PassQueen
a.k.a dumpsqueen.com

HCIP-Cloud Computing V4.0

Huawei H13-527

Version Demo

Total Demo Questions: 42

Total Premium Questions: 425

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing Trends	3
Topic 2, Cloud Computing Core Technologies	89
Topic 3, Huawei Cloud Computing Solution	231
Topic 4, Huawei Cloud Computing Solution O&M	102
Total	425

QUESTION NO: 1

In Huawei desktop cloud, FusionAccess management monitoring interface can view the desktop virtual machine: (multiple choice)

A. Running status statistics

B. Login status statistics

C. Distribution status statistics

D. Performance statistics

ANSWER: A B C

Explanation:

FusionAccess provides centralized desktop-service monitoring so that administrators can quickly understand desktop availability, user connection activity, and how desktop resources have been assigned. **Running status statistics** are available because the management console needs to present the operational state of virtual desktops, supporting day-to-day service checks and fault handling. **Login status statistics** are also available, allowing administrators to observe desktop-user session and login conditions across the managed environment. **Distribution status statistics** are part of desktop management monitoring because administrators must be able to track the allocation or delivery state of desktop resources to users.

These status views are management-level, aggregate operational statistics intended to help administrators assess the current desktop-cloud service condition. They align with FusionAccess's role as the access and desktop-service management component in a Huawei desktop-cloud solution, while detailed infrastructure metrics are typically handled through the relevant virtualization and management components. Huawei's product documentation and support portal provide the official documentation resources for FusionAccess administration and monitoring: Huawei FusionAccess product documentation and Huawei Enterprise Support.

QUESTION NO: 2

Which of the following statements about linked clone virtual machines in Huawei desktop cloud is correct?

A. The creation speed is slow and it takes up little space

B. The creation speed is slow and takes up a lot of space

C. Fast creation speed and small space occupation

D. The creation speed is fast and takes up a lot of space

ANSWER: C

Explanation:

Fast creation speed and small space occupation is correct. In Huawei desktop-cloud deployments, a linked-clone desktop is created from a common parent or gold-image virtual machine rather than by making a complete independent copy of that image for every desktop. The clone retains a reference to the shared base disk and stores only its own changed data in a separate differential disk. Because the system does not need to copy the entire operating-system disk and application image for each new virtual desktop, provisioning is substantially faster than creating full-clone desktops. The shared base image also avoids duplicating the same operating-system and application files across every desktop, which greatly reduces initial storage consumption and is especially valuable when large numbers of similar desktops are deployed. Administrators can centrally prepare and maintain the base image, then rapidly deliver desktops based on that standardized image. The storage benefit applies to the common portions of each desktop image; each user desktop still consumes additional capacity as writes and personalized changes accumulate in its differential disk. Huawei's desktop-cloud portfolio is designed around centralized desktop delivery and image-based desktop provisioning, as described on the Huawei Desktop Cloud product page. For the underlying linked-clone principle, see the Horizon virtual desktop documentation, which describes desktops created from a shared parent image with per-desktop delta disks.

QUESTION NO: 3

The Nova modules deployed in load mode include:

(Multiple choice)

A. nova-api

B. nova-scheduler

C. nova-compute

D. nova-conductor

ANSWER: A B D

Explanation:

In a highly available Nova control-plane deployment, **nova-api**, **nova-scheduler**, and **nova-conductor** are service roles that can be deployed as multiple active instances behind a load-balancing mechanism. nova-api is the client-facing REST API endpoint, so load balancing distributes incoming API requests and provides endpoint availability. nova-scheduler is a stateless control-plane service that selects an appropriate compute host for a requested instance; multiple scheduler processes can consume scheduling requests through the messaging infrastructure. nova-conductor is also a control-plane service that mediates database access and coordinated instance operations for compute services, and it supports running multiple service instances for resilience and capacity. The duplicated *nova-compute* wording has been corrected to *nova-scheduler*, and the conductor service name has been corrected to *nova-conductor*. This matches the standard Nova service architecture and its high-availability design, where API, scheduler, and conductor services are scalable control-plane components. See the OpenStack Nova high-availability guide and the Nova services documentation.

QUESTION NO: 4

Which of the following options belong to the services included in the role of sys-server in Huawei FusionSphere OpenStack?

A. cps

B. apecheproxy

C. ntp

D. haproxy

ANSWER: A B C D

Explanation:

In FusionSphere OpenStack, **sys-server** is a foundational management role that deploys shared system services required for reliable operation of the cloud platform. **cps** belongs to this role because Cloud Provisioning Service components support the platform's management and provisioning functions. **apecheproxy** (the service name is commonly presented as Apache proxy in technical documentation) is included to provide HTTP reverse-proxy capabilities for platform service access. **ntp** is also a sys-server service because consistent time synchronization among controllers, compute nodes, storage components, and management services is essential for authentication, logging, certificates, and distributed-service coordination. **haproxy** is deployed in this role to provide load balancing and high-availability access to management interfaces and backend services. Together, these services form part of the common system-service layer rather than tenant-facing workload services. Huawei's FusionSphere OpenStack documentation and product-support resources describe the platform's role-based deployment architecture and its supporting management infrastructure: Huawei FusionSphere OpenStack support portal and Huawei FusionSphere OpenStack product information.

QUESTION NO: 5

Fusionsphere can realize the horizontal integration of heterogeneous computing, storage, and network resources of multiple manufacturers, and provide open IT service interfaces to the outside world

A. Yes

B. Wrong

ANSWER: A

Explanation:

“Yes” is correct. Huawei FusionSphere is a cloud operating-system and virtualization platform designed to pool and manage data-center infrastructure resources, including compute, storage, and networking. A core value of this type of platform is horizontal resource integration: it abstracts underlying physical infrastructure so that resources from different vendors can be incorporated into a unified cloud resource pool and delivered consistently. This enables enterprises to evolve existing data centers rather than requiring every infrastructure component to come from one manufacturer.

FusionSphere also exposes cloud capabilities through standardized, open interfaces, enabling external systems, operational tools, and service consumers to request and manage IT resources programmatically. These interfaces support the delivery of infrastructure services such as virtual machines, virtual networks, and storage through self-service and automated workflows. Consequently, the statement accurately describes both FusionSphere’s heterogeneous resource-integration role and its ability to provide open IT service interfaces. Huawei describes its cloud-computing portfolio and infrastructure-cloud capabilities at Huawei Enterprise Cloud Computing; additional product and technical documentation is available through Huawei Enterprise Support: FusionSphere OpenStack.

QUESTION NO: 6

If the user accidentally disables the virtual machine’s network, the virtual machine will not be able to log in through the HDP protocol, but it can be

VNC

Log in to this VM,

Then enable the network

A. Yes

B. Wrong

ANSWER: A

Explanation:

Yes is correct. VNC console access is an out-of-band console capability provided through the virtualization or cloud management platform. It presents the virtual machine’s display and accepts keyboard and mouse input through the platform’s console channel, rather than requiring the guest operating system’s production network interface to be reachable. Therefore, disabling the VM’s virtual NIC or incorrectly configuring its IP settings prevents normal network-based remote access, including HDP access, but does not prevent an administrator from opening the VNC console from the management plane.

After connecting through VNC, the administrator can sign in locally to the guest operating system and restore connectivity by enabling the network adapter, correcting network configuration, or restarting the relevant network service. This makes VNC especially valuable for recovery from guest-network misconfiguration, firewall mistakes, or accidental adapter disablement. The management platform itself must remain available to the administrator, and the VM must be powered on, but the guest’s own network path is not a prerequisite for console access. Huawei Cloud likewise documents VNC login as a console-based mechanism for accessing an ECS when ordinary remote login is unavailable: Huawei Cloud ECS User Guide: Logging In to a Windows ECS Using VNC. General VNC protocol information is available from RFC 6143.

QUESTION NO: 7

When adding DNS to FusinSphere OpenStack OM service, which of the following statements is correct?

A. The domain name format is: local-oam.domainname.com

B. For single-node deployment, the IP address is the floating IP address of FusionSphere OpenStack OM on the Extemal_API plane.

site

C. For single-node deployment, the IP address is FusionSphere OpenStack OM floating IP on the Extemal_OM plane

address

D. Single node deployment, the IP address is the FusionSphere OpenStack OM management IP address

ANSWER: A D

Explanation:

When a DNS entry is added for the FusionSphere OpenStack OM service, the OM-domain naming convention uses the `local-oam` prefix. Therefore, `local-oam.domainname.com` is the required domain-name format. This naming convention provides a stable, recognizable endpoint for access to the operations and maintenance service and is used consistently when the deployed environment resolves OM-service addresses through DNS.

In a single-node deployment, the appropriate address for this DNS record is the FusionSphere OpenStack OM management IP address. A single-node architecture does not rely on the same service-plane floating-IP arrangement used by highly available multi-node designs; the management IP is the address that identifies and provides access to the OM node in this deployment model. Configuring the DNS record with the management IP ensures that requests to `local-oam.domainname.com` resolve to the node hosting the OM service.

These settings follow the FusionSphere OpenStack deployment and O&M networking approach, in which DNS records must correspond to the service role and deployment architecture. See Huawei's FusionSphere OpenStack support documentation and the Huawei enterprise documentation portal for deployment documentation and configuration guidance.

QUESTION NO: 8

Which of the following is not the template type used by desktop cloud?

A. Full copy

B. Linked clone

C. Fast packaging

D. Quick copy

ANSWER: D

Explanation:

Quick copy is the correct choice because it is not a standard template type or template-creation mechanism in Huawei FusionAccess desktop-cloud terminology. FusionAccess desktop provisioning uses established template approaches that distinguish how a desktop image is copied, shared, or prepared for rapid delivery. The naming is important in this context: administrators select the supported template/provisioning type when creating and distributing virtual desktops, and *Quick copy* is not one of the recognized labels used by the platform.

This distinction matters operationally because the selected template method determines such characteristics as image storage consumption, the relationship between a desktop and its base image, lifecycle management, and the speed at which desktops can be delivered. A term that is not part of the supported template taxonomy cannot be selected as a valid desktop-cloud template type in the relevant Huawei management workflow. Huawei positions FusionAccess as its virtual desktop infrastructure solution and provides product and documentation resources through its enterprise portal. See the

Huawei FusionAccess product page and the Huawei Enterprise Support documentation portal for version-specific FusionAccess administration documentation.

QUESTION NO: 9

The storage types that can be used to create virtualized data storage in fusioncompute are: (multiple choice)

A. FC SAN

B. NAS

C. IP SAN

D. Fusionstorage

ANSWER: A B C D

Explanation:

FusionCompute virtualized datastores can be created by connecting hosts to both block-storage and file-storage back ends, as well as to Huawei's distributed storage platform. **FC SAN** is supported as a Fibre Channel block-storage back end, allowing hosts to access LUNs through an FC SAN fabric. **IP SAN** is also supported for block storage delivered over IP networking, such as iSCSI-based storage. **NAS** is supported for file-based shared storage, enabling datastores based on network file-sharing protocols. **Fusionstorage** is likewise a supported storage type: Huawei FusionStorage integrates with FusionCompute to provide distributed storage resources for virtual machines and their virtual disks. These choices represent the principal external/shared storage categories available when creating virtualized storage resources in the FusionCompute environment. Huawei describes FusionCompute as its virtualization platform and documents its integration with storage resources in its cloud-computing portfolio; Huawei also identifies FusionStorage as a distributed storage product designed for cloud data-center scenarios. See the Huawei FusionCompute product page and the Huawei FusionStorage product page.

QUESTION NO: 10

For computing virtualization, which of the following descriptions is correct (multiple choice)

A. Computing virtualization can realize the reuse of computing resources

B. Virtual machines can run operating systems and applications like physical machines

C. With virtualization, physical machines are no longer needed

D. With virtualization technology, one physical machine can run multiple virtual machines at the same time

ANSWER: A B D

Explanation:

Computing virtualization abstracts the CPU, memory, storage, and network resources of a physical server and presents them as configurable virtual hardware. This abstraction permits computing resources to be allocated, adjusted, and reused among workloads more efficiently than dedicating an entire server to each workload. Huawei Cloud describes Elastic Cloud Servers as virtual servers with configurable computing resources, illustrating this resource-pooling model. See Huawei Cloud Elastic Cloud Server.

Virtual machines provide an isolated virtual hardware environment in which a guest operating system and its applications can be installed and operated in much the same way as on a physical computer. A hypervisor manages the virtual machines and shares the underlying physical host resources while maintaining separation between individual guests. As a result, one physical host can concurrently support multiple virtual machines, subject to the host's available capacity and the resource allocations configured for those virtual machines. This consolidation capability is a core benefit of server virtualization and supports improved infrastructure utilization, flexible provisioning, and workload isolation. Huawei's cloud-computing overview also identifies virtualization as the foundation for delivering pooled and elastic computing services; see Huawei Cloud: What Is Cloud Computing?.

QUESTION NO: 11

In OpenStack Swift, which internal data structure is used to map an account, container, or object partition to the storage devices that hold its replicas?

A. Ring

B. Proxy server

C. Swift database

D. Object server

ANSWER: A

Explanation:

Ring is correct. Swift uses rings to map partitions to devices across storage nodes. Separate account, container, and object rings allow the relevant Swift services to determine where metadata and object replicas should be stored and retrieved. The ring design supports data distribution and replica placement across failure domains.

A proxy server receives client requests and forwards them to the appropriate storage services, but it does not replace the ring as the placement map. The Swift database stores account or container metadata, while the object server stores object data on storage devices.

See the OpenStack Swift ring overview.

QUESTION NO: 12

The purpose of virtualization is to virtualize and run one or more virtual machines on the indoor machine through the use of a virtual manager

A. Yes

B. Wrong

ANSWER: A

Explanation:

Yes is correct. In this statement, “indoor machine” is evidently intended to mean the physical host machine. Virtualization uses a hypervisor, also called a virtual machine monitor or virtual manager, to abstract the host’s physical compute, memory, storage, and network resources. The hypervisor creates and manages isolated virtual hardware environments, allowing one physical server to host and run one or more independent virtual machines. Each virtual machine can run its own operating system and applications while sharing the underlying hardware in a controlled manner. This abstraction is a foundational cloud-computing capability because it improves hardware utilization, enables flexible resource allocation, simplifies workload deployment, and provides isolation between workloads. Huawei Cloud Elastic Cloud Server services similarly present virtual server instances to users while the underlying infrastructure supplies and manages the physical resources. Huawei’s cloud-service materials describe ECS as virtual computing resources, and its cloud architecture relies on virtualization to pool and deliver infrastructure resources efficiently. See Huawei Cloud Elastic Cloud Server and Huawei Cloud Image Management Service.

QUESTION NO: 13

When OpenStack adopts 2+1 deployment, what services must be deployed on the third node?

A. zookeeper

B. swift

C. compute

D. image

ANSWER: A B

Explanation:

In the Huawei OpenStack 2+1 deployment model, the third node is assigned the supporting services **zookeeper** and **swift**. ZooKeeper provides the distributed coordination capability required by the platform's clustered components. Placing it on the dedicated third node separates this coordination workload from the two primary nodes and supports stable service operation in the compact deployment architecture. Apache ZooKeeper describes its role as a centralized service for configuration, synchronization, and group services used by distributed applications: ZooKeeper documentation.

Swift is the OpenStack Object Storage service. It provides distributed object storage for data such as images, backups, and other object-based content used by cloud services. Deploying Swift on the third node is part of the prescribed 2+1 role allocation, ensuring that object-storage functionality is available without assigning this dedicated storage-service role to the primary nodes. The OpenStack project documents Swift as a highly available, distributed object-storage system at OpenStack Swift documentation. Therefore, the required third-node services are zookeeper and swift.

QUESTION NO: 14

Traditional virtualization technology is the underlying support technology of OpenStack, but OpenStack can also be used without virtualization technology.

outer

Provide cloud services.

A. Correct

B. Error

ANSWER: A

Explanation:

Correct is the right answer. OpenStack commonly delivers compute services by managing virtual machines through hypervisors such as KVM, which is why traditional virtualization is a major foundational technology in many OpenStack deployments. However, OpenStack is not limited to virtual-machine-based cloud services. Its Bare Metal service, called Ironic, enables OpenStack to provision, manage, and reclaim physical servers directly. This allows an OpenStack cloud to provide Infrastructure as a Service on dedicated hardware without placing a traditional hypervisor and virtual-machine layer on those servers. Bare-metal provisioning is particularly useful for workloads that need direct hardware access, predictable performance, specialized accelerators, or operating-system configurations that are not appropriate for virtual machines. OpenStack services such as identity, image, networking, orchestration, and block storage can still be integrated with such physical-server deployments, so the cloud-management capabilities remain available even where compute instances are bare metal. Therefore, virtualization is widely used by OpenStack but is not an absolute prerequisite for OpenStack to provide cloud services. The OpenStack Bare Metal service documentation describes Ironic as a service for provisioning physical machines: OpenStack Ironic documentation. The OpenStack project also identifies Ironic as its bare-metal provisioning component: OpenStack Bare Metal project page.

QUESTION NO: 15

The following descriptions of the design and development ideas of OpenStack are wrong:

A. Must use Python language for development

B. Extensive use of plug-in methods for architecture design and implementation

C. Composed of multiple independent projects, each project contains multiple independent service components

D. No center, stateless architecture

ANSWER: A

Explanation:

“Must use Python language for development” is the incorrect description. Python has historically been the primary implementation language for many OpenStack services and is central to much of the project’s tooling, libraries, and contributor workflow. However, OpenStack’s architecture and open-source governance do not impose Python as an absolute requirement for every implementation, integration, client, driver, tool, or associated project. OpenStack is designed around interoperable services exposed through documented APIs, allowing software written in different languages to consume or extend cloud capabilities where appropriate. The essential concern is compatibility with OpenStack APIs, project requirements, security practices, testing, and community development processes—not a universal, mandatory programming-language rule. The OpenStack Contributor Guide describes the community’s development resources and workflows, while the Project Team Guide explains the project structure and governance model used across the OpenStack community. These materials support the distinction between Python’s major practical role in OpenStack and an unconditional requirement that all development use Python. See the OpenStack Contributor Guide and the OpenStack Project Team Guide.

QUESTION NO: 16

Which of the following descriptions about ManageOne Servicecenter is wrong?

A. Provide customers with cloud data center management solutions

B. Support heterogeneous and multi-resource pool management and orchestration

C. Customized services for customer scenarios

D. Provide customers with data center operation and operation and maintenance services

ANSWER: D

Explanation:

“Provide customers with data center operation and operation and maintenance services” is the wrong description of ManageOne ServiceCenter. ManageOne ServiceCenter is positioned as the cloud-management and service-delivery component of the ManageOne solution. Its focus is to expose cloud resources and services to users, support service catalogs and request fulfillment, coordinate resource provisioning across supported resource pools, and enable customers to tailor cloud services to their own business scenarios. In this role, it helps a cloud provider or enterprise IT organization deliver manageable, consumable cloud services rather than acting as an outsourced team that performs the customer’s daily data-center operations and maintenance work. Operational activities such as monitoring, alarm handling, routine maintenance execution, and operational process management are organizational or operations-platform responsibilities; they are not described as a customer-facing O&M service delivered by ServiceCenter itself. Therefore, the statement incorrectly characterizes the product as directly providing data-center operation and O&M services, instead of providing the management and orchestration capabilities that support cloud service delivery. Huawei presents ManageOne as its cloud management platform for managing cloud infrastructure and services; see Huawei ManageOne and Huawei ManageOne technical support resources.

QUESTION NO: 17

The entire mirroring process is:

1 Create a new virtual machine 2 Make a template

3 Install the virtual machine control system 4 Export template

The correct order is ()

A. 1 to 2 to 3 to 4

B. 2 3 1 4

C. 1 3 2 4

D. 1 4 2 3

ANSWER: C

Explanation:

The correct sequence is “1 3 2 4.” A template or image must be built from a usable, configured virtual machine, so the process begins by creating the new virtual machine. The required virtual-machine operating/control system is then installed, providing the software environment that will be captured for later reuse. Once the operating system installation and any baseline configuration are complete, the virtual machine can be converted into a template. This template represents a standardized source from which consistent new virtual machines can be provisioned. The final step is exporting that completed template, allowing it to be retained externally, transferred to another environment, or imported when needed for future deployments. This workflow preserves the dependency between stages: the template is created only after the VM has been prepared, and export occurs only after there is a completed template to export. Huawei Cloud documentation similarly describes images as reusable VM templates created from configured server environments and managed or exported after creation. See Huawei Cloud Image Management Service documentation and Huawei Cloud documentation on creating private images.

QUESTION NO: 18

In fusioncompute, the two virtual machines in the unified cluster are running normally, but they cannot communicate with each other. The possible reasons are:

A. The administrator adjusted the disk capacity of one of the VMs

B. The administrator adjusted the memory size of one of the VMs

C. The administrator adjusted the port group of one of the VMs

D. The administrator adjusted the description of one of the VMs

ANSWER: C

Explanation:

The administrator adjusted the port group of one of the VMs is correct because a FusionCompute virtual NIC obtains its Layer 2 network connectivity from the port group to which it is attached. A port group defines the logical network path and typically associates the VM NIC with a specific VLAN or network segment. If one VM is moved to a different port group, its traffic can be placed in a different broadcast domain or VLAN from the other VM. Both virtual machines can therefore remain powered on and otherwise operate normally, while direct communication between them fails because their packets no longer reach the same logical network. This is especially relevant after network reconfiguration, migration, or manual NIC changes: the administrator should verify that both VM NICs are connected to the intended port groups, that the configured VLAN IDs are consistent with the physical-switch configuration, and that the relevant virtual switch uplinks are available. Huawei FusionCompute documentation describes virtual networking and port groups as the connection mechanism between VM virtual NICs and logical networks; see the Huawei FusionCompute support documentation and Huawei’s FusionCompute product documentation portal.

QUESTION NO: 19

Which component or software cannot be connected to ManageOne OperationCenter?

A. ServiceCenter

B. FusionInsight Manager

C. eBackup

D. VMware vCenter

ANSWER: C

Explanation:

eBackup is the correct answer. ManageOne OperationCenter is the centralized O&M component of the ManageOne platform. Its interconnection framework is intended to integrate supported management systems so that OperationCenter can centrally monitor resources, collect alarms and performance data, and provide related operational views. eBackup is Huawei backup-management software, but it is not a supported direct interconnection target for ManageOne OperationCenter in this context. Backup operations are handled through the relevant backup-service integration and management path rather than by registering eBackup as an OperationCenter-connected management system. Therefore, eBackup cannot be connected directly to OperationCenter as described in the question. Huawei's ManageOne product documentation and support portal provide the authoritative compatibility and component-interconnection information for the applicable ManageOne release; these details can vary by version, so deployment planning should always use the documentation matching the installed release. See the Huawei ManageOne support documentation and the Huawei ManageOne product page.

QUESTION NO: 20

When designing the Huawei desktop cloud, which of the following should be considered when planning storage? (Multiple choice)

A. Type of storage media

B. Purpose of storage

C. Storage capacity

D. File system used for storage

ANSWER: A B C

Explanation:

Storage planning for a Huawei desktop cloud must begin with the **Type of storage media**, because the selected media tier determines the available IOPS, latency, throughput, reliability characteristics, and cost. Desktop workloads commonly include boot storms, login storms, and concurrent application activity, so matching media performance to the workload is essential. The **Purpose of storage** must also be identified, such as hosting desktop system disks, user disks, shared user data, images, backups, or logs. Each purpose has different persistence, availability, performance, and protection requirements. Finally, **Storage capacity** must be calculated from the number of desktops, disk sizes, user-data growth, image copies, snapshots, backups, redundancy overhead, and planned expansion. Proper capacity planning ensures that the platform can sustain both initial deployment and future growth without compromising desktop availability. Huawei's cloud-storage guidance describes selecting storage based on workload performance and capacity requirements, while its Workspace documentation provides context for desktop-cloud service architecture and resource planning. See Huawei Elastic Volume Service overview and Huawei Cloud Workspace overview.

QUESTION NO: 21 - (SIMULATION)

Which of the following descriptions about the logical network planning in FusionSphere OpenStack is correct: A. Limited to the network integrity requirements of the control node, the business network that the control node cannot deploy

B. The management network needs to be interoperable with the external-om and external-api network, and the external-om needs to communicate with each other.

externalapi network layer three interoperability

C. The control node supports access to IPSan and FCSan

D. Several physical deployments of the database, and its network card is connected to the external-om plane

ANSWER: See the explanation for the answer

Explanation:

Correct answers: A, B, and C.

A: Correct. To preserve the integrity and stability of the control-node networking, the business/tenant network is not deployed on the control node.
B: Correct. The management network must have Layer 3 connectivity with the `external-om` and `external-api` networks. The `external-om` and `external-api` networks also require Layer 3 interoperability.
C: Correct. FusionSphere OpenStack control nodes support connection to both IP SAN and FC SAN storage environments.
D: Incorrect. When the database is deployed on physical servers, its NIC is connected to the management network plane, not the `external-om` plane.

QUESTION NO: 22

Which of the following descriptions of CPS in FusionSphere OpenStack is correct? (Multiple choice)

A. MPS is responsible for Host management and service deployment management

B. SWH is responsible for package management

C. DS is responsible for Host OS configuration loading

D. NPC is responsible for Host information collection and reporting

ANSWER: A B C D

Explanation:

All four statements describe cooperating components of the CPS (Cloud Provisioning Service) framework used by FusionSphere OpenStack for automated host and service deployment. “MPS is responsible for Host management and service deployment management” is correct because MPS centrally coordinates the provisioning workflow and manages deployment-related host and service tasks. “SWH is responsible for package management” is correct because the software warehouse provides controlled storage and management of the software packages required during deployment. “DS is responsible for Host OS configuration loading” is correct because the deployment service delivers and loads the required operating-system configuration onto hosts as part of provisioning. “NPC is responsible for Host information collection and reporting” is also correct: the node-side provisioning client gathers host information and reports it to the management side, enabling the deployment process to identify and handle the target hosts. Together, these components separate orchestration, package storage, OS configuration delivery, and node information reporting, which supports consistent large-scale installation and expansion. Huawei product documentation and support materials for FusionSphere OpenStack are available through the Huawei FusionSphere OpenStack support page and the Huawei Enterprise Support portal.

QUESTION NO: 23

Which of the following descriptions of Huawei application virtualization and application scenarios are correct?

A. Soft terminals cannot use application virtualization. Xu Yaoh uses application virtualization in the scenario of using thin terminals

B. Application virtualization is recommended for mobile office scenarios to support multiple smart phone terminals to access the same application across platforms

C. It is not recommended to use application virtualization in the safe surfing scene

D. In a simple office scenario, when no user needs an independent desktop, application virtualization cannot be used

ANSWER: B

Explanation:

Application virtualization is recommended for mobile office scenarios to support multiple smart phone terminals to access the same application across platforms. In Huawei's virtual workspace architecture, applications are centrally installed, maintained, and executed in the data center or cloud workspace environment. The endpoint primarily provides the user interface and sends keyboard, mouse, touch, display, and audio data across the network. This centralized model enables an organization to deliver a consistent application service to users working from different endpoint types and locations, including mobile devices.

For mobile-office users, the key value is that the application environment remains centrally controlled instead of requiring the organization to install, patch, and protect a separate local copy on every smartphone. Users can connect to the published application through a supported client, while enterprise data and application processing remain in the managed virtual environment. This supports cross-platform access, simplifies application lifecycle management, and helps apply consistent access-control and data-security policies. Huawei's Cloud Workspace documentation describes the service as providing cloud desktops and applications for secure, flexible access, and Huawei's FusionAccess materials describe centralized virtual desktop and application delivery capabilities. See Huawei Cloud Workspace Product Description and Huawei FusionAccess.

QUESTION NO: 24

In the full-memory desktop memory plan, factors that need to be considered when calculating the memory size occupied by DomainU include: (multiple choice)

A. Single virtual machine memory

B. Memory reuse factor

C. IO Tailor running memory

D. Template parent volume size

ANSWER: A B C D

Explanation:

The DomainU capacity calculation for a full-memory desktop deployment must combine the capacity consumed by each desktop with the platform-level capacity required to operate the desktop environment. **Single virtual machine memory** is the baseline consumption per provisioned desktop and therefore directly determines the aggregate desktop footprint. The **memory reuse factor** is applied during sizing to account for the planned reuse or overcommit model, so the design reflects the effective memory demand rather than simply multiplying an idealized value without considering the deployment policy.

IO Tailor running memory must also be reserved because this service requires runtime memory while providing the I/O optimization functions used by the desktop solution. Finally, the **template parent volume size** is part of the required DomainU capacity planning because full-memory desktop provisioning depends on the parent/template resources from which desktops are created. Including both per-desktop demand and shared platform/template requirements produces a capacity plan that can support desktop creation and steady-state operation. Huawei Desktop Cloud documentation describes the platform architecture and sizing considerations, while Huawei support resources provide product documentation and deployment guidance: Huawei Desktop Cloud and Huawei FusionAccess Support.

QUESTION NO: 25

Which configuration and maintenance interface must be used in the process of desktop cloud construction and maintenance?

A. FusionManager

B. FusionAccess

C. FusionCompute

D. FusionStorage

ANSWER: B

Explanation:

FusionAccess is the required configuration and maintenance interface because it is Huawei's virtual desktop infrastructure and desktop-cloud solution platform. In a FusionAccess desktop-cloud deployment, administrators use the FusionAccess management functions to plan and configure desktop services, create and manage desktop pools, associate users with desktops, deliver virtual applications or desktops, and perform routine operational maintenance. It is the management layer designed specifically for the desktop-cloud service lifecycle, from initial construction through ongoing administration and troubleshooting.

FusionAccess also coordinates the desktop access and service components required to provide end users with secure virtual desktop sessions. This makes it the operational entry point for tasks such as desktop provisioning, user access management, policy configuration, and monitoring the desktop-cloud environment. Huawei describes FusionAccess as a desktop-cloud solution intended to centrally manage and deliver virtual desktops, which directly aligns with the construction and maintenance activities named in the question. See Huawei's FusionAccess product overview and the FusionAccess support documentation portal for product and administration resources.

QUESTION NO: 26

Available partitions in OpenStack are logical partitions of physical resources (computing, storage, network), and are user-oriented resource sets

合 , can be divided

When using partitions, the principles of division are:

(Multiple choice)

A. All resource clusters in an availability zone must originate from a virtualized environment

B. Resource clusters in an availability zone can originate from multiple virtualization environments

C. The VLANs associated with the DVSs of all resource clusters in an availability zone must be the same

D. The VLANs associated with the DVS of all resource clusters in an available zone can be different

ANSWER: A C

Explanation:

"All resource clusters in an availability zone must originate from a virtualized environment" and "The VLANs associated with the DVSs of all resource clusters in an availability zone must be the same" express the required consistency boundaries for an availability-zone resource partition. An availability zone is intended to expose a coherent, selectable pool of compute, storage, and network resources to tenants. Keeping its resource clusters within one virtualization environment ensures that the zone is managed through a consistent virtualization domain, with compatible resource-management behavior and operational policies. Maintaining the same VLAN association for the distributed virtual switches across those clusters ensures that workloads placed in the same zone have a uniform network foundation and can use the zone's network resources consistently. This design supports predictable instance scheduling, network provisioning, tenant isolation, and fault-domain planning. In OpenStack, availability zones are administrative groupings used to control where instances and related resources are deployed; their purpose is to provide clear infrastructure boundaries rather than an arbitrary mix of independently configured environments. See the OpenStack Nova availability-zone administration guidance at OpenStack Nova Availability Zones and the OpenStack networking overview at OpenStack Networking Administration Guide.

QUESTION NO: 27

The thin terminal does not contain an operating system, but is just a hardware device used to present remote virtual desktop display content and intervene peripherals.

Prepare

A. Yes

B. Wrong

ANSWER: B

Explanation:

Wrong is correct because the statement describes a zero client rather than a thin client. In a VDI architecture, a thin client is an endpoint device with a lightweight operating system or embedded firmware/software stack. That local software is needed to boot the device, establish a network connection, authenticate the user, run the remote-display protocol client, process peripheral redirection, and apply endpoint configuration or security policies. The user's primary desktop operating system and applications run remotely on the virtual desktop, but this does not mean that the thin terminal itself has no operating system at all.

Huawei Cloud Workspace documentation describes the service as providing cloud desktops that users access through endpoint devices and clients, with the endpoint responsible for connecting to and displaying the remote desktop session. This aligns with the thin-client model: local endpoint software supports access while compute resources and the user desktop are centralized. Huawei's product description is available at Huawei Cloud Workspace Product Introduction. For further Huawei Cloud Workspace documentation on client access and endpoint use, see Huawei Cloud Workspace User Guide.

QUESTION NO: 28

Which of the following storage subsystem parameters can affect the experience of Huawei desktop cloud users?

A. Hard disk performance

B. Read and write I/O lower disk ratio

C. Hard disk capacity

D. Read and write I/O ratio

ANSWER: A B D

Explanation:

Hard disk performance, Read and write I/O lower disk ratio, and Read and write I/O ratio directly influence the latency and throughput available to virtual desktops. Desktop-cloud workloads are highly sensitive to storage response time because operating-system booting, application launches, profile loading, paging, antivirus scans, and simultaneous logons all generate frequent I/O. The performance tier and capability of the underlying disks therefore determine whether those requests can be serviced promptly during both normal operation and peak periods.

The Read and write I/O lower disk ratio is a storage-planning parameter used to account for the amount of I/O demand placed on the disk subsystem. It helps ensure that the configured storage resources can sustain the aggregate desktop workload rather than merely meeting nominal virtual-machine allocation figures. The Read and write I/O ratio is equally important because read-heavy and write-heavy workloads exercise storage differently; accurate characterization is needed to select suitable media, cache behavior, and performance capacity. Huawei describes its Cloud Desktop solution as centrally delivering desktop resources and emphasizes the importance of the supporting infrastructure: Huawei Cloud Desktop. Huawei's storage documentation also discusses performance concepts and workload planning for storage systems: Huawei OceanStor Documentation.

QUESTION NO: 29

Which of the following descriptions about the disk configuration in FusionSphere OpenStack is correct?

A. Cinder can be connected to FC SAN storage to provide block storage services

B. Disk expansion has an impact on Cinder components

C. Remote storage only supports FCSAN

D. Disk configuration can be modified uniformly

ANSWER: A

Explanation:

Cinder can be connected to FC SAN storage to provide block storage services is correct. In FusionSphere OpenStack, Cinder is the block-storage component. It manages the lifecycle of volumes and connects compute instances to backend storage through supported Cinder drivers. Fibre Channel SAN is a supported enterprise block-storage access method: the Cinder driver communicates with the storage array to create, map, attach, detach, and delete volumes, while the compute host uses Fibre Channel connectivity to access the mapped block device. This model enables virtual machines to use centrally managed SAN LUNs as persistent disks, rather than relying only on local disks. Fibre Channel is especially common in production environments because it provides dedicated storage-network connectivity, multipathing capabilities, and integration with enterprise storage arrays. The OpenStack Cinder documentation describes Huawei's Cinder driver and its support for storage-protocol configurations, including Fibre Channel attachment scenarios. Huawei's FusionSphere OpenStack product information also identifies the platform as an OpenStack-based cloud solution providing enterprise storage and virtualization capabilities. See the OpenStack Cinder Huawei driver documentation and Huawei FusionSphere OpenStack product support page.

QUESTION NO: 30

Which of the following health checks cannot be performed by FAUSIONcare?

A. Business Virtualizer

B. FusionManager

C. FusionAccess

D. FusionStorage

ANSWER: A

Explanation:

Business Virtualizer is the correct answer. FusionCare is Huawei's maintenance and technical-support service framework for supported ICT products and solution components. Its health-check activities assess the operating status, configurations, alarms, capacity indicators, and risks of products within the supported solution scope. FusionManager, FusionAccess, and FusionStorage are Huawei platform products that can form part of Huawei cloud-infrastructure, desktop-cloud, and storage solution environments, allowing their relevant health information to be collected and assessed through the applicable FusionCare service processes. Business Virtualizer is not a Huawei product component covered by that FusionCare health-check scope, so it cannot be selected as a FusionCare health-check target. In practical service delivery, the exact inspection items and available automated checks also depend on the purchased service level, product version, deployment architecture, connectivity, and authorization. Huawei's Enterprise Support portal is the authoritative location for product documentation, service information, and support resources: Huawei Enterprise Support. Huawei's enterprise services overview also describes the company's support and lifecycle service approach: Huawei Enterprise Services.

QUESTION NO: 31

In the full-memory desktop memory plan, factors that need to be considered when calculating the memory size occupied by DomainU include: (multiple choice)

A. Single virtual machine memory

B. Memory reuse factor

C. I/O Tailor running memory

D. Template parent volume size

ANSWER: A B C D

Explanation:

In a full-memory desktop deployment, DomainU memory sizing must account for both each desktop's configured runtime requirement and the additional memory consumed by the full-memory architecture. **Single virtual machine memory** is the baseline capacity required for every desktop instance. The **memory reuse factor** is also essential because identical memory pages among desktops can be shared or deduplicated; applying the reuse factor produces a more realistic aggregate memory estimate than simply summing all configured VM memory. **I/O Tailor running memory** must be reserved because the I/O Tailor component runs as part of the desktop solution and consumes host memory independently of the user desktop workload. Finally, the **template parent volume size** is considered because full-memory desktops load template-related data into memory, so parent-image capacity affects the memory reservation calculation. Together, these inputs represent the desktop workload, memory-sharing efficiency, required supporting service overhead, and template-image footprint. This is consistent with Huawei FusionCompute's focus on sizing compute resources according to VM configuration and platform service requirements. See Huawei's Enterprise Support documentation portal and the FusionCompute product overview.

QUESTION NO: 32

In the overall deployment of OpenStack single-machine room modules, which of the following components is in the correct proportion: (multiple choice)

A. The relationship between AC and Open Stack is 1: N

B. The relationship between AC and Fabric is 1: N

C. The relationship between gateway group and gateway switch is 1:N

D. Fabric and Spine/Leaf have a 1:1 relationship

ANSWER: B C

Explanation:

In a single-data-center OpenStack deployment using Huawei's controller-based fabric architecture, an AC is able to centrally manage multiple network fabrics. This 1:N relationship enables a common controller domain to apply consistent network policies, automate provisioning, maintain fabric inventory, and coordinate service deployment across separate fabric instances. This model supports centralized control while allowing the physical data-center network to scale through additional fabrics.

A gateway group is a logical service and redundancy grouping that can contain multiple gateway switches. The 1:N relationship provides scalable north-south connectivity and allows gateway resources to be organized as a group for service forwarding and high-availability design. Associating several gateway switches with one gateway group lets the deployment expand gateway capacity without changing the logical service grouping used by the cloud-network design. These relationships align with Huawei CloudFabric's centralized fabric-management approach and with OpenStack networking's use of logical network abstractions over physical infrastructure. See Huawei Enterprise Support and OpenStack Neutron documentation.

QUESTION NO: 33

Fusioncompute data storage expansion allows one data storage to manage multiple physical LUN spaces.

When data storage needs to be expanded,

By adding another physical LUN to the data storage or expanding the physical LUN and then expanding the data storage, the data storage can be flexibly expanded and the data storage scalability can be effectively

improved:

A. Yes

B. Wrong

ANSWER: A

Explanation:

Yes is correct. In FusionCompute, a datastore is a logical storage resource that can use physical storage capacity supplied through LUNs. Its capacity can be increased without replacing the datastore or rebuilding the virtual-machine storage layout. An administrator can either attach an additional physical LUN to the datastore, allowing the datastore to aggregate more available capacity, or first increase the capacity exposed by an existing physical LUN on the storage array and then perform the corresponding datastore expansion operation in FusionCompute. Both approaches support online, planned capacity growth and let the platform continue managing storage as one logical datastore. This capability is particularly useful when virtual-machine disks, templates, snapshots, or other datastore-resident resources consume available space over time. The essential operational requirement is that the underlying storage-side change is completed and visible to the host environment before the datastore is expanded in the virtualization management system. Huawei presents FusionCompute as its virtualization platform for managing compute and storage resources; see the Huawei FusionCompute product page and the Huawei Enterprise Support portal for product documentation and operational guidance.

QUESTION NO: 34

In FusionAccess R6 version, which of the following infrastructure component virtual machine words should use the linux operating system (multiple choice)

A. WI

B. HDC

C. DB

D. ITA

ANSWER: A B C D

Explanation:

In FusionAccess R6, WI, HDC, DB, and ITA are deployed as infrastructure component virtual machines using the Linux operating-system image specified by the FusionAccess deployment documentation. WI provides the web access and connection-entry function for desktop users. HDC is the central control component responsible for desktop lifecycle management, user connections, and resource scheduling. DB provides the database service required to store the platform's configuration and operational data. ITA provides the infrastructure integration and management functions needed by the FusionAccess platform. Deploying these components on the prescribed Linux operating system ensures that their software packages, service dependencies, security settings, and maintenance procedures match the supported FusionAccess R6 architecture. The operating-system choice is therefore made according to the role of each infrastructure VM rather than according to whether a user's delivered virtual desktop runs Windows or Linux. Huawei's FusionAccess product information and support portal provide the product documentation, version-specific release material, and deployment guidance used to validate this component architecture: Huawei FusionAccess product page and Huawei FusionAccess support portal.

QUESTION NO: 35

In Huawei's computing virtualization technology, memory reservation specifies the highest physical memory reserved by the virtual machine

A. Yes

B. Wrong

ANSWER: B

Explanation:

Wrong is correct because a memory reservation represents a guaranteed *minimum* amount of host physical memory made available to a virtual machine, rather than the highest amount it can use. This guarantee helps protect the VM from memory contention on the host: the hypervisor reserves the configured quantity so that the VM can maintain access to that baseline memory capacity when other workloads compete for resources. The VM may consume additional memory beyond its reservation when unreserved host memory is available and the applicable scheduling and memory-management policies permit it.

The upper boundary of a VM's memory consumption is controlled by its configured virtual-memory size and, where supported and configured, a memory limit. These are different resource-control concepts from reservation. Therefore, describing memory reservation as the "highest physical memory reserved" reverses its purpose. In Huawei virtualization resource management, administrators use reservations to provide a minimum resource guarantee for workloads with important performance or availability requirements. Huawei's FusionCompute product documentation is available through the Huawei FusionCompute support portal; Huawei also describes VM resource configuration concepts in its Elastic Cloud Server user documentation.

QUESTION NO: 36

Which of the following conditions are generally required before performing live migration of a virtual machine between hosts in a FusionCompute cluster? (Multiple choice)

A. The destination host has sufficient available resources

B. The participating hosts meet CPU compatibility requirements

C. The migration network between the hosts is available

D. The virtual machine must be powered off before migration

ANSWER: A B C

Explanation:

Live migration transfers a running virtual machine from a source host to a destination host with minimal interruption. The destination host must have sufficient available CPU and memory resources to run the virtual machine after migration. The source and destination hosts must also have compatible CPU capabilities according to the configured CPU compatibility policy. In addition, the migration network must be available so that memory pages and execution state can be transferred between the hosts.

A virtual machine does not need to be powered off for live migration. In fact, shutting down the virtual machine changes the operation into a cold migration or an offline relocation. Storage accessibility requirements depend on the migration mode and storage architecture, but a running VM is not normally powered off as a prerequisite for live migration. See the Huawei FusionCompute support portal for version-specific migration requirements.

QUESTION NO: 37

ManageOne Servicenter can be used in private cloud or public cloud.

A. Correct

B. Error

ANSWER: A

Explanation:

ManageOne ServiceCenter can be used in both private-cloud and public-cloud scenarios. It is the ManageOne component oriented toward cloud-service users and tenants, providing a unified self-service portal through which users can request, provision, view, and manage cloud resources and services. This service-management role is applicable regardless of

whether the underlying cloud is operated as an enterprise's private cloud or delivered as a public-cloud service. In a private cloud, ServiceCenter enables internal departments or tenants to consume pooled IT resources through standardized service catalogs and workflows. In a public-cloud scenario, it provides the tenant-facing service experience needed for self-service resource consumption and lifecycle management. Huawei describes ManageOne as a cloud management platform supporting cloud-service operations and management across cloud environments, with ServiceCenter serving the service-consumption side of the platform. Therefore, the statement accurately reflects the intended deployment scope of ManageOne ServiceCenter. See Huawei's ManageOne product information and the ManageOne documentation portal for product and documentation resources.

QUESTION NO: 38

In Huawei virtualization, what is the correct statement about virtual machine live migration (multiple choice)

A. Ordinary users are not aware of the process of virtual machine hot migration

B. Only virtual machines running on virtualized storage can be live migrated

C. Virtual machine hot migration uses memory alternative migration technology

D. Virtual machine hot migration uses memory compression technology to improve migration efficiency

ANSWER: A C D

Explanation:

Live migration is designed to move a running virtual machine from one host to another with service interruption minimized to a very short, generally imperceptible interval. Therefore, the statement that ordinary users are not aware of the process of virtual machine hot migration correctly describes its intended service-continuity outcome. Huawei FusionCompute live migration transfers the VM execution context while the VM remains running. It uses iterative memory-copy processing: memory pages are copied to the destination repeatedly while the source VM continues to execute, and only the remaining changed pages and CPU state are switched during the final cutover. This is the meaning of "memory alternative migration technology" in the question. Memory compression can also reduce the amount of memory data sent over the migration network, improving transfer efficiency and helping reduce migration duration, especially where bandwidth is constrained. These mechanisms support maintenance activities, host load balancing, and fault-avoidance operations while preserving application availability. Huawei's virtualization product and documentation resources are available through the Huawei FusionCompute support portal; Huawei also describes VM migration concepts in its cloud documentation.

QUESTION NO: 39

FusionSphere Open adds a FusionCompute resource pool. What are the parameters that must be configured? (Multiple choice)

A. FusionCompute service IP address

B. User name and password for docking

C. VXLAN DVS name

D. Glance service IP address

ANSWER: A B

Explanation:

Adding a FusionCompute resource pool requires the FusionCompute service IP address and the user name and password for docking. The service IP address identifies the FusionCompute management endpoint that FusionSphere Open must contact to register the virtualization resource pool, obtain inventory information, and perform lifecycle operations on compute, storage, and network resources. The docking credentials provide the authenticated connection between the cloud-management platform and FusionCompute. They must belong to an account with the permissions needed to access and manage the intended FusionCompute resources. Together, these settings establish both the network reachability and

authorization required for the resource-pool integration to be created successfully. After the platform validates the endpoint and credentials, it can synchronize the available resources and make them available to cloud services. Huawei documentation for FusionSphere OpenStack and its related cloud-computing products is available through the Huawei FusionSphere OpenStack support portal. General product documentation and version-specific deployment guides can also be located in the Huawei Enterprise documentation center.

QUESTION NO: 40

The following network types supported by Neutron are:

(Multiple choice)

A. Flat

B. VLAN

C. VXLAN

D. GRE

ANSWER: A B C D

Explanation:

OpenStack Neutron supports several Layer-2 network segmentation mechanisms through its Modular Layer 2 plug-in framework. **Flat** networks provide unsegmented connectivity on a physical network, normally for cases where tenant separation is not required at the Layer-2 level. **VLAN** networks use IEEE 802.1Q tags to create isolated logical networks across shared physical switching infrastructure. **VXLAN** is an overlay network technology that encapsulates Layer-2 traffic over an IP underlay and uses a 24-bit virtual network identifier, allowing substantially more tenant segments than traditional VLAN identifiers. **GRE** is also supported as an overlay network type in Neutron deployments and encapsulates tenant traffic across an IP network between tunnel endpoints. These network types are configured by the Neutron ML2 type drivers; the available choices in a specific cloud depend on the enabled driver configuration and supporting physical or tunnel infrastructure. Huawei Cloud Computing exam objectives commonly use the standard OpenStack Neutron network-type model, in which all four listed types are recognized network types. OpenStack's networking administration documentation describes ML2 network type drivers and their segmentation models: Neutron ML2 configuration guide. The Neutron configuration reference also documents the flat, VLAN, VXLAN, and GRE type-driver settings: ML2 configuration reference.

QUESTION NO: 41

Which of the following descriptions about OpenStack Cinder volume types are correct? (Multiple choice)

A. Extra specifications can be associated with a volume type.

B. Cinder Scheduler can use volume type information to select a storage backend.

C. A volume type can be associated with QoS specifications.

D. A volume type defines the vCPU and memory specifications of an instance.

ANSWER: A B C

Explanation:

A Cinder volume type defines a category of block-storage service. Administrators can associate extra specifications with a volume type to describe capabilities required from a storage backend, such as performance tier, storage protocol, or replication capability. The Cinder scheduler uses these specifications together with backend capabilities to select an appropriate backend for volume creation.

Volume types can also be associated with QoS specifications to apply storage performance controls where the selected driver and backend support them. A volume type is not a Nova flavor: flavors define instance compute specifications, whereas volume types define block-storage service characteristics.

See the OpenStack Cinder volume types documentation.

QUESTION NO: 42

Which of the following statements about the network in the VPC is correct:

A. Apply for a floating IP directly on the network and bind it to the virtual switch port to achieve access to the virtual machine from the private network

B. A routing network within a VPC and an external network can be in the same vlan

C. After the routing network is created, no other configuration is required, and the virtual machines in the network can directly access the external network

D. VMs on a direct network can directly access the external network

ANSWER: D

Explanation:

VMs on a direct network can directly access the external network is correct because a direct network provides Layer 2 connectivity from a VM's virtual NIC to the associated external physical network. The direct-network VLAN is mapped to an external network, allowing VM traffic to leave the VPC through that connected network without requiring the VPC routing-network gateway/NAT path. In practical deployments, the VM must still be configured with valid network parameters, such as an IP address, subnet mask, and an appropriate default gateway supplied by or reachable on the external network. External routing and security policies must also permit the intended traffic. This design is useful when workloads need direct attachment to a physical network segment, for example to communicate with existing enterprise resources or to receive externally routable addressing through that segment. Huawei VPC documentation describes VPC networking and connectivity concepts at Huawei Cloud VPC User Guide. Huawei's VPC service overview also explains that VPCs provide isolated virtual networks and connectivity capabilities for cloud resources: Huawei Cloud Virtual Private Cloud.