

HCIE - Cloud (Huawei Certified Internetwork Expert-Cloud)

Huawei H13-531

Version Demo

Total Demo Questions: 17

Total Premium Questions: 176

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

In a FusionCompute environment, which of the following technologies allows a virtual machine to directly use a physical PCI device installed on the host?

- A. PCI pass-through
- B. Memory multiplexing
- C. Virtual machine snapshot
- D. Thin provisioning

ANSWER: A

Explanation:

PCI pass-through is correct. PCI pass-through maps a physical PCI device, such as a network adapter, GPU, or other supported accelerator, directly to a virtual machine. The guest operating system can then use the device with reduced virtualization overhead and near-native performance.

Because the device is dedicated to the virtual machine, pass-through also introduces lifecycle restrictions. For example, features such as live migration may not be available for a VM using a pass-through device, depending on the device type and platform configuration. Hardware compatibility, BIOS virtualization settings, and host driver support must be verified before deployment. Refer to the [Huawei FusionCompute support documentation](#) for version-specific device pass-through requirements.

QUESTION NO: 2

Which of the following description about OpenStack/FusionSphere OpenStack are wrong? (Multiple choice)

- A. OpenStack is a community version of the virtualization hypervisor, can be based on it and configure server equipment to realize the virtual environment building.
- B. Huawei FusionSphere OpenStack cloud operating system products are based on the community version of OpenStack for commercial enhancement products.
- C. OpenStack and virtualization have obvious differences and boundaries, is the open source version of the cloud operating system.
- D. OpenStack is the cloud, the same concept and connotation.

ANSWER: A D

Explanation:

The statements "OpenStack is a community version of the virtualization hypervisor" and "OpenStack is the cloud, the same concept and connotation" are the incorrect descriptions. OpenStack is not a hypervisor. It is an open-source cloud-platform project that provides software services for managing compute, networking, storage, identity, images, and other cloud resources. A cloud deployment using OpenStack can integrate hypervisors such as KVM to run virtual machines, but the OpenStack control plane and a hypervisor serve fundamentally different roles. The OpenStack project describes its software as cloud infrastructure components that control large pools of compute, storage, and networking resources through APIs and dashboards, rather than as virtualization-hypervisor software. See the [OpenStack software overview](#) and [OpenStack Installation Guide overview](#). Likewise, "the cloud" is a broader service and architectural concept, encompassing delivery models, operational processes, applications, infrastructure, and user-facing services. OpenStack is one possible cloud operating-system platform used to build and manage an IaaS cloud; it is not synonymous with every cloud or with cloud computing itself. This distinction is essential when separating virtualization technology from cloud resource orchestration and service management.

QUESTION NO: 3

In Huawei cloud computing solution, in order to enhance the virtual machine's graphics processing power; you can configure the GPU to the virtual machine, which of the following configuration can be included? (Multiple Choice)

- A. GPU pass-through
- B. GPU virtualization
- C. GPU sharing
- D. GPU aggregation

ANSWER: A B

Explanation:

GPU pass-through is a supported GPU attachment approach for graphics- or compute-intensive virtual machines. It maps a physical GPU directly to a virtual machine, allowing the guest to use the GPU with near-native performance. This approach is appropriate where a workload requires dedicated accelerator capacity, such as high-end visualization, engineering graphics, or GPU computing.

GPU virtualization is also a supported configuration approach. It virtualizes a physical GPU into virtual GPU resources that can be allocated to virtual machines, enabling controlled use of GPU acceleration in a virtualized environment. This makes it possible to deliver graphics acceleration while managing GPU resources through the cloud platform's virtualization layer. Huawei Cloud documents GPU-accelerated ECS offerings for workloads such as graphics rendering, AI, and scientific computing, and its Elastic Cloud Server documentation describes the platform capabilities and instance families used to provide accelerated compute resources. See the [Huawei Cloud Elastic Cloud Server product page](#) and [Huawei Cloud ECS User Guide](#).

QUESTION NO: 4

In a VPC, what is the main function of a route table?

- A. To determine the forwarding path for traffic leaving a subnet
- B. To allocate IP addresses to virtual machines
- C. To store virtual machine image files
- D. To create user accounts and roles

ANSWER: A

Explanation:

To determine the forwarding path for traffic leaving a subnet is correct. A route table contains destination routes and corresponding next hops. When an instance sends traffic to an address outside its local subnet, the VPC uses the applicable route entry to determine whether traffic should be forwarded to a virtual router, VPN connection, peering connection, NAT gateway, or other network target.

Route tables do not allocate IP addresses and do not replace security groups or network ACLs. Address allocation is performed by subnet and IP address management functions, while security policies control whether traffic is permitted. Huawei Cloud VPC documentation describes route tables as a mechanism for controlling traffic forwarding within and outside VPC networks. See the [Huawei Cloud VPC User Guide](#).

QUESTION NO: 5

What are the advantages of deploying on a virtual server comparing to the busniess deploy on a physical server?

- A. High reliability
- B. High performance
- C. High cost
- D. Green energy saving

ANSWER: D

Explanation:

Green energy saving is the correct advantage. Virtualization allows multiple independent virtual servers and workloads to run on a smaller number of physical hosts. Rather than assigning each application to a separately powered, often underutilized physical server, compute, memory, storage, and network resources can be pooled and allocated dynamically to virtual machines. This server consolidation reduces the total number of devices that must be powered, cooled, and maintained. As a result, organizations can lower data-center electricity consumption and cooling demand while improving average hardware utilization. When demand decreases, workloads can be consolidated further and unused hosts may be placed into lower-power states, producing additional energy savings. This is a core operational benefit of virtualized cloud infrastructure and aligns with Huawei Cloud Elastic Cloud Server concepts, in which compute capacity is delivered as flexible virtual server instances rather than dedicated physical machines. See the Huawei Cloud [Elastic Cloud Server product overview](#) and Huawei's [green data center information](#) for further context on virtualized compute services and energy-efficient data-center operations.

QUESTION NO: 6

In FusionCompute, which of the following is the main purpose of a virtual machine snapshot?

- A. Preserving the point-in-time state of a virtual machine
- B. Balancing CPU load between hosts
- C. Automatically restarting a virtual machine after host failure
- D. Providing remote data replication between sites

ANSWER: A

Explanation:

Preserving the point-in-time state of a virtual machine is correct. A snapshot records the state of a virtual machine and its virtual disks at a specific time. It can be used before a configuration change, software installation, or maintenance activity to provide a short-term rollback point.

A snapshot is not a replacement for an independent backup or disaster-recovery solution. It does not provide host resource scheduling, virtual-machine high availability, or long-distance data replication. Snapshot retention must also be controlled because long-lived snapshots can consume storage capacity and affect disk performance. Refer to the [Huawei FusionCompute support documentation](#).

QUESTION NO: 7

Which of the following types of service migration are supported by the Huawei Virtualization Migration Solution? (Multiple Choice)

- A. Windows Block-Level
- B. AIX Block-Level
- C. Linux Block-Level
- D. Solaris Block-Level

ANSWER: A C**Explanation:**

The Huawei Virtualization Migration Solution supports block-level service migration for both Windows and Linux workloads. Block-level migration copies disk blocks from the source system to the destination environment and supports incremental synchronization after the initial replication. This approach minimizes the service interruption required during cutover, because data changes can be synchronized repeatedly while the source workload remains in service. It is therefore appropriate for moving supported physical or virtual Windows and Linux servers into the Huawei virtualization environment with controlled downtime.

Windows Block-Level and Linux Block-Level are the supported migration categories in this solution's service-migration scope. The operating-system-specific migration agent and its associated compatibility requirements must be observed during implementation, including supported OS versions, disk layouts, network connectivity, and cutover procedures. Huawei's migration documentation describes the agent-based migration workflow, including full replication, synchronization, and cutover, and identifies Windows and Linux as supported source operating-system families. See the [Huawei Cloud Server Migration Service User Guide](#) and Huawei's [Virtualization Migration support resources](#) for migration guidance and compatibility information.

QUESTION NO: 8

Which layers as the following are included in the Huawei Cloud Data Center Service Architecture Hierarchical Model? (Multiple Choice)

- A. The infrastructure layer, such as servers, storage, networking products and so on.
- B. Virtualization layer, the realization of virtualization resource pool.
- C. Cloud service layer, blending virtualization resource pool, to realize the centralized management and allocation of resources.
- D. Application layer, support all types of cloud services.

ANSWER: A B C D**Explanation:**

The Huawei cloud data-center service architecture is conventionally organized as a layered model that progresses from physical resources to cloud-facing applications and services. "The infrastructure layer, such as servers, storage, networking products and so on." is the physical foundation: compute, storage, and network devices supply the underlying capacity on which the cloud is built. "Virtualization layer, the realization of virtualization resource pool." is included because virtualization abstracts those physical devices and aggregates them into shareable pools of compute, storage, and network resources.

"Cloud service layer, blending virtualization resource pool, to realize the centralized management and allocation of resources." represents the cloud-management capability above virtualization. It coordinates pooled resources, provides centralized management, and enables resource provisioning and allocation for tenants and services. "Application layer, support all types of cloud services." is also part of the model because applications and cloud services consume the managed resources and expose usable business capabilities to users.

This layered approach is consistent with Huawei Cloud's description of cloud infrastructure services, where virtualized compute resources are delivered from underlying infrastructure through cloud management capabilities. See the [Huawei Cloud ECS product overview](#) and the [Huawei Cloud overview](#).

QUESTION NO: 9

Which of the following scenarios is used in the Huawei business continuous disaster recovery product BCManager typical disaster recovery protection solution?

- A. Database Scene, such as Oracle, DB2, SQL Server and Exchange,
- B. Virtualized application environments, such as FusionSphere and VMware.

C. Local file system disaster recovery protection scenarios.

D. LUN and NFS disaster recovery protection scenarios.

ANSWER: A B C D

Explanation:

Huawei BCManager is designed to orchestrate and manage end-to-end disaster recovery across heterogeneous application, virtualization, and storage environments. Therefore, the listed database workloads are valid protection scenarios: BCManager supports coordinated disaster-recovery processes for enterprise applications and databases, including Oracle, DB2, Microsoft SQL Server, and Exchange-related workloads. Virtualized application environments are also a core use case because BCManager can provide centralized recovery orchestration for cloud and virtual-machine environments, including Huawei FusionSphere and VMware platforms.

Local file system protection is included because disaster recovery solutions must protect file-based application data and enable controlled recovery according to recovery plans. Storage-level protection scenarios using LUNs and NFS are likewise typical: block and file storage data can be replicated and incorporated into a recovery workflow, allowing BCManager to coordinate the recovery of the application, host, and associated storage resources. These capabilities reflect BCManager's role as a business-continuity management platform rather than a product limited to one hypervisor, database engine, or storage access method. Huawei describes BCManager as a solution for centralized disaster-recovery management and orchestration across applications and infrastructure. See [Huawei BCManager product information](#) and [Huawei BCManager documentation and support resources](#).

QUESTION NO: 10

B companies need to enable the access control on the cluster's virtual machine resource, the following description about the virtual machine resources, which are the wrong? (Multiple choice)

A. The network QoS policy provides bandwidth configuration control ability and does not support traffic throttling between virtual machines on the same host.

B. If the host memory is sufficient, the virtual machine's memory requests can be unrestricted; if the host memory is low, then allocate according to the memory level and share

C. You can assign CPU individual physical cores to a single virtual machine.

D. You can set the resources' usage limit on Virtual machines, including CPU, memory, network, disk IOPS and so on.

ANSWER: A D

Explanation:

The statements "The network QoS policy provides bandwidth configuration control ability and does not support traffic throttling between virtual machines on the same host." and "You can set the resources' usage limit on Virtual machines, including CPU, memory, network, disk IOPS and so on." are the incorrect descriptions. FusionCompute network QoS is intended to control virtual-machine network traffic through bandwidth-related policies, including traffic control for VM communication paths; it is not limited merely because communicating VMs reside on the same host. Therefore, claiming that same-host VM traffic cannot be throttled is inaccurate.

VM resource QoS also provides differentiated controls according to resource type rather than one universal "usage limit" for every listed resource. CPU, network, and disk performance can be constrained using the applicable QoS parameters, while memory management is based on mechanisms such as reservation, shares, and memory priority rather than a general VM memory-limit setting in the sense asserted by the statement. Accurate resource-control configuration must therefore use the supported control model for each resource. Huawei's FusionCompute product and documentation resources provide the relevant virtualization and resource-management guidance: [Huawei FusionCompute technical support](#) and [Huawei FusionCompute product information](#).

QUESTION NO: 11

When you enable Huawei Guest NUMA function, there are some constraints, which of the following is not a constraint to enable this function?

- A. The number of cores for a virtual machine CPU must be an integer multiple of the number of host CPUs or an integer multiple of the number of host CPU threads
- B. Enable cluster memory multiplexing function or host CPU resource isolation mode
- C. Turn on the NUMA Support option in the host's advanced BIOS settings
- D. After setting the cluster GuestNUMA policy, you need to restart the virtual machine on the host

ANSWER: B

Explanation:

Enable cluster memory multiplexing function or host CPU resource isolation mode is the correct choice because enabling either of these cluster or host resource-management features is not an inherent prerequisite for enabling Guest NUMA. Guest NUMA is intended to expose a virtual NUMA topology to a virtual machine so that the guest operating system and NUMA-aware applications can make placement decisions that better align with the physical host's NUMA architecture. Its core enablement considerations concern the host's NUMA capability, the virtual machine's CPU topology, and applying the changed topology to the running environment. Memory multiplexing and CPU resource-isolation settings are separate resource-management controls. They can affect scheduling, capacity planning, and operational policy, but simply enabling one of them is not a required condition for Guest NUMA activation. Administrators should evaluate such settings according to workload performance and consolidation requirements rather than treating them as a mandatory Guest NUMA configuration step. Huawei's FusionCompute documentation and product resources provide the relevant virtualization configuration and resource-management guidance: [Huawei Enterprise Support](#) and [Huawei FusionCompute](#).

QUESTION NO: 12

In order to ensure the normal operation of the company's business, project design need to consider which of the following factors? (Multiple choice)

- A. Reasonable investment
- B. Business Continuity
- C. Future evolution
- D. Flexible expansion

ANSWER: A B C D

Explanation:

Project design that supports normal business operation must balance immediate service requirements with resilience, capacity planning, and long-term sustainability. **Reasonable investment** is necessary because the architecture should deliver required availability and performance without unnecessary cost, while reserving budget for essential reliability measures. **Business Continuity** is a core design objective: workloads, data, and operational processes need protection against faults or disruptions so that critical services can continue or recover within business requirements. **Future evolution** is equally important because enterprise applications, service models, and technical requirements change over time; the solution should accommodate upgrades, migration, and new capabilities without disrupting existing operations. **Flexible expansion** ensures that compute, storage, network, and related cloud resources can be increased as business demand grows, preventing capacity limitations from affecting production services. Together, these considerations reflect the lifecycle approach expected in cloud solution design: economically appropriate deployment, continuity planning, scalable architecture, and an evolution path. Huawei Cloud guidance describes elastic resource scaling and reliability-oriented architecture as foundational cloud capabilities; see [Huawei Cloud Elastic Cloud Server](#) and [Huawei Cloud Enterprise Solutions](#).

QUESTION NO: 13

In FusionSphere solutions, VSA ManageDVS network is in the plane of the network with management can not be exchanged.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because the statement inaccurately describes the relationship between the VSA ManageDVS network and management-plane communication. In a FusionSphere deployment, a distributed virtual switch is used to provide the required logical network connectivity for the relevant management services and components. Management traffic must have properly designed, reachable paths between the platform components that administer, monitor, and coordinate cloud resources. Network segmentation, VLANs, security groups, and access-control policies can be used to isolate management traffic from service, storage, or external traffic, but segmentation does not mean that management communication is inherently impossible. Rather, the management network is planned so authorized management components can exchange the traffic necessary for normal operation while unnecessary access is restricted. Therefore, it is not accurate to state that the VSA ManageDVS network is on a management-related network plane in which communication cannot occur. FusionSphere network design relies on logical isolation together with controlled connectivity for management services. Huawei's FusionSphere product information is available at [Huawei FusionSphere](#), and official documentation resources can be accessed through the [Huawei Enterprise Support](#) portal.

QUESTION NO: 14

In the Huawei desktop cloud solution, about the description of the full memory virtual machine, which is incorrect?

- A. Full memory virtual machines do not support hot migration
- B. Full memory virtual machines do not support snapshots
- C. Deployed full memory virtual machine templates allow converting to virtual machines
- D. Full memory virtual machine support shutdown reduction

ANSWER: C

Explanation:

"Deployed full memory virtual machine templates allow converting to virtual machines" is the incorrect statement. In Huawei Desktop Cloud, a full-memory virtual machine is a specialized VM type designed for scenarios where desktop performance benefits from retaining its disk data in host memory. Because this VM type has specific deployment and lifecycle constraints, a template that has been deployed as a full-memory VM template cannot subsequently be converted into an ordinary virtual machine. This restriction preserves the resource model and operational characteristics assigned when the full-memory template is deployed.

Administrators therefore need to select the full-memory VM type deliberately during template planning and deployment, rather than expecting to repurpose an already deployed full-memory template as a standard VM later. This is particularly relevant to capacity planning, because full-memory desktops consume memory resources according to their special operating model. Huawei product documentation and support resources for Desktop Cloud and related cloud-management features are available through the [Huawei Enterprise Cloud Computing support portal](#) and the [Huawei Enterprise documentation portal](#).

QUESTION NO: 15

In a virtualized cloud computing platform, which component is responsible for abstracting physical hardware resources and providing virtual CPU, memory, storage, and network resources to virtual machines?

- A. Hypervisor

- B. DNS server
- C. Database server
- D. Load balancer

ANSWER: A

Explanation:

Hypervisor is correct. The hypervisor runs on the physical host and abstracts underlying hardware resources so that multiple isolated virtual machines can share the same server. It allocates virtual CPUs, virtual memory, virtual disks, and virtual network interfaces to guest operating systems while controlling access to physical devices.

In FusionCompute, the virtualization layer provides this resource abstraction and allows physical server resources to be pooled and centrally managed. The management platform configures and monitors virtual machines, but the hypervisor is the component that performs the fundamental virtualization and resource-isolation functions on the host. Huawei FusionCompute product information is available at [Huawei FusionCompute](#).

QUESTION NO: 16

B company plans to build a cloud data center, switch the existing office system to the cloud computing environment to run, when planning the total storage capacity, which of the following factors need to be considered? (Multiple choice)

- A. RAID
- B. LUN
- C. IOPS
- D. Hard disk capacity

ANSWER: A B D

Explanation:

Storage-capacity planning must begin with the physical capacity supplied by the selected disks. Disk count and individual hard disk capacity determine the raw-capacity pool available to the storage system. This raw value is not the capacity that can be presented to cloud workloads, however. RAID must be included because its protection scheme consumes capacity for parity or mirrored copies; therefore, the selected RAID level determines the conversion from raw capacity to usable capacity. Spare-disk and system-reserved capacity should also normally be allowed for in a production design.

LUN planning is equally relevant because LUNs are the logical storage resources presented to hosts or virtualization platforms. The required LUN sizes, their expected growth, and whether thin provisioning is used establish the usable capacity that services require. In a cloud data center, the plan should aggregate present LUN demand, reserve growth headroom, and ensure that the resulting logical allocation can be safely supported by the RAID-protected physical disk pool. Huawei storage documentation describes LUNs as the logical units used to provide storage to applications and identifies RAID as a mechanism for data protection and capacity organization. See [Huawei OceanStor Documentation](#) and [Huawei Storage Products](#).

QUESTION NO: 17

Huawei FusionSphere solutions, which is wrong?

- A. CNA can manage virtual machines on compute nodes.
- B. VRM can dynamically adjust the resources within the cluster.
- C. VSA can provide virtual firewall, DHCP function of the system service components, do not rely on FusionManager services for virtual machine deployment.

D. FusionManager northbound interface, provide common resource management, virtual machine management, backup management, disk management, network management and other business functions.

ANSWER: C

Explanation:

VSA can provide virtual firewall, DHCP function of the system service components, do not rely on FusionManager services for virtual machine deployment. is the incorrect statement because it combines VSA service capabilities with an inaccurate assertion about the management and deployment relationship. In a FusionSphere cloud solution, FusionManager is the cloud-management layer used to orchestrate and provision cloud resources and services, including virtual-machine lifecycle operations exposed to tenants and administrators. System service components such as a Virtual Security Appliance are integrated into that managed cloud environment; their deployment and lifecycle are coordinated through the platform's management and orchestration mechanisms rather than being treated as independent of FusionManager when cloud services are deployed. FusionCompute provides the underlying virtualization and resource-management functions, while FusionManager supplies the higher-level cloud service management, provisioning, operations, and northbound integration needed in a FusionSphere solution. Therefore, the claim that virtual-machine deployment does not rely on FusionManager services makes the statement unsuitable as a description of the complete FusionSphere solution architecture. Huawei describes FusionSphere as a cloud platform integrating virtualization and cloud-management capabilities; see [Huawei FusionSphere](#) and Huawei's [FusionSphere OpenStack documentation portal](#).