are deployed on VIMS shared storage. The VM live migration technology can allow the system to automatically migrate running VMs to other nodes in the same computing cluster without service interruption. This function is the prerequisite for implementing automatic load balancing and energy saving.
- Operators can also manually live migrate VMs.

Storage Live Migration

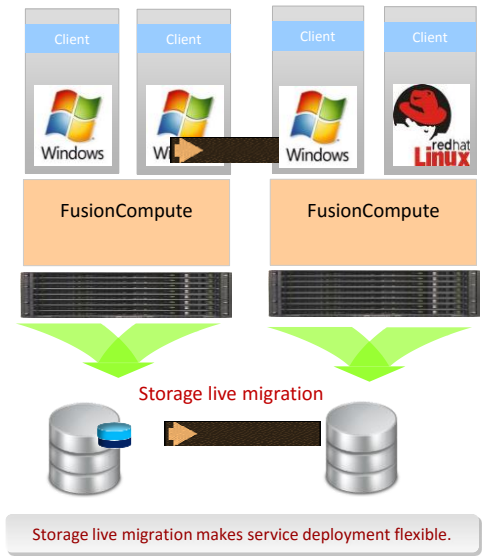

Storage live migration makes service deployment flexible.

- Technical features
  - Allows VMs to be migrated between hosts that use different storage media online.
  - Make the migration bandwidth controllable, preventing adverse impacts on services.
- Application scenarios
  - Host maintenance in the scenario where local disks are used.
  - Cluster migration
  - Storage capacity efficiently managed (releasing storage space)

HUAWEI

- Storage live migration can prevent service downtime due to storage system maintenance.
- When storage devices need to power off due to maintenance demands or device replacement, storage live migration allows VM resources to be migrated from one data store to another without service interruption.
- Storage live migration allows a new data store to be attached to a computing cluster and allows VM resources to be migrated from the original data store to the new data store.

Storage Thin Provisioning that Dramatically Reduces Storage Investment

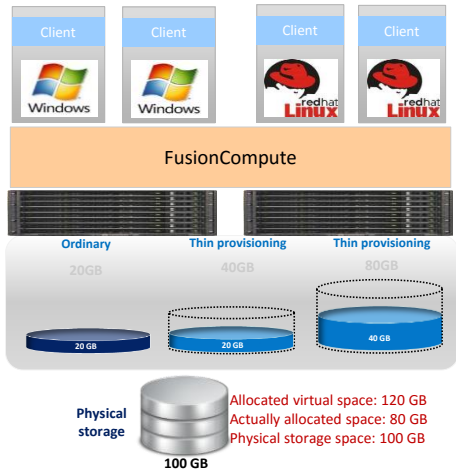

The storage usage is dramatically increased and the storage cost is dramatically reduced.

- Technical features
  - A VM user can view the configured capacity only. However, the actual space occupied by physical disks is adjusted based on the used space.
  - The space recycling is supported.
  - The storage cost is reduced by 40%. This ratio is related to application models.
  - The IOPS on a single VM is reduced compared with the traditional storage method.
- Application scenarios
  - The storage capacity is large and the IOPS requirement is low

Automatic Load Balancing (One of the DRS Functions)

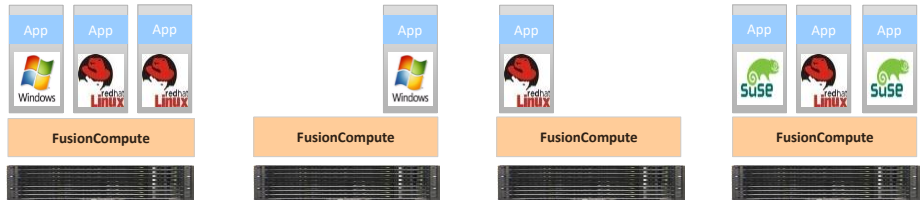

- Technical features
  - In the same cluster, the system automatically balances VM workloads based on policies.
  - The load balancing algorithm is optimized to prevent invalid VM migration.
- Application scenarios
  - The load balancing function ensures service performance.
  - The load balancing function staggers the busy and idle periods to avoid network congestion in peak hours.

- DRS: Dynamic Resource Scheduling
- FusionCompute computing clusters are deployed on VIMS shared storage devices. The DRS monitors resource usage for each computing node in the same cluster in real time, and uses the VM live migration function to intelligently migrate VMs with high workloads to other nodes with sufficient resources, thereby balancing resource use and ensuring sufficient resources.

Automatic Energy Saving Provided by DPM

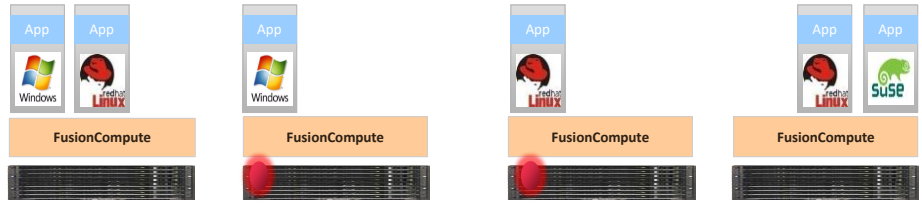

- Technical features
  - The system automatically powers on or off proper physical servers, therefore reducing the number of migrated VMs.
  - Some physical servers can be hibernated to meet the requirements of new services.
- Application scenarios
  - In the evening, only a few services are processed on physical servers. In this case, VMs can be automatically merged and idle physical machines can be powered off.
  - In the daytime, services to be processed increase. Physical servers automatically power on to take over the VMs on heavy-load physical servers, which ensures the QoS.

- DPM: Dynamic Power Management
- FusionCompute computing clusters are deployed on VIMS shared storage devices. The DPM module migrates VMs on multiple light-load hosts to light-load hosts and powers off idle hosts to save energy using the DRS function based on the service load and physical resource usage.

Anti-Virus Virtualization

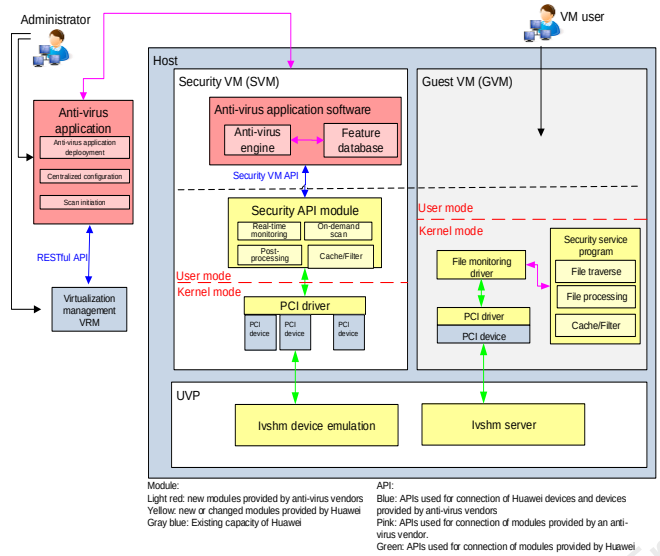

- Compared with the traditional antivirus function, the VM antivirus function requires no complete antivirus software for user VMs, and does not require computing resources on user VMs, leaving sufficient storage and computing resources for user VMs.
- Compared with the online antivirus function, the VM antivirus function provides higher performance because the system uses host physical memory swapping to scan, monitor, and remove viruses, and does not require network resources on user VMs.

- Apply to scenarios where the services require high security.

Distributed Virtual Switch

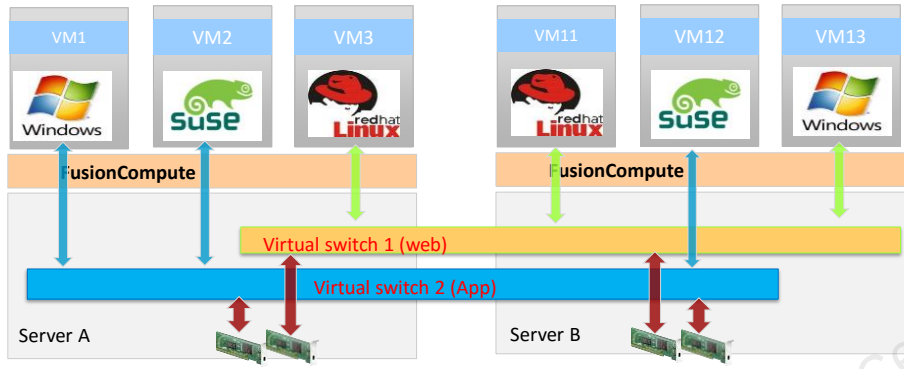

- **Technical features**
  - Provides a unified interface for configuration and management.
  - Supports VLANs, layer 2 security policies, and bandwidth limiting.
- **Application value**
  - Supports flexible access control for networks.

VXLAN

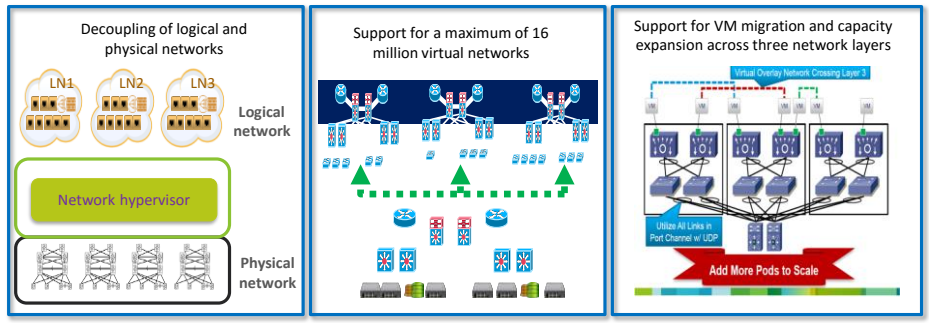

- Application value
  - The virtual network is decoupled from the physical network by configuring physical switches.
  - The backbone network does not need to be reconstructed for supporting TRILL or due to MAC addresses insufficiency. Existing devices on the backbone network can be reused.
  - A maximum of 16 million virtual networks are supported to meet multi-tenant requirements. (The VLAN technology supports only about 4000 virtual networks.)
  - VMs can be migrated across three physical networks, which facilitates remote DR and backup that requires the VIS remote active-active DR storage device.

- In the hypervisor, VXLAN is important for creating the large layer 2 network and solves the problem that VXLAN IP address pools are not available under the multi-tenant public cloud condition.
- This document mainly focuses on product training, including product functions, configuration and test. For details about VXLAN, see the Huawei cloud computing certification course or online information.
- VXLAN can enable the VM to migrate across data centers.
- Upon completion of this course, you must have a brief understanding for FusionSphere and FusionCompute. Next, we will learn how to install and deploy FusionCompute.
- Now, trainees must know what FusionCompute is, what functions it can provide, and where it is located in the FusionSphere solution. However, trainees still cannot know how to use these functions in practice. Therefore, trainees will need to further learn FusionCompute functions, application scenarios, and configuration and deployment in terms of computing virtualization, storage virtualization, and network virtualization.

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. **FusionCompute Deployment**
   - Solution Overview and Process
   - Environment Preparations
   - Manual Installation

Solution Overview

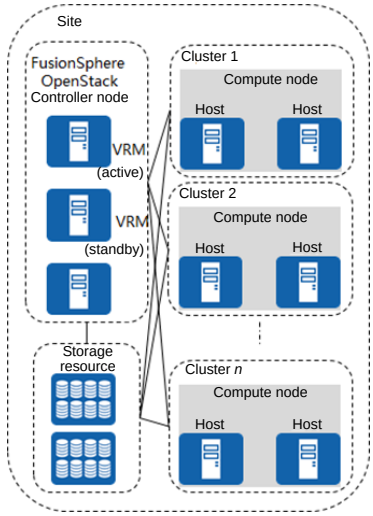

- A VRM node is a FusionCompute management node. It provides a management portal for users to centrally manage virtual resources.
- Hosts are also called physical servers and provide computing resources for FusionCompute. When local disks are used, hosts also provide storage resources. Certain hosts can also provide storage resources when local hard disks are used as storage media.

Node Type	Deployment Mode	Deployment Principle
Host	Physical deployment	Multiple hosts are deployed based on customer requirements on computing resources to provide virtual computing resources. The hosts also provide storage resources if local hard disks are used to provide storage resources. In virtualization deployment, a specified host is required to create a VRM VM.
VRM	Virtualization deployment	In the scenario where VRM is deployed on VMs, two hosts need to be selected from FusionSphere OpenStack management nodes to deploy VRM VMs in active/standby mode.

Installation Process

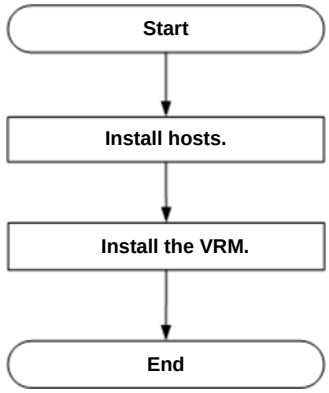

- Installation methods
- Installation using tools
  - Hosts can be installed.
  - VRM is not allowed to be installed on physical servers.
  - VRM is not allowed to be installed on FusionSphere OpenStack management nodes.
- Manual installation
  - Hosts and VRM need to be independently installed.
  - Host installation: Manually mount an ISO image to the host to be installed and set host parameters.
  - VRM installation: After logging in to the FusionSphere OpenStack web client, when you deploy VRM on FusionCompute, manually upload the VRM installation file and manually create the VRM VM.
- This slide uses the manual installation method.

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. **FusionCompute Deployment**
   - Solution Overview and Process
   - Environment Preparations
   - Manual Installation

Software Packages

- Software packages required for installation

Software Package	Software Description	How to Obtain
FusionCompute V100R006C10_Installer.zip (Not required by manual installation)	FusionCompute installer	• For enterprises: Visit http://support.huawei.com/enterprise and choose Support > Cloud Computing > FusionCloud > FusionCompute > V100R006C10. • For carriers: Visit http://support.huawei.com and choose Product Support > Carrier IT > Cloud Computing > FusionCloud > FusionSphere > FusionCompute > V100R006C10.
FusionCompute V100R006C10_CNA.iso	FusionCompute host OS	• For enterprises: Visit http://support.huawei.com/enterprise and choose Support > Cloud Computing > FusionCloud > FusionCompute > V100R006C10. • For carriers: Visit http://support.huawei.com and choose Product Support > Carrier IT > Cloud Computing > FusionCloud > FusionSphere > FusionCompute > V100R006C10.
FusionCompute V100R006C10_VRM.zip	VRM VM template file	

FusionCompute Installation Data

- Data to be prepared for installing FusionCompute

Type	Parameter	Description	Example Value
Host	Host name	Identifies a server in the system.	CNA01
	Management plane IP address	Specifies the IP address of the host management network port.	10.1.0.205
	Management plane subnet mask	Specifies the subnet mask of the host management plane.	255.255.255.0
	Management plane gateway	Specifies the gateway address of the host management plane.	10.1.0.1
VRM	VRM node name	Identifies a VRM node in the system.	VRM01
	Management plane IP address	Specifies the IP address of the VRM management interface when VRM nodes are deployed on physical servers. Specifies the IP address of the VM on which the VRM is deployed.	N/A
	Management plane subnet mask	Specifies the subnet mask of the VRM management plane.	N/A
	Management plane gateway	Specifies the gateway address of the VRM management plane.	10.1.0.1
	Floating IP address	Specifies the floating IP address of the VRM management plane.	10.1.0.101
Shared storage Note: This data is not required if FC SAN storage devices are used.	Storage management IP address	Specifies the management IP address of the storage device.	192.168.60.4
	Storage port IP address	Specifies the storage port IP address of the storage device.	172.20.100.100
	Storage device vendor	The value can be Huawei or Non-Huawei.	Huawei

Huawei Training & Certification

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. **FusionCompute Deployment**
   - Solution Overview and Process
   - Environment Preparations
   - Manual Installation

Installing a Host Using an ISO Image (1)

- Log in to the iBMC web portal of a server.
  - Open the browser on the local PC, enter http://BMC IP address of the host to be installed or https:// BMC IP address of the host to be installed in the address bar of the browser, and press Enter.
  - For the BMC system, the default username and password are root and Huawei12#$, respectively.

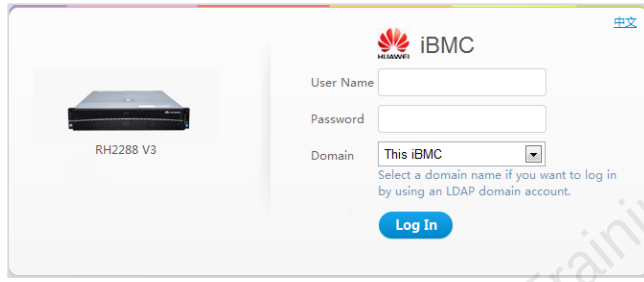

Installing a Host Using an ISO Image (2)

- Select Remote Virtual Console.

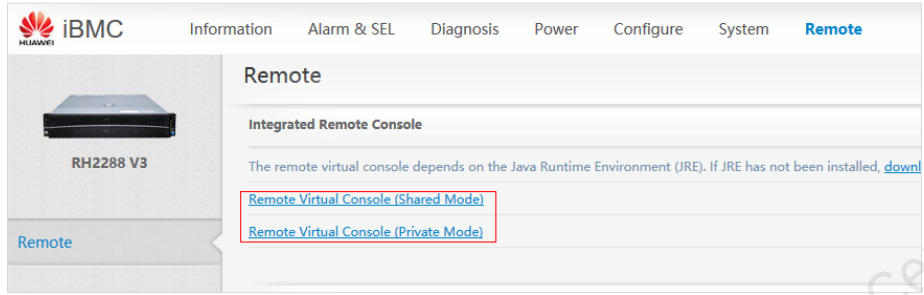

Installing a Host Using an ISO Image (3)

- Log in to the KVM virtual console, manually mount the ISO image package, click connect, restart the server, repeatedly press F11, and select DVD boot.

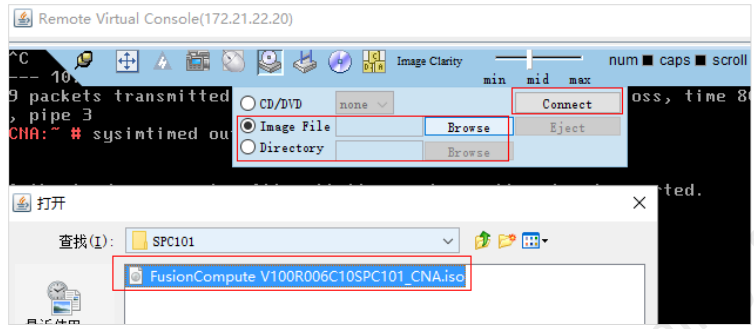

Installing a Host Using an ISO Image (4)

- Select **Install**.

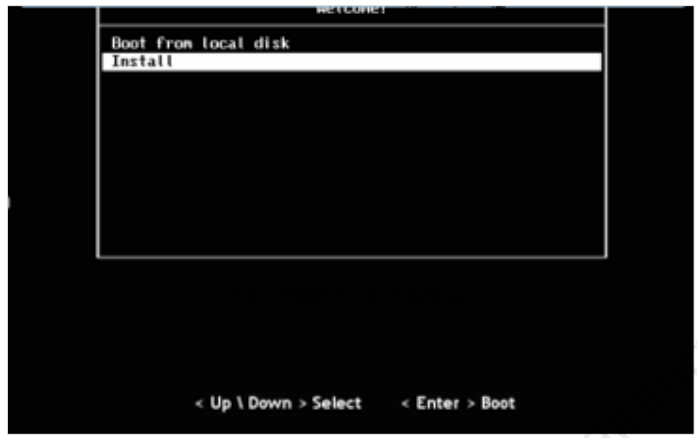

Installing a Host Using an ISO Image (5)

- Configure the host.
  - Modify the hard driver, network, and host name as planned.

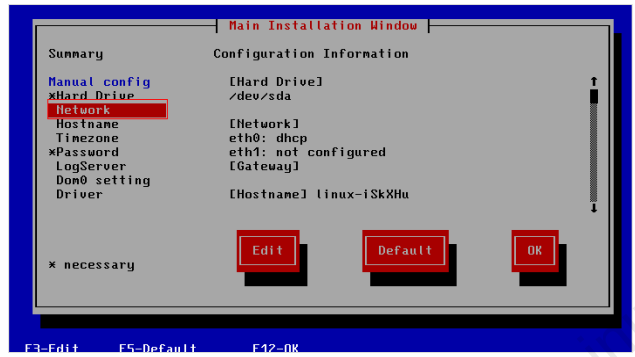

Installing a Host Using an ISO Image (6)

- Modify the network information, as shown in the figure.
- After the information is modified, select **OK**.

Installing VRM (1)

- Log in to the FusionSphere OpenStack web client using the reverse proxy IP address.
  - https://x.x.x.x:8890/#/login

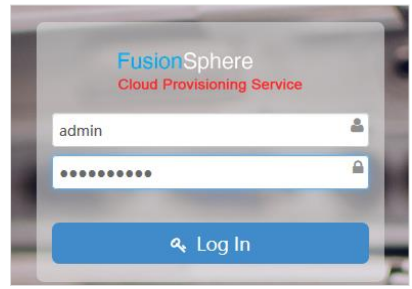

Installing VRM (2)

- Upload the VRM VM template software package.
  - Choose **Cloud Service > FusionCompute**.

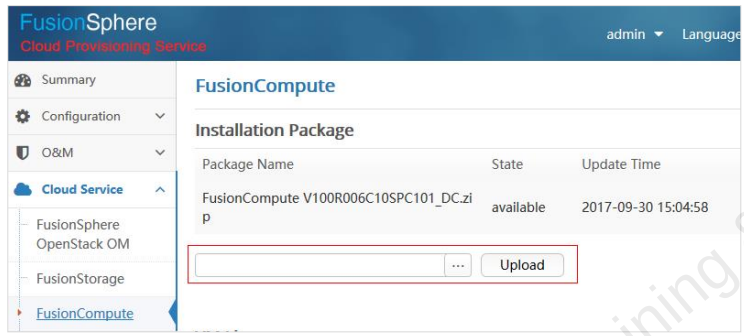

Installing VRM (3)

- After the VM template software package is uploaded, click **Create** and set VRM VM parameters.

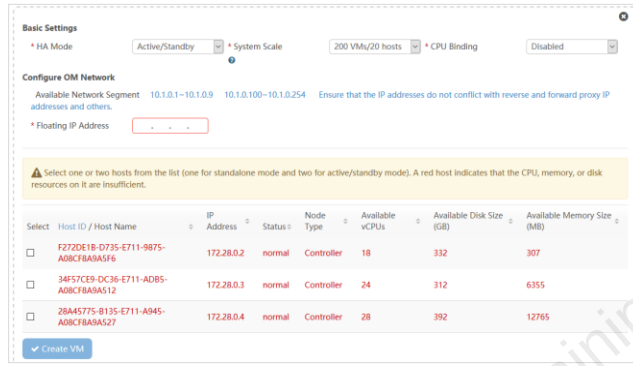

- Set VRM VM parameters as follows:
  - HA: Select Active/Standby. You are advised to use the active/standby deployment mode.
  - Management scale:
  - 200 VMs or less than or equal to 20 hosts: 8 CPUs, 16 GB memory, and 80 GB disk capacity
  - 1000 VMs or less than or equal to 50 CNA hosts: 12 CPUs, 32 GB memory, and 80 GB disk capacity
  - CPU core-binding: Select Disable.
  - Floating IP address: 10.1.0.101. The IP address must reside on the available network segment of the External OM network. The fixed IP addresses of active and standby VRM nodes are automatically assigned by the External OM network plane.
- Select two hosts for deploying VRM VMs.

Completing the Installation

- After the installation progress reaches 100%, click the link to switch to the VRM configuration page.

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. FusionCompute Deployment
4. **FusionSphere OpenStack Interconnection Configuration**
5. Initial Configuration of FusionCompute
6. Interconnection of FusionCompute and FusionSphere OpenStack
7. FusionCompute O&M

Interconnection of FusionCompute and FusionSphere OpenStack (1)

- Data to be prepared

Parameter	Description	Example Value
Keystone URL	Keystone is a rights management service provided by FusionSphere OpenStack. Keystone URL is in format of https://identity.Keystone Domain name:Port number/identity/v3.	https://identity.az1.dc1.domainname.com:443/identity/v3
Keystone IP	Log in to the FusionSphere OpenStack web client, choose Configuration > System > Reverse Proxy, and locate the system proxy service IP address on the External OM plane.	192.168.150.100
Region Name	This parameter indicates the local domain name prefix for connecting FusionSphere OpenStack.	az2.dc2
neutron password	neutron is the FusionSphere OpenStack Neutron service authentication account. The default password is FusionSphere123.	FusionSphere123
MQ IP	The RabbitMQ service IP address needs to be configured when you configure network information.	
MQ password	This parameter indicates the password of user rabbit in FusionSphere OpenStack. The default password is FusionSphere123.	FusionSphere123
Neutron MQ IP	This parameter indicates the RabbitMQ IP address used by Neutron.	128.6.35.125
Neutron MQ password	This parameter indicates the password of user rabbit used by Neutron. The default password is FusionSphere123.	FusionSphere123
FNM password	This parameter indicates the password of user FNM in FusionSphere OpenStack. The default password is FusionSphere123.	FusionSphere123

Huawei Training & Certification

Interconnection of FusionCompute and FusionSphere OpenStack (2)

- Data to be prepared

Parameter	Description	Example Value
Nova URL	The format is https://compute.Local domain name:Port number/v2.1.	https://compute.az2.dc2.domainname.com:443/v2.1
Nova IP	Log in to the FusionSphere OpenStack web client, choose **Configuration > System > Reverse Proxy**, and locate the system proxy service IP address on the External OM plane.	192.168.150.100
Neutron URL	The format is https://network.Local domain name:Port number/v2.0.	https://network.az2.dc2.domainname.com:443/v2.0
Neutron IP	Log in to the FusionSphere OpenStack web client, **choose Configuration > System > Reverse Proxy**, and locate the system proxy service IP address on the External OM plane.	192.168.150.100
FNM IP	FNM is a network O&M service in FusionSphere OpenStack and used to collect the network information of the FusionCompute hosts in the scenario where FusionCompute interconnects with FusionSphere OpenStack.	192.168.151.60

Interconnection of FusionCompute and FusionSphere OpenStack (3)

- On FusionCompute, choose System > Connect To > OpenStack.
- Select OpenStack and configure related information as planned.

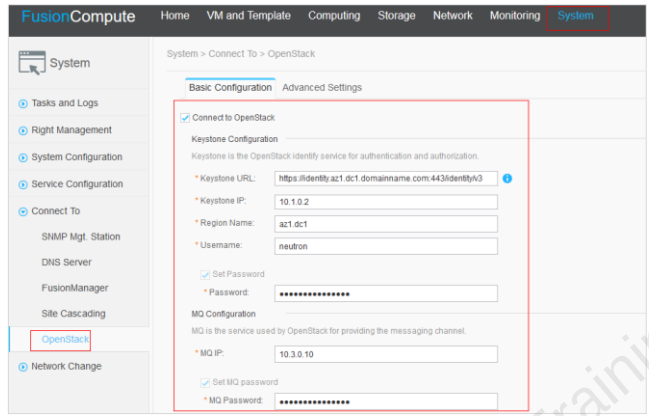

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. FusionCompute Deployment
4. FusionSphere OpenStack Interconnection Configuration
5. **Initial Configuration of FusionCompute**
6. Interconnection of FusionCompute and FusionSphere OpenStack
7. FusionCompute O&M

Initialization Configuration Process

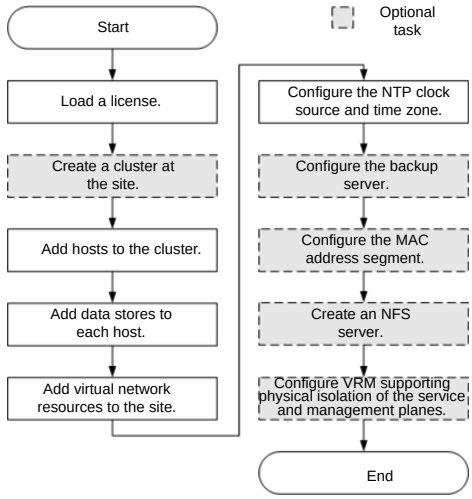

- Note
  - During the project implementation period, you can apply for a license for commissioning. The project without a license file is allowed for 90 days of trial use.
  - When you configure a host, bind its ports in active/standby or load sharing mode and configure port aggregation on the switch.
  - FusionSphere supports multiple storage devices, including IP SAN, FC SAN, NAS, FusionStorage, and local disks.
  - IP SAN storage requires host storage ports. Multipathing is recommended rather than NIC binding.
  - The Network Time Protocol (NTP) clock source is mandatory. If domain names are required, deploy a DNS server.

Uploading a FusionCompute License

- Log in to FusionCompute.
  - On FusionCompute, choose System > System Configuration > License Management.
  - Click Load License File.
- Load a new license.
  - Select License.
  - Click Obtain ESN (Make a note of the ESN.)
  - Apply for a license file using the ESN.
  - Click Browse to the right of Upload.
  - In the displayed dialog box, select the prepared license file and click OK.
  - In the displayed dialog box, click OK. The new license is loaded.

Host and Cluster Configuration

- Host and cluster configuration involves the following operations: create clusters, add hosts to them, and adjust host or cluster configurations.

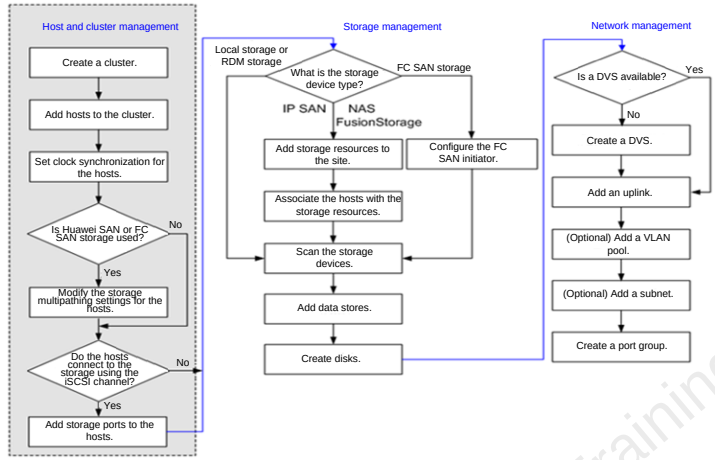

- In the FusionCompute hypervisor, resources pools are created in a cluster way.
- VMs created on computing clusters can take the advantage of multiple functions of FusionCompute.

Host and Cluster Configuration - Creating a Cluster (1)

- Method 1: Right-click the site and select **Create Cluster**.
- Method 2: On the **Getting Started** page of the site, click **Create Cluster**.

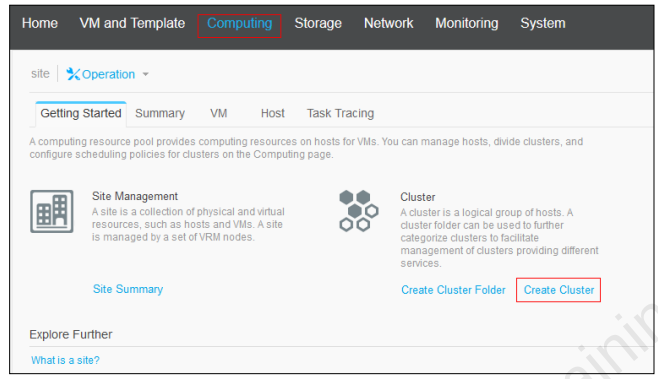

Host and Cluster Configuration - Creating a Cluster (2)

- Configure the basic information for the cluster and set the resource scheduling policy.

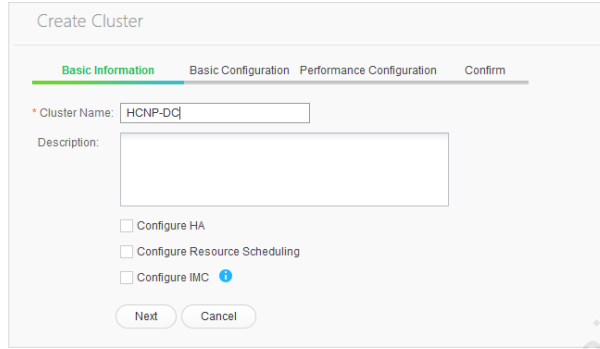

Host and Cluster Configuration - Adding Hosts (1)

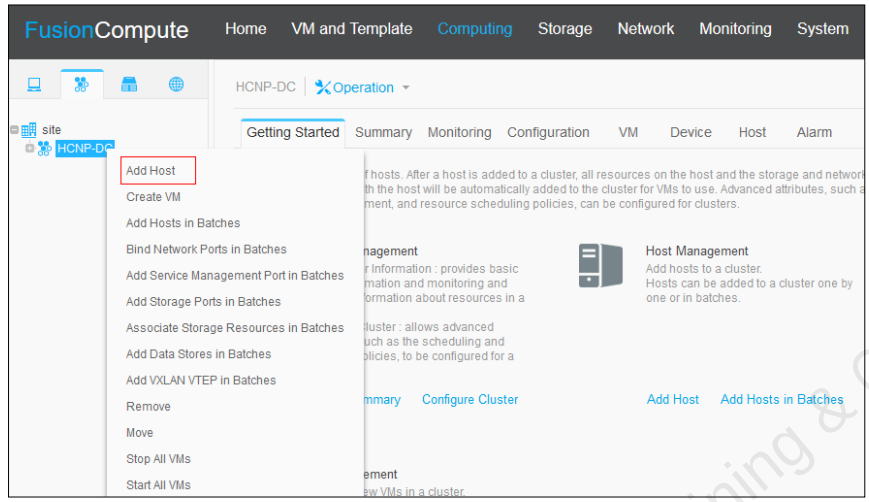

Host and Cluster Configuration - Adding Hosts (2)

- To add a host, set the following parameters: name, IP address, description, BMC IP address, username, and password.
- Manually add hosts one by one or add hosts in batches using a template.
- Manually add hosts: Only one host can be added at a time. This method is easy to operate and is recommended when a small number of hosts are to be added.
- Add hosts in batches: All information about the hosts can be set in a template and imported at a time. This method is recommended when a large number of hosts are to be added.

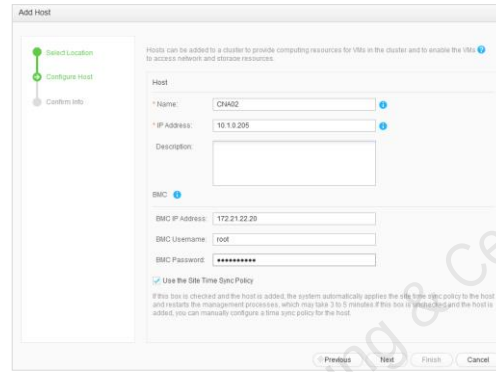

- BMC: Baseline Management Center
- BMC is a server management module developed by Huawei. If an IP address is configured for BMC network ports, administrators can remotely control and manage the server.

Host and Cluster Configuration - Network Port Binding

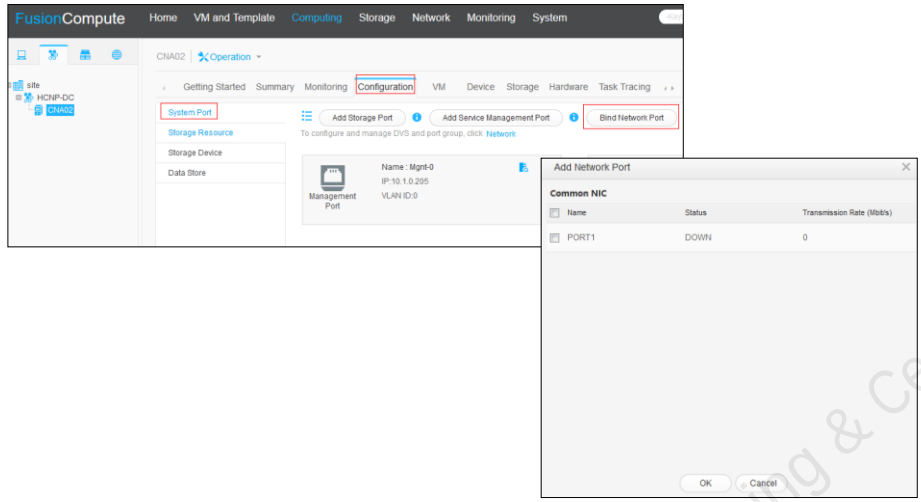

- To improve network reliability, you are advised to bind network ports for the management plane and the service plane, respectively.

- The following binding modes are available for common NICs: Active/Standby: applies to scenarios where two network ports are to be bound. This mode provides high reliability and the bound port bandwidth is equivalent to that of a member port.
- Source and destination MAC addresses-based load sharing: applies to scenarios where two or more network ports are to be bound. The bandwidth of the bound port in this mode is higher than that of a member port because the member ports share workloads based on the MAC addresses of the source and destination ports.
- This mode is recommended when most network traffic is on the layer 2 network. This mode allows network traffic to be evenly distributed based on MAC addresses.
- Polling: applies to scenarios where two or more network ports are to be bound. The bandwidth of the bound port in this mode is higher than that of a member port, because the member ports share workloads in sequence.
- This mode may result in data packet disorder because traffic is evenly sent to each port. Therefore, MAC-based load sharing prevails over Polling in load sharing modes.

Host and Cluster Configuration — Adding a Storage Interface to a Host (1)

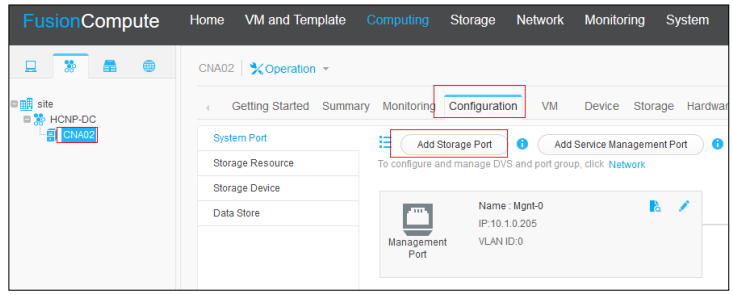

- Adding storage interfaces to the host: Interconnect the host with the storage device. Multiple storage interfaces added to the host can achieve multipathing transmission.
- Modes of adding storage interfaces: Manually add storage interfaces and add storage interfaces in batches.
- Parameters: includes the name, description, IP address, subnet mask, VLAN ID, bandwidth configuration, and upper bandwidth limit.

Host and Cluster Configuration — Adding a Storage Interface to a Host (2)

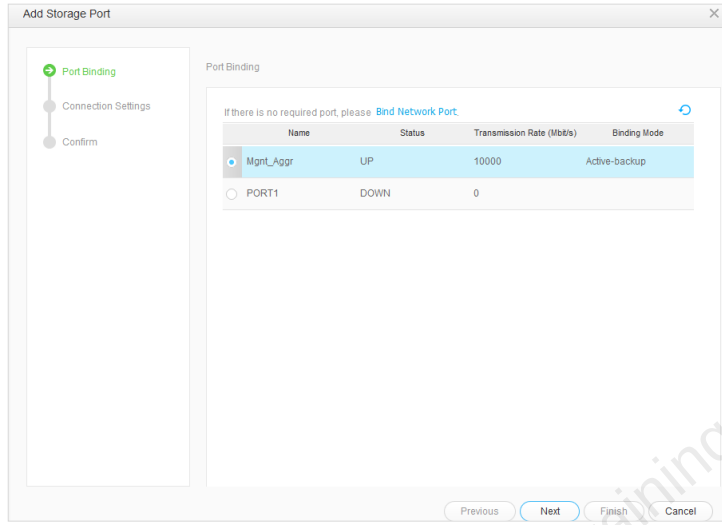

Host and Cluster Configuration — Adding a Storage Interface to a Host (3)

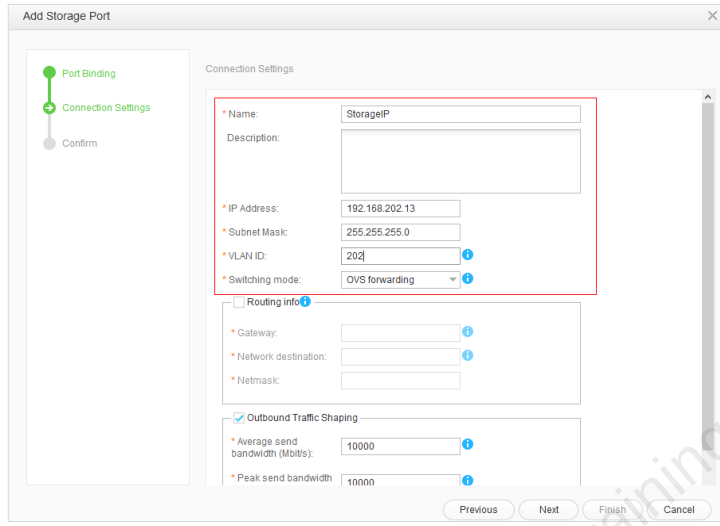

Initial Configuration of the Storage

- You can manage storage resources, data stores, and disks.

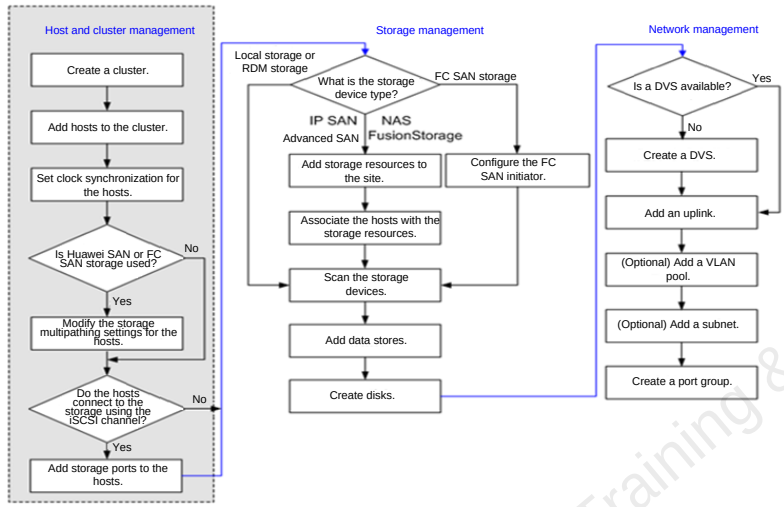

- In the FusionCompute hypervisor, FusionCompute provides computing clusters with shared storage to support functions, such as VM live migration, dynamic resource scheduling (DRS), and dynamic power management (DPM).

Storage Configuration — Configuring Storage Multipathing

- The system uses the universal multipathing type by default. If non-Huawei storage is in use, skip this task.
- If Huawei storage is in use, such as Huawei iSCSI and FC SAN storage, set the multipathing type to Huawei to improve storage multipathing performance.

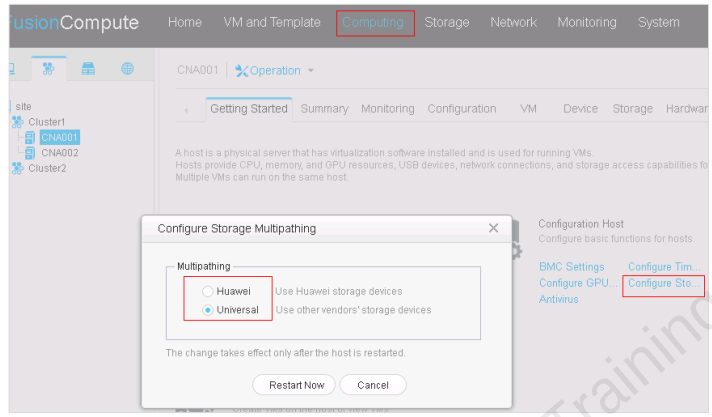

Storage Configuration — Adding a Storage Resource (1)

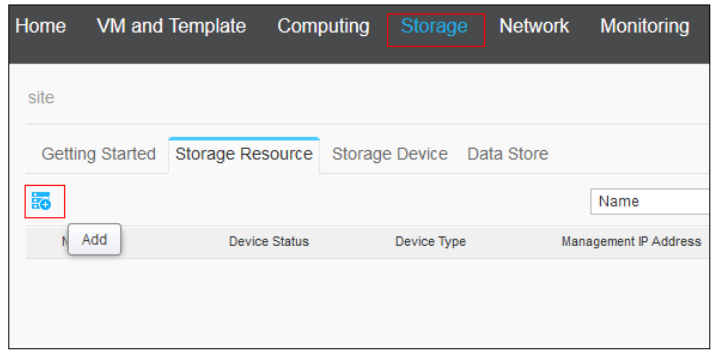

- If all hosts in the system use FC SAN storage or local disks they each provide, skip this task.

Storage Configuration — Adding a Storage Resource (2)

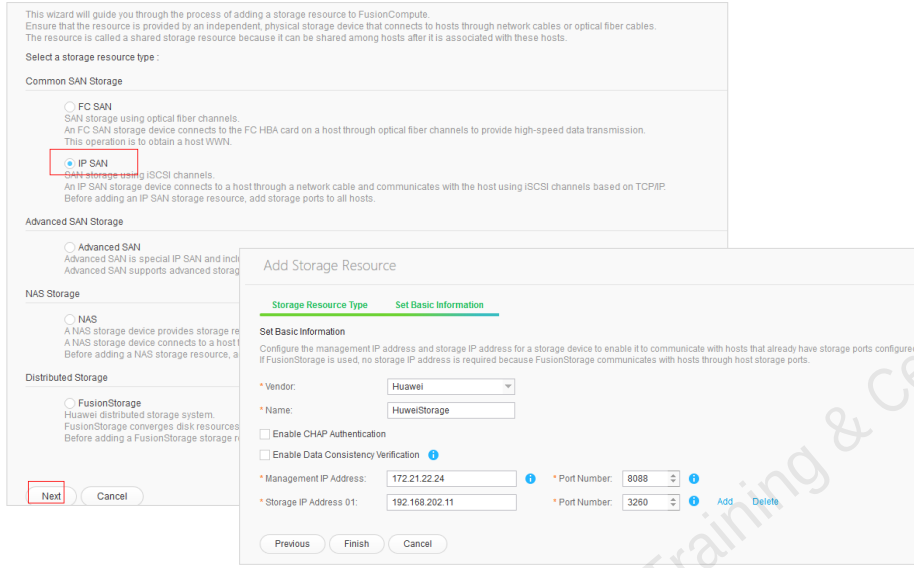

Storage Configuration — Associating a Storage Device with a Host

- Right-click the target storage device and select Associate Host.

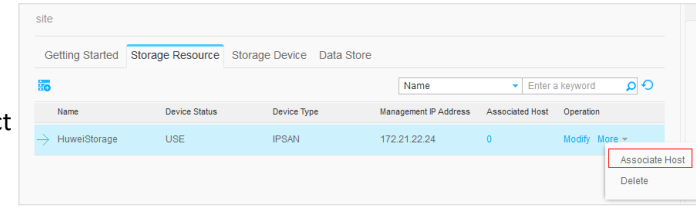

- Select the host.

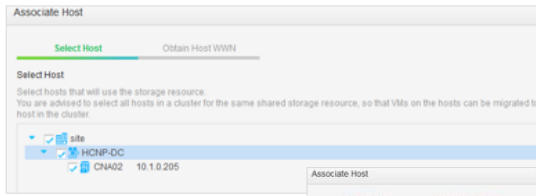

- Obtain the WWN.

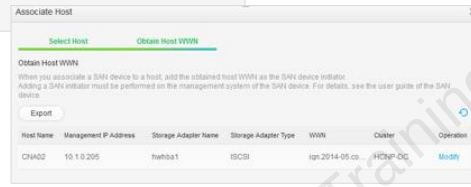

Creating LUNs on a Storage Device

- Log in to OceanStor DeviceManager.
- Choose **Provisioning > LUN**, create LUNs, set the FusionCompute capacity to 500 GB, and select **StoragePool002**.

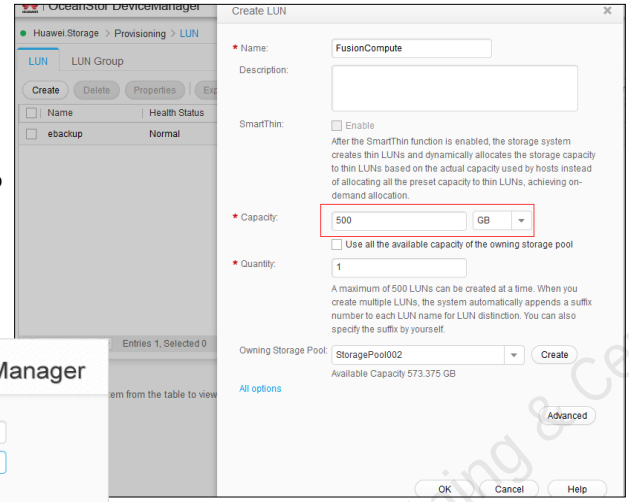

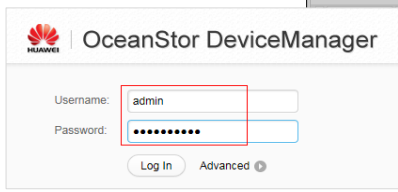

Creating a LUN Group on a Storage Device

- Click LUN Group to create LUN group FusionCompute.

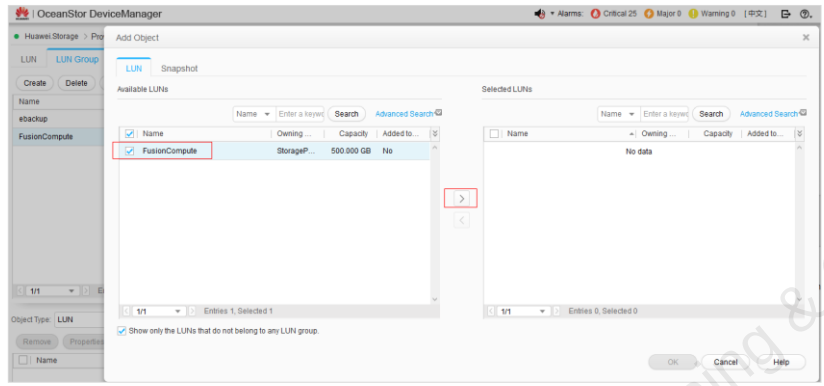

Creating a Host on a Storage Device

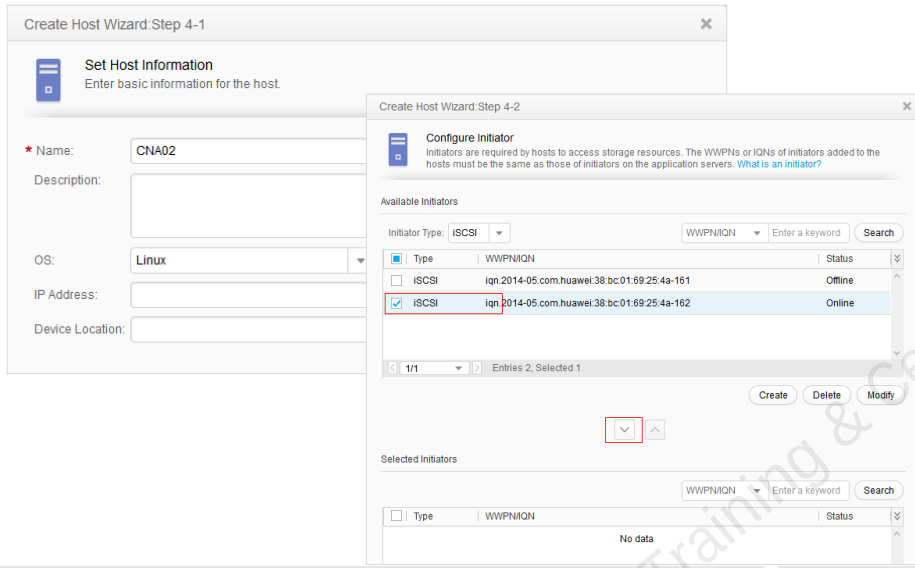

Creating a Host Group on a Storage Device

- Click Host Group to create host group FusionCompute.

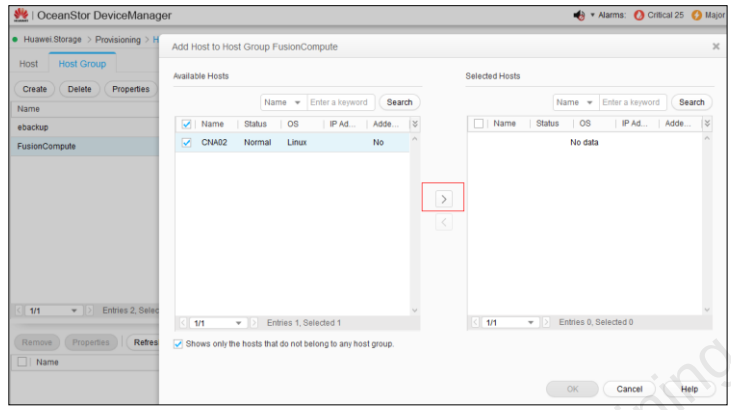

Creating a Mapping View on a Storage Device

- On the Mapping View page, select the LUN group and host group.

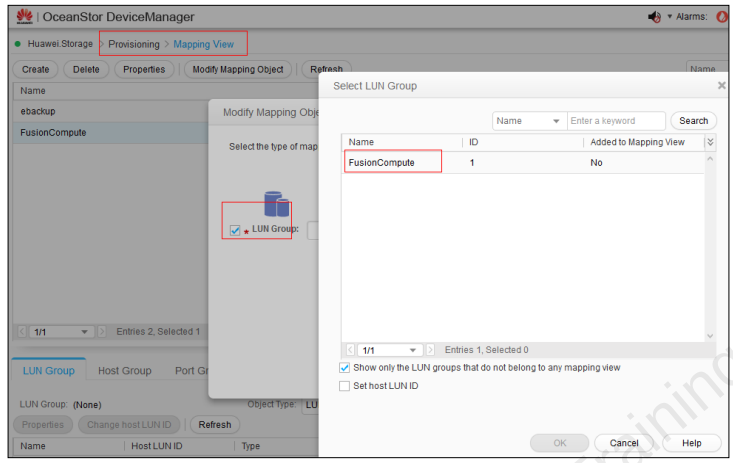

Storage Configuration — Scanning Storage Devices

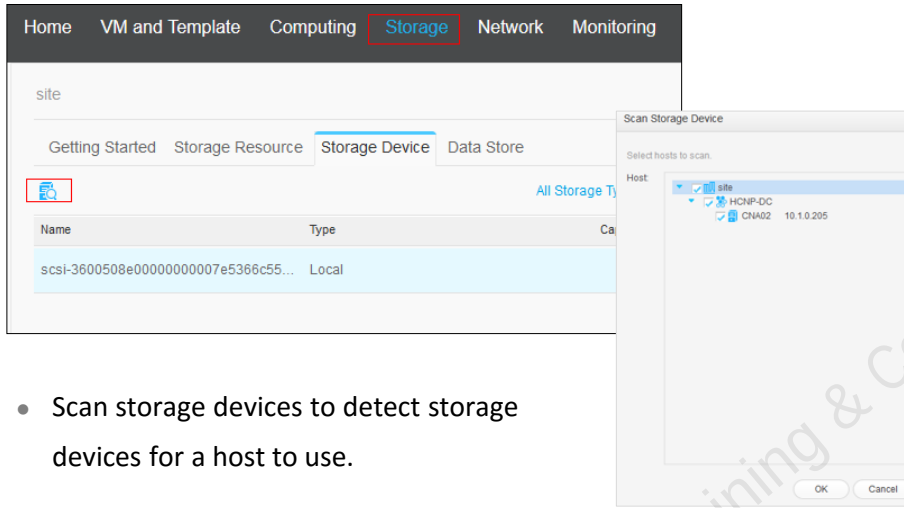

- Scan storage devices to detect storage devices for a host to use.

Storage Configuration — Adding Data Stores

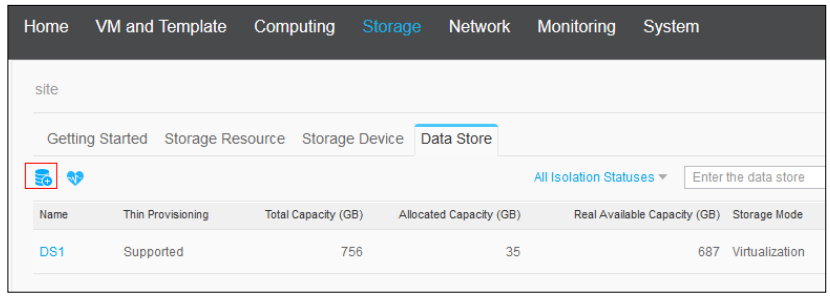

HUAWEI

- Add data store: Add a data store to a host so that VM disks can be created on the added data store. One host can carry multiple data stores, and one data store can be added to multiple hosts. A VM can be migrated between two hosts only when the data store providing the disks of the VM is added to both the hosts.
- Data store type: Virtualization, non-virtualization, and RDM storage (shared storage)
  - A virtualized data store supports advanced storage features, such as thin provisioning disks, snapshots, and live storage migration, and provides high resource utilization. However, common disk creation on it is slow.
  - A non-virtualized data store does not support advanced storage features. However, disks can be quickly created on non-virtualized data stores and provide high performance.
  - A RDM data store is a physical LUN on a SAN storage resource and is directly attached to a VM as a disk, which improves performance of the SAN storage and supports the SCSI. Only one disk can be created on an RDM data store and cannot be divided into partitions. Therefore, An RDM data store does not support advanced storage features.

Storage Configuration — Configuring Data Stores

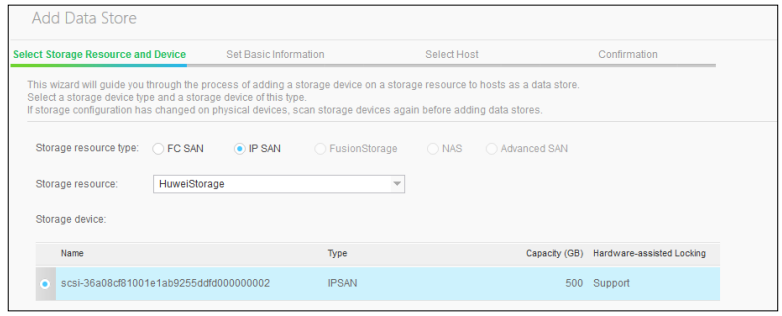

- The data store is classified into virtualization storage, non-virtualization storage, and raw device mapping (RDM) storage.
- Multiple data stores can be added to a host, and a data store can be added to multiple hosts. A VM can be migrated between two hosts only when the data store providing the disks of the VM is added to both hosts.
- If the data is stored in virtualization mode, it cannot be switched to non-virtualization mode.

Initial Configuration of the Network

- The network management provides network resources, such as creating a distributed virtual switch (DVS) and port group. You can configure and adjust network resources using the network management.

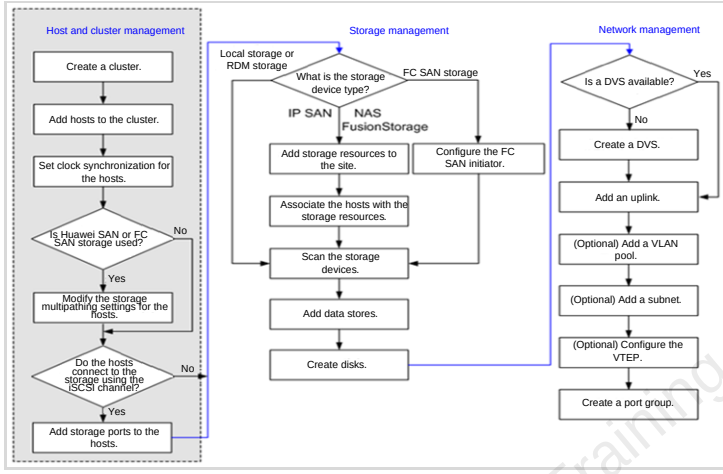

- In the FusionCompute hypervisor, a VM has networks added to provide network functions.

Network Configuration — Creating a DVS (1)

- In the downlink direction, the DVS connects to VMs through virtual ports. In the uplink direction, the DVS connects to physical Ethernet adapters on hosts where VMs reside. VMs connect to outside devices through DVSs using the physical NICs of the hosts accommodating the VMs.

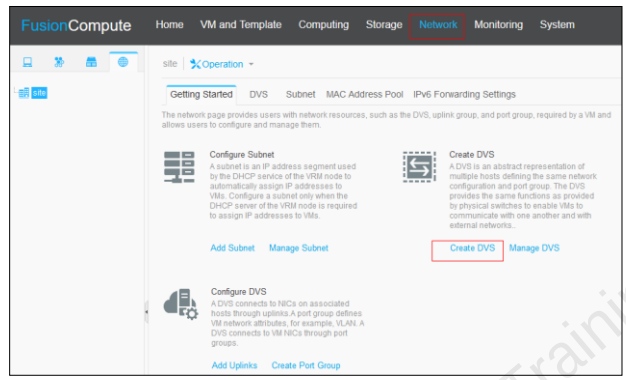

Network Configuration — Creating a DVS (2)

- Create DVS **physnet1**, associate the DVS with physical network **physnet1**, and select **Add Uplink**.

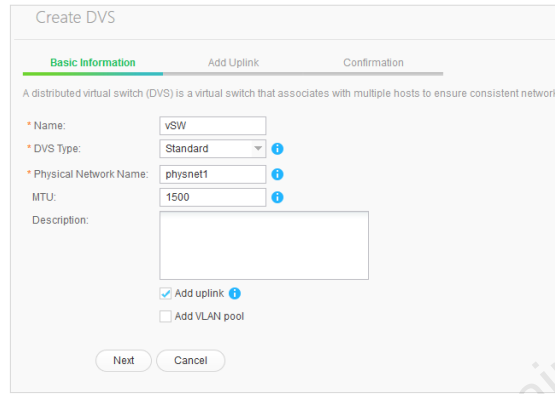

Network Configuration - Adding an Uplink

- After the uplink is added, click Next. Create a DVS.

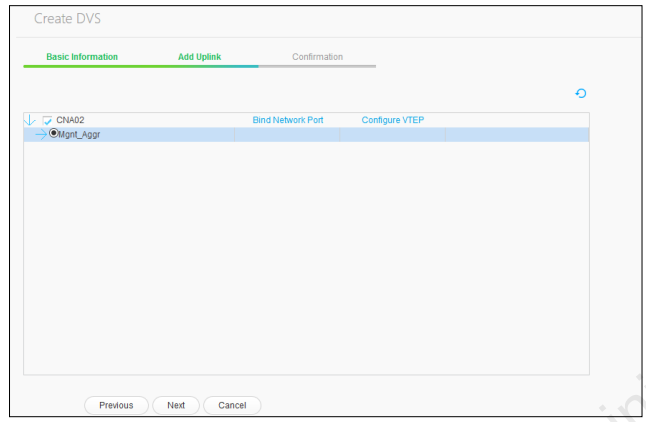

- An uplink is added to a DVS to connect the DVS to a physical switch so that VMs connected with the DVS can access external resources through the Internet.

Configuring NTP

- This operation interrupts services on the host temporarily because the host service process must be restarted after the time synchronization function is configured. Therefore, exercise caution when performing this operation.
- You are advised to preferentially configure an external Network Time Protocol (NTP) clock source. If no external NTP clock source exists, configure the clock source as that for VRM hosts or VRM VMs depending on VRM deployment mode.
- If the system time synchronization is not configured, you must manually configure the time of each host, ensuring that the time difference between all hosts and VRM nodes is within 1 minute. The FusionCompute cannot monitor hosts with the time difference longer than one minute.
- The configured system time zone is the zone to which the alarm information display time belongs. The time displayed in FusionCompute is the time of the zone configured for the browser.

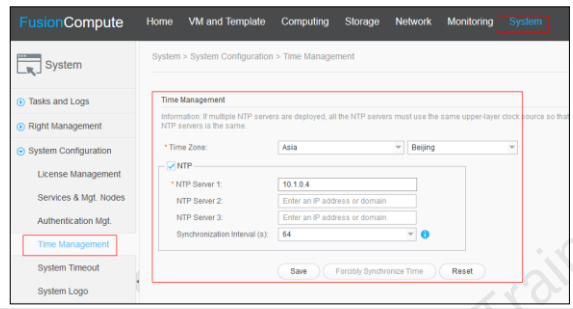

- The NTP IP address must be obtained in advance.

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. FusionCompute Deployment
4. FusionSphere OpenStack Interconnection Configuration
5. Initial Configuration of FusionCompute
6. **Interconnection of FusionCompute and FusionSphere OpenStack**
7. FusionCompute O&M

Creating a User for Interconnecting with FusionSphere OpenStack

- In FusionCompute, create a user for interconnecting with FusionSphere OpenStack.
  - On Fusion**System > Right Management > Role Management**Compute, choose.
  - Click **Add Role**.
  - Select **Interface interconnection user**.
  - Set the username and password.
  - Select **administrator** from the Subrole.
  - Confirm information and click **OK**.

Interconnection of FusionCompute and FusionSphere OpenStack — Preparing Data

- Data prepared for interconnecting FusionCompute with FusionSphere OpenStack

Parameter	Description	Example Value
FusionCompute service IP address	Specifies the management IP address of the VRM node If VRM is deployed in standalone mode or the floating IP address of the VRM nodes if VRM is deployed in active/standby mode. If this parameter is modified on FusionCompute, configure this parameter on FusionSphere OpenStack again.	192.168.60.10
FusionCompute interconnection user	Specifies the interface interconnection user used to connect FusionCompute to FusionSphere OpenStack. This user must be assigned the administrator role and is used for configuring the Cinder and Nova plug-ins for the interconnection. If this parameter is modified on FusionCompute, configure this parameter on FusionSphere OpenStack again.	VRM01_HW
FusionCompute interconnection password	Specifies the password used for connecting FusionCompute to Cinder and Nova plug-ins. If this parameter is modified on FusionCompute, configure this parameter on FusionSphere OpenStack again.	Huawei@123
Confirm interconnection password	Specifies the confirm password of the interconnection username.	-

Configuring FusionCompute to Interconnect with FusionSphere OpenStack (1)

- Connect to the FusionCompute resource pool.
  - On the FusionSphere OpenStack web client, choose Configuration > Resource Pool.
  - In the Add Resource Pool area in the right pane of the Resource Pool page, click + to the right of FusionCompute to enter the add FusionCompute Resource Pool page.
  - Configure FusionCompute-related parameters.

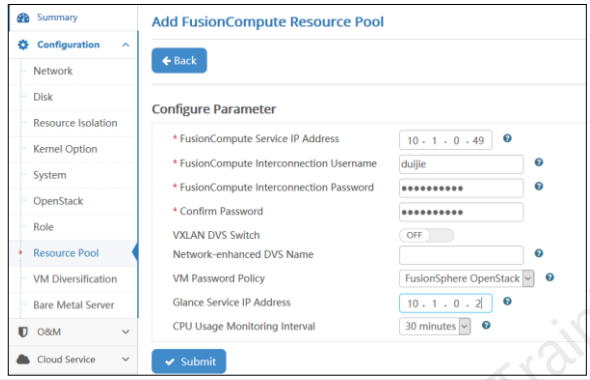

Configuring FusionCompute to Interconnect with FusionSphere OpenStack (2)

- Add a FusionCompute computing cluster.
  - Click + to the right of the newly added FusionCompute resource pool to enter the Cluster Management page.
  - Click  n the Compute Cluster area to enter the Add Compute Cluster page.
  - Set the cluster name in Configuration Parameter.
  - Select two hosts in Select Hosts for deploying the fc-nova-compute role in active/standby mode.
    - The fc-nova-compute role is used to manage the configured cluster. An fc-nova-compute role can manage one or multiple FusionCompute clusters and is presented as one host.
    - You can add FusionCompute clusters to an AZ multiple times. An fc-nova-comput role is automatically generated each time. These roles are named starting from fc-nova-compute001 in ascending numerical order, for example, fc-nova-compute001, fc-nova-compute002, and fc-nova-compute003. The host name presented to users are the same as the role name.
  - Click 

Configuring FusionCompute to Interconnect with FusionSphere OpenStack (3)

- Click ✚ in the Storage Cluster area to enter the Add Storage Cinder page.
- Configure the AZ in Configuration Parameter.
- Click + to the right of Backend storage configurationto set the data store name to the name of the data store used by the FusionCompute cluster.

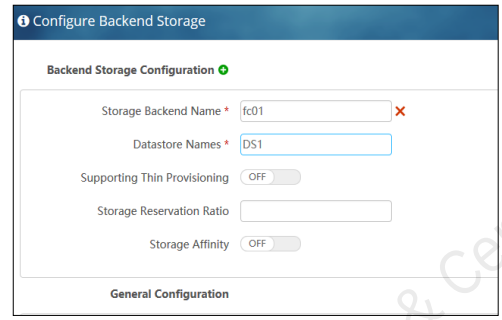

Configuring FusionCompute to Interconnect with FusionSphere OpenStack (4)

- Optional configurations of backend storage

Configuration Item	Description	Remarks
Storage thin provisioning	Specifies the thin mode of the FusionCompute data store. That means that the data store can be reused.	Enabled
Thin provisioning percentage	Sets the storage overcommitment ratio in thin provisioning mode. The default value is 1.0. The value needs to be greater than 1.0 when overcommitment is required. Both thin provisioning percentage and overcommitment ratio need to be configured.	For example, 2.0
Storage reservation rate	Specifies the space reservation rate (%) of the data store.	For example, 10
Storage affinity	Enables or disables the storage affinity.	Enabled
Storage affinity rate	In multi-storage backend scenarios, when multiple data stores corresponding to multiple storage backends are used, the data stores are prioritized based on their available storage space. You can set the affinity rate for some data stores. When you calculate the data store priority, multiply available storage space by the affinity rate to change the data store priority. The value is an integer greater than 1.	For example, 5
Affinity list	Specifies the data stores used for setting the application storage affinity rate.	ds1, ds2

Configuring FusionCompute to Interconnect with FusionSphere OpenStack (5)

- After the configuration is complete, select the host in Select Hostsfor deploying Cinder. Select three hosts for interconnecting with a storage cluster, and click submit.
- Configure the process of downloading an image. The image is downloaded only on FusionCompute.
  - Use PuTTY to log in to the controller host in the cascaded FusionSphere OpenStack system (Ensure that the External OM IP address and username fsp are used to establish the connection). Run the su – root command to switch to user root.
  - Import environment variables.
  - Run the following command to add configuration item volume_download_image=False:
  - cps template-ext-params-add --service cinder cinder-volume-vrm001 --parameter DEFAULT.volume_download_image=False
  - cps commit
- After the preceding operations are performed, FusionCompute will be interconnected with FusionSphere OpenStack. After ensuring that FusionSphere is completely installed, configure and allocate resources (including host group and AZ division and scheduling policy configuration) by following operations in Operation and Maintenance > Administrator Guide > Initial Configuration in the FusionSphere Product Documentation.

Contents

1. FusionCompute Architecture
2. FusionCompute Features
3. FusionCompute Deployment
4. FusionSphere OpenStack Interconnection Configuration
5. Initial Configuration of FusionCompute
6. Interconnection of FusionCompute and FusionSphere OpenStack
7. **FusionCompute O&M**

VM Management — VM Creation Methods

Method	Description	Recommended Application Scenario
Creating a bare VM	A bare VM is like a bare physical server without an OS installed. When creating a bare VM, you can specify VM specifications, such as CPUs, memory, disks, and network interface cards (NICs). After the bare VM is created, you need to install an OS on it.	The FusionCompute system is initially installed. The OSs or specifications of existing VMs or VM templates in the system do not meet user requirements. A bare VM is required to be converted or cloned to a template for VM creation. Before cloning or converting it to a template, install an operating system (OS) on it.
Creating a VM using a VM template	A template is used to create a VM that has similar specifications to the template.	The OSs and specifications of templates in the system meet user requirements. The OS and specifications of a template in a peer site meet user requirements.
Creating a VM using a VM	An existing VM is used as the parent VM to clone a VM that has similar specifications to the parent VM.	Multiple similar VMs are required. You can install different software on one VM and clone it for multiple times to obtain required VMs.

- VMs can be created using OVF templates or templates converted from VMs.

VM Management - Creating a VM

- A VM is created based on the following procedure:
  - Choose a creation position.
  - Set the VM attributes.
  - Set NICs and disks.
  - Create a VM.

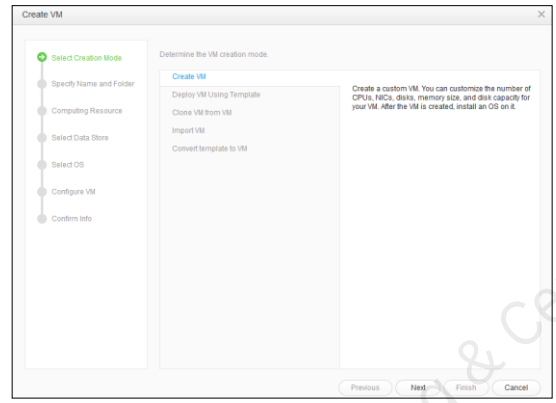

VM Management — Adjusting the VM Specifications

- You can modify CPU attributes, adjust memory specifications, expand disk capacity, attach a disk to a VM, detach a disk from a VM, add a NIC, and delete a NIC.

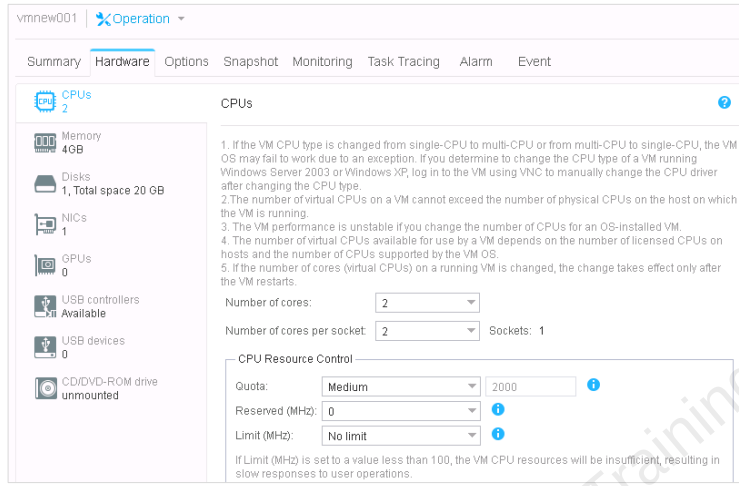

VM Management — VM Operation Management

- You can start, hibernate, wake up, stop, migrate, or delete a VM as well as bind a VM to a host.

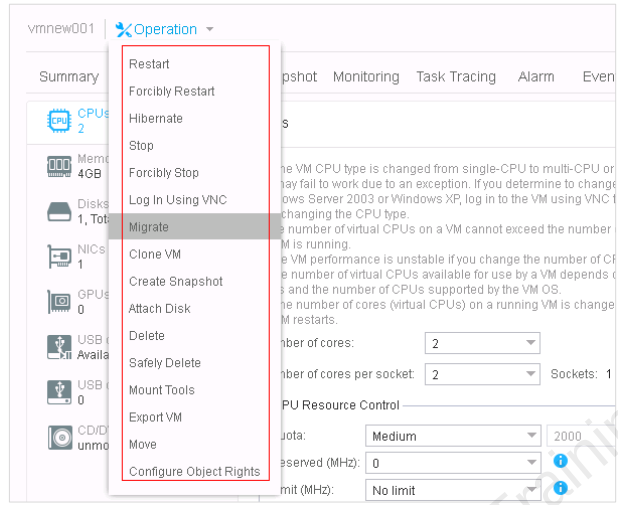

VM Management — Creating a VM Template

- Methods for creating a VM template: VM to template conversion, VM to template cloning, and template to template cloning
- Creating process
  - Ensure that a VM meets requirements for creating a template.
  - For details about Windows VMs, see VM SID configuration and VM specification configuration.
  - For details about Linux VMs, see VM specification configuration and VM NIC configuration.
  - Perform operations for template conversion.

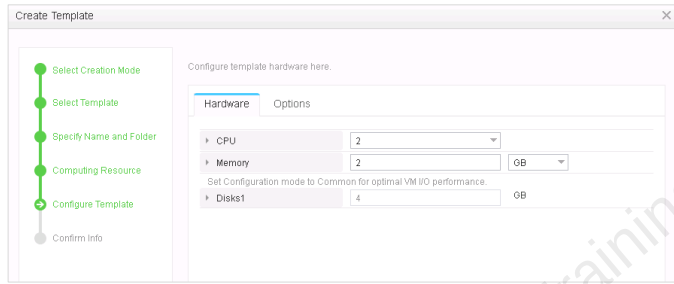

- Convert a VM to a VM template: All parameters of the template are from the VM. The VM is deleted after the converting is complete.
- Clone a VM to a VM template: All parameters of the template are from the VM, but you can customize some parameters during the cloning process. The VM is still available after the cloning is complete.
- Clone a VM template to a VM template: All parameters of the new template are from the original VM template, but you can customize some parameters during the cloning process. The parent template is still available after the cloning is complete.

VM Management — Creating a VM Snapshot

- VM snapshot: A snapshot is taken to save information of all disks on a VM at a certain point for restoring VM data. Multiple snapshots can be created for one VM. If a VM is being restored using a snapshot, other snapshots of the VM are not affected.

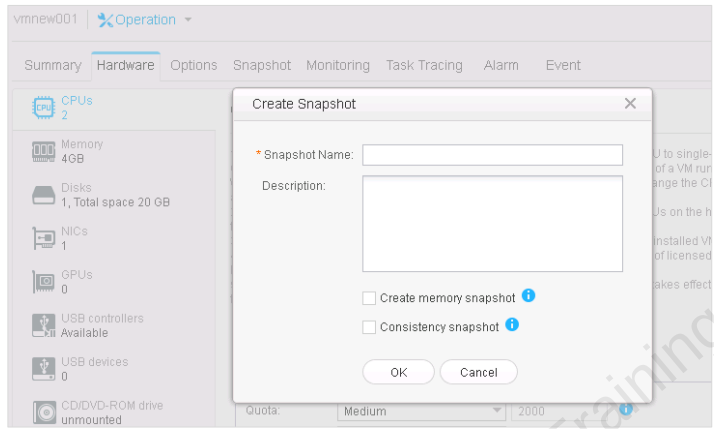

- The restrictions for creating a snapshot are as follows:
- Multiple snapshots can be created for one VM. If a VM is being restored using a snapshot, other snapshots of the VM are not affected.
- A maximum of 32 snapshots are supported by a VM. The created snapshots can be only manually deleted.
- Virtualization management nodes do not support snapshot creation, for example, VRM VMs.
- The VM snapshot function is supported only when the VM disk mode is set to Advanced SAN storage, FusionStorage, virtual local disks, virtual SAN storage, or NAS storage.
- A snapshot cannot be created because disks on the VM use multiple types of storage devices.
- You are not allow to create a memory snapshot for the VM that is bound with a sharing disk. Shared disks do not support snapshot creation during the creation of common and consistency snapshots.
- If a disk on a VM is set to not affected by snapshots, when you create snapshots, the system does not create snapshots for the disk.

VM Management — Restoring a VM Using a VM Snapshot

- Precautions for restoring a VM using a snapshot are shown as follows:
  - If you restore a VM using a snapshot, the VM must be stopped or hibernated.
  - If a new disk is attached to a VM after a snapshot is taken for the VM, the disk will be automatically detached from the VM after the VM is restored using this snapshot. If necessary, you can attach this disk to the VM again.
  - If a disk is not affected by snapshots, the disk data will not be restored when the VM is restored using a snapshot. If a disk is detached from a VM after a snapshot is created for a VM, the disk will be attached to the VM again after the VM is restored using the snapshot, however, the disk data cannot be restored.
  - If a disk is deleted after a snapshot is created for the VM, the disk will not be attached to the VM after the VM is restored using the snapshot.

- A hypervisor can be created based on the preceding content.
- This document relates to resource addition and initial configuration after the hypervisor is created. As the operation process involves many choices for enabling related functions, trainees need to understand the content based on the computing, storage, and network function introduction.

Quiz

1. Which of the following components are included in FusionCompute? ( )
   A. CNA
   B. DPM
   C. VRM
   D. DRS
2. True or false

   A CNA node consists of the UVP and VNAs and connects to the VRM node through a VNA node. ( )

- Answers:
  1. AC
  2. True

Recommendations

- Huawei E-Learning website:
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei support case library:
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Cloud Data Center FusionStorage Block Distributed Storage Principles and Architecture

www.huawei.com

Foreword

- This course introduces the functions and features, logical architecture, and installation and deployment of FusionStorage Block.

Objectives

- Upon completion of this course, you will be able to:
  - Describe Huawei FusionStorage Block.
  - Understand the functions and features of FusionStorage Block.
  - Master and describe the logical architecture of FusionStorage Block.
  - Master the FusionStorage Block installation in different scenarios.
  - Master the FusionStorage Block configurations.

Contents

1. **FusionStorage Block Product Overview**
2. FusionStorage Block Functions and Features
3. FusionStorage Block Logical Architecture
4. FusionStorage Block Software Installation
5. FusionStorage Block Configuration

Getting to Know Server SAN

Concept

- Server SAN is a storage resource pool that is made up of resources provided by multiple independent servers, and the Server SAN resource pool integrates both computing and storage resources.

Characteristics

- General-purpose devices replacing dedicated devices
- Linear expansion of computing and storage resources
- Simplified management and low TCO

- Software and hardware decoupling: Traditional storage software is highly-integrated with hardware devices. Server SAN products are separated with the hardware devices and not bound to any specified hardware devices.
- Converged storage and computing: Server SAN products are deployed on generic x86 servers, achieving computing and storage convergence.

Huawei Server SAN Product - FusionStorage Block

- Distributed block storage software.
- Employs distributed technologies to organize hard disk drives (HDDs) and solid-state drives (SSDs) of generic x86 servers into large-scale storage resource pools.
- Provides standard Small Computer System Interface (SCSI) and Internet Small Computer System Interface (iSCSI) interfaces for upper-layer applications and VMs.
- Open APIs.

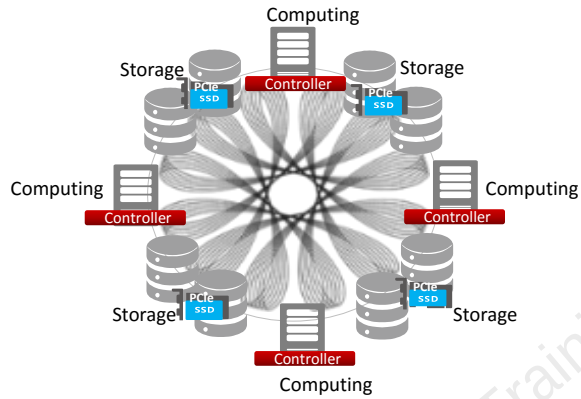

- FusionStorage Block supports the typical traditional block storage application scenarios. For example, various service applications, such as SQL, Oracle RAC, Web, and industry applications.
- FusionStorage Block can be integrated with mainstream cloud platforms, including Huawei FusionSphere, VMware vSphere, and open source OpenStack, allocating storage resources on demand.
- FusionStorage Block is the only commercial Server SAN product that supports PB-level capacity.

Traditional SAN Architecture

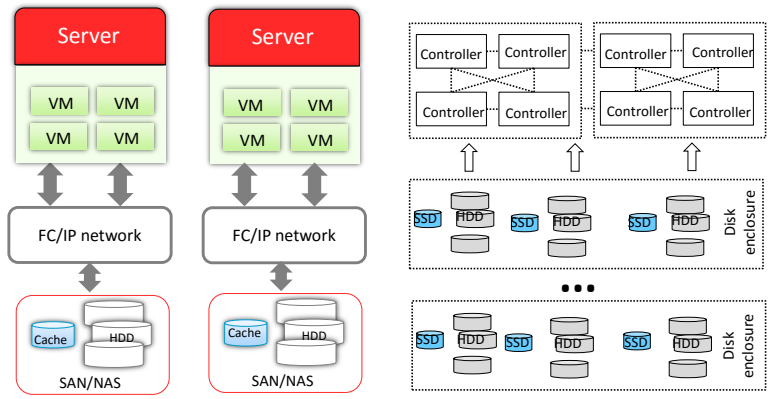

Isolated storage resources

Storage devices that support scale-up only

- Engine bottlenecks: 2 to 16 controllers, linear expansion not supported
- Cache bottlenecks: Normally in the unit of GB
- Network bottlenecks: 10GE and 8 GB FC networks

- Lack of resource sharing: Traditional storage devices and resources are usually provided by different vendors. Therefore, storage resources cannot be shared, and are isolated from each other in the data center.
- Centralized metadata management: Traditional storage usually employs a centralized metadata management mechanism, which records the distributions of data offsets in all Logic Unit Numbers (LUNs) in disks. For example, a 4-KB data identified by an address starting with LUN1+LBA1 is stored on LBA2 of the 32nd disk. On a traditional storage device, each I/O operation initiates a query request to the metadata service. As the system scale grows, the metadata size also increases. However, the concurrent operation capability of the system is subject to the capability of the server running the metadata service. As a result, the metadata service eventually becomes a performance bottleneck of the system.

Distributed Server SAN Architecture

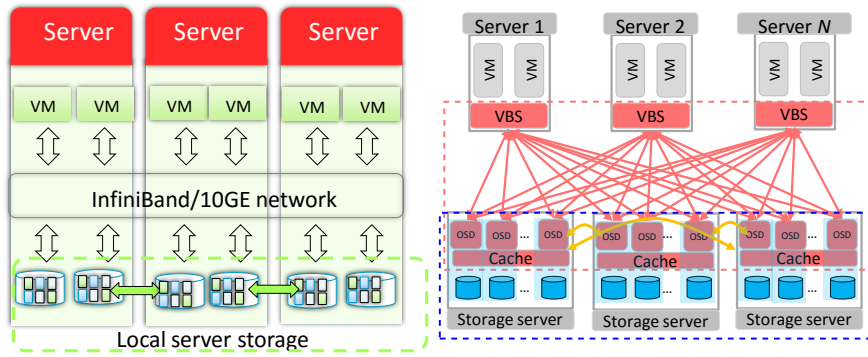

- Distributed controller: Supports linear expansion to up to 4096 nodes.
- Distributed cache: Supports capacity expansion to the TB level.
- Network: Supports point-to-point (P2P) high-speed InfiniBand network and 56 GB InfiniBand RDMA network.

- Data center level resource sharing: One huge storage resource pool can be built in one data center to meet the capacity, performance, and reliability requirements for various applications and implement resource sharing and centralized management.
- Storage investment options in cloud data centers
- VBS and OSD will be described in later slides.
- The DHT algorithm has the following characteristics:
  - Balance: Data is evenly distributed to all nodes as possible, thus balancing load among all nodes.
  - Monotonicity: When new nodes are added to the system, the system will redistribute the data. In this process, data will only be migrated to new nodes, and data on the existing nodes will not be greatly adjusted.
- FusionStorage Block uses a distributed architecture to organize the dispersedly distributed, low-efficiency SATA or SAS disks into an efficient storage pool that provides higher I/O throughput than SAN devices, maximizing the storage performance and improving performance with 1 to 3 folds.
- FusionStorage Block uses distributed stateless engines. These engines are deployed on each server that requires access to FusionStorage Block, thereby preventing performance bottlenecks that may be caused by traditional, centrally deployed engines. Moreover, these distributed engines deployed on standalone servers consume much fewer CPU resources but provide much higher IOPS and throughput than centrally deployed engines do.
  - An example is provided as follows: A system has 20 servers that need to access the storage resources provided by FusionStorage Block, and the bandwidth that each server provides for the storage plane is 2 x 10 Gbit/s. One VBS module (storage engine) is deployed on each server, so that the total throughput can reach 400 Gbit/s (20 x 2 x 10 Gbit/s). With the growth of the cluster scale, storage engines can be linearly added, thus eliminating the performance bottlenecks that may be caused by centralized engines in traditional dual-controller or multi-controller storage systems.

Two Main Application Scenarios of Huawei FusionStorage Block

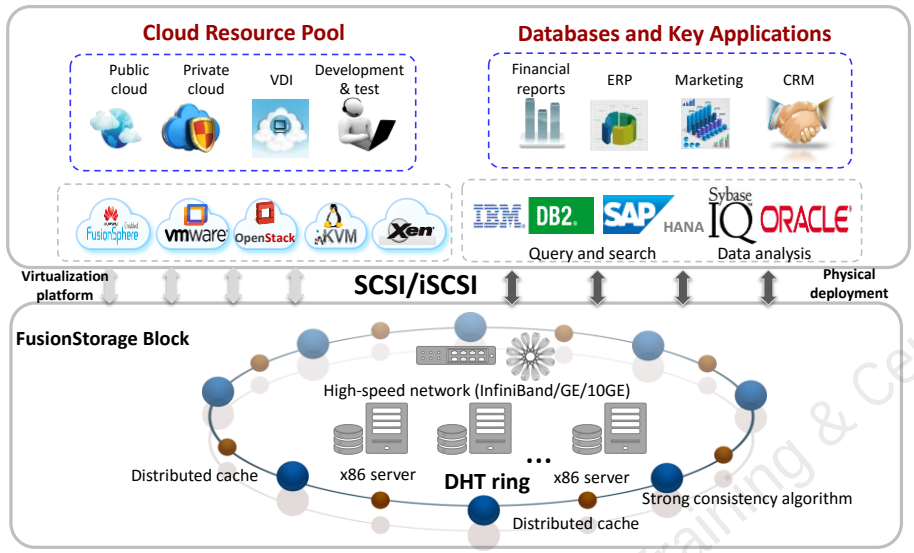

HUAWEI

- Open and compatibility: FusionStorage Block is compatible with mainstream databases, virtualization platforms, and servers.
- Converged deployed: FusionStorage Block supports converged deployment of virtualization platform and database resource pool, that is, one data center shares the resource pool provided by FusionStorage Block.
- FusionStorage Block uses SSDs other than HDDs as a high-speed storage device, and uses InfiniBand network other than GE/10GE networks to provide high bandwidth, meeting requirements of real-time processing of massive data and million-level IOPS.

Contents

1. FusionStorage Block Product Overview
2. **FusionStorage Block Functions and Features**
3. FusionStorage Block Logical Architecture
4. FusionStorage Block Software Installation
5. FusionStorage Block Configuration

FusionStorage Block Distributed Software Architecture

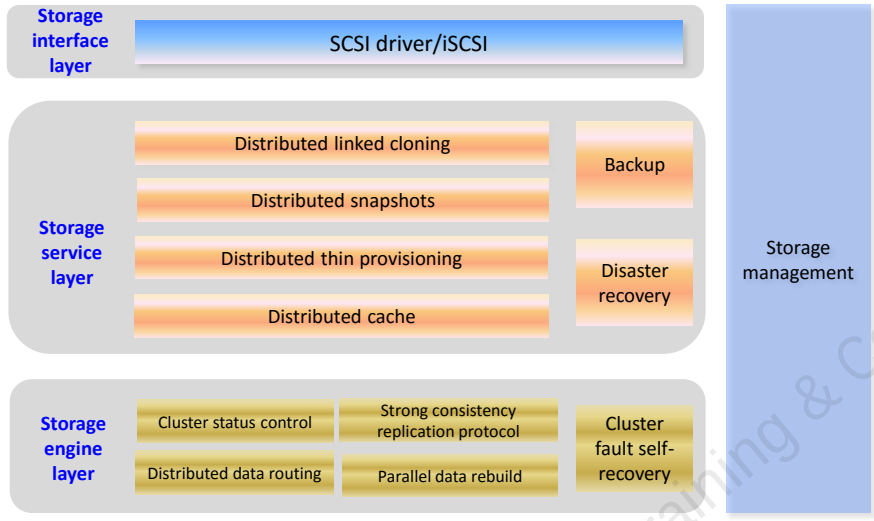

FusionStorage Block Storage Capabilities - SCSI & iSCSI Interfaces

- FusionStorage Block provides block storage interfaces using the VBS over the SCSI or iSCSI protocol.
  - SCSI: Used to deploy VBS physically, and install FusionSphere OpenStack or KVM.
  - iSCSI: Used to install the VMs or hosts, such as VMware, Microsoft SQL Server clusters except VBS, that provide accesses to storage resources.

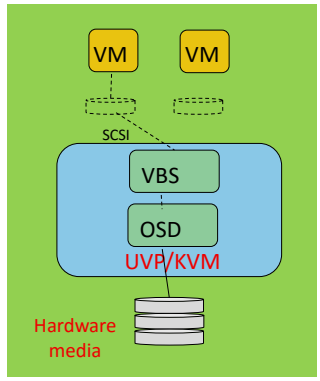

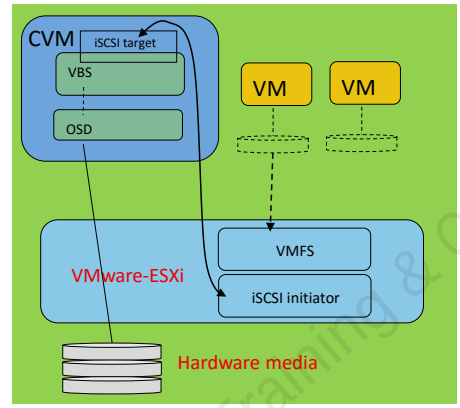

- Only the SCSI storage interfaces support snapshots, snapshot backups, and linked cloning functions. iSCSI storage interfaces do not support such functions.
- Users access the storage by connecting the local Initiator to the iSCSI target provided by the VBS.
- FusionStorage Block supports the following security access standards:
  - CHAP identity authentication
  - LUN Masking to authorize hosts to access LUNs.

FusionStorage Block Thin Provisioning

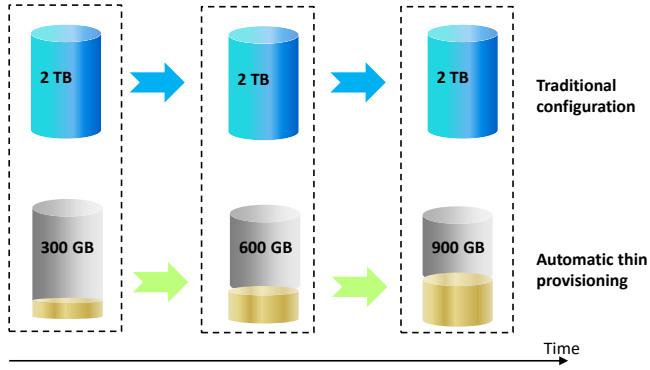

- Thin provisioning significantly improves storage utilization than the traditional storage allocating method can do.
- FusionStorage Block supports automatic thin provisioning, which delivers almost the same performance as traditional SAN storage.

FusionStorage Block Snapshot

- The snapshot mechanism of FusionStorage Block stores the status of data on logic volumes at certain points in time for possible exporting and restoration.
- FusionStorage Block uses the Redirect-On-Write (ROW) mechanism to create snapshots, which impose no adverse impact on volume performance.
- Infinite snapshots: Snapshot metadata is dispersedly stored and supports horizontal expansion, which eliminates bottlenecks caused by centralized storage. Therefore, the number of snapshots that can be created is not limited theoretically.
- Rapid volume restoration: Using snapshots, volumes can be restored within 1s without data migration. In contrast, restoration takes several hours on traditional SAN devices.

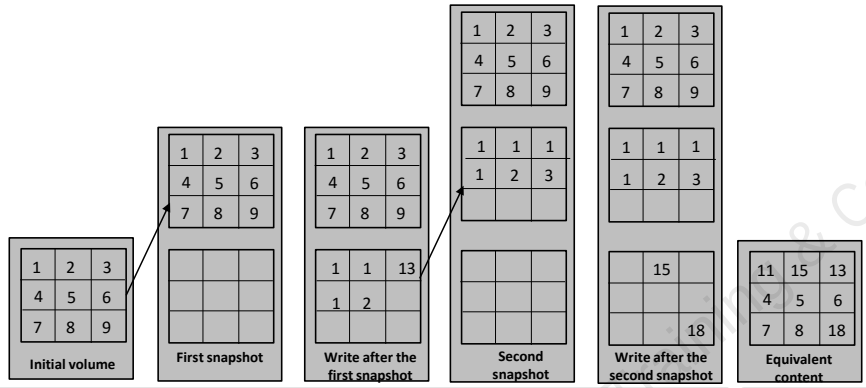

- In the preceding figures, empty cells indicate that no physical disk space is allocated.

FusionStorage Block Linked Cloning

- FusionStorage Block supports multiple cloned volumes to be created for one snapshot. Changes made to a cloned volume do not affect the snapshot and other cloned volumes.
- A cloned volume has all the functions of a common volume. You can create snapshots for a cloned volume, use a snapshot to restore the data in the cloned volume, and clone using a cloned volume as the base volume.
- FusionStorage Block supports batch volume creation for VMs and can create hundreds of volumes in a second.
- FusionStorage Block supports a linked clone ratio of 1:2048, which greatly improves storage space utilization.

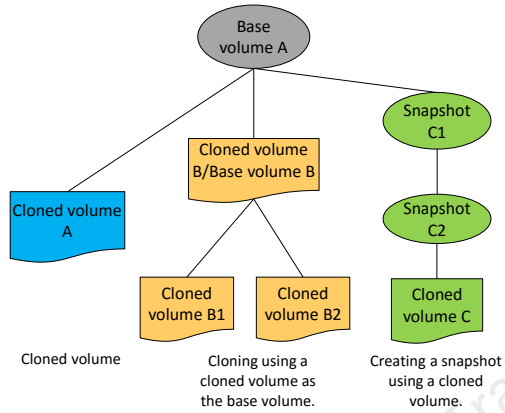

FusionStorage Block User-defined Storage SLA

- Users can customize resource pool SLA specifications based on service requirements.
- Data volumes of a VM can belong to different storage pools.

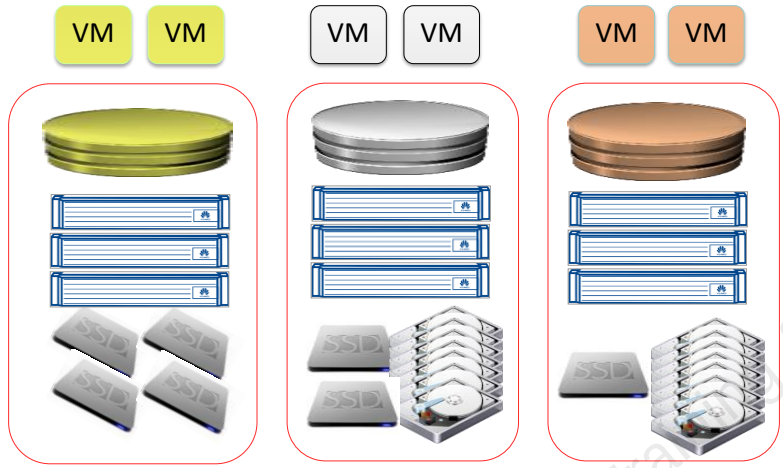

FusionStorage Block Cross-Resource Pool Volume Cold Migration

- Cross-Resource Pool Volume Cold Migration: The volume data can be migrated from the source resource pool to the target resource pool. The process is as follows: Create a target volume, replicate the volume data, delete the source volume, and change the name of the target volume.
- During this process, you cannot write data into the source volume. For this reason, the migration is called a cold migration.
- You can also use tools to invoke the replication interface to perform the cold migration.
- Scenario 1: Load balancing between resource pools, such as migrating volumes from a full pool to an idle pool.
- Scenario 2: Volume migration between resource pools of different performance, such as migrating volumes from a low-performance pool to a high-performance pool.

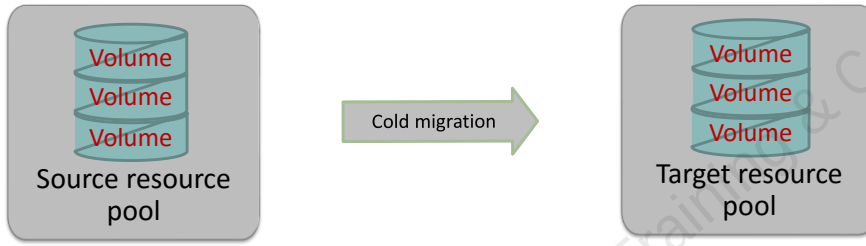

Contents

1. FusionStorage Block Product Overview
2. FusionStorage Block Functions and Features
3. **FusionStorage Block Logical Architecture**
4. FusionStorage Block Software Installation
5. FusionStorage Block Configuration

FusionStorage Block Logical Architecture (1)

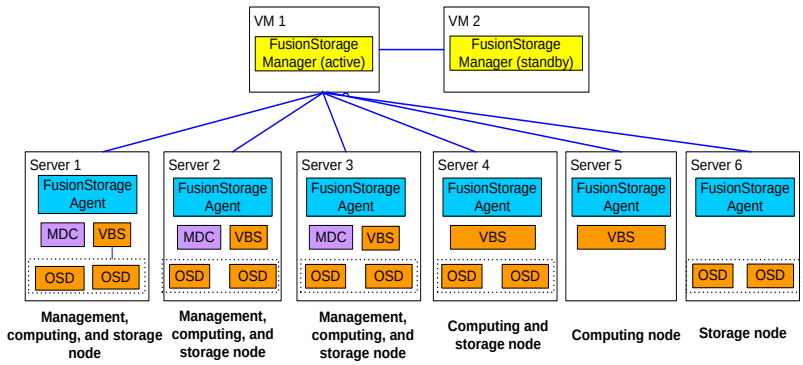

- FusionStorage Manager (FSM): The FusionStorage Block management module, which provides O&M functions including alarm management, service monitoring, operation logging, and data configuration. In most cases, FSM is deployed working in active/standby mode.
- FusionStorage Agent (FSA): The FusionStorage Block management agent module, which is deployed on each node to enable the node to communicate with FSM. FSA contains MDC, VBS, and OSD processes, which can be freely combined to implement specified functions based on system requirements.

FusionStorage Block Logical Architecture (2)

- Metadata Controller (MDC): The metadata control component, which controls distributed cluster status, data distribution rules, and data rebuild rules. By default, MDC processes are deployed on ZooKeeper disks of three servers to form the MDC cluster.
- Virtual Block System (VBS): The virtual block storage management component, which manages volume metadata, provides access points for distributed clusters, and enables computing resources to access distributed storage resources. By default, one VBS process is deployed one node, and all the VBS processes form the VBS cluster. Multiple VBS processes can also be deployed on one node to improve the node I/O performance.
- Object Storage Device (OSD): The object-based storage device, which performs I/O operations. Multiple OSD processes can be deployed on one server. Normally, one OSD process is deployed for one disk. When SSD cards are used as the main storage, multiple OSD processes can be deployed on one SSD card to maximize the SSD card usage and performance. For example, one 2.4 TB SSD card supports a maximum of six OSD processes, each of which manages the I/O operations for 400 GB space.

- Main storage: Disks on the servers that are used to store data in FusionStorage Block.

FusionStorage Block Deployment Mode

- **Converged Deployment**
  - VBS and OSD are installed on the same server.
  - Converged deployment is recommended for virtual applications.
- **Separated Deployment**
  - VBS and OSD are installed on the different servers.
  - Separated deployment is recommended for high-performance databases.

Contents

1. FusionStorage Block Product Overview
2. FusionStorage Block Functions and Features
3. FusionStorage Block Logical Architecture
4. **FusionStorage Block Software Installation**
5. FusionStorage Block Configuration

FusionStorage Block Installation Process

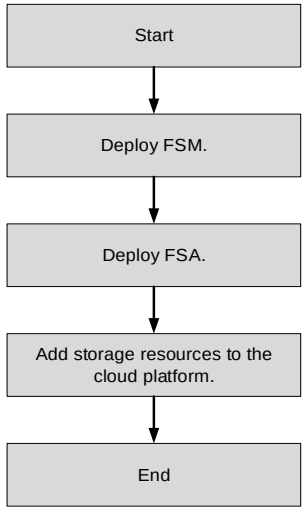

- Deploy FSM.
  - On the FusionSphere OpenStack web client, create two FSM VMs. Then log in to the FusionStorage Block web client and complete subsequent configurations.
- Deploy FSA.
  - Add servers to FusionStorage Block and selectively configure MDC, OSD, and VBS processes based on the roles of the servers in the system, implementing the storage and management functions of FusionStorage Block.
- Connect FusionStorage Block to the cloud platform.
  - Configure Cinder to enable Cinder to invoke the FusionStorage Block storage resources.

Obtaining the Software Package

- Required software package

Item	Description	How to Obtain
FusionStorage Block V100R006C00SPC100_UVP.tar.gz	Unified Virtualization Platform (UVP)-based FusionStorage Block software package	For enterprise users: Log in at http://support.huawei.com/enterprise, choose Support > IT > Cloud Computing > FusionCloud > FusionStorage Block, click the Downloads tab and select V100R006C00SPC100. For carrier users: Log in at http://support.huawei.com, choose Support > Product Support > Carrier IT > Cloud Computing > FusionCloud > FusionSphere > FusionStorage Block, click the Software tab and select V100R006C00SPC100.

Preparing for the Deployment

- Before installing FusionStorage Block, prepare the virtualization environment and ensure the following requirements are met:

Item	Description	Reference
Operating system (OS)	The hosts used to deploy FSA must be configured based on the FusionStorage Block requirements to ensure that the system provides sufficient computing resources for the FSA processes. The hosts synchronize time from the NTP server for consistency.	For the host Domain 0 specification and memory requirements, see the FusionStorage Block product documentation.
Configuring the storage network plane	The FusionStorage Block storage network plane is configured in Cinder on the FusionSphere OpenStack web client.	See the section describing the FusionSphere OpenStack configuration in the FusionSphere product documentation.

Configuring the Cinder Storage Network

- Log in to the FusionSphere OpenStack web client, choose Configuration > OpenStack and click Cinder.

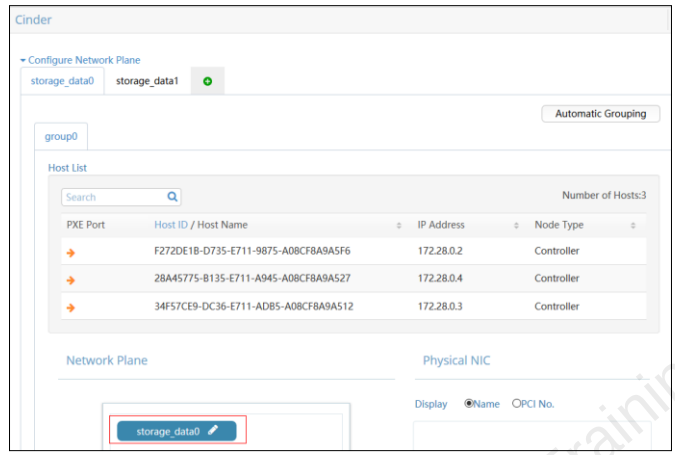

Installing FSM (1)

- Log in to the FusionSphere OpenStack web client over the External API network. The login address format is https://FusionSphere OpenStack Reverse proxy IP address:8890.
- Choose Cloud Service > FusionStorage Block. In the right pane, upload the software package.

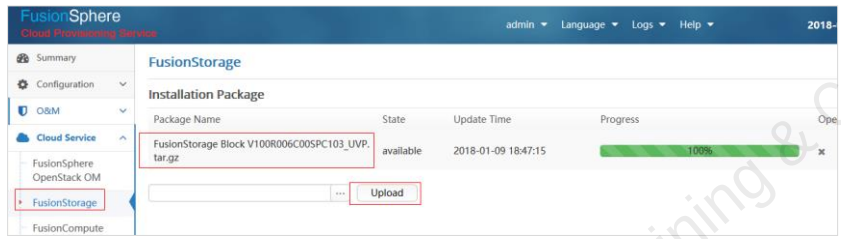

Installing FSM (2)

- Create VMs.

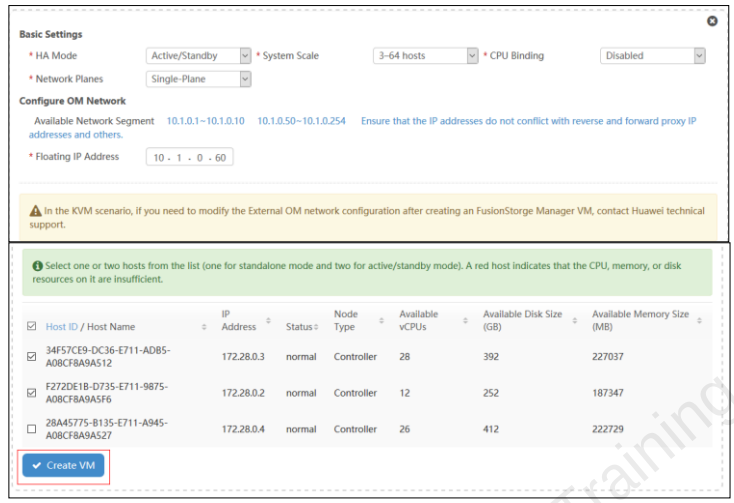

Installing FSM (3)

- After the installation is complete, click the link to log in to the FusionStorage Block web client.
  - The default username is admin, and its default password is **Huawei@CLOUD8!**.

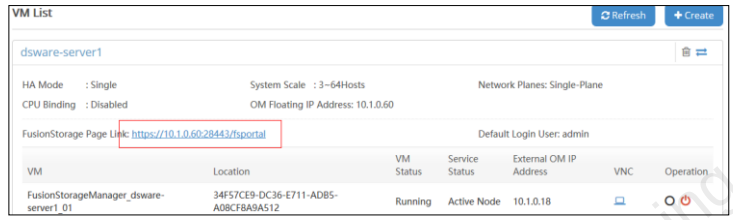

Installing FSA (1)

- Log in to the FusionStorage Block web client.

Installing FSA (2)

- If you log in to the FusionStorage Block web client for the first time, a message will be displayed indicating that servers have not been added to the system.

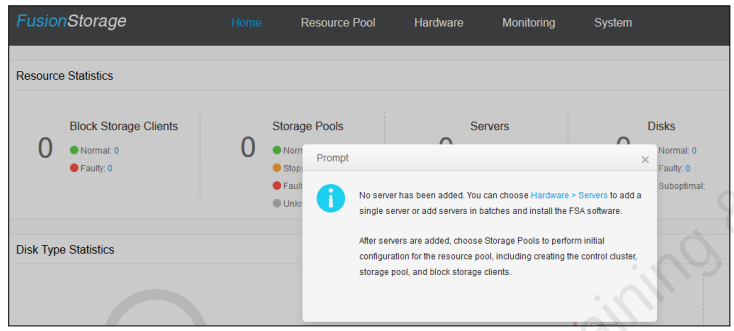

Installing FSA (3)

- Choose Hardware > Servers. In the right pane, click Single Add. In the displayed dialog box. enter the server management IP address. This server will be used as the converged deployment node.

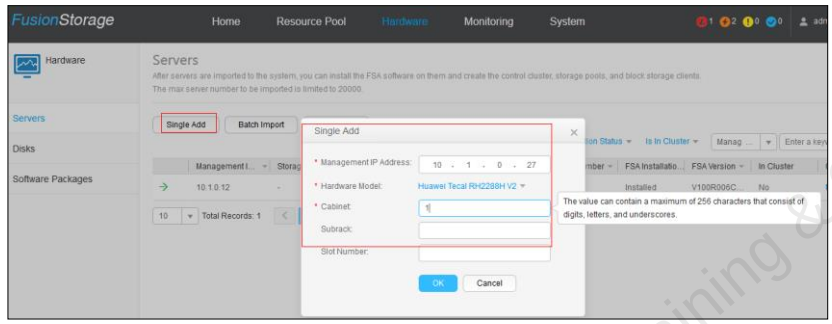

Installing FSA (4)

- The server management IP address is not the server BMC IP address. In FusionSphere OpenStack scenarios, you can query the server management IP addresses on the homepage of the FusionSphere OpenStack web client. Add servers to FusionStorage Block.

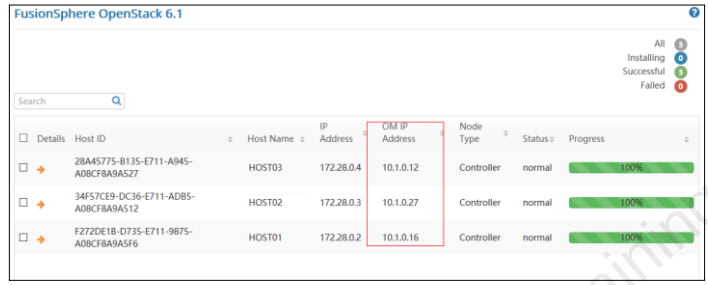

Installing FSA (5)

- In the right pane, click Install FSA. On the displayed Select Software Package page, click Select. Select DSwareAgent.tar.gz. Click Add.

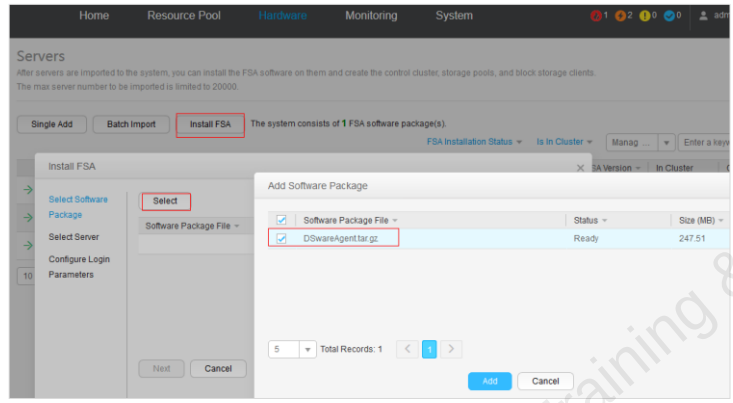

Installing FSA (6)

- Click Next to select the servers where you want to deploy FSA.

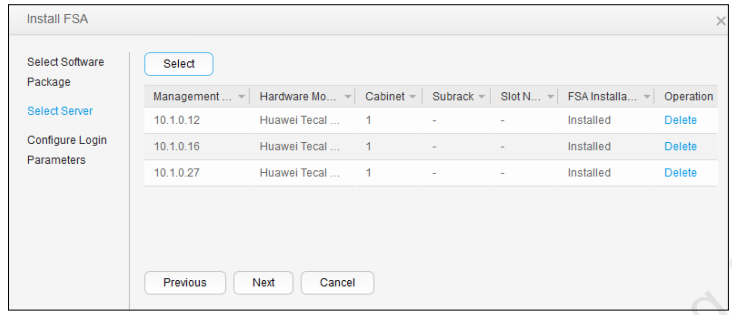

Installing FSA (7)

- Select Password Authentication, then set the authentication parameters. Set fsp for Server Login Username, set Huawei@CLOUD8 for Password, and set Huawei@CLOUD8! for Root password. Click Install to start the FSA installation.

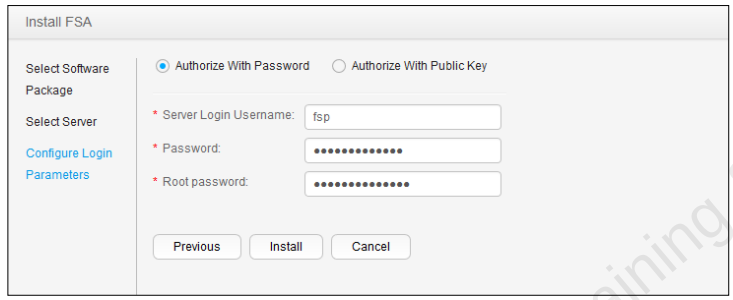

Installing FSA (8)

- The FSA installation starts and takes about 5 minutes to complete.

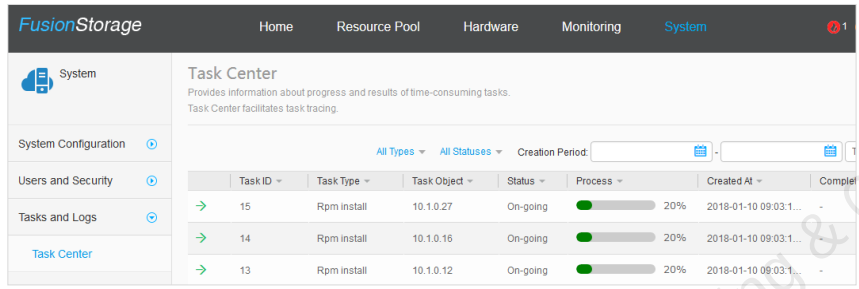

Assigning the blockstorage Role (1)

- Log in to the FusionSphere OpenStack web client and choose Configuration > Role. In the Resource Configuration Workspace for Specific Roles area on the displayed page, click the button next to the blockstorage role. In the displayed dialog box, configure resource isolation for this role.
- Configure the memory and vCPU resources for FusionStorage Block, respectively. For details, see the memory and vCPU requirements in section Host Domain 0 Requirements.

Assigning the blockstorage Role (2)

- Domain 0 specification requirements:

Item	Specification
CPU	The number of vCPUs required for running FusionStorage Block is as follows: • Converged deployment: 8 Separated deployment: • Computing node: 4 • Storage node: 6
Memory	Memory used by FusionStorage Block = Memory used by MDC processes + Memory used by VBS processes + Memory used by OSD processes + Memory used by caches • Memory used by MDC processes: 5 GB. Each storage node must reserve 5 GB of memory for the MDC process. • Memory used by VBS processes: In this example, common NICs are used. Therefore, VBS processes take up 8 GB of memory. • Memory used by VBS processes: In this example, disks are used as the main storage, the memory used by OSD processes is calculated as follows: • Memory used by OSD processes = N x Number of disks on the server (when the disk capacity is less than 2 TB, N is 2 GB) • Memory used by caches: In this example, SSDs or SSD cards are used as the cache, the memory used by caches is calculated as follows: Memory used by caches = (Total size of SSD cache devices on the server x 8)/1024 An example is provided as follows: In a 10GE networking scenario, two storage pools will be created, the server OS takes up 8 GB of memory, four 2 TB SATA disks are used as data disks, and one 2.4 TB SSD card is used as the cache. Memory used by FusionStorage Block = 8 GB (OS) + 5 GB (MDC) + 8 GB (VBS) + 2 GB x 4 (OSD) + (2.4 x 1024 GB x 8)/1024 (caches) = 48.2 GB

Huawei Training & Certification

Assigning the blockstorage Role (3)

- Select the target FSA node and drag the blockstorage role from System Role Management Workspace to Role Configuration Workspace.

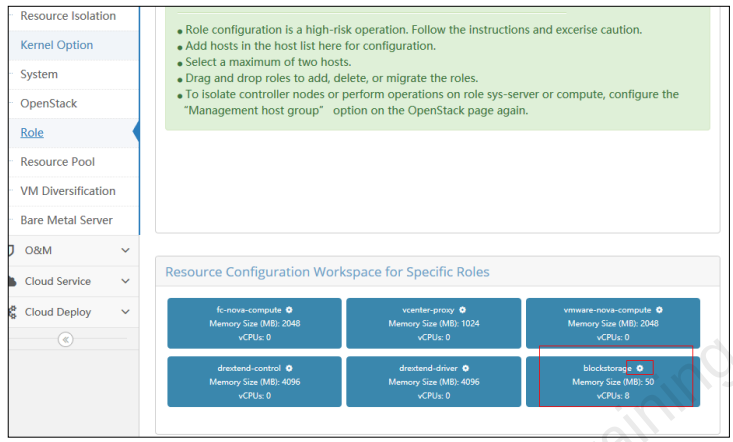

Contents

1. FusionStorage Block Product Overview
2. FusionStorage Block Functions and Features
3. FusionStorage Block Logical Architecture
4. FusionStorage Block Software Installation
5. **FusionStorage Block Configuration**

Configuring Alarm Reporting (1)

- Log in the FSM VM using Virtual Network Computing (VNC).
  - The default password of user dsware is **Huawei@CLOUD8**.

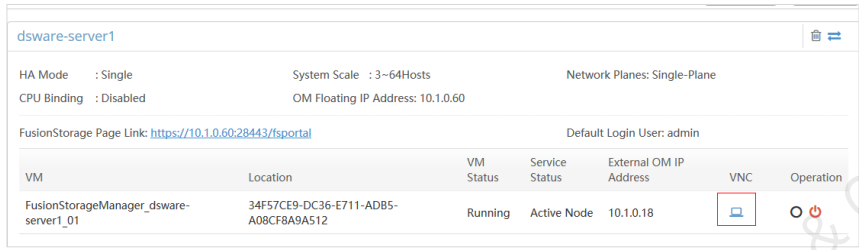

Configuring Alarm Reporting (2)

- Run the following command to configure the alarm reporting mechanism:
  - sh /opt/dsware/client/bin/dswareTool.sh --op configAlarmInfoForOMM -address 10.1.0.200 -addressPort 8043 -userNameOmmRest -authenticationType general
- After the command is executed, enter y and the password of the authentication user as prompted. The default password is **Omm@1234**.
- Since the dswareTool command is used, you also need to enter authentication username and password of command dswareTool as prompted.
- The default authentication username is cmdadmin, and its default password is **cmdHuawei@123**.

Configuring Alarm Reporting (3)

- Configure FusionStorage Block alarm reporting to FusionSphere OpenStack OM.
  - After alarm reporting is configured, connect FusionStorage Block to FusionSphere OpenStack OM on the FusionSphere OpenStack OM web client. FusionStorage Block floating IP address, username gmadmin and default password gmHuawei@123 are used to add FusionStorage Block.

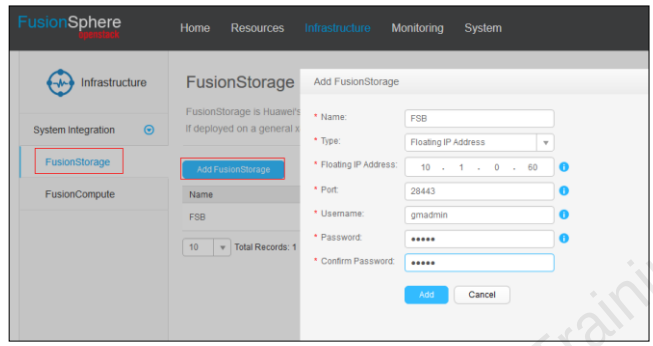

FusionStorage Block Basic Configurations

- On the FusionStorage Block web client, choose Resource Pool. Click Configuration Wizard.

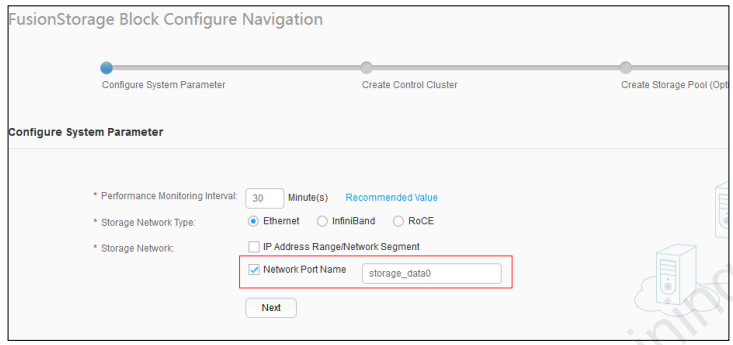

Creating the Control Cluster

- Configure the cluster name and metadata disks.

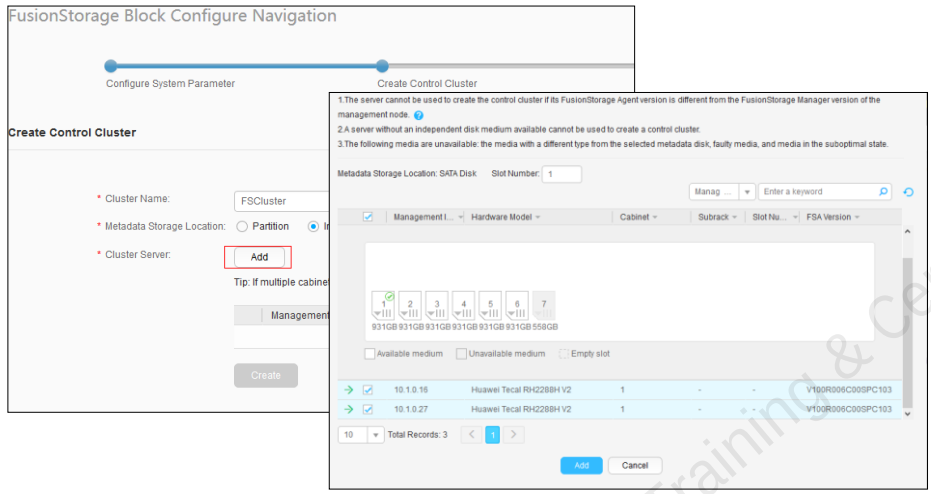

- If independent disks are used to deploy metadata, specify the metadata disk slot numbers.
- You are advised to set the disk in the last slot except the slots housing RAID disks to the metadata disk. Metadata disk slot numbers on servers can be different. However, you are advised to deploy metadata on the disk with the same slot number for each server to facilitate management and monitoring.
- Select three servers and click Add.

Creating a Storage Pool

- Enter the storage pool name, select 3-copy mode, select SATA disks to be used as the main storage, and select the server-level security.

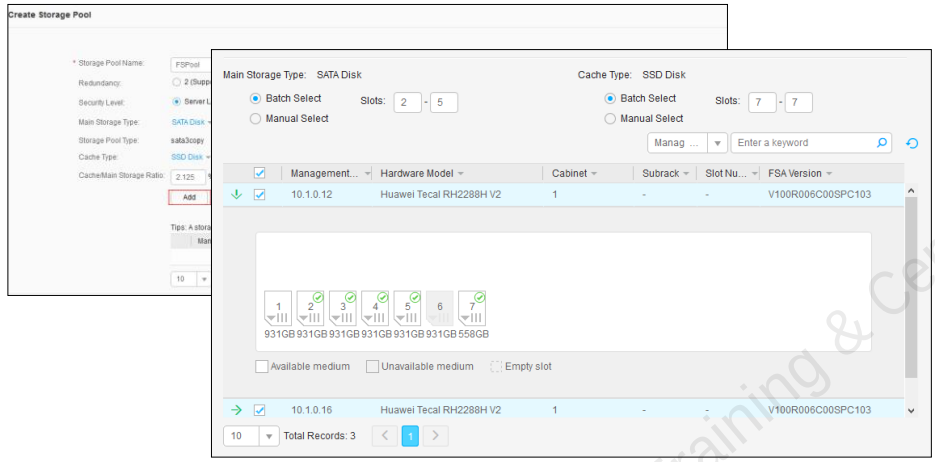

- Creating a Storage Pool
  - Support configurations of number of data replicas and types of main storage and cache. The main storage types include SAS, SATA, SSD card, and SSD. The cache types include NVDIMM, SSD card, and SSD. SSD cards or SSDs cannot be configured for both the main storage and cache.
  - In a resource pool, the main storage quantity on each server can be different. However, the difference can be 2 at most, and the disk slot numbers can be inconsecutive.
  - Batch selection means that your specified start slot and end slot on the target server are already inserted with disks, the disks are used for the storage pool. With manual selection, you can configure disks in different slots on different servers to be used for the storage pool. Deploy one OSD process for one disk.
  - Since such operations are managed as tasks in the system, users can check whether an operation is successful on the Task Management page.
  - One FusionStorage Block system supports up to 128 storage pools.

Creating Block Storage Clients

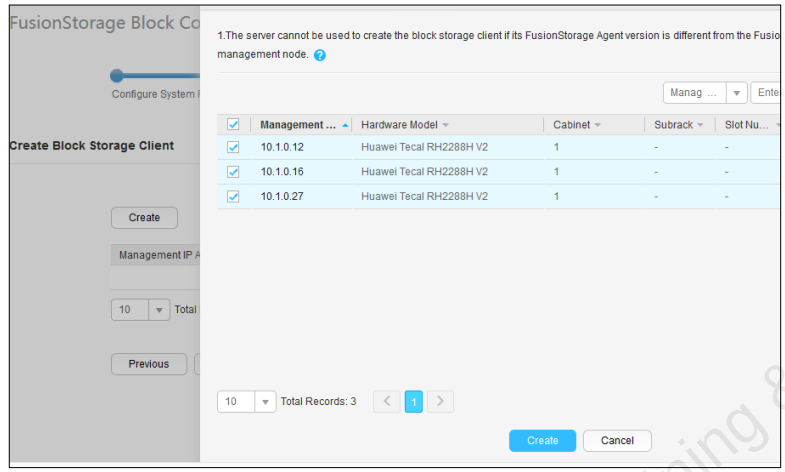

- Creating Block Storage Clients
  - The block storage clients need to be created for all the servers (VBS) that need to access the FusionStorage Block resources.
  - In one FusionStorage Block system, servers with VBS processes deployed can access all the storage pools in the FusionStorage Block system.
  - Block storage clients can be created in batches.

Connecting FusionStorage Block to FusionSphere OpenStack (1)

- On the FusionSphere OpenStack web client and choose Configuration > Resource Pool. In the right pane, click the icon next to the KVM resource pool. On the displayed Storage Cluster page, click the related icon to configure the parameters of the backend storage.

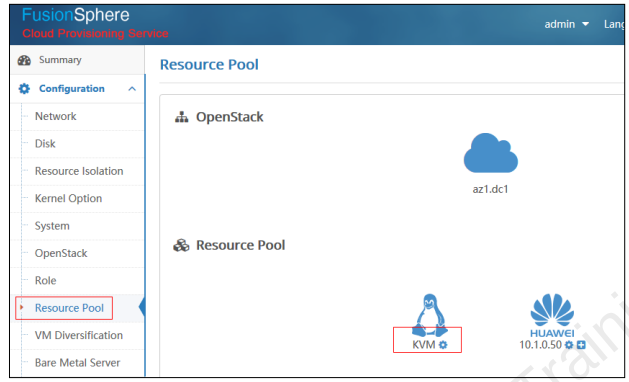

Connecting FusionStorage Block to FusionSphere OpenStack (2)

- Click the plus sign to add the storage cluster.

Connecting FusionStorage Block to FusionSphere OpenStack (3)

- Configure the connection parameters.
  - Select FusionStorage Block for the storage type, set FS as the backend storage name. Query the storage pool ID on the FusionStorage Block web client and set it for the resource pool ID. Set the FusionStorage Block floating IP address for the network floating IP address. Query the block storage client IP address on the FusionStorage Block web client and set it for the block storage client management IP address.

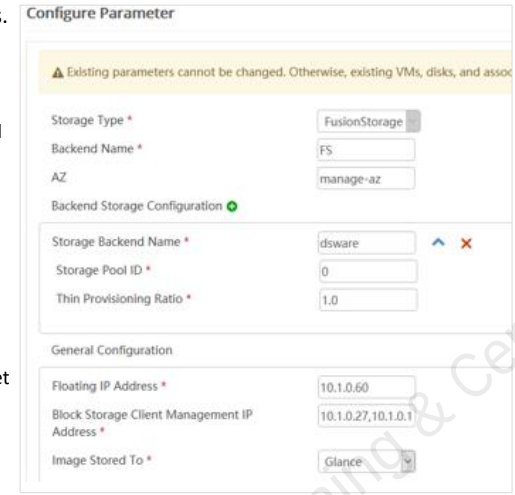

Verifying the Connection

- Log in to the FusionSphere OpenStack OM web client. The login username is cloud_admin, and its default password is FusionSphere123. On the backend storage page, view the FusionStorage Block status. The connection is successful if the FusionStorage Block status is displayed.

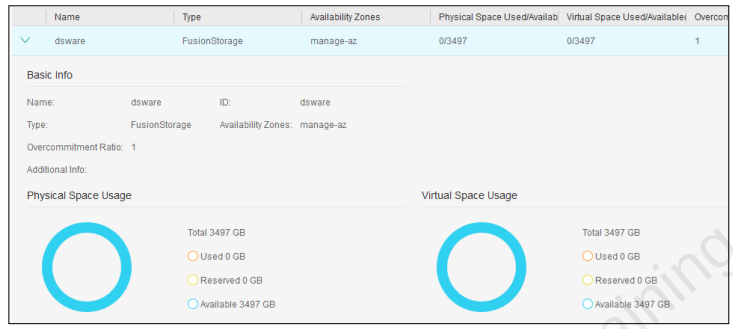

Quiz

1. How to Connect FusionStorage Block to FusionSphere OpenStack?

Summary

- Key points in this document:
  - FusionStorage Block architecture
  - FusionStorage Block functions: snapshots, linked cloning, and thin provisioning
  - FusionStorage Block software installation process: FSM and FSA
  - FusionStorage Block configuration

Recommendations

- Huawei Learning website:
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei support cases:
  - http://support.huawei.com/enterprise/servicecenter?lang=zh

Thank You

www.huawei.com

ManageOne Principles and Architecture

www.huawei.com

Foreword

- This course describes ManageOne in general, including architectures, functions, and features of ServiceCenter and OperationCenter, helping you get a better knowledge of ManageOne.

Objectives

- Upon completion of this course, you will be able to:
  - Describe the architecture of ManageOne.
  - Understand principles of ServiceCenter.
  - Understand principles of OperationCenter.

Contents

1. **ManageOne Architecture**
2. Operation Management
3. O&M Management

ManageOne Solution Overview

- Huawei ManageOne functions as a one-stop solution for the operation and O&M of cloud DCs. The ManageOne solution integrates physically dispersed multi-DC resources into a logically unified resource pool and provides infrastructure resources (such as computing, storage, and network resources) as cloud services to users. In this way, services are self-maintained by users and physical and virtual resources are centrally scheduled, automatically controlled and deployed, enabling the unified, process-based, and standard monitoring on cloud services.

- Huawei ManageOne brings the following benefits to customers:
- Higher resource utilization and lower total cost of operation (TCO)
  - The ManageOne solution integrates physically dispersed multi-DC resources into a logically unified resource pool. In the resource pool, ManageOne centrally allocates and schedules resources in different DCs by service resource usage. In addition, the ManageOne solution sorts resource fragments in physical resource pools and resource zones to prevent unnecessary resource expansion.
- Agile services and fast service rollout
  - The service catalog and self-service features allow service departments to quickly apply for resources by themselves. The what you see is what you get (WYSIWYG) service orchestration defines the templates for required computing, network, and storage resources to enable automatic control of resources and efficient allocation of required resources.
- Innovative service mode
  - Resources are metered and resource statistics are collected by department or service unit to provide the resource cost data for each department and service unit.

Future DC Management System Required by Enterprises

Global and Strategic Resource Scheduling

- Real-time service views of multiple DCs
- Elastic resource scheduling, improving resource usage and service reliability

Rapid and Flexible Service Provisioning

- Highly efficient and standard operation platform
- Automatic deployment and rapid service provisioning

E2E DC Monitoring

- Aware of service status and user experience
- Monitors infrastructure performance and facilitates fault locating

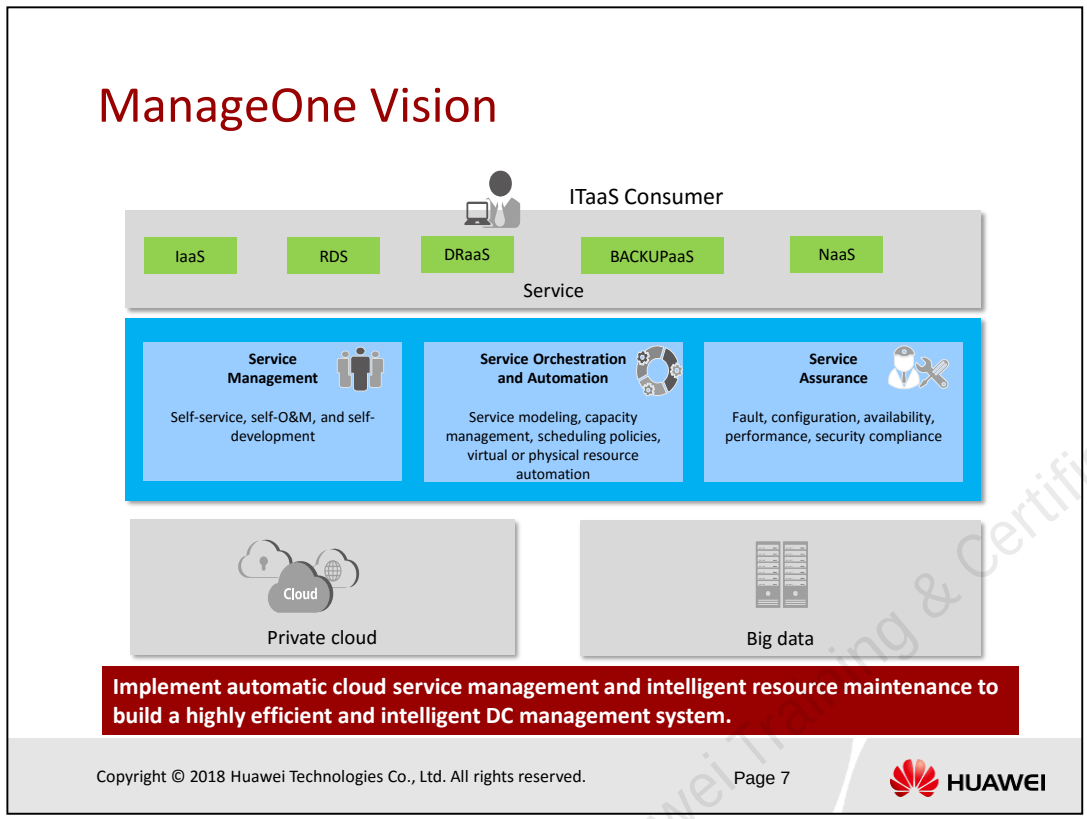

- Huawei ManageOne solution consists of ServiceCenter and OperationCenter and is used to manage cloud DCs.
- ServiceCenter
  - Provides customized DC services and unified service orchestration and automatic management capabilities based on cloud and non-cloud resources provided by resource pools.
- OperationCenter
  - Provides scenario-based O&M operations and visualized status, risk, and efficiency analysis capabilities for DC services, and works with ServiceCenter to implement DC self-optimization and self-healing based on the analysis results.
  - The ManageOne solution typically applies to three scenarios: converged resource pool, NFVI, and public cloud. In the NFVI and public cloud scenarios, only ManageOne OperationCenter is used to implement system O&M.

ManageOne Operation Architecture

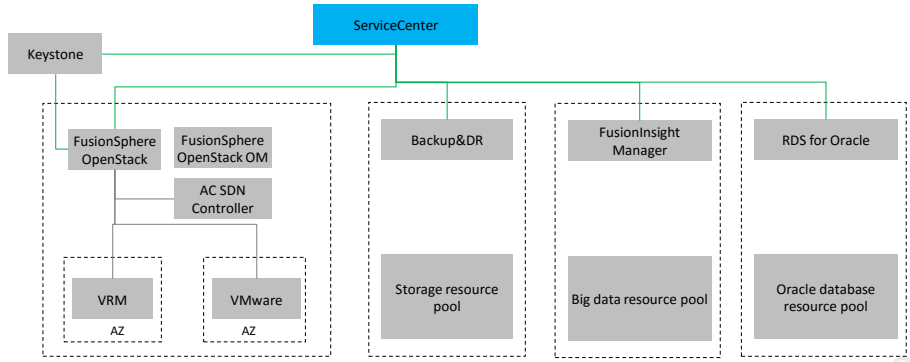

- ManageOne provides ServiceCenter and OperationCenter. The former is responsible for service provisioning and the latter provides the O&M and monitoring function.
- Provided by FusionSphere OpenStack, the IaaS resource pool supports KVM, VRM, and VMware. In multi-FusionSphere OpenStack scenarios, ServiceCenter can connect to multiple FusionSphere OpenStack systems through one Keystone.
- The backup and DR products in the DR solution provide VM backup and DR services. Only disk-based backup is supported. Backup can be automatically executed. VM DR is supported. Online application and offline provisioning of VMs are supported.
- The FusionInsight resource pool provides big data services, including HDFS, HBase, Spark, Hive, and MapReduce. These services are automatically provisioned.

ManageOne O&M Architecture

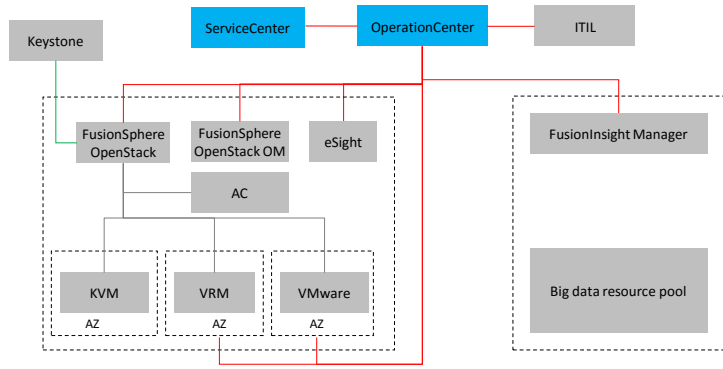

- ManageOne provides ServiceCenter and OperationCenter. The former is responsible for service provisioning and the latter provides the O&M and monitoring function.
- eSight monitors physical devices in DCs. FusionSphere OpenStack, FusionSphere OpenStack OM, VRM, and VMware monitor virtual devices in DCs. OperationCenter converges monitoring information about physical and virtual devices and provides a unified view.
- The FusionInsight resource pool provides big data services.

Contents

1. ManageOne Architecture
2. **Operation Management**
3. O&M Management

ServiceCenter Introduction - Architecture

- Architecture

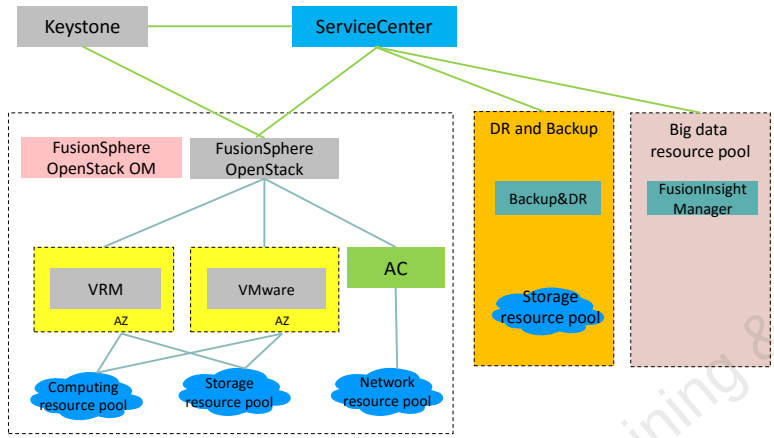

- ManageOne is a cloud DC management solution provided for customers. The cloud service platform of ManageOne is devoted to cloud service standardization and provides customized services for customers.
- The cloud service platform is loosely coupled with FusionSphere, which can connect to third-party management software. The cloud service platform supports flexible customization capabilities and provides DC operation and O&M services for customers. Among which there are:
  - ServiceCenter customizes DC services and implements unified service arrangement and automatic management based on cloud and non-cloud resources provided by resource pools. It provides customized heterogeneous support capabilities and multiple resource pool policies and arrangement, customized enterprise service integration, and customized resource pool management system capabilities supplemented by third-party components, especially automatic provisioning capabilities for traditional resources.
  - FusionSphere uses OpenStack as the benchmark and is the Huawei enterprise edition of OpenStack. FusionSphere provides pooled DC facilities for ManageOne.

User Management

- Function Description
  - Role

Role Name	Permission
System super administrator	System super administrators have the permissions of system configuration, resource configuration, service management, system super administrator management, organization administrator management, organization management, and alarm query. Services created by a system super administrator are visible to all organizations.
Organization administrator	Organization administrators have the permissions of VDC management, service management, VDC application order management, organization administrator management, and VDC administrator management. An organization administrator has the rights of all VDC administrators in the organization. If this organization has no VDCs, the organization administrator cannot log in to ManageOne ServiceCenter in tenant view. Services created by an organization administrator are visible to all VDCs in the organization.
VDC administrator	VDC administrators have the permissions to manage and use resources, applying for services, manage applications, tenant application orders, VDC administrators, VDC operators, VPCs, and tags, view monitoring data, and view resource metering data in their VDCs.
VDC operator	VDC operators have the permissions to apply for services, use resources, view services, user information, and monitoring data, manage tags, and view resource metering data.
System super administrator & Organization administrator	Users attached to this role are granted all permissions of the system administrator and organization administrator.

User Management

- Typical scenario: two-level tenant
  - For a multinational carrier or a cross-province company that consists of multiple provincial companies or subsidiaries, the position of each role is shown in the following figure.

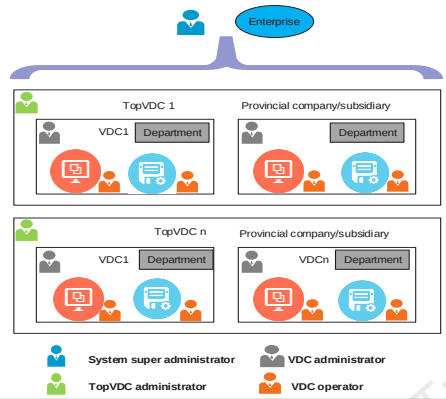

User Management

- Typical scenario: one-level tenant
  - For small- and medium-sized enterprises that do not consist of provincial companies or subsidiaries, the position of each role is shown in the following figure.

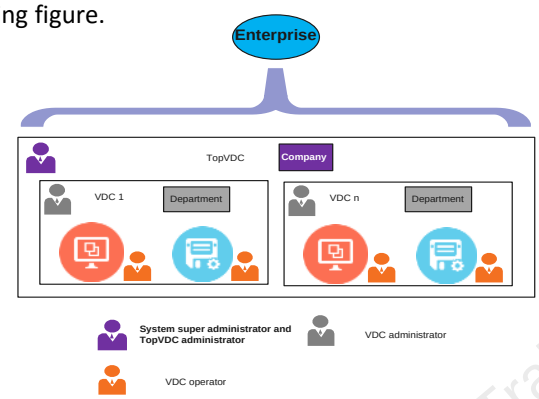

Service Management

- Function Description
  - Services are resource templates that users can apply for. Services are created by system administrators or organization administrators and published in service catalogs. Services created by system administrators are visible to all users or the specified organizations, whereas those created by organization administrators are visible to users in the organization to which the organization administrator belongs or the specified VDCs.
  - Preconfigured services include cloud host, cloud hard disk, elastic IP address, dedicated server, vAPP (POC), VLB (requiring F5 load balancers), backup, DR (online application and offline provisioning), and big data services.
  - A system administrator or organization administrator can define a new service by modifying parameters of preconfigured services.
  - The system provides the one-level approval process by default. A system administrator can define two- to five-level approval processes to replace the default approval process of a service with the defined one.

- You can install preconfigured services when installing ServiceCenter for convenience. In the service list, the services whose IDs start with default are preconfigured services. The language of name and description of preconfigured services is specified during installation and does not vary by the language of the environment. You can cancel the publishing of the service to change its name and description.
- If active/standby switchover occurs on ServiceCenter when an application order is being implemented, the application order will be implemented again after the switchover.

Service Management

- Feature Configuration

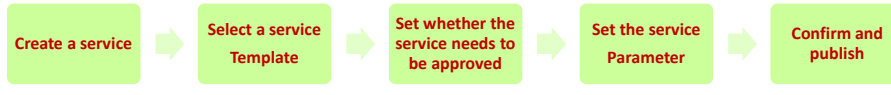

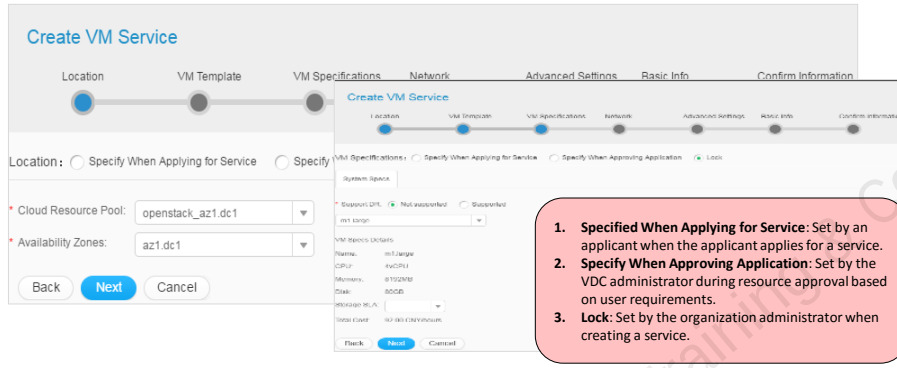

1. **Specified When Applying for Service**: Set by an applicant when the applicant applies for a service.
2. **Specify When Approving Application**: Set by the VDC administrator during resource approval based on user requirements.
3. **Lock**: Set by the organization administrator when creating a service.

Service Application

- Function Description

1. A VDC user submits an application for a service in the service catalog.
2. The system implements the service application and automatically generates a service application order.
3. A VDC administrator approves the service application after receiving the to-be-approved application order.
4. The system invokes the computing, storage, and network APIs of the resource pool to provision resources.
5. The system notifies the service application result to the user.

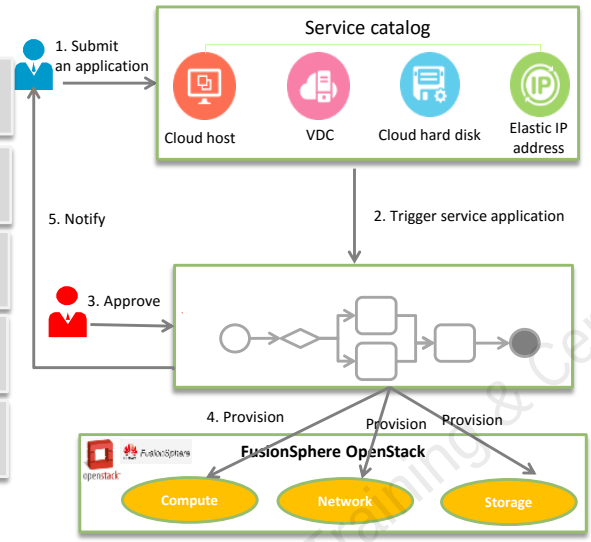

Organization Management

- Function Description
  - An organization is the unit used by ServiceCenter to allocate resources and is used in multi-level operation scenarios. The organization administrator manages resources in the entire organization while each VDC administrator manages only resources in the corresponding VDC, thereby achieving layered rights-based management.
  - If multi-level operation is not required, only one organization needs to be created. The organization quota is not limited. The system administrator is also the organization administrator and works as the global administrator to manage all resources.

Organization Management

- Function Description

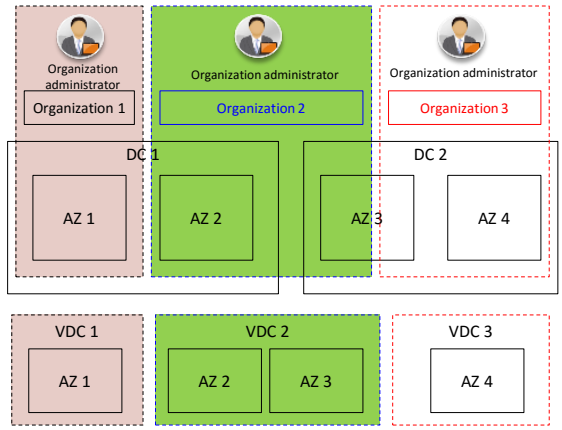

- Application scenario: If a department needs to be divided into multiple sub-departments for resource management, users can create an organization (which corresponds to a department) and multiple VDCs (which correspond to the sub-departments) in the organization.
- For example, Municipal Commission of Economy and Information Technology establishes resource pools for its ministries and allocates resources to the ministries, and each ministry needs to allocate resources to its departments. In this case, each ministry corresponds to an organization, and each department corresponds to a VDC.
- An organization includes quotas, users, resources, service catalogs, VDCs, and templates (vAPP templates). From the perspective of resource management, organizations are similar to VDCs. Resources can be provisioned in a VDC other than organization.

- Organization 1 uses resources of AZ 1 in DC 1, and resources, such as VMs and disks, in Organization 1 are provisioned in AZ 1.
- Organization 2 uses resources of AZ 2 in DC 1 and AZ 3 in DC 2, and resources, such as VMs and disks, in Organization 2 are provisioned in AZ 2 and AZ 3.
- Organization 3 uses resources of AZ 3 and AZ 4 in DC 2, and resources, such as VMs and disks, in Organization 3 are provisioned in AZ 3 and AZ 4.
- Note:
  - Organization 1 is exclusive to AZ 1, which needs to be configured by the system administrator because ServiceCenter does not provide exclusive AZ occupation.
  - The system administrator sets the quotas of organization 2 in AZ 2 and AZ 3.

Organization Services

- Configuration Process
  - If an organization is created on ServiceCenter, a domain is created on FusionSphere OpenStack.
  - The association relationship between organizations and AZs is only saved on ServiceCenter, and is not synchronized to FusionSphere OpenStack.

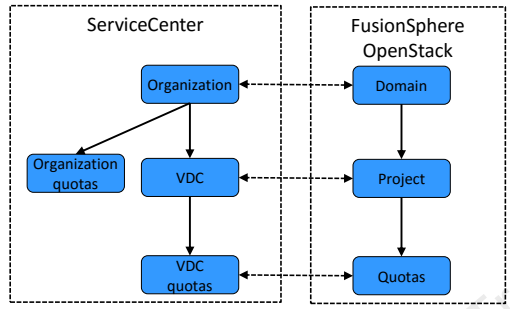

- The organization is analogous to a domain on FusionSphere OpenStack. Organization quotas are controlled by ServiceCenter, and the total quotas of VDCs in an organization cannot be greater than the organization quotas.
- The VDC is analogous to a project on FusionSphere OpenStack, and the quotas of a VDC and a project corresponding to the VDC are the same.

Organization Management

- Feature Configuration
  - Quota management: Users can select Limited and Not limited for resource quotas in an organization. If Limited is selected for quotas of a resource, users must select Limited for the quotas of the resource in all VDCs of the organization.

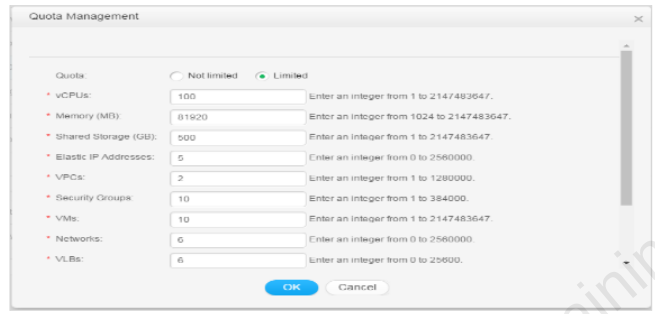

Organization Management

- Feature Configuration
  - AZ management: After an organization is created, users can add or delete cloud resource pools and AZs associated with the organization. If an AZ associated with the organization includes resources that have been applied for, the AZ cannot be deleted.

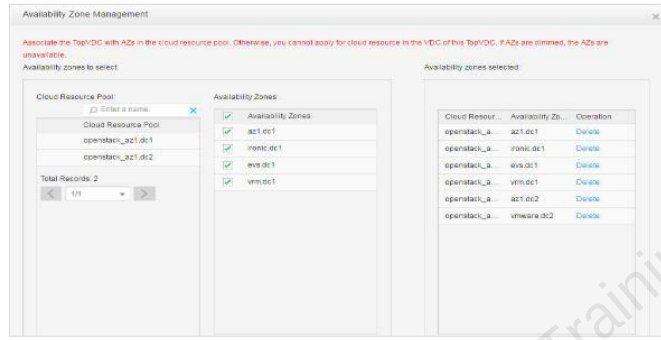

VDC Service

- Principles
  - A VDC is a collection of virtual resources, including computing, storage, and network resources, allocated to a department. A VDC is the representation of a physical DC at the virtualization layer.
  - The following figure shows the mappings between VDCs and physical DCs.

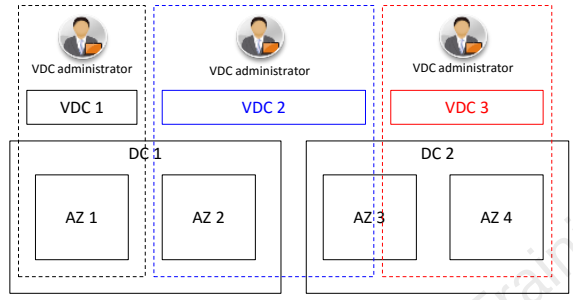

- VDC 1 uses resources of AZ 1 in DC 1, and resources, such as VMs and disks, in VDC 1 are provisioned in AZ 1.
- VDC 2 uses resources of AZ 2 in DC 1 and AZ 3 in DC 2, and resources, such as VMs and disks, in VDC 2 are provisioned in AZ 2 and AZ 3.
- VDC 3 uses resources of AZ 3 and AZ 4 in DC 2, and resources, such as VMs and disks, in VDC 3 are provisioned in AZ 3 and AZ 4.
- Note:
  - 1. VDC 1 is exclusive to AZ 1, which needs to be configured by the system administrator because ServiceCenter does not provide exclusive AZ occupation.
  - 2. The quantities of resources in AZ 3 used by VDC 2 and VDC 3 depends on the quotas of VDC 2 and VDC 3.

VDC Service

- The following figure shows the mappings between the VDC and FusionSphere OpenStack.
  - A VDC corresponds to a project on Keystone, a VDC administrator or user corresponds to a user on Keystone, and a user is associated with a project through a user group.
  - Quotas of resources in VDCs correspond to those managed by the Nova, Cinder, and Neutron components of FusionSphere OpenStack. All of these three components provide project-specific quota management interfaces.

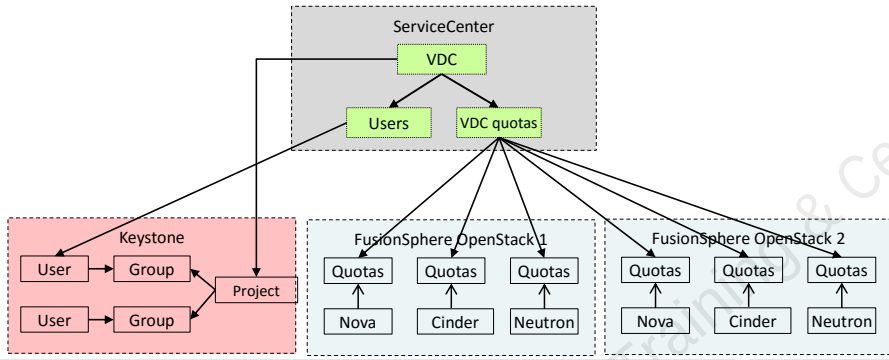

- There is a unique project in multiple FusionSphere OpenStack systems because the systems use the same Keystone.

VDC Service

- A VDC model includes quotas, users, resources, service catalogs, networks, and templates.
  - Quota: A quota indicates the maximum quantity of resources available to a VDC. It is specified by the VDC administrator or organization administrator during VDC creation.
  - Users: A VDC includes VDC administrators and VDC operators. VDC administrators manage users in the VDC and approve service applications, and VDC operators use resources.
  - Resources: Resources refer to services that users can apply for, including cloud hosts, cloud hard disks, elastic IP addresses, and other resources. Resources can be maintained and monitored by users.
  - Templates: The application template is a POC feature.

- Service catalog: A service catalog includes services that can be applied for by users in a VDC. System administrators can customize service catalogs for VDCs based on requirements of corresponding departments.
- VPC: A VPC indicates a logically isolated network environment providing virtual routers, networks, access control list (ACL), and virtual private networks (VPNs).
- Template: Application templates (POC feature) and VM templates are provided. VDC users can use global templates and also create private application templates and VM templates.

VDC Service

- Configuration Process

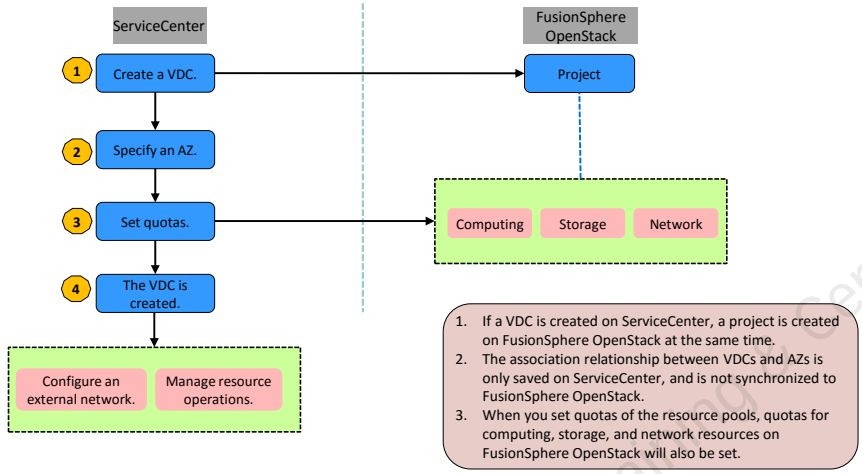

VDC Service

- Feature Configuration

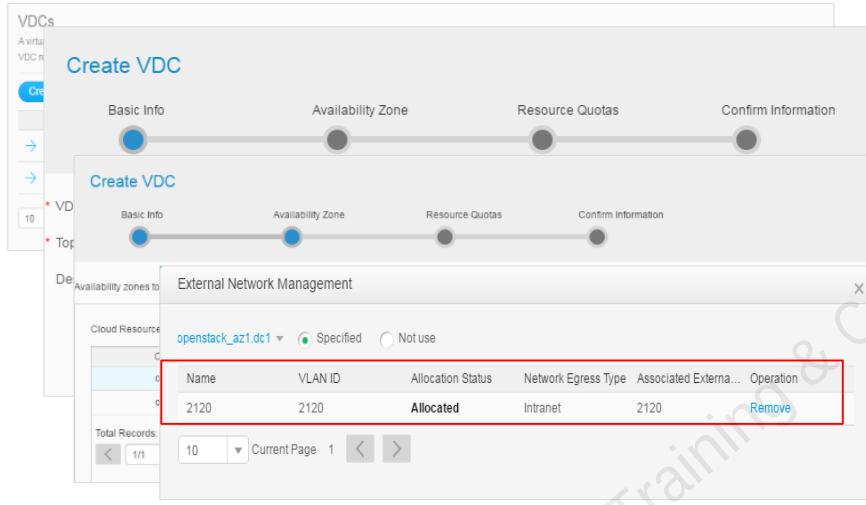

- A VDC can be associated with AZs that have been associated with the organization to which the VDC belongs.
- If resource quotas of an organization are limited, resource quotas of VDCs in the organization must be limited.
- After a VDC is created, users can add or delete cloud resource pools and AZs associated with the VDC. If an AZ associated with the organization includes resources applied for, the AZ cannot be deleted.
- External network management: After a VDC is created, all external networks that belong to the cloud resource pool and have the network egress type configured are displayed in the VDC by default. Users can specify external networks that can be used by the VDC.

VDC Service

- Feature Configuration
  - VDC metering: ServiceCenter meters the used and total quota of resources in a VDC and collects accumulated metering results for settlement with departments. By default, ServiceCenter meters resources in a VDC based on the used quota.

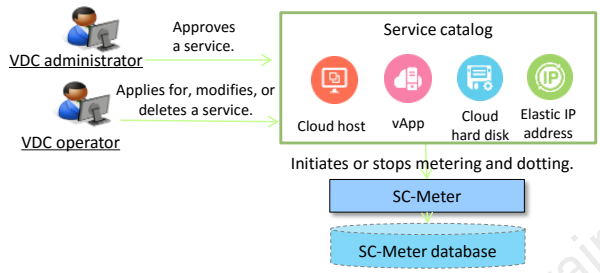

VDC Service

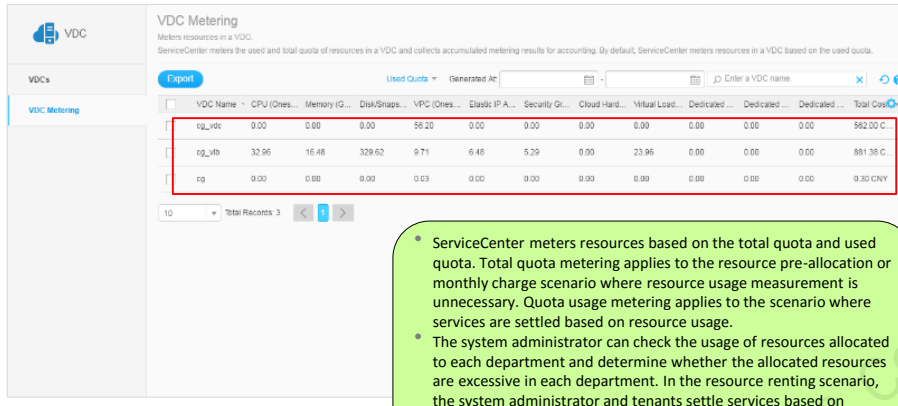

- ServiceCenter meters resources based on the total quota and used quota. Total quota metering applies to the resource pre-allocation or monthly charge scenario where resource usage measurement is unnecessary. Quota usage metering applies to the scenario where services are settled based on resource usage.
- The system administrator can check the usage of resources allocated to each department and determine whether the allocated resources are excessive in each department. In the resource renting scenario, the system administrator and tenants settle services based on metering lists.

- ServiceCenter provides a metering list that records application time and resource quantity of each service instance.

Computing Services - Cloud Host

- Function Description
  - The cloud host service provides usable-upon-application VM services. Users can apply for cloud hosts with specified specifications and OS types.
  - A cloud host runs on a host and obtains CPUs, memory, graphics cards, universal serial bus (USB) devices, network resources, and storage access capabilities of the host. Multiple cloud hosts can concurrently run on a host.

Computing Services - Cloud Host

- Function Description

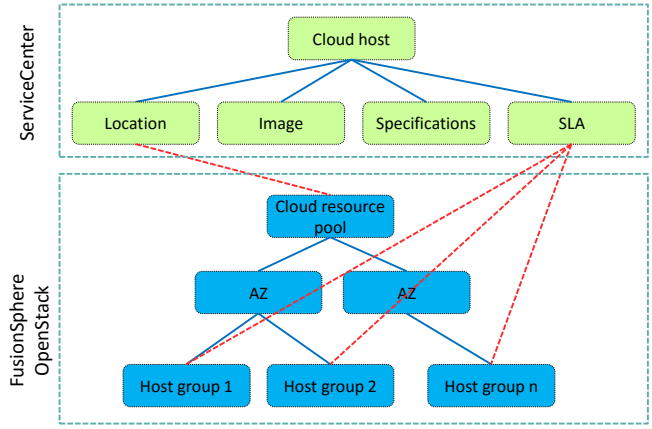

SLA: FusionSphere OpenStack OM provides the tag management for host groups. Based on the attributes and characteristics of the host in the host group, set a tag for the host group. When applying for a VM, users can select different host groups based on SLAs for resource allocation to meet user requirements.

- FusionSphere OpenStack OM provides the tag management for host groups. Based on the attributes and characteristics of the host in the host group, set a tag for the host group. After the configuration, users can use the tag to configure a tag for cloud host specifications. When you use the cloud host specifications to create a cloud host, the system creates the cloud host on the host in the host group meeting tag requirements. For example, if the tag configured for host group 1 is computing_performance with value High, and the same tag is configured for the specifications, the system selects a host from host group 1 to create the instance when users use the specifications to create a compute instance.

Computing Services - Cloud Host

- Feature Configuration

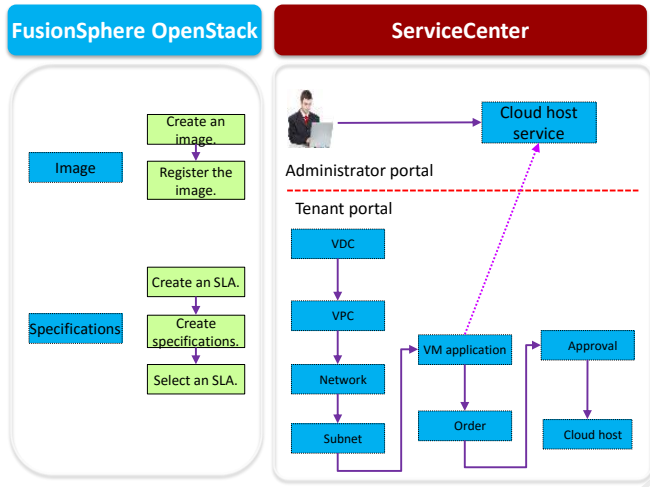

Before applying for a VM, the following conditions must be met:

1. FusionSphere OpenStack OM: The system administrator registers an image, create VM specifications, and define SLA tags on FusionSphere OpenStack OM.
2. ServiceCenter: The system administrator publishes the cloud host service, a VDC operator applies for a cloud host, a VDC administrator approves the application, and ServiceCenter automatically provisions the cloud host by invoking the FusionSphere OpenStack interface.

- Image: You can upload an image on FusionSphere OpenStack OM. You are advised to select Windows and Linux for the OS type, instead of Other, when uploading an image.

Computing Services - Cloud Host

- Feature Configuration

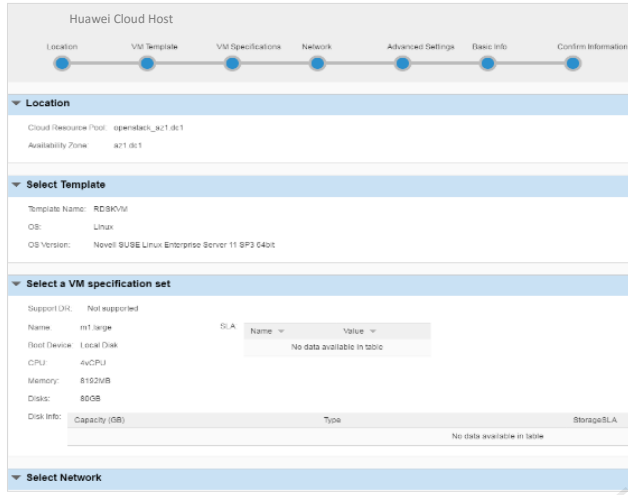

1. The setting mode for parameters of the cloud host service can be Specify When Applying for Service, Specify When Approving Application, or Lock.
2. Specify When Applying for Service: An applicant sets the parameters such as Location, Image, and Specifications.
3. Specify When Approving Application: A VDC administrator set the parameters during approval.
4. Lock: A system administrator sets the parameters when creating the service.

Storage Services - Cloud Hard Disk

- Function Description
  - Cloud hard disks provide data storage space for cloud hosts and dedicated servers and are independent from cloud hosts and dedicated servers. A cloud hard disk can be mounted to different cloud hosts or dedicated servers and supports data migration among cloud hosts and dedicated servers.
  - In addition, the cloud hard disk service provides the snapshot function. VDC operators can create snapshots for cloud hard disks. A faulty cloud hard disk can be restored using its snapshot.

Storage Services - Cloud Hard Disk

- Application Scenarios
- The cloud hard disk service provides on-demand, scalable block storage for users' cloud hosts. The cloud host OSs access the block storage space by volume.
- The cloud hard disk service applies to the following scenarios:
  - When applying for a cloud host, select an existing cloud hard disk.
  - When the disk space of a cloud host is insufficient, mount a cloud hard disk on the cloud host.
  - If a cloud host is faulty, unmount the cloud hard disk from the cloud host, and mount it on a new cloud host.

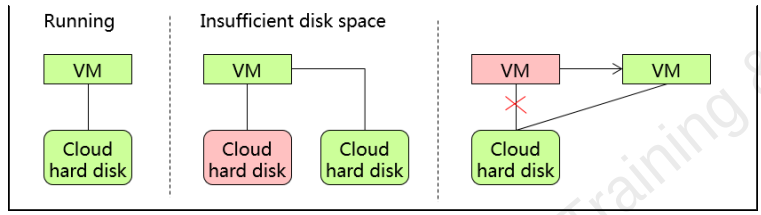

Storage Services - Cloud Hard Disk

- Feature Configuration

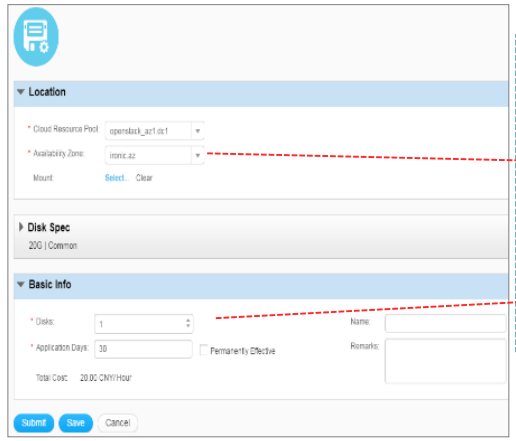

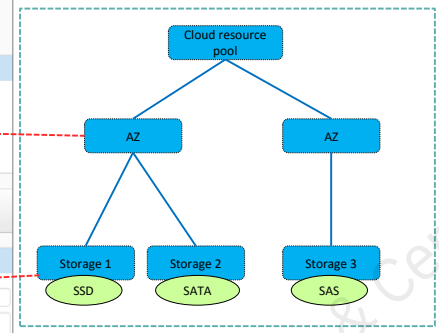

1. Creation Source: Image, Snapshot, or Cloud Hard Disk
2. Type: Common or Shared
3. Storage SLA: For example, SATA, SAS, or SSD

- Creation source
  - Image: Apply for a cloud hard disk using an image. After the cloud hard disk is successfully applied for, it will include the image data, which can be used to restore the system or compare system files.
  - Snapshot: Apply for a cloud hard disk using a snapshot. After the cloud hard disk is successfully applied for, it will include the snapshot data, which can be used to restore the system.
  - Cloud hard disk: Apply for a new cloud hard disk.
- Storage SLA: The disk type or tag can be selected during disk creation. A disk type defines a set of capabilities of the backend storage to be used for the disk.

Network Services - VPC

- Function Description
  - A VPC provides a secure, isolated network environment in a VDC. Users can customize virtual networks that provide the same functions as traditional networks in a VPC. The custom virtual networks can also provide advanced network features, such as elastic IP address and secure network address translation (SNAT) to keep up with service deployment requirements.
  - The following figure shows the mappings between objects in a VPC and FusionSphere OpenStack.

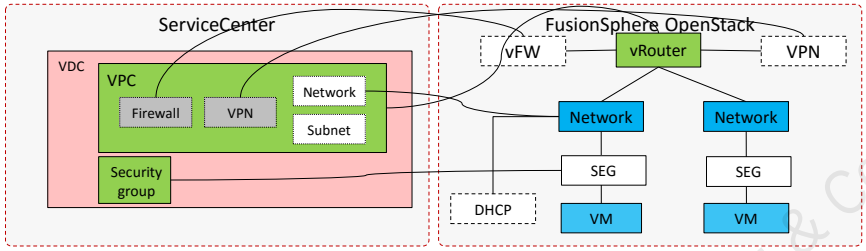

1. A VDC administrator can create VPCs and security groups in the VDC to which the VDC administrator belongs.
2. When users create a VPC, a vRouter object without an external network is created on FusionSphere OpenStack.
3. After applying for a router in a VPC, users can apply for a VPN, firewall, network, and subnet, and enable the Dynamic Host Configuration Protocol (DHCP) when creating a subnet.

- When users create a VPC, a router without an external network is created on FusionSphere OpenStack. The router is only an object. It cannot provide router functions.
- The subnet IP address can be allocated using either of the following methods:
  - Automatically allocated by the DHCP process
  - Manual configuration
- The following rules are set for the default security group:
  - One IPv4 inbound security group rule (specifies that cloud hosts in the security group are allowed only to access each other), one IPv4 outbound security group rule (specifies that cloud hosts in the security group are allowed to access cloud hosts in other subnets or security groups).
  - One IPv6 inbound security group rule (specifies that cloud hosts in the security group are allowed only to access each other), one IPv6 outbound security group rule (specifies that cloud hosts in the security group are allowed to access cloud hosts in other subnets or security groups).

Network Services - Networks

- Function Description
  - Internal network: It is only for internal use and cannot communicate with external networks. For example, it is recommended that databases be deployed on internal networks.
  - Routed network: provides VLANs and Layer 3 gateways. ServiceCenter automatically enables routes between all the routed networks in a VPC. That is, all the routed networks in a VPC can communicate with each other by default. A routed network is used to communicate with external networks. For example, it is recommended that the portal components of web applications be deployed on routed networks.
  - Direct network: A direct network connects VMs directly to external networks. The gateway and router of a direct network are not on the network planes managed by VDCs.

Networks - Internal Network

- Internal Network

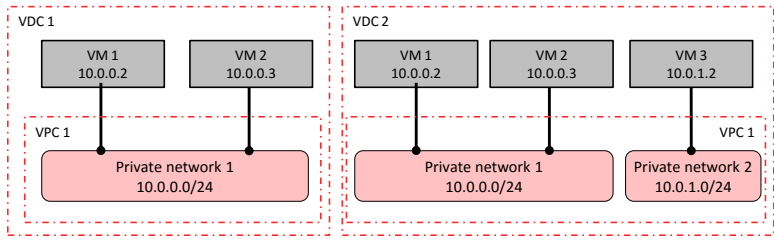

A private network (for example, 10.0.0.0/24) can be created in each VPC. It is recommended that private networks of different VPCs be not overlapped. If they are overlapped, VPC interworking will be affected. A private network is an internal network. For example, in the previous figure, Private network 1 and Private network 2 are internal networks in VDC 1 and VDC 2, respectively.
Users can select **DHCP** or **Manual** for the IP address allocation mode when creating a network.

- The process of creating a network by a tenant is as follows:
  - Create a private network and a subnet for VDC 1. Different private networks need to be isolated using VLANs.
  - Create an external network and a subnet. The external network logically functions as a physical network in the DC, and can be used to access the Internet.
  - Create a router in a VPC to be used for accessing the external network,
  - Add the created private network to the router.
  - The virtual network can communicate with the physical network.

Network Services - Routed Network

- Example

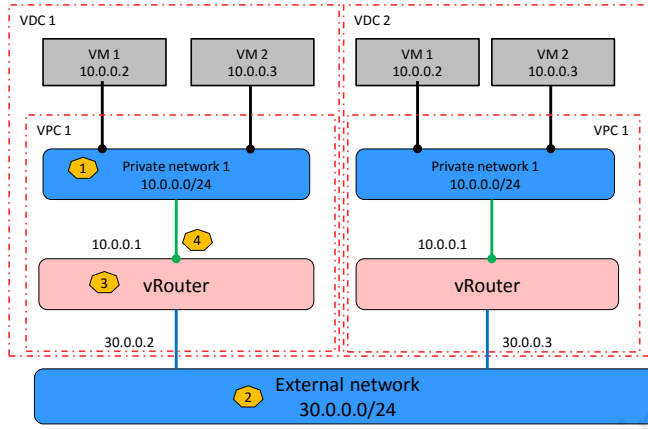

Users can create a router for a VPC and associate an internal network in the VPC with the router. Then, the internal network becomes a routed network. For example, in the previous figure, Private network 1 in VDC 1 and that in VDC 2 become routed networks after being associated with the vRouters.

- The process of creating a routed network is as follows:
  - Create a private network and a subnet in a VPC. Different private networks need to be isolated using VLANs.
  - Create an external network and a subnet. The external network logically functions as a physical network in the DC, and can be used to access the Internet.
  - Create a router in the VPC, and add the private network to the router for accessing the external network.

Network Services - Direct Network

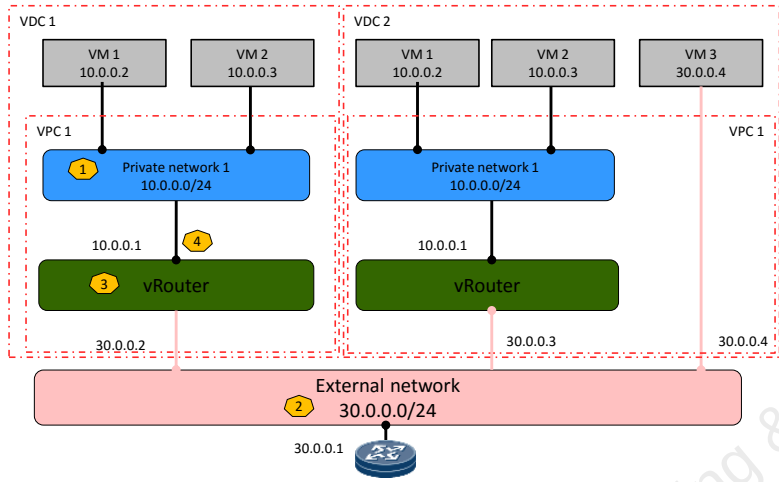

If users apply for a VM on a direct network, the VM can directly access the Internet.

Network Service - Router

- Function Description
  - Users can apply for a router to create routed networks. Routed networks in a VPC can communicate with each other. Cloud hosts in a routed network can access public networks using elastic IP addresses or the SNAT function.

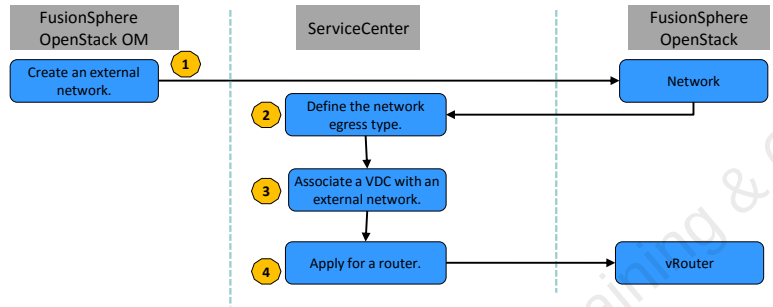

Network Services - Elastic IP Address

- Function Description
- An elastic IP address is a public IP address that is bound to a cloud host on a routed network, making services in a VPC accessible to external services. It allows service cloud hosts in a VPC to access external services using a fixed public IP address.

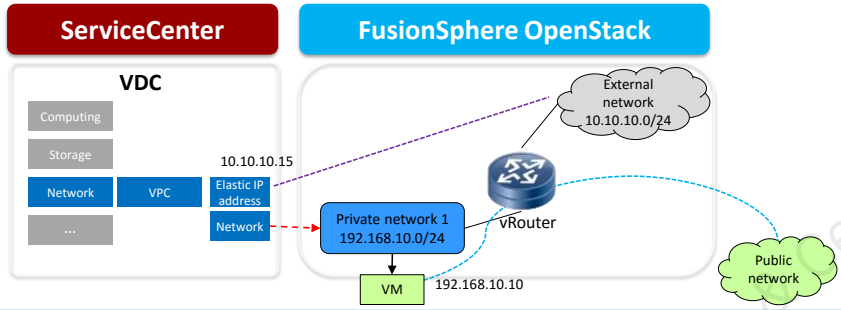

Apply for a vRouter in a VPC and associate the vRouter with an external network whose IP address is 10.10.10.0/24.
Create an internal network Private network 1 (whose IP address is 192.168.10.0/24) in the VPC, and associate the internal network with a vRouter.
Apply for an elastic IP address in the VPC. 10.10.10.15 is allocated.
Apply for a cloud host in the VPC. 192.168.10.10 is allocated to the cloud host. Then, bind the cloud host with an elastic IP address 10.10.10.15 so that the cloud host can access Internet using 10.10.10.15.

Network Services - Elastic IP Address

- Application Scenarios
- When a cloud host or virtual load balancer (VLB) in a VPC needs to access external networks, the cloud host or the VLB can be bound to an elastic IP address.
- Prerequisites
  - A cloud resource pool has been connected to ServiceCenter and the elastic IP address function has been set to Supported in the advanced settings of the cloud resource pool.
  - VPCs have been created in a VDC.
  - A router has been applied for in a VPC, and the Intranet network associated with the router has been associated with an Internet network.

Network Services - SNAT

- Function Description
- If an internal IP address initiates a connection to the services on the public network, the gateway on the router or firewall translates the private IP address into a public IP address. This process is known as SNAT, which applies to access to public networks using shared internal IP addresses.

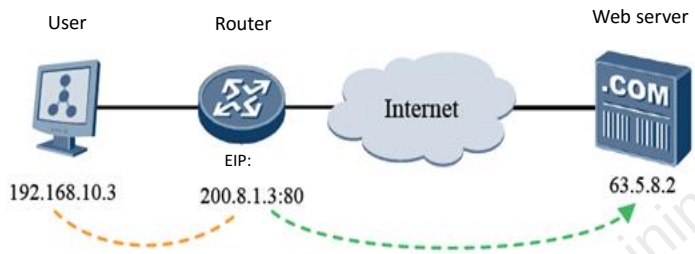

Network Services - SNAT

- Application Scenarios
  - When one or more routed networks in a VPC need to access external networks, users can configure the SNAT function.
- Prerequisites
  - A cloud resource pool has been connected to ServiceCenter and the SNAT function has been set to Supported in the advanced settings of the cloud resource pool.
  - VPCs have been created in a VDC.
  - A router has been applied for in a VPC, and the Intranet network associated with the router has been associated with an Internet network.

Network Services - Firewall

- Function Description
- A firewall is logically divided into multiple virtual firewalls. Each virtual firewall can provide independent security insurance for enterprises, maximizing resource utilization of firewalls. Virtual firewalls can be provided by physical or software firewalls.

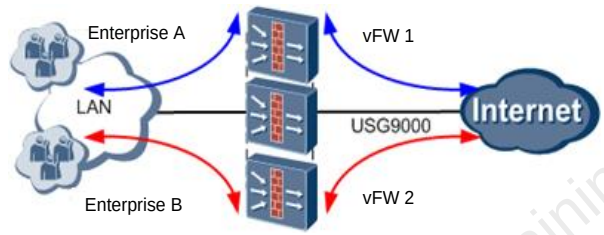

Network Services - Security Group

- Function Description
- A security group functions as the virtual firewall of a cloud host, controlling the incoming and outgoing network messages. Only authorized messages are allowed to pass.

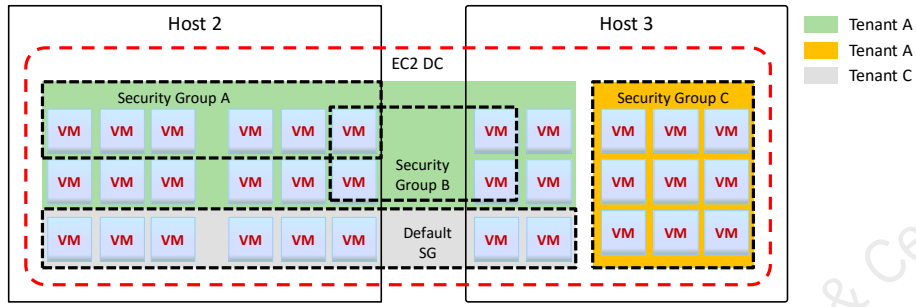

VMs in the same security group can be distributed on different physical servers. The VMs in a security group can communicate with each other, while the VMs in different security groups are not allowed to communicate with each other. If the VMs in different security groups need to communicate with each other, configure security group rules for them.

- A security group provides access control policies for incoming and outgoing traffic of a port. Users can configure packets that can be sent or received by a port using a security group.
- A security group itself does not contain any access control policy. The access control policy is defined in security group rules.
- qbr is a Linux bridge. The TAP device connects VMs with the qbr bridge and corresponds to a virtual network adapter on VMs.
- Security groups are implemented by the Layer 2 agent. That is, the Layer 2 agent, such as neutron_openvswitch_agent and neutron_linuxbridge_agent, will convert security group rules to iptables rules. This happens to all computing nodes in normal cases.
- Firewalls are implemented by the Layer 3 agent. Firewall rules will be converted to iptables rules at the Layer 3 agent node where the tenant router resides.
- A security group can be applied to control any traffic that flows into and out from a VM, whereas a firewall can only monitor cross-network-segment traffic.
- Firewalls are advanced network services and can be applied to service chains, but security groups cannot. Both firewalls and security groups can be configured for the Neutron component to achieve double protection.
- Firewalls can filter out unauthorized external access, saving resources of security groups on computing nodes.
- If traffic passes a firewall, a security group can filter it out to protect the VM. In addition, security groups can prevent internal malicious access.

Network Services - Security Group

- Application Scenarios
- A security group implements access control for VMs within a security group or in different security groups, enhancing VM security. VDC administrators can define different access rules for a security group, and these rules take effect for all VMs added to this security group.
- Prerequisites
- If the security group status of FusionSphere OpenStack is set to supported, turn on the security group switch when configuring the Neutron component. The security group status of ServiceCenter needs to be consistent with that of FusionSphere OpenStack.

Contents

1. ManageOne Architecture
2. ServiceCenter Management
3. **OperationCenter Management**

O&M Architecture

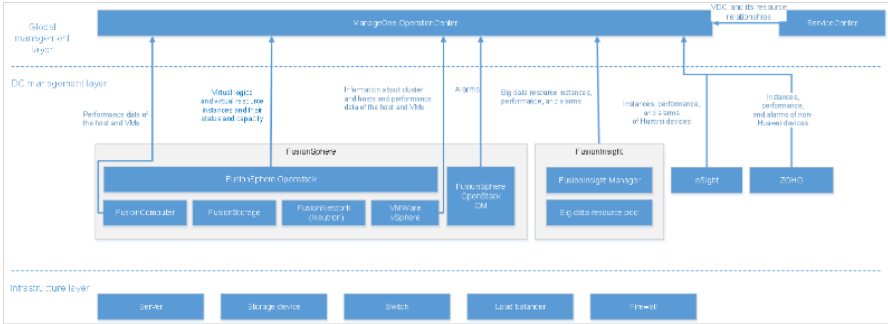

- OperationCenter: DC O&M software, ManageOne O&M center, and unified O&M center of cloud and non-cloud resources of multiple DCs
- FusionSphere: a cloud OS solution developed by Huawei based on the FusionSphere OpenStack architecture for customers from multiple industries
- FusionInsight Manager: OM component of Huawei big data solution - FusionInsight
- eSight: Huawei DC device monitoring software, which is mainly used for monitoring and collecting O&M data of Huawei devices

O&M Scenarios - Roles and Responsibilities

System administrator

Objectives

- Function as the super administrator to configure the system.
- Manage the system.

Concerns

- Security management
- Log analysis

Monitoring personnel

Objectives

- Identify faults in a timely manner.
- Perform preliminary analysis and fault locating.
- Dispatch work orders to the maintenance personnel for troubleshooting.
- Trace the troubleshooting process.

Concerns

- Routine monitoring
- Fault detection
- Risk alarm
- Fault demarcation
- Order dispatching

Maintenance personnel

Objectives

- Rectify faults as soon as possible.
- Periodically learn DC health status and provide optimization suggestions.

Concerns

- Troubleshooting
- Historical performance/ report analysis
- Asset management
- Risk prediction
- Optimization analysis
- Routine maintenance

O&M Scenarios - Initial Configuration

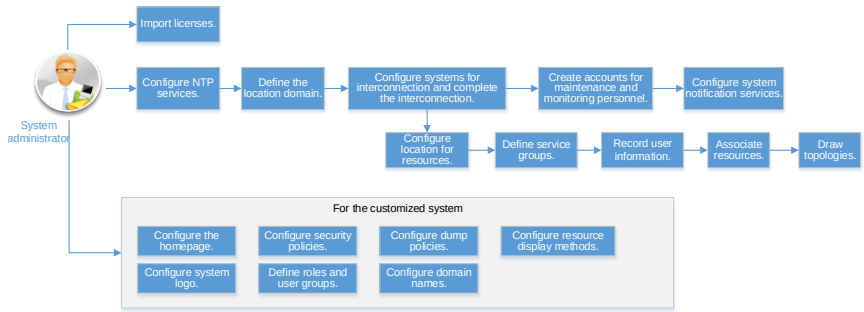

- After initial OperationCenter installation, a three-month default license is provided. Users need to import a license for commercial use.
- Users can customize system functions based on site requirements.
- It is recommended that a unified NTP service be provided onsite to ensure normal system services.
- Based on administrators' location information, OperationCenter controls the scope of resources that administrator of each level can view and perform operations on. Therefore, administrators must define their location information.

O&M Scenarios - Routine O&M

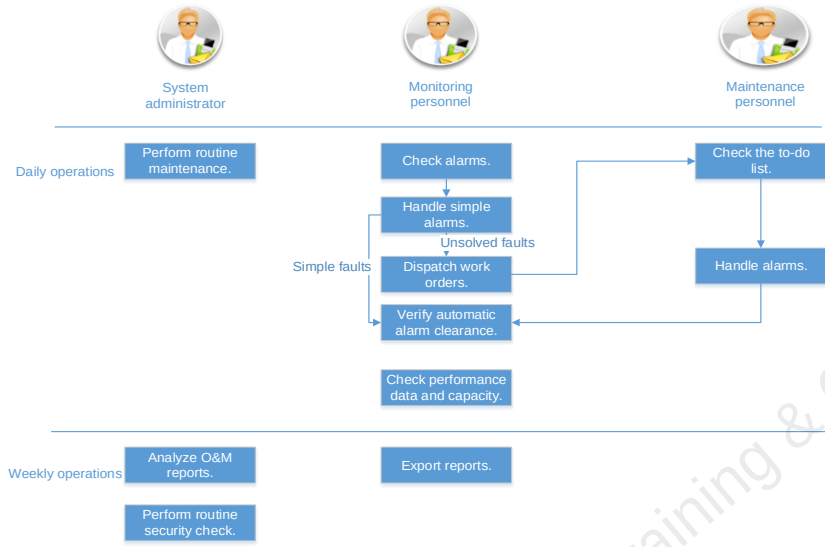

O&M Scenarios - Troubleshooting

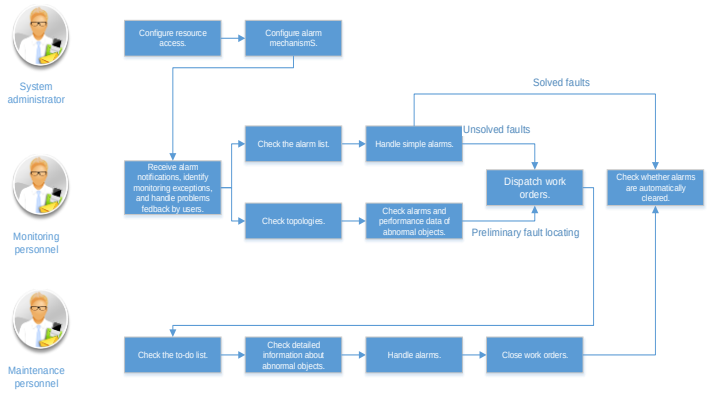

- OperationCenter notifies monitoring personnel of alarms through SMS message or email.
- OperationCenter can interconnect with the work order system to transfer alarms to work orders.
- If no work order system is used to trace alarms online, monitoring personnel need to communicate with maintenance personnel offline to assign the alarm clearance task.

O&M Scenarios - Real-Time Monitoring

Monitoring and maintenance personnel

Configure real-time monitoring views. → Check monitoring results. → Export data.

- Users can obtain performance data of major devices or servers in a timely manner through real-time monitoring.
- Users can also export the performance data for further analysis.

O&M Scenarios - Viewing Reports

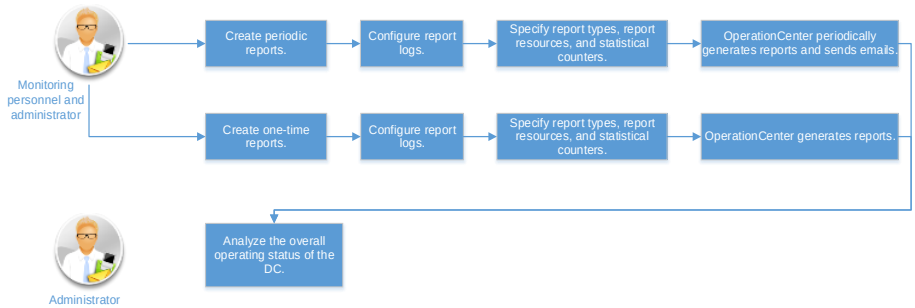

- Monitoring personnel and administrators need to obtain the overall system situation on a scheduled basis. OperationCenter provides abundant report features.

O&M Scenarios - Log Audit

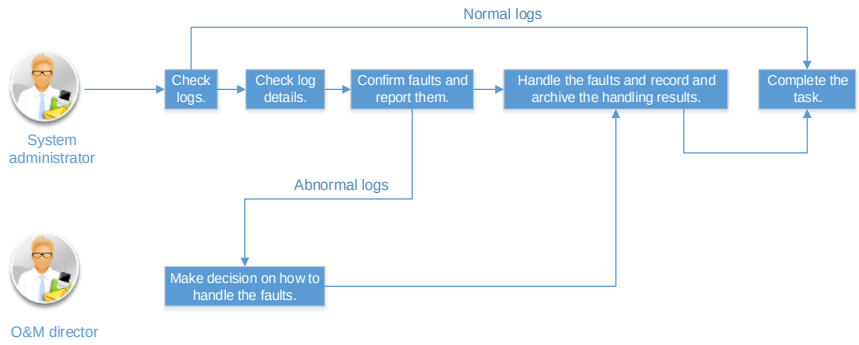

- When security audit is performed based on security policies, the system administrator checks logs on OperationCenter to identify exceptions. If no exception occurs, no further action is required.
- If exceptions occur, the system administrator confirms the exceptions and reports them to the O&M supervisor. After the O&M supervisor makes handling decisions, the system administrator clears the problems and records handling results.

O&M Scenarios - Capacity Expansion

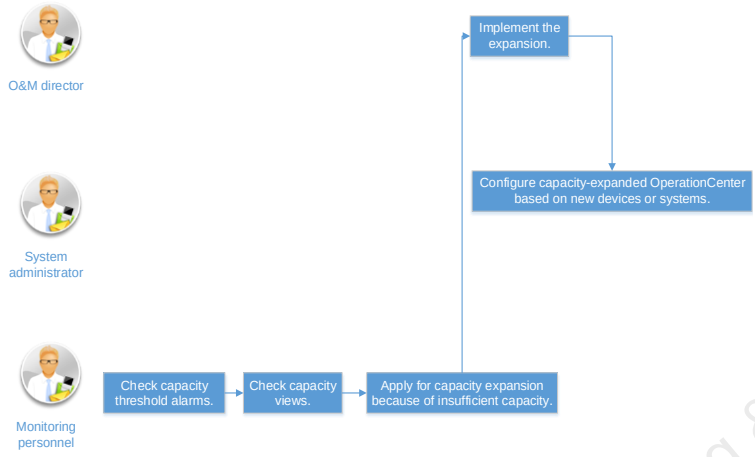

- Monitoring personnel can identify the capacity expansion requirements of the current system in a timely manner using the capacity threshold alarms and the capacity monitoring function.

O&M Operations - Homepage (1/3)

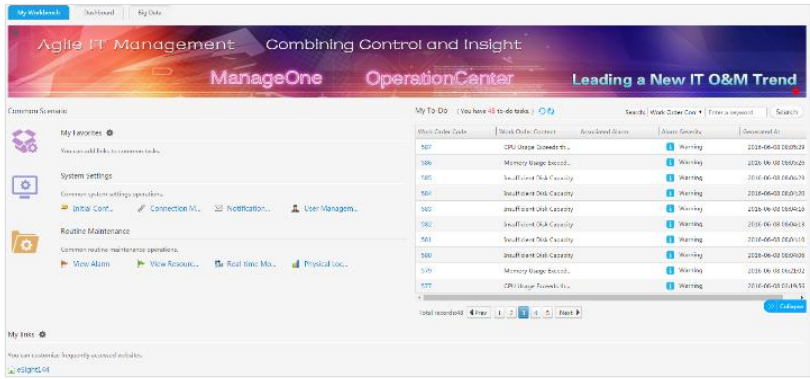

- Users can customize the banner.
- Common Scenario: provides common function portals. Users can customize their own common scenarios.
- My links: functions as the main portal for all subsystems after SSO configuration.
- My To Do: allows users to quickly query their to-do tasks in the work order system.

O&M Operations - Homepage (2/3)

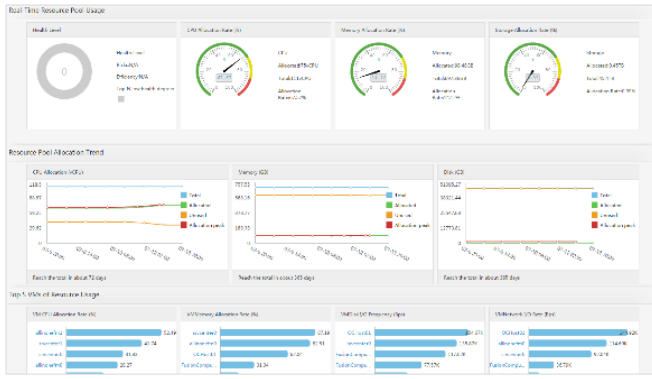

Top *N* Virtual Resource Pools with Highest Usage

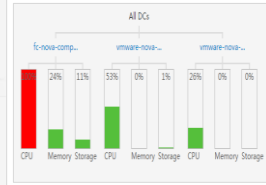

- Dashboard: provides resource pool usage, top 5 VMs, resource pool allocation and usage trend, top N virtual resource pools with the highest usage, and top 5 physical devices with the highest usage.
- Users can choose to display or hide the preceding items.
- OperationCenter can collect statistics of the preceding items by DC.

O&M Operations - Homepage (3/3)

- After customizing the big screen on the Dashboard page, users can switch to the big-screen mode.
- Items displayed in big-screen mode are the same as those on the Dashboard page.

O&M Operations - Alarm Management (1/5)

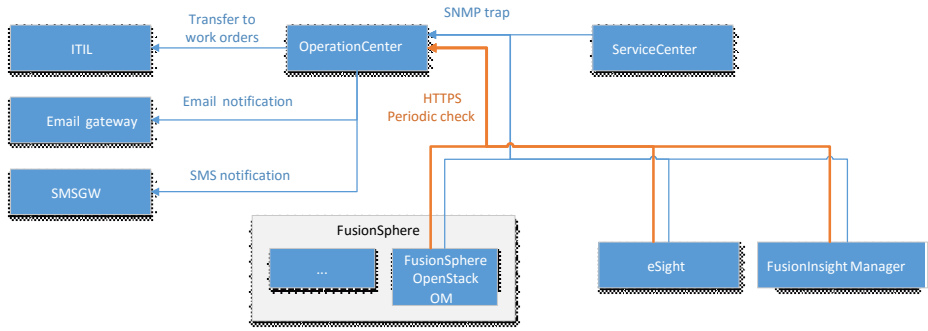

- OperationCenter can obtain all alarms generated on eSight, FusionSphere OpenStack OM, and FusionInsight Manager.
- OperationCenter synchronizes alarm information from lower-layer management systems using SNMP traps and HTTPS. SNMP traps can be used to identify alarms in a timely manner, and HTTPS can be used to avoid loss of alarm information.
- OperationCenter interacts with the email gateway using SMTP, and interacts with the SMSGW using SMPP, CMPP, SMGP, or SIP to send alarm information to monitoring personnel in a timely manner.
- OperationCenter uses the CA SDM system as the ITIL system by default so that alarms can be automatically or manually transferred to work orders.

O&M Operations - Alarm Management (2/5)

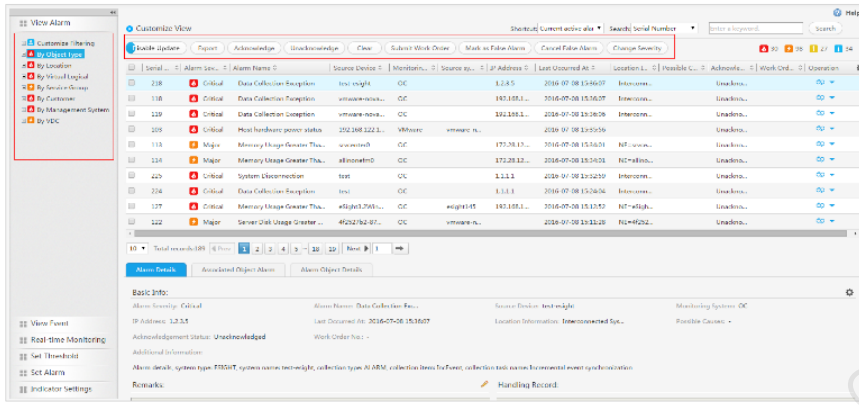

- Query alarm from multiple dimensions.
- Users can export, acknowledge, and clear alarms, and manually update alarm severities, transfer alarms to work orders, and mark and cancel false alarms.
- Allowed event-related operations are similar to those related to alarms.

O&M Operations - Alarm Management (3/5)

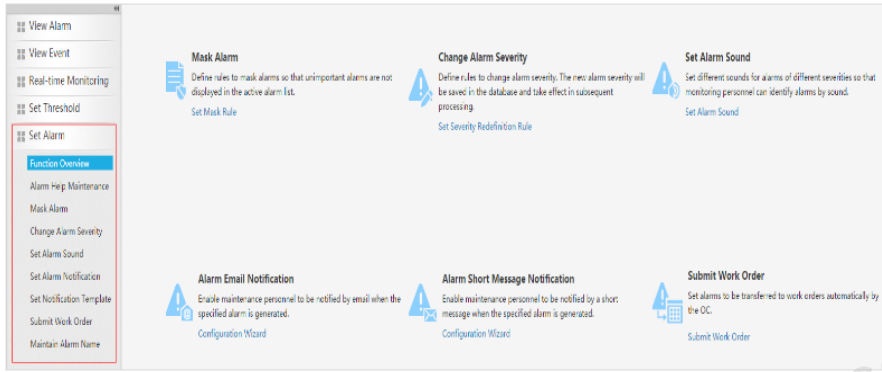

- Alarm settings:
- Include alarm help and maintenance, alarm masking rules, alarm severity redefinition, alarm sounds, notification rules, notification template, automatic alarm-transfer-to-work order rules, and alarm name maintenance.

O&M Operations - Alarm Management (4/5)

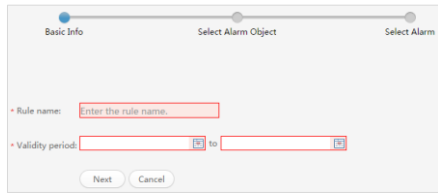

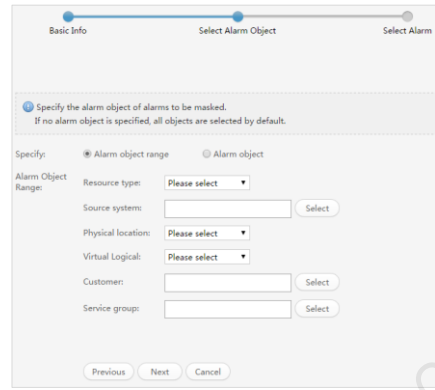

- Alarm masking rule:
- Users can mask unnecessary alarms by setting alarm masking rules. Masked alarms will not be displayed in the current alarm list.
- Alarm masking rules including rule validation period, involved resources, and alarm type.

O&M Operations - Alarm Management (4/5)

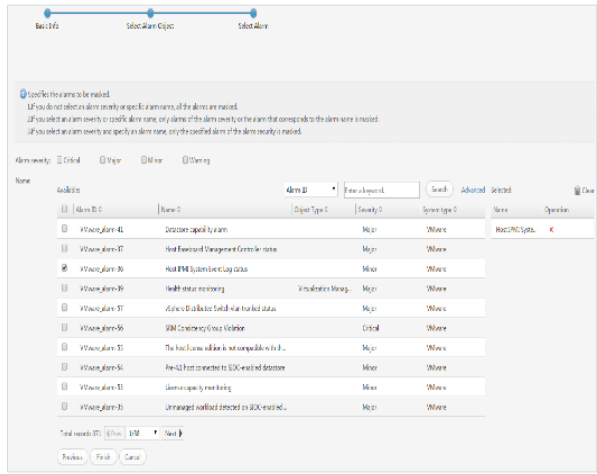

- The relationship between alarm severities and alarm IDs is AND.

O&M Operations - Resource Management (1/5)

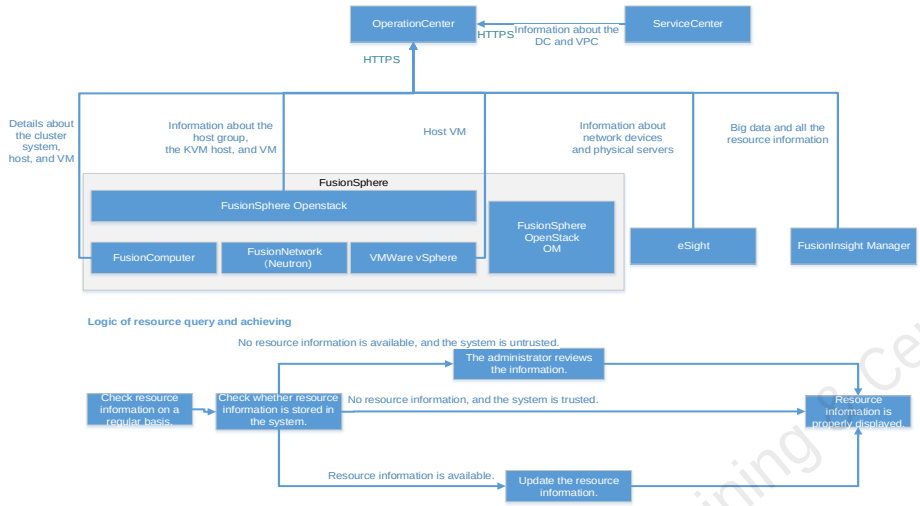

O&M Operations - Resource Management (2/5)

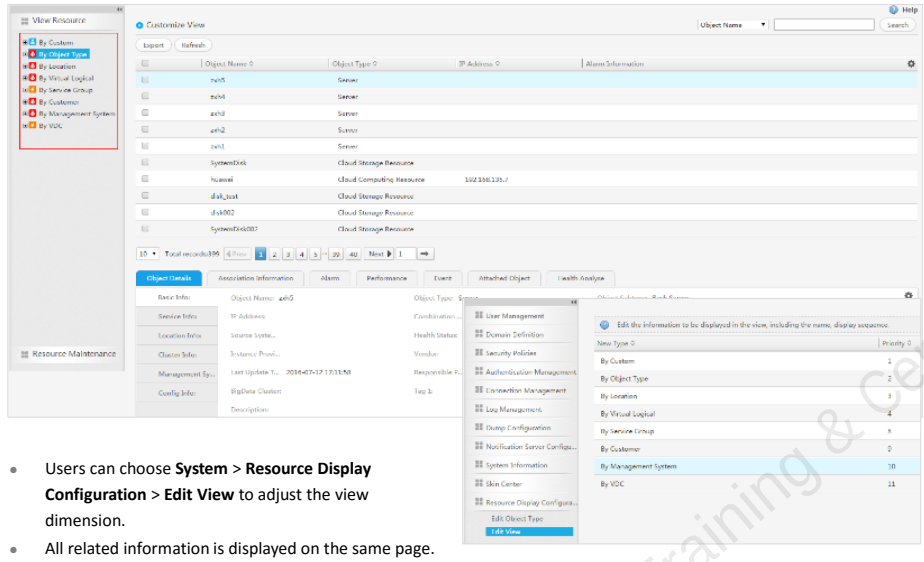

- Users can choose **System** > **Resource Display Configuration** > **Edit View** to adjust the view dimension.
- All related information is displayed on the same page.

O&M Operations - Resource Management (3/5)

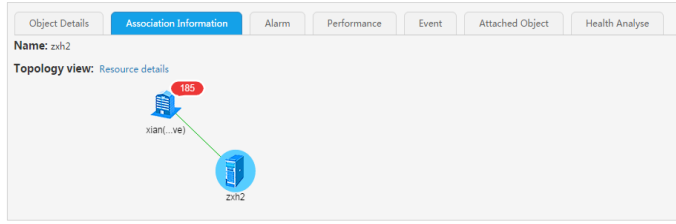

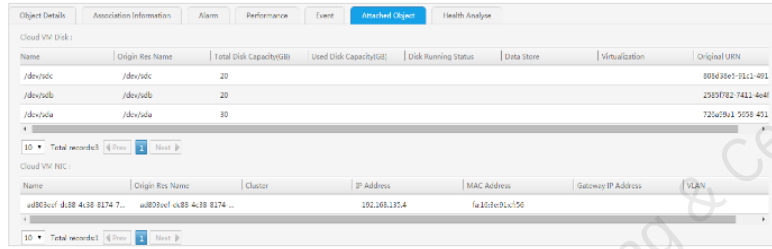

- **Association Information**: displays relationships between the current resource and other items.
- **Attached Object**: displays details about the current resource.

O&M Operations - Resource Management (4/5)

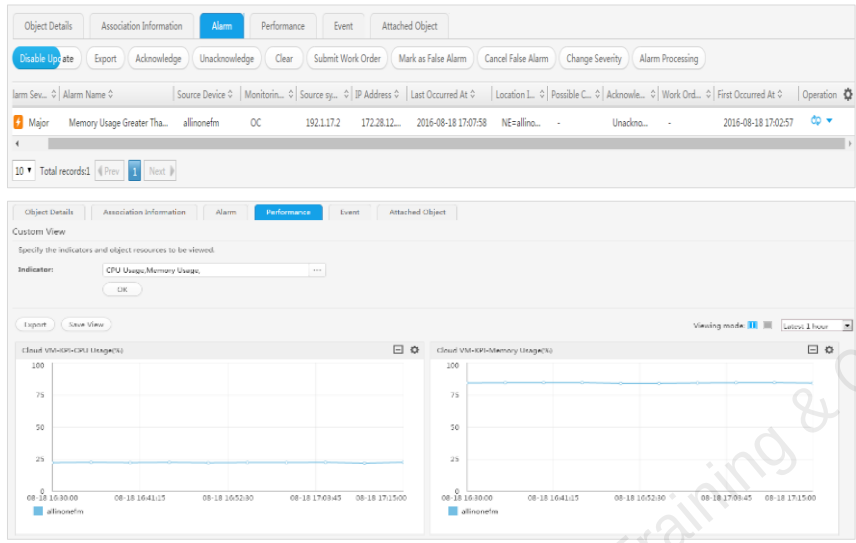

O&M Operations - Resource Management (4/5)

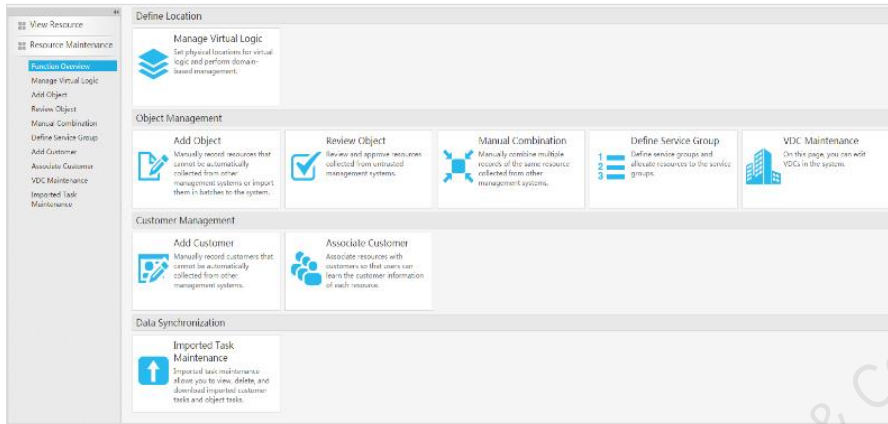

- OperationCenter supports multiple resource maintenance methods. Users can perform operations, such as defining virtual logics, manually importing objects, maintaining VDC and customer information.

O&M Operations - Capacity Monitoring (1/3)

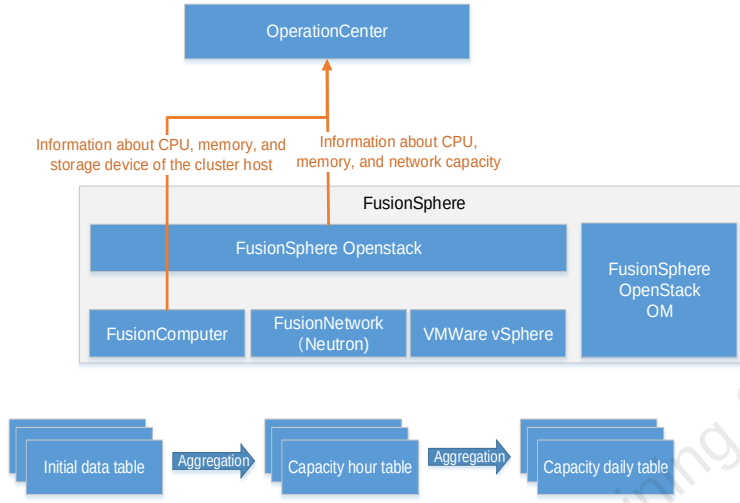

O&M Operations - Capacity Monitoring (2/3)

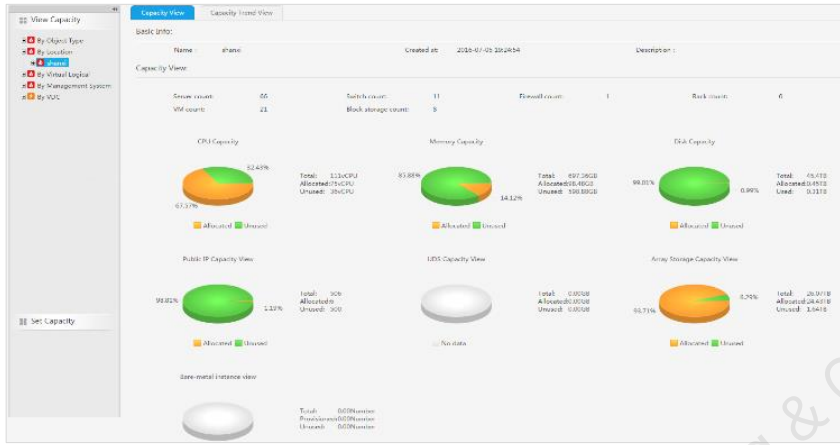

- Users can view capacity from different dimensions (By Object Type, By Location, By Virtual Logical, By Management System, or By VDC).

O&M Operations - Capacity Monitoring (2/3)

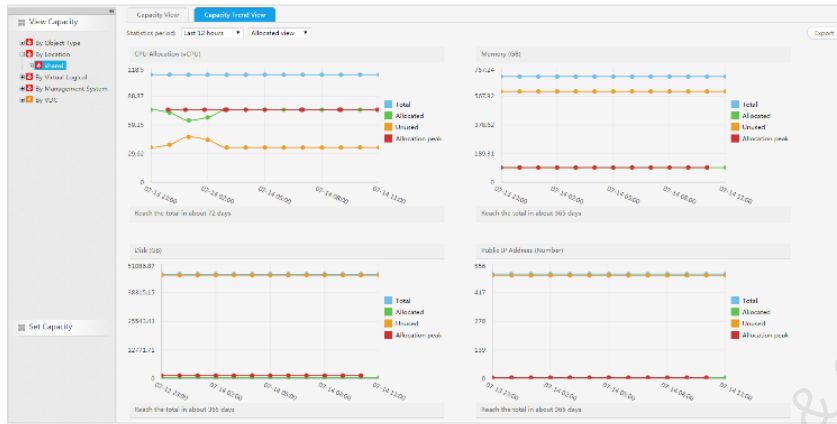

- Capacity trend can be displayed in a non-global granularity (By Object Type, By Location, By Virtual Logical, By Management System, or By VDC). Capacity of the latest 12 hours, or the latest 1, 2, 3, or 7 days can be displayed.

O&M Operations - Topology Navigation (1/4)

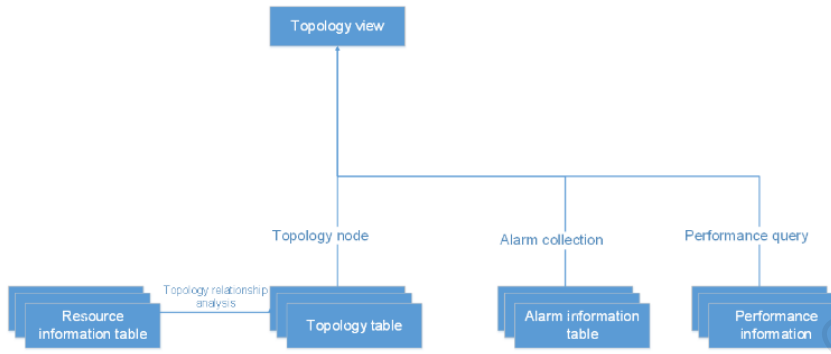

- The topology view displays information about topologies, alarms, and performance in a unified manner.

O&M Operations - Topology Navigation (2/4)

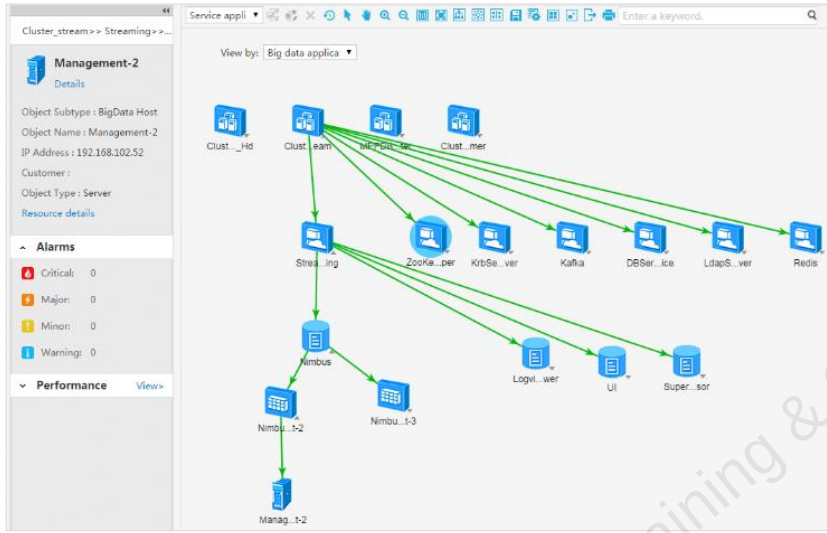

O&M Operations - Topology Navigation (3/4)

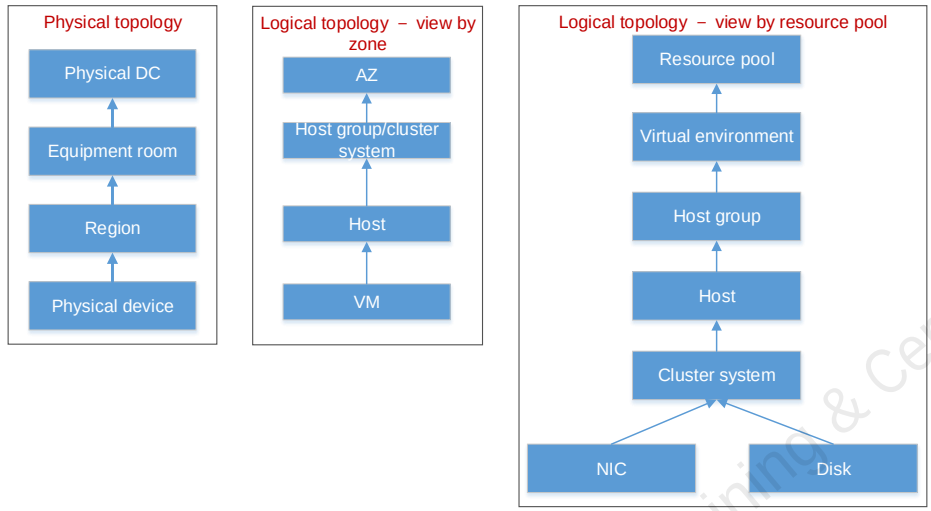

O&M Operations - Topology Navigation (4/4)

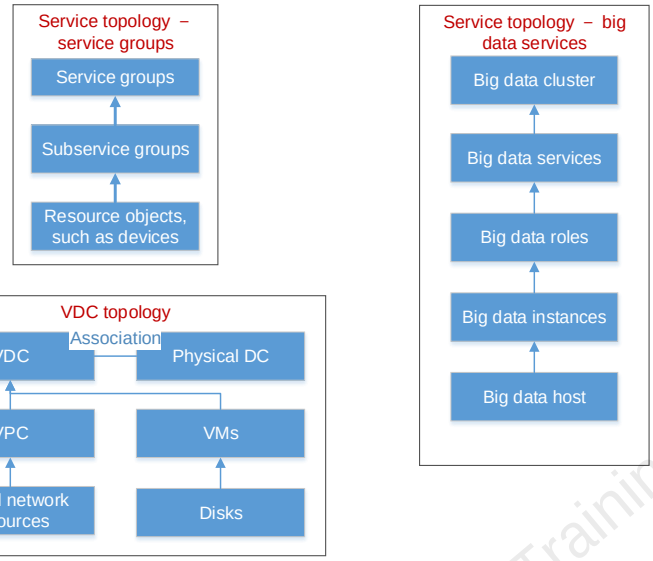

O&M Operations - Performance Report (1/2)

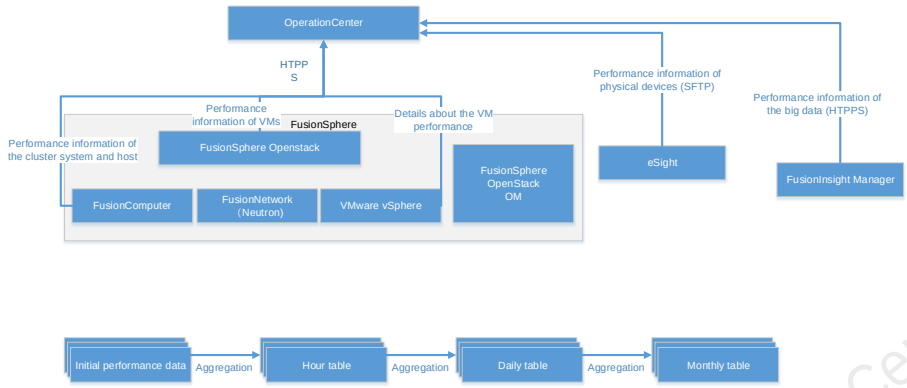

- OperationCenter queries performance from lower-layer management systems on a scheduled basis.
- OperationCenter aggregates the queried data on a scheduled basis to relieve pressure on the database and improve query efficiency.

O&M Operations - Performance Report (2/2)

- OperationCenter provides multiple types of manual reports and periodic reports.
- Users can quickly view the DC status through the sky map.

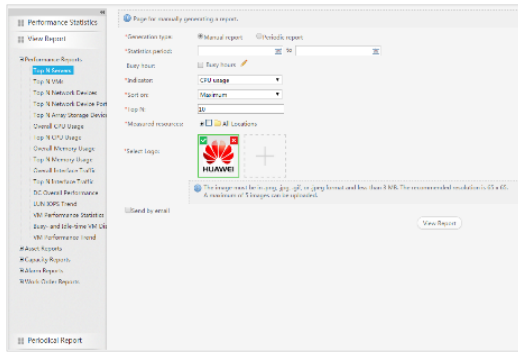

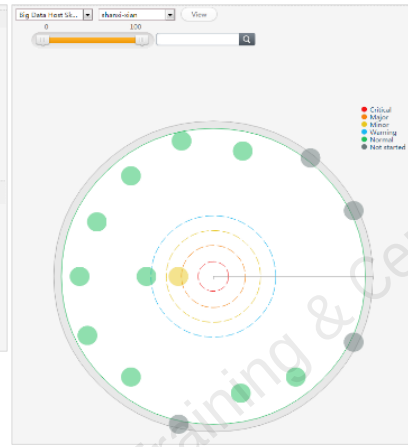

- The VM sky map displays overall VM usage.
  - VM usage decreases from the center of the circle.
  - A sector indicates a cluster.
  - A point indicates a VM.
  - The color of the point indicates the status of the VM.
  - The size of the point indicates the specifications of the VM.
  - Concentric cells indicate different threshold ranges. Drag the arrow above the circle to display only VM with high usage.
  - Users can locate a VM by its name.

O&M Operations - Intelligent Search

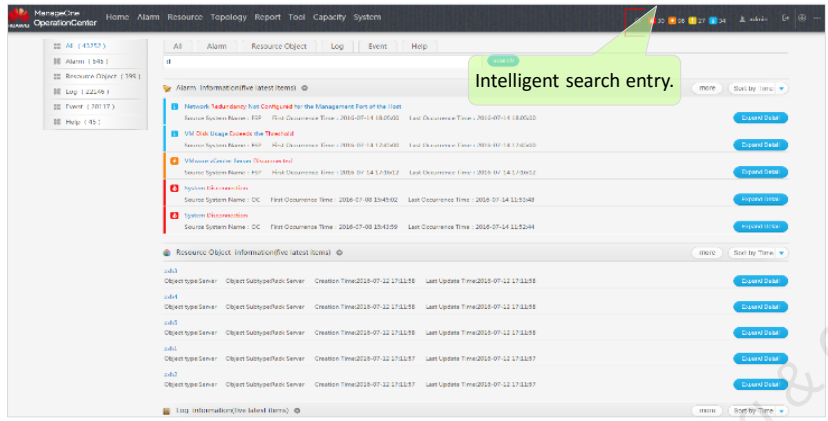

O&M Operations - System Configuration (1/5)

- The following figure shows the process of interconnecting other systems with OperationCenter. Operations in blue are performed on OperationCenter, and those in orange are performed on the corresponding system by its administrator.

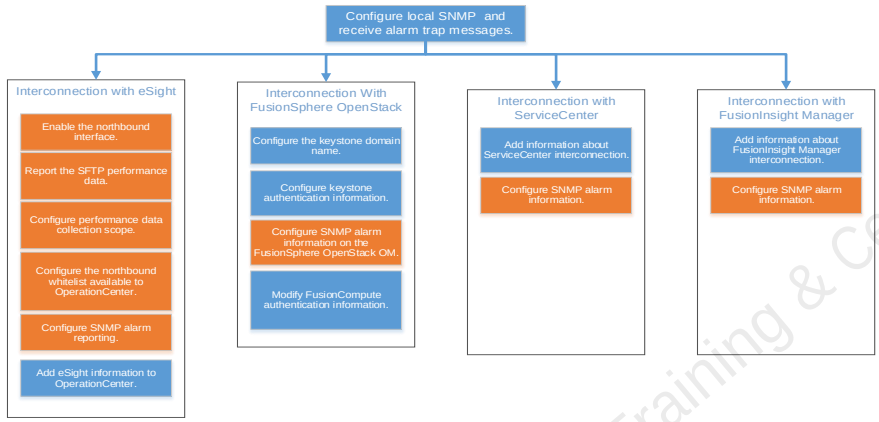

O&M Operations - System Configuration (2/5)

- When users interconnect FusionSphere OpenStack with OperationCenter:
  - OperationCenter automatically detects managed FusionCompute and VMware environments from FusionSphere OpenStack. Users do not need to and cannot manually add interconnected systems. Manual system addition will cause data obtaining errors.
  - If OperationCenter interconnects with a cascaded FusionSphere OpenStack system, users need to choose System > Authentication Management > Configure Keystone to add the domain information about the multiple lower-layer FusionSphere OpenStack systems.
- Select a correct version for interconnection:
  - If the version for interconnection is incorrect, the interconnection seems to be correct, but some data cannot be obtained.
- If users need to delete an interconnected system, delete the information about collected objects and reported alarms about this system as well. Retain such information only when necessary.
  - If such information is not deleted, redundant data may exist on OperationCenter. If the data needs to be deleted later, users need to delete the data by type.

O&M Operations - System Configuration (3/5)

- When creating a user, users need to specify the role and domain rights using either of the following methods:
  - Directly specify the role and domain rights.
  - Add the user to a user group.

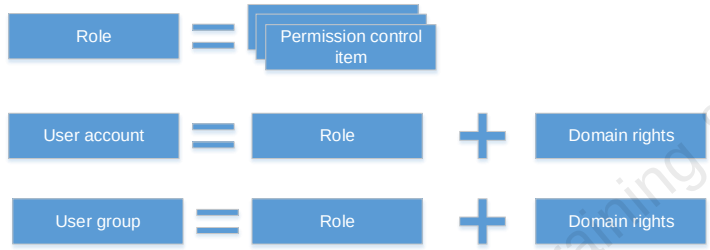

O&M Operations - System Configuration (4/5)

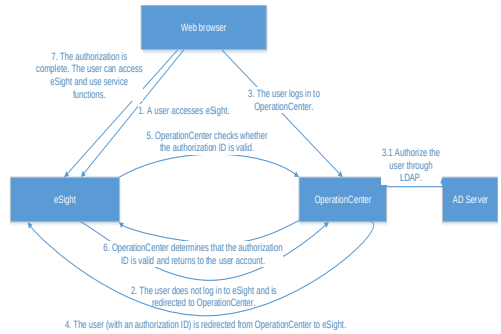

- As an SSO web client, OperationCenter supports SSO from itself to eSight or FusionInsight Manager. Additionally, OperationCenter can use an ADS as the third-party user library (optional).
- SSO requires the same OperationCenter account on eSight or FusionInsight Manager. If SSO is configured, when users create an account on OperationCenter, OperationCenter creates the same account on eSight or FusionInsight Manager by default to support SSO.
- OperationCenter can use an ADS as the third-party user library. When users configure interconnection between OperationCenter and an ADS, ensure that the configured user is the administrator and is under the User directory.

O&M Operations - System Configuration (5/5)

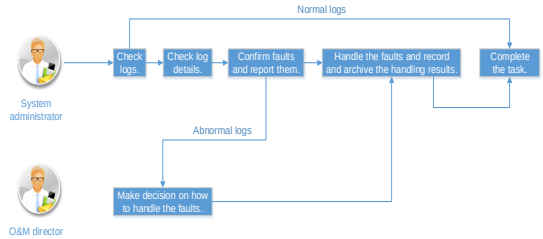

- Log management
  - Security logs record:
    - Users' management operations, such as addition, deletion, and modification.
    - User login and logout
    - Security policies
    - System-detected attacks
  - System logs record:
    - Automatic system processing procedures, for example, dumping
  - Operation logs record:
    - Non-query operations performed by users using a browser, for example, system interconnection and report creation

O&M Operations - Security Policies

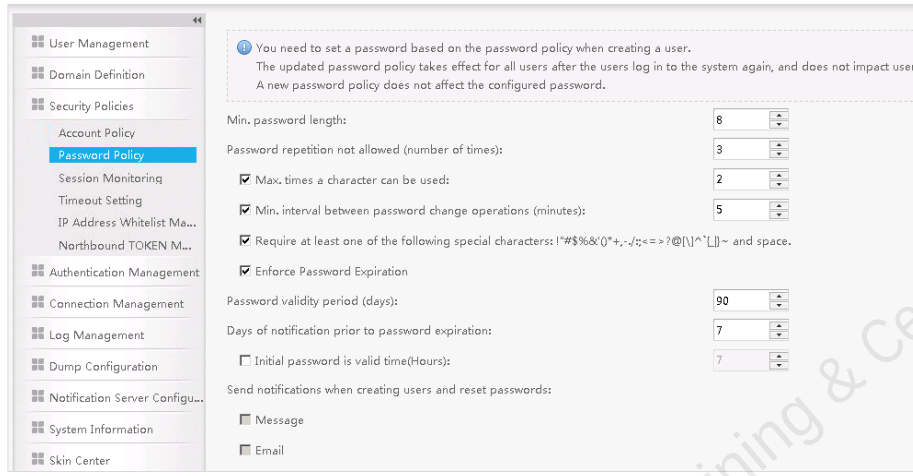

O&M Operations - Health (1/10)

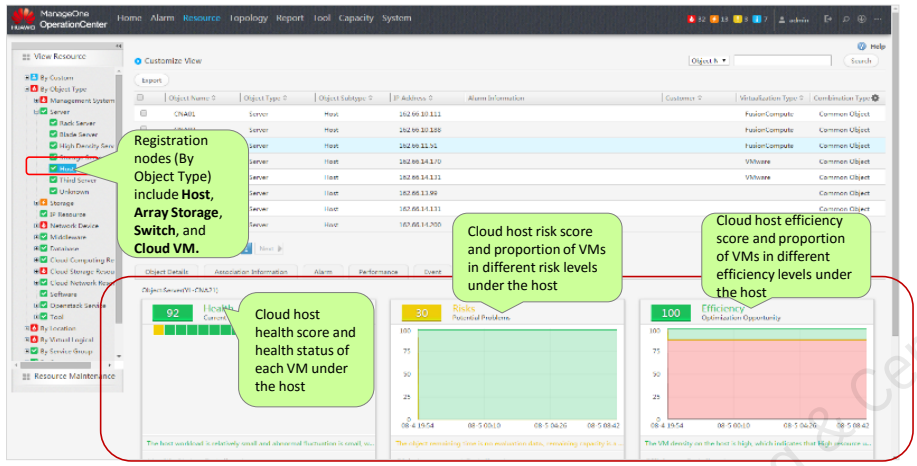

- Resource navigation tree
  - Register health analysis from different dimensions (such as By Object Type, By Location, By Virtual Logical, or By Service Group).

O&M Operations - Health (2/10)

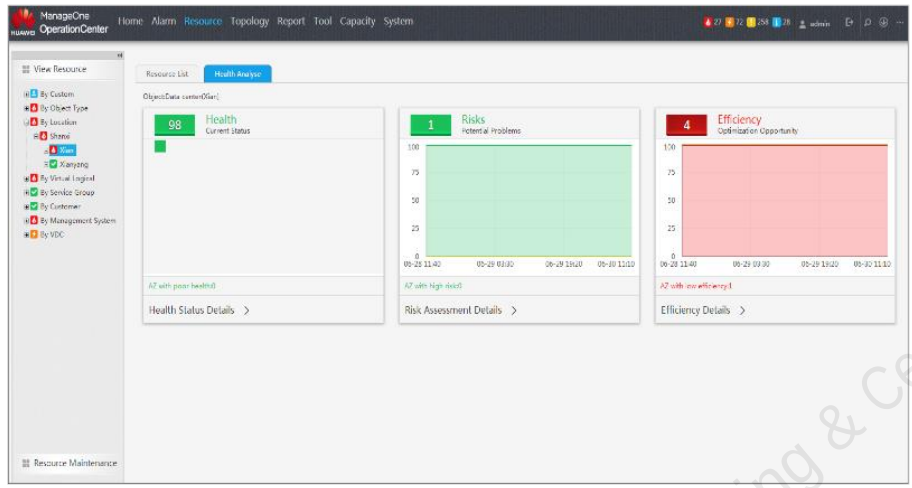

- DC
  - On the **Health Analyse** tab page, the health status of a specified DC and AZs in the DC is displayed.

O&M Operations - Health (3/10)

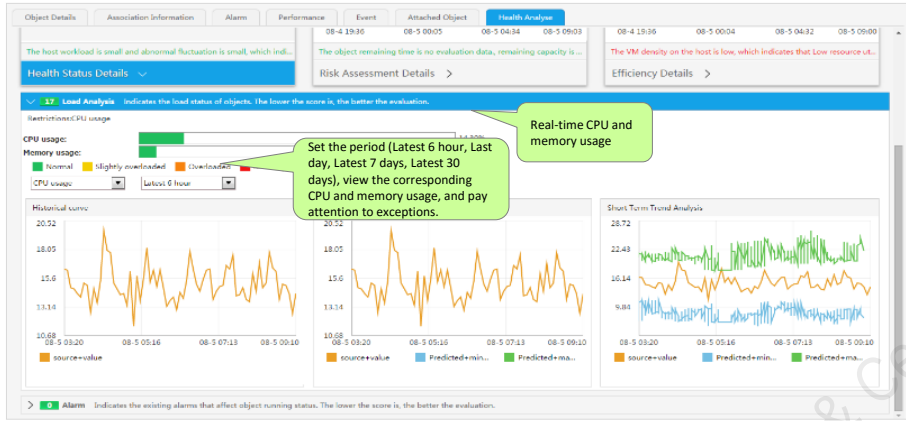

- Health score - load
  - Load score is a sub-item of health score. Performance data of the latest period is collected to score CPU and memory usage. The load score is standard deviation of the doubled average value of CPU and memory usage scores.

O&M Operations - Health (4/10)

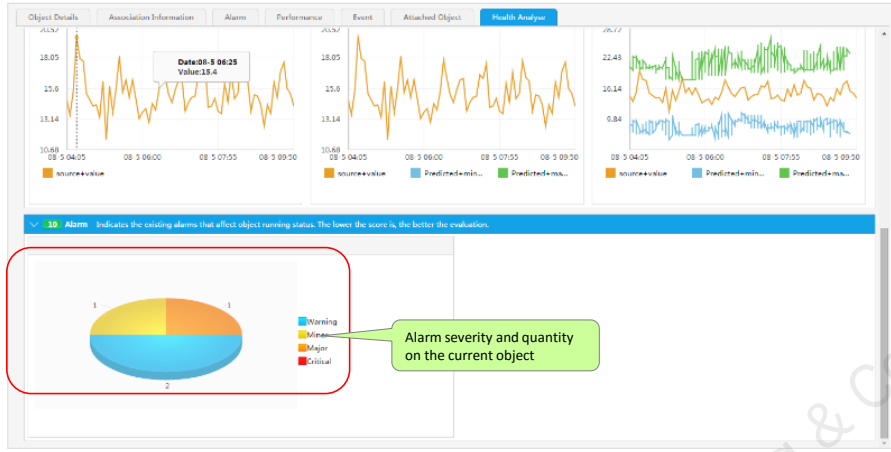

- Health score - alarm
  - Alarm score is obtained based on the number and severities of alarms on the current object.

O&M Operations - Health (5/10)

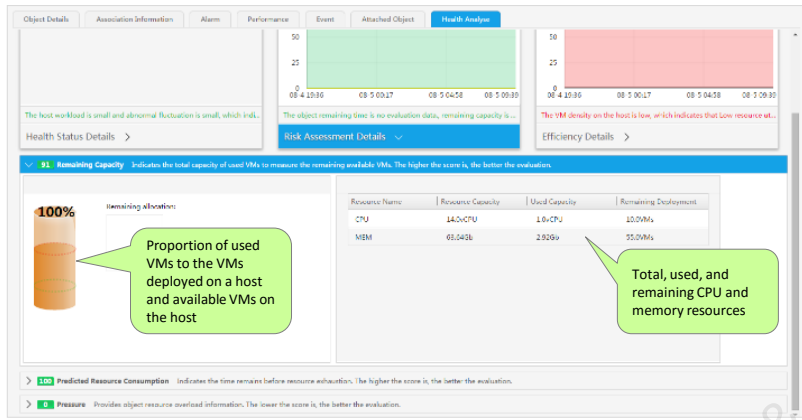

- Risk score - remaining capacity
  - Remaining capacity ratio is the proportion of available VM capacity to the total VM capacity. With this data, you can properly deploy and allocate VMs resources.

O&M Operations - Health (6/10)

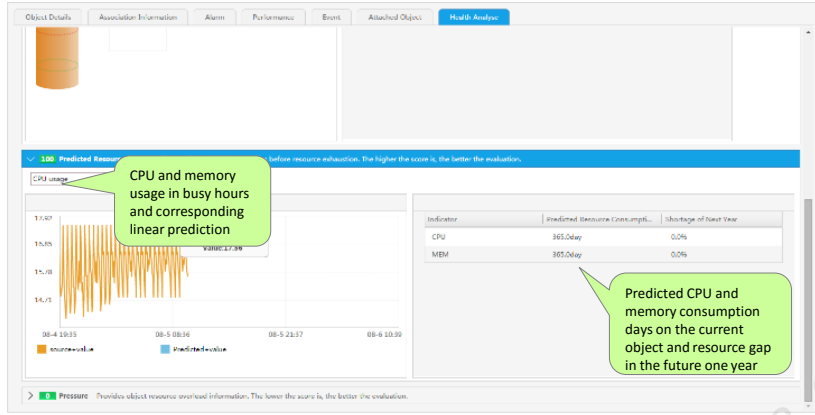

- Risk score - resource consumption prediction
  - Resource consumption prediction focuses on CPU and memory usage, which indicates remaining time and estimated resource gap and warns a resource risk.

O&M Operations - Health (7/10)

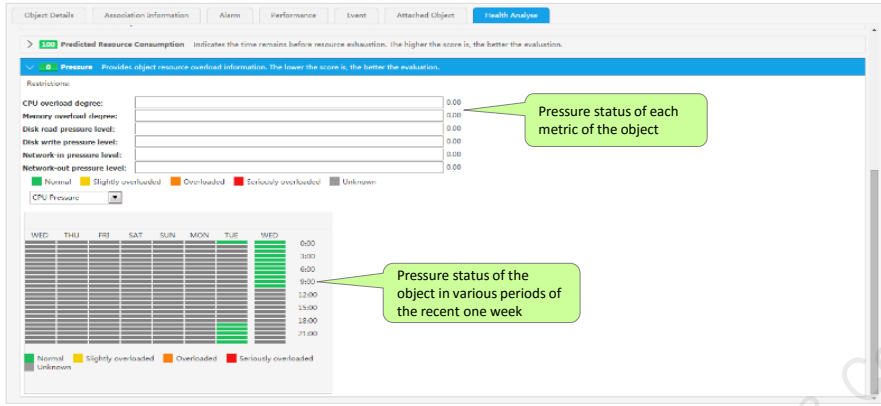

- Risk score - pressure
  - Pressure ratio is the proportion of overload area exceeding the pre-set pressure threshold to the total pressure area, which can be obtained by analyzing load data, such as CPU, memory, disk I/O inbound and outbound traffic rates, and network I/O inbound and outbound traffic rates.

O&M Operations - Health (8/10)

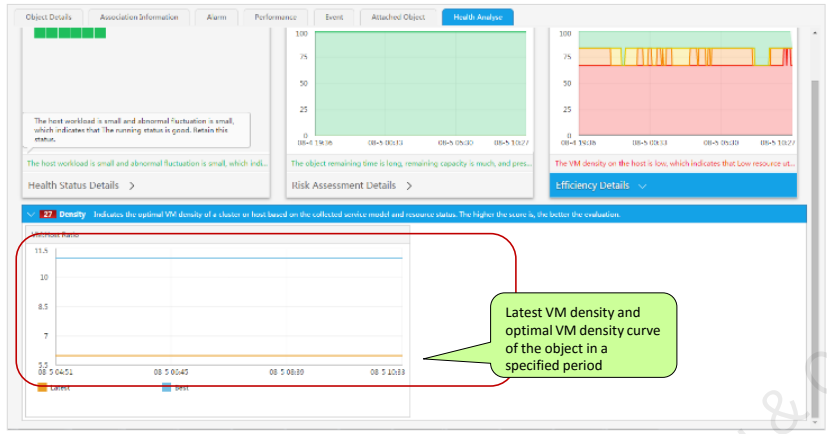

- Efficiency score - density
  - Density indicates the optimal VM density of a cluster or host to maximize resource utilization.

O&M Operations - Health (9/10)

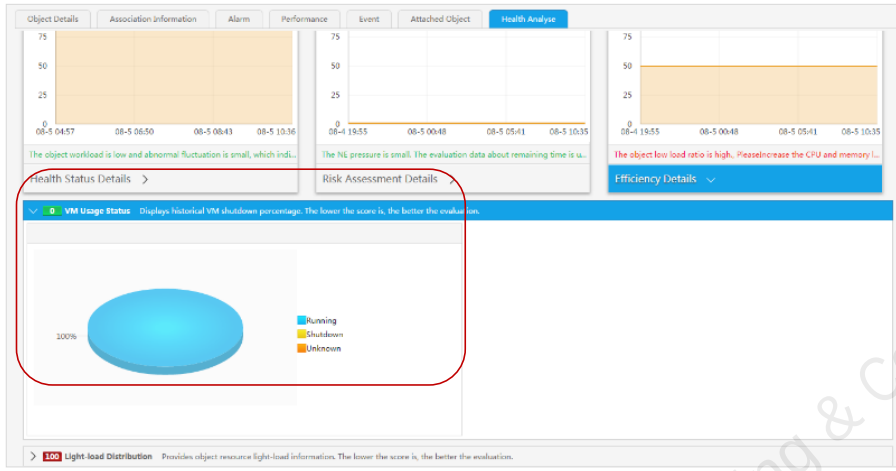

- Efficiency score - VM usage status
  - VM usage status indicates the running status of VMs in the most recent period.

O&M Operations - Health (10/10)

- Efficiency score - low-load distribution
  - Low-load ratio is the proportion of low-load area under the pre-set threshold to the total low-load area, which can be obtained by analyzing CPU and memory load information. A high low-load ratio indicates insufficient resource usage and low efficiency.

Summary

- This course introduced:
  - ManageOne solution architecture
  - ServiceCenter principles and features
  - OperationCenter principles and features

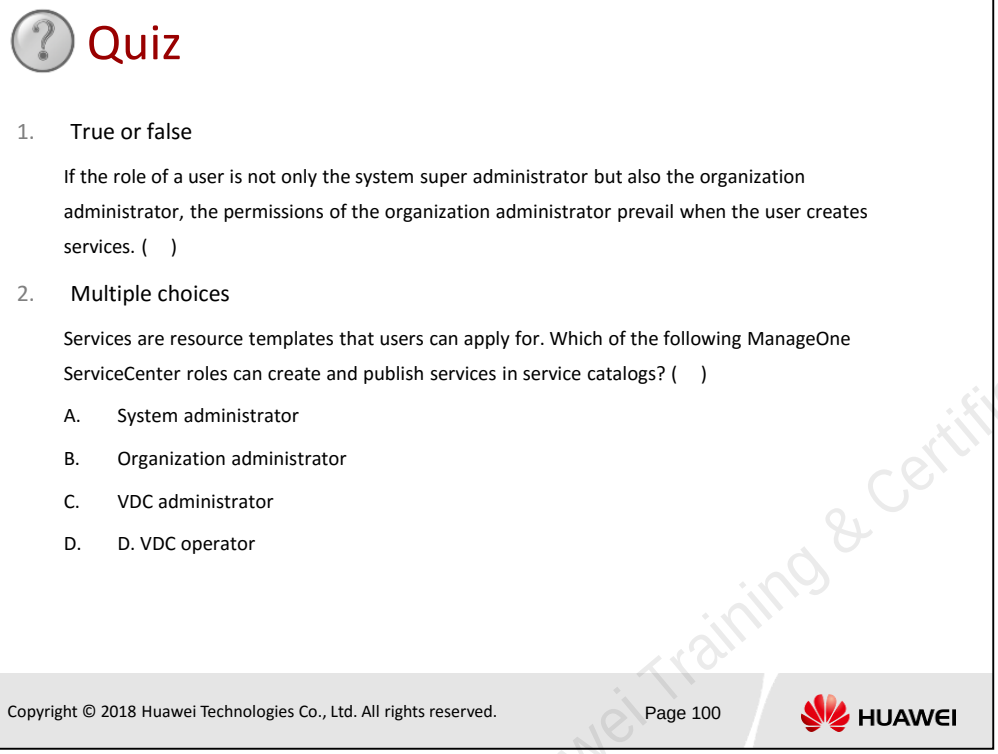

- Answers:
  - True or false: 1. F
  - Multiple choice: 1. AB

More Information

- Huawei official website
  - Enterprise business: http://e.huawei.com/us/
  - Technical support: http://support.huawei.com/enterprise/en/cloud-computing/manageone-sc-pid-21270651?category=product-documentation
  - Online learning: http://learning.huawei.com/en/
- Documentation tool
  - HedEx Lite

Recommendations

- Huawei Learning website:
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei case library:
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Principles and Architecture of VDCs in Cloud DCs

www.huawei.com

HUAWEI

Foreword

- FusionSphere is a Huawei-developed OpenStack-based cloud operating system designed for cloud applications of customers from a wide range of industries. It offers powerful virtualization and resource pool management functions, rich cloud infrastructure components and tools, and open APIs. FusionSphere helps customers horizontally consolidate physical and virtual resources in the DC and vertically optimize service platforms, thereby facilitating construction and usage of enterprise cloud computing platforms.
- This course introduces VDC and VPC service models in the FusionSphere cloud DCs, as well as the computing, storage, and network models and their interconnection configurations.

Objectives

- On completion of this course, you will be able to know:
  - Service models of VDCs and VPCs
  - Computing and storage models and interconnection configurations
  - Network models and interconnection configurations

Contents

1. **Service Models and Concepts**
2. Computing and Storage Models and Interconnection Configurations
3. Network Models and Interconnection Configurations

Huawei Training & Certification

VDC Service Principles (1)

- A virtual data center (VDC) defines the encapsulation and boundary of virtual resources used by departments, and indicates a set of virtual resources, including computing, storage, and network resources. Similar to a physical DC, a VDC is a representation of a physical DC at the virtualization layer.
- The resource mapping relationships between VDCs and physical DCs are as follows:

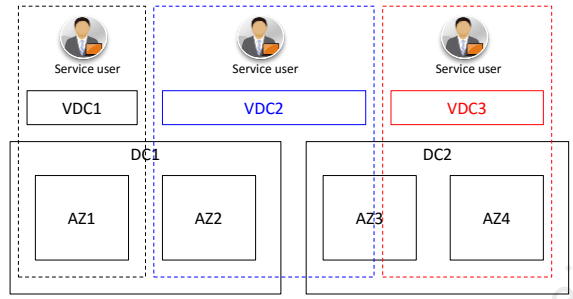

- VDC1 maps to AZ1 in DC1. Virtual machines (VMs) and disks are provisioned to AZ1.
- VDC2 maps to AZ2 of DC1 and AZ3 of DC2. VMs and disks are provisioned to AZ2 and AZ3.
- VDC2 maps to AZ3 and AZ4 in DC2. VMs and disks are provisioned to AZ3 and AZ4.
- Notes:
- 1. The system administrator can set AZ1 exclusive to VDC1. The exclusive function cannot be set on ServiceCenter.
- 2. The quotas allocated to VDC2 and VDC3 determine how many resources in AZ3 they can use.

VDC Service Principles (2)

- The logical mapping relationships between VDCs and OpenStack are as follows:
  - VDCs map to projects in the Keystone and VDC administrators or users map to users in the Keystone. Users correlate with projects through groups.
  - VDC quotas map to quotas for Nova, Cinder, and Neutron. The three components provide project-based quota management APIs.

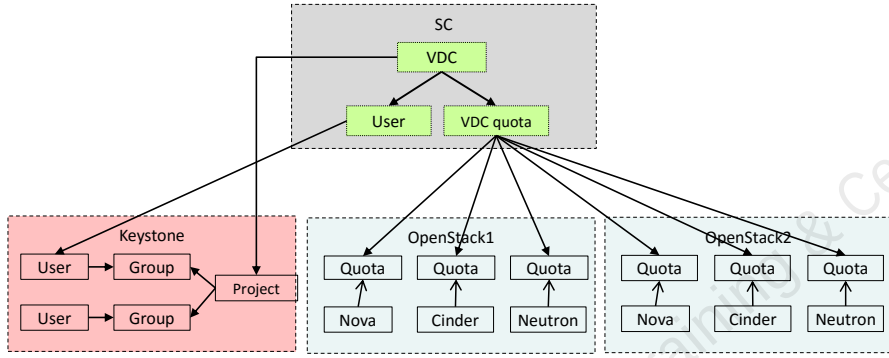

- Multiple OpenStack systems share one Keystone, and thus the project is unique to those OpenStack systems.

VDC Service Principles (3)

- The following figure shows a complete VDC that includes quotas, users, resources, service catalogs, networks, and templates.
  - Quota: A quota indicates the maximum quantity of resources available to a VDC. It is specified by the VDC administrator or organization administrator during VDC creation.
  - User: A VDC involves two user roles: VDC administrator and service user. VDC administrators manage users in VDCs and approve services. Service users are end users of the resources.
  - Resources: Users can apply for service resources, including cloud hosts, cloud disks, and flexible IP addresses. Users can maintain and monitor the resources.
  - Template: The application template is a POC feature.

Quota	
Region	Shenzhen, Beijing
CPU	100
Memory	200 GB
...	...

User	
Administrator	Tom
Server user	Jack
Service user	Jenney
...	...

Resource	
Cloud host	WebInstance
Cloud host	DBInstance

Service Catalog
Cloud host
Cloud disk
Flexible IP address
Virtual DC

VPC


Template
App template
VM template

HUAWEI

- Service catalog: A service catalog lists the services that can be applied for by VDC users. The system administrator can customize service catalogs for different departments.
- Virtual Private Cloud (VPC): A VPC is a logically isolated network environment and provides sub-functions such as vRouter, ACL, and VPN.
- Template: Templates can be categorized into application templates (POC feature) and VM templates. VM templates can be discovered in OpenStack and application templates can be created on SC. VDC users can use global templates or create VDC private application templates.

VDC Service Configuration (1)

- VDCs are allocated by department. One department is allocated one VDC. A department can apply for a VDC to obtain computing, storage, and network resource quota in batches. Within the resource quota, the VDC administrator can allocate the computing, storage, and network resources freely.

Organization Name	VDC Name	Description	AZ	Quota
South Base	Production VDC	Production VDC	az1.dc1	vCPU: 100 cores; memory: 81920 MB; storage: 500 GB
South Base	Development and testing VDC	Development and testing VDC	az2.dc1	vCPU: 100 cores; memory: 81920 MB; storage: 500 GB

- VDC1 maps to AZ1 in DC1. VMs and disks are provisioned to AZ1.
- VDC2 maps to AZ2 of DC1 and AZ3 of DC2. VMs and disks are provisioned to AZ2 and AZ3.
- VDC2 maps to AZ3 and AZ4 in DC2. VMs and disks are provisioned to AZ3 and AZ4.
- Notes:
- 1. The system administrator can set AZ1 exclusive to VDC1. The exclusive function cannot be set on ServiceCenter.
- 2. The quotas allocated to VDC2 and VDC3 determine how many resources in AZ3 they can use.

VDC Service Configuration (2)

- Configuration process

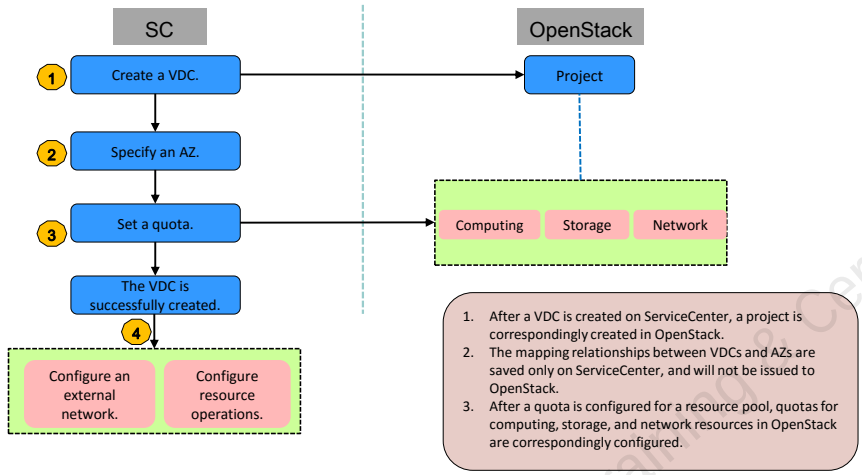

HUAWEI

VDC Service Configuration (3)

- Feature configuration

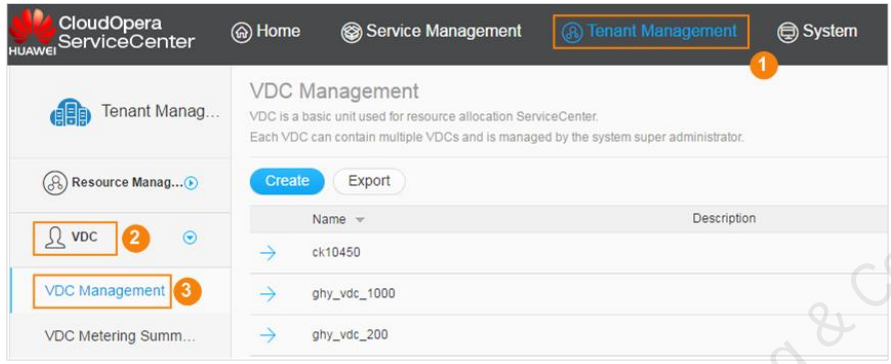

- The associated AZs in a VDC are the same as those associated in an organization.
- After a VDC is created, you can add or delete associated AZs. However, the associated AZs that have applied for resources cannot be deleted.
- External network management: All external networks that have the network egress type configured default to being displayed in the cloud resource pool. You can specify which external network a VDC can use.

VDC Service Configuration (4)

- VDC metering: meters resources in a VDC based on the used and total quota of resources and collects accumulated metering results for settlement between departments. By default, the system meters resources in a VDC based on the used quota.

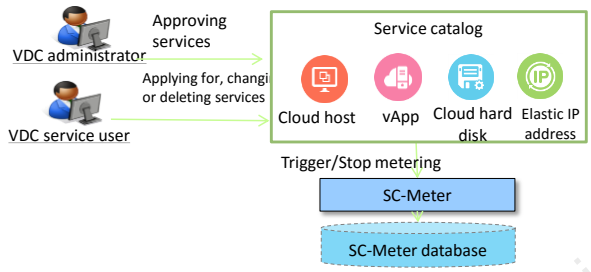

- Used quota:
- Total quota:

VDC Service Configuration (5)

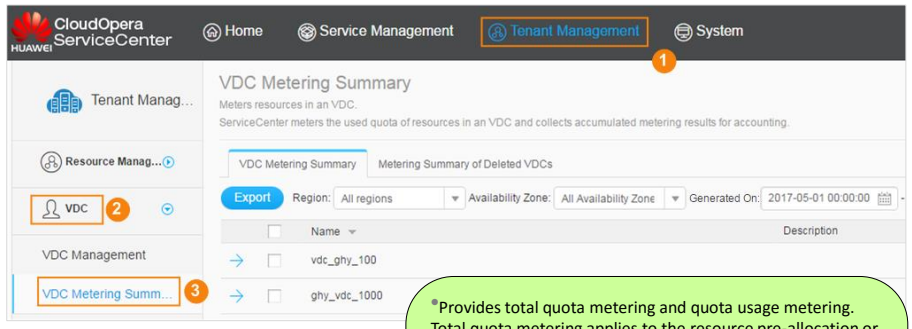

- Provides total quota metering and quota usage metering. Total quota metering applies to the resource pre-allocation or monthly subscription scenario where there is no need to meter specific resource usage. Quota usage metering applies to the scenario where services are settled based on resource usage.
- The system administrator can check the usage of the resources allocated to each department and determine whether the allocated resources are excessive in each department. In the resource renting scenario, the system administrator and tenants settle services based on metering lists.

- ServiceCenter provides usage details in the list for future reference. The list records each service instance request time and amount applied for.

VDC Service Configuration (6)

- Restrictions

No.	Restriction	Remarks
1	Max. number of VDCs supported by the system	System specifications: Small scale: 64 Standard specifications: 128 Large scale: 256
2	If the quota of an organization where a VDC resides is set to limited or unlimited, the quota of the VDC must be set to limited or unlimited accordingly.	
3	The sum of quotas for VDCs in an organization cannot be greater than the total quota for the organization.	
4	VDC associating with Internet: applying for VMs VDC associating with Intranet: applying for routers and VMs VDC associating with Internet and Intranet: applying for routers, VMs, and EIPs	

Components of a VPC

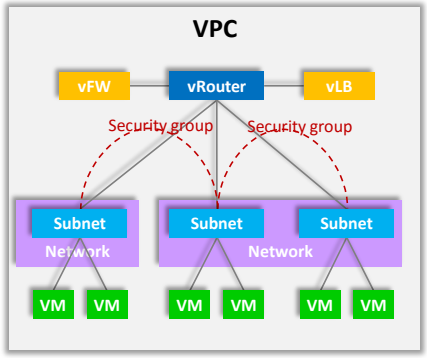

Component	Description
vRouter	• vRouter acts as a service subnet gateway for Layer 3 communication between subnets. • A VPC can have only one vRouter.
Network	• The network is defined as a Layer 2 network. Generally, a network contains only one subnet. In share mode, it can contain more than one subnet. (In share mode, a network is mapped to a switch and one network port enables multiple secondary IP addresses to create multiple network segments. For example, multiple small-sized tenants can share one VLAN.) • Note: FusionSphere does not support the share mode. Therefore, the network component is absent, whereas the component is present in the native OpenStack.
Subnet	• A subnet is used for isolating Layer 2 broadcast domains and corresponds to one subnet segment. • The Layer 3 gateways of different subnets in one VPC are located on the same vRouter. • VMs no matter in one subnet or cross subnets communicate with each other by default. They are isolated by security group.
vFW	• As the VPC edge, vFW can control access from extranet to VPC to protect VPC resources, and provide access service from extranet to VPC. • Available features: FW, EIP, SNAT and IPsec VPN.
vLB	• vLB is used to balance loads between internal servers for external users. • A vLB can carry multiple listeners. Users can apply for different listeners for different services.

Huawei FusionSphere OpenStack Service Model

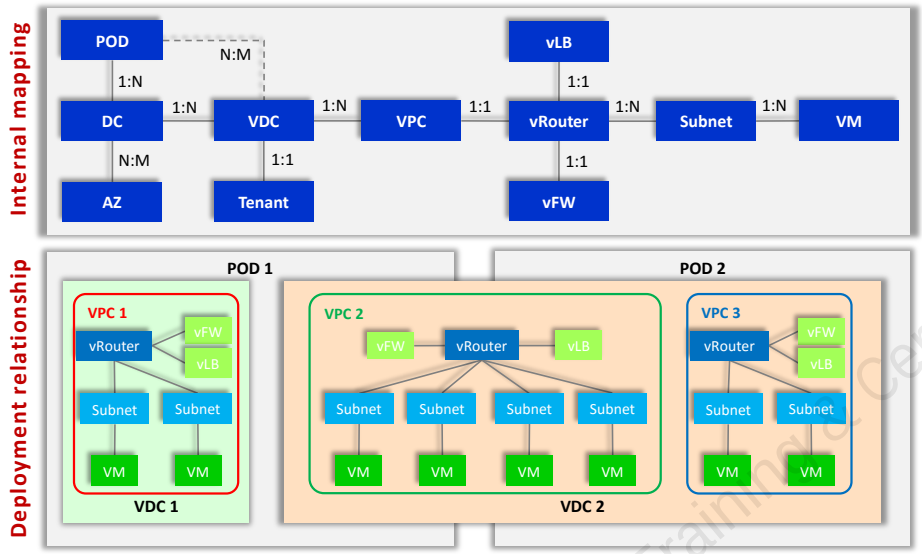

- POD is a physical unit. For example, one Fabric network can be regarded as a POD. One set of resources can be deployed on one or more PODs and one POD can carry multiple sets of resources.
- VDC is a resource unit. It supports cross-POD deployment. Tenants implement resource leasing based on VDCs.
- VPC is a service unit. It supports cross-POD deployment. Different VPCs of a tenant can be deployed on different PODs.

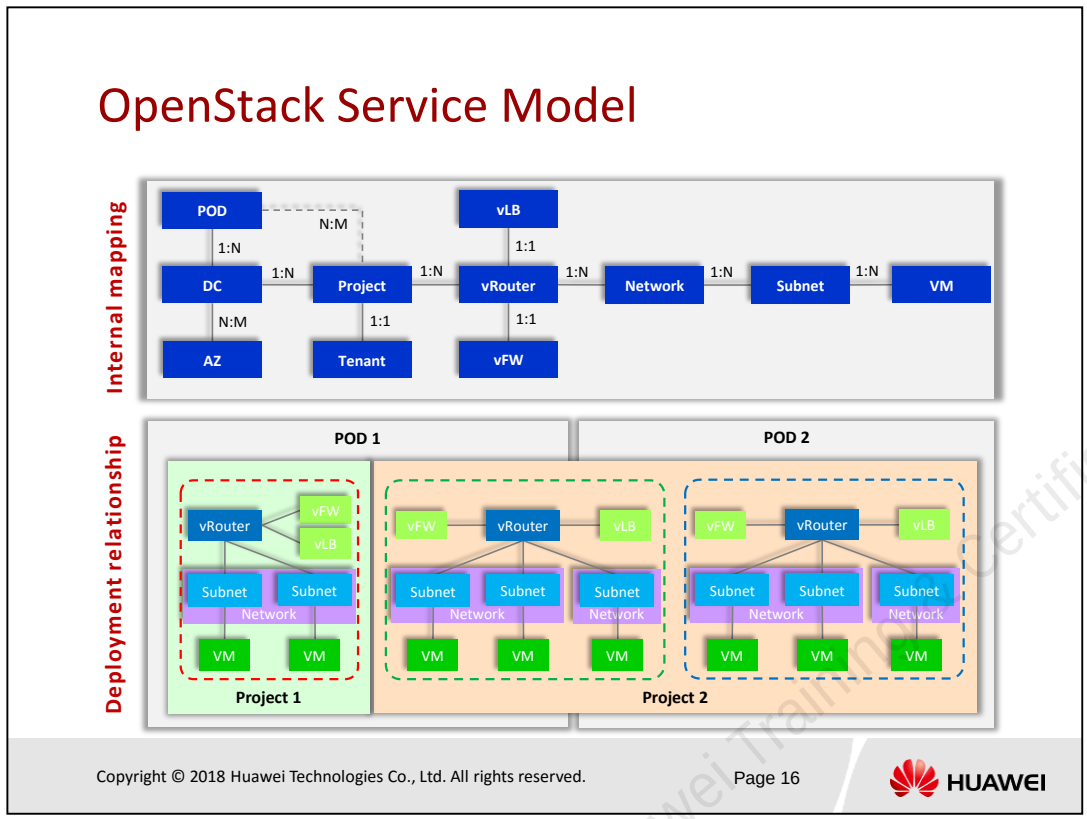

- POD is a physical unit. For example, one Fabric network can be regarded as a POD. One set of resources can be deployed on one or more PODs and one POD can carry multiple sets of resources.
- Project is a resource unit. It supports cross-POD deployment. Tenants implement resource leasing based on projects. In a project, deploy service units with vRouter as the core. The service unit can implement across-POD deployment. Different service units of a tenant can be deployed on different PODs.

Hierarchical Mapping of the OpenStack, FusionSphere, and AC Service Models

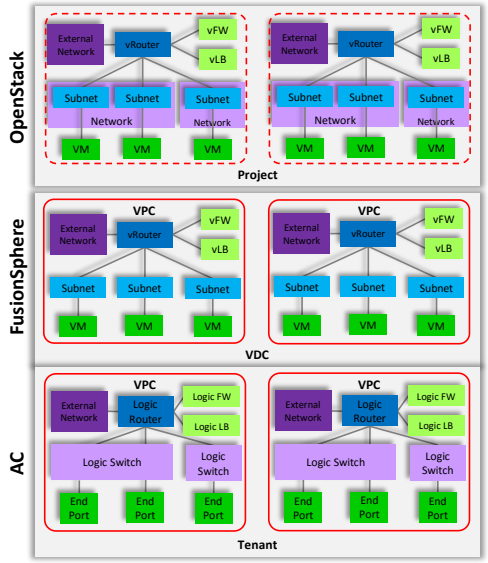

	Resource Mgmt. Layer	Logical Organization Layer	Network Entity Layer	Computing Entity Layer
Description	This layer distributes DC resources based on the tenant granularity, and specifies the corresponding tenant administrator. It is a basic resource unit.	This layer is a logical organization of network and computing entities. It is a basic service unit. The networks between different service units are isolated from each other.	This layer displays various network entities included in a service unit.	This layer displays all computing entities included in a service unit.
OpenStack	• Project/Tenant	• None	• External Network • vRouter • Network/Subnet • vFW/vLB	• VM
FusionSphere	• VDC/Tenant	• VPC	• External Network • vRouter • Subnet • vFW/vLB	• VM
AC	• Tenant	• VPC	• External Network • Logic Router • Logic Switch • Logic FW/Logic LB	• End Port

Contents

1. Service Models and Concepts
2. **Computing and Storage Models and Interconnection Configurations**
3. Network Models and Interconnection Configurations

Mapping Model of Computing and Storage Resources

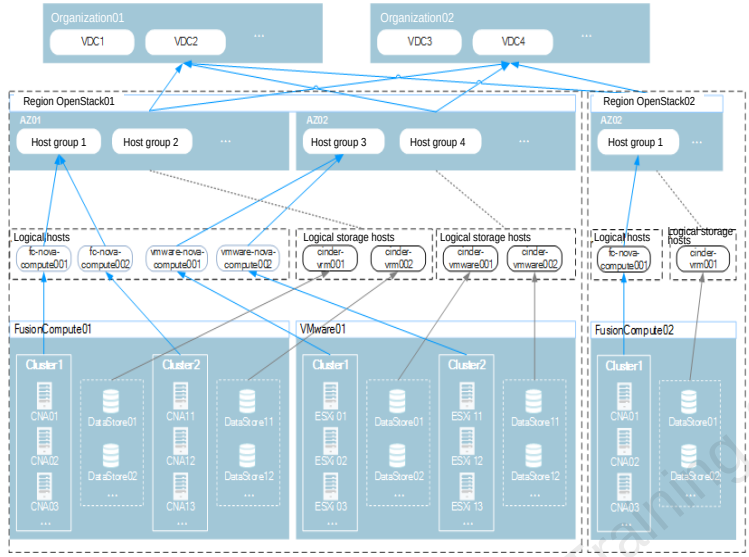

Mapping Model of Computing and Storage Resources

- An OpenStack can contain multiple AZs. The AZs can use different virtual environments.
- Hosts in OpenStack and computing clusters in FusionCompute are one-to-one correspondence. The cluster names set in OpenStack must be exactly the same as those in FusionCompute.
- OpenStack back-end storage and FusionCompute data storage are one-to-one correspondence. The data storage name is set to be exactly the same as that used by the FusionCompute cluster.
- A host group can contain multiple hosts. The hosts must use the same DVS and storage.

Mapping Model of OpenStack and ServiceCenter Resources

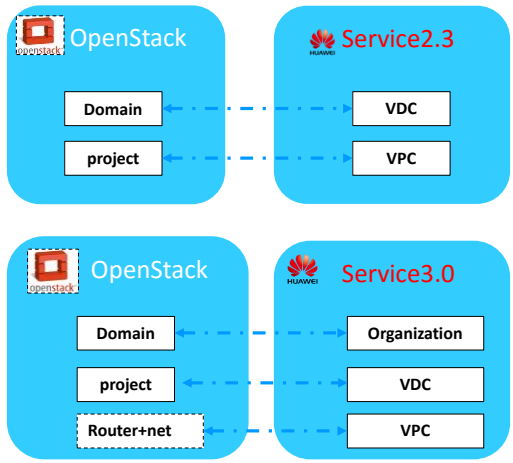

Different from SC 2.3, SC 3.0 has the following mapping relationships with OpenStack:

1. An organization corresponds to a domain.
2. A VDC corresponds to a project.
3. A VPC corresponds to Router+Net.

1. Log in to the database Keystone.

2. View DOMAIN: "SELECT * from DOMAIN;".

3. View PROJECT: "SELECT * from PROJECT".

```
controller03:/home/fsp # su - gaussdba
gaussdba@controller03:~> gsql keystone
```

Region

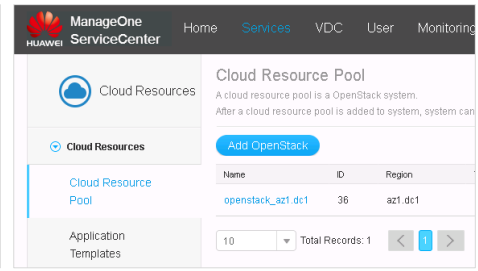

- One OpenStack instance corresponds to one region (that is specifically used in OpenStack). A region name is in the format of a default AZ + a period (.) + a DC name. To check for the region, log in to the REGION table under the Keystone of GaussDB (see the preceding figure).
- Multiple OpenStack instances interconnect with one Keystone for authentication.

AZ

- After OpenStack is installed, a management AZ named manage-az is automatically created. This AZ contains three controller nodes.
- To interconnect with FusionCompute or VMware, create new AZs.
- If a single DC contains multiple sets of OpenStack, AZs need to be planned based on site requirements.

Command for querying the AZ:

```
controller03:/home/fsp # nova availability-zone-list
+-----------------------------------+-----------------------------------------------------+
| Name                              | Status                                              |
+-----------------------------------+-----------------------------------------------------+
| manage-az                         | available                                           |
| |- controller01                   |                                                     |
| | |- nova-compute                 | enabled :-) 2016-06-06T19:50:18.351638              |
| |- controller02                   |                                                     |
| | |- nova-compute                 | enabled :-) 2016-06-06T19:50:18.874063              |
| |- controller03                   |                                                     |
| | |- nova-compute                 | enabled :-) 2016-06-06T19:50:22.985738              |
| internal                          | available                                           |
| |- controller01                   |                                                     |
| | |- nova-conductor               | enabled :-) 2016-06-06T19:50:21.794237              |
| | |- nova-consoleauth             | enabled :-) 2016-06-06T19:50:22.090675              |
| | |- nova-scheduler               | enabled :-) 2016-06-06T19:50:24.222939              |
| | |- nova-console                 | enabled :-) 2016-06-06T19:50:19.128686              |
| |- controller02                   |                                                     |
| | |- nova-conductor               | enabled :-) 2016-06-06T19:50:25.378475              |
| | |- nova-scheduler               | enabled :-) 2016-06-06T19:50:20.543459              |
| | |- nova-consoleauth             | enabled :-) 2016-06-06T19:50:24.315675              |
| | |- nova-console                 | enabled :-) 2016-06-06T19:50:25.954936              |
| |- controller03                   |                                                     |
| | |- nova-conductor               | enabled :-) 2016-06-06T19:50:25.465263              |
| | |- nova-consoleauth             | enabled :-) 2016-06-06T19:50:26.053310              |
| | |- nova-scheduler               | enabled :-) 2016-06-06T19:50:23.513739              |
| | |- nova-console                 | enabled :-) 2016-06-06T19:50:24.955499              |
| vrm001                            | available                                           |
| |- fc-nova-compute001             |                                                     |
| | |- nova-compute                 | enabled :-) 2016-06-06T19:50:24.609704              |
| vmware001                         | available                                           |
| |- vmware-nova-compute001         |                                                     |
| | |- nova-compute                 | enabled :-) 2016-06-06T19:50:25.700819              |
```


Huawei Training & Certification

Interconnecting OpenStack with FusionCompute Resource Pools

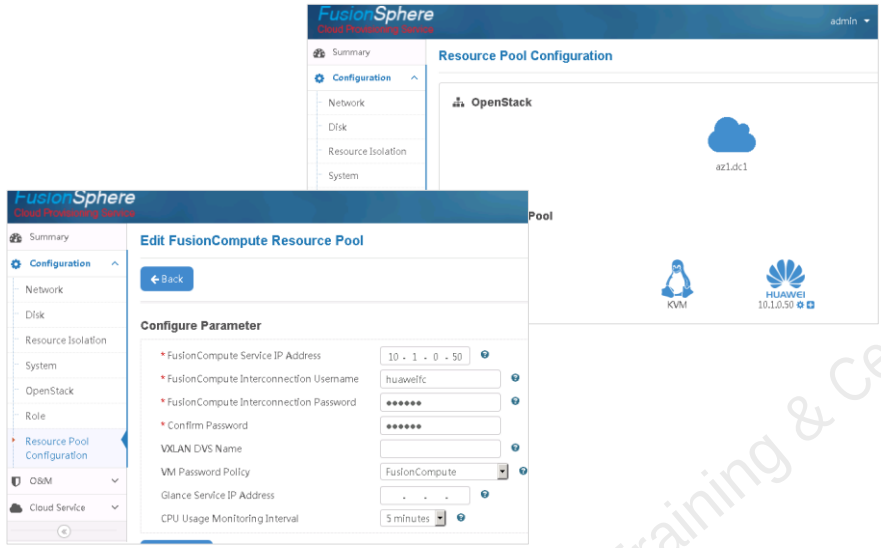

Mapping FusionCompute Clusters to OpenStack

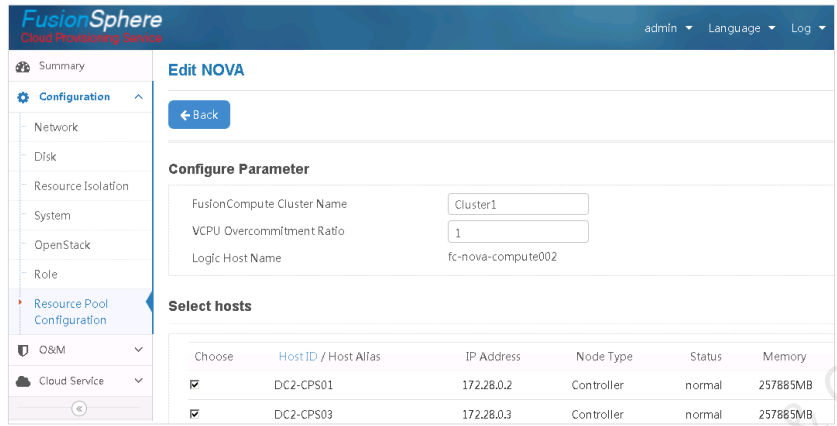

- The FusionCompute cluster name must be exactly the same as that configured on the FusionCompute page.
- After the computing cluster is successfully configured, the system automatically generates the corresponding logical host name fc-nova-compute00x.

Mapping FusionCompute Datastores to OpenStack

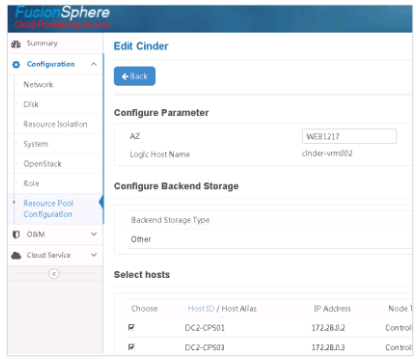

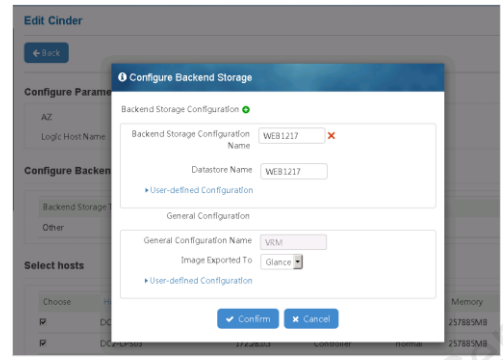

- Each AZ is planned in advance. After Submit is clicked, the system automatically creates it on the background.
- When configuring interconnection of back-end storage with a FusionCompute datastore, ensure that the datastore name must be the same as that configured on the FusionCompute page.
- When configuring storage clusters, add a pair of logical hosts cinder-vm00X to manage storage cluster resources.

Creating AZs and Host Groups on OM

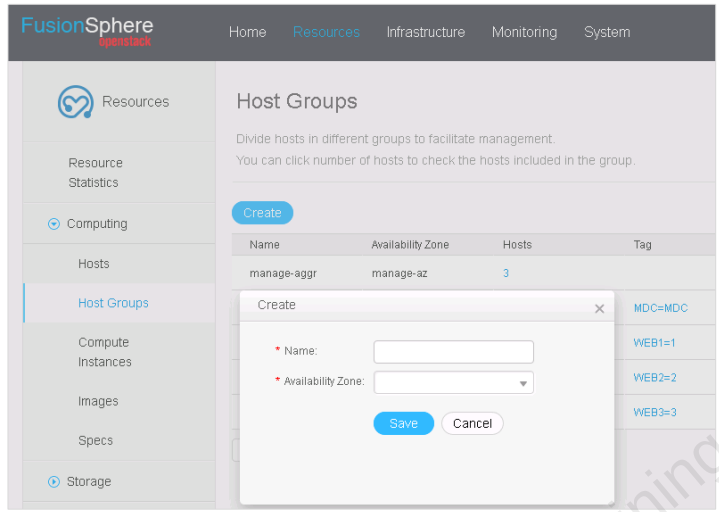

- On OM, create a host group, select an AZ, and add hosts to the host group.

Contents

1. Service Models and Concepts
2. Computing and Storage Models and Interconnection Configurations
3. **Network Models and Interconnection Configurations**

Model of Interconnecting OpenStack with FusionCompute Network Resources

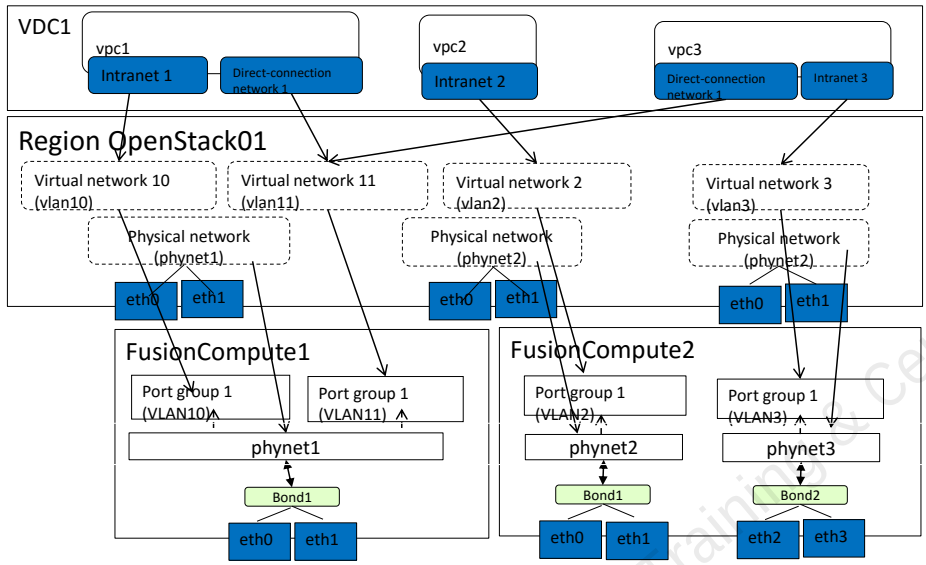

HUAWEI

- The physical network of OpenStack and the DVS of FusionCompute are one-to-one correspondence and share the same name.
- A DVS in a service VM network corresponds to a physical network in OpenStack.
- For example, if FusionCompute has two service planes, service1 (phynet1 for DVS) and service2 (phynet2 for DVS), the OpenStack physical network also has two service planes named phynet1 and phynet2. If the DVS (manageDVS for example) to which the FusionCompute management plane corresponds is used to create a service network, OpenStack must have the corresponding physical network called manageDVS.

Model of Interconnecting OpenStack with VMware Network Resources

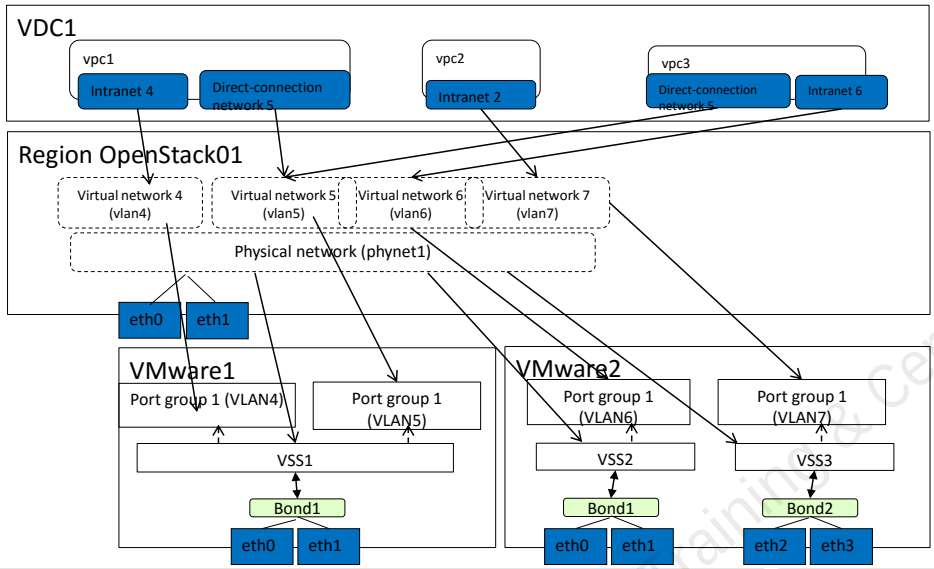

- OpenStack only needs to have one physical network to communicate with VSS of VMware. The physical network can correspond to multiple VSSs and they can have different names.

Interconnecting OpenStack with FusionCompute Network Resources (1)

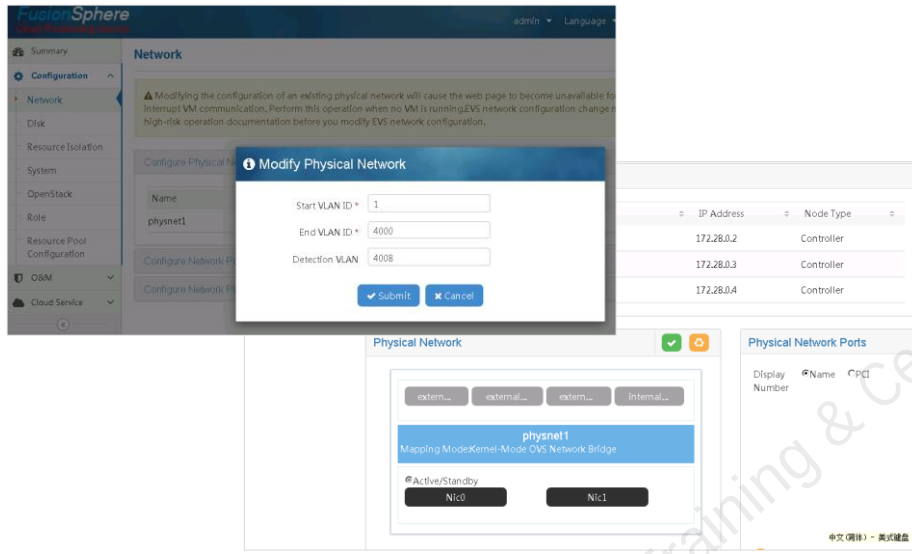

HUAWEI

- On the CPS page, configure network resources, add physical networks, and bind physical network ports to the logical physical network. The physical network name is the same as the DVS name created in FusionCompute. If VMware is interconnected, the physical network name must be the same as the VSS name of VMware.
- For example, If FusionCompute has three planes, management (phynet1 for DVS), service (phynet2 for DVS), and FusionCompute storage (DVS not needed), the physical network of OpenStack must have two planes: phynet1 and phynet2. For reliability purposes, the two physical network ports on each physical network are aggregated into one bond.

Interconnecting OpenStack with FusionCompute Network Resources (2)

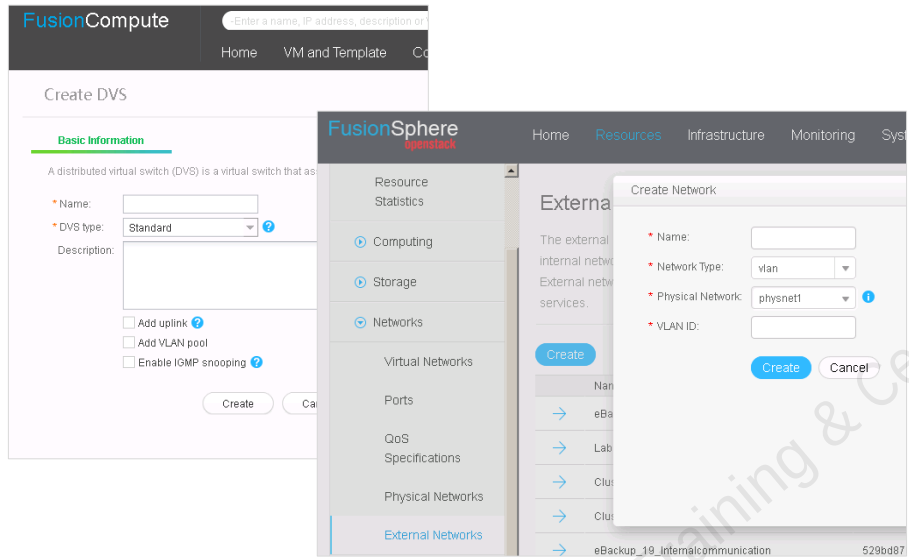

HUAWEI

- Create a DVS on FusionCompute and associate it with an uplink of CNA. The physical network name of OpenStack is the same as the DVS name created in FusionCompute.
- The physical network can be updated and synchronized on OM. A newly created external network must be associated with the physical network.

Creating VPCs and Cloud Hosts on ServiceCenter

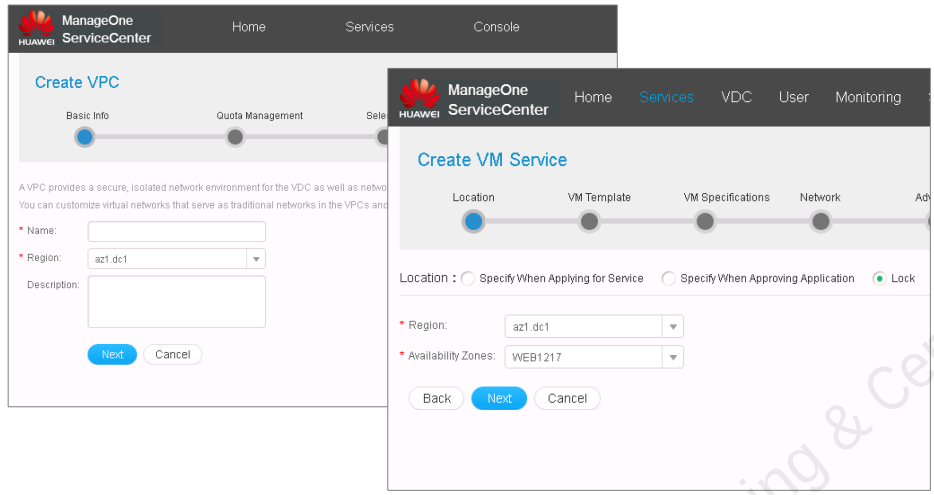

- When creating VPCs on ServiceCenter, select a cloud resource pool.
- When creating a VM, select a cloud resource pool and AZ.

Summary

- Service models and concepts
- Computing and storage models and interconnection configurations
- Network models and interconnection configurations

Recommendations

- Huawei E-Learning
  - http://support.huawei.com/learning/en/newindex.html
- Huawei Support Case Library
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Huawei Cloud Data Center eBackup Principles, Deployment, Operation, and Maintenance

www.huawei.com

Foreword

- This course describes how to install, deploy, initialize, and configure Huawei eBackup and how to back up data to the cloud with eBackup.

Objectives

- Upon completion of this course, you will be able to:
  - Describe the major functions of eBackup.
  - Know how to install and deploy eBackup.
  - Know eBackup Backup and Restore.

Contents

1. **Introduction to eBackup**
2. eBackup Installation and Deployment
3. eBackup Backup and Restore

Huawei Virtual Backup Software – eBackup

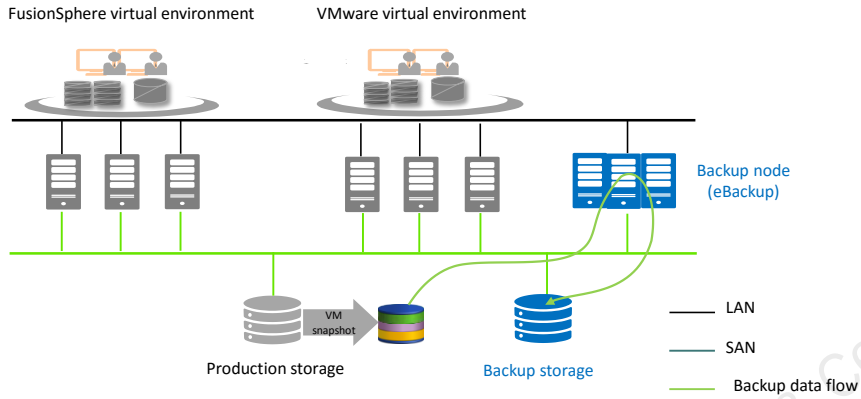

- eBackup is a backup software product designed for FusionSphere and VMware vSphere. Employing the snapshot and change block tracing (CBT) technologies, eBackup provides comprehensive protection for VMs.
- With highlights such as ease-of-use and cost-effectiveness, eBackup meets backup requirements of a large number of VMs.

Distributed Architecture with High Reliability and Scalability

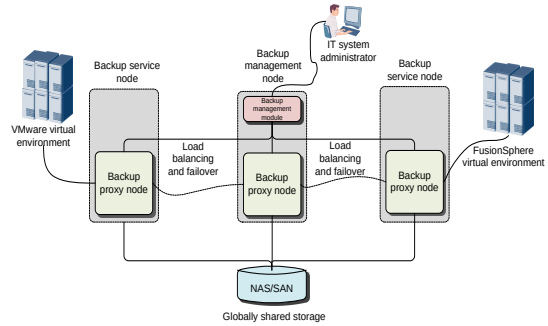

- Distributed architecture: eBackup consists of backup management nodes and backup service nodes. You can add backup service nodes as required. A maximum of 64 backup service nodes are supported, enabling you to back up or restore 10,000 VMs.
- High reliability: Enables you to balance loads among backup nodes during backup jobs. When a node becomes faulty, its backup jobs can quickly fail over to other nodes within one minute.
- High scalability: Enables you to add backup service nodes, improving the backup performance in a linear manner. A maximum of 200 Gbit/s backup speed is supported.

eBackup Software Architecture

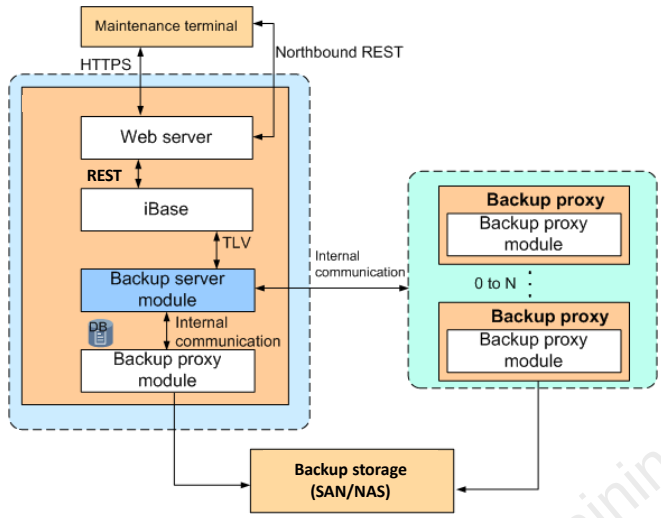

eBackup Network Architecture

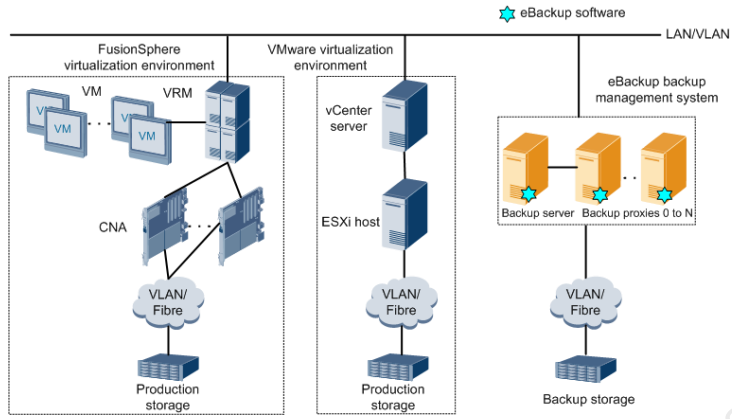

- The eBackup Virtual Backup Solution consists of the following:
  - Backup object: Protects data of VMs on Huawei FusionSphere and VMware vSphere virtualization platforms.
  - eBackup software: Uses a distributed architecture and enables you to add backup proxies to achieve smooth expansion.
  - Backup storage: Supports SAN or NAS, improving backup performance and reliability.
  - Backup policy: Creates backup policies for VMs and one or multiple VM disks.

eBackup Backup Storage Management

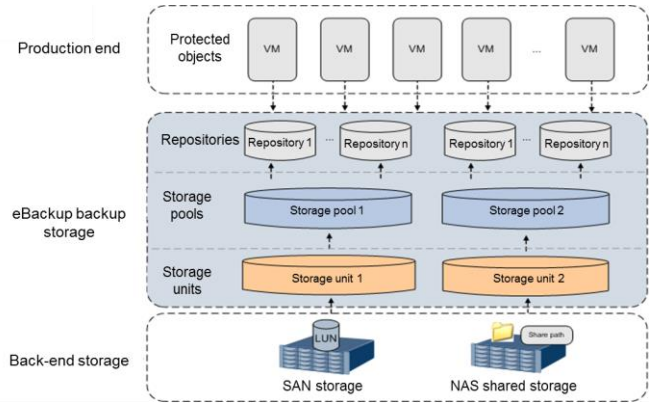

- Uses SAN or NAS as backup media.
- Manages backup storage by layers (storage unit, storage pool, and repository) and allows capacity expansion for business growth by changing the repository capacity (if storage pools have sufficient space), requiring no adjustment of the backup plan and reducing maintenance workload.

eBackup Backup Policy Management

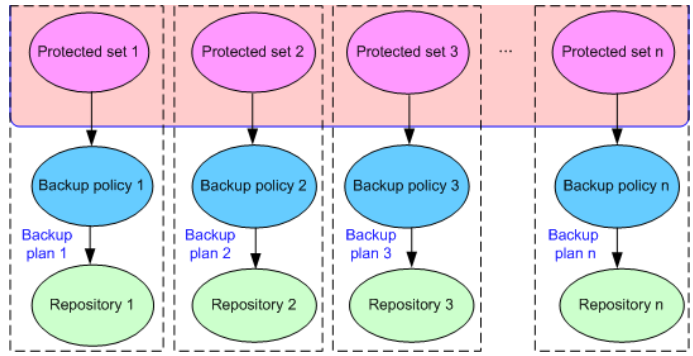

- A backup plan consists of a protected set (backup object), a backup policy (backup type, backup period, and retention period), and a repository (backup target).

Contents

1. Introduction to eBackup

2. **eBackup Installation and Deployment**

   - Installing eBackup
     - Initializing eBackup
     - Running Startup Checks

3. eBackup Backup and Restore

Huawei eBackup Software Installation Flowchart

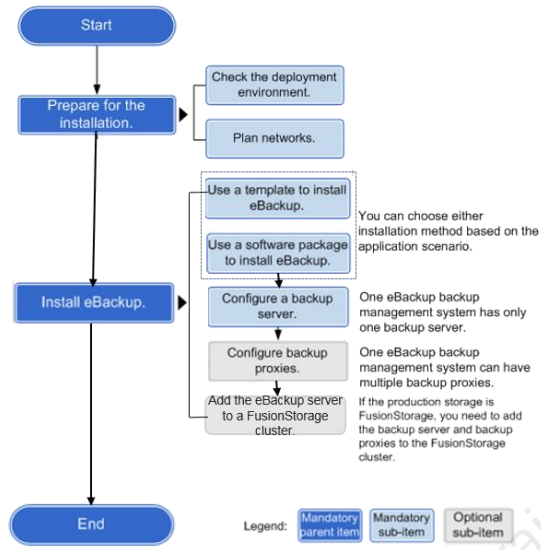

Requirements for Installing eBackup

Item	Specification
Operating system	You can deploy the eBackup software on physical servers or VMs using a software package, or deploy the eBackup software on VMs using a template. • Template-based deployment: FusionSphere V100R005C10 • The supported OS is SUSE Linux Enterprise Server 11 SP3 64-bit, which must be upgraded to 3.0.101-0.47.52-default kernel version. **NOTE:** • The eBackup software supports common Internet file system (CIFS) share storage as backup storage. Before installing the eBackup software, install the CIFS tool (You need to manually install the tool only when you install the software using a software package.) • If you install the eBackup software on VMware VMs using a software package, comply with the following requirements: The vSphere version must be 5.1 or 5.5 and the VM hardware version must be 7 or later. • When a Fiber Channel network is used to provide LAN-free backup, the eBackup software must be deployed on physical servers using the software package.
CPU	4 core x 2.0 GHz
Memory	12 GB
Free disk capacity	120 GB

Network Preparation (1)

Item	Description
Management node and host node at the production end	The management nodes of FusionSphere and VMware are VRM and vCenter Server, respectively, and the host nodes are CNA and ESXi, respectively.
Production storage	All accessible storage resources at the production end, including local disks of hosts or dedicated storage systems. eBackup supports only two types of production storage systems: • Virtual storage • FusionStorage
Backup storage	All physical storage devices that are used for saving backup data and must be shared by each backup server and backup proxy, such as SAN or NAS
Number and roles of network planes	The following network planes must be configured so that the eBackup backup software can work properly and safely: management plane, internal communication plane, and storage plane. The management plane is divided into two sub planes: backup management plane and production management plane.

Network Preparation (2)

Item	Description
Backup management plane	Backup management plane is a communication plane for eBackup to provide external services. Maintenance terminals or third-party systems can log in to the eBackup backup management system using the IP address of this plane to configure and manage backup and restore jobs. Production management plane is a communication plane between the backup server/backup proxy and the management plane at the production end. It is also a plane for transferring backup and restore data in a LAN-base manner.
Internal communication plane	Communication plane between the backup server and backup proxy
Storage plane	• Communication plane between the backup server/backup proxy and the storage plane at the production end • Plane for transferring backup and restore data in a LAN-free manner • Data transmission plane based on file-level restore • Communication plane between the backup server/backup proxy and backup storage

Installing eBackup Using a Template (1)

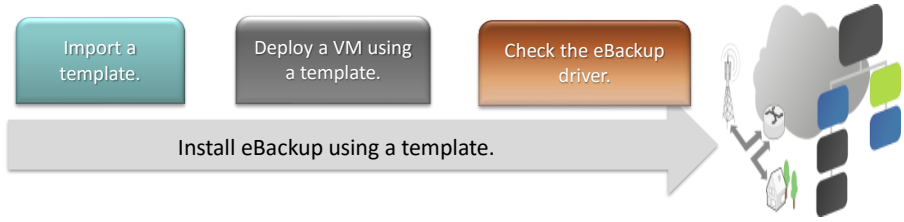

- For details about the installation process, see pre-learning virtualization courses.

Installing eBackup Using a Template (2) – Checking the Version of Tools

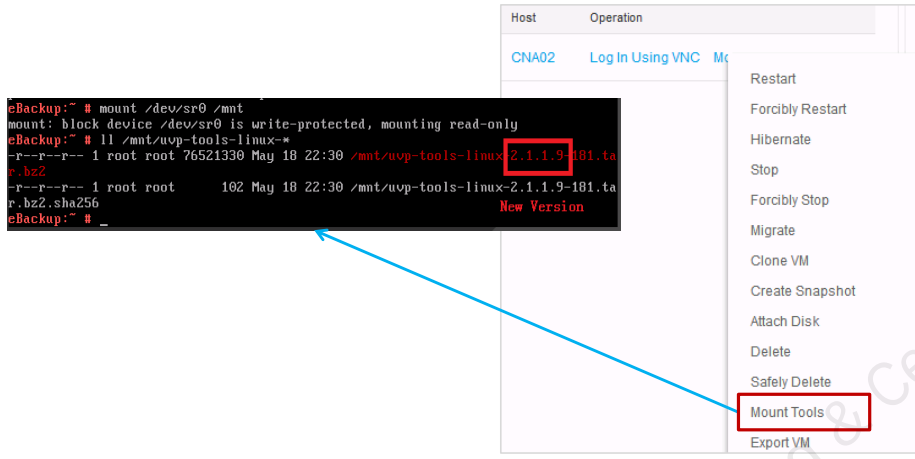

- Check the version of eBackup Tools.
- Check whether FusionCompute is the latest.

Installing eBackup Using a Software Package

- Install a standard SUSE 11 SP3 OS on the server.
- Install a SUSE 11 SP3 kernel patch specified for eBackup.
- Install eBackup.

Installing Standard SUSE 11 SP3 on the Server

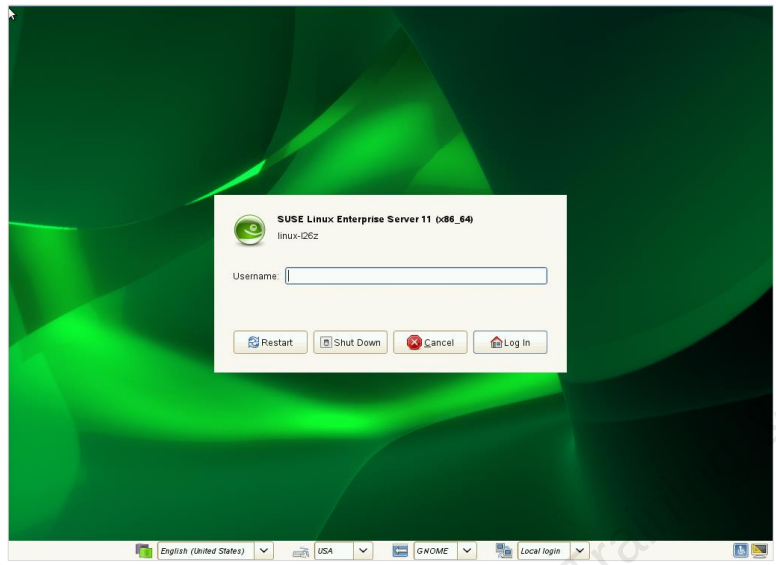

Installing a Kernel Patch

- Check the kernel version. -> Obtain the kernel patch specified for eBackup. -> Install the kernel patch.
- Restart the OS. -> Check the kernel patch.

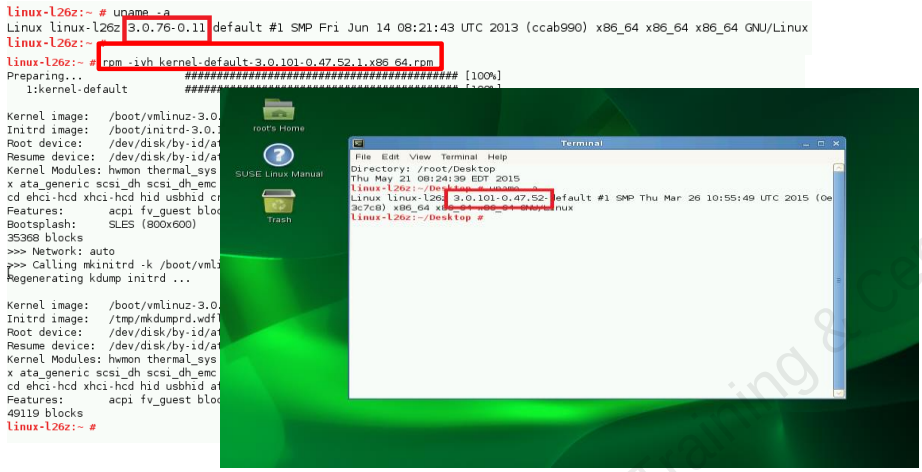

- Run the uname –a command to view the internal version of the operating system.

Installing the eBackup Software

- Obtain the installation package. -> Decompress the package. -> Install the eBackup software.

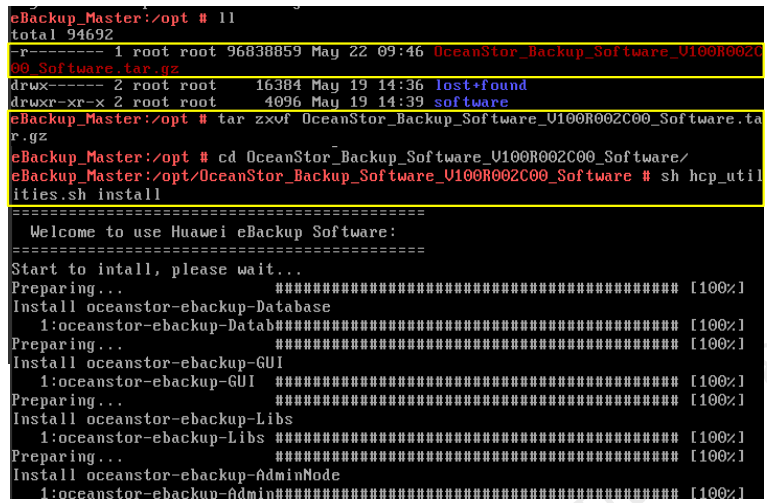

- Run the uname –a command to view the internal version of the operating system.

Contents

1. Introduction to eBackup

2. **eBackup Installation and Deployment**

   - Installing eBackup
   - Initializing eBackup
   - Running Startup Checks

3. eBackup Backup and Restore

eBackup Configuration Items and Process

- **Mandatory Backup Server Configuration Items**
  - Server role
  - Network plane
  - External shared storage (storing database backup files)
  - Time zone
  - Network time protocol (NTP) server
- **Other Backup Server Configuration Items**
  - Adding a backup server to a FusionStorage cluster
- **Mandatory Backup Proxy Configuration Items**
  - Server role
  - Network plane
  - Leader IP (IP address of the internal communication plane of the backup server)
  - Time zone
- **Other Backup Proxy Configuration Items**
  - Adding a backup proxy to a FusionStorage cluster

eBackup Configuration Path and Related Commands

- Run the following command to enter the installation directory:
  - cd /opt/OceanStor_Backup_Software_V100R002C00_Software
- Run the following command to configure eBackup:
  - sh hcp_utilities.sh config

```
eBackup_Master:~ # cd /opt/OceanStor_Backup_Software_V100R002C00_Software/
eBackup_Master:/opt/OceanStor_Backup_Software_V100R002C00_Software # sh hcp_utilities.sh config
```

Configuring the eBackup Backup Server (1)

- Select role Backup server.
- Configure the backup management plane, internal communication plane, production management plane, and storage plane in sequence.

```
eBackup_Master:/opt/OceanStor_Backup_Software_V100R002C00_Software # sh hcp_utilities.sh config
Please select a role for this machine:
 1.Backup server
 2.Backup proxy
1
================================================================================
Note:
In the following steps you will be required to configure four network planes for eBackup.
The definition of each network plane is as follows:
Backup management plane: the communication plane for eBackup to provide external services.
Internal communication plane: the communication plane between backup server and backup proxy.
Production management plane: the communication plane between eBackup and the management plane of the production end.
Storage plane: the communication plane between eBackup and the storage plane of the production end and communication plane between eBackup and backup
storage.
--------------------------------------------------------------------------------
[1] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Backup management
[2] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Internal communication
[3] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Production management
[4] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Storage
Set network adapter for 'Backup management' network plane:
[1] eth0 MAC=28:5E:D5:88:C6:F1 IP=10.158.199.7 MASK=255.255.254.0
[2] eth1 MAC=28:5E:D5:88:C6:F2 IP=172.16.199.7 MASK=255.255.0.0
[3] eth2 MAC=28:5E:D5:88:C6:F3 IP=192.168.0.7 MASK=255.255.0.0
[4] eth3 MAC=28:5E:D5:88:C6:F4 IP=192.168.5.7 MASK=255.255.0.0
Which network adapter from the above list would you like to bind to the 'Backup management' network plane?
1
[1] MAC:28:5E:D5:88:C6:F1       IP=10.158.199.7 plane name=Backup management
[2] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Internal communication
[3] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Production management
[4] MAC:00:00:00:00:00:00       IP=127.0.0.1    plane name=Storage
```

Configuring the eBackup Backup Server (2)

- Configure the NAS or SAN shared storage based on your networking requirements.

```
Configuring shared storage which is used to save backup data of eBackup database.
Please select a type of shared storage:
 1.NFS
 2.SAN
1
Please input the remote path(eg:192.168.1.100:/nfs/shared):
10.183.23.96:/nfs_shared
```

```
Please select a type of shared storage:
 1.NFS
 2.SAN
2
Please input disk path(eg:/dev/disk/by-id/scsi-360022a11000d6ee02e878b2400000009):
/dev/disk/by-id/scsi-360022a11000d6ee32e878b2400000189
```


- Before configuring the shared storage, you need to complete the NAS configuration or mount LUNs of the storage array to eBackup.

Configuring the eBackup Backup Server (3)

- Configure the time zone and NTP.

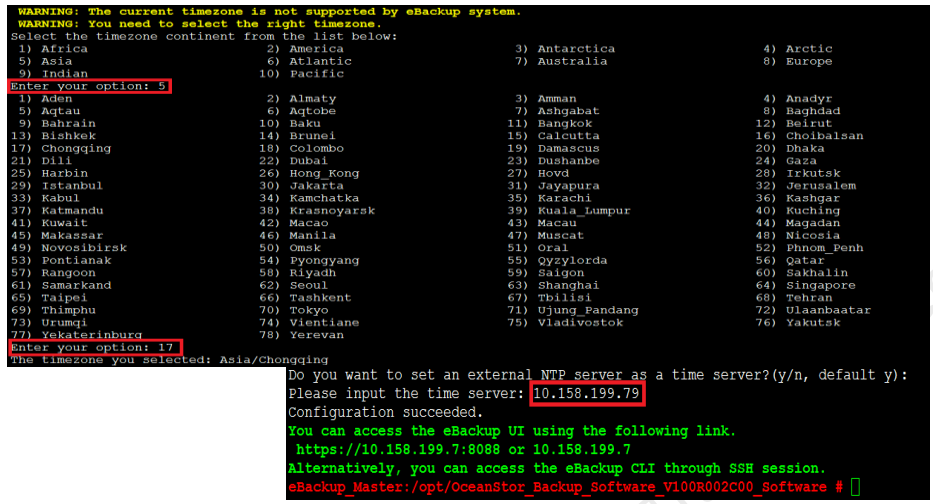

- Before configuring the shared storage, you need to complete the NAS configuration or mount LUNs of the storage array to eBackup.

Configuring the eBackup Backup Proxy (1)

- Select role **Backup proxy**.
  - Configure the backup management plane, internal communication plane, production management plane, and storage plane in sequence.
- Configure the backup proxy. For details, see the backup server configuration method.
  - Configure the leader IP address.

- Before configuring the shared storage, you need to complete the NAS configuration or mount LUNs of the storage array to eBackup.

Configuring the eBackup Backup Proxy (2)

- Configure the time zone.

```
Select the timezone continent from the list below:
 1) Africa                      2) America                     3) Antarctica                  4) Arctic
 5) Asia                        6) Atlantic                    7) Australia                   8) Europe
 9) Indian                     10) Pacific
Enter your option: 5
 1) Aden                        2) Almaty                      3) Amman                       4) Anadyr
 5) Aqtau                       6) Aqtobe                      7) Ashgabat                    8) Baghdad
 9) Bahrain                    10) Baku                       11) Bangkok                    12) Beirut
13) Bishkek                    14) Brunei                     15) Calcutta                   16) Choibalsan
17) Chongqing                  18) Colombo                    19) Damascus                   20) Dhaka
21) Dili                       22) Dubai                      23) Dushanbe                   24) Gaza
25) Harbin                     26) Hong_Kong                  27) Hovd                       28) Irkutsk
29) Istanbul                   30) Jakarta                    31) Jayapura                   32) Jerusalem
33) Kabul                      34) Kamchatka                  35) Karachi                    36) Kashgar
37) Katmandu                   38) Krasnoyarsk                39) Kuala_Lumpur               40) Kuching
41) Kuwait                     42) Macao                      43) Macau                      44) Magadan
45) Makassar                   46) Manila                     47) Muscat                     48) Nicosia
49) Novosibirsk                50) Omsk                       51) Oral                       52) Phnom_Penh
53) Pontianak                  54) Pyongyang                  55) Qyzylorda                  56) Qatar
57) Rangoon                    58) Riyadh                     59) Saigon                     60) Sakhalin
61) Samarkand                  62) Seoul                      63) Shanghai                   64) Singapore
65) Taipei                     66) Tashkent                   67) Tbilisi                    68) Tehran
69) Thimphu                    70) Tokyo                      71) Ujung_Pandang              72) Ulaanbaatar
73) Urumqi                     74) Vientiane                  75) Vladivostok                76) Yakutsk
77) Yekaterinburg              78) Yerevan
Enter your option: 17
The timezone you selected: Asia/Chongqing
Configuration succeeded.
You can access the eBackup UI using the following link.
 https://backup server's backup management plane:8088 or backup server's backup management plane
Alternatively, you can access the eBackup CLI through SSH session.
eBackupProxy:/opt/OceanStor_Backup_Software_V100R002C00_Software #
```


- Before configuring the shared storage, you need to complete the NAS configuration or mount LUNs of the storage array to eBackup.

Contents

1. Introduction to eBackup
2. **eBackup Installation and Deployment**
   - Installing eBackup
   - Initializing eBackup
   - Running Startup Checks
3. eBackup Backup and Restore

eBackup Initial Startup Check and Login Process

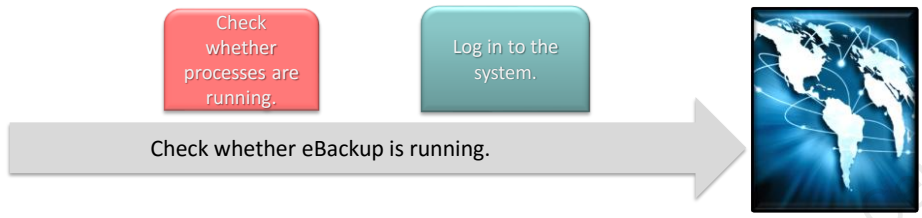

- Online backup: hitless backup without suspending services on VMs.
- Snapshot-based backup: backs up data using snapshots of the file system or main storage. No backup agents are installed on VMs, minimizing the impact on VMs during the backup.

eBackup Initial Startup Check

- Check whether backup server processes are running.
- Check whether backup proxy processes are running.

```
eBackup:/opt/OceanStor_Backup_Software_V100R002C00_Software # service hcp status
Checking for Huawei eBackup Service
gaussdb is running
AdminNode is running
BackupNode is running            eBackup Master
hcplogrotate is running
apache/iBase is running
upgradeDaemon is running
HCPProcessMonitor is running
                                                                                running
```

```
eBackupProxy:/opt/OceanStor_Backup_Software_V100R002C00_Software # service hcp status
Checking for Huawei eBackup Service
BackupNode is running
hcplogrotate is running              eBackup Proxy
upgradeDaemon is running
HCPProcessMonitor is running
                                                                                running
```

eBackup Login

Contents

1. Introduction to eBackup
2. eBackup Installation and Deployment
3. **eBackup Backup and Restore**
   - Backup Principle
   - Backup Application & Configuration
   - Restore Principle
   - Restore Application & Configuration

Backup Principle

- Backup is a process of replicating data from the production-end storage to the back-end storage with which the backup system is associated. If the production system fails or data in the production system is lost, the backup data can be used to restore the system or data. The eBackup backup management system can back up FusionSphere and VMware VM data to its associated back-end storage.

Basic Concepts

Term	Description
Snapshot	A snapshot is a point-in-time copy of source data. It can be quickly generated and occupies a small storage footprint. Creating a snapshot is like taking a picture. The picture can record the status of the object for which you took a picture at a certain point in time. Snapshots can be used to back up and restore protected objects.
CBT	CBT is an incremental backup technology for VMs. It uses bitmaps to track and record the data blocks changed since last backup. When you back up data, CBT enables you to identify the changed data blocks and back up the incremental data. The production end achieves the backup function using CBT.
Full backup	• For virtual disks, all data on the virtual disks is backed up. • For VMs, all disk data and VM configuration information are backed up. No matter when data was changed or backed up last time, all data can be backed up. Full backup provides the most complete backup protection; however, full backup takes a long time and occupies large space.
Incremental backup	• For disks, the data changed since last full or incremental backup for disks is backed up. • For VMs, the disk data and VM configuration information changed since last full or incremental backup for VMs are backed up. The amount of incremental backup data is small and the backup time is short.

Deduplication

- Deduplication is a data compression technology that searches for duplicate data, saves only one copy of the data, and uses pointers that point to the unique copy to replace the duplicate copies. The eBackup backup management system uses deduplication to minimize the occupation of back-end storage space and address the increasing demand for storage capacity. The following figure shows the deduplication working process. The system checks for duplicate data blocks and uses deduplication to delete redundant data blocks (The data in the dotted box implies the duplicate data).

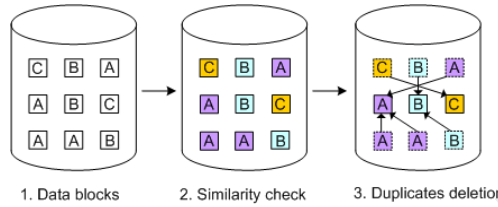

Backup Modes and Application Scenarios

Backup Mode	Application Scenario	Description
CBT	FusionSphere (storage type: virtual storage) **NOTE: CBT is automatically enabled during backup.**	Snapshots are not reserved at the production end, so they do not occupy too much storage space. Using the CBT backup mode, the eBackup backup management system obtains changed data in a fast manner.
Snapshot comparison	FusionSphere (storage type: FusionStorage)	Snapshots are reserved at the production end, so they occupy large storage space. Using the snapshot backup mode, the eBackup backup management system obtains changed data in a slow manner.

Backup Principle

Backup Mode	Backup Principle (Used for Disk Backup)
CBT	• Initial backup (full backup) The eBackup backup management system starts the initial full backup job for disks to back up all data on the disks. • Subsequent backup (incremental backup) The eBackup backup management system starts an incremental backup job for disks. By default, the system uses CBT to obtain the data changed since the full backup and incremental backup from the disks and backs up the changed data.
Snapshot comparison	• Initial backup (full backup) The eBackup backup management system starts the initial full backup job for disks to back up all data on the disks. • Subsequent backup (incremental backup) The eBackup backup management system starts an incremental backup job for disks. By default, the system compares the two snapshots created for disks at the production end to obtain the data changed since the full and incremental backup from the disks and backs up the changed data.

Contents

1. Introduction to eBackup
2. eBackup Installation and Deployment
3. **eBackup Backup and Restore**
   - Backup Principle
   - Backup Application & Configuration
   - Restore Principle
   - Restore Application & Configuration

Backup Application & Configuration

- eBackup backup flowchart

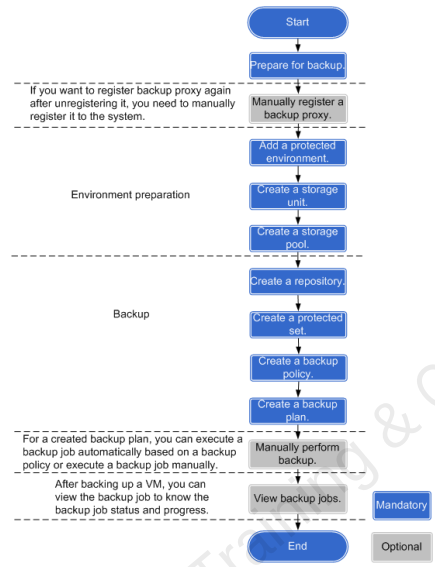

Preparing for Backup (1)

- Before a backup job, prepare the data required for service configuration so that you can find related data quickly and accurately.

Data	Description	Value or Default Value
A1: IP address of the backup server backup management plane	IP address of the backup management plane used to log in to the eBackup backup management system where you can configure and manage backup or restore services	Obtain the IP address of the backup server backup management plane from the eBackup backup management system administrator.
A2: user name and password of the eBackup backup management system super administrator	The eBackup backup management system provides a default super administrator. The super administrator has all operation rights and can manage all resources.	The user name of the super administrator is **admin** and is unchangeable. The default password is **Huawei@CLOUD8**!.
B1: vCenter/ESXi IP address	When adding a VMware Protected Environment, if multiple hosts need to be managed, enter the vCenter Server IP address.	Obtain the vCenter/ESXi IP address from the VMware administrator.
B2: user name and password of VMware vSphere Web Client	User name and password of VMware vSphere Web Client	Obtain the user name and password of VMware vSphere Web Client from the VMware administrator.

Preparing for Backup (2)

Data	Description	Value or Default Value
C1: VRM IP address	• If FusionCompute is deployed in single node mode, the VRM IP address is the management IP address of FusionCompute. • If FusionCompute is deployed in active/standby mode, the VRM IP address is the floating IP address of FusionCompute.	Obtain the VRM IP address from the FusionCompute administrator.
C2: name and password of the user for interconnecting FusionCompute	Name and password of the user for interconnecting FusionCompute. With C2, the eBackup backup management system can access FusionCompute, and then performs backup and restore jobs on FusionCompute.	Obtain the user name and password from the FusionCompute administrator.
D1: path	The eBackup backup management system supports NAS and SAN storage. If the storage type is NAS, enter the path for accessing the NAS shared storage when creating the storage units.	Obtain the path from the eBackup backup management system administrator.

- When using NFS share storage as the back-end storage, the format of path is IP(or domain name):/share path. IP and domain name indicate the service IP address and domain name of the NAS storage system, respectively. The share path consists of 1 to 255 characters and can contain letters, digits, and special characters, including .#%-=@{_}~/:. The name cannot end with a slash (/). The following provides an example: /192.168.10.10:/home/storage.
- When using CIFS share storage as the back-end storage, the format of path is //IP(or domain name)/share name. IP and domain name indicate the service IP address and domain name of the NAS storage system, respectively. The share path consists of 1 to 255 characters and can contain letters, digits, and special characters, including .#%-=@{_}~/:. The name cannot end with a slash (/). The following provides an example: //192.168.10.10/home/storage.
- If the storage type is SAN, enter the path to access the space mapped from the SAN storage when creating the storage units.

Adding a Protected Environment

- On the navigation bar, choose FusionSphere.
- Click in the Protected Environment area.

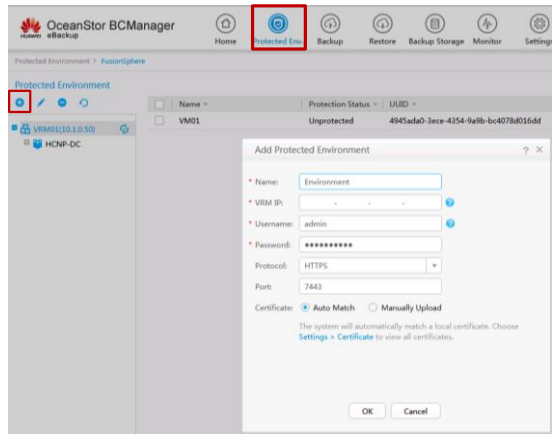

Note: A maximum of four protected environments can be added to one eBackup system.

Creating a Storage Unit

- On the navigation bar, choose Storage Unit and click Create.

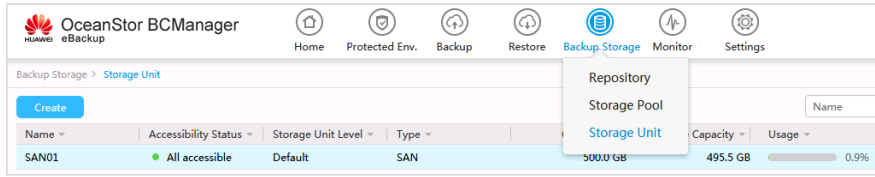

- Enter the basic information of a storage unit.

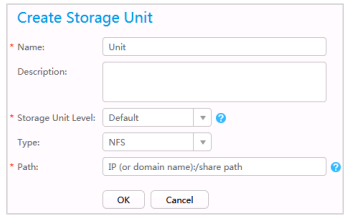

Note: A maximum of 350 storage units of different types can be created in one eBackup system.

A maximum of 20 VIMS storage units can be created for a file system.

Creating a Storage Pool

- On the navigation bar, choose Storage Pool and click Create.

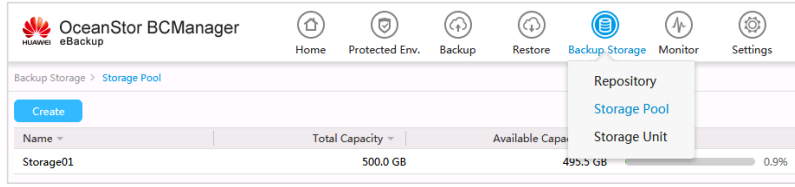

- Enter the basic information of a storage pool.

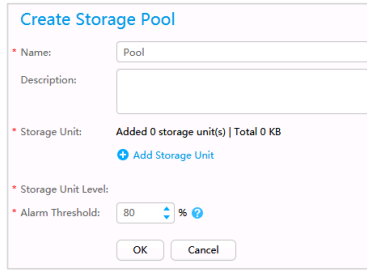

Note: A maximum of 350 storage pools can be created in one eBackup system.

Creating a Repository

- On the navigation bar, choose > **Repository** and click **Create**.

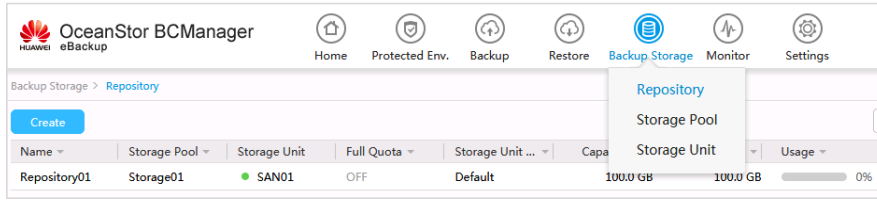

- Enter the basic information of a repository.

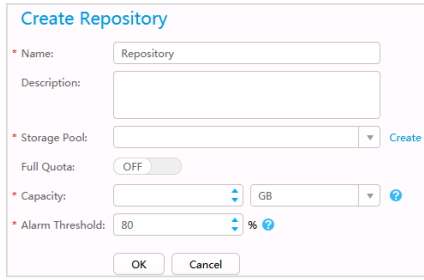

Note: A maximum of 1500 repositories can be created in one eBackup system.

Creating a Protected Set

- On the navigation bar, choose **Protected Set** and **click** Create.

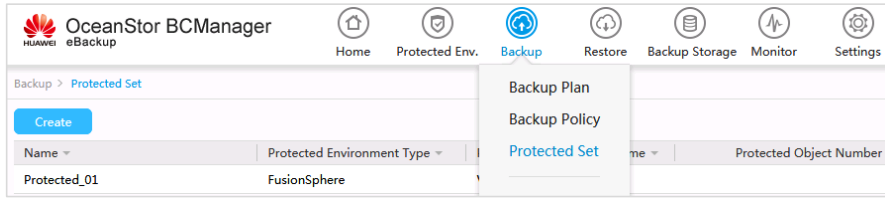

- Enter the basic information of a protected set.

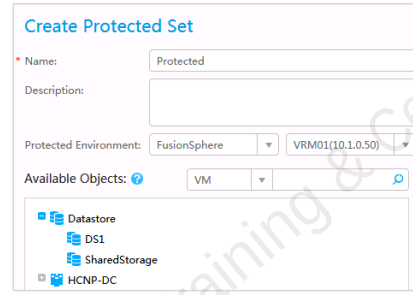

Note: A maximum of 10000 protected sets can be created in one eBackup system.

Creating a Backup Policy

- On the navigation bar, choose **Backup Policy** and click **Create**.

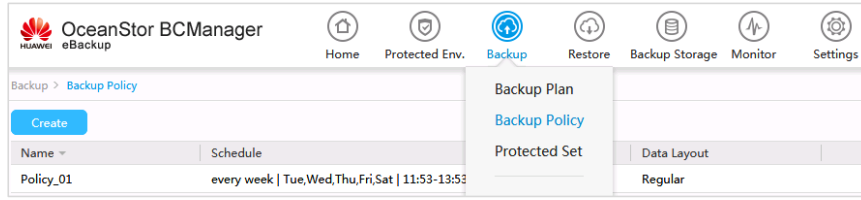

- Enter the basic information of a protected set.

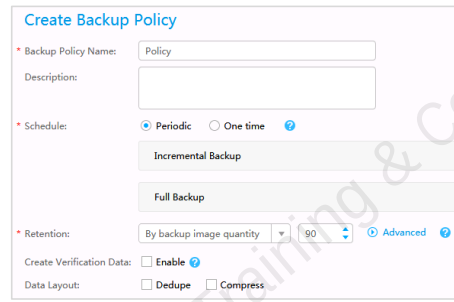

Note: A maximum of 10000 backup policies can be created in one eBackup system.

Creating a Backup Plan

- A backup plan consists of a repository, protected set, and backup policy. After a backup plan is created, you can perform backup jobs for the backup plan.
- On the navigation bar, choose Backup Plan or choose Create Backup Plan in the backup management system, and click Create.

- Enter information on the General Info, Protected Set, Backup Policy, and Repository pages in sequence.

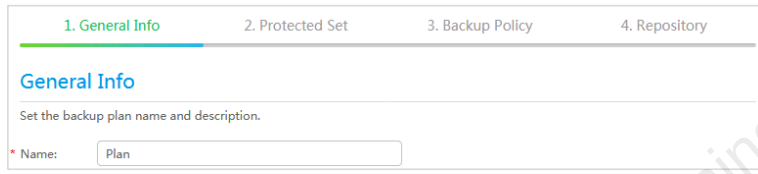

Note: A maximum of 10000 backup plans can be created in one eBackup system.

(Optional) Performing a Backup Job Manually

- For a created backup plan, you can back up data automatically based on a backup policy or perform a backup job manually.
- On the navigation bar, choose > Backup Plan.
- You can manually back up data using any of the following methods:

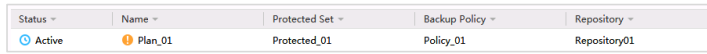

  - Place the cursor on the desired backup plan, and then click in the button area on the right.
  - Click the desired backup plan, and then click in the Preview area on the right.
  - If backup jobs have been performed based on the backup plan, the system automatically performs incremental backup after you click .
  - If no backup jobs have been performed based on the backup plan, the system automatically performs full backup after you click .
  - You can click the drop-down arrow beside and select Run Full Backup Job. Then the system will perform a full backup as required.
- **The Info dialog box is displayed.**
- **Click Yes**.

(Optional) Viewing a Backup Job

- After backing up a VM, you can view the backup job to know the backup job status and progress.
- On the navigation bar, choose **Job**.
- Optional: In the search field in the upper right corner, set the search condition; (set **Type** to **Backup**), or click **Advanced**, set conditions, and click **OK** to search for desired backup jobs.
- View backup jobs.

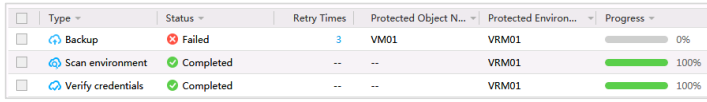

- Click a backup job. In the preview area on the right, view its details.
- Optional: Manually stop a backup job in either of the following ways when an error occurs due to misoperation or when you do not need to perform the backup job:
  - Click the backup job and click in the preview area on the right.
  - Place the cursor on the backup job and click the button area on the right.

Contents

1. Introduction to eBackup
2. eBackup Installation and Deployment
3. **eBackup Backup and Restore**
   - Backup Principle
   - Backup Application & Configuration
   - Restore Principle
   - Restore Application & Configuration

Restore Function

- When the VM data is damaged or lost at the production end and needs to be restored, the eBackup backup management system restores the backup data on the back-end storage to the production end using backup images generated at a specific point in time.

Basic Concepts

Item	Description
Backup image	A backup image is the backup data generated by a backup plan in one backup job according to a backup policy. Each backup image is marked with the creation time. Users can select a backup image to restore data according to a timestamp.
CBT	CBT is an incremental backup technology for VMs. It uses bitmaps to track and record the data blocks changed since last backup. The production end achieves the backup function using CBT. When VM data needs to be restored, you can use CBT to compare data changes between the restore execution point in time and backup point in time and restore only the changed data to achieve incremental restore.
Full restore	• For disks, all disk backup data to which a backup image corresponds are restored. • For VMs, all disk backup data and VM configuration information to which a backup image corresponds are restored. A full restore job takes a long time and restores a large volume of data.
Incremental restore	• For disks, disk data changed between the backup point in time and restore execution point in time is restored. • For VMs, disk data changed between the backup point in time and restore execution point in time and all VM configuration information is restored. An incremental restore job takes a short time and restores a small volume of data.

Restore Principle

Restore Mode	Restore Principle
Full restore	When the eBackup backup management system starts a restore job for disk M, if the original disk on the target VM exists with capacity and location unchanged, obtain all the backup data from the backup image and write it onto disk M; if the original disk on the target VM does not exist, create new disk M1 on the production end based on the metadata of disk M in the backup image, obtain all the backup data from the backup image and write it onto disk M1, and mount disk M1 to the target VM.
Incremental restore	The eBackup backup management system starts a restore job for disk M and uses CBT to obtain the data changed between the restore execution point in time of disk M and the backup point in time of the backup image. After receiving the changed data from the backup image, the system writes it to disk M at the production end.

Contents

1. Introduction to eBackup
2. eBackup Installation and Deployment
3. **eBackup Backup and Restore**
   - Backup Principle
   - Backup Application & Configuration
   - Restore Principle
   - Restore Application & Configuration

Restore Application & Configuration

- The eBackup backup management system supports VM restore, batch VM restore, VM disk restore, batch VM disk restore, and file-level restore in FusionSphere.

Restoring a VM

- You can restore a VM that has been backed up to a specified location.
- On the navigation bar, choose FusionSphere.
- In the Backed Up Environment area, click the protected environment where the VM that you want to restore resides.
- Select the VM that you want to restore by performing either of the following operations.
  - Click the VM in the list.
  - Find the VM by querying it on the upper right of the list. Then click the VM.
- In the preview area on the right, place the pointer on the backup image required for the restore and click .
- Restore the VM to a new location.
  - Click Restore VM to New Location.
  - Set the basic information about Restore VM to New Location.
  - (Optional) Select Start VM after restore.
  - Click OK.

Restoring VMs in Batches

- You can restore VMs that have been backed up to specified locations in batches.
- On the navigation bar, choose FusionSphere.
- In the Backed Up Environment area, click the protected environment where the VMs that you want to restore reside.
- Select the VMs that you want to restore by performing either of the following operations.
  - Click the VMs in the list.
  - Find the VMs by querying them on the upper right of the list. Then click the VMs.
- Click Batch Restore VM to New Location.
- Set the basic information about Batch Restore VM to New Location.
- (Optional) Select Start VM after restore.
- Click OK.

Restoring a Disk on a VM

- You can restore a disk on a VM that has been backed up to the original VM or a specified VM.
- On the navigation bar, choose **FusionSphere**.
- In **the Backed Up Environment** area, click the protected environment where the VM disk that you want to restore resides.
- Select the VM disk that you want to restore by performing either of the following operations.
  - Click the VM where the VM disk resides in the list.
  - Find the VM where the VM disk resides by querying the VM on the upper right of the list. Then click the VM.
- Select the backup image required for restoration.
  - In the preview area on the right, place the pointer on the backup image required for the VM restoration and click .
  - In the preview area on the right, place the pointer on the backup image required for the VM disk restoration and click .
- You can restore a VM disk to the original VM or a specific VM.

Restoring a VM Disk to the Original VM

- Click **Restore VM Disk to Original VM**.
- Select **the VM disk** that you want to restore.
- Select **Start VM after restore**.
- Click **OK**.

Restoring a VM Disk to a New VM

- Click **Restore VM to New Location**.
- Set the basic information about **Restore VM to New Location**.
- Select **Start VM after restore**.
- Click **OK**.

Restoring Disks on Multiple VMs in Batches

- You can restore disks on multiple VMs that have been backed up to the original VMs in batches.
- On the navigation bar, choose **FusionSphere**.
- In the **Backed Up Environment** area, click the protected environment where the VM disk that you want to restore resides.
- Select the VM disks that you want to restore by performing either of the following operations.
  - Click the VMs where the VM disks reside in the list.
  - Find the VMs by querying them on the upper right of the list.
- **Click Batch Restore VM Disk**.
- Select the VM disks that you want to restore in batches.
- (Optional) Select the datastores that you want to restore the VM disks to.
- (Optional) **Select Start VM after restore**.
- Click **OK**.

Restoring a File in a Disk (1)

- You can restore a file in a disk that has been backed up. The restore is achieved through the file level fine-granularity restore function. The disk where the file resides is mounted to the target VM (the original VM or a new VM). Then, you can log in to the VM to copy the file to a specified location.
- On the navigation bar, choose > **FusionSphere**.
- In the **Backed Up Environment** a, click the protected environment where the file resides.
- Select the VM where the file resides by performing either of the following operations.
  - Click the VM where the file resides in the list.
  - Find the VM where the file resides by querying the VM on the upper right of the list. Then click the VM.
  - You can restore one or all disks in the VM to the target VM. In the preview area on the right, place the pointer on the backup image required for the restore and click .
  - **NOTE**: The total number of disks on the target VM and d to be restored cannot exceed 11.
  - Select **Restore to Original VM or Restore to New VM**, and then click **OK**.

Restoring a File in a Disk (2)

- On the navigation bar, choose **Job**.
- View the **Status** and **Progress** of the file level restore job. When **Status** is In **progress** and **Progress** is 100%, you can go to the next step.
- Restore files.
- After file replication completes, you can perform either of the following operations to abort the job.
  - Click the job. In the preview area on the right, click
  - Place the pointer on the job. In the operation button area on the right, click .

Viewing a Restore Job

- In the navigation tree, choose · **Job**.
- (Optional) On the search bar in the upper right corner, set the search criteria (set **Type** to **Restore machine**, **Restore VM disk**, or **File-level restore**). Alternatively, click **Advanced** to set search criteria and click **OK** to initiate a search.
- View the information about restore jobs.

- Click the job whose information you want to view. In the preview area on the right, view the details about the job.
- (Optional) You can manually stop the job by performing either of the following operations if you have an incorrect operation or do not need the job.
  - Click the job. In the preview area on the right, click .
  - Place the pointer on the job. In the operation button area on the right, click .

Quiz

1. (Fill in the blanks) Backup modes include ( ) and ( ).
2. (Fill in the blanks) Restore modes include ( ) and ( ).
3. (True or false) A backup policy can be executed manually.
4. (True or false) A maximum of 30 VIMS storage units can be created for a file system.
5. (True or false) Full backup provides the most complete backup protection, but takes a long time, and occupies small space.
6. (True or false) The eBackup backup management system does not support file-level restore in batches.

Summary

- The eBackup backup management system supports VM restore, batch VM restore, VM disk restore, batch VM disk restore, and file-level restore in FusionSphere.
- The eBackup backup management system supports full restore and incremental restore.
- A full restore job takes a long time and restores a large volume of data.
- An incremental restore job takes a short time and restores a small volume of data.

Thank You

www.huawei.com

Principle and Deployment of Huawei Cloud Data Center Agile Controller-DCN

www.huawei.com

Huawei Training & Certification

Foreword

- This section describes basic concepts of data center SDN, and architecture, installation, and deployment of Huawei cloud data center SDN controller, Agile Controller-DCN.

Objectives

- Upon completion of this course, you will be able to:
  - Understand basic concepts of data center SDN
  - Learn architecture of the Agile Controller-DCN
  - Learn how to install and deploy the Agile Controller-DCN

Contents

1. **Agile Controller-DCN Overview**
2. Agile Controller-DCN Architecture
3. Agile Controller-DCN Networking Mode
4. Agile Controller-DCN Deployment Process

Basic Concepts of Cloud Data Center

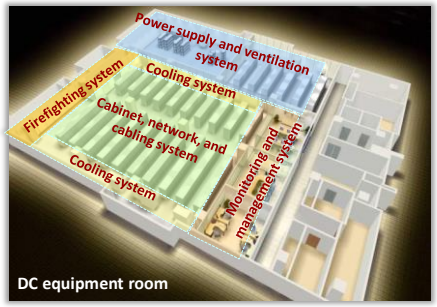

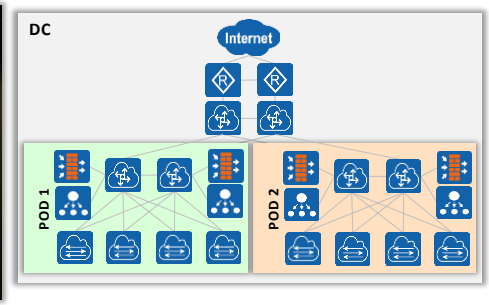

- Data center (DC): DC is a physical concept, implementing centralized data processing, storage, transmission, exchange, and management in a physical space (such as an equipment room). Key devices in a DC include servers, network devices, and storage devices. Key auxiliary facilities in a DC include the power supply, cooling, firefighting, and monitoring infrastructures.
- Point of delivery (POD): To facilitate DC resource pooling, a DC can be divided into one or multiple physical partitions, and each physical partition is a POD. POD is also a physical concept and is the basic deployment unit of a DC. A physical device can belong to only one POD.
- Availability Zone (AZ) is a logical concept and represents a logical fault isolation zone. For example, hosts in the zone share one set of power and network devices. If the devices fail, all resources in the zone become unavailable. Mapping between AZ and DC can be flexibly planned based on the site requirements. For example, in a large-scale public cloud, an AZ can contain multiple DCs. In a small- and medium-scale private cloud, a DC can contain multiple independent AZs. A DC can also be mapped only to one AZ. In this case, they are equivalent.

Data Center Network (DCN)

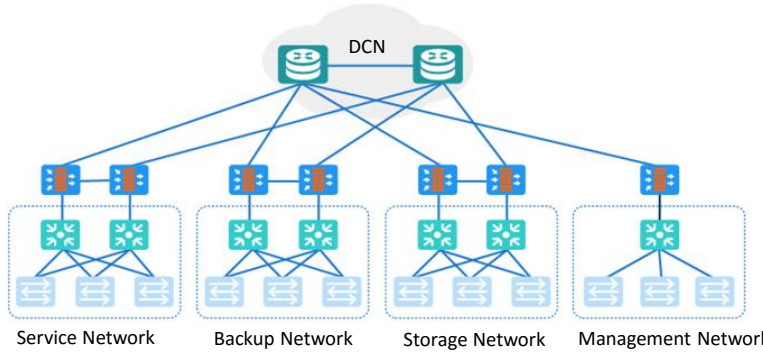

- DCNs are the infrastructure carrying data center services and consists of computer systems, storage systems, communication systems, corresponding communication connections, and management and monitoring systems. Enterprises or organizations can store, manage, and advertise information in data centers.
- Multiple DCNs can connect branches of enterprises or organizations in different areas. In addition, DCNs connect to the Internet, so that enterprises or organizations can access the Internet.

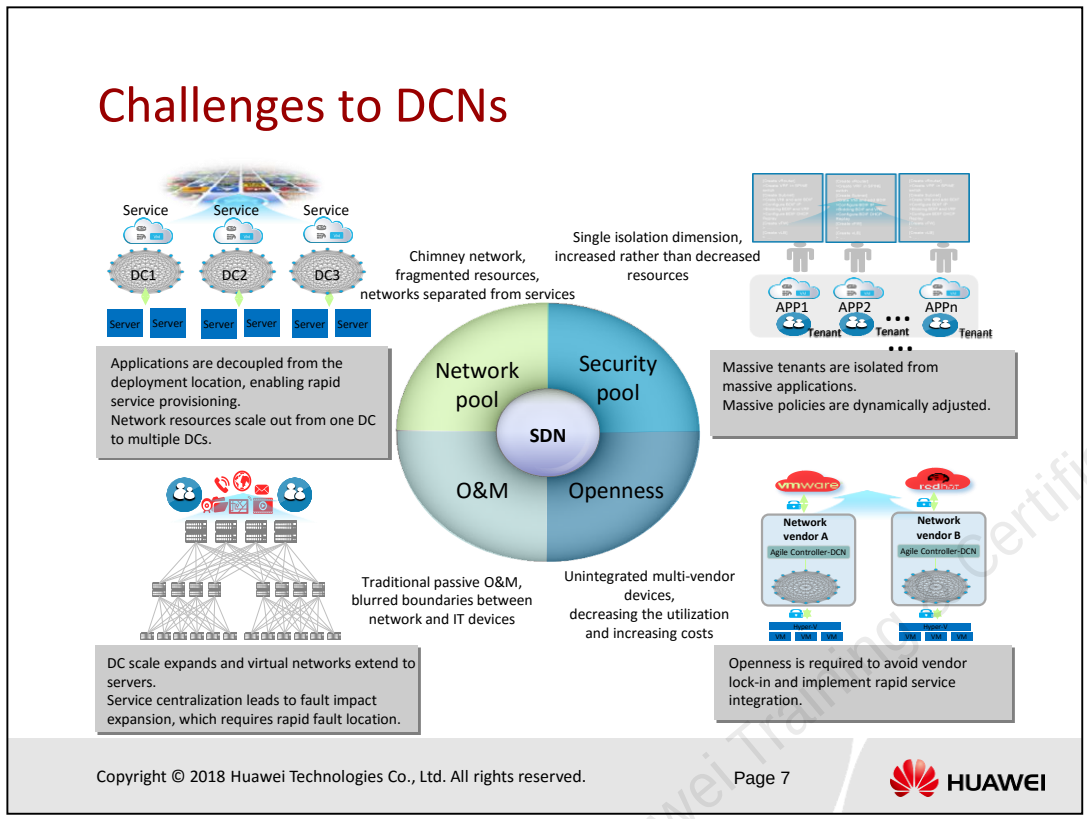

- In a data center, data exchange is centralized and east-west traffic increases, posing the following requirements on the DCN: large scale, high scalability, enhanced robustness, low configuration costs, high bandwidths between servers, efficient network protocols, flexible topology and link capacity control, energy conservation, traffic isolation between services, and low costs. The traditional three-layer architecture cannot meet the requirements. A flattened, virtualized, programmable, and definable network architecture becomes a new trend of DCN development.
- To meet the preceding requirements, DCNs begin to apply the VXLAN and SDN technologies. The technologies allow cloud platforms to implement adaptation and association between networks.

SDN Advantages

- Make networks agile from the architecture.
  - The Agile Controller-DCN centrally manages the underlay network.
  - Network resources are associated with computing resources.
  - Networks change based on demands. The Agile Controller-DCN dynamically delivers network policies.
  - The large layer 2 network accesses computing resources in a unified mode.
- Build a simple, efficient, and open cloud DCN to support long-term evolution of enterprise cloud services.

SDN

- An SDN has three network layers:
  - SDN application
  - Agile Controller-DCN
  - Physical network

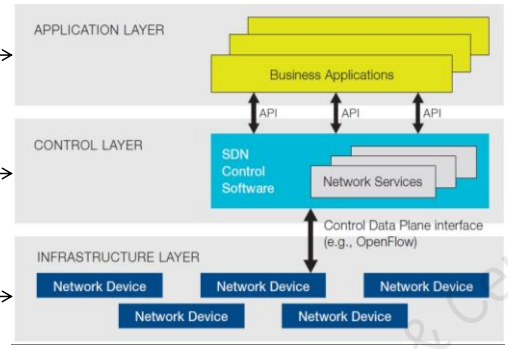

Agile Controller-DCN

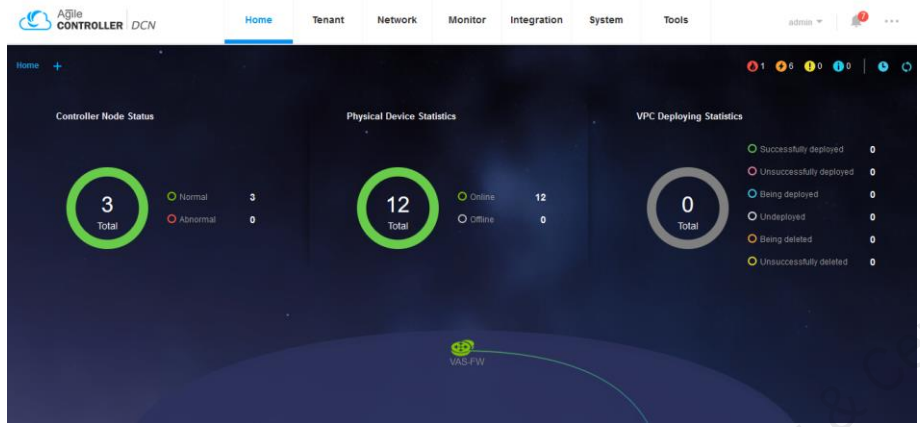

- Huawei Agile Controller-DCN constructs the overlay network to implement cloud-network integration.

Cloud Data Center Solution

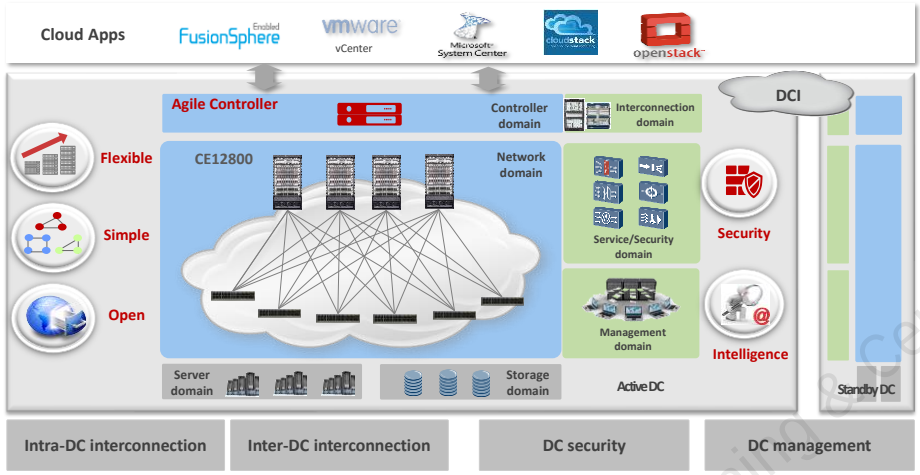

HUAWEI

- Cloud platform: Implements centralized management of service resources and service ingresses, and collaborative scheduling of computing or storage resources and network resources.
- The Agile Controller-DCN: Implements centralized management of physical and virtual networks, and standard interconnection with cloud platforms through northbound interfaces. It also provides an independent GUI for service policy orchestration, network modeling and automatic network service deployment. It also manages and controls physical and virtual network devices through standard southbound interfaces including the OpenFlow, OVSDB, and NETCONF interfaces.
- Network: Fabrics are based on the network or hybrid overlay VXLAN and used for high-speed intra-DC data exchange.
- Server: Provides various computing resources, including VMs, physical machines, and bare metal servers. VM management platforms are supported.

Cloud Data Center SDN Architecture

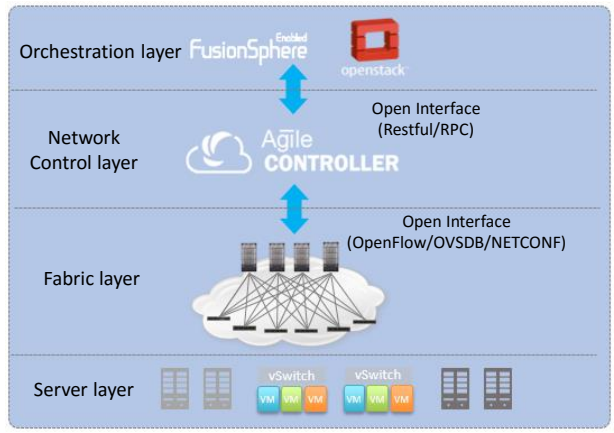

- Cloud data center SDN can be divided into the service presentation (collaboration) layer, network control layer, network service layer, and computing access layer.

- The service presentation/orchestration layer implements network orchestration, service provisioning, and collaborative scheduling of computing or storage resources and network resources. Collaborative service scheduling and provisioning can also be implemented through interconnection with the cloud platform, customer portal, and the Agile Controller-DCN GUI.
- The network control layer implements centralized control of physical and virtual networks. It supports interconnection with cloud platforms through northbound interfaces, direct service orchestration, network modeling, and automatic network service deployment. It also supports management and control of physical and virtual network devices through southbound interfaces including the OpenFlow, OVSDB, and NETCONF interfaces.
- The network service layer supports fabrics based on the network, hybrid, or host overlay VXLAN, network device stacking, M-LAG networking, active-active VXLAN gateways, and Layer 4 to Layer 7 service automation.
- The computing access layer is responsible for computing and storage resource pools. It also supports mainstream virtualization platforms and automatic provisioning of bare metal servers.

Network Virtualization and Cloud-Network Integration

- Network virtualization and cloud-network integration are scenarios of cloud data center SDN.
- In the network virtualization scenario, no cloud platform is deployed. However, in the cloud-network integration scenario, a cloud platform is deployed.
- In the cloud-network integration scenario, the Agile Controller-DCN, network, the virtual machine management (VMM) platform, and the cloud platform are deployed. The Agile Controller-DCN and the VMM platform respectively interconnect with the cloud platform. The cloud platform provides the administrator with a unified computing and network management page for service provisioning and management.
- The network virtualization scenario is divided into hosting and computing scenarios.
  - Network virtualization - hosting: The Agile Controller-DCN and network are deployed. No cloud platform or interconnection with the VMM platform is supported. The Agile Controller-DCN provides the network administrator with an independent management GUI for unified network management.
  - Network virtualization - computing: The Agile Controller-DCN, network, and the VMM platform are deployed. No cloud platform is deployed. The Agile Controller-DCN provides the network administrator with an independent management GUI for unified management of physical and virtual networks. In addition, the network system associates with the computing system which is managed by the computing administrator.

Contents

1. Agile Controller-DCN Overview
2. **Agile Controller-DCN Architecture**
3. Agile Controller-DCN Networking Mode
4. Agile Controller-DCN Deployment Process

VM Deployment View

- Components:
  - Physical server: RH1288V3, E9000
  - Layer I: FusionCompute 6.0/VMware (including Host OS)
  - GuestOS: SUSE11SP3/SUSE12SP2
  - Controller software: Agile Controller-DCN V300R001C00
- Typical reliability deployment:
- 3 physical machines/computing nodes (standard)
- 5 service VMs
  - 3 service processing VMs (cluster, expansible)
  - 2 network optimization VMs (in active/standby mode, optional)

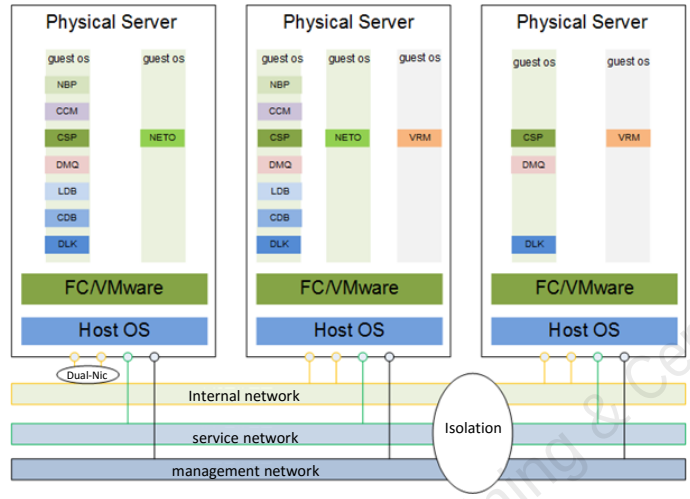

PM Deployment View

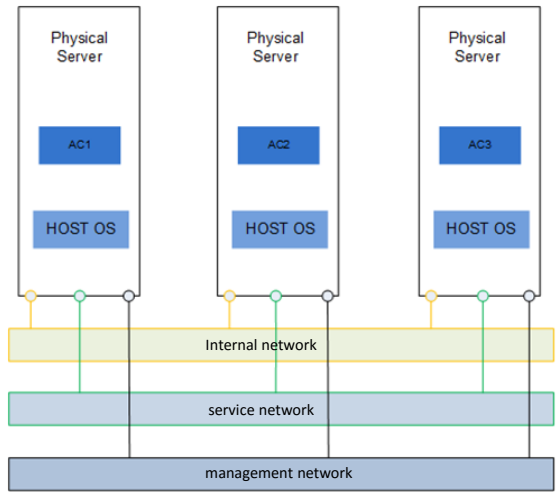

- Components:
- Physical server: RH1288V3
- Host OS: SUSE/RedHat/CentOS
- Controller software: Agile Controller-DCN V300R002C00

- Typical reliability deployment:
- 3 physical machines/computing nodes (basic configuration)
- 2*CPU @2.4 GHz, 8-core
- DDR4 128 GB
- HDD 4×600 GB RAID10
- 4*GE + 4*GE

Component Description

Type	Function	Deployment Mode
(NorthBound Proxy (NBP) Northbound proxy (Ngnix)	Uses the floating IP address to provide a unified northbound ingress, and load balances northbound requests to different service cluster nodes.	In a cluster, only two nodes provide the northbound proxy function. To improve the system reliability, the component is deployed on two cluster nodes in active/standby mode.
Controller Cluster Management (CCM) Cluster management (Puppet)	Uses the floating IP address to manage all Agile Controller-DCN cluster nodes and implements configuration and maintenance inside the cluster, such as cluster configuration, startup, stop, and process protection.	In a cluster, only two nodes provide the cluster management function. The component is deployed on the same node in active/standby mode on which northbound proxy is deployed.
Controller Service Processing (CSP) Service processing (Karaf)	Processes the Agile Controller-DCN services, such as NE management, topology management, and service provisioning, and sends the processing result to the southbound forwarding device.	At least three nodes are required in a cluster. If the cluster performance reaches the bottleneck, you can add nodes to improve the cluster performance.
Distributed Database (DDB) Distributed database GuassDB (PostgreSQL)	Saves all service configurations for real-time service query and recovery after restart.	The component supports multiple groups of active/standby node deployment. By default, there are two groups of instances: local database (LDB) and common database (CDB for active/standby backup in the redundancy scenarios).
Distributed Message Queue (DMQ) -Kafka	Uses a publish/subscribe-based message system to provide transmission channels for massive data exchanges.	The components can be deployed in an independent cluster or on the same cluster node where the service processing component is deployed. A maximum of five nodes can be deployed.
Distributed Lock (DLK) Distributed lock (ETCD)	Provides capabilities of locking global resources for services.	
Network Optimization (NETO) Network optimization node (SNC)	Collects the entire network topology information, optimizes the topology paths, and delivers optimization results to southbound forwarding devices.	In a cluster, only two nodes in active/standby mode support topology optimization. The cluster node deployed in exclusive mode cannot be deployed on the same node where other components are deployed. The component is optional. You do not need to install it if the PCE+/EVPN function is not required.

- The network optimization node connects to the forwarding device through southbound interfaces. It collects the entire network topology information and takes over the network port to improve the package sending and receiving performance. Topology optimization requires high computing performance, so the network optimization node needs to be exclusively deployed on a node. Except the network optimization node, other components can be deployed independently on a node or on the same nodes where other components are deployed. The deployment location can be adjusted based on data access relations during installation.

Network Plane

Name	Function	Recommended Number of NICs
Internal network plane	• The plane is isolated physically or by VLAN, and is mainly used for internal communication of the Agile Controller-DCN cluster. • The plane requires high reliability and security, and is isolated from external components. • When there are SNCs, it is divided into two sub-planes: the internal communication plane for inter-node communication, and internal management plane for the Agile Controller-DCN nodes to deliver NETCONF configurations to SNCs.	Double NICs are recommended for higher reliability.
Service network plane	• The plane uses southbound protocols to connect to external forwarders. • The plane is divided into two sub-planes: the southbound protocol plane and southbound management plane.	You can choose single NIC or double NICs based on NIC resources of physical machines and networking reliability.
Management network plane	• Northbound management is mainly used for O&M, management, and service provisioning. • The presentation mechanisms of external interfaces for northbound service provisioning and northbound management are unified. The usage scope can be controlled based on user rights and the management network can be shared. • If northbound service interfaces need to be open to the public untrusted zone, security protection, such as NAT on the hardware firewall, whitelist, and packet cleaning, is required.	

- The network optimization node(SNC) connects to forwarding devices through southbound interfaces. It collects the entire topology information and takes over network ports to improve the packet sending and receiving performance. So during the installation with SNCs, the network port name must be ETH0/1/2/11.
- ETH0 is used for inter-node communication in a cluster, ETH1 is used for the Agile Controller-DCN to deliver configurations to the SNC, ETH2 is used to connect to forwarders through southbound protocols, ETH3 to ETH10 are reserved, ETH11 is used for northbound management.
- In actual deployment, VMs are used to avoid special requirements for network port names of physical machines. The vNICs of VMs can be bound to any network port of physical machines through vSwitches.
- Southbound protocols include control protocols such as OF/PCE/BGP/IGP, and management protocols such as NETCONF/SNMP. Usually the management port of a forwarder can be configured with management protocols only. Service ports can be configured with both control and management protocols.
- So there are two networking modes: Service ports can be configured with control protocols and implements in-band management; Or the control protocols are configured on service ports, and management protocols are configured on management ports. In this mode two networks are isolated. The mode is mainly for carriers.
- When the Agile Controller-DCN interconnects with the southbound forwarder network, ETH2 is used to connect the control protocols. Management protocols are connected with IP addresses discovered by devices. Any network port can be reused. To ensure security, the configuration items of the dedicated southbound management network plane are added. ETH2 is used by default. It can be modified to other ports such as ETH11.

Northbound Request

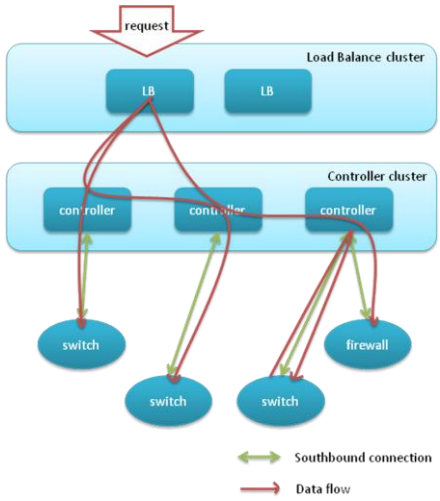

- After a northbound request enters the LB (Nginx), the LB (Nginx) forwards the request to an Agile Controller-DCN node for processing based on the polling policy. The Agile Controller-DCN node performs calculation and determines the network devices to which configurations will be delivered. If the network devices are not managed by the Agile Controller-DCN node, the Agile Controller-DCN node distributes the configuration to another Agile Controller-DCN node which will then deliver the configuration to the southbound network device.

Network Optimization Node

- As the southbound routing module, the network optimization node collects network topology information, optimizes routes on the entire network, and delivers control information to NEs.
- The dual-host hot backup deployment mode is adopted.

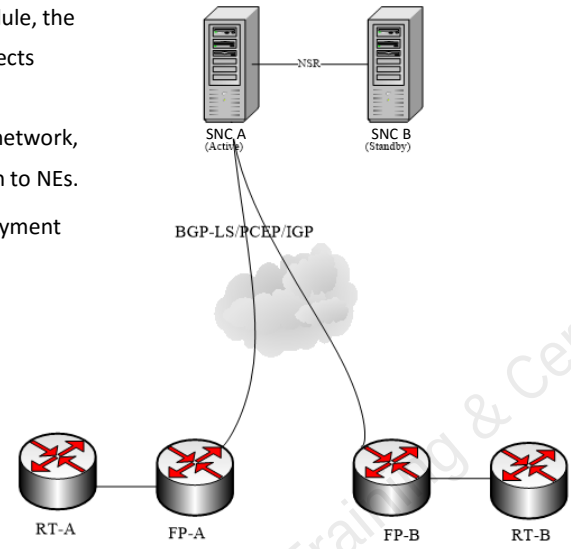

- Non-stop routing(NSR) mechanism: The network control node supports the NSR mechanism. The standby component in the Agile Controller-DCN back up the data status of the active component in real time to ensure data consistency. If a component, process, node, or board on the active network control node fails, services can be automatically switched to the standby network control node without interrupting traffic.
- Warm reset mechanism: The network control node supports the warm reset mechanism. The service components back up the resource data to the database. If the network control node fails and restarts, it can restore the resource data that is requested before the fault occurs from the database.
- The figure shows the dual-host hot backup process. The network optimization node implements backup of configuration and dynamic data using the NSR mechanism. If the active network optimization node fails and restarts, the standby network optimization node (implementing NSR) becomes the active node. Connection with the forwarder has been backed up to the standby node, and the standby node can thereby establish connection with the forwarder.

Cluster Management Service

- The cluster management service is a function on the management plane of the Agile Controller-DCN cluster. It provides functions such as cluster node status monitoring, alarming, and configuration management for administrators.
- The cluster management service includes two parts: Web service presented to users and Puppet responsible for configuration management.
- The Agile Controller-DCN uses Puppet to save and deliver configuration files, start and stop the processes, and expand and reduce capacity.
- The Agile Controller-DCN cluster monitoring page displays cluster node information, including the host name, IP address, NIC, operating system, software version, as well as CPU, memory, and disk usage. The cluster module running status is also monitored for the administrator to locate faults in time.

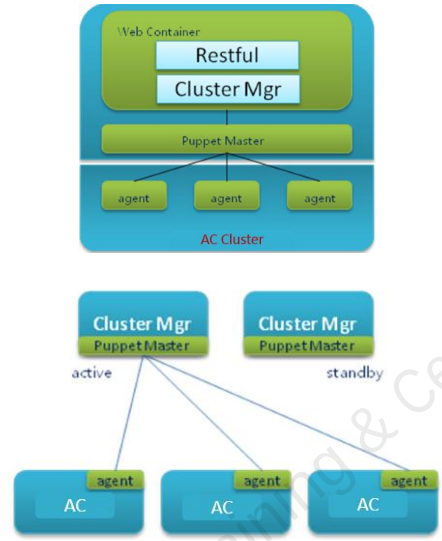

- The Web Container provides interactive web pages, RESTful interfaces, and backend services. It provides functions such as cluster monitoring, cluster capacity expansion and reduction, and QoS configuration of cluster packets.
- Puppet is an open-source and centralized system configuration management framework. It can manage configuration files of the Agile Controller-DCN components running on the server, operating system users, software packages, and system services. The Puppet agent node regularly communicates with the Puppet master to ensure that the local configuration information is consistent with that on the Puppet master node.

Single-Cluster Database Deployment

- Databases consist of the local database and service database. The local database only stores cluster-specific data which will not be synchronized to the database node in the disaster recovery scenario. The service database stores service data which will be synchronized to the disaster recovery cluster for backup. In this way, data that needs to be synchronized is reduced, improving system performance.
- Non-DR scenario:
  - The databases in the production cluster work in active/standby mode. The active database provides external services.

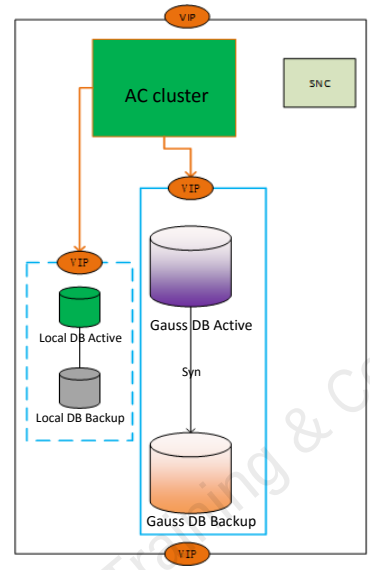

Contents

1. Agile Controller-DCN Overview
2. Agile Controller-DCN Architecture
3. **Agile Controller-DCN Networking Mode**
4. Agile Controller-DCN Deployment Process

Deployment Using a VM (6 Ports - Recommended)

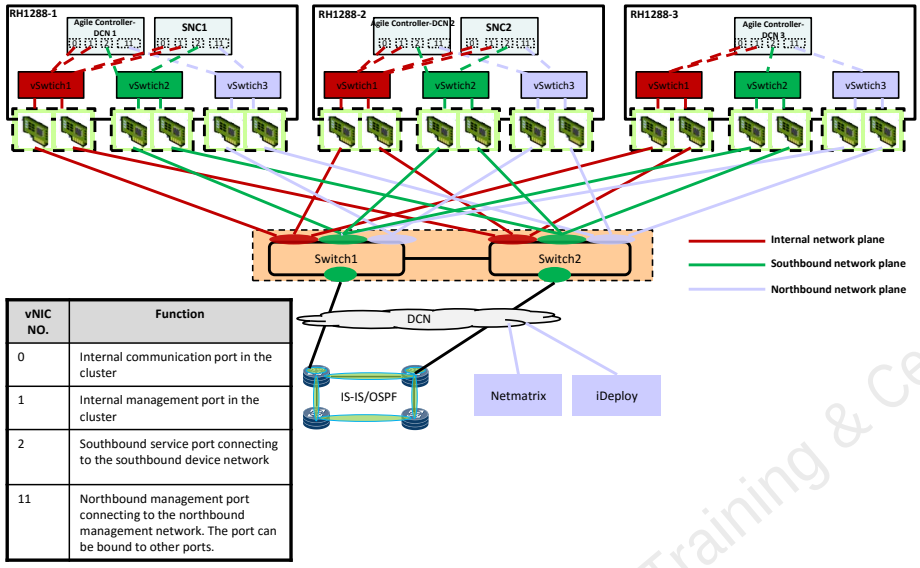

vNIC NO.	Function
0	Internal communication port in the cluster
1	Internal management port in the cluster
2	Southbound service port connecting to the southbound device network
11	Northbound management port connecting to the northbound management network. The port can be bound to other ports.

- After infrastructure-layer software is installed, a vSwitch is created and a physical uplink interface is bound to it.
- After creating a VM template, a specified virtual NIC can be bound to the vSwitch.
- For the VM mode, networks can be isolated physically or by VLANs. The number of network ports can be 1 to 6 when VLAN isolation is adopted. The figure shows a reliable networking where three network planes are all configured with two ports for backup.

VM Deployment (4 Ports - Reliability First)

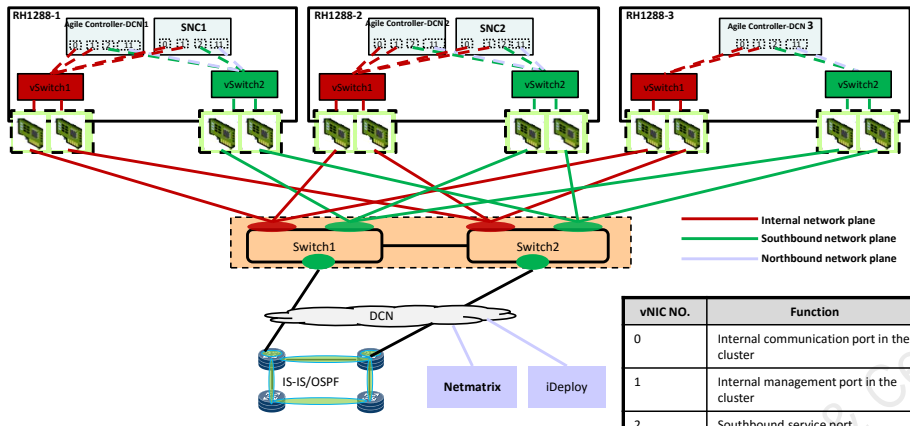

vNIC NO.	Function
0	Internal communication port in the cluster
1	Internal management port in the cluster
2	Southbound service port connecting to the southbound device network
11	Northbound management port connecting to the northbound management network. The port can be bound to other ports.

- After infrastructure-layer software is installed, a vSwitch is created and a physical uplink interface is bound to it.
- After creating a VM template, a specified virtual NIC can be bound to the vSwitch.
- The northbound plane and the internal plane can share vSwitch1 and be isolated by VLAN. The southbound plane is configured with an independent vSwitch and is physically isolated (recommended).
- The internal plane can be configured with a single vSwitch and be physically isolated. The northbound and southbound planes share the same vSwitch and are isolated by VLAN.

VM Deployment (4 Ports - Security First)

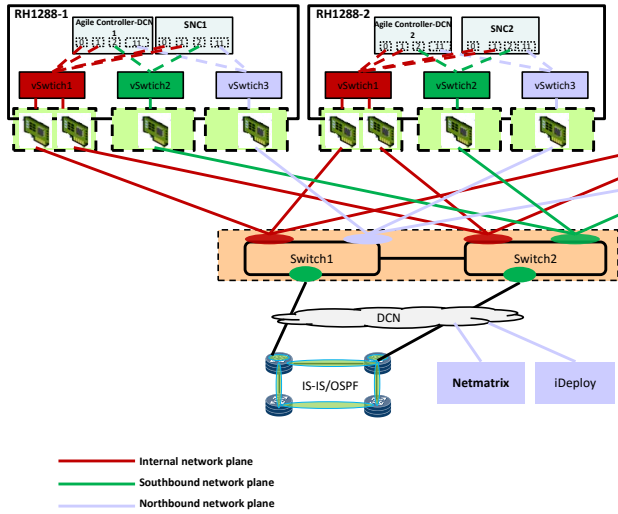

vNIC NO.	Function
0	Internal communication port in the cluster
1	Internal management port in the cluster
2	Southbound service port connecting to the southbound device network
11	Northbound management port connecting to the northbound management network. The port can be bound to other ports.

- After infrastructure-layer software is installed, a vSwitch is created and a physical uplink interface is bound to it.
- After creating a VM template, a specified virtual NIC can be bound to the vSwitch.
- Three planes are physically isolated. The internal cluster communication plane is configured with two ports for backup and other communication planes are configured with a single port.

Physical Machine Deployment (6 Ports)

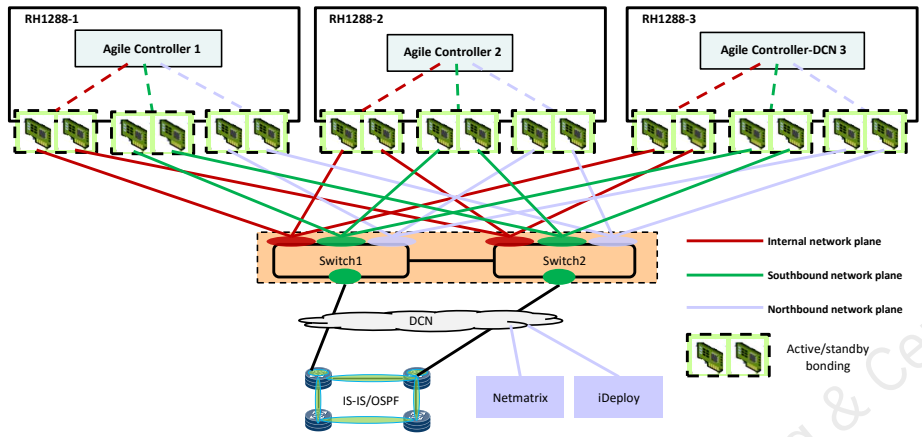

- The internal, southbound, and northbound network planes can be dynamically bound to separate network ports.
- BOND ports can be bound to implement the dual-port backup for high reliability (BOND ports need to be manually configured first).
- Each plane supports 1 to 6 ports. It is recommended that 6 ports be used to implement physical isolation and high reliability.

Layer 3 Deployment (Physical Machine Deployment in Multiple Equipment Rooms - OSPF)

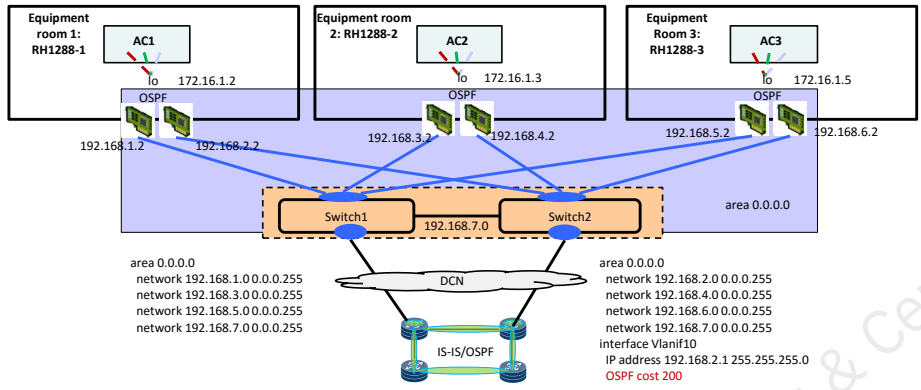

- Some components of the Agile Controller-DCN use the active/standby mechanism where a floating IP address is required. The floating IP address cannot be directly deployed for different segments of different equipment rooms and must be advertised by routing protocols.
- Layer 3 network deployment is complex, which requires routing protocols for interconnection with the peer Layer 3 switch. 7 VLANs, 8 network segments, and 17 basic IP addresses are needed.
- Intra-cluster communication, northbound, and southbound planes share the IP address of the loopback interface, and implement reliability through two routing protocols (value of COST).

Layer 3 Deployment (Physical Machine Deployment in Multiple Equipment Rooms - BGP)

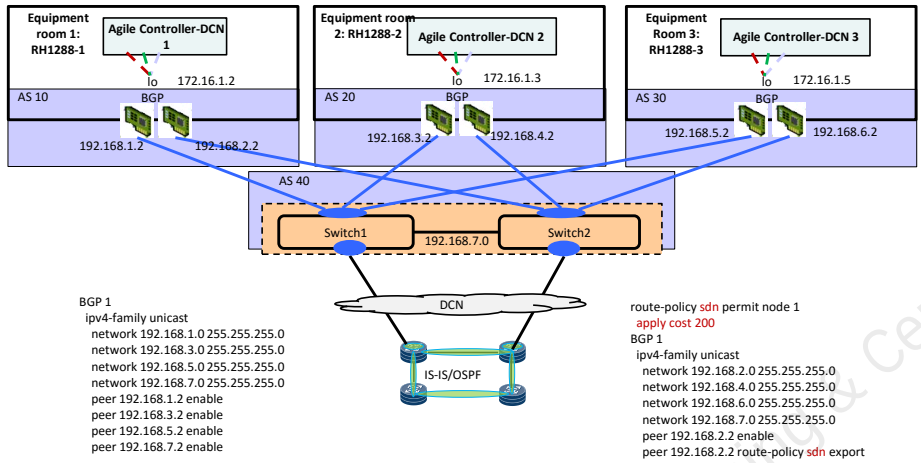

- Some components of the Agile Controller-DCN use the active/standby mechanism where a floating IP address is required. The floating IP address cannot be directly deployed for different segments of different equipment rooms and must be advertised by routing protocols.
- Layer 3 network deployment is complex, which requires routing protocols for interconnection with the peer Layer 3 switch. 7 VLANs, 8 network segments, and 17 basic IP addresses are needed.
- Intra-cluster communication, northbound, and southbound planes share the IP address of the loopback interface, and implement reliability through two routing protocols (value of COST).

Contents

1. Agile Controller-DCN Overview
2. Agile Controller-DCN Architecture
3. Agile Controller-DCN Networking Mode
4. **Agile Controller-DCN Deployment Process**

Overview

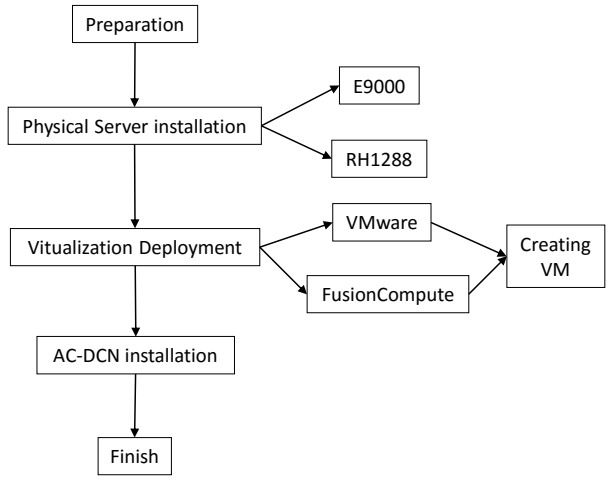

- Procedure (4 steps):
  - Make preparations (environment, packages, and tools).
  - Configure physical servers (RAID and NIC).
  - Install the operating system or virtualization software.
  - Install and deploy the Agile Controller-DCN.
- Pre-installation:
  - Including physical server configuration, installation of the operating system or virtualization software (only for FusionCompute), and installation of the Agile Controller-DCN software
- Deployment using a physical machine or a VM
  - For deployment using a physical machine, a disk is provided for one-click operating system installation, including the dependency package and security hardening package.
  - For deployment using a VM, FusionCompute and VMware templates are provided to shorten the VM creation and deployment time.

Network Plan

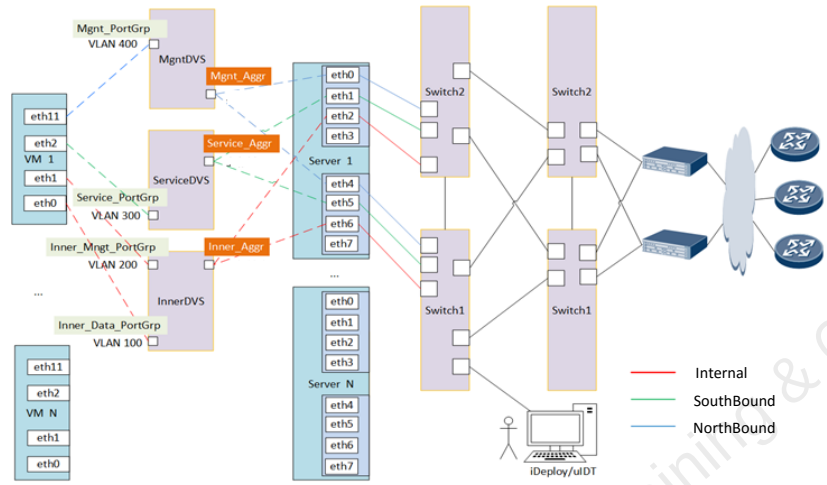

- The overall networking consists of two parts:
- vSwitches
  - Create a distributed virtual switch (DVS) on each network plane (MgntDVS is created by default on the FusionCompute).
  - Create an uplink on the vSwitch and bind it to the physical network port (one network port is selected for each of the two physical NICs to back up each other).
  - Create a port group (VLANs can be configured). You can bind the virtual port to the port group when creating a VM.
- Physical switches
  - Six physical ports of a physical server are divided into three groups, which are respectively dual-homed to two access physical switches.
  - Configure an access port on the access switch to allow VLANs to pass through. The uplink link is dual-homed to the aggregation switch.
  - The aggregation switch functions as a gateway. Configure a subinterface or VLANIF to terminate VLANs. The Layer 3 routes can communicate with the firewall and forwarder.

- For deployment using a physical machine, the Agile Controller-DCN is deployed on a physical machine. Configure the physical ports with two NICs to BOND ports in pairs. Choose these ports when installing the Agile Controller-DCN to improve reliability.

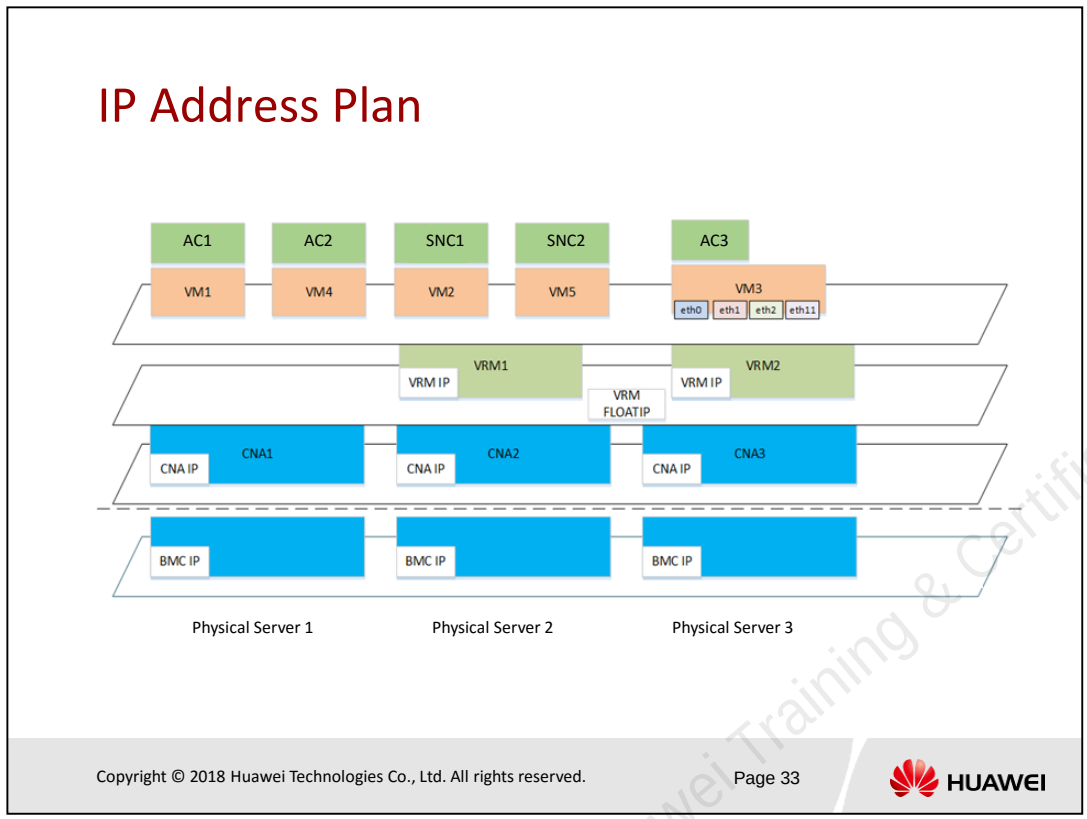

- Physical servers provide the web page of iBMC for users to view the version information and running status, change the password and IP address, and restart (E9000 is the active/standby management board).
- Each node of the virtualization layer requires an IP address. Active/standby virtualization resource management (VRM) nodes need to be installed for FusionCompute. Three extra IP addresses are required and cannot be the same as the IP address of iBMC. The IP addresses can be on the same network segment.
- Four IP address segments are required for three network planes and four ports of VMs. ETH0 and ETH1 for the internal communication plane are for internal subnets, ETH2 is connected to the forwarder, and ETH11 is used for VM management and northbound plane.
- A total of five IP address segments are required, two for internal subnets and three for services or management. ETH11 can also share the same network segment with iBMC and the virtualization layer for management. In this case, only four network segments are required.
- For deployment using a physical machine, three network segments, three ports (which can be BOND ports), and three network segments are recommended. You can also configure the same port for different network planes to reduce the number of ports and network segments.

Required Software

Solution for I-Layer	Software	Description
FusionCompute solution	FusionCompute V100R006C00RC1_CNA.iso	An ISO image file used for the FusionCompute CNA (Compute Node Agent) installation
	FusionCompute V100R006C00RC1_VRM.zip	A software package. It is decompressed to install FusionCompute VRM (Virtual Resource Manager) through the corresponding program.
	ac_template_for_fc_xxx.ova	Huawei FusionCompute template for VM creation. The matched Guest OS is also provided.
VMware solution	VMware-VMvisor-Installer-version.x86_64.iso (supporting ESXi 5.5 Update 2 and later versions)	An image file used for VMware ESXi installation
	VMware-viclient-all-version-299632.exe (supporting ESXi Client 5.5 Update 2 and later versions)	An executable program for VMware vSphere Client installation and VMware ESXi configuration
	ac_template_for_VMware_xxx.ova	Huawei FusionCompute template for VM creation. The matched Guest OS is also provided.

Project	Software	Description
One-click installation disk of the SUSE operating system	An ISO image file: Agile Controller-DCN version SUSE11-SP3.iso	Used for installing the 64-bit SUSE11 SP3 operating system (enterprise) on physical machines and third-party VMs, including the dependency package, disk partition, and security packages
iDeploy	Installation package for the Windows operating system: Breeze_iDeploy_version_win.zip Installation package for the Linux operating system: Breeze_iDeploy_version_suse.zip	Web tool for the Agile Controller-DCN installation and deployment, which can be executed after decompression
Agile Controller-DCN installation software	Software package: AgileController-DCNversion_install_pkg.zip	Used for installing the Agile Controller-DCN on hosts
Agile Controller-DCN configuration software	Software package: AgileController-DCNversion_install_pkg.zip	Used for configuring the Agile Controller-DCN. Services are deployed on specified hosts based on the networking.

- VM templates are provided to shorten the VM creation and deployment time.
- A one-click installation disk for the SUSE operating system is provided for quick installation.
- The signature files can be downloaded together with the software packages for validity verification.

Physical Server Installation

Server	Default IP address	Login Mode
E9000	Active: 192.168.2.1/24 Standby: 192.168.2.2/24 Floating IP address: 192.168.2.3/24	https:// floating IP address of MM910 root/Huawei12#$
RH1288V3	192.168.2.100/24	https://default IP address root/Huawei12#$

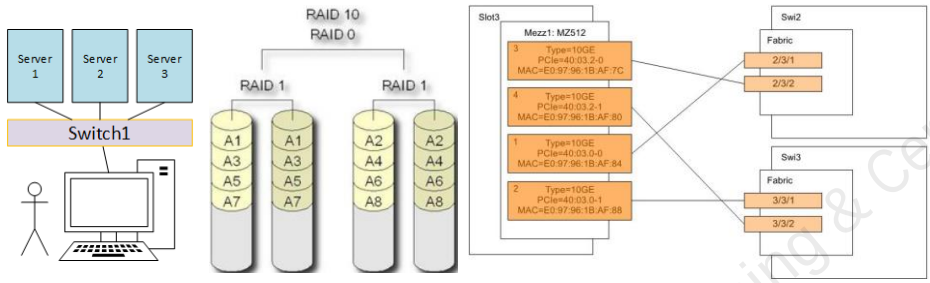

HUAWEI

- The server provides the web management page. You can log in to the server using default mode. After the login, you can change the password and IP address, and configure the hardware and NIC on the remote virtual console page.
- It is recommended that hard disk be configured with RAID 10 (or RAID 1). Divide four hard disks into two groups and back them up by using RAID 1, and then implement rapid reading and writing through RAID 0. Delete the existing RAIDs first.
- For E9000, use four ports or channelize two ports of them based on requirements. Each port can be channelized into 4 PFs (the PFs need to be configured with different PVIDs from those of vSwitches, which makes the networking more complex).

Operating System Installation (Deployment Using a Physical Machine)

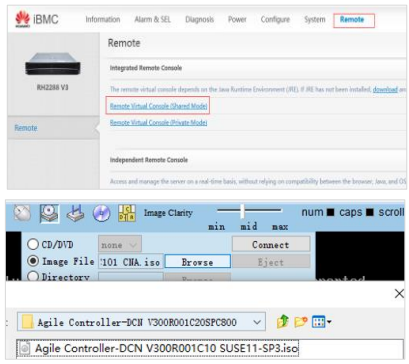

- Log in to the web page of iBMC, choose **Configure** > **Boot Option**, select **Boot from Hard Disk**, and save the settings.
- Select **Remot** > **Remote Virtual Console(Shared Mode)** to open the virtual console.
- Set the drive icon, select **ISO**, click **Browse**, and select **Agile Controller-DCN version SUSE11-SP3.iso**. Click **Connect**. After the image file is successfully connected, click the triangle icon and click **ced System Reset** to restart the physical host.
- After the host is restarted, the startup page of SUSE Linux Enterprise Server 11 SP3 is displayed. Click **Installation** to install the operating system.

Virtualization Software Installation (VM Deployment - FusionCompute)

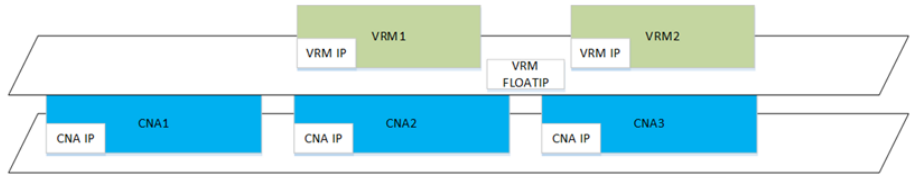

- FusionCompute manages virtual resources by using VRM which uses Compute Node Agent (CNA) to manage PM resources.
- Before installing the virtualization software, follow the OS installation steps to install the CNA on the physical machine. Mount the ISO image file through the virtual remote console to install the software. Configure the management IP address and root password.
- Decompress the VRM package on the Windows operating system. Double-click the executable program to install VRM to CNA 2 and CNA 3 respectively. The active/standby mode is recommended.
- After the installation, log in to the VM management page through http://*VRM floating IP address*. The default user name and password is **admin** and **Huawei@CLOUD8!**, and can be changed upon the first login.

Virtualization Software Installation (Adding Resources)

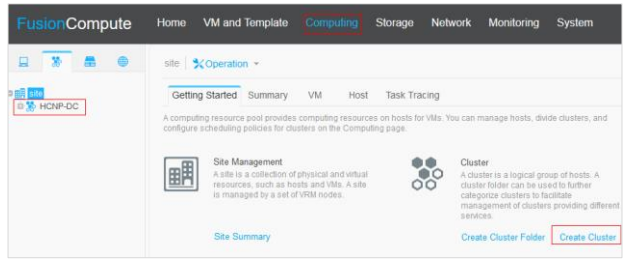

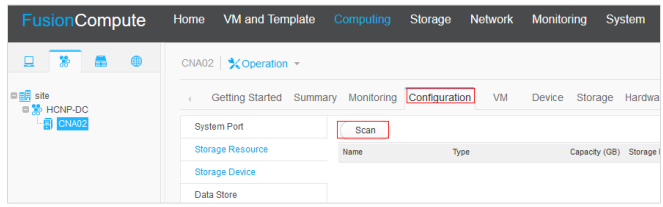

- The CNA nodes where the active/standby VRMs are located will be automatically added to the ManagementCluster computing pool. Other CNA nodes need to be manually added.
- On the Configuration tab page, click Scan. Hard disks added to the new CNA nodes are added to DataStore.

Virtualization Software Installation (Creating a DVS)

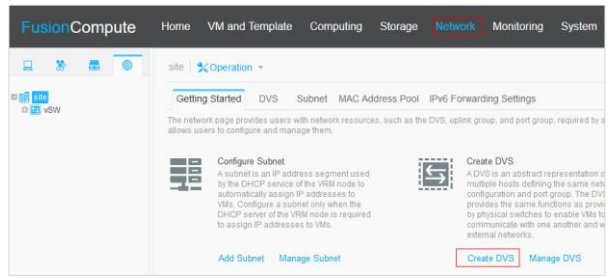

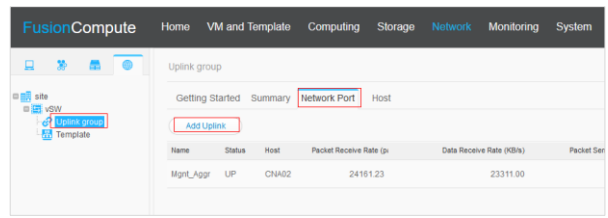

- Create a DVS and add an uplink.

HUAWEI

- After the VM management software is installed, ManagementDVS is created by default for the management plane. Two more DVSs need to be created respectively for the internal data plane and service plane.
- After the creation, click the vSwitch and select **上行链路组**. On the **网口** page, click **添加上行链路** to add uplinks to the three vSwitches in turn.
- Only one uplink can be added between each vSwitch and each host. However, users can bind the physical ports to logical ports(BOND) and add the BOND to the uplink.
- The management switch has been bound to a CNA management port by default. Users can add another physical port to improve reliability.

Virtualization Software Installation (Creating a Port Group)

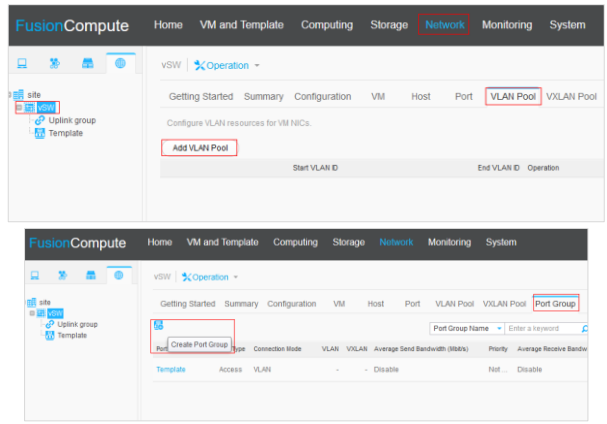

- If the VM networks need to be isolated by VLANs, you need to add a VLAN pool to the vSwitch, and set the start and end VLAN ID.
- Create a port group for the vSwitch and specify a VLAN for vNIC binding when creating a VM.

Virtualization Software Installation (Importing a VM Template)

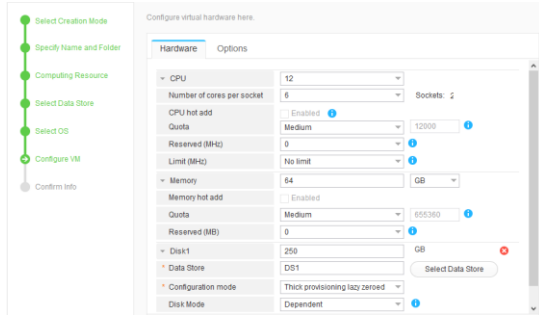

- On the FusionCompute page, select VM and Template, click the Getting Started tab, click Import Template, and select Local PC to import the VM template.
- A window will be displayed when the VM template is imported.
- When deploying the template, set Configuration mode to Thick provisioning lazy zeroed and select port groups for NICs. Add the port groups and NICs to vSwitches.

Virtualization Software Installation (Creating a VM)

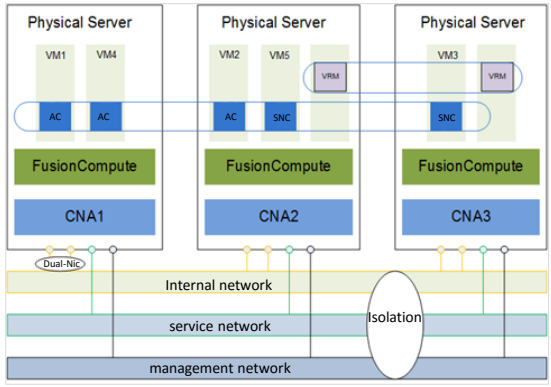

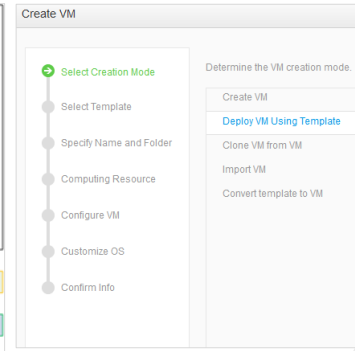

- After the template is imported, click the Template tab. Click more and select Deploy VM Using Template.
- Create SNCs for CNA 1 and CNA 2 based on the plan. You can modify CPU specifications when configuring the VM. Create the Agile Controller VMs for CNA1, CNA2, and CNA3 with default specifications.
- If there are other FusionCompute instances on the Layer 1 and Layer 2 networks, you need to choose Network > MAC Address Pool and modify the MAC address scope first to prevent VM MAC address conflict.

Virtualization Software Installation (Creating a VM)

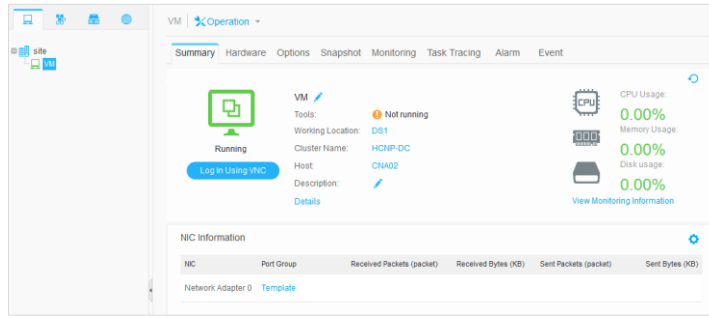

- When creating a VM, select start VM. Choose Options > HA and set HA to Enable, so that VMs can start with the physical server. After the VM is started, you can log in to the VM through VNC and change the passwords of the root user and swift user.
- Both the service management node and the network optimization node need to be deployed across servers to prevent cluster failure due to fault of a single physical machine. Therefore, configure IP addresses for VMs in the ascending order. At the same time, update the external switch configuration to ensure communication between VMs.

Preparing the Tool (1)

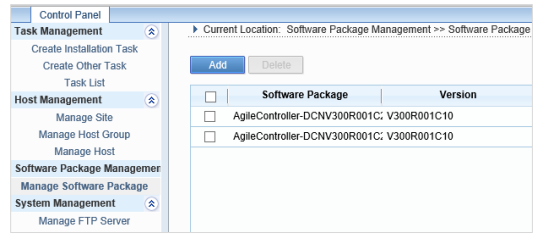

- Use iDeploy to install and deploy the Agile Controller-DCN.
- Upload the installation and configuration packages locally.

Preparing the Tool (2)

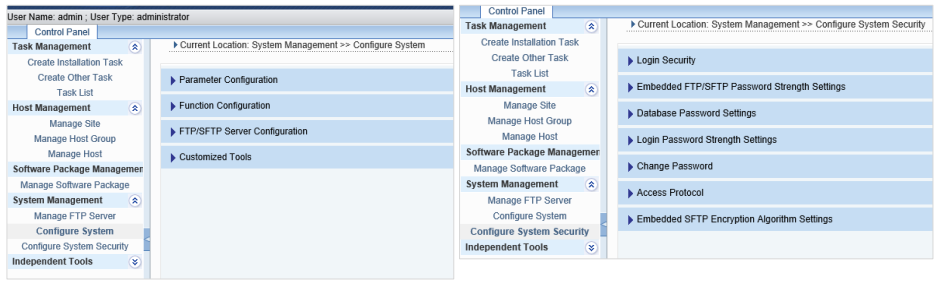

- An FTP or SFTP server is built in iDeploy for the installation host to download the uploaded software package.
- On the left-side control panel, choose **System Management** > **Configure System**. On the configuration page, choose **FTP/SFTP Server Configuration** to view and modify the settings.
- Upon the first login, a dialog box is displayed for you to choose the listening IP address of the FTP or SFTP server for the installation host to connect to.

Preparing the Tool (3)

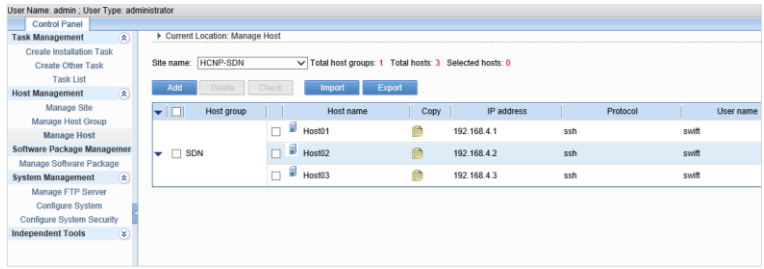

- On the host management page of the iDeploy control panel, add the site, host group, and host in turn. The host is the physical machine or VM where the Agile Controller-DCN software is to be installed.
- Ensure that the iDeploy client, iDeploy server, and each node of the installation host can communicate with each to avoid installation and deployment interruption due to network faults.

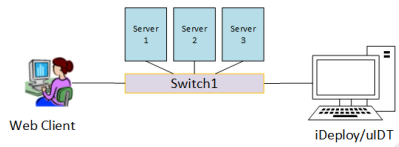

Agile Controller-DCN Software Installation (Creating a Task)

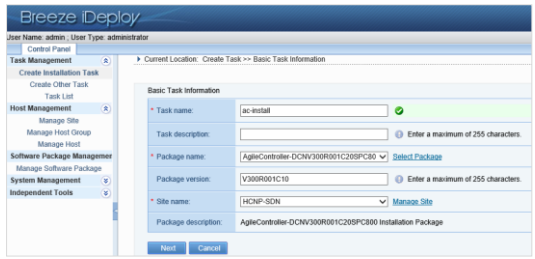

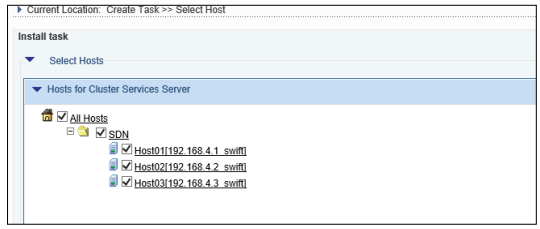

- The Agile Controller-DCN is installed step by step using iDeploy. The whole process contains installation and configuration which correspond to the installation task and configuration task respectively.
- The released software package is also divided into the installation package and configuration package. The former is used to install the Agile Controller-DCN components based on the installation task. The latter is used for the initial configuration and startup of the Agile Controller-DCN components based on the configuration task.
- 1. Create an installation task, and select the installation package and the corresponding site.
- 2. Select the host group to be installed.

Agile Controller-DCN Software Installation (Executing a Task)

- After creating a task, you can view the status, information, and installation progress it in the task list (as is shown above). Click the task name. The detailed task execution page is displayed.
- The lower left figure shows the detailed task page. The upper-left area displays user-defined steps, and the middle part shows the installation progress of each host. Select a host. The lower part of the page displays the output information during the installation. The execution failure output is displayed using red keywords.
- The software package uploading takes most of the installation time, which is about five minutes in total. After the installation is completed, you can click 卸载 in the lower left corner of a specific installation task to uninstall the software.

Agile Controller-DCN Cluster Deployment (Creating a Task)

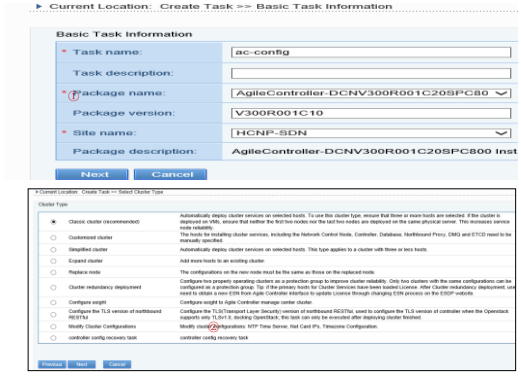

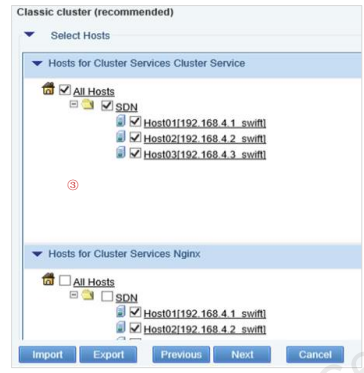

- 1. Create an installation task, and select the installation package and the corresponding site.
- 2. Choose a deployment mode.
  - Classic cluster deployment: You only need to select a host, and do not need to select and assign the internal components of the Agile Controller-DCN. The components will be automatically selected and assigned during the installation.
  - User-defined cluster deployment: If you are familiar with the internal framework of the Agile Controller-DCN and specific component requirements, you can select cluster nodes based on component features to install and deploy the components.
- 3. Select the host where the components are deployed.

Agile Controller-DCN Cluster Deployment (Configuring the Network Plane)

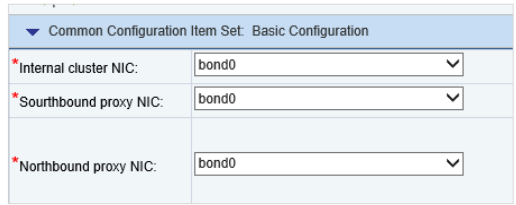

- Three scenarios of basic network plane configuration:
- 1. Deployment with SNCs: The port name is from eth0 to eth11. The northbound proxy network port is eth11 by default. It can also share a port with other planes.
- 2. Non-Layer 3 deployment without SNCs: Three planes can all be dynamically configured, and can share a port. The port can be a bare physical port or preconfigured BOND port.
- 3. Layer 3 deployment without SNCs: Two ports are used as route backup, and cannot be shared as one port. OSPF and BGP protocols are supported. If OSPF is used, AERA is 0.0.0.0 by default. You only need to fill in three virtual IP addresses. If BGP is used, you need to set the virtual IP address, local AS, and IP address and AS of the remote switch.

Agile Controller-DCN Cluster Deployment (Configuring a Database)

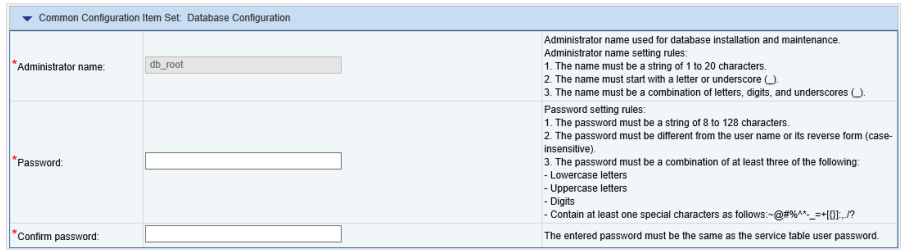

- Two instances of basic database configurations:
  - 1. General database (data that requires remote backup in the DR scenario)
  - 2. Local database (local data that does not require remote backup)
- The two database instances use floating IP addresses to provide external services.

Agile Controller-DCN Cluster Deployment (Configuring Floating IP addresses)

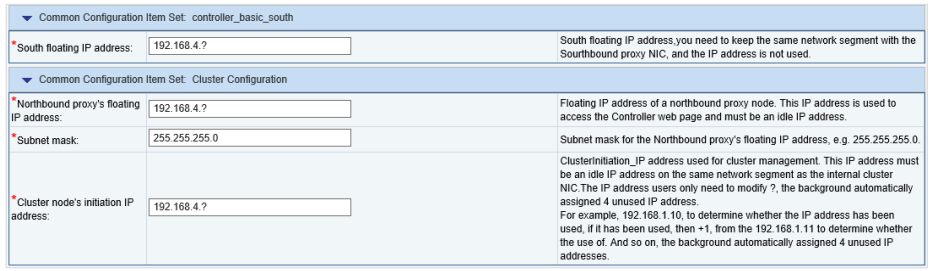

- Floating IP address of the northbound proxy NIC (mask and gateway): It is used for all northbound information and configured on the northbound proxy NIC. You need to configure the mask and gateway if this IP address is not on the same segment as that of the northbound proxy NIC.
- Internal management floating IP address of the northbound proxy NIC: It is used to access the SNC RR+ service and is configured on the internal management NIC of the cluster.
- Intranet floating IP address of the northbound proxy NIC: It is used for intranet NAT address translation and is configured on the internal communication NIC of the cluster.
- Cluster management floating IP address: It is used for cluster management and is configured on the internal communication NIC of the cluster.
- Southbound management floating IP address: It is used for all southbound information and configured on the southbound service NIC.

Agile Controller-DCN Cluster Deployment (Configuring SNCs)

- Floating IP address: SNCs are deployed in active/standby mode and use the floating IP address to provide external services.
- Account name: It is used for the Agile Controller-DCN NETCONF to access the northbound SNCs. The password is dynamically generated during installation and can be modified.
- Instance ID: It is used to identify multiple sets of the Agile Controller-DCN deployed on the same Layer 2 network.
- Internal network segment: It is the sub-segment used for communication between the active and standby SNCs. If multiple sets of the Agile Controller-DCN are deployed on the same Layer 2 network, they need to be distinguished from each other.

Agile Controller-DCN Cluster Deployment (Configuring the System)

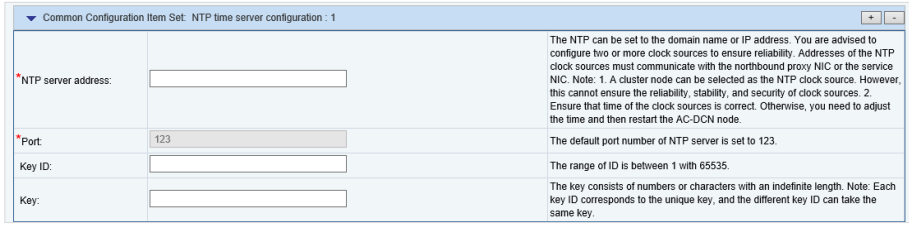

- NTP server: Three external NTP servers or one internal NTP server is supported. You need to set the NTP server IP address or domain name.
- Port: Indicates the port through which the NTP server provides services.
- Key ID: Indicates the NTP authentication key ID.
- Key string: Indicates the NTP authentication key.
- Time zone: Indicates the system time zone, which can be selected from the drop-down list menu.

Agile Controller-DCN Cluster Deployment (Executing a Task)

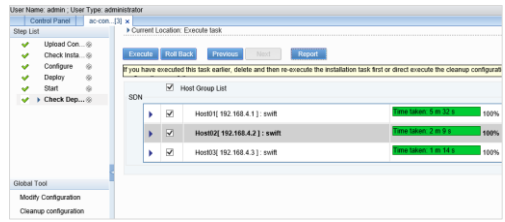

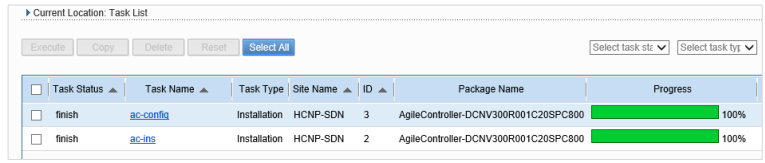

- After creating a task, you can view it in the task list. Select a configuration package task and click **Execute** to start cluster deployment.
- The second figure shows the detailed task page. The upper left corner displays specific sub-steps. The middle part shows the progress and status of the execution host. If the installation fails, click **Cleanup configuration** at the lower left corner and then execute the task again.

Agile Controller-DCN Cluster Deployment (Completing the Installation)

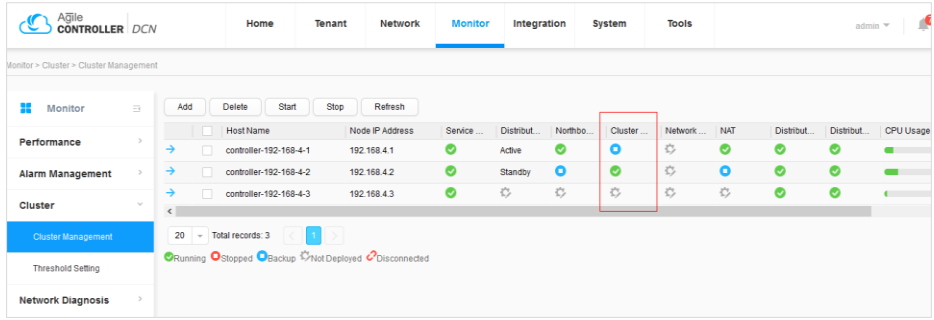

- After the cluster is deployed, log in to the Agile Controller-DCN through https://*IP address of northbound proxy*:18002/index.html. The default user name and password is **admin** and **Changeme123**. You need to change the password upon the first login.
- Choose **Monitor** > **Cluster Management** to view the running status of the host components and the resource usage.

Quiz

1. In a typical Agile Controller-DCN deployment scenario, how many nodes does a cluster contain? ( )
   A. 1
   B. 3
   C. 5
   D. 7
2. What are functions of the southbound and northbound floating IP addresses of the Agile Controller-DCN?

- Answer:
  - 3
  - The southbound floating IP address is used for interconnection with network devices. The northbound floating IP address is used for interconnection with cloud platforms and user access.

Summary

- Agile Controller-DCN Overview
- Agile Controller-DCN Architecture
- Agile Controller-DCN Networking Mode
- Agile Controller-DCN Deployment Process

More Information

- Huawei Support website: http://support.huawei.com
- Agile Controller-DCN Product Documentation

More Information

- Huawei Learning:
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei Support Case Library:
  - http://support.huawei.com/enterprise/servicecenter?lang=zh

Thank You

www.huawei.com

Cloud Data Center SDN Controller Working Principles

www.huawei.com

Foreword

- This chapter describes the basic working principles of AC-DCN 3.0.

Objectives

- After completing this course, you will be able to learn:
  - Master the addition of network device on AC-DCN
  - Understanding the creation of Fabric networks on AC-DCN
  - Familiar with adding VAS resource pool on AC-DCN
  - Understanding Internet resources and global resources

Contents

1. **SDN Basic Consepts**
2. CloudFabric Application Scenarios
3. Deployment Procedures
4. Service Provisioning Procedure
5. Typical Deployment

DCN Architecture Development Trends

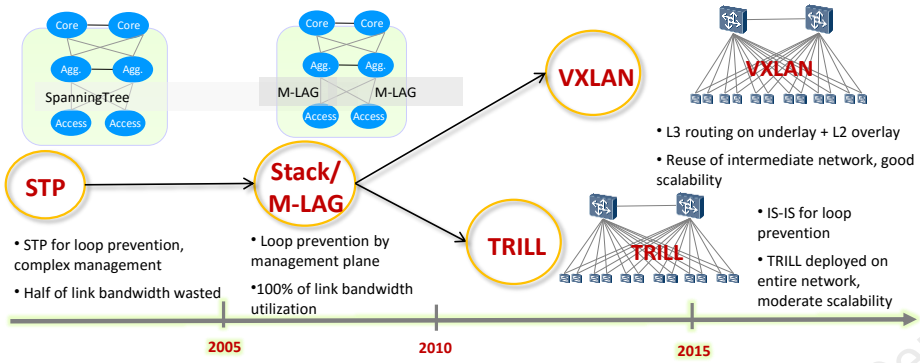

• TRILL technology removes the need for STP and allows for a large L2 network. It uses a mature link state routing algorithm and IS-IS extensions to establish a loop-free network and implements load balancing on multiple equal-cost paths.

• VXLAN technology provides better usability and scalability and is easier to manage and maintain. It has become a mainstream technology advocated by IT and CT vendors.

HUAWEI

- IT+CT consolidation
- TRILL is a revolutionary technology, whereas VXLAN is an improvement technology.

Overlay and Underlay

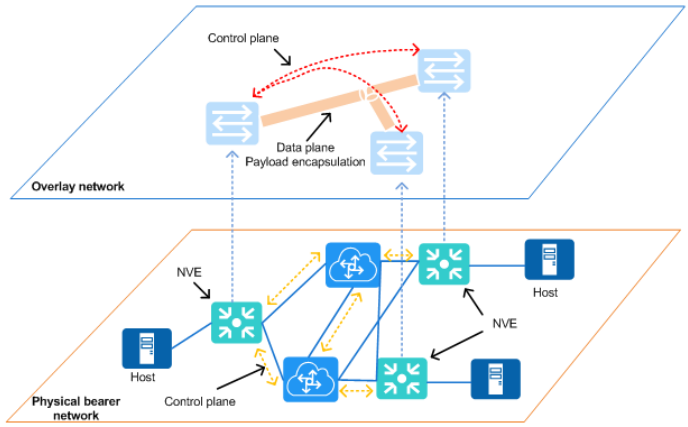

- An overlay network is a virtual network established over an existing underlay network, consisting of logical nodes and logical links.
- An overlay network has independent control and forwarding planes. The physical network is transparent for the terminals connecting to the overlay network virtualization edge (NVE).

- In an overlay network model, a software-defined logical network is added to a physical network without changing the physical network. Service logic of the logical network is defined by software to solve problems of data center networks. The overlay technology transmits Layer 2 packets (of services) to Layer 3 or Layer 4 (of a traditional network).
- The overlay technology includes VXLAN, Network Virtualization using Generic Routing Encapsulation (NVGRE), and Geneve. The principle of the overlay technology is as follows: Overlay encapsulates a Layer 2 packet in a tunnel, transmits the encapsulated packet transparently, and then decapsulates the packet to obtain the raw packet after the packet arrives its destination. In this way, the current network is overlaid with a large Layer 2 network.
- The underlay network is a bearer network consisting of physical devices, such as TOR switches, aggregation switches, core switches, LBs, and firewalls.
- The overlay technology can be implemented to build a logical network over an underlay network.

VXLAN-based Overlay Network

- An overlay network is established using the VXLAN technology. Service packets are transmitted on the VXLAN overlay network and decoupled from the physical bearer network.
- VXLAN-based overlays are classified into network overlay, host overlay, and hybrid overlay according to the deployment mode of VXLAN NVEs.
- Hardware overlay: All NVEs are deployed on physical switches.
- Software overlay: All NVEs are deployed on vSwitches.
- Hybrid overlay: Some NVEs are deployed on physical switches and others are deployed on vSwitches.

VDC

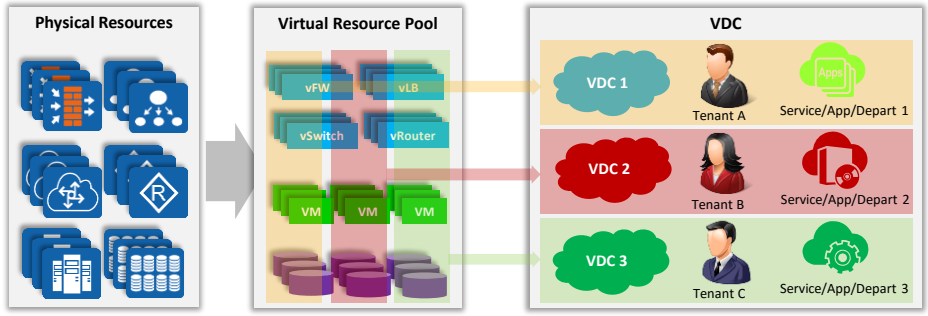

- VDC: A virtual data center is a pool of certain quotas of computing, storage, and network resources designed specifically for only one organization.
- A tenant applies for DCN, storage, and computing resources. Any user, enterprise, or unit can be a tenant when applying for resources.

Tenant

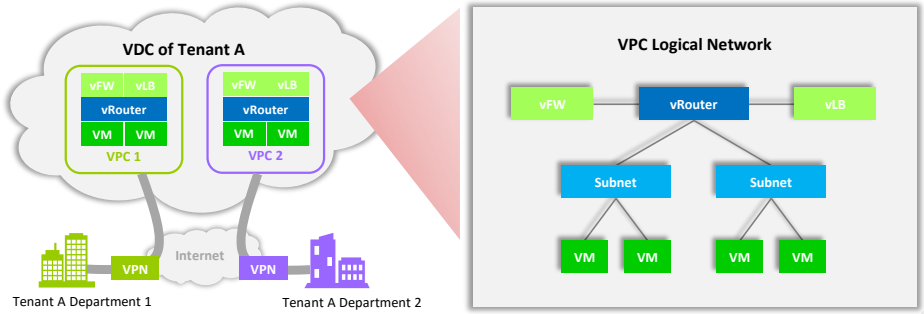

- VPC: A virtual private cloud (VPC) is a security zone for one type of services or a department. VPCs use resources in VDCs. Each VPC belongs to one VDC, and each VDC can have multiple VPCs.
- VPCs have the following advantages:
  - Isolated environments: VPCs provide isolated VMs and network environments for different departments.
  - Diverse services: Each VPC can provide independent virtual firewall (vFW), elastic IP address, security group, super VLAN, Internet Protocol Security (IPSec) VPN, and NAT gateway services (some of the functions are provided by the vFW).
  - Flexible networking: Multiple networking modes are supported, including directly connected networks and routing networks.

VPC Elements

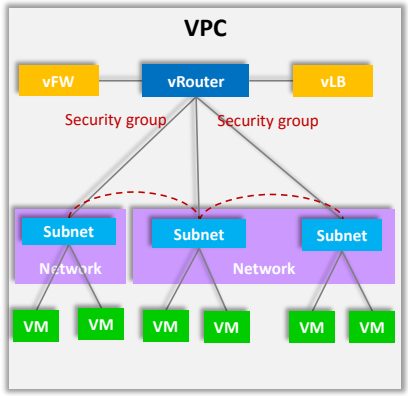

Element	Description
vRouter	• A vRouter acts as the gateway to enable L3 communication between subnets. • One VPC can have only one vRouter.
Network	• A network in a VPC represents an L2 network, which contains only one subnet in common mode and supports multiple subnets in shared mode. (In shared mode, a switch interface is assigned multiple secondary IP addresses to support multiple subnets. An example of this scenario is that multiple small tenants share the same VLAN.) • Note: The FusionSphere service model does not support the shared mode and therefore does not show networks. The native OpenStack service model shows this element.
Subnet	• Subnets are used to isolate L2 broadcast domains. • The same vRouter acts as the L3 gateway for different subnets in a VPC. • Hosts in the same subnet and different subnets can communicate with each other by default. Subnets can be isolated using security groups.
vFW	• A vFW is the border of a VPC that provides access to the VPC from external users and ensures access security. • A vFW provides services such as FW, EIP, SNAT, and IPsec VPN.
vLB	• A vLB distributes traffic sent from external users to internal servers to implement load balancing. • A vLB can have multiple listeners for different services.

SDN in Cloud Data Center

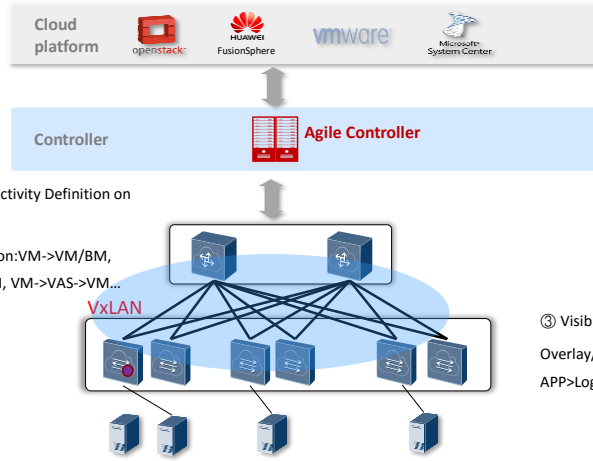

- ① Connectivity Definition on Demand:
- Connection:VM->VM/BM, VM->VPN, VM->VAS->VM…

② Connectivity Multi-Cloud & Non-Cloud, Do what CMS can not do:

- Connection between physical and virtual resources
- Connection between different services

③ Visibility across overlay & underlay:

Overlay/Underlay network mapping, APP>Logic>Physical network

Fabric Network

- A fabric network consists of one set of interconnected spine and leaf nodes and provides undifferentiated communication among all connected nodes. The fabric network allows multiple tenants to use one physical device and one tenant to use multiple physical devices. This feature slashes costs and increases resource usage.

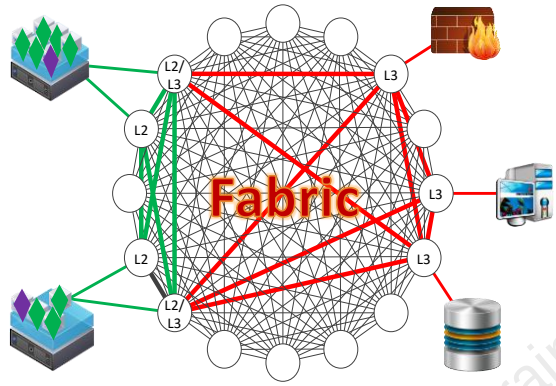

Topology Design - Overview

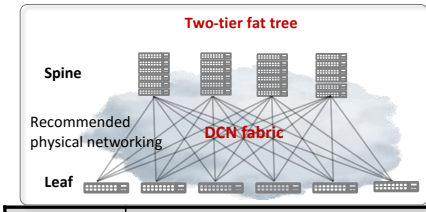

Role in Physical Networking	Definition and Functions
Fabric	Data center basic physical network topology formed by a group of interconnected spine and leaf nodes.
Spine	Core node on a VXLAN fabric network, which provides high-speed IP forwarding and connects to functional leaf nodes using high-speed interfaces.
Leaf	VXLAN fabric network access node, which connects various devices to the VXLAN network. Leaf nodes are classified into service leaf, server leaf, and border leaf.
Service leaf	A leaf node that connects L4-L7 VAS devices such as FWs or LBs to the VXLAN fabric network.
Server leaf	A leaf node that connects computing resources including physical and virtual servers to the VXLAN fabric network.
Border leaf	A leaf node that forwards traffic sent from outside the data center to the VXLAN fabric network, used to connect routers or transmission devices.

Physical fabric networking

- Physical fabric features:
  - High bandwidth, large capacity
  - Indifference between access nodes
  - Flat architecture
  - Flexible networking
  - Elastic scaling
  - Same technology using in one fabric
- Recommended physical fabric networking:
  - Spine-leaf architecture composed of CE series hardware switches
  - L3 network running an IGP protocol
  - Load balancing and link backup using ECMP
- Fabric service
  - Undifferentiated communication between access nodes

DCN Fabric Uses VXLAN

Cloud Data Center Network Requirements		VXLAN	CSS/SVF	TRILL
> 4K tenants	Resource releasing data centers must support a large number of tenants.	16M	4K	4K (16M supported by the latest version)
Investment protection	New fabrics can be deployed on existing networks.	No specific requirements on existing networks, as long as they have L3 capabilities.	New network construction	New network construction
SDN support	The current network can migrate to an SDN network smoothly.	Overlay is one of the four major SDN development routes.	Not supported	Not supported
Standard protocol	Multi-vendor interoperability must be supported.	Standard	Proprietary	Standard
Cross-DC networking	A large L2 network can be established across an IP WAN.	Supported	Not supported	Not supported

- VXLAN is the most popular technology because of better support for SDN and multi-tenancy.

Various VXLAN Solutions to Meet Different Customer Requirements

Network Overlay

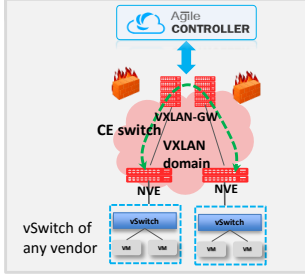

- VXLAN network formed by physical switches
- **High forwarding performance and service reliability**
- Decoupled from VM management platform

Host Overlay

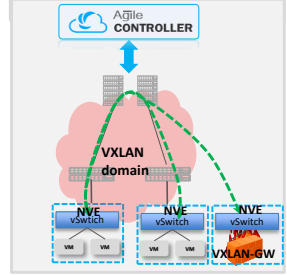

- VXLAN network formed by virtual switches
- Low forwarding performance
- No need to change current physical devices

Hybrid Overlay

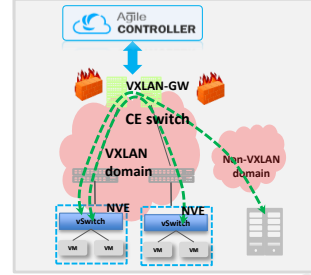

- VXLAN network formed by virtual and physical switches
- Moderate forwarding performance
- VXLAN-capable physical gateway devices added

Leaf and Spine Nodes

- A leaf node is an access node on a VXLAN fabric network, which connects various network devices to the VXLAN network.
- A spine node is a core node on a VXLAN fabric network, which provides high-speed IP forwarding and connects to leaf nodes using high-speed interfaces.

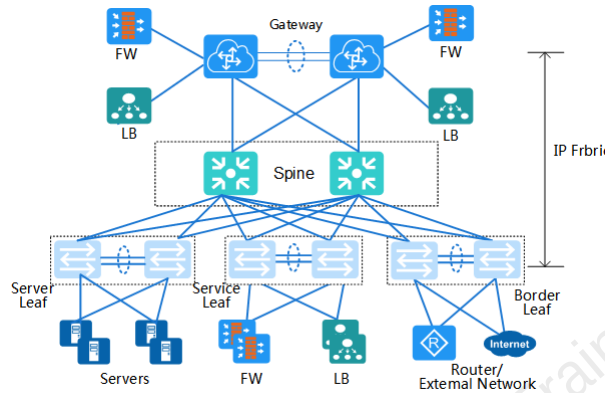

- Leaf and spine nodes are fully meshed so that equal-cost multiple-path routing (ECMP) paths are available to ensure high availability of the network.
  - Generally, spine nodes are two or more large-capacity switches, and connect to leaf nodes at Layer 3 to establish an IP fabric network.
  - Access switches can be deployed at leaf nodes and fully meshed to the spine nodes. Leaf nodes are the boundary between Layer 2 and Layer 3 on the underlay network, and connect to the spine nodes at Layer 3. If many TOR switches are deployed at the leaf nodes, it is recommended that the TOR switches be configured to set up an iStack, a multi-chassis link aggregation group (M-LAG), or a super virtual fabric (SVF) system for server access.
- Leaf nodes are access nodes in the physical networking. They connect servers to the VXLAN network. Leaf nodes can be classified into the following types according to the access objects.
  - Border leaf: connects to routers or transmission devices to forward traffic from external networks to the VXLAN fabric network of a data center.
  - Server leaf: connects virtual and physical servers to the VXLAN fabric network.
  - Service leaf: connects Layer 4 to Layer 7 value-added services (VASs), such as firewalls and LBs, to the VXLAN fabric network.

L4 - L7 Pool

- L4-L7 originally refer to the transport layer, session layer, presentation layer, and application layer. In this solution, L4-L7 refer to services, for example, ACL, load balancing, NAT, QoS, and IPSec VPN, provided by firewalls and load balancers. These services are called VASs.

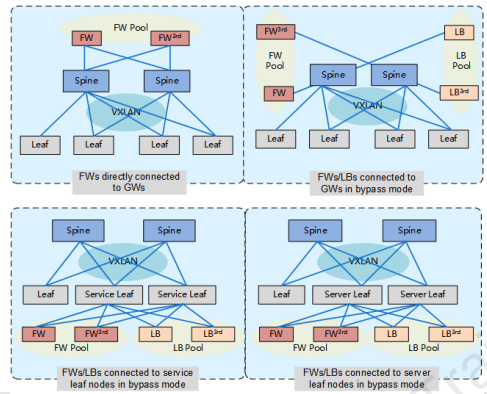

Tenant Network

- Logic Port
  - Logic ports provide access services to virtual machines (VMs), bare metal servers (BMs), andL4-L7 devices.
- Logic Switch
  - Logic switches connect different VMs and enable them to communicate with each other at Layer 2.
  - Logic switches provide L2 switching services among logic ports.
- Logic Router
  - Logic routers connect VMs on different network segments and enable them to communicate with each other.

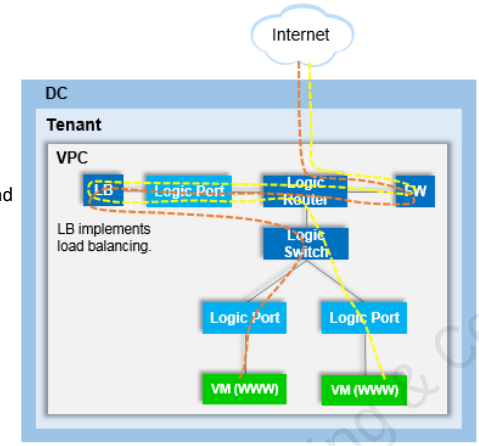

HUAWEI

- Logic Port
  - One physical port on a network device can be virtualized into multiple logic ports for different tenants to use. That is, multiple tenants can access the network through one physical port. For each tenant, a logic port functions as an independent and real port.
- Logic Switch
  - Logic switches connect different VMs and enable them to communicate with each other at Layer 2.
  - One network device can be virtualized into multiple logic switches for different tenants to use. That is, multiple tenants can share one network device. For each tenant, a logic switch functions as an independent and real switch, and has independent software and hardware resources and running space. Services on different logic switches do not affect each other.
  - Logic switches provide L2 switching services among logic ports.
- Logic Router
  - Logic routers are virtualized from network devices running virtualization software (for example, the Virtual System on CE series switches). Logic routers connect VMs on different network segments and enable them to communicate with each other.
  - One network device can be virtualized into multiple logic routers for different tenants to use. That is, multiple tenants can share one network device. For each tenant, a logic router functions as an independent and real router, and has independent software and hardware resources and running space. Services on different logic routers do not affect each other.
  - One VPC has only one logic router. Logic routers provide L3 routing services and gateway services among logic ports.

VRF

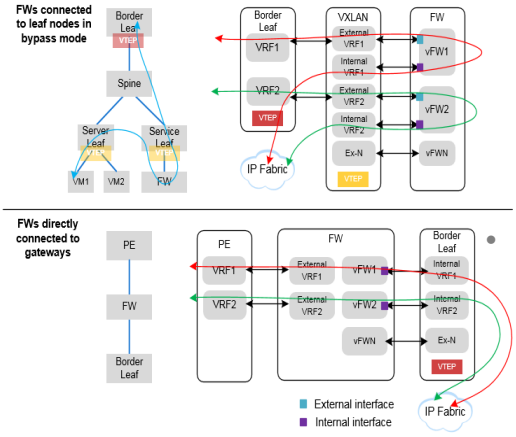

- Virtual Routing and Forwarding (VRF) allows multiple independent instances to be saved in the routing table of a router. Different instances can use the same IP address without conflicts occurring between addresses.
- When there are multiple tenants, the Agile Controller-DCN implements security isolation based on VRFs on VXLAN gateways and firewalls to allow tenants to communicate with each other or access external networks independently.

Contents

1. SDN Basic Consepts
2. **CloudFabric Application Scenarios**
   - Network Virtualization - Hosting
   - Network Virtualization - Computing
   - Cloud-Network Integration
3. Deployment Procedures
4. Service Provisioning Procedure
5. Typical Deployment

Network Virtualization - Hosting

- In the network virtualization – hosting scenario, the Agile Controller-DCN is used to realize service provisioning without a cloud platform. In this scenario, the provider (usually the carrier) leases equipment rooms and cabinet spaces to tenants and offers the following types of services:
  - Basic services: basic network access services including cabinet space/rack leasing, access bandwidth, and outbound bandwidth
  - VASs: L4-L7 VASs including security service, load balancing, virtual private network (VPN), and network address translation (NAT)
- The network virtualization – hosting scenario is further divided into two application scenarios based on the gateway provider:
  - Provided by tenant: Hosted devices include servers, Layer 2 switches, gateways, and firewalls. Tenant devices connect to the carrier network in Layer 3 mode.
  - Provided by carrier: Hosted devices include servers and Layer 2 switches, but not gateways and firewalls. Tenant devices connect to the carrier network in Layer 2 mode. Gateways and VASs are provided by the carrier.

- Racking leasing is the main service of an Internet Data Center (IDC) carrier. A tenant leases resources, including equipment room spaces, racks, ingress and egress bandwidths, security protection, and load balancing, from a carrier, so that the tenant can connect its services to networks and provide services for external users without building its own data equipment room.

Typical Deployment

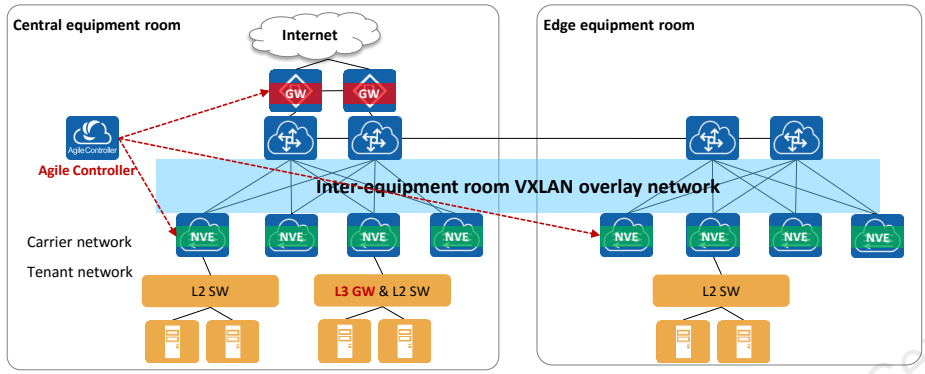

- Provide rack hosting through VXLAN networks and support multi-tenant.
- Decouple services and physical networks through VXLAN so that physical networks do not need to be changed when services are changed. This flexibly expands Layer 2 networks and facilitates network automation.
- Use VXLAN to interconnect equipment rooms, without occupying fiber resources. This consolidates rack resources of multiple equipment rooms.

Contents

1. SDN Basic Consepts
2. **CloudFabric Application Scenarios**
   - Network Virtualization - Hosting
   - Network Virtualization - Computing
   - Cloud-Network Integration
3. Deployment Procedures
4. Service Provisioning Procedure
5. Typical Deployment

Huawei Training & Certification

Network Virtualization - Computing

- Network virtualization – computing scenario realizes service provisioning without a cloud platform. The Agile Controller-DCN is responsible for service delivery.
- Service provisioning consists of two parts:
  - Network service provisioning: Administrators use the Agile Controller-DCN to allocate network resources to specified services or applications. The Agile Controller-DCN automatically delivers service configurations and VMM configurations to overlay networks.
  - Computing service provisioning: Administrators create, delete, and migrate computing and storage resources on the VMM. The Agile Controller-DCN can automatically detect operation changes on the VMM.

- If a unified cloud platform cannot be built for a customer due to the complexity of the computing service management system or inadequate convergence between computing management and network management, the network virtualization – computing scenario is recommended. In this scenario, the Agile Controller-DCN interconnects with a computing virtualization platform to implement collaborative provisioning of computing and network resources and network configuration automation. Network virtualization – computing is applicable for construction of non-cloudified data centers.

Typical Deployment

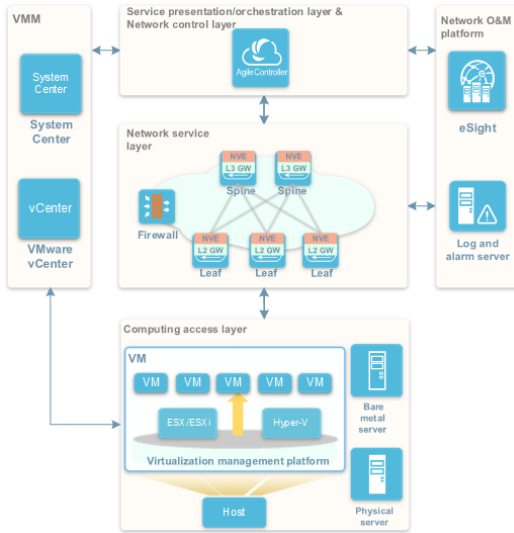

- Service presentation layer
- Network control layer
- Network service layerComputing access layer
- Computing access layer
  - Virtualization server
  - Physical server

- Service presentation layer: This is oriented at data center users. The AC-DN at this layer provides GUIs for service, network, and tenant administrators, implementing service orchestration, policy delivery, automatic deployment, and O&M.
- Network control layer: This layer is the core in the solution. The Agile Controller-DCN implements network modeling and instantiation, works with physical and virtual networks to provide network resource pooling and automation, and provides the view of the entire network to uniformly control and deliver service flow tables. The Agile Controller-DCN is the key component for separating SDN network control and forwarding.
- Network service layer: This layer is the infrastructure of a data center network that provides high-speed channels for services, for example, L2-L3 basic network services and L4-L7 VASs.
- Computing access layer: This layer provides access to virtualization servers and physical servers.
  - Virtualization server: Virtualization technology is used to virtualize one physical server into multiple VMs and vSwitches. VMs connect to the fabric network through vSwitches. This solution is compatible with mainstream server virtualization products including VMware ESXi and Microsoft Hyper-V.
  - Physical server: Physical servers connect to the fabric network through logical ports.

Contents

1. SDN Basic Consepts
2. **CloudFabric Application Scenarios**
   - Network Virtualization - Hosting
   - Network Virtualization - Computing
   - Cloud-Network Integration
3. Deployment Procedures
4. Service Provisioning Procedure
5. Typical Deployment

Cloud-Network Integration - FusionSphere

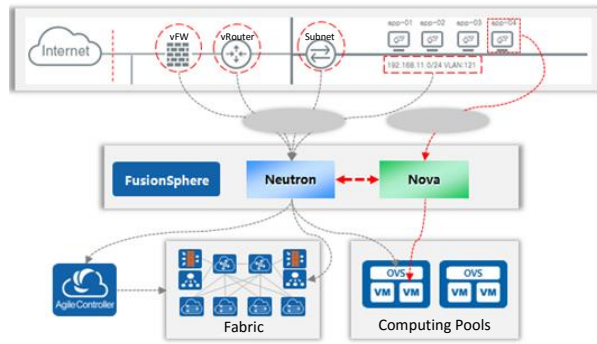

- In the cloud-network integration , cloud platform is used to uniformly provision computing and network resources
- Cloud platform provides graphical user interfaces (GUIs) for unified management of computing and network resources.
- The Agile Controller-DCN interconnects with cloud platform to implement flexible resource management, service provisioning, and service migration.

Typical Deployment

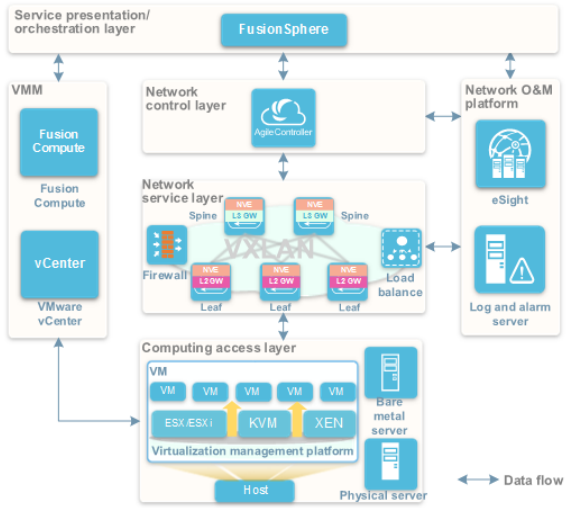

- Service presentation/orchestration layer
- Service presentation layer
- Service orchestration layer
- Network control layer.
- Network service layerComputing access layer
- Computing access layer
  - Virtualization server
  - Physical server
  - Bare metal server

- Service presentation/orchestration layer
  - Service presentation layer: is oriented at data center users. The cloud platform at this layer provides GUIs for service, network, and tenant administrators, implementing service management, automatic service provisioning, as well as resource and service guarantee.
  - Service orchestration layer: consists of Nova, Neutron, and Cinder components of the cloud platform. These components interoperate to realize virtualization, pooling, and collaborative scheduling of computing, storage, and network resources in data centers.
- Network control layer: This layer is the core in the solution. The Agile Controller-DCN implements network modeling and instantiation, works with physical and virtual networks to provide network resource pooling and automation, and provides the view of the entire network to uniformly control and deliver service flow tables. The Agile Controller-DCN is the key component for separating SDN network control and forwarding.
- Network service layer: This layer is the infrastructure of a data center network that provides high-speed channels for services, for example, L2-L3 basic network services and L4-L7 VASs.
- Computing access layer: This layer provides access to virtualization servers, physical servers, and bare metal servers.
  - Virtualization server: Virtualization technology is used to virtualize one physical server into multiple VMs and vSwitches. VMs connect to the fabric network through vSwitches. This solution is compatible with mainstream server virtualization products including Kernel-based Virtual Machine (KVM), XEN, and VMware ESXi.
  - Physical server: Physical servers connect to the fabric network through logical ports.
  - Bare metal server: FusionSphere directly interacts with the bare metal servers to manage them as instances.

Cloud-Network Integration - OpenStack

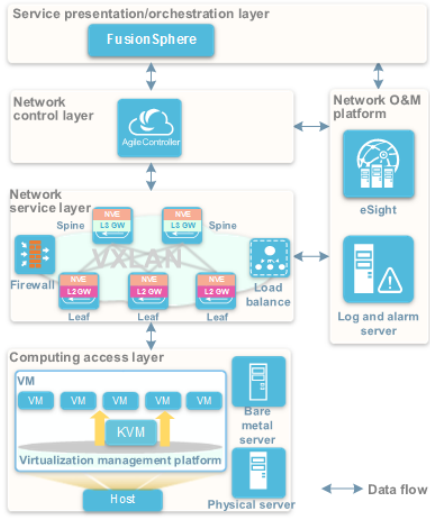

- Agile Controller-DCN is controlled by native OpenStack

Contents

1. SDN Basic Consepts
2. CloudFabric Application Scenarios
3. **Deployment Procedures**
4. Service Provisioning Procedure
5. Typical Deployment

Integrated Service Operations in the Network Virtualization - Hosting Scenario

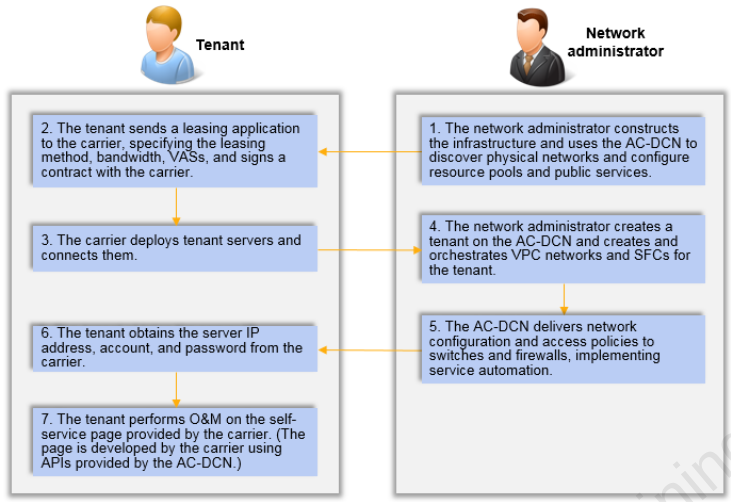

Deployment Procedures in Hosting Scenarios

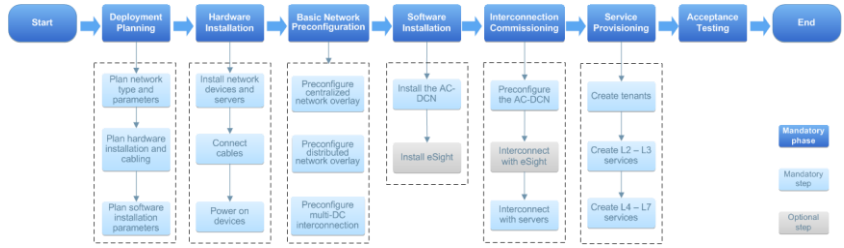

- In hosting scenarios, the computing virtualization platform is deployed by users and does not need to be interconnected with the Agile Controller-DCN

Collaborative Provisioning of Computing and Cetwork Services

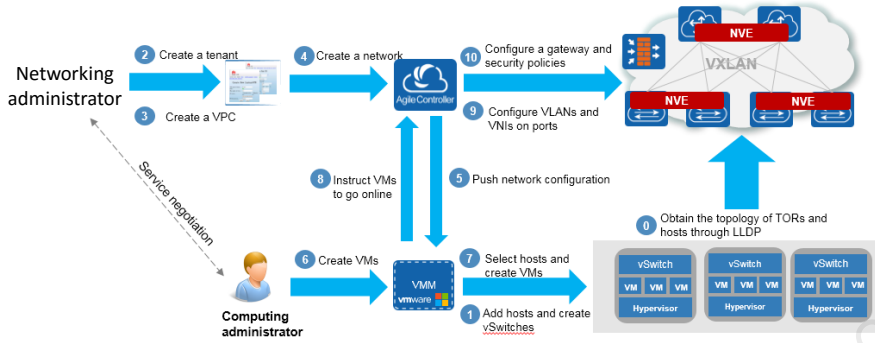

- Network service provisioning: Administrators use the Agile Controller-DCN to allocate network resources to specified services or applications. The Agile Controller-DCN automatically delivers service configurations and VMM configurations to overlay networks.
- Computing service provisioning: Administrators create, delete, and migrate computing and storage resources on the VMM. The Agile Controller-DCN can automatically detect operation changes on the VMM

Deployment Procedures in Computing Scenarios

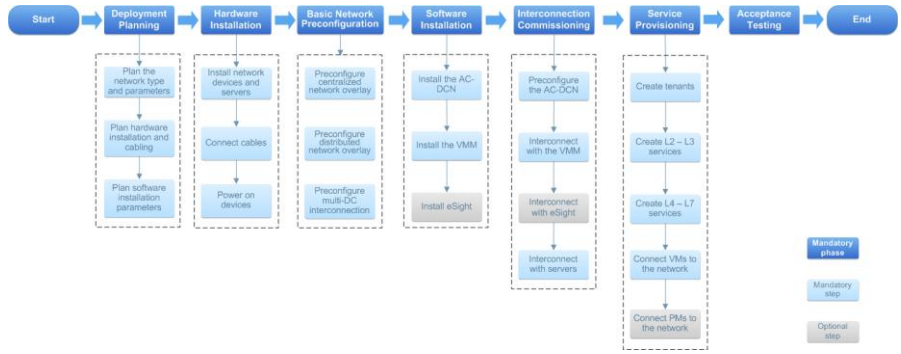

- In computing scenarios, the Agile Controller-DCN supports interconnection with the vCenter and the System Center
- Agile Controller-DCN keeps heartbeat line with VMM.

Cloud-Network Integration - FusionSphere

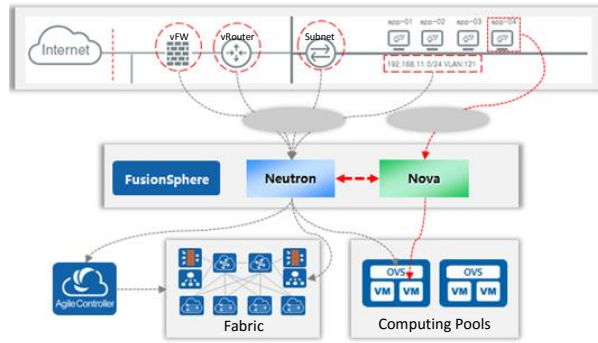

- A cloud platform centrally manages data center service resources and provides a user interface (UI) for configuring services. It connects to northbound Application Programming Interfaces (APIs) of Huawei's Agile Controller and delivers service configurations to the Agile Controller for orchestration and deployment. The cloud platform and Agile Controller work together to schedule computing, storage, and network resources.
- Some customers may develop their own application portals and connect them to cloud platforms to manage resources, deliver services, and monitor the services.

Deployment Procedures in FusionSphere Scenarios

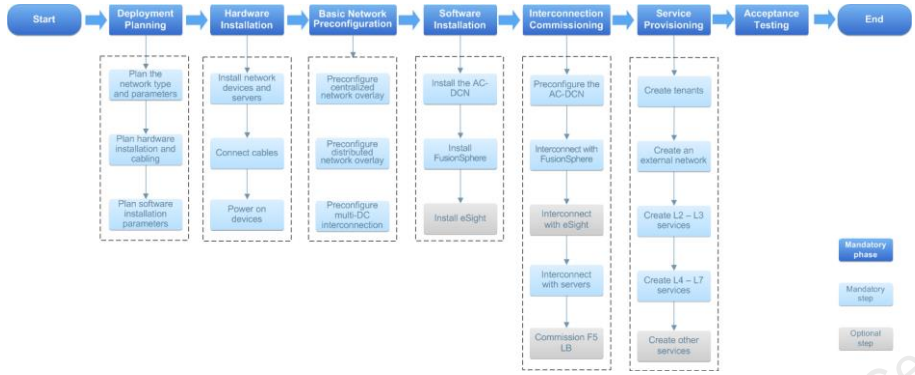

- Common cloud platforms include the open-source OpenStack, Huawei's FusionSphere, HP's Helion, Red Hat's OpenStack, and Mirantis.

Deployment Procedures in OpenStack Scenarios

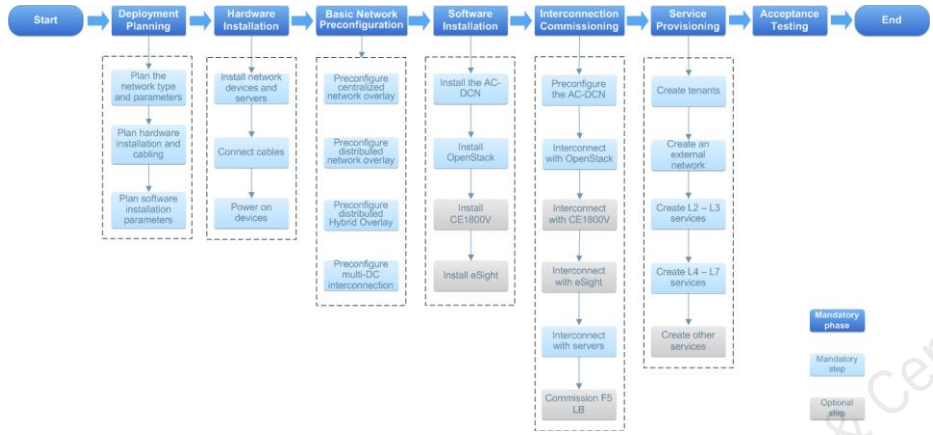

HUAWEI

Contents

1. SDN Basic Consepts
2. CloudFabric Application Scenarios
3. Deployment Procedures
4. **Service Provisioning Procedure**
   - Network Virtualization
   - Cloud-Network Integration
5. Typical Deployment

Network Model Mapping

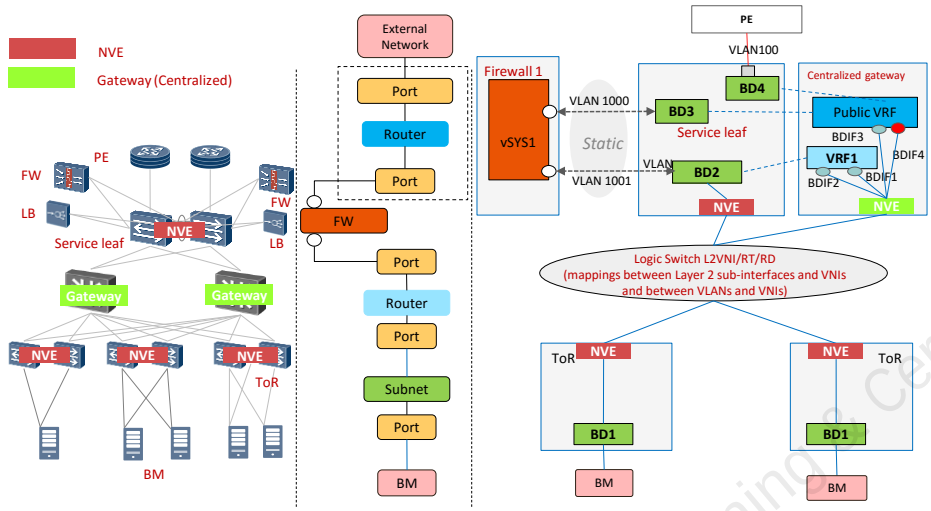

Service Provisioning Process

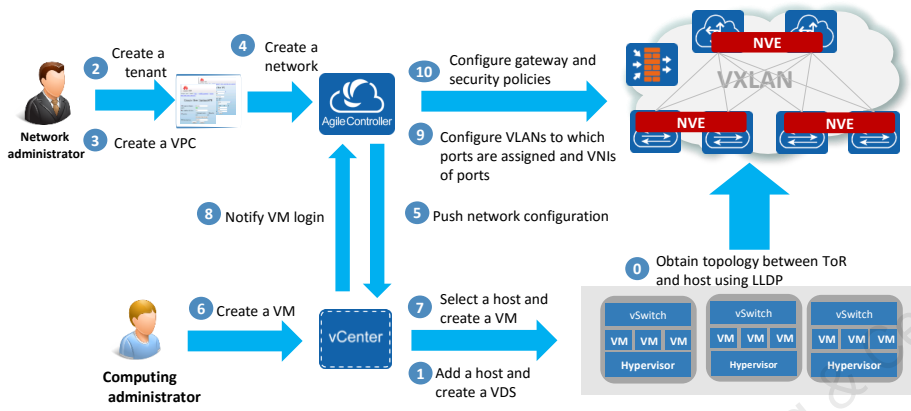

(Note: This example uses vCenter to describe the service provisioning process, which is similar with System Center.)

Network Resource Provisioning Process

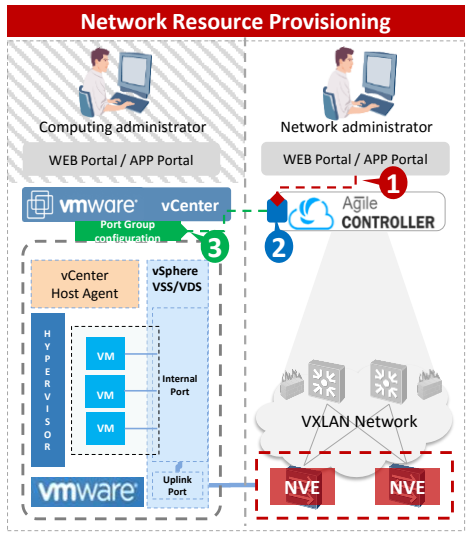

- The network administrator first delivers the logical network required by services on the Agile Controller UI and notifies the Agile Controller through RESTful interfaces.
- The Agile Controller assigns VNIs of the overlay network based on configuration and assigns Local VLANs used by the virtual network based on the P+V or 1:1 rules preset during POD creation and saves mappings between VLANs and VNIs.
- The Agile Controller interoperates with vCenter through the WebService interface, creates the Port Group required by the Local Network on the VDS, and binds Local VLAN to Port Group.
- Note: The network administrator can name the service network on the Agile Controller UI. When the Agile Controller pushes the Port Group configuration to the VDS, it can configure the network name to the Port Group so that the computing administrator can bind VMs to the network.

Computing Resource Provisioning Process

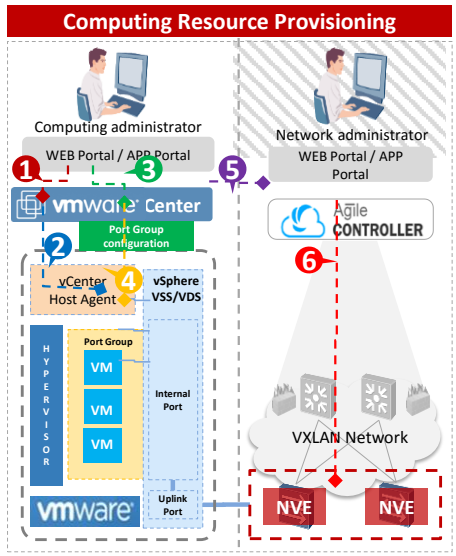

① The computing administrator sets the computing resource quota (vCPU, RAM, ROM, and operating system) based on service requirements on vCenter.

② The computing administrator creates a VM. vCenter evenly schedules computing resources dynamically based on configuration, selects an ESXi host, and loads the VM to the host based on configuration.

③ The computing administrator can view the Port Group configuration delivered by the network administrator on vCenter and manually binds the VM to Port Group.

④ vCenter pushes the Port Group configuration to the ESXi host and binds the Port Group to the VM.

⑤ The Agile Controller detects the VM login and Port Group binding event through the WebService interface and obtains the VM login location.

⑥ The Agile Controller examines the database for indexes and obtains the mapping between Local VLAN and VNI. It queries LLDP neighbor information to obtain the connection between the ESXi and ToR ports and delivers the mapping between Local VLAN and VNI to ToR ports through the NETCONF interface.

Device Discovery

- Auto discover devices

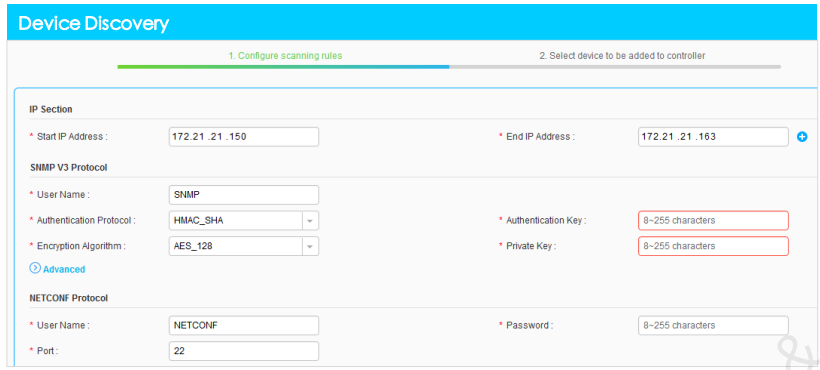

Link Discover

- Support inter-device links and host access links auto discovery by LLDP and manually adding

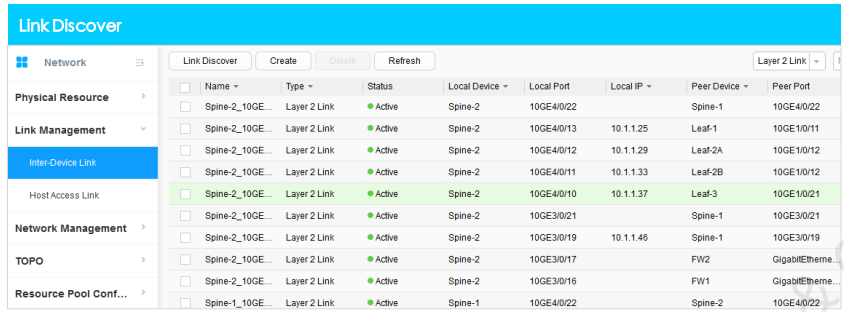

Create Fabric

- Create a Fabric and configure the basic information and network features

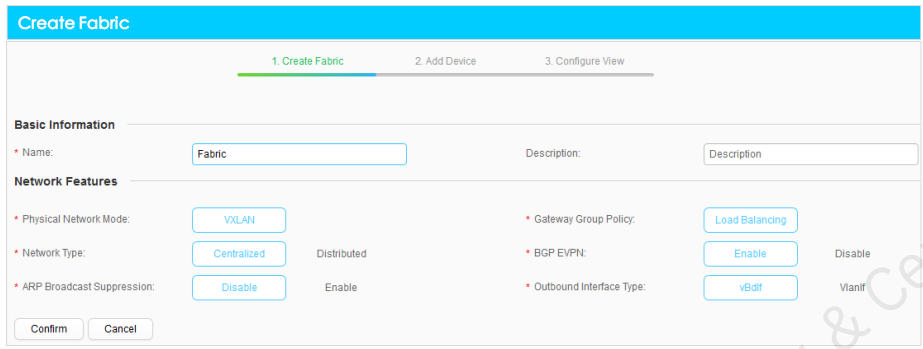

Create L4 - L7 Pool

- Create a L4-L7 pool. Configure the basic information and network features.

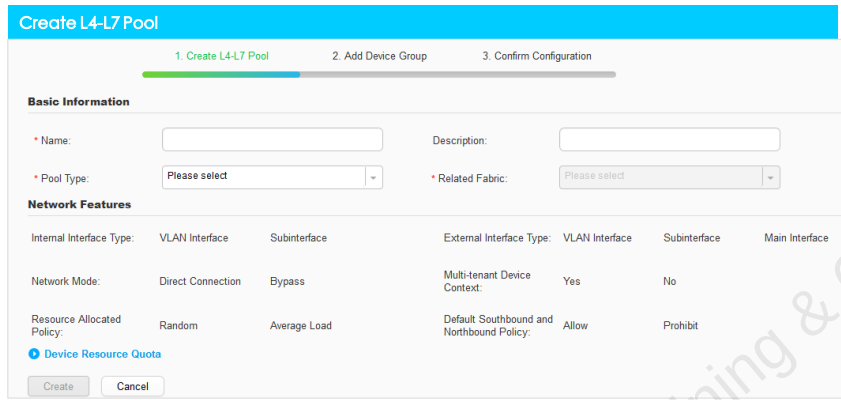

Configure Interconnection Resource

- Configure VLAN and IP segments which are used between gateways and firewalls.

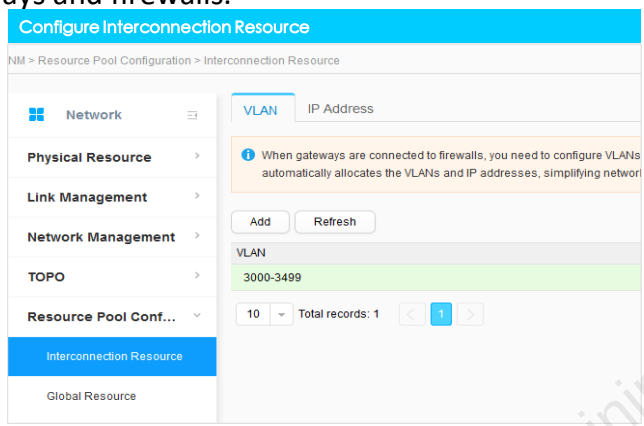

Configure Global Resources

- Configure BD, VLAN, VNI, Public IP Addresses and so on which are used when Agile Controller-DCN delivery services.

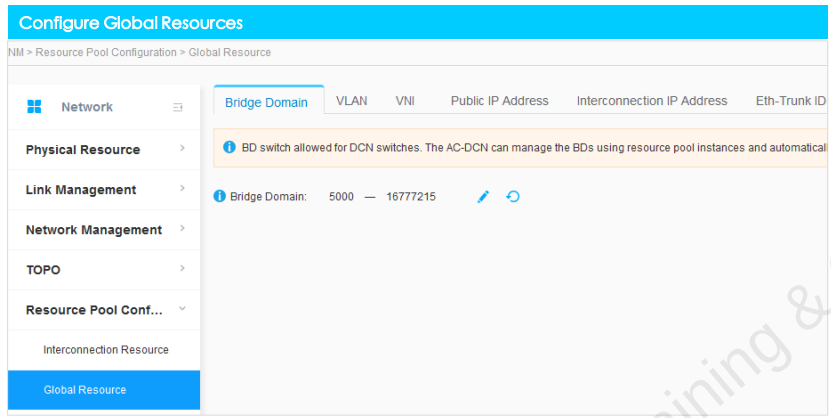

Configure DHCP Service

- Configure DHCP server IP address and parameters.

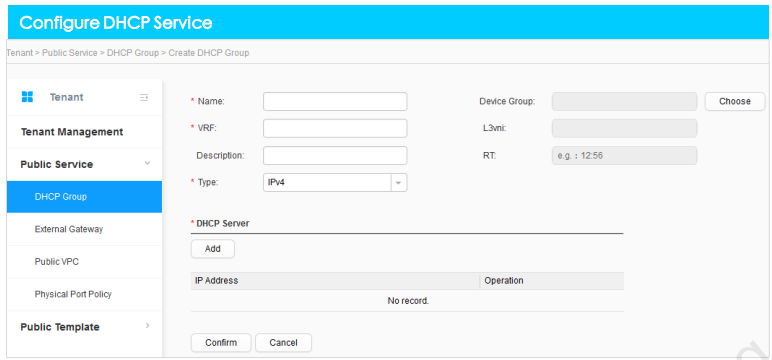

Configure External Gateway

- Configure parameters of external gateway.

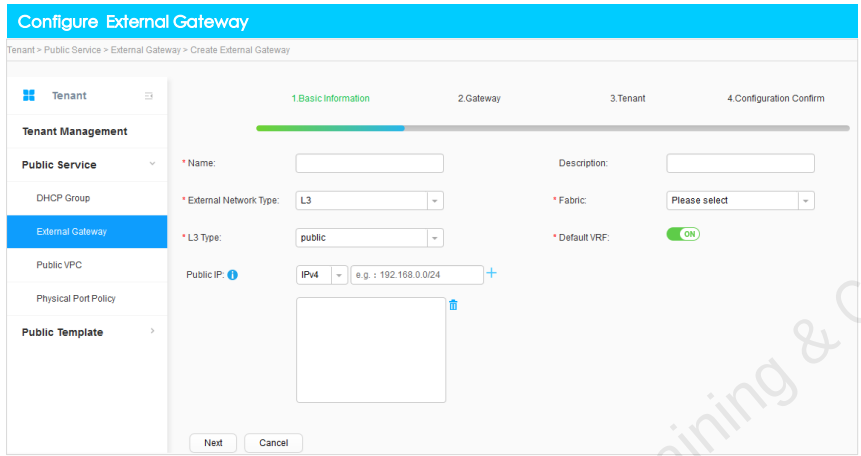

Connect to VMM

- Create VMM which Agile Controller-DCN connects to.

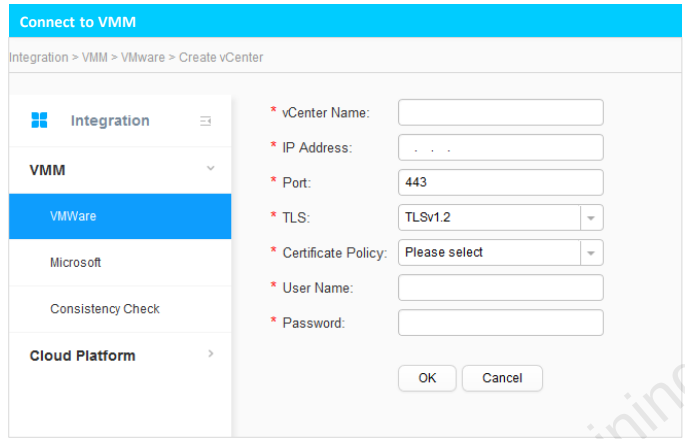

Create Tenants

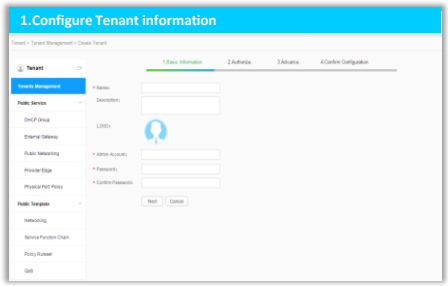

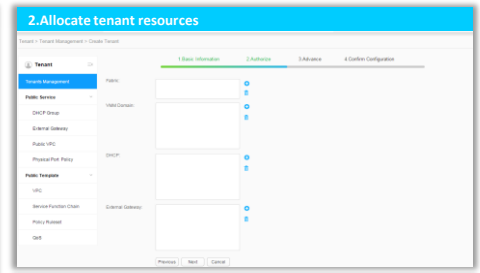

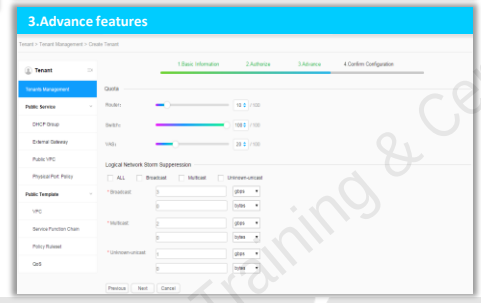

- Create a tenant and allocate the network resources.

Create Public Templates

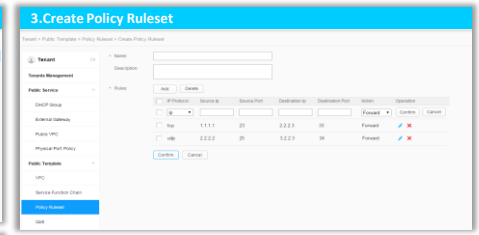

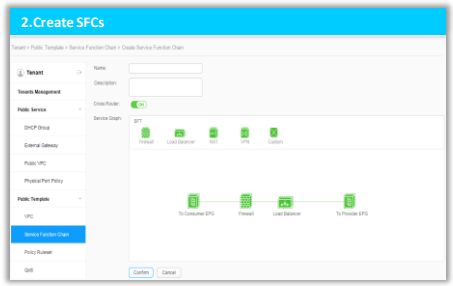

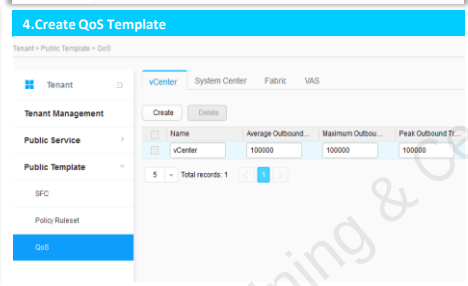

- Create a VPC network profile, service link profile, policy profile, and QoS profile to facilitate service delivery

Create VPC

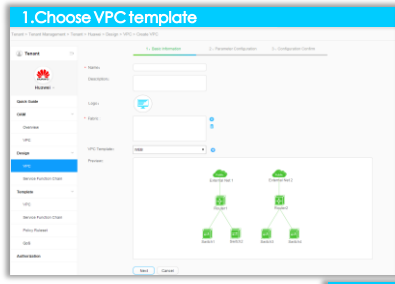

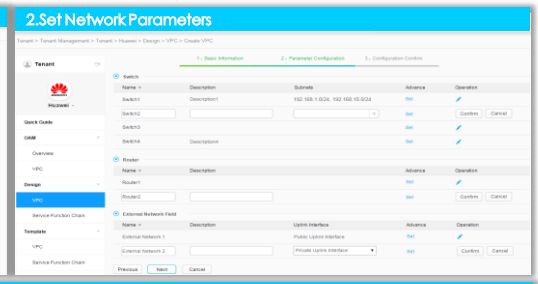

- Choose a VPC template. Configure network parameters and use NE icons to create logical network.

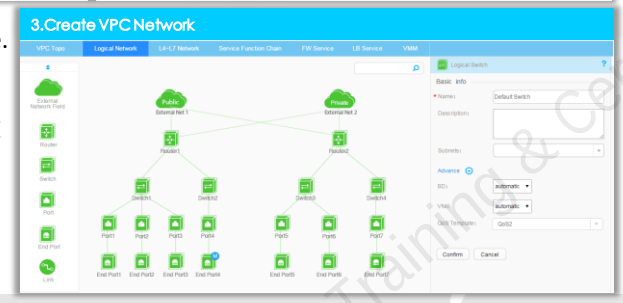

Configure SFCs

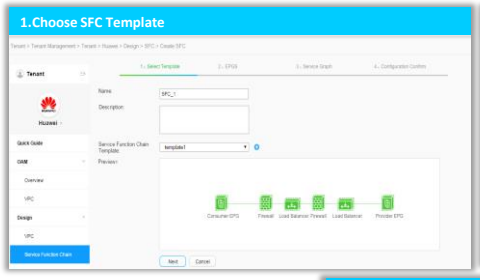

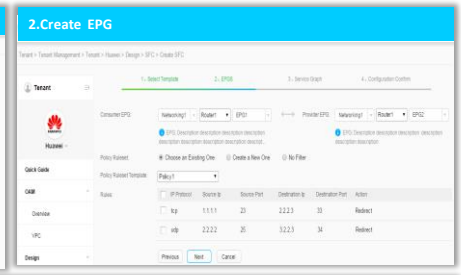

- Choose a SFC template. Configure the policy ruleset for the Service Function Chain.

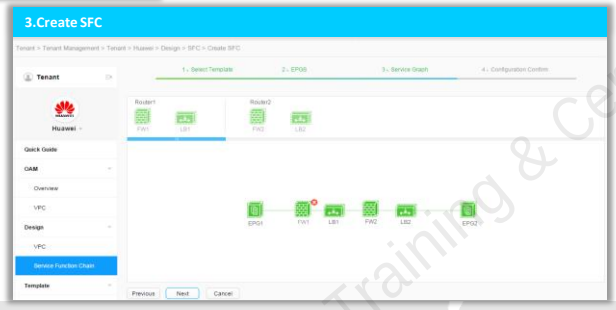

Topology

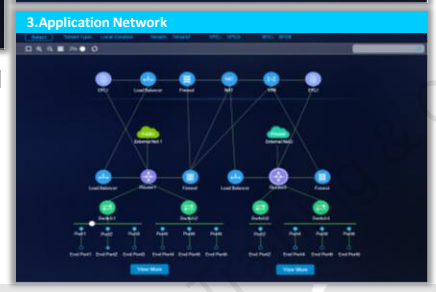

- The application, logical, and physical network topologies are mutually visible. Phycial network, Logical Network and Application Network.

Monitor

- The Agile Controller-DCN can collect and display statistics about flow performance of logical switches and L4-L7 network service. Users can query the current and historical performance data.

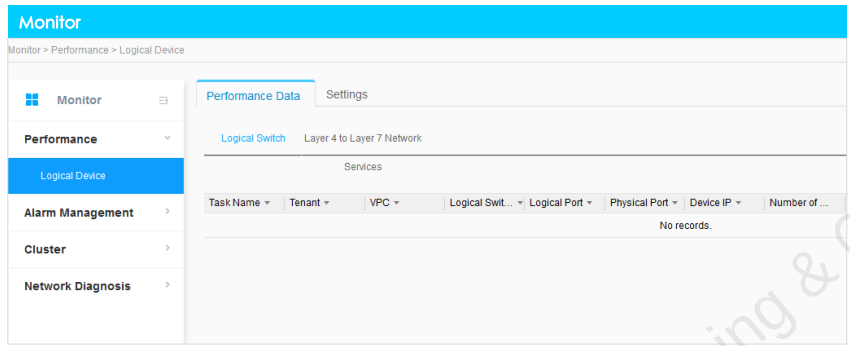

Alarm Management

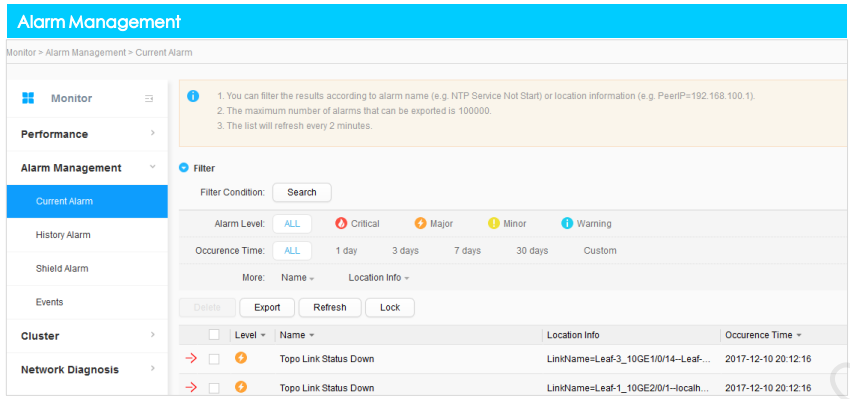

- Monitor alarm information in real time and troubleshoot faults based on the alarm information and handling suggestions in a timely manner, ensuring normal operation of services

Service Path Detection

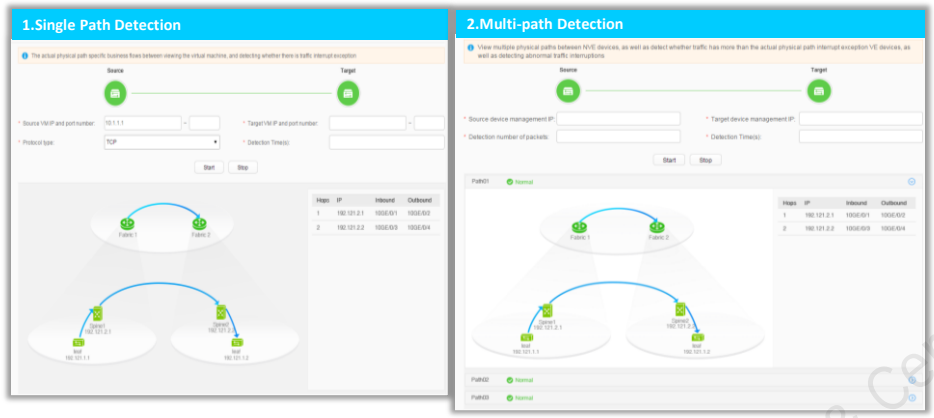

- Support single path detection, multi-path detection, loop detection and Ping to help maintenance.

Tenant Resource Statistics

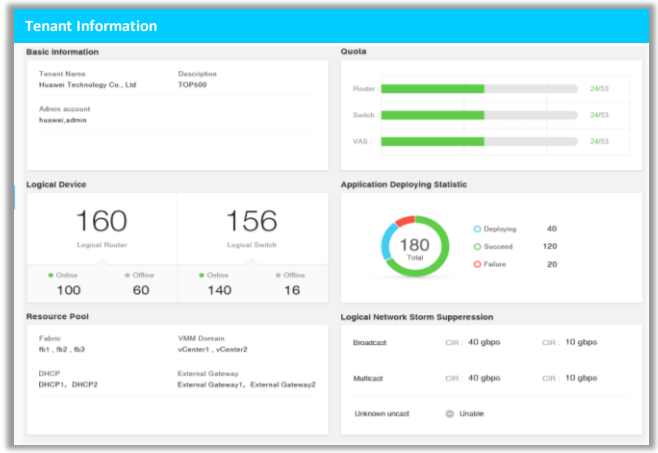

- Tenant resource statistics show in the overview interface.

VPC

- Tenant administrator can see VPCs, including logic networks and service function chains.

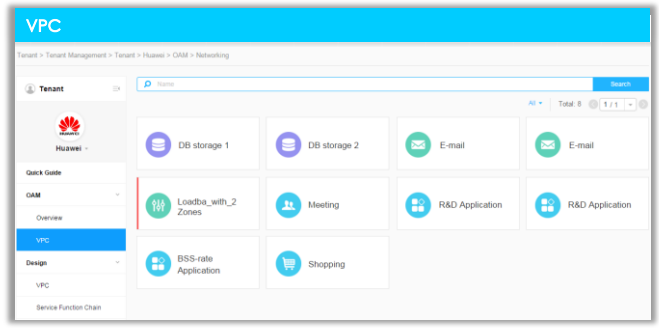

Contents

1. SDN Basic Consepts
2. CloudFabric Application Scenarios
3. Deployment Procedures
4. **Service Provisioning Procedure**
   - Network Virtualization
   - Cloud-Network Integration
5. Typical Deployment

HUAWEI

Huawei FusionSphere OpenStack Service Model

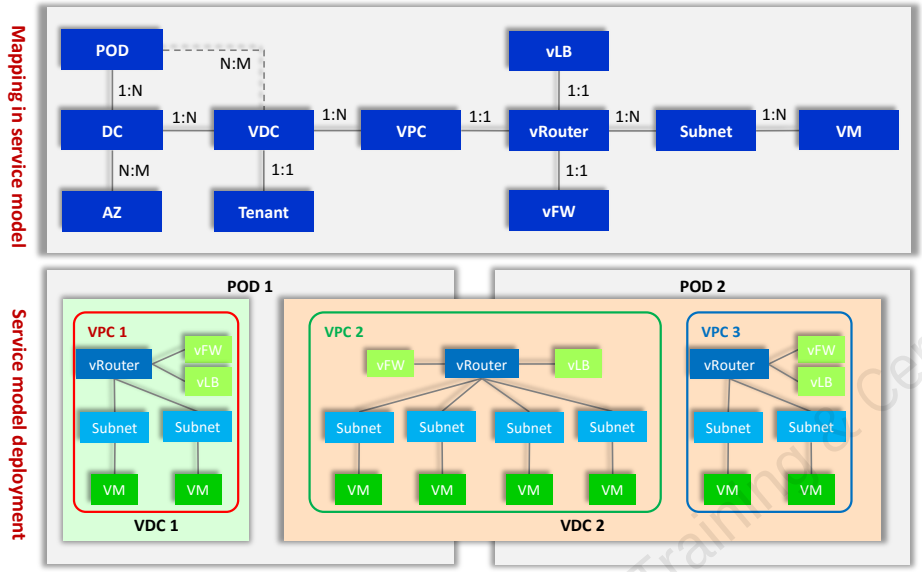

HUAWEI

- POD is a physical unit. For example, a fabric network can be considered a POD. A POD is a basic service deployment unit. One set of resources can be deployed in one or more PODs, and one POD can accommodate multiple sets of resources.
- VDC is a resource unit and can span multiple PODs. Tenants rent resources with granularity of VDC.
- VPC is a service unit and can span multiple PODs. A tenant can have multiple VPCs deployed in different PODs.

Cloud Service Operation Workflow

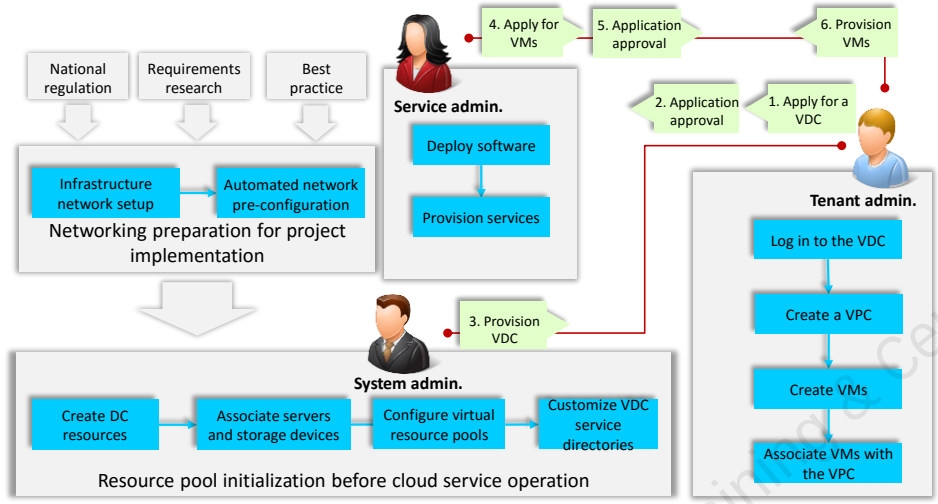

Cloud Service Provisioning Process

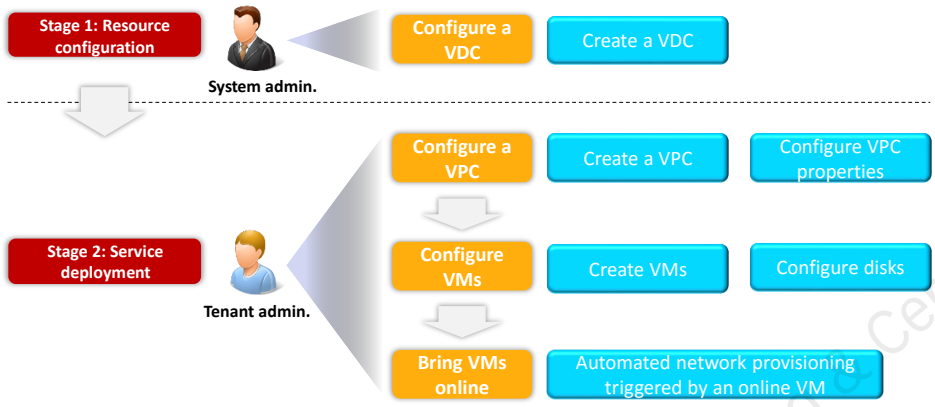

(Note: FusionSphere is used as an example. The operations vary on different cloud platforms)

Create VDC

Sys admin

1. Set VDC name

2. Choose VDC AZs

3. Set VDC Quota

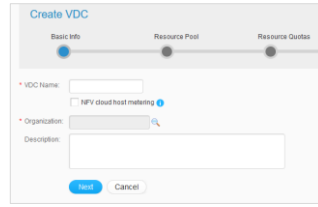

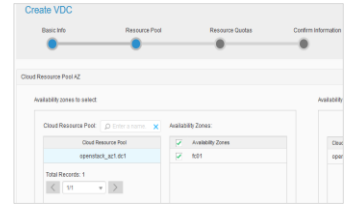

Create VPC

Tenant admin

1. Set VPC name

2. Apply for vRouter

3. Create Subnet and bind vRouter

4. Configure vFW

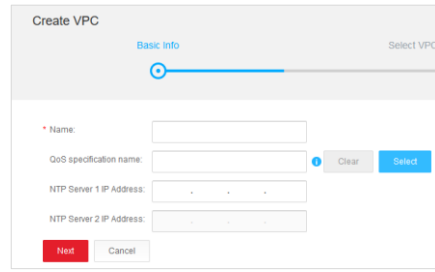

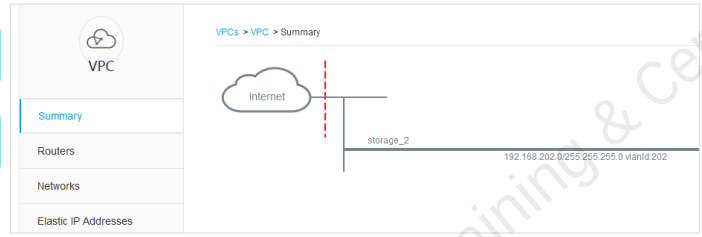

Configure VPC Parameters

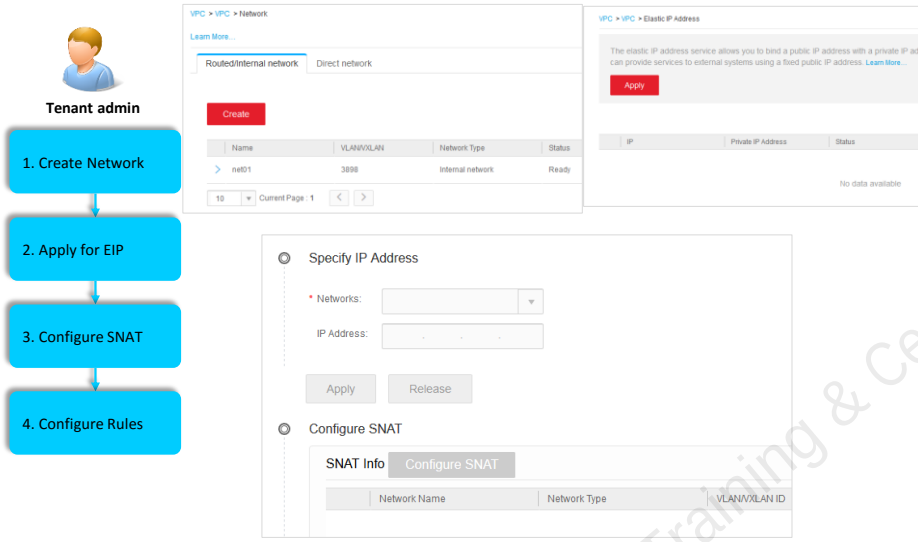

Create VM

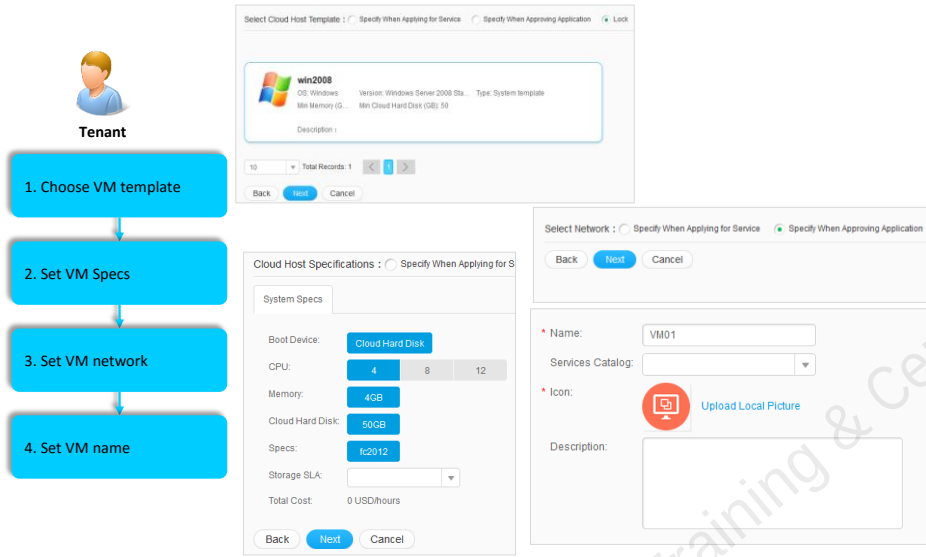

Configure Disk

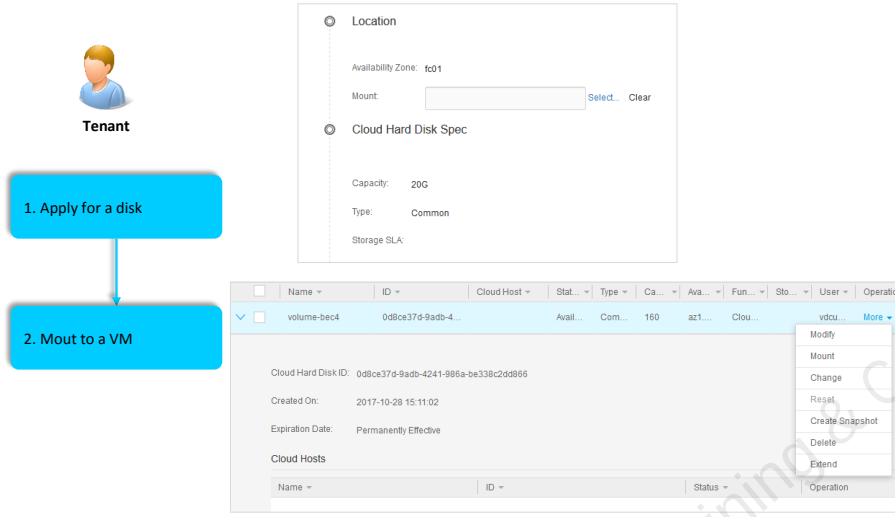

Automated Network Provisioning Triggered by an Online VM

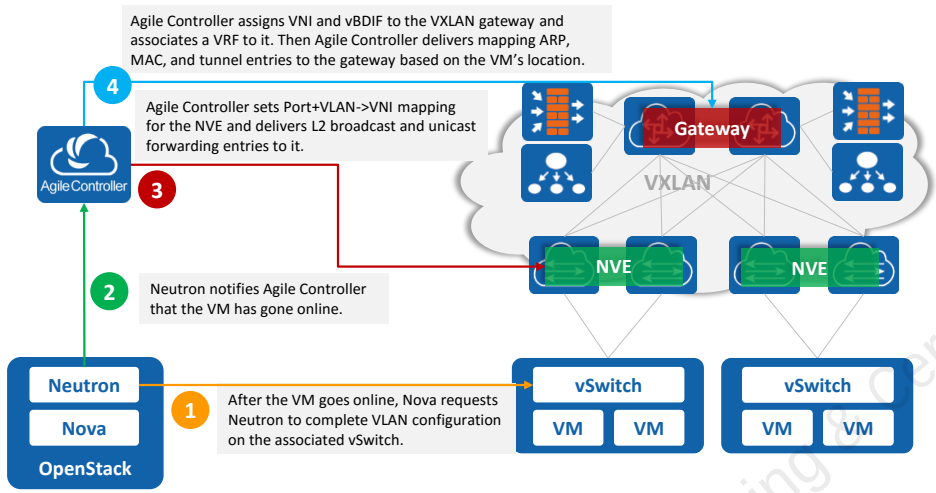

(Note: This example shows the network provisioning process in a network overlay scenario where Neutron manages OVSs directly.)

Contents

1. SDN Basic Consepts
2. CloudFabric Application Scenarios
3. Deployment Procedures
4. Service Provisioning Procedure
5. **Typical Deployment**

Typical Scenario

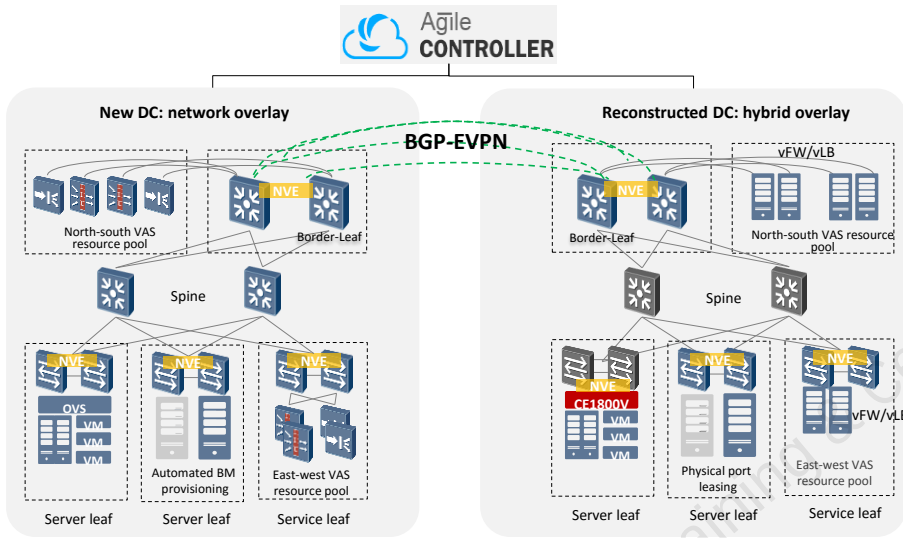

HUAWEI

- Converged network resource pool
  - One suite of Agile Controller to manage multiple types of fabric networks (centralized network overlay, distributed network overlay, hybrid overlay)
  - Stretched controller cluster to manage multiple DCs (telecommunications equipment room consolidation, RTT<10ms, active-standby/independent egresses)
  - Hybrid overlay solution compatible with third-party switches
  - Support for diverse computing resources, including VMs, physical servers, bare metal servers, and ports
- Elastic VAS resource pool
  - Elastic north-south/east-west VAS resource pools, which can be independently deployed or combined
  - Third-party VAS devices supported in service chains
  - Huawei NGFW, vNGFW, and vLB managed by Agile Controller, F5 LBs managed by the OpenStack platform
  - Two service chain modes supported: PBR and NSH

Quiz

1. What are Network Virtualization and Cloud-Network Integration?
2. Which of following is right? ( )
   A. Tenant can share a physical port
   B. A VPC can contains several Logic Router
   C. VPC must have vFW
   D. Logic Switch provide L2 switching services

- 参考答案:
  1. Different CloudFabric application scenarios.
  2. AD

Summary

- SDN Basic Consepts
- CloudFabric Application Scenarios
- Deployment Procedures
- Service Provisioning Procedure
- Typical Deployment

More Information

- CloudFabric DCN Solution
  - http://support.huawei.com/onlinetoolsweb/NetSolution/DataCenterNetwork/en/index.html

Recommendations

- Huawei Learning Website
  - http://support.huawei.com/learning/en/newindex.html
- Huawei Support Knowledge Base
  - http://support.huawei.com/enterprise/servicecenter?lang=en

Thank You

www.huawei.com

Cloud DC eSight ICT Management System Principle and Deployment

www.huawei.com

Foreword

- eSight is a next-generation network management system (NMS) for campus and branch networks of enterprises. It uniformly manages enterprise resources, services, and users and implements intelligent association. eSight supports unified management of IT&IP and non-Huawei devices. In addition, eSight provides a flexible and open platform for enterprises to build their own intelligent management systems.
- This course introduces the positioning, architecture, new management functions, and system capabilities of eSight. In addition, this course also describes the management features and evolution, function changes, and how to install eSight on the Windows Server 2008 R2.

Objectives

- After completing this course, you will be able to understand:
  - Positioning, architecture, and functions of eSight
  - New management functions, evolutions, and changes of eSight
  - Major reference documents and tools for network O&M using eSight
  - Help channels for eSight
  - eSight details
  - Procedures for eSight installation and uninstallation
  - License function and application method of eSight

Contents

1. **eSight Positioning**
2. Main eSight Services
3. eSight Architecture
4. Typical Applications of eSight
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. eSight Installation Process

Product Positioning

- eSight is an ICT full life cycle management system built for enterprises.

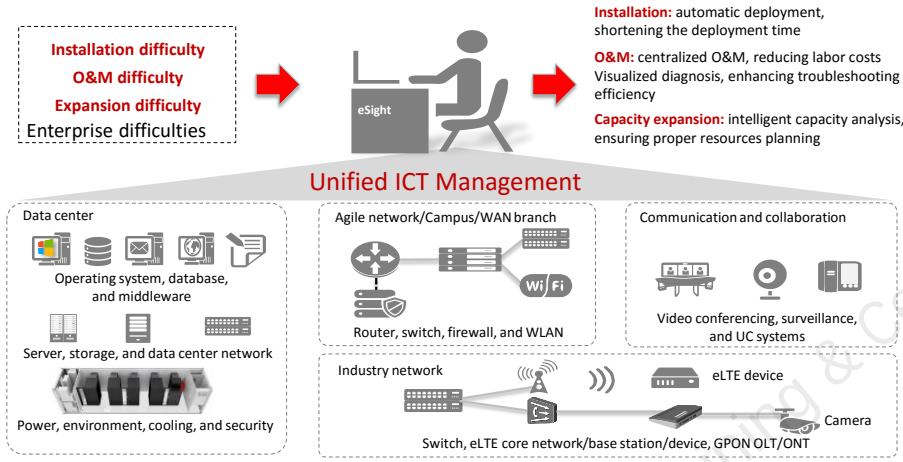

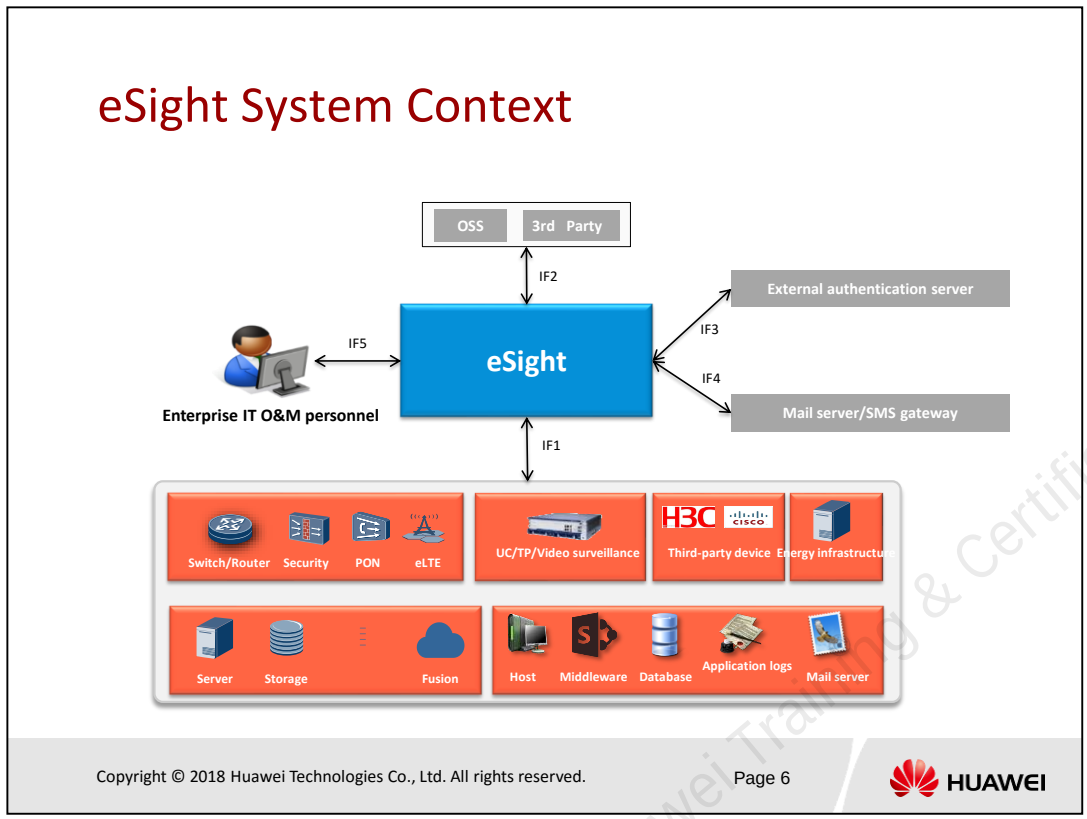

- eSight external interfaces and users are described as follows:
  - IF1: southbound interface used to interact with managed objects, including data collection and command delivery. eSight supports the following southbound interface types: SNMP, Telnet/STelnet, FTP/SFTP/FTPS, TR-069, and Huawei Man-Machine Language (MML).
  - IF2: northbound interface used to integrate with upper-layer management systems. Third-party systems can obtain network resources, alarms, and performance data from eSight through the SNMP, HTTP, or FTP interfaces.
  - IF3: interface used to integrate with external authentication servers. eSight can interconnect with authentication servers deployed by enterprises themselves through LDAP and RADIUS.
  - IF4: message notification interface. eSight can interconnect with existing mail servers and SMS gateways of enterprises so that eSight can send important alarms, events, and reports to customers through emails or SMS messages.
  - IF5: user interface, that is, the GUI. System administrators and operators use functions provided by eSight through this interface.

Contents

1. eSight Positioning
2. **Main eSight Services**
3. eSight Architecture
4. Typical Applications of eSight
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. eSight Installation Process

Unified Discovery of Multiple Types of Devices

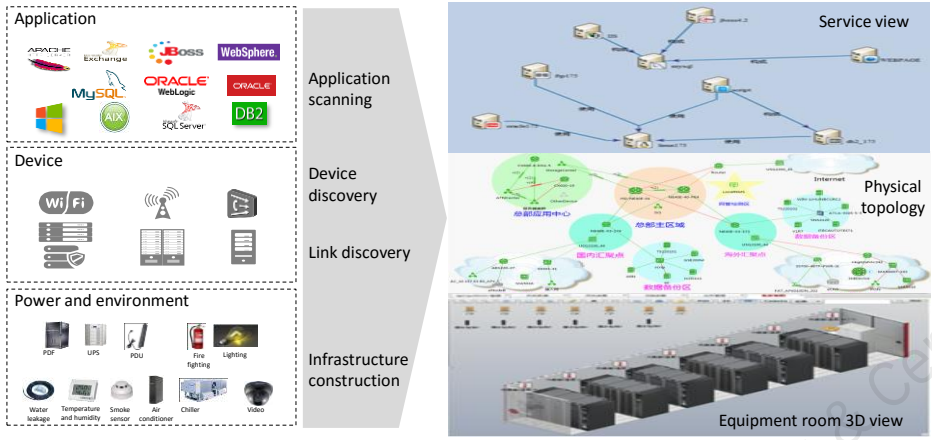

Unified discovery: switch, router, PON, WLAN, security gateway, UC, video conferencing, video surveillance, eLTE, server, virtual resource, and storage device

Batch Device Configuration Delivery: Accelerating Service Provisioning

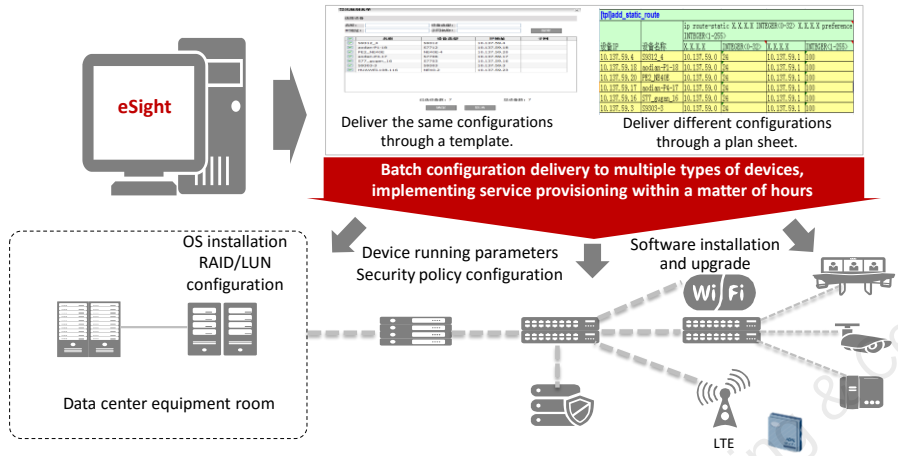

Unified Device Display: Greatly Improving O&M Efficiency and Reducing Costs

Unified alarm for all devices | Unified performance | Resource group management

Visual infrastructure view | Integrated topology monitoring

Holistic Security Management: Building a Reliable Enterprise Working Environment

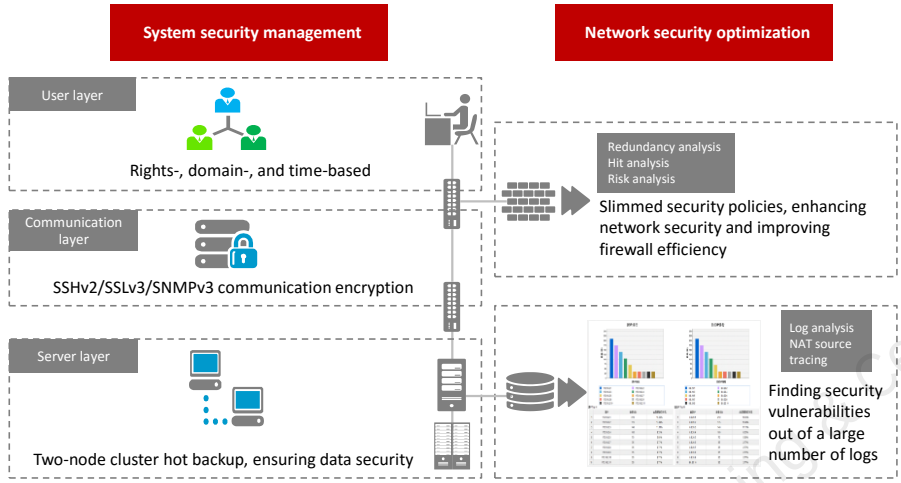

Diverse Troubleshooting Means: Ensuring Reliable Operation of ICT Systems

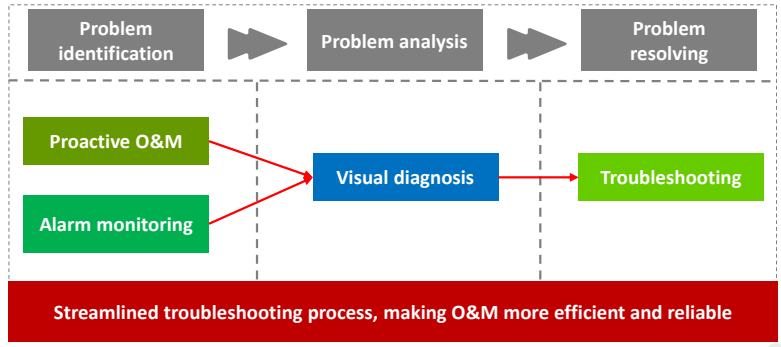

Active Hierarchical Prevention: Detecting and Rectifying Risks in Advance

Device
Analyze device performance as well as device and interface traffic to detect risky devices.

Service
Analyze host and session traffic to detect risky apps.

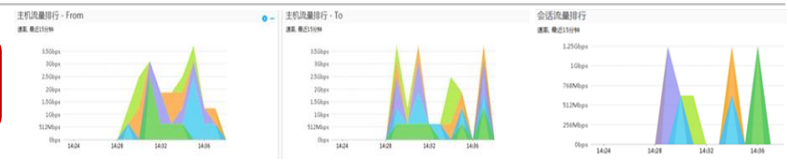

User
Audit user Internet behavior to detect risky users.

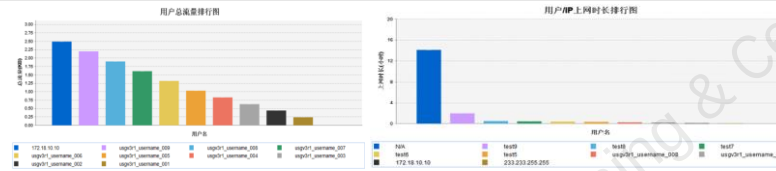

Convenient Troubleshooting: Quickly Recovering Services

Re-adjust the AP coverage scope to alleviate AP conflicts.

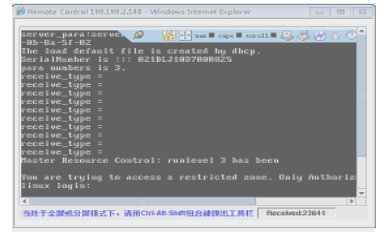

Integrate the keyboard, video and mouse (KVM) to remotely control devices.

Service restoration without going outside

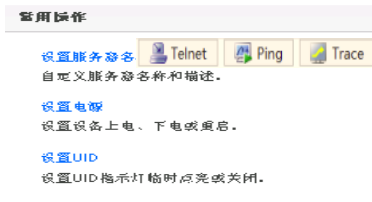

Power on/off and reset devices remotely.

The fault indicator precisely locates the frame, board, and disk to help O&M personnel quickly locate the faulty device.

Contents

1. eSight Positioning
2. Main eSight Services
3. **eSight Architecture**
4. Typical Applications of eSight
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. eSight Installation Process

eSight Solution

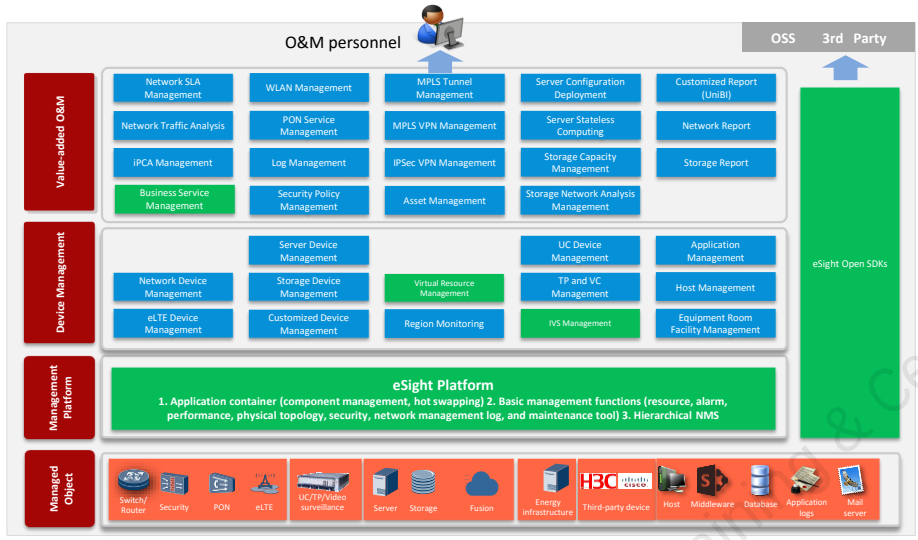

- Enhance the Intelligent Business Service Manager, Virtualization Manager, Intelligent Video Surveillance Device Manager (device inspection, VCM), capability of the Management Platform (low cost, two-node cluster), and Open SDKs (remote O&M support)

eSight Technical Architecture

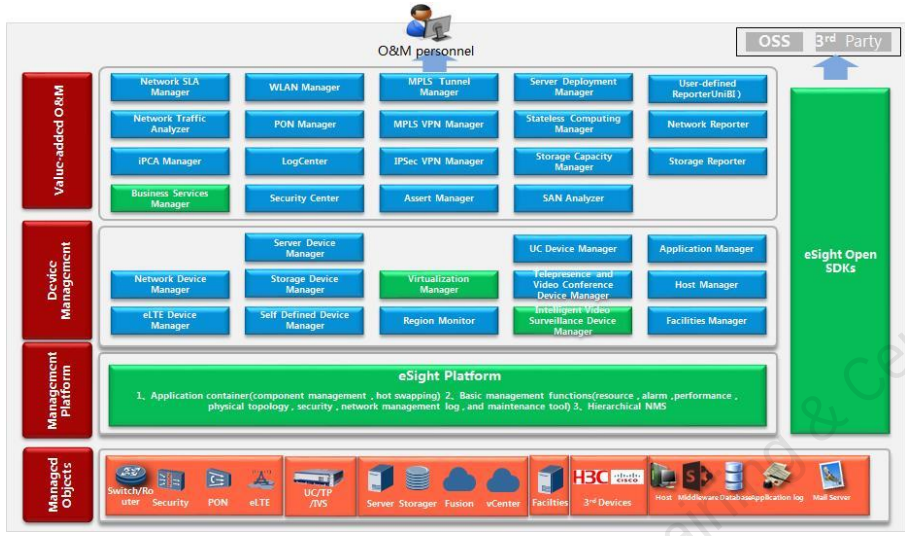

Contents

1. eSight Positioning
2. Main eSight Services
3. eSight Architecture
4. **Typical Applications of eSight**
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. eSight Installation Process

Private Cloud/DC Management

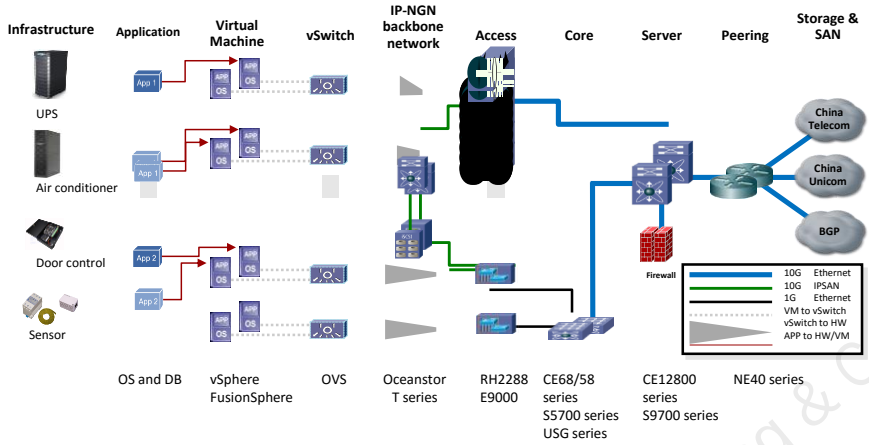

- Compared with V3R5, this version has the following changes:
  - The virtualization management capability is enhanced to uniformly manage resources, alarms, and performance of the FusionSphere.
  - The business service management streamlines the application and device connection relationships, implementing business service-centered monitoring.

Private Cloud/DC Typical Application Scenario

- According to Huawei IT O&M experience, the DC network has about 5 or 6 1/2-level faults, of which 50% are hardware faults (board or module damage). The other faults are mainly caused by software bugs. Relatively, there are more 3/4-level faults. However, these faults do not affect services.
  - -- In this situation, the management system is required to report alarms as soon as possible after a device is faulty, avoiding delay.
- During fault processing, the device and link connectivity are generally checked. After a fault occurs or a fault is rectified, the device and link connectivity need to be checked.
  - -- The topo module of the management system must be able to display the connectivity of devices and links, and display connectivity changes on the topo map.
- The performance is optimized as follows: (1) Optimize the bandwidth performance of important convergence links. (2) Optimize the performance of important single device (CPU/memory/interface traffic capability/scalability).
  - -- The management system must be able to monitor the bandwidth usage of devices and links, set thresholds for generating threshold alarms, and support performance reports for a large number of devices/links.

Public Cloud Management

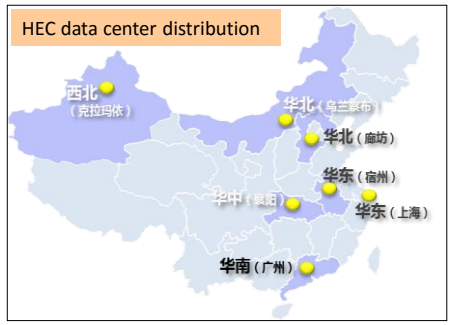

1. Generally, the public cloud consists of multiple data centers.
2. There are a large number of devices. Generally, the number of servers exceeds 10 thousands.
3. Data centers are connected to each other through dedicated lines and VPNs.
4. Generally, there is a centralized O&M center.

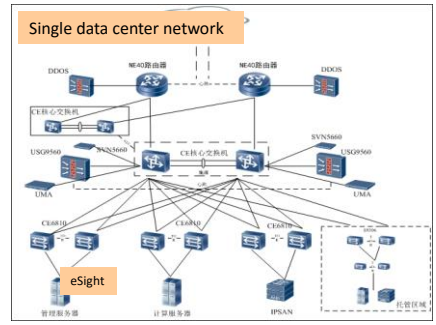

1. Each server is connected to two CE6810 access switches through two 10 GE ports.
2. Each CE6810 switch provides two 40 GE ports for connecting to two CE12808 core switches.
3. Each of the two core switches provides a 40 GE port for connecting to the router NE40.

Public Cloud Typical Application Scenarios

Scenario	Typical Application
Getting online	1. Batch configuration for getting servers and storage devices online in batches 2. Network device configuration and backup 3. Network device intelligent configuration
O&M	1. Unified management of physical and virtual facilities (unified alarm and topology) 2. Region-based monitoring (cluster-level monitoring threshold setting and related alarms) 3. E2E all-link monitoring in term of application (application faults caused by lower-layer device faults can be quickly located) 4. Tenant-level network traffic monitoring (network traffic components)
Optimization	Report (weekly resource usage report, storage capacity trend and prediction report)

Contents

1. eSight Positioning
2. Main eSight Services
3. eSight Architecture
4. Typical Applications of eSight
5. **Networking Modes of eSight**
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. eSight Installation Process

Single-Node Server Mode

- eSight uses the browser/server mode and allows simultaneous access by multiple web browsers.

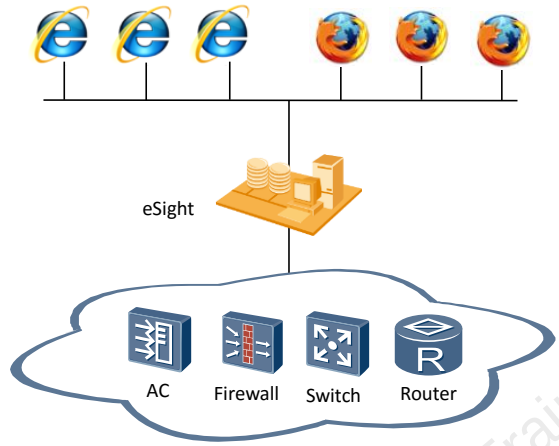

HUAWEI

- The eSight single-node system deployment solution is applicable to scenarios where the network scale is small and reliability requirements are not high.
- The single-node system deployment solution uses multiple web clients and a single server. The web client and eSight server are connected through a LAN or WAN.
- The eSight single-node system can be deployed on Huawei FusionSphere 6.1 (XEN+KVM) and VMWare ESXI 5.5 VMs. However, it does not support components such as the NTA, WLAN RTLS, Infrastructure Manager, Application Manager, Asset Manager, and Business Service Manager. Distributed deployment and HA deployment do not support VM deployment.
- If the Oracle database is deployed, the database and eSight can be deployed on the same server or different servers. (Customers need to prepare the server to install the database.)
- For a single-node system, the NTA Manager can manage 0 to 10 nodes and process 2000 flows per second.

Distributed Mode

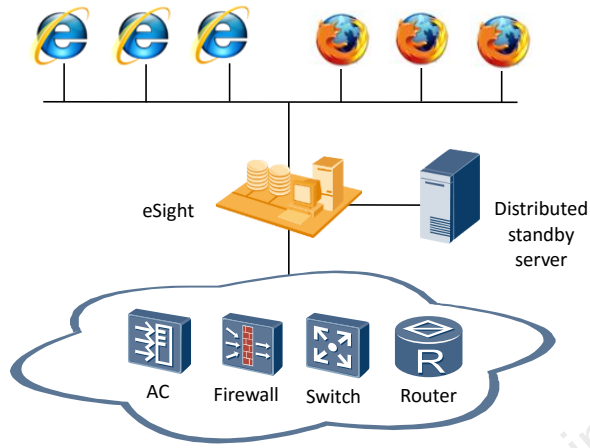

- The eSight distributed deployment solution is applicable to scenarios where the network scale is large, cross-region management is required, and reliability requirements are not high.
- eSight supports distributed deployment solution in scenarios where the network scale is large and cross-region management is required.

- Distributed servers refer to the NTC and WLAN RTLS.
  - eSight supports only one NTC.
  - eSight supports only one WlAN RTLS.

- If the Oracle database is deployed, the database and eSight can be deployed on the same server or different servers. (Customers need to prepare the server to install the database.)
- If more than 10 devices needs to be managed, the NTC must be deployed on an independent server. When the NTC is deployed in distributed mode, it can manage a maximum of 10 devices and can process 30000 flows per second.
- When over 1000 APs are involved, the WLAN RTLS must be deployed on an independent server. When the WLAN RTLS is deployed in distributed mode, it can support a maximum of 2000 APs.

HA System Mode - Local HA System

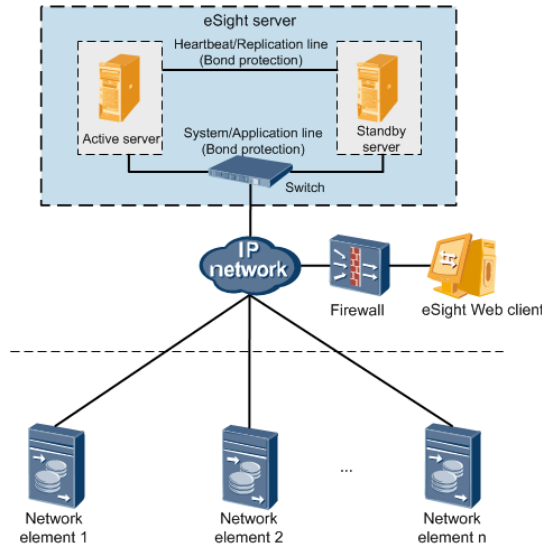

- The HA deployment mode is applicable to scenarios with high reliability requirements. eSight can be deployed in a local HA system or a remote HA system.
- The eSight local HA system consists of one active server and one standby server. Both the active and standby servers have eSight installed. Data between the active and standby servers is synchronized through the dedicated replication line. When the active server is faulty, the system automatically switches to the standby server to ensure that the eSight system runs properly.
- A floating IP address is set between the active and standby servers in the local HA system. After the active and standby servers are switched over, devices do not need to connect to eSight again.

HA System Mode - Remote HA System

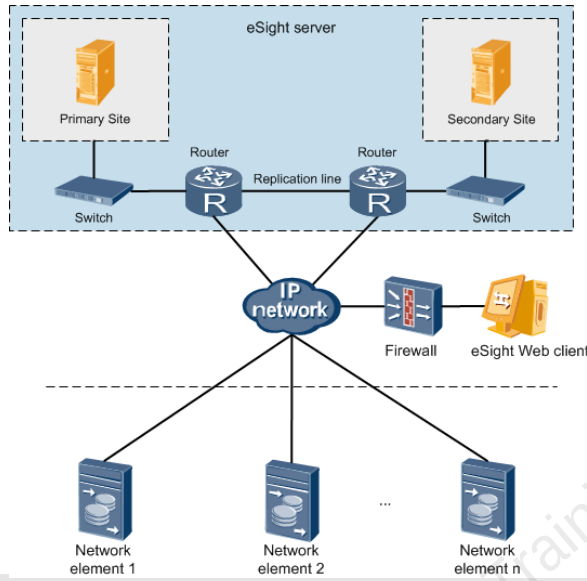

- The eSight remote HA system consists of one active server and one standby server. Both the active and standby servers have eSight installed. The two servers can be deployed at different locations. When the active server is faulty, the system automatically switches to the standby server. Data between the active and standby servers is synchronized through the dedicated replication line to ensure that the eSight system runs properly.
- The two eSight servers in the remote HA system use different IP addresses. After the remote HA system is deployed, the IP addresses of both the active and standby servers must be set on the managed devices. After the active and standby servers are switched, information such as the alarm information on the devices is automatically sent to the standby server to ensure normal device monitoring and management.

Hierarchical Deployment Mode

- In the hierarchical deployment mode, the upper-layer NMS can add the lower-layer NMS to the system and provide a link for accessing the lower-layer NMS GUI.
- When a user clicks the link to the lower-layer NMS, a new browser window is displayed and the login page of the lower-layer NMS is opened in the new browser window.

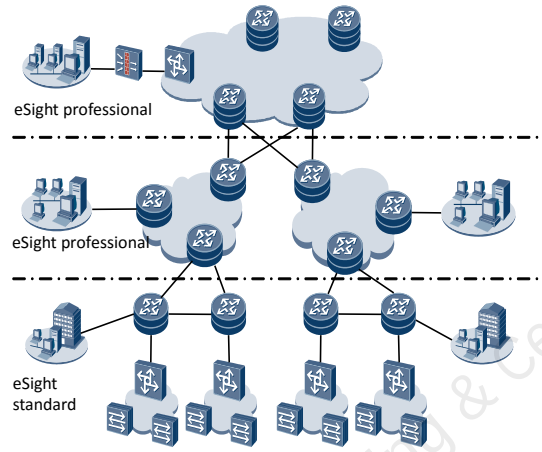

- When the upper-layer NMS accesses the home page of the lower-layer NMS, the SSO mechanism provided by the OMS is used. No manual login is required.

eSight and OSS Networking Application

- eSight can be integrated with the upper-layer OSS. eSight reports network alarms and connects to the OSS alarm system through the SNMP protocol.
- Integration with the OSS has the following advantages:
  - Improve the network management capability through the OSS system.
  - Isolate NE management from network management.
  - Meet requirements of the enterprise O&M mechanism.

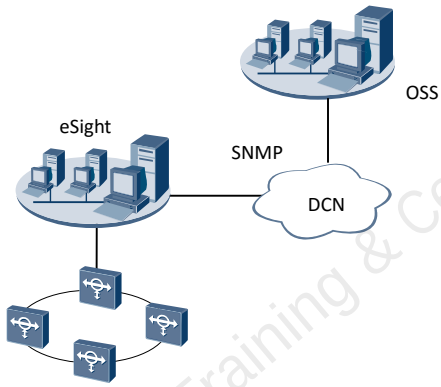

- Operating Support System (OSS)
  - The OSS is the network management software running on the workstation of the network management center. Network administrators operate the OSS to send requests to managed devices, so that they can monitor and configure these devices.
- eSight connects to the upper-layer OSS through the northbound interface. By default, eSight provides a northbound interconnection with the upper-layer NMS through the SNMP, Restful Open API, and FTP protocols.

Contents

1. eSight Positioning
2. Main eSight Services
3. eSight Architecture
4. Typical Applications of eSight
5. Networking Modes of eSight
6. **Technical Specifications of eSight**
7. eSight Installation Preparation
8. eSight Installation Process

Technical Specifications of eSight

Item	Subitem	Compact	Standard	Professional
Management capability	Number of managed NEs	40	5000	20000
Resource occupation	CPU usage	---	The usage exceeds 30% within a period of 15 minutes.	The usage exceeds 30% within a period of 15 minutes.
Storage capacity	Number of current alarms	20,000	20,000	20,000
	Number of historical alarms	---	1,500,000	1,500,000
	Number of logs	1,000,000	1,000,000	1,000,000
	Number of performance records	---	60,000,000	60,000,000
Processing capability	Alarm response speed	The interval from the time when an alarm is generated on the device to the time when the alarm is displayed on eSight does not exceed 30 seconds.	The interval from the time when an alarm is generated on the device to the time when the alarm is displayed on eSight does not exceed 30 seconds.	The interval from the time when an alarm is generated on the device to the time when the alarm is displayed on eSight does not exceed 30 seconds.
	Response speed of performance processing		About 30,000 performance records are collected within 15 minutes.	About 30,000 performance records are collected within 15 minutes.
	Maximum response time to an operation on the GUI	3 seconds	3 seconds	3 seconds
Status update time	Device status	---	No more than 300 seconds	No more than 300 seconds
	Link status	---	No more than 35 seconds	No more than 35 seconds

- eSight can manage a maximum of 20,000 NEs.

Hardware and Software Configuration Requirements

Version	Management Scale	Minimum Hardware Configuration	VM	System Configuration
Standard	200-500	CPU: 2 x dual-core, 2 GHz or higher Memory: 4 GB Hard disk space: 60 GB	The usage exceeds 30% within a period of 15 minutes.	Configuration 1: Windows Server 2012 R2 Standard Edition (64-bit) + MySQL 5.6/Microsoft SQL Server 2012 Standard Edition Configuration 2: Novell SUSE Linux Enterprise Server-Multi-language version-Enterprise-11.0 SP3 (Simplified Chinese or English version) + Oracle Database Standard Edition 11g R2
	500-2000	CPU: 2 x quad-core, 2 GHz or higher Memory: 4 GB Hard disk space: 120 GB	CPU: 2 x quad-core, 2 GHz or higher Memory: 12 GB Hard disk space: 600 GB	
Professional	2000-5000	CPU: 2 x quad-core, 2 GHz or higher Memory: 16 GB Hard disk space: 250 GB	CPU: 4 x quad-core, 2 GHz or higher Memory: 32 GB Hard disk space: 1024 GB	Configuration 1: Windows Server 2012 R2 Standard Edition (64-bit) + MySQL 5.6/Microsoft SQL Server 2012 Standard Edition Configuration 2: Novell SUSE Linux Enterprise Server-Multi-language version-Enterprise-11.0 SP3 (Simplified Chinese or English version) + Oracle Database Standard Edition 11g R2
	5000-20000	CPU: 4 x quad-core, 2 GHz or higher Memory: 32 GB Hard disk space: 320 GB	No more than 300 seconds	Configuration 1: Novell SUSE Linux Enterprise Server-Multi-language version-Enterprise-11.0 SP3 (Simplified Chinese or English version) + Oracle Database Standard Edition 11g R2

- For more information about the hardware and software configuration of the server, see the product documentation.

Client Environment

- The web-based eSight client has no special requirements on the operating system. However, it has the following requirements on the web browser version and memory:
  - Browser versions: Internet Explorer 9.0/10.0/11, Firefox 31/38 ESR/45.3 ESR, and Chrome 29/41/43/52. Internet Explorer 11, Firefox45.3 ESR, and Chrome 52 are recommended.
  - Recommended resolution: 1280 x 1024 pixels
  - Memory: 2 GB or above
- eSight Mobile:
  - Android 4.2 and later versions. The supported phones include Huawei Honor 6, Mate7, P7, and Honor 3C.
  - iOS 8.0 and later versions. The supported phones include iPhone5s, iPhone6, iPhone6 plus, iPhone6s, iPhone6s plus, and later.

Contents

1. eSight Positioning
2. Main eSight Services
3. eSight Architecture
4. Typical Applications of eSight
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. **eSight Installation Preparation**
8. eSight Installation Process

Installation Plan

- Before installing eSight, ensure that the installation information such as the IP address, host name, and password is correctly configured so that eSight can be installed quickly and correctly.
  - User name and password planning: The initial user name and password of eSight are admin and Changeme123, respectively. When you log in to eSight for the first time, eSight prompts you to change the password.
  - Host name and IP address planning: You can change the host name and IP address of the server based on the site requirements.
  - Hard disk partition planning: Drive C is used for installing the operating system, and Drive D is used for installing the database and eSight.
  - Installation path planning: The installation path is D:\eSight, which can be changed.
  - Time zone planning: When eSight is delivered to the site, change the time zone and time based on the site requirements.

Software Obtaining

- eSight can be installed in either of the following modes:
  - Using an installation CD-ROM: You need to obtain the desired CD-ROM.
  - Using a software package: You need to obtain the related software package.
- The software can be obtained from the following website:
  - Huawei enterprise network website:http://e.huawei.com/cn/.
- Perform the following operations to obtain the software:
  - Visit http://e.huawei.com/cn/.
  - Click Log In. In the dialog box that is displayed, enter the account name and password, and click Login.
  - In the search box, enter eSight.
  - Click the path displayed under the search box to go to the product documentation page.
  - Click the Software Download tab and download the required software.

Contents

1. eSight Positioning
2. Main eSight Services
3. eSight Architecture
4. Typical Applications of eSight
5. Networking Modes of eSight
6. Technical Specifications of eSight
7. eSight Installation Preparation
8. **eSight Installation Process**

Installation Process

- Preinstallation: The operating system, database, antivirus software, and eSight have been preinstalled on servers delivered by Huawei.
- Brand-new installation: If a self-purchased server is used or eSight needs to be reinstalled, follow the flowchart on the right to install eSight.

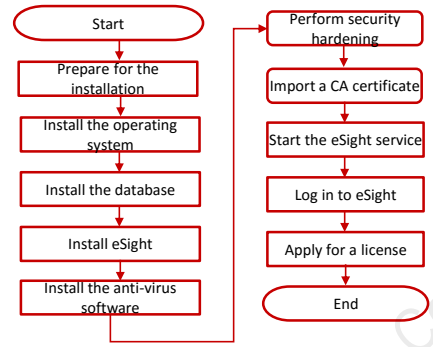

- Installation description:
  - Support the Windows and SUSE Linux operating systems.
  - Support silent installation.
  - Support English and Chinese.
  - Adopt the Browser/Server architecture, which is free from client installation.
  - The installation process is irrelevant of the license. The license must be imported by users after installation.
  - The size of installation CD-ROM of each edition does not exceed 1 GB.
  - The installation can be completed in 10 minutes.

Install the Database

- Windows Server 2012 R2 supports the following two types of databases:
  - MySQL:
    - The MySQL database is automatically installed along with eSight.
  - SQL Server 2012:
    - If the SQL Server 2012 database is used, manually install it following the installation manual delivered with the database before installing the eSight server.

Start eSight Installer

- Run the setup.bat into the installation process, the first step to select a language, which determines the language of the installation disk and network management language, the default is English.

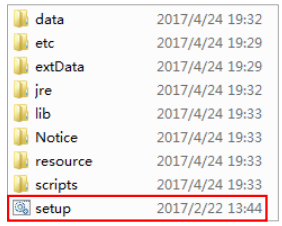

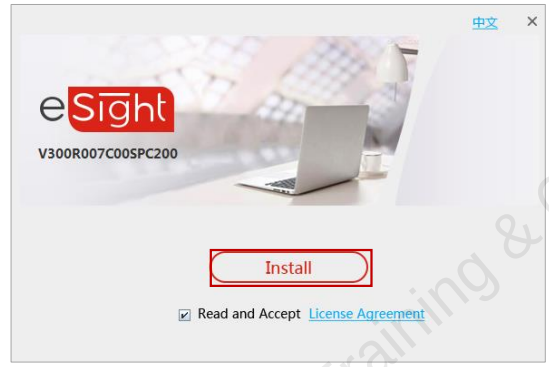

Installation Parameters

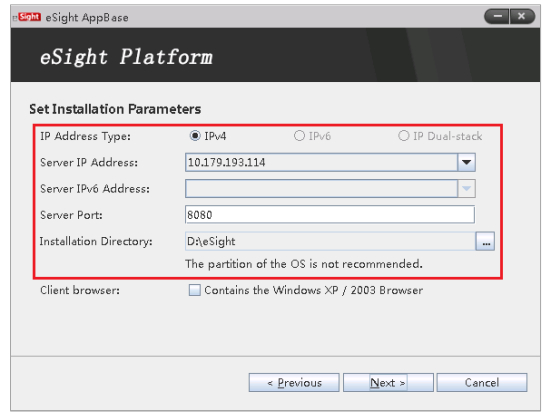

- IP address: the current server default IP address. If the server has multiple IP addresses in the drop-down list, select the IP address to be announced.
- Port: the default 8080, if the port is occupied, then modify the port number.
- Installation path: the eSight Software installation path, the user can manually modify the path.

Database Parameters

- If the Database Type value is SQLServer, enter the sa account password.
- If the Database Type value is MySQL, the system will automatically enter the database password. For the dbadmin account, the password is Changeme123; for the commonuser account, the password is Changeme_123.
- During incremental installation, the password text box is dimmed and you do not need to enter the commonuser account password.

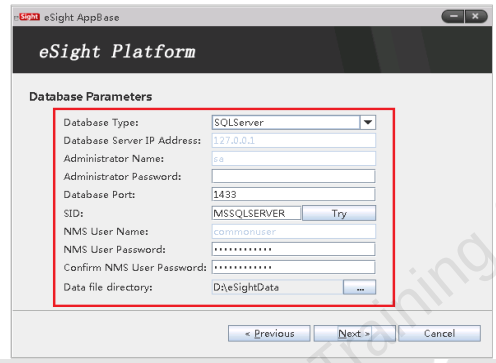

Select Software Components

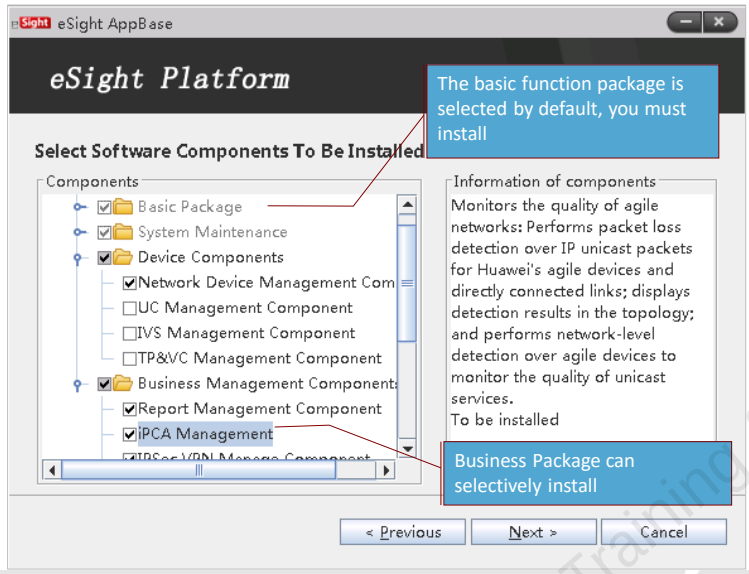

HUAWEI

- Components under Basic Package are dimmed by default, and must be installed.
- Select required components under Business Component. eSight components support incremental installation. If a component is not installed during the first eSight server installation, it can be installed during the next eSight server installation.

Show Installation Progress

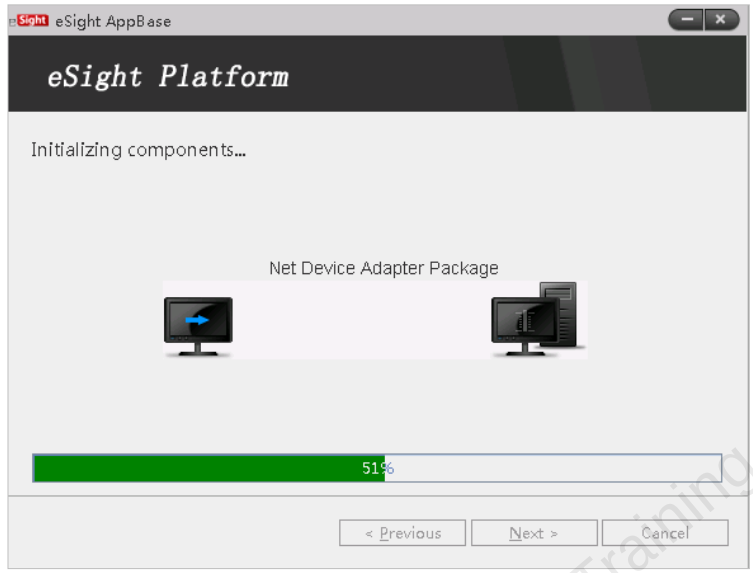

Installation Completed

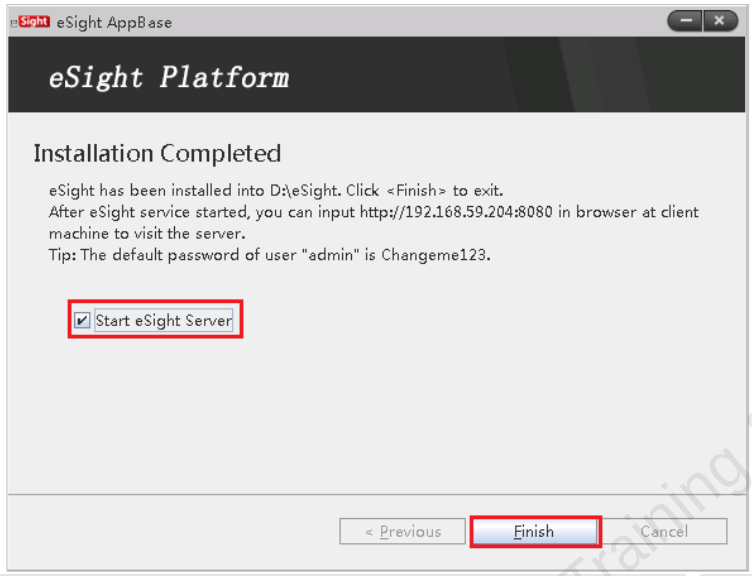

Integrated Deployment with LogCenter

- eSight and LogCenter Deployed on Different Servers
  - When eSight and LogCenter are deployed on different servers, LogCenter can be deployed on the Linux operating system, and eSight can be deployed on the Windows operating system.
  - LogCenter can be deployed in centralized mode, or collectors can be deployed in distributed mode.
- eSight and LogCenter Deployed on Same Server
  - When eSight and LogCenter are deployed on the same server, the server must run the Windows operating system.
  - In single-server integrated installation is complete, LogCenter and eSight are started, stopped, or uninstalled together.In distributed deployment mode, collectors can be installed, started, stopped, or uninstalled independently.

Start eSight

1

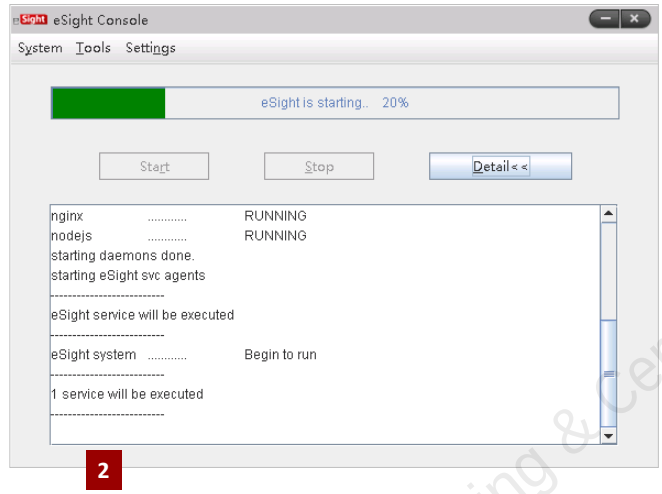

2

Install the certificate

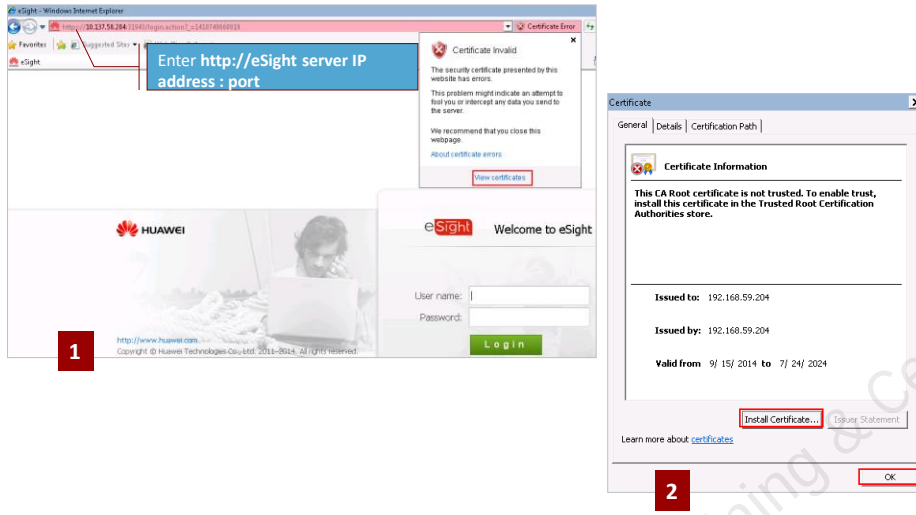

HUAWEI

- Log in to the eSight server from the Firefox.
- If you log in to the eSight server for the first time, the message "There is a problem with this website's security certificate" is displayed. Click Continue to this website (not recommended).
- Click View Certificates.
- Select Install Certificate and click OK.

Log in to the eSight server from the IE9

- In the address box, enter http://eSight server IP address:port number, and press Enter.
- Optional: If you log in to the eSight server for the first time, the message "There is a problem with this website's security certificate." is displayed. Click Continue to this website (not recommended).

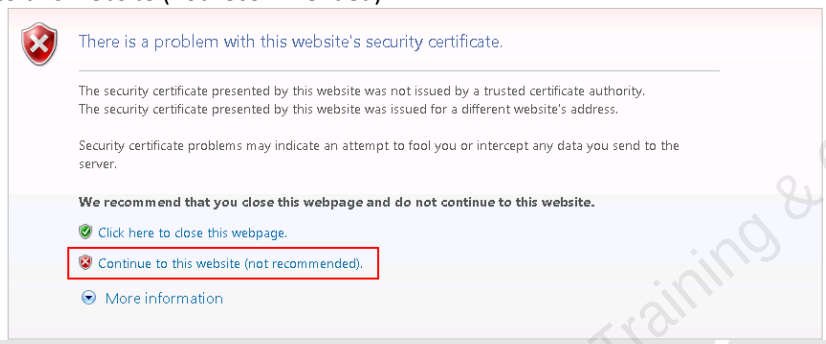

HUAWEI

- Why there is always prompt certificate error on first logon to eSight?
  - Answer: To understand this issue, please understand the basic principles of encryption:
  - Symmetric encryption & Asymmetric encryption: Symmetric encryption is used with a password between A and B, It uses the same password to encrypt and decrypt, the rate is relatively high.
  - Asymmetric encryption is divided into public and private keys, the private key can be saved by the user, publish your public key to a public network, data can only be used by the public key encryption private key to decrypt data encrypted by the private key can only by the public key to decrypt.
- HTTPS Secure Hypertext Transfer Protocol using SSL protocol (Secure Socket Layer), the following is a brief interaction:
  - The client sends a request to the server-side.
  - The server sends a message with the certificate (including the server side of the certificate and public key) to the client.
  - Client through the encryption algorithm to generate a key A, and then use the certificate (which is server-side public key) to encrypt, send to the server side.
  - Server-side use their own private key to decrypt the data sent and received the key A.

Login-Input User Name & Password

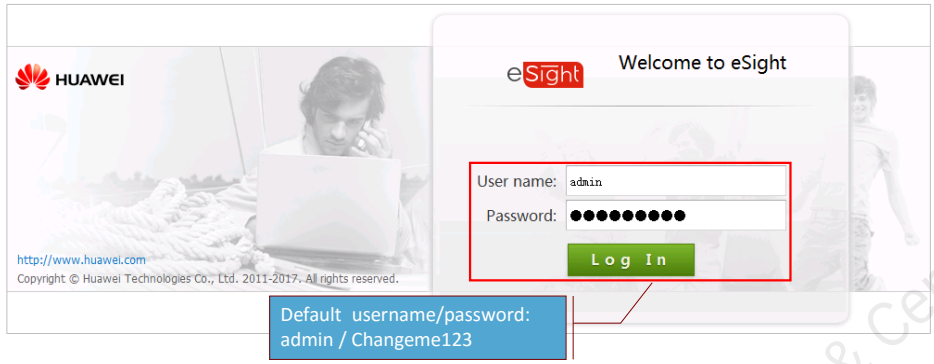

Login Successful

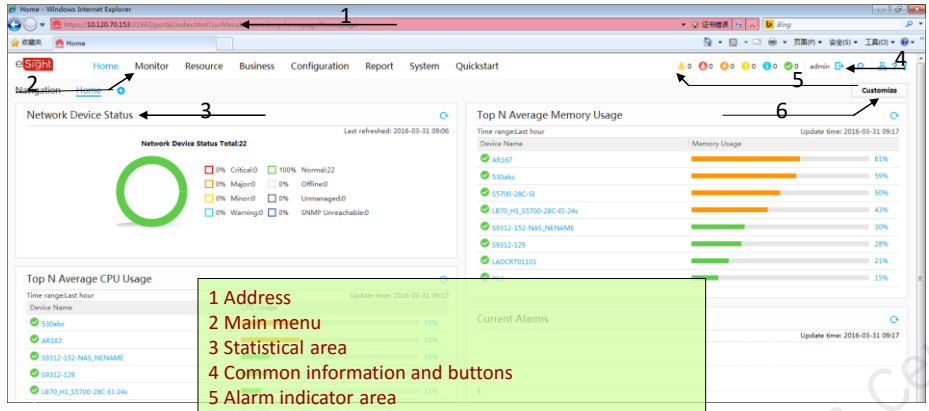

1 Address
2 Main menu
3 Statistical area
4 Common information and buttons
5 Alarm indicator area
6 Button for adding a Portal

Introduction to License Files

- License files are authorized files generated by special encryption tools.
- License files information include contracts signed by customers and Huawei and information about the server where eSight is installed.
- When customers get the license file, then they manually load their license files to eSight to obtain the permission to use eSight.

Applying for an eSight License

- Step1: Obtain the contract information.
  - The soft copy of Huawei license authentication certificate is delivered along with the eSight. Obtain the contract number, license authentication code (LAC), product name, and product model from Huawei license authentication certificate.
- Step2: Obtain the server equipment serial number (ESN).
  - ESN(Equipment serial number) : An ESN is a string that uniquely identifies a device. It is used to ensure that the license is granted to the specified device.
- Step3: Apply for an eSight license.
  - Login to the Huawei enterprise FNO website to apply for licenses : http://app.huawei.com/isdp/
  - Problem encountered in the license application process, please contact the Huawei local eSight products manager or call the Huawei Service Hotline.

- ESN（Equipment serial number）：the device serial number is a string that uniquely identifies the device, used to ensure that the License authorization key to the specified device, also known as the "device fingerprint.

Obtaining Contract Information

- Obtain the contract number, product name, and product version from the license certificate.

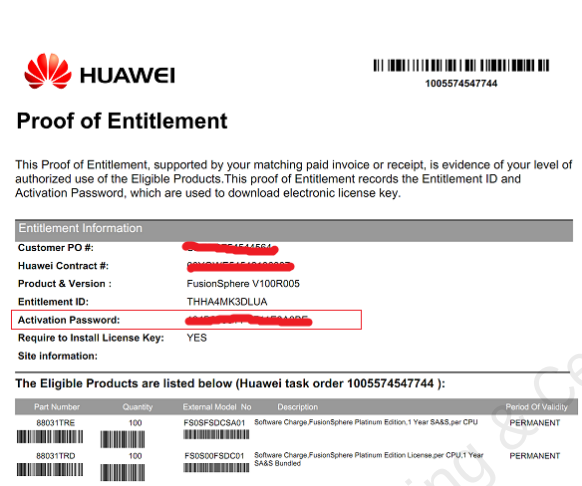

HUAWEI

1005574547744

Proof of Entitlement

This Proof of Entitlement, supported by your matching paid invoice or receipt, is evidence of your level of authorized use of the Eligible Products.This proof of Entitlement records the Entitlement ID and Activation Password, which are used to download electronic license key.

Entitlement Information	
Customer PO #:	[redacted]
Huawei Contract #:	[redacted]
Product & Version :	FusionSphere V100R005
Entitlement ID:	THHA4MK3DLUA
Activation Password:	[redacted]
Require to Install License Key:	YES
Site information:	

The Eligible Products are listed below (Huawei task order 1005574547744):

Part Number	Quantity	External Model No	Description	Period Of Validity
88031TRE	100	FS0SFSDCSA01	Software Charge,FusionSphere Platinum Edition,1 Year SA&S,per CPU	PERMANENT
88031TRD	100	FS0S00FSDC01	Software Charge,FusionSphere Platinum Edition License,per CPU,1 Year SA&S Bundled	PERMANENT

Applying for a License

- Step 1: Choose "License Activation >Password Activation" to open the Password Activation page, enter the activation code.

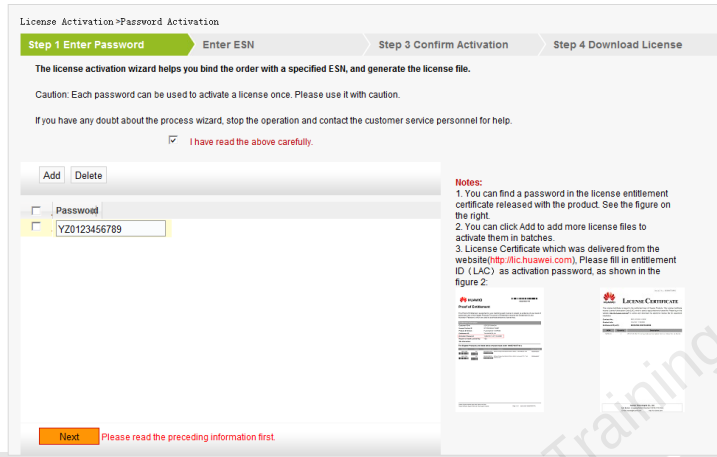

- Login to the Huawei enterprise FNO website to apply for licenses: http://app.huawei.com/isdp/ Users can login Huawei enterprise FNO website by Activation ID in certificate, activate and download the license. Users can click the link of ‘with Activation ID:’ in left pf page, or click the link in certificate email to login Huawei enterprise FNO website.

Applying for a License

- Step 2: Enter the ESN of the server to be bound.

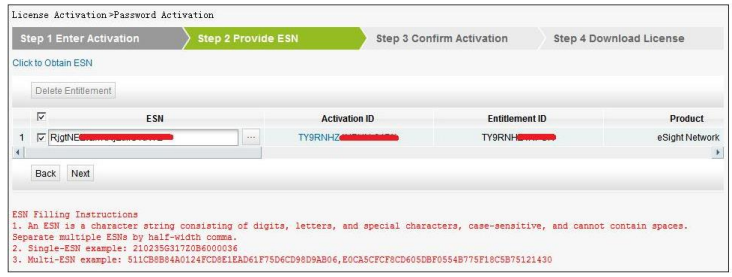

- Step 3: Confirm activation.

- Users can click the button of ‘active’ and perform the operations of activation.

Loading a License File

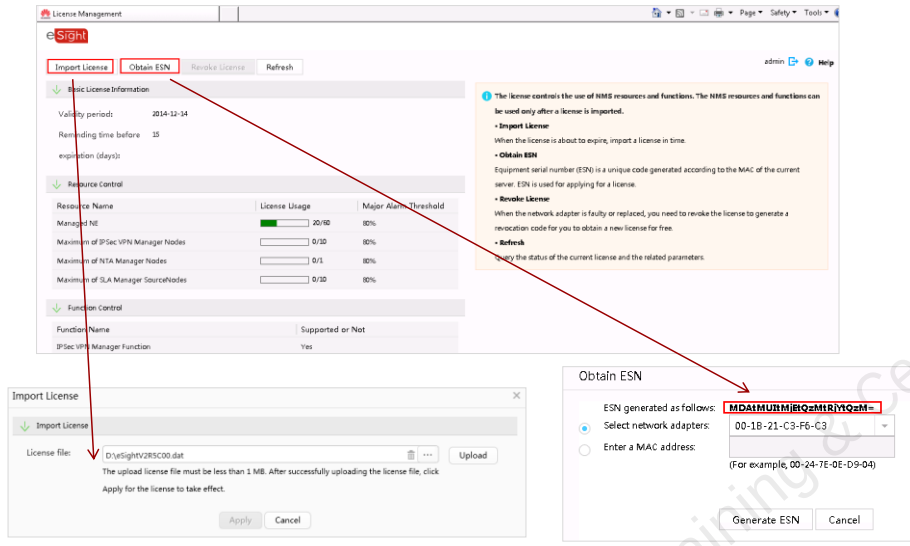

- After the success of the License load, use the administrator user login eSight system. Select System> License Management to view the items and resources of the License expiration date, and control applications License checking whether correct.

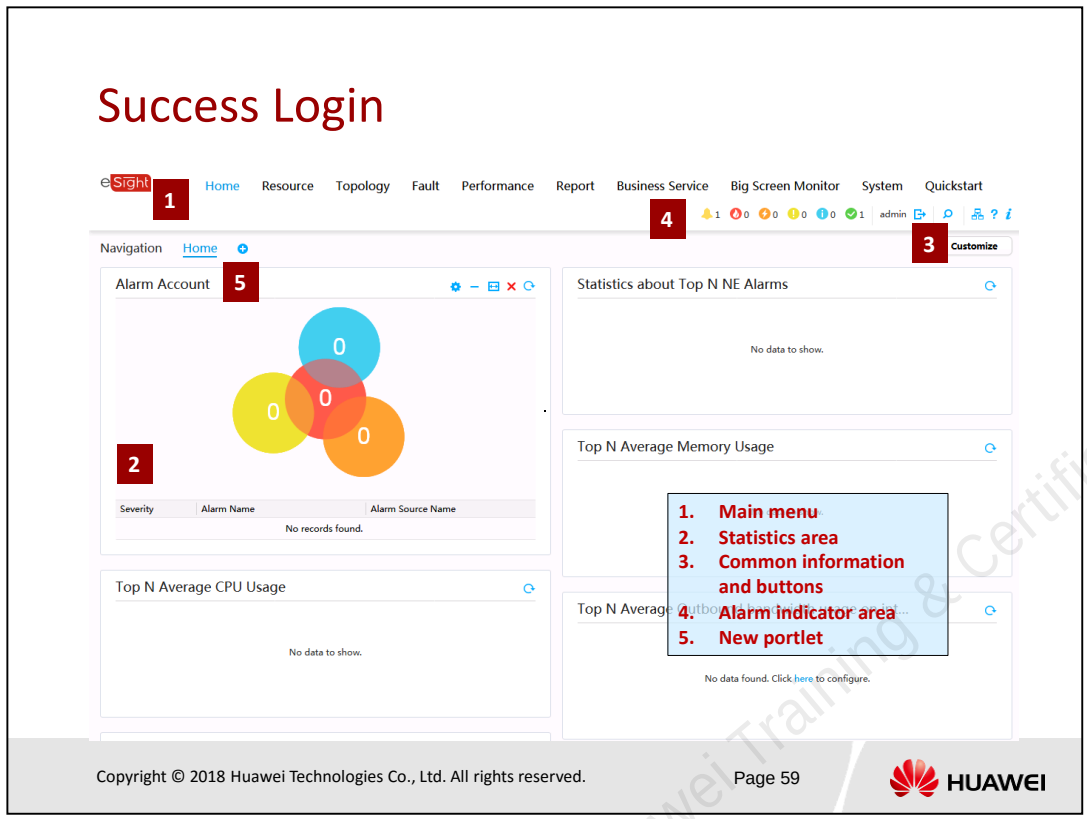

- If the browser is not equipped with a flash player, download and install Adobe Flash Player provided by eSight. Otherwise, some pages cannot be displayed properly. When you upgrade eSight or install a patch on eSight, verify the digital signature to ensure software validity. This operation is recommended to prevent security risks caused by software tampering or replacement.

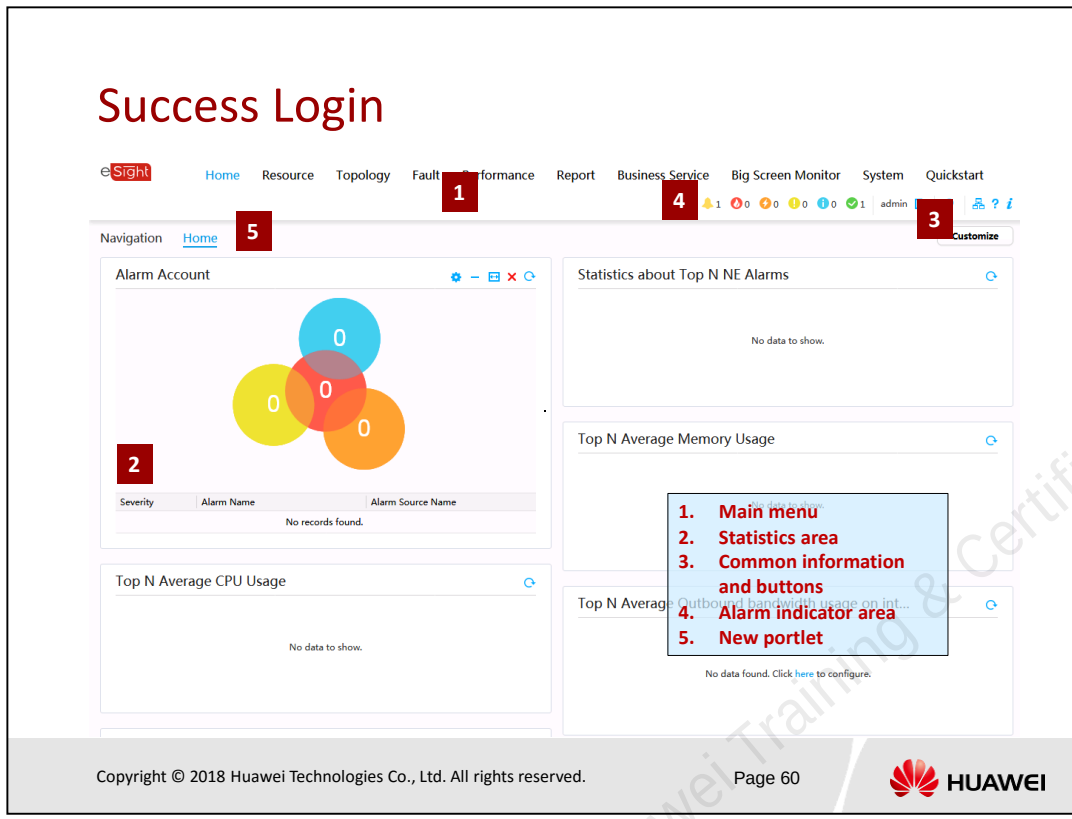

- If the browser is not equipped with a flash player, download and install Adobe Flash Player provided by eSight. Otherwise, some pages cannot be displayed properly. When you upgrade eSight or install a patch on eSight, verify the digital signature to ensure software validity. This operation is recommended to prevent security risks caused by software tampering or replacement.

Uninstall - Show Steps

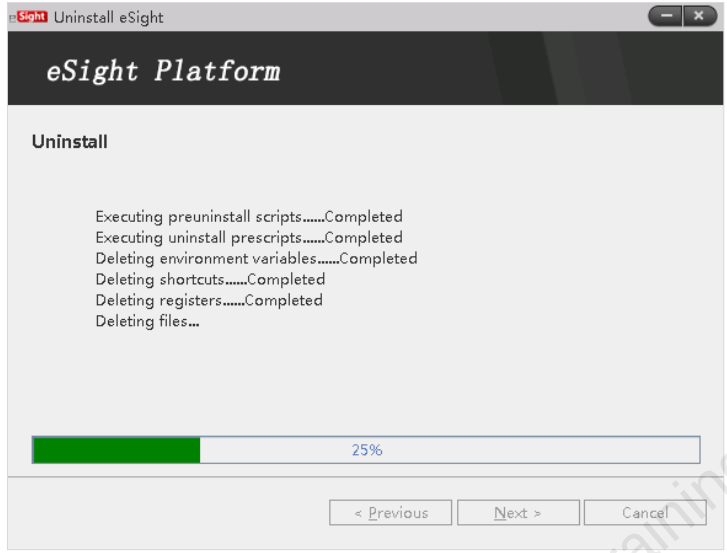

Uninstall Completed

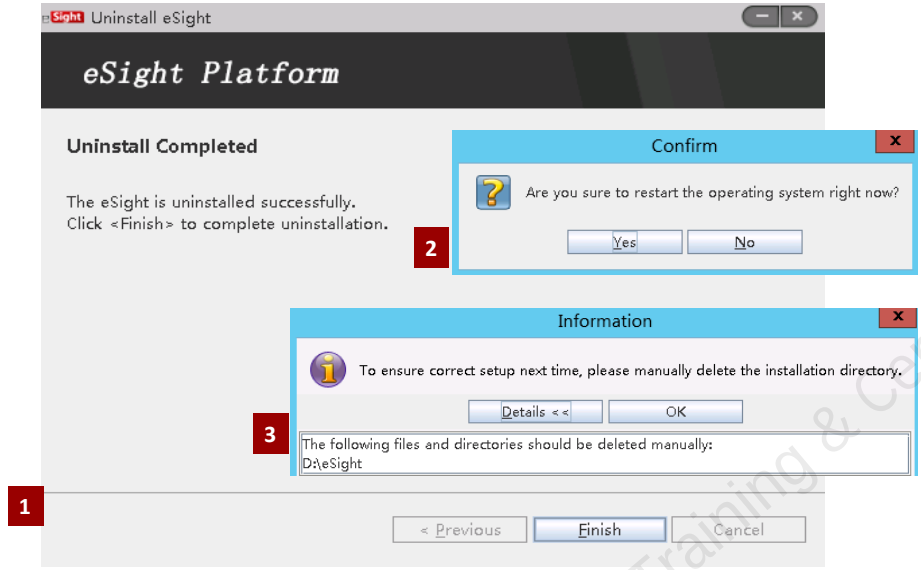

Summary

- eSight Positioning
- Main eSight Services
- eSight Architecture
- Typical Applications of eSight
- Networking Modes of eSight
- Technical Specifications of eSight
- eSight Installation Preparation
- eSight Installation Process

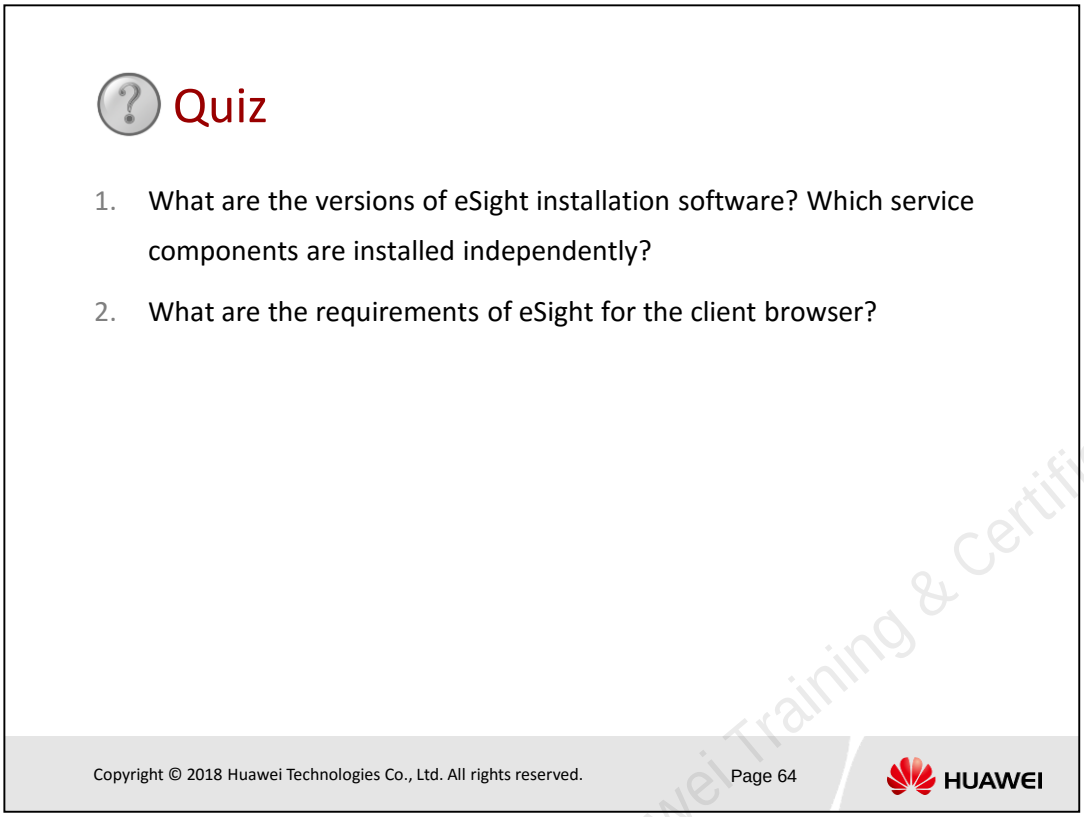

- What are the versions of eSight installation software? Which service components are installed independently?
  - Supports the two sets of CD-ROM for the Windows and SUSE Linux operating systems.
- What are the requirements of eSight for the client browser?
  - The web-based eSight client has no special requirements on the operating system. However, it has the following requirements on the web browser version and memory:
  - Browser version: Internet Explorer 9.0/10.0, Mozilla Firefox 27.0/30.0/31.0, and Chrome 29/30/31
  - Recommended resolution: 1024 x 768 pixels
  - Memory: 1 GB or above

Recommendations

- Huawei learning website:
  - http://support.huawei.com/learning/Index!toTrainIndex
- Huawei support case library:
  - http://support.huawei.com/enterprise/servicecenter?lang=zh

Thank You

www.huawei.com

Recommendations

- Huawei Learning Website
  - http://learning.huawei.com/en
- Huawei e-Learning
  - https://ilearningx.huawei.com/portal/#/portal/ebg/51
- Huawei Certification
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_31&lang=en
- Find Training
  - http://support.huawei.com/learning/NavigationAction!createNavi?navId=_trainingsearch&lang=en

More Information

- Huawei learning APP