

HCIE-Cloud Computing (Written) V2.0

Huawei H13-531 V2.0

Version Demo

Total Demo Questions: 78

Total Premium Questions: 786

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing Technologies	236
Topic 2, Huawei Cloud Stack	153
Topic 3, Cloud Computing Solution Design	54
Topic 4, Cloud Computing Solution Deployment	128
Topic 5, Cloud Computing Solution O&M	133
Topic 6, Cloud Computing Solution Troubleshooting	82
Total	786

QUESTION NO: 1

(Multiple Choice) In Huawei Desktop Cloud, which of the following are the factors in the performance triangle model that affects virtual machine density?

- A. capacity
- B. density
- C. workload
- D. QoS

ANSWER: A C D

Explanation:

The performance triangle used for Huawei Desktop Cloud sizing balances **capacity**, **workload**, and **QoS** when determining how many desktop virtual machines can be placed on a host. Capacity represents the available physical compute, memory, storage, and network resources. These resources establish the baseline from which desktop instances can be provisioned and supported. Workload describes the resource-consumption profile of the desktop users and applications, including factors such as processor intensity, memory usage, graphics demand, storage input/output operations, and concurrent activity. A host can support more lightly used task desktops than resource-intensive engineering or multimedia desktops.

QoS represents the service experience that must be maintained for users, such as responsiveness, acceptable login time, application performance, and display-protocol experience. The required QoS determines how much resource headroom must remain available during peak demand. Consequently, desktop density is derived by balancing the host's capacity against the workload characteristics while preserving the required QoS. Huawei's cloud-desktop documentation and product information provide the relevant context for desktop-resource planning and user-experience requirements: [Huawei Cloud Workspace User Guide](#) and [Huawei Desktop Cloud](#).

QUESTION NO: 2

【Multiple Choice】 Which of the following scenarios cannot be migrated using Rainbow?

- A. Migrating Alibaba Cloud's ECS to HUAWEI CLOUD
- B. Migrate the running virtual machine to FusionCompute.
- C. Migrate the virtual machine running in the Beijing data center to the virtual machine running in the Shanghai data center through the Internet on FusionCompute.
- D. The Rainbow server can access the source and destination virtual machines, but cannot access the management network of FusionCompute network.

ANSWER: A C D

Explanation:

Rainbow is a FusionCompute-oriented VM migration tool and requires a supported source environment, a reachable target FusionCompute environment, and the required management-plane connectivity for orchestration. Therefore, "Migrating Alibaba Cloud's ECS to HUAWEI CLOUD" is not a Rainbow migration scenario: cross-public-cloud ECS migration should use the relevant Huawei Cloud migration services and workflows rather than a FusionCompute Rainbow task. "Migrate the virtual machine running in the Beijing data center to the virtual machine running in the Shanghai data center through the Internet on FusionCompute" is also outside Rainbow's supported deployment model because an Internet-based, inter-data-center migration does not provide the required controlled migration connectivity and target environment assumptions. Finally, "The Rainbow server can access the source and destination virtual machines, but cannot access the management network of FusionCompute network" cannot be migrated because Rainbow must communicate with FusionCompute management components to obtain target resources, create and configure the destination VM, and control the migration workflow. Huawei's FusionCompute product information is available at [Huawei Enterprise Support: FusionCompute](#). For Huawei Cloud migration service guidance, see [Huawei Cloud ECS Migration documentation](#).

QUESTION NO: 3

Which of the following functions does CPU virtualization in Fusioncompute support? (multiple choice)

- A. Support heterogeneous servers
- B. CPU resource over-allocation
- C. QOS
- D. Support Windows and Linux virtual machines to increase CPU online

ANSWER: A B C D

Explanation:

FusionCompute CPU virtualization abstracts physical processor resources and presents configurable virtual CPUs to virtual machines. This abstraction allows a virtualization cluster to use supported heterogeneous server hardware, so hosts with differing processor generations or configurations can participate subject to FusionCompute compatibility requirements. It also enables CPU resource over-allocation: the total virtual CPU capacity assigned to virtual machines may exceed the physical CPU capacity, allowing workloads with noncoincident peaks to use hardware more efficiently. CPU QoS is supported to control and protect CPU-resource allocation through scheduling and resource-management policies, helping administrators prioritize workloads and prevent contention from affecting critical services. FusionCompute additionally supports online CPU addition (CPU hot add) for supported Windows and Linux guest operating systems. This permits the vCPU count of an eligible running virtual machine to be increased without shutting down the guest, improving operational continuity when additional processing capacity is needed. These capabilities are part of the platform's core compute virtualization and VM resource-management model. See Huawei's [FusionCompute product overview](#) and the [FusionCompute technical documentation portal](#) for product documentation and release-specific support details.

QUESTION NO: 4

Multiple choice) In FusionCompute, which of the following storage resources can be used to create virtualized data storage?

- A. FC-SAN
- B. IP-SAN
- C. FusionStorage Block
- D. NAS

ANSWER: A B C

Explanation:

FC-SAN, **IP-SAN**, and **FusionStorage Block** can be used to create virtualized data storage in FusionCompute. FusionCompute virtualized storage is designed around block-storage resources that are presented to hosts as storage LUNs and then abstracted into storage pools/datastores for virtual-machine disks. FC-SAN supplies those block devices through Fibre Channel, while IP-SAN supplies them over IP-based block-storage protocols such as iSCSI. FusionStorage Block is Huawei's distributed block-storage product and integrates with FusionCompute as a supported block-storage backend. This architecture allows FusionCompute to manage capacity and provide VM storage independently of the underlying supported block-storage implementation. Huawei describes FusionStorage Block as distributed block storage for cloud and virtualization workloads, and its FusionCompute documentation distinguishes block-storage resources used by virtualized storage from file-storage access methods. See Huawei's [FusionStorage product information](#) and the [FusionCompute technical documentation portal](#).

QUESTION NO: 5

(Multiple choice questions) FusionaccessWhich of the following logs are collected during system failure?

- A. User VM logs
- B. User terminal log
- C. FusionCompute running log
- D. Linux infrastructure virtual machine running log

ANSWER: A B C D

Explanation:

During a FusionAccess system failure, troubleshooting requires end-to-end evidence from every layer involved in virtual desktop delivery. **User VM logs** record the guest operating system and desktop-side events that affect application execution, agent behavior, session initialization, and user experience. **User terminal logs** provide endpoint-side information, including client connection attempts, protocol negotiation, and local device or client conditions. **FusionCompute running logs** are required to investigate the underlying virtualization platform, including host, virtual-machine, network, and storage operations that can affect desktop availability. **Linux infrastructure virtual machine running logs** capture service-side events from the FusionAccess infrastructure components deployed on Linux virtual machines, such as connection-broker, authentication, management, and supporting service activity. Collecting these log sets together provides the time-correlated evidence needed to trace a failure across the terminal, desktop VM, access platform, and virtualization infrastructure. Huawei's enterprise support portal provides product documentation and log-collection guidance for fault diagnosis: [Huawei Enterprise Support](#) and [Huawei FusionAccess Support](#).

QUESTION NO: 6

FusionCompute can use templates to deploy virtual machines, or convert templates into virtual machines. Deploy virtual in batches When using a virtual machine, the template is generally converted into a virtual machine.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct. In FusionCompute, a template is a reusable golden-image resource designed to standardize and accelerate virtual-machine provisioning. When virtual machines must be deployed in batches, the platform normally creates multiple new virtual machines from the template while retaining the template for future use. This lets administrators consistently provision the same operating system, application baseline, and configuration across many instances. Template-to-virtual-machine conversion serves a different operational purpose: it changes the template itself into a virtual machine, rather than using it as a continuing source for repeated deployments. Consequently, conversion is not the general method for batch deployment, because consuming the template would remove the reusable base image needed to create the remaining instances. Huawei's image-management guidance likewise distinguishes reusable image-based creation from operations that change an image's state or role. See Huawei's [Image Management Service documentation](#) for image concepts and VM creation usage, and the Huawei [FusionCompute documentation search portal](#) for product-specific administration documentation.

QUESTION NO: 7

Single-choice question] Regarding the correspondence between the Fusion Access virtual group and the virtual machine type, which of the following is impossible of?

- A. VM Group (Express Encapsulation) - Virtual Machine Type (Express Encapsulation)
- B. Virtual Machine Group (Hosted Machine) - Virtual Machine Type (Hosted Machine)
- C. VM Group (Full Memory) - VM Type (Full Memory)

ANSWER: A

Explanation:

VM Group (Express Encapsulation) - Virtual Machine Type (Express Encapsulation) is the impossible correspondence. In FusionAccess, express (quick) encapsulation describes the desktop-group provisioning and delivery mechanism rather than a virtual-machine type that can be assigned to a VM. The express-encapsulation workflow rapidly provisions and personalizes desktops from the applicable source configuration, while the resulting virtual machines are managed as the supported underlying VM type, such as full-memory VMs, according to the desktop deployment design. Therefore, selecting an express-encapsulation virtual group does not create a corresponding VM type named "Express Encapsulation." The virtual-group category and VM-type field must be interpreted as separate, but compatible, configuration dimensions: the group identifies how desktops are organized and provisioned, whereas the VM type identifies the actual desktop VM implementation. Huawei's FusionAccess documentation and product materials describe FusionAccess as a desktop-cloud platform that manages desktop provisioning and lifecycle operations through these defined desktop and VM mechanisms. See the [Huawei FusionAccess product support page](#) and the [Huawei Desktop Cloud product overview](#) for FusionAccess documentation and product information.

QUESTION NO: 8

A school plans to use Huawei FusionAccess to build a desktop environment that can support 100 people to do experiments at the same time. Needle

What is the correct configuration for this scenario? (multiple choice)

- A. Use the linked clone desktop and set the shutdown to restore
- B. Use the same master image to make 100 linked clone desktops
- C. Configure the desktop group as a dynamic pool
- D. Close the port redirection of USB storage devices to prevent viruses from invading the system

ANSWER: A B C D

Explanation:

For a school laboratory in which many users perform temporary experiments, linked-clone desktops provide an efficient way to deliver a large number of standardized virtual desktops while minimizing storage consumption. A linked clone is created from a common master image, so using the same master image for the 100 desktops ensures that each student receives a consistent operating system, application set, and experiment baseline. A dynamic desktop pool is appropriate because it assigns desktops to users on demand and supports a nonpersistent, shared-use lab model rather than maintaining a permanently assigned desktop for every student.

Configuring restoration on shutdown returns each linked-clone desktop to its defined baseline after a student finishes. This preserves the environment for the next class, removes experiment changes, and reduces operational effort associated with manually rebuilding desktops. Disabling USB-storage-device redirection is also an appropriate security control in a shared lab: it limits introduction of malware through removable media and helps keep the standardized desktop environment trustworthy. These measures combine capacity efficiency, consistent delivery, rapid recovery, and endpoint security, which are central goals of a FusionAccess teaching-lab deployment. See Huawei's [FusionAccess product information](#) and the [Huawei Enterprise Support documentation portal](#).

QUESTION NO: 9

HDP is a new-generation cloud access desktop protocol developed by Huawei. Compared with traditional desktop protocols it has. Which of the following features are available?

- A. Support 64 virtual channels, each virtual channel can carry different upper-layer application protocols

B. Rich protocol management strategies

C. Different compression algorithms can be used according to different application types, flexible use of server rendering and local Fast rendering

D. Support lossless compression algorithm

ANSWER: A B C D

Explanation:

Huawei Desktop Protocol (HDP) is designed for virtual desktop delivery and uses a channel-based architecture to transport multiple types of desktop traffic concurrently. The feature “Support 64 virtual channels, each virtual channel can carry different upper-layer application protocols” reflects HDP’s ability to separate and carry services such as display, audio, USB redirection, printing, and other redirected peripheral or application traffic. “Rich protocol management strategies” is also a core HDP capability because desktop-cloud deployments need policies that adapt protocol behavior to network quality, endpoint type, and user scenarios.

HDP optimizes the desktop experience through application-aware encoding and rendering. Therefore, “Different compression algorithms can be used according to different application types, flexible use of server rendering and local Fast rendering” correctly describes its adaptive use of compression and rendering mechanisms. Lossless compression is likewise supported where visual fidelity is required, such as office documents, images, and detailed graphical content. These combined capabilities help deliver usable desktop sessions across varying bandwidth and latency conditions while balancing image quality, resource consumption, and responsiveness. Huawei’s Desktop Cloud product information is available at [Huawei Desktop Cloud](#), and Huawei’s official enterprise documentation portal is available at [Huawei Enterprise Support](#).

QUESTION NO: 10

(Multiple choice) Regarding the functions of each service in FusionSphere OpenStack, which of the following descriptions is incorrect?

A. hostVBSprocess withOSDProcess communication exception

B. GlanceProvides query, upload and download of virtual machine images

C. MDCProcess exception

D. hostVBSinternal error

ANSWER: A D

Explanation:

hostVBSprocess withOSDProcess communication exception and **hostVBSinternal error** are incorrect as descriptions of FusionSphere OpenStack service functions. These phrases describe hostVBS process fault or alarm conditions: one concerns abnormal communication with an OSD process, and the other identifies an internal hostVBS error. An alarm condition reports an abnormal operating state that may require diagnosis, recovery, or service maintenance; it does not define the responsibility or externally available capability of an OpenStack service. Service-function descriptions should instead identify a component’s operational role, such as managing compute resources, networking, identity, block storage, or image lifecycle operations. In particular, the OpenStack image service is responsible for image discovery, registration, storage, retrieval, and delivery, illustrating the difference between an actual service capability and an exception message. Huawei’s enterprise documentation portal provides the product documentation and alarm references used to operate FusionSphere-related platforms, while the upstream OpenStack documentation defines the Image service role and interfaces. See [OpenStack Glance documentation](#) and [Huawei Enterprise documentation](#).

QUESTION NO: 11

For CVMs with shared volumes mounted on FusionCloud, it is not supported to apply for CVM backup CSBS service instances and cloud server high-availability CSHA service instances.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. A shared volume is attached to more than one CVM and is therefore used in a clustered or shared-storage arrangement. In this arrangement, the volume's data and consistency state can be affected by multiple hosts, which prevents service-level protection from treating one CVM as an independent, self-contained backup or high-availability unit. FusionCloud consequently does not support applying for Cloud Server Backup Service (CSBS) instances or Cloud Server High Availability (CSHA) instances for CVMs that have shared volumes mounted. This restriction protects data consistency and prevents an unsupported recovery or failover workflow from being created for storage that is concurrently visible to multiple servers. Before planning backup or HA for a workload using shared volumes, administrators should use a protection design that is compatible with the application's cluster and shared-storage architecture, rather than applying CVM-level CSBS or CSHA protection to those servers. Huawei's cloud backup documentation describes backup support and constraints for server and disk resources, while its high-availability documentation presents the service scope and deployment prerequisites: [Cloud Backup and Recovery Service overview](#) and [Huawei cloud high-availability service overview](#).

QUESTION NO: 12

(Multiple choice questions) FusionaccessmiddlevdeskWhich of the following features are supported?

- A. Collect user VM logs
- B. Cooperate with HDP to control user virtual machines
- C. Optimize user virtual machine configuration
- D. Collect TC runtime logs.

ANSWER: A C

Explanation:

"Collect user VM logs" and "Optimize user virtual machine configuration" are supported functions of the FusionAccess vDesk middleware. vDesk operates within the virtual desktop environment and provides desktop-side maintenance support for administrators. Its log-collection capability gathers diagnostic information from a user's virtual machine, helping operations personnel investigate desktop faults, application issues, and abnormal operating conditions without relying solely on manually captured evidence. This capability is particularly useful when desktop images and user sessions are centrally managed at scale.

vDesk also provides virtual-machine configuration optimization capabilities. These capabilities help ensure that the desktop operating environment is configured appropriately for virtual desktop delivery, improving manageability and the user experience after a VM is provisioned from a template. Together, centralized log acquisition and configuration optimization reduce the operational effort required to maintain virtual desktops and support consistent service delivery. Huawei's enterprise support portal provides product documentation, patches, and knowledge resources for FusionAccess deployments, while Huawei's enterprise site provides the broader desktop-cloud product context: [Huawei Enterprise Support](#) and [Huawei Enterprise](#).

QUESTION NO: 13

When creating a port group in Fusioncompute, the following operations are incorrect?

- A. Set the VLAN ID to "5000".
- B. Add "this is the test port" to the description

- C. Set the port type to "Normal"
- D. Set the name of the port group to "test".

ANSWER: A

Explanation:

Set the VLAN ID to "5000" is the incorrect operation. A FusionCompute port group uses an IEEE 802.1Q VLAN identifier to associate virtual-machine traffic with a Layer 2 VLAN. The usable VLAN-ID range is limited by the 12-bit VLAN Identifier field defined by IEEE 802.1Q: valid user VLAN values are normally 1 through 4094. The remaining encodings are reserved, including VLAN 0 for priority-tagged frames and 4095 as a reserved value. Consequently, 5000 is outside the supported identifier range and cannot represent a valid 802.1Q VLAN for the port group. The configuration must instead use a VLAN ID that is available in the physical network and lies within the valid range, while ensuring that the connected switches and uplinks permit that VLAN. This preserves consistent traffic tagging between the FusionCompute virtual switch, the host's physical NICs, and the upstream network. The VLAN identifier format and reserved values are defined by the [IEEE 802.1Q standard](#); Huawei's enterprise support portal also provides product documentation for network and virtualization configuration at [Huawei Enterprise Support](#).

QUESTION NO: 14

Compared with traditional physical desktops, which of the following are the characteristics of cloud desktops?

- A. desktop standardization
- B. mobile office
- C. Terminal and data isolation
- D. automated scheduling

ANSWER: A B C D

Explanation:

Cloud desktops centralize desktop images, applications, and user data in cloud or data-center infrastructure rather than binding them to individual physical PCs. This architecture supports **desktop standardization** because administrators can create and maintain standardized desktop images and apply policies consistently to many users. It also enables **mobile office**: authorized users can access their assigned desktop remotely from supported endpoint devices, while the desktop environment remains centrally hosted. **Terminal and data isolation** is a core security characteristic, since business data and processing can remain in the hosted desktop environment instead of being stored persistently on endpoint terminals; this reduces exposure from lost or unmanaged devices. Finally, cloud desktop platforms use pooled compute, storage, and network resources, allowing **automated scheduling** and centralized lifecycle operations to allocate desktop resources, start or stop instances, and scale capacity according to service policies and demand. Huawei Cloud Workspace describes cloud desktops as centrally managed virtual desktops that provide secure access from multiple terminals and simplify desktop operations. Huawei Cloud's Workspace product information is available at [Huawei Cloud Workspace](#), with related service documentation at [Huawei Cloud Workspace Documentation](#).

QUESTION NO: 15

ManageOneThe homepage of the O&M plane displays information such as data center resources, alarms, capacity, and network status through different O&M map pages. The O&M map page is used to uniformly manage O&M functions and services and provide

unified O&M. Entrance

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Huawei ManageOne, the O&M plane provides centralized operational visibility for the data center. Its homepage uses O&M map views to present key operational information, including resource status, alarms, capacity-related information, and network conditions. These visual views allow operations personnel to understand the overall health and utilization of the managed environment from a unified console rather than having to enter separate management systems for each domain.

The O&M map is also intended to be an integrated entry point for operations and maintenance activities. From the map and its associated dashboards, administrators can access relevant O&M functions and services, investigate alarms, view resource or capacity conditions, and navigate to domain-specific management capabilities. This supports ManageOne's centralized O&M model across cloud infrastructure and data-center services. Therefore, the statement accurately describes both the information-distribution role of the O&M homepage and the unified-management role of O&M map pages. Huawei's ManageOne product documentation and support portal provide the documentation entry points for these O&M capabilities: [Huawei ManageOne User Manuals](#) and [Huawei Enterprise Support: ManageOne](#).

QUESTION NO: 16

Which of the following items will cause the failure to create a VM in FusionCloud R6? (many)

- A. The image_cache partition space of the computing node is insufficient
- B. Disk type configuration error
- C. VMTools is not installed in the mirror
- D. The label of the flavor does not match the label of the host group

ANSWER: A B D

Explanation:

The image_cache partition space of the computing node is insufficient is a VM-creation failure condition because the compute host must cache or prepare the selected image locally before the virtual machine can be provisioned. If the image cache filesystem lacks sufficient capacity, image preparation cannot complete and the creation task fails. Disk type configuration error is also a valid cause because the requested disk type must correspond to storage resources and capabilities configured in the cloud platform. An invalid, unavailable, or incompatible disk type prevents successful volume allocation and attachment during provisioning. The label of the flavor does not match the label of the host group is likewise a scheduling failure condition. Flavor and host-group labels are placement constraints: the scheduler must locate an eligible host group that satisfies the flavor's required label. When no matching host group is available, the platform cannot select a compliant compute host for the VM. These outcomes follow the general Huawei cloud provisioning model, where image availability, volume specifications, and host-placement constraints must all be satisfied before an instance can be created. See Huawei's [Image Management Service documentation](#) and [Elastic Cloud Server User Guide](#) for the related image, disk, and instance-provisioning requirements.

QUESTION NO: 17

Which of the following descriptions about the cloud disk backup service VBS in Fusioncloud 6.x is wrong?

- A. Support file or directory level recovery
- B. Within the same Region, support replication and recovery between different Azs
- C. Support consistent backup of multiple cloud disks
- D. Support concurrent multiple backups of the same cloud hard disk at the same time

ANSWER: A D

Explanation:

The incorrect descriptions are “Support file or directory level recovery” and “Support concurrent multiple backups of the same cloud hard disk at the same time.” VBS is a cloud-disk protection service that creates backups at the disk/block-device level. Its recovery workflow restores backup data to a cloud disk, whether by restoring the original disk or creating/restoring to a target disk as supported by the platform; it is not an in-guest file-system backup product that browses and restores individual files or directories. File- and directory-granular recovery requires a backup solution with guest OS or file-system awareness.

Backup processing for one source cloud disk is also serialized. A backup operation establishes a consistent point-in-time copy of the disk, and the service does not permit simultaneous backup jobs against that same disk. A further backup must wait until the current backup task has completed, failed, or otherwise left the active state. This control prevents overlapping read operations and protects the validity and manageability of the disk backup chain. Huawei’s VBS documentation describes VBS as a cloud disk backup and restoration capability and provides its operational constraints in the service documentation: [Huawei Cloud VBS Product Description](#) and [Huawei Cloud VBS User Guide](#).

QUESTION NO: 18

Which of the following types of backend storage is not supported for FusionSphere server virtualization?

- A. object storage
- B. IP SANstorage
- C. FC SANstorage
- D. NASstorage

ANSWER: A**Explanation:**

object storage is the correct answer. In FusionSphere server virtualization, VM disks require storage presented through the virtualization platform as persistent datastores that support the block- or file-oriented virtual-disk access model used for VM creation, running, snapshots, migration, and high availability. Object storage uses an API-based object model: data is stored and retrieved as whole objects identified by keys and accompanied by metadata, rather than being exposed as a VM datastore with the required virtual-disk semantics. It is therefore not a supported backend datastore type for FusionSphere server virtualization. Object storage can still have valuable cloud roles, such as storing unstructured data, archives, backups, images, or application objects, but that does not make it a direct VM disk backend in this context. Huawei describes FusionSphere as a cloud operating system and virtualization platform in its [FusionSphere product information](#). Huawei’s enterprise support portal provides the associated product documentation and compatibility guidance at [Huawei Enterprise Support](#).

QUESTION NO: 19

(Single-choice questions) FusionCloud, which of the following can be achieved VDC Isolation of internal resources?

- A. quota
- B. tenant
- C. Project
- D. available partition

ANSWER: B**Explanation:**

tenant is correct because a FusionCloud Virtual Data Center (VDC) is designed to provide logically isolated cloud environments for different tenants. A tenant represents an independent organizational entity or customer boundary. Within that boundary, resources such as compute, storage, networks, images, and service permissions can be allocated and

managed separately from those of other tenants. This isolation enables multiple departments or customers to use a shared cloud infrastructure while maintaining distinct resource ownership, administration, access-control boundaries, and usage accounting.

VDC resource isolation is therefore fundamentally based on the tenant model: administrators can assign resources and users to a tenant, apply tenant-level governance, and ensure that operations performed by one tenant do not affect the resources assigned to another. This aligns with the multi-tenant cloud-management model used by Huawei cloud platforms, where tenants provide the top-level logical scope for resource and identity management. Huawei's IAM documentation describes tenants as the organizational boundary for cloud resources and identities, while Huawei's cloud IAM overview explains the role of account-level isolation and authorization in multi-tenant environments. See [Huawei Cloud IAM User Guide](#) and [Huawei Cloud IAM overview](#).

QUESTION NO: 20

If a storage device is configured with 8 open 1 Gb/s iSCSI storage cards, 15% of the network bandwidth consumption, which of the following methods can calculate the effective bandwidth of the storage device (Mb/s) ?

- A. $8 \times 1024 \times 0.15$
- B. $8 \times 1024 \times 1024 \times 0.15$
- C. $8 \times 1024 \times (1 - 0.15)$
- D. $8 \times 1024 \times 1024 \times (1 - 0.15)$

ANSWER: C

Explanation:

The correct calculation is **$8 \times 1024 \times (1 - 0.15)$** . The storage device has eight iSCSI network cards, each rated at 1 Gb/s. Using the binary bandwidth conversion convention used by this question, each 1 Gb/s link is represented as 1024 Mb/s. Therefore, the combined nominal bandwidth is 8×1024 Mb/s. Network bandwidth consumption of 15% means that 15% of the available aggregate bandwidth is already used or unavailable for the effective storage service. The remaining usable proportion is consequently $1 - 0.15$, which is 0.85. Multiplying the total nominal bandwidth by this remaining proportion gives $8 \times 1024 \times 0.85 = 6963.2$ Mb/s. This calculation follows the standard capacity-planning principle of aggregating the bandwidth of parallel links and then applying the available-bandwidth ratio. iSCSI transports SCSI block-storage commands over IP networks, so network throughput and utilization directly determine the effective bandwidth available to storage traffic. See the [IETF iSCSI Protocol specification](#) and the [NIST reference on binary prefixes](#) for relevant terminology and unit context.

QUESTION NO: 21

If FusionCompute is deployed in the three-member separation mode, the deployment of FusionAccess must also adopt the three-member separation mode.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct because FusionCompute's deployment topology and FusionAccess's deployment topology are designed and sized separately. A three-member separation deployment in FusionCompute describes how the virtualization platform's management and service roles are arranged to provide platform availability and operational separation. It does not impose a mandatory matching topology on an upper-layer virtual desktop infrastructure solution.

FusionAccess is deployed according to its own component roles, availability requirements, user scale, and resource planning. Its architecture includes services such as access gateways, portal/web access components, desktop management components, and supporting infrastructure. These roles can be arranged in a topology appropriate to the desktop-cloud

design, provided that the selected FusionAccess version is compatible with the underlying FusionCompute environment and that network, storage, compute, and high-availability prerequisites are met.

Therefore, an administrator can use FusionCompute in three-member separation mode while selecting a different supported FusionAccess deployment architecture. The deciding factors are FusionAccess capacity and reliability requirements rather than a requirement to duplicate FusionCompute's role-separation model. Huawei's product information for [FusionAccess](#) and [FusionCompute](#) describes them as distinct products with independently planned solution roles and deployment considerations.

QUESTION NO: 22

When using Rainbow migration tool, if the migration is interrupted due to network failure, online file migration supports resuming the upload And block migration is not supported.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct. Rainbow is designed to make long-running data migrations resilient to transient connectivity problems. Its migration mechanisms maintain migration progress information so that, after a network interruption is resolved, the migration task can continue from the recorded checkpoint rather than requiring all previously transferred data to be sent again. This resumable-transfer capability applies to block migration as well as online file migration. For block migration, the tool tracks the status of transferred blocks and uses that status when the task is restarted or resumed, allowing it to continue with outstanding data. This is particularly important for large virtual-machine disks, where restarting a full disk copy after a temporary network failure would create excessive migration time and network consumption. Therefore, it is inaccurate to state that only online file migration supports upload resumption while block migration does not. Consult the official Huawei Cloud documentation portal for the applicable Rainbow version and its migration-task recovery guidance: [Huawei Cloud Support Documentation](#) and [Huawei Cloud Migration services](#).

QUESTION NO: 23

In HuaweiFusionCloud6.3In the solution, the following aboutLVSWhich of the component descriptions are correct?

- A. Work on the four layers of the network, strong anti-load ability
- B. no flow,LVSDistribute requests only
- C. yeshttp, database for load balancing
- D. LVSDependency ratio for the networkNginxLow

ANSWER: A C

Explanation:

LVS is appropriate for high-performance, transport-layer load balancing. It is implemented through Linux IP Virtual Server (IPVS), which operates on IP packets and TCP/UDP connections at the network/transport boundary commonly described as Layer 4. Because forwarding and scheduling are performed in the kernel rather than by a user-space application proxy, LVS can sustain very high connection and packet-processing rates. This makes "Work on the four layers of the network, strong anti-load ability" correct. The Linux kernel IPVS documentation describes IPVS as a transport-layer load-balancing facility and documents its schedulers and forwarding methods: [Linux kernel IPVS documentation](#).

"yeshttp, database for load balancing" is also correct when understood as Layer 4 balancing of the underlying TCP services. LVS can distribute TCP connections destined for HTTP service ports as well as TCP connections used by common databases. It does not need to interpret HTTP requests or database SQL in order to distribute those connections; it selects a

real server according to the configured virtual service and scheduling algorithm. The Linux Virtual Server project also documents IPVS virtual services and the supported TCP/UDP-based forwarding model: [Linux Virtual Server IPVS overview](#).

QUESTION NO: 24

The functions of FusionCare health check and information collection can be triggered when needed, or a period can be set Trigger, suitable for real-time information collection scenarios.

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. FusionCare's health-check and information-collection capabilities are designed to support both immediate operational needs and recurring maintenance activities. An engineer can initiate a collection task on demand when troubleshooting an active issue, validating the current state of a system, or preparing data for support analysis. This on-demand mode is appropriate where current, real-time operational information is needed rather than waiting for the next planned task. FusionCare also supports configuring periodic execution, allowing health information and diagnostic data to be gathered at scheduled intervals. This provides a regular operational baseline and helps maintain continuous visibility into the environment without requiring manual intervention each time. Using both trigger modes lets operations teams balance responsive incident handling with routine preventive maintenance, which is consistent with Huawei's service and technical-support approach to proactive operations and system lifecycle management. Huawei's enterprise support resources are available through the [Huawei Enterprise Support Community](#), and Huawei's technical-services overview describes its operational and maintenance service model at [Huawei Technical Services](#).

QUESTION NO: 25

Which of the following descriptions about disaster recovery backup are correct? (many)

- A. The backup system not only protects data, but also ensures business continuity
- B. Backup is mainly aimed at factors such as human misoperation, virus infection, logic errors, etc. Data recovery, data backup is generally carried out in the same data center
- C. The disaster recovery system only protects the recoverable version data at different time points
- D. Disaster recovery is mainly aimed at major natural disasters such as fires and earthquakes. Therefore, there must be a connection between the production site and the disaster recovery site. A certain safety distance must be maintained

ANSWER: B D

Explanation:

"Backup is mainly aimed at factors such as human misoperation, virus infection, logic errors, etc. Data recovery, data backup is generally carried out in the same data center" is correct because backup creates recoverable copies or versions of data. It is chiefly used to restore data after logical corruption, accidental deletion, application errors, malware infection, or similar data-level incidents. Backup repositories are commonly deployed within the same data center or region for efficient operational recovery, while further offsite copies may be added to meet stronger resilience requirements.

"Disaster recovery is mainly aimed at major natural disasters such as fires and earthquakes. Therefore, there must be a connection between the production site and the disaster recovery site. A certain safety distance must be maintained" is also correct. Disaster recovery addresses site-level failures that can render an entire production location unavailable. A separate recovery site at an appropriate geographic distance reduces exposure to the same physical disaster and enables workloads and services to be recovered there. Huawei's disaster-recovery

services use replication between production and DR sites to support this recovery model. See Huawei Cloud's [Server Disaster Recovery Service overview](#) and its [Cloud Backup and Recovery overview](#).

QUESTION NO: 26

Which of the following is true about backup and disaster recovery?

A. As long as data transmission, data storage and data exchange occur, data failure may occur, which will affect the use of services.

Therefore, data backup is required to ensure data security.

B. Data backup requires the production storage and backup storage to be placed in the same computer room

C. Disaster recovery can be achieved by placing data in different regions

D. Disaster recovery or backup can ensure that data will not be lost

ANSWER: A

Explanation:

“As long as data transmission, data storage and data exchange occur, data failure may occur, which will affect the use of services.

Therefore, data backup is required to ensure data security.” is correct because business data is exposed to a range of risks throughout its lifecycle, including accidental deletion, logical corruption, software faults, hardware failures, operational errors, and security incidents. A backup creates a recoverable copy or recovery point, allowing an organization to restore required data after an incident and thereby reduce the impact on service continuity and data availability. Backup planning is a fundamental data-protection control: it should define what is protected, how frequently copies are created, how long they are retained, where they are stored, and how restoration is verified. Huawei Cloud’s backup services are designed around this principle, providing policy-based backup and restoration capabilities for cloud resources and data. Huawei Cloud also describes disaster recovery as a broader continuity capability that uses recovery strategies and protected copies to restore workloads following disruptive events. Thus, maintaining backups is necessary to protect data and support dependable recovery when failures affect normal service use. See [Huawei Cloud Backup Service Overview](#) and [Huawei Cloud Disaster Recovery Service](#).

QUESTION NO: 27

The security management mechanism provided by FusionAccess includes access security, user information security and desktop management.

Manage component security. Deploy antivirus software on user virtual machines to prevent user virtual machines from being attacked by viruses,

This belongs to the security mechanism - security.

A. Desktop management

ANSWER: A

Explanation:

Desktop management is the applicable FusionAccess security mechanism because antivirus deployment on user virtual machines is an endpoint and virtual-desktop operating-management control. In a virtual desktop infrastructure, the desktop-management layer is responsible for maintaining the guest operating system environment and its installed software, including security agents that detect, prevent, and remediate malware. Centrally deploying and maintaining antivirus software on user virtual machines helps ensure that every provisioned desktop receives consistent protection, including desktops created from a standard image or recovered through desktop maintenance processes. This directly protects the virtual desktop workload where users run applications and handle data, rather than merely controlling login access or safeguarding user-account information. Huawei positions FusionAccess as a virtual-desktop solution with centralized desktop operations and management capabilities; endpoint security software is therefore part of securing the managed

desktop environment. See Huawei's [FusionAccess product information](#) and the [Huawei Enterprise Support documentation portal](#) for product and configuration documentation.

QUESTION NO: 28

In FusionCompute, when a virtual machine is migrated between hosts in the same cluster, which of the following conditions is required for hot migration when the VM uses virtualized storage?

- A. The source and target hosts can access the same datastore
- B. The virtual machine must be powered off
- C. The source and target hosts must use the same management IP address
- D. The virtual machine must use a local disk on the source host

ANSWER: A

Explanation:

The source and target hosts can access the same datastore is correct. During hot migration, the running VM is moved to the destination host while retaining access to its virtual disks. Therefore, the destination host must be able to access the datastore on which the VM disks reside. This enables the VM to continue reading and writing the same disk files after its execution state has been transferred.

A common datastore alone does not remove all migration requirements: the target host must also be available, have sufficient resources, and provide the required VM network connectivity. However, shared datastore access is the applicable storage prerequisite in this scenario. Huawei FusionCompute product documentation is available from the [Huawei FusionCompute support portal](#).

QUESTION NO: 29

Which of the following processing policies does FusionCompute support when a Windows virtual machine blue screen occurs? (multiple choice)

- A. do not handle
- B. Restart the virtual machine
- C. Perform HA for the virtual machine
- D. Shut down the virtual machine

ANSWER: A B C D

Explanation:

FusionCompute can detect a Windows guest operating system blue-screen event and apply a configured blue-screen handling policy. The available actions include *do not handle*, which leaves the virtual machine in its current state for investigation; *Restart the virtual machine*, which initiates a guest/virtual-machine restart action to restore service; and *Shut down the virtual machine*, which powers the affected workload off in a controlled policy-driven response. FusionCompute also supports triggering high availability processing through *Perform HA for the virtual machine*. This allows the platform's HA mechanism to recover the workload according to the configured availability policy and available host resources. These choices let administrators balance rapid service recovery, forensic preservation of the failure state, and HA-based automated recovery according to the importance of the application. Blue-screen handling is therefore configured as part of virtual-machine reliability and operations management rather than being limited to one mandatory recovery action. Huawei identifies FusionCompute as its virtualization platform for virtual-machine lifecycle, availability, and resource management capabilities; see the [Huawei FusionCompute product page](#) and the [Huawei FusionCompute technical support portal](#).

QUESTION NO: 30

In the Huawei Desktop Cloud Solution, which of the following is incorrect about the description of linked-clone desktop features?

- A. Linked clones offer the advantages of fast virtual machine creation and a small storage footprint
- B. Linked-clone desktops have higher computing power than ordinary virtual machines
- C. Linked clone is a technology that groups and maps a parent volume and a differential volume into a linked clone volume
- D. Linked-clone desktop pools provide desktop maintenance functions such as unified software update and unified system restore, greatly reducing maintenance costs

ANSWER: B

Explanation:

“Linked-clone desktops have higher computing power than ordinary virtual machines” is the incorrect statement. A linked clone is primarily a storage and lifecycle-management mechanism, not a mechanism for increasing a desktop’s CPU, memory, or graphics capacity. The clone is created from a common parent or base image and uses its own differential data area for changes made after deployment. This architecture makes it practical to provision many standardized desktops quickly while reducing the amount of duplicated operating-system and application data that must be stored.

Computing performance for a linked-clone desktop is determined by the virtual hardware and resources assigned to it, such as vCPU quantity, memory, GPU configuration where applicable, storage performance, and host resource availability. Sharing a base image does not inherently give a linked-clone desktop more compute power than a conventional virtual machine with comparable resource specifications. Huawei’s desktop-cloud offerings emphasize centralized desktop delivery, image-based management, and efficient resource use; these are distinct from increasing the processing capability allocated to an individual virtual desktop. See Huawei’s [Desktop Cloud product information](#) and [Huawei Cloud Workspace overview](#) for the platform’s desktop-cloud capabilities.

QUESTION NO: 31

OpenStackAs the most popular open source cloud operating system platform, which of the following cloud computing service models belong to?

- A. DaaS
- B. PaaS
- C. IaaS
- D. SaaS

ANSWER: C

Explanation:

IaaS is correct because OpenStack is principally an open-source infrastructure cloud platform. It provides the core building blocks used to create and operate virtualized data-center resources: compute instances through Nova, networking through Neutron, block storage through Cinder, object storage through Swift, image management through Glance, and identity services through Keystone. Cloud administrators deploy and manage these services to expose programmable pools of processing, storage, and networking resources to users and tenants.

This aligns with the Infrastructure as a Service model, in which consumers provision fundamental computing infrastructure on demand and retain control over operating systems, deployed applications, and many configuration choices within their allocated virtual resources. OpenStack’s dashboard, APIs, command-line clients, scheduling, tenancy, quota, and orchestration capabilities support this self-service infrastructure consumption model. Although organizations can build higher-level platform or application services on an OpenStack-based cloud, those services are layered on top of OpenStack rather than defining its primary service-model classification. The OpenStack project itself describes OpenStack as software

for building and managing cloud-computing platforms for public and private clouds, with services delivering compute, networking, and storage infrastructure. See the [OpenStack software overview](#) and the [OpenStack documentation](#).

QUESTION NO: 32

(Multiple choice) Which of the following IaaS service usage FusionCloud Used when building a forum website?

- A. RDS
- B. ECS
- C. EVS
- D. VPC

ANSWER: A B C D

Explanation:

A forum website normally requires compute, persistent storage, isolated networking, and a relational data store, so RDS, ECS, EVS, and VPC are all applicable FusionCloud IaaS services. ECS supplies the virtual-server resources on which the forum's web, application, and supporting components can run. EVS provides attachable block storage for persistent operating-system, application, uploaded-content, log, or other workload data. VPC provides the logically isolated virtual network in which the forum resources communicate and in which subnets, routing, security controls, and external-access design can be configured. RDS supplies a managed relational database capability appropriate for persistent structured forum information such as user accounts, posts, topics, comments, permissions, and transactional metadata. In a practical deployment, ECS instances are placed in a VPC, use EVS volumes as needed, and connect privately to the RDS database service. This combination maps directly to the fundamental infrastructure capabilities required to deploy a scalable, maintainable forum application. Huawei's service descriptions identify ECS as elastic compute, EVS as elastic volume/block storage, VPC as an isolated virtual network, and RDS as a managed relational database service; see [Huawei Cloud ECS overview](#) and [Huawei Cloud RDS overview](#).

QUESTION NO: 33

Which of the following operations is feasible when using rainbow for V2V migration?

- A. Use the NAT function of the router to open up the network between the source and destination virtual machines.
- B. Configure the routing function on the rainbow to connect the network between the source and destination virtual machines.
- C. Configure nat on rainbow; connect the network between the source and destination virtual machines.
- D. Use the routing function of the Layer 3 switch to open up the network between the source and destination virtual machines.

ANSWER: D

Explanation:

Use the routing function of the Layer 3 switch to open up the network between the source and destination virtual machines. This is feasible because Rainbow V2V migration requires reachable network paths between the source environment and the destination environment so that migration-related traffic can be transmitted reliably. A Layer 3 switch can provide IP routing between distinct subnets or VLANs, provided that the necessary interfaces, routes, and security policies are configured in both directions. This makes it possible to establish the required end-to-end connectivity without changing the virtual machines' workloads solely to accommodate the migration process. In practice, the migration design should confirm bidirectional reachability, correct gateway and route-table entries, sufficient bandwidth, and firewall or security-group rules that permit the required migration communication. Layer 3 routing is a standard network design method for connecting separated IP networks and is appropriate where the source and destination VM networks are not in the same Layer 2 broadcast domain. Huawei Cloud networking documentation describes the role of route tables in forwarding traffic between

network destinations and the configuration principles for virtual private cloud connectivity. See the [Huawei Cloud Support Documentation portal](#) and the [Huawei Cloud VPC product documentation overview](#).

QUESTION NO: 34

(Multiple choice) When FusionAccess Which of the following virtual machine types is not included when the desktop group type is "Static Pool"?

- A. linked clone
- B. Fast packaging
- C. full copy
- D. full memory

ANSWER: B

Explanation:

Fast packaging is the virtual-machine type not included for a FusionAccess static-pool desktop group. A static pool is designed to assign desktops persistently, so that a user is associated with the same desktop after initial allocation. Its supported desktop creation models include persistent VM forms such as full-copy, linked-clone, and full-memory desktops, depending on the FusionAccess version and deployment configuration. These models support the persistent assignment and lifecycle characteristics required by static desktops.

Fast packaging is instead associated with rapid, template-based desktop delivery rather than the static-pool VM-type set. It is therefore not selected as a VM type when creating a desktop group whose type is Static Pool. This distinction matters operationally because the desktop-group type determines allocation persistence and available provisioning models; administrators must choose a supported persistent VM type when configuring the pool. Huawei's FusionAccess product information describes its virtual desktop infrastructure and centralized desktop provisioning capabilities, while the Huawei Enterprise Support portal provides the version-specific FusionAccess product documentation and configuration guides used to verify supported desktop-group settings. See [Huawei FusionAccess](#) and [Huawei Enterprise Support](#).

QUESTION NO: 35

In FusionAccess, when creating a virtual machine template, if required ----- group members can also log in

If you want to record a virtual machine, you need to check "Configure user login" when creating a virtual machine template

- A. Users

ANSWER: A

Explanation:

Users correctly completes the blank. In FusionAccess desktop-template creation, the "Configure user login" setting is used to prepare login-related configuration in the template. This allows the template's intended user configuration to be retained when desktops are provisioned or recorded from it. Where access is needed for a group rather than only a single individual, members of the configured user group can authenticate and log in to the resulting virtual desktop, subject to the usual desktop-assignment and access-control configuration. This is consistent with FusionAccess's role as a virtual desktop infrastructure platform: desktop resources are delivered to authorized users, and user or user-group membership is the basis on which a shared desktop-access requirement is expressed. Selecting "Configure user login" during template creation is therefore important when the virtual machine must support the relevant user-login configuration before it is used for recording or subsequent desktop delivery. Huawei's FusionAccess product information is available at [Huawei FusionAccess](#), and Huawei documentation and product support resources can be accessed through the [Huawei Enterprise Support](#) portal.

QUESTION NO: 36

Which data on the FusionCompute host is usually forwarded using the management port? (single choice)

- A. Business virtual machine live migration command
- B. Business virtual machine data disk data
- C. Business virtual machine system disk data
- D. Business virtual machine memory access data

ANSWER: A

Explanation:

Business virtual machine live migration command is the correct answer because the management port is used for the control-plane communication required to administer FusionCompute hosts and virtual machines. A live migration operation must first be initiated, coordinated, authorized, and monitored through management communication between the cloud management components and the involved compute hosts. The migration command is therefore management traffic: it carries operational instructions and status signaling rather than the workload's bulk data path. Keeping this control traffic on the management network enables administrators to manage host resources and execute lifecycle operations reliably, including migration orchestration. In a standard FusionCompute networking design, management-plane connectivity is separated from service and storage-related connectivity so that control operations remain available and are not affected by normal workload traffic patterns. Huawei's FusionCompute product documentation and technical resources describe the platform's host and VM management functions, including network-planning considerations for management communication. See the [Huawei FusionCompute support portal](#) and the [Huawei Enterprise Support documentation portal](#) for product documentation and deployment guidance.

QUESTION NO: 37

Which of the following cloud services are supported by the three TYPE scenarios of FusionCloud? (many)

- A. OBS
- B. EIP
- C. VAPP
- D. AS

ANSWER: A C D

Explanation:

FusionCloud's typical cloud-service scenarios include object storage, virtual application delivery, and elastic scaling capabilities. **OBS** provides object-based storage for unstructured data and supports scalable, durable cloud data storage. **VAPP** represents virtual application services, enabling application templates and their associated virtual resources to be delivered and managed as a logical application unit. **AS** provides auto-scaling functionality, allowing the platform to adjust application compute capacity according to defined policies and workload demand. Together, these services reflect storage, application-oriented resource delivery, and elastic resource management capabilities within FusionCloud scenarios. Huawei's cloud service documentation describes OBS as an object storage service and Auto Scaling as a service for automatically adjusting resources in response to workload changes. See the [Huawei Cloud OBS service overview](#) and the [Huawei Cloud Auto Scaling service overview](#).

QUESTION NO: 38

FusionAccessProvisioning of linked-clone virtual machines fails. Which of the following reasons cannot be the cause?

- A. Fusion Access the number of concurrent users has reached the maximum limit
- B. The virtual machine template was deleted by mistake

- C. FusionComputePlatform failure
- D. ADServer exception

ANSWER: A

Explanation:

Fusion Acces the number of concurrent users has reached the maximum limit cannot itself cause linked-clone virtual-machine provisioning to fail. A concurrent-user limit is an access-control or licensing threshold that applies when users attempt to establish desktop sessions. It governs whether additional users can log in after the permitted concurrent-session capacity is exhausted; it is not a prerequisite in the linked-clone creation workflow. Linked-clone provisioning is performed by the FusionAccess and FusionCompute components using the configured desktop template, storage and virtualization resources, and directory-related configuration needed to create and register desktop instances. Consequently, the relevant provisioning process can continue independently of the current number of users who are connected to virtual desktops. An administrator may need to address concurrent-session capacity before allowing another user to access a desktop, but reaching that threshold does not prevent the platform from creating a linked-clone VM. Huawei identifies FusionAccess as its virtual desktop infrastructure platform and provides its product documentation and support materials through the [Huawei FusionAccess support portal](#). Additional platform context is available on the [Huawei FusionAccess product page](#).

QUESTION NO: 39

In FusionSphere OpenStack, which network plane addresses can be used to log in to the CPS interface?

- A. Internal Base
- B. External API
- C. External OM
- D. External Base

ANSWER: A B

Explanation:

The CPS interface can be accessed through the **Internal Base** network plane and the **External API** network plane. Internal Base is the internal service-network path used by cloud-platform components and authorized operations personnel within the management environment. Therefore, when an administrator has network reachability to that plane, its address provides a valid path to the CPS web interface. External API is the northbound-facing plane through which designated external clients and administrative tools access cloud-platform service interfaces. CPS publishes its access endpoint on this plane as well, allowing authorized users outside the internal service network to reach the interface according to the deployed security policy and routing design.

Using these two planes provides both an internal management access path and a controlled external-facing service/API access path. In practice, administrators must also ensure that the relevant security groups, firewalls, routes, certificates, and browser access permissions permit the connection. Huawei's FusionSphere OpenStack product resources are available through the [Huawei FusionSphere OpenStack support page](#); Huawei's cloud-computing documentation portal is available at [Huawei Enterprise Cloud Computing Support](#).

QUESTION NO: 40

In Fusioncloud, which of the following descriptions about virtual private network VPN are correct? (many)

- A. Currently VPN only supports IPsec VPN
- B. One VPN gateway in the VPN can correspond to multiple VPN connections
- C. VPN is a secure and encrypted communication tunnel used to connect the local data center and the network on the cloud
- D. VPN can also be used to connect two VPCs under the same region

ANSWER: A B C

Explanation:

“Currently VPN only supports IPsec VPN” is correct in the FusionCloud/Huawei Cloud VPN service context: the service establishes standards-based IPsec tunnels, using IKE for security association negotiation and IPsec protocols to protect traffic. “One VPN gateway in the VPN can correspond to multiple VPN connections” is also correct. A VPN gateway is the cloud-side access point, and multiple VPN connections can be created on that gateway to provide connectivity to different customer gateways or networks, subject to the applicable service quotas and design requirements. “VPN is a secure and encrypted communication tunnel used to connect the local data center and the network on the cloud” accurately describes the primary hybrid-cloud use case. The IPsec tunnel encrypts traffic transmitted over public networks between an on-premises customer gateway and a cloud VPN gateway, allowing private network resources to communicate securely without exposing application traffic in plaintext. Huawei’s VPN documentation describes VPN as an encrypted communication tunnel between an on-premises data center and a VPC, and documents VPN gateways and VPN connections as separate service objects. See [Huawei Cloud VPN Service Overview](#) and [Huawei Cloud VPN User Guide](#).

QUESTION NO: 41

The sequence design of business migration is very important. The following description of the sequence of business migration is correct

What are the items? (multiple choice)

- A. According to the impact on the business from small to large
- B. According to the business interruption time from short to long
- C. According to business complexity, from easy to difficult
- D. According to the migration risk from low to high

ANSWER: A B C D

Explanation:

A phased migration sequence should deliberately begin with workloads that have the smallest business impact, the shortest acceptable interruption window, the simplest dependencies, and the lowest migration risk. Therefore, “According to the impact on the business from small to large,” “According to the business interruption time from short to long,” “According to business complexity, from easy to difficult,” and “According to the migration risk from low to high” are all appropriate sequencing principles. Starting with these lower-exposure workloads lets the migration team validate connectivity, replication, cutover procedures, rollback processes, monitoring, and operational ownership under controlled conditions. Each completed wave produces practical experience and reusable runbooks, reducing uncertainty before moving to increasingly critical, complex, and risk-sensitive services. This wave-based approach also limits the blast radius of early issues and supports continuous improvement of the migration plan based on real results. Huawei Cloud migration guidance emphasizes assessment, planning, phased execution, and risk control as key parts of a successful migration program. See [Huawei Cloud Migration](#) and [Huawei Cloud Documentation](#).

QUESTION NO: 42

In the FusionSphere private cloud solution, users can apply for cloud hard disk backup to realize local backup of cloud hard disks backup and recovery.

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. FusionSphere private-cloud environments provide a cloud hard disk backup capability through their backup services, enabling tenants to protect cloud hard disks by creating backup copies and restoring data when required. A user first obtains or applies for the relevant backup service and then can create backups according to operational needs, such as before a major application change, maintenance activity, or data migration. The resulting backup data provides a recovery point from which the protected cloud hard disk can be restored, supporting recovery from accidental deletion, logical corruption, or other data-loss events. In a private-cloud deployment, this capability is delivered within the customer's cloud environment and is commonly described as local backup, distinguishing it from a separately deployed off-site disaster-recovery design. Backup and restoration are therefore core lifecycle operations of cloud hard disk protection, rather than merely manual copying by the tenant. Huawei's Cloud Backup Service documentation describes the service as providing backup protection and restoration for cloud resources, including disk-related data, while Huawei's cloud-storage documentation describes elastic cloud disks as persistent block storage designed for use by cloud servers. See [Huawei Cloud Backup and Recovery documentation](#) and [Huawei Elastic Volume Service documentation](#).

QUESTION NO: 43

The security management of accounts in FusionAccess is also an important part of the operation and maintenance work. The following

What is wrong about the description of changing the password of the Linux infrastructure server "gandalf" account?
item?

- A. The "gandalf" account is the local ordinary user account of the Linux server
- B. The "gandalf" account is used to start and stop VAG services on the VAG server
- C. If the "gandalf" account of the linux infrastructure server where the vAG and HDC components are located is modified at the same time password, you need to reconfigure the HDC component
- D. "gandalf" account can be used for SSH login to Linux server

ANSWER: C

Explanation:

The statement "If the "gandalf" account of the linux infrastructure server where the vAG and HDC components are located is modified at the same time password, you need to reconfigure the HDC component" is the incorrect description. The `gandalf` account is a FusionAccess infrastructure service account, and its password change must be performed through the prescribed maintenance procedure so that dependent FusionAccess service configurations can be updated consistently. However, changing that operating-system account password does not, by itself, mean that the HDC component must be reconfigured merely because VAG and HDC are deployed on the same Linux infrastructure server. Reconfiguring a component is an operation associated with component deployment, connection parameters, or service configuration changes; it is not the normal mandatory consequence stated here for a coordinated `gandalf` password update. Administrators should follow the version-specific FusionAccess operations guide, use the supported account-management workflow, and verify service status after the password-maintenance task. Huawei's desktop-cloud product information is available at [Huawei FusionAccess](#), and official product documentation and maintenance materials can be accessed through the [Huawei Enterprise Support portal](#).

QUESTION NO: 44

(multiple choice)Manage One the operation and maintenance plane can provide three-dimensional panoramic monitoring function. Which of the following monitoring objects are supported?

- A. resource pool
- B. Tenant application
- C. physical device
- D. Cloud service instance

ANSWER: A B C D

Explanation:

ManageOne's operations and maintenance capabilities are designed for end-to-end, layered monitoring across the cloud data center rather than for a single infrastructure layer. Therefore, the supported three-dimensional panoramic monitoring scope includes the **resource pool**, which provides visibility into aggregate cloud resource status and capacity; the **Tenant application**, which connects monitoring to tenant-facing application workloads; the **physical device**, which covers the underlying hardware infrastructure; and the **Cloud service instance**, which provides operational visibility into instantiated cloud services. Together, these objects provide the resource, infrastructure, service, and application perspectives needed for panoramic O&M. This consolidated monitoring model enables operators to correlate alarms and health information across layers, assess service impact, and locate faults more efficiently from tenant applications down to physical devices. Huawei describes ManageOne as cloud management software that provides unified management, operations, and maintenance for data-center cloud resources and services. Its product materials also emphasize centralized monitoring and operational management across heterogeneous cloud infrastructure. See Huawei's [ManageOne product page](#) and the [ManageOne documentation portal](#).

QUESTION NO: 45

The Neutron project provides network services in OpenStack. In native scenarios, virtual

The internal communication of a single node and the communication of virtual machines across nodes need to be supported by the virtual machine network inside the physical host.

Among them, the virtual network bridge (or virtual machine switch) is a kind of important homogeneous element in the virtual network, and carries the control

A series of flow tables issued by the level, such as

The virtual bridge is implemented by --- open source software. (Please fill in the English abbreviation)

A. OVS

ANSWER: A

Explanation:

OVS is the correct abbreviation for Open vSwitch, the open-source multilayer virtual switch commonly deployed on OpenStack compute hosts. It provides the virtual bridge functionality needed to connect virtual-machine interfaces, local host networking, and physical-network uplinks. In Neutron deployments, Open vSwitch can be controlled by Neutron agents and is able to apply forwarding behavior through OpenFlow flow tables. These flows implement network segmentation and connectivity policies for tenant networks, including local virtual-machine communication and traffic that must traverse between compute nodes through overlay or provider networks. Open vSwitch is specifically designed to operate in virtualized environments and supports integration with hypervisors such as KVM, making it a standard component of traditional OpenStack networking architectures. The OpenStack Neutron Open vSwitch agent documentation describes the agent as configuring Open vSwitch bridges and flow rules for Neutron networking. The official Open vSwitch documentation further identifies Open vSwitch as a production-quality, multilayer virtual switch intended for virtual-machine environments. References: [OpenStack Neutron documentation](#) and [Open vSwitch documentation](#).

QUESTION NO: 46

In the FusionCloud private cloud scenario, the following descriptions about NAT gateway related concepts are correct. What are the items? (many)

A. There are four types of NAT gateways: small, medium, large and extra large

B. To create a NAT gateway, you must specify the VPC, subnet, and NAT gateway specifications where the NAT gateway is located

C. The specifications of the NAT gateway will affect the maximum number of connections and the number of new connections per second of the SNAT function, and the data throughput
The amount is determined by the bandwidth of the Elastic IP

D. The set composed of five elements: destination IP address, destination port, and transport layer protocol is regarded as a DNAT connect

ANSWER: A B C

Explanation:

“There are four types of NAT gateways: small, medium, large and extra large” is correct because FusionCloud NAT Gateway capacity is provided through these gateway specifications. The selected specification establishes the service capacity available to NAT processing, rather than merely acting as a descriptive label. “To create a NAT gateway, you must specify the VPC, subnet, and NAT gateway specifications where the NAT gateway is located” is also correct: a NAT gateway is deployed into a particular VPC and subnet, and its specification must be selected during provisioning so that it can provide services to resources in that virtual network.

“The specifications of the NAT gateway will affect the maximum number of connections and the number of new connections per second of the SNAT function, and the data throughput
The amount is determined by the bandwidth of the Elastic IP” correctly distinguishes gateway processing capacity from public-network bandwidth. The NAT gateway flavor controls SNAT connection scale and new-connection rate, while the associated elastic public IP bandwidth constrains external traffic throughput. This separation is important when sizing a gateway: sufficient EIP bandwidth does not replace the connection capacity supplied by the selected NAT gateway specification. Huawei’s NAT Gateway documentation describes the service model and its SNAT/DNAT capabilities, while its EIP documentation explains bandwidth-based public-network access: [Huawei Cloud NAT Gateway Product Description](#) and [Huawei Cloud Elastic IP Product Description](#).

QUESTION NO: 47

Backing up different components in the FusionAccess system requires different backup settings. by

Which of the following descriptions about the backup settings corresponding to each component is correct? (many)

- A.** When backing up DNS, you need to configure DNS information in the "Alarm Component" of FusionAccess
- B.** When backing up VLB components, you need to configure VLB information in FusionAccess "Monitoring Components"
- C.** When backing up DHCP, you need to configure DHCP information in the "Alarm Component" of FusionAccess
- D.** When backing up AD, in FusionAccess "System Management > Initial Configuration > Domain/OU", the "Enable Backup" is set to "Yes"

ANSWER: A C D

Explanation:

FusionAccess backup preparation depends on the component type because the system must have the required connection and management information before it can collect and protect that component’s data. For DNS backup, the DNS server information is configured under the FusionAccess alarm-component configuration. This provides FusionAccess with the component details required for backup-related management and alarm handling. DHCP uses the same alarm-component configuration approach: DHCP information must be configured there before DHCP data can be backed up.

Active Directory backup is enabled differently because it is associated with the configured domain environment. In *System Management > Initial Configuration > Domain/OU*, setting *Enable Backup* to *Yes* enables backup for the AD domain configuration. These settings ensure that the FusionAccess backup process includes the infrastructure services and directory information on which desktop authentication and service operation depend. Huawei’s enterprise support portal provides the applicable FusionAccess product documentation and maintenance guides: [Huawei Enterprise Support](#). Product documentation can also be accessed through the [Huawei Enterprise Documentation portal](#).

QUESTION NO: 48

(Multiple choice) The company needs to FusionCompute. The virtual machine resources in the cluster are subject to permission control. Which of the following is incorrect about

the derogation of virtual machine resources?

- A. You can set resource usage limits for virtual machines, including CPU Memory, Network and Disk IOPS Wait
- B. network QoS Policies provide bandwidth configuration control capabilities and do not support traffic restrictions between virtual machines on the same host
- C. If the host memory is sufficient, the memory request of the virtual machine can be unlimited, if the host memory is insufficient, it will be allocated according to the memory level and share
- D. can CPU a single physical core is assigned to a single virtual machine

ANSWER: A C

Explanation:

The statements "You can set resource usage limits for virtual machines, including CPU Memory, Network and Disk IOPS Wait" and "If the host memory is sufficient, the memory request of the virtual machine can be unlimited, if the host memory is insufficient, it will be allocated according to the memory level and share" are incorrect because they overstate how FusionCompute controls virtual-machine resources. Resource control is implemented through distinct mechanisms for compute, memory, network, and storage rather than as one unrestricted, universally applicable limit covering the listed items. In particular, storage quality-of-service controls use defined disk-performance parameters and policies; they should not be expressed as a generic "Disk IOPS Wait" resource-limit setting. Memory management is likewise bounded by the virtual machine's configured memory and the host's available physical resources. Memory priority and shares participate in contention handling, but they do not make a virtual machine's memory request unlimited when capacity appears available. Administrators must plan reservations, limits, shares, and overcommitment according to supported FusionCompute settings and host capacity. Huawei's enterprise support portal provides the official FusionCompute product documentation and configuration guidance at [Huawei FusionCompute Support](#). Additional official product and documentation searches are available through [Huawei Enterprise Support Search](#).

QUESTION NO: 49

When FusionCompute installs a virtual machine operating system, which of the following methods are supported for mounting the optical drive? (many)

- A. File method
- B. to share
- C. local way
- D. remote

ANSWER: A B C

Explanation:

FusionCompute supports three relevant virtual optical-drive mounting modes during virtual-machine operating-system installation: **File method**, **to share**, and **local way**. The file method uses an ISO image available through the platform's supported image or storage access mechanism, allowing the VM to boot from installation media without a physical disc. The shared method enables installation media made available through a shared resource, which is useful when centrally managed installation images are provided to virtual machines. The local method maps the administrator's local CD/DVD device or locally selected installation media through the management client to the VM's virtual CD/DVD drive. These mounting approaches let an administrator present bootable operating-system media to a VM and then select that virtual optical drive as the startup source for installation. Huawei's FusionCompute documentation describes VM lifecycle and device-management operations, including the use of virtual CD/DVD devices and ISO media, in its product documentation resources. See the [Huawei FusionCompute product documentation portal](#) and the [FusionCompute documentation library](#).

QUESTION NO: 50

HuaweiManageOne OperationCenterpassHTTPSandSNMP TrapThere are two ways to synchronize alarm data from the lower management system. The former ensures that alarms can be detected in time, and the latter ensures that alarms will not be lost.

- A. True
- B. False

ANSWER: B

Explanation:

False is correct. In ManageOne OperationCenter alarm synchronization, SNMP Trap is the event-driven, near-real-time reporting mechanism: a lower-level management system proactively sends an alarm notification when the event occurs. This supports prompt alarm awareness, but SNMP notifications are normally transported using UDP and are not inherently acknowledged or retransmitted; therefore, a Trap transmission alone cannot guarantee that every alarm will be received. HTTPS synchronization is used as a reliable complementary mechanism for obtaining alarm information from the lower-level system. By using request/response communication and synchronization or reconciliation of alarm data, it helps recover information that may not have arrived through real-time Trap reporting and thus prevents alarm records from being missed. Consequently, the statement reverses the purposes of the two mechanisms: real-time detection is associated with SNMP Trap, while preventing alarm loss is associated with HTTPS-based synchronization. This behavior is also consistent with the standard SNMP notification model: RFC 3413 distinguishes unacknowledged Traps from acknowledged Inform notifications, whereas HTTPS uses reliable HTTP exchanges over TCP. See [RFC 3413, SNMP Applications](#) and [RFC 9110, HTTP Semantics](#).

QUESTION NO: 51

The following aboutFusionAccessIn the description of the virtual machine disk configuration mode during fast provisioning, which options are correct?

- A. When the virtual machine template is a full copy template, the configuration mode must be normal delay zero
- B. When the virtual machine template is a full copy template, the configuration mode is recommended to be normal delay zero.
- C. When the virtual machine template is a linked clone template, the configuration mode is recommended to be thin
- D. When the virtual machine template is a linked-clone template, the configuration mode must be condensed Correct

ANSWER: B D

Explanation:

During FusionAccess fast provisioning, the disk configuration must match the desktop-creation method because full-copy and linked-clone desktops consume and manage storage differently. "When the virtual machine template is a full copy template, the configuration mode is recommended to be normal delay zero." is correct: full-copy provisioning creates an independent virtual disk for each desktop, and the normal delayed-zero configuration is the recommended balance between creation speed and normal disk initialization behavior. "When the virtual machine template is a linked-clone template, the configuration mode must be condensed Correct" is also correct in substance: linked clones depend on the parent/template disk and snapshots, so the condensed disk configuration is required to support the linked-clone provisioning model and its storage-efficient behavior. These settings are therefore not merely interchangeable storage-allocation choices; they are tied to the selected template type in FusionAccess fast provisioning. Huawei positions FusionAccess as a virtual desktop infrastructure solution integrated with its cloud and virtualization platforms, where template and storage configuration are fundamental to scalable desktop delivery. See the [Huawei FusionAccess product page](#) and the [Huawei Enterprise Cloud Computing support portal](#) for product documentation and support resources.

QUESTION NO: 52

Cloud resource monitoring mainly real-time data center cloud resource usage from five aspects of computing, storage, network, application and security, collect invigilation

indicators of each cloud resource module, and detect the availability of resource modules, so that administrators can fully understand cloud resources. The current situation, and

then analyze the operating status and health of the business, and respond to abnormal alarms in a timely manner to ensure the smooth operation of the application.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. Cloud resource monitoring is intended to provide continuous visibility into the operational condition and consumption of cloud infrastructure and cloud-hosted services. In a data center cloud environment, this includes collecting metrics across compute resources, storage systems, network components, applications, and security-related services. Typical monitored indicators include CPU and memory utilization, virtual-machine status, storage capacity and performance, network traffic and connectivity, application response or service health, and alarm events.

By aggregating these indicators, monitoring enables administrators to assess whether resource modules are available and whether workloads are operating within expected thresholds. It also supports health analysis by correlating infrastructure conditions with business-service behavior. When monitored values breach configured alarm thresholds or a resource becomes unavailable, alarm notifications enable operations personnel to investigate and respond promptly. This proactive monitoring-and-alarm cycle helps reduce service interruption risk and supports stable application operation.

Huawei Cloud documentation describes Cloud Eye as a service for monitoring cloud resources, reporting alarms, and providing operational visibility. Its monitoring model covers resource metrics and alarm management, which directly supports the purpose stated in the question. See [Huawei Cloud Eye product description](#) and [Huawei Cloud Eye alarm documentation](#).

QUESTION NO: 53

Ebackup only supports data backup in virtualization scenarios, eReplication supports both physical machines and virtual machine disaster recovery

A. Correct

B. False

ANSWER: B

Explanation:

“False” is correct because the statement incorrectly limits Huawei eBackup to virtualized environments. eBackup is an enterprise data-protection product designed to protect more than virtual-machine workloads. Its protection scope includes physical-server application data and commonly used enterprise workloads, as well as virtualized infrastructure. Depending on the deployed version, licenses, and integration components, eBackup can provide centralized backup and recovery for physical servers, databases, file systems, and virtual machines. Therefore, support for physical workloads is an intended part of its data-protection capability rather than an exclusion.

Huawei eReplication is used for disaster-recovery and replication scenarios and can protect both physical and virtualized environments through appropriate replication and recovery configurations. However, because the question joins the eBackup claim and the eReplication claim into one statement, the inaccurate description of eBackup means the overall statement must be evaluated as false. Huawei’s enterprise data-protection portfolio information is available at [Huawei eBackup](#) and [Huawei eReplication](#).

QUESTION NO: 54 - (SIMULATION)

If a linked clone virtual machine template is 30G, when using this template to create a linked clone virtual machine, The size of the parent volume is 30G, the differential volume cannot be less than 3G, and the maximum cannot exceed - G.

ANSWER: See the explanation for the answer

Explanation:

Answer: 30G.

For a linked-clone VM, the differential volume size must be between **10%** and **100%** of the parent volume size.

Parent volume: 30G

Minimum differential volume: $30\text{G} \times 10\% = 3\text{G}$

Maximum differential volume: $30\text{G} \times 100\% = 30\text{G}$

Therefore, the differential volume cannot exceed **30G**.

QUESTION NO: 55

In useRainbowMigrating a special virtual machine toHCSplatform, if the operating system of the source virtual machine isWindows 2016 Server, it is necessary toHCSA corresponding version of the image is indicated on the image for creating the destination virtual machine during migration

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct. Rainbow migration creates the destination virtual machine using the migration workflow and the source virtual machine's collected configuration and disks; it does not require an operator to select a matching Windows Server 2016 image from the HCS image catalog solely because the source guest runs that operating system. The key migration prerequisites are compatibility of the source workload with the migration scenario, correct target resources and networking, and the guest OS preparation required by the migration tool. The target system disk is produced from the migrated source disk data, rather than being replaced by a newly selected base operating-system image. Consequently, Windows Server 2016 on the source does not itself introduce a mandatory corresponding-image selection step while creating the destination VM. Huawei's cloud documentation describes images as templates used when independently provisioning new ECS instances, whereas migration workflows transfer an existing workload to the target environment. See Huawei's [Image Management Service documentation](#) for the role of images in instance creation, and the [Migration Center documentation](#) for Huawei migration-service workflow information.

QUESTION NO: 56

There are two indicators for disaster recovery, RTO and RPO, which are correct (multiple choice)

A. A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, and the RTO is 30 minutes

B. A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, RPO for 30 minute

C. A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, and the RPO is 5 minutes

D. A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, and the RTO is 25 minutes

ANSWER: C D

Explanation:

“A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, and the RPO is 5 minutes” is correct because the recovery point objective measures the maximum acceptable amount of data loss, expressed as a time interval. The recovered data represents the state at 7:00, while the failure occurs at 7:05. Therefore, the recovery point is five minutes before the failure, producing an RPO of five minutes. This indicates that changes made between 7:00 and 7:05 are not present after recovery.

“A service fails at 7:05, the service is restored at 7:30, the restored service data is the data at 7:00, and the RTO is 25 minutes” is also correct. The recovery time objective measures the allowable or actual time for restoring service availability following an interruption. The service is unavailable beginning at 7:05 and becomes available again at 7:30, an elapsed interval of 25 minutes. RTO concerns restoration of service operation; it is independent of the timestamp represented by the recovered data. These definitions align with the disaster-recovery terminology in the [NIST RTO glossary entry](#) and the [NIST RPO glossary entry](#).

QUESTION NO: 57

Which of the following is the backup file storage path for local backup?

- A. /backup
- B. /var/backup
- C. /var/vdesktop/backup
- D. /var/huawei/vdesktop/backup

ANSWER: C**Explanation:**

The correct local-backup storage path is `/var/vdesktop/backup`. In a FusionAccess deployment, this directory is the predefined location used by the platform's local backup mechanism to retain generated backup files on the relevant management node. Administrators performing backup inspection, capacity planning, retention management, or controlled backup-file export should identify this directory as the local destination. Because backup content can grow over time, the filesystem hosting `/var/vdesktop/backup` should be monitored for available space and protected according to the organization's operational and security procedures. Backup operations should be performed through the supported FusionAccess management and maintenance workflows so that backup metadata, file consistency, and restoration requirements are maintained. Before moving, deleting, or copying files from this location, administrators should confirm the backup status and follow the applicable product-version documentation, since active backup jobs and retention policies may affect file handling. Huawei publishes FusionAccess product documentation and lifecycle resources through its Enterprise Support portal, available at [Huawei FusionAccess Enterprise Support](#). General product information and supported solution context are also available at [Huawei FusionAccess](#).

QUESTION NO: 58

UltraVR supports centralized deployment and distributed deployment. The following describes the two deployment methods

What are the wrong options? (many)

- A. In the distributed deployment scenario, UltraVR Server is deployed on a separate server at the disaster recovery site. Disaster recovery management of current production resources and disaster recovery resources.
- B. In the centralized deployment scenario, UltraVR Server is deployed on a separate server at the disaster recovery site to achieve disaster recovery management of production resources and disaster recovery resources.
- C. In the distributed deployment scenario, UltraVR Server is deployed on the production site and the service of the disaster recovery site respectively
On the server, the disaster recovery management of the site resources is realized.

D. In the centralized deployment scenario, the UltraVR Server deployed at the production site and the disaster recovery site are required to be deployed network connectivity between UltraVR Servers.

ANSWER: A D

Explanation:

The incorrect descriptions are the ones that interchange the UltraVR Server placement and connectivity requirements of centralized and distributed deployment. "In the distributed deployment scenario, UltraVR Server is deployed on a separate server at the disaster recovery site.

Disaster recovery management of current production resources and disaster recovery resources." describes a single disaster-recovery-site server managing resources at both sites. This is the centralized management model, not the distributed model. In a distributed architecture, both the production site and disaster-recovery site have an UltraVR Server, with each server managing the local site's resources.

Likewise, "In the centralized deployment scenario, the UltraVR Server deployed at the production site and the disaster recovery site are required to be deployed network connectivity between UltraVR Servers." describes the inter-server connectivity characteristic of a distributed deployment, because it presupposes UltraVR Servers at both locations. Centralized deployment uses one UltraVR Server to manage production and disaster-recovery resources, so it does not require an UltraVR Server pair to communicate across sites. This distinction is central to selecting the incorrect statements. Refer to the [Huawei Enterprise Support UltraVR documentation search](#) and the [Huawei Enterprise Support portal](#) for the applicable UltraVR product documentation and deployment guides.

QUESTION NO: 59

[Multiple Choice] In the FusionCloud troubleshooting, which of the following paths are FusionCloud running?

What is the specific path of Nova log storage for logs?

- A. /var/log/fusionsphere/component/nova-api
- B. /var/log/fusionsphere/component/novaControl
- C. /var/log/fusionsphere/component/nova-conductor
- D. /var/log/fusionsphere/component/nova-scheduler

ANSWER: A C D

Explanation:

The Nova service log directories are /var/log/fusionsphere/component/nova-api, /var/log/fusionsphere/component/nova-conductor, and /var/log/fusionsphere/component/nova-scheduler. FusionCloud organizes component logs below the common /var/log/fusionsphere/component/ hierarchy and uses service-specific directory names. Consequently, troubleshooting an API request, an instance-build workflow, or a scheduling decision begins by collecting logs from the directory that matches the responsible Nova service. The Nova API service records REST API requests and related processing activity in the nova-api directory. Nova Conductor handles operations that coordinate database access and compute-side processing, so its diagnostic output is held in the nova-conductor directory. Nova Scheduler records host-selection and placement decision information in the nova-scheduler directory. These logs are particularly useful when tracing an instance lifecycle across control-plane services: correlate timestamps, request identifiers, instance UUIDs, and user or project identifiers among the three service logs. Nova is the OpenStack Compute service, whose deployment and configuration model includes these independently operating API, conductor, and scheduler services; see the [OpenStack Blue administration guide](#) and the [Nova configuration documentation](#).

QUESTION NO: 60

HuaweiFusionSphere OpenStackimprovedOpenstackAvailability of infrastructure, which of the following components are usedActive-ActivetField Department Chang way? (multiple choice)

- A. nova-api
- B. KeyStone
- C. HAProxy
- D. RabbitMQ

ANSWER: A B

Explanation:

nova-api and **KeyStone** are appropriate for active-active deployment because they are stateless OpenStack API services. Multiple service instances can process requests concurrently when placed behind a load-balancing layer, enabling continued API access if one instance fails and allowing request traffic to be distributed across healthy instances. This design improves control-plane availability without requiring client requests to remain bound to one API process.

For Nova, nova-api exposes the Compute API and can be horizontally scaled with multiple workers or nodes. Keystone similarly provides the Identity API and supports multiple concurrently operating API endpoints that share the required backend state, such as the identity database and token-related configuration. In a highly available OpenStack architecture, these API endpoints are commonly deployed as redundant, concurrently serving instances, with a virtual endpoint or load balancer directing clients to available nodes.

This follows the OpenStack high-availability model for stateless services: API services are scaled out and made highly available through multiple active service instances, while their persistent state is stored in shared backends. See the [OpenStack High Availability Guide: Stateless services](#) and the [Nova Administration Guide](#).

QUESTION NO: 61

In the FusionSphere server virtualization backup scenario, the following describes the eBackup backup process:

Which one is wrong? (one)

- A. When using snapshot comparison technology for incremental backup, it is necessary to retain snapshot data on the production side
- B. Use CBT technology for incremental backup, which saves storage space compared to snapshot comparison.
- C. CBT technology will not be used during the first full backup
- D. When using CBT technology for incremental backup, snapshot technology will not be involved

ANSWER: D

Explanation:

“When using CBT technology for incremental backup, snapshot technology will not be involved” is the incorrect description. In an eBackup virtual-machine backup workflow, Changed Block Tracking (CBT) is used to identify the disk blocks that have changed since the previous successful backup point, allowing the incremental job to transfer only those changed blocks rather than the entire virtual disk. However, CBT does not eliminate the need for a consistent point-in-time view of the virtual machine’s disks. The backup process still uses the virtualization platform’s snapshot mechanism to preserve a stable data view while the source virtual machine can continue running and changing. eBackup can then read the required changed data from that protected point-in-time view, after which the temporary snapshot is handled according to the backup workflow. Thus, CBT and snapshots have complementary roles: CBT optimizes changed-block selection, while snapshots support online, consistent backup acquisition. This distinction is important when planning production storage capacity and monitoring snapshot lifecycle during backup windows. Huawei’s FusionSphere Virtualization documentation and enterprise support resources provide the relevant platform and backup-operation guidance: [Huawei FusionSphere Virtualization support documentation](#) and [Huawei FusionSphere Virtualization overview](#).

QUESTION NO: 62

What should I do if the user thinks about the pre-symptom of the desktop card when using the FusionAccess desktop cloud? (multiple choice)

- A. Restart the virtual machine and log in to the virtual desktop again.
- B. Disabled in service listHDPRelated Services.
- C. deleteHDA
- D. Close processes with high CPU and memory usage.

ANSWER: A D

Explanation:

The appropriate actions for a FusionAccess virtual desktop that is experiencing lag or intermittent freezing are to restart the virtual machine and then log in to the virtual desktop again, and to close processes consuming excessive CPU or memory. Restarting the desktop VM clears transient operating-system, application, and virtual-desktop-agent conditions that can cause an unresponsive or degraded session. It also establishes a fresh desktop connection, which can restore normal interaction when the issue is limited to the current VM or session.

Resource contention inside the guest operating system is another common source of perceived desktop lag. A process that continuously uses substantial processor time or exhausts available memory can delay user-interface rendering, application responses, and remote-display updates. Identifying and closing unnecessary high-CPU or high-memory processes releases resources for the operating system, the FusionAccess desktop components, and the user's workload. These steps align with standard virtual-desktop troubleshooting practice: validate guest-VM health and eliminate resource pressure before escalating the issue. Huawei's FusionAccess product and support resources are available at [Huawei FusionAccess](#) and [Huawei Enterprise Support for FusionAccess](#).

QUESTION NO: 63

An enterprise uses the cloud server backup service CSBS to ensure the security and accuracy of data.

After the data is damaged, use the cross-Az off-site backup copy to restore, generally need to manually update or switch DNS.

into business traffic switching. Manual switching is relatively time-consuming and can be deployed—the service realizes the above-mentioned business traffic.

switch automatically. (Fill in uppercase English abbreviations)

- A. CNAME
- B. GTM

ANSWER: B

Explanation:

GTM is the appropriate abbreviation because Global Traffic Management provides DNS-based traffic scheduling and disaster-recovery switchover across endpoint groups or sites. In this scenario, the enterprise restores workloads from a cross-AZ off-site backup copy after a failure. The restored environment may have different service endpoints, and changing DNS records manually delays recovery and risks operational errors. GTM can be configured with primary and standby service addresses, health checks, and routing policies. When the primary endpoint becomes unavailable, GTM detects the unhealthy status and returns the standby endpoint according to the configured policy, directing client requests to the recovered service automatically. This turns the otherwise manual DNS traffic cutover into an automated disaster-recovery process and helps reduce the recovery time experienced by users. GTM is therefore used for business access and traffic steering, while the backup service protects and restores the underlying server data. Huawei Cloud presents Global Traffic Management as a Cloud DNS capability for intelligent, health-aware traffic distribution and cross-region or multi-site disaster recovery. See the [Huawei Cloud DNS service page](#) and the [Huawei Cloud backup and recovery service page](#).

QUESTION NO: 64

(True or False)ebackupIt supports comprehensive protection of user data in four application scenarios: virtualization, public cloud, private cloud and hybrid cloudD.

- A. True
- B. False

ANSWER: A

Explanation:

The correct answer is **True**. Huawei eBackup is designed as an enterprise data-protection solution that addresses backup and recovery requirements across multiple deployment models rather than being limited to a single on-premises environment. Its protection scope includes virtualized workloads, where virtual machines and their application data require centralized backup and restoration; public-cloud workloads; private-cloud infrastructure; and hybrid-cloud environments that combine local and cloud resources. This broad coverage is important because enterprise services are commonly distributed across virtualization platforms and multiple cloud locations, while administrators still need consistent backup policies, retention management, recovery workflows, and operational visibility. eBackup therefore supports comprehensive protection for these four application scenarios, allowing data-protection capabilities to follow workloads as they are deployed or migrated between private infrastructure and cloud environments. Huawei presents its enterprise data-protection portfolio as supporting cloud-oriented, virtualized, and multi-environment backup and recovery needs. See Huawei's [data protection portfolio](#) and the [eBackup product documentation](#) for product and deployment information.

QUESTION NO: 65

FusionSphereIn a private cloud scenario, which of the following conditions may cause the cloud hard disk to fail to be deleted?

- A. The cloud disk has been mounted to the cloud host
- B. The cloud hard disk is the system disk
- C. The cloud hard disk has been transferred
- D. Cloud Drive is created for the application

ANSWER: A

Explanation:

The cloud disk has been mounted to the cloud host is correct because an attached cloud hard disk is actively associated with a virtual machine and may be in use by its operating system or applications. FusionSphere protects this attachment relationship during normal lifecycle operations so that storage is not removed while the host can still issue read or write requests to it. Before deleting an attached data disk, the administrator should first ensure that workloads have stopped using the disk, unmount any file systems from within the guest as appropriate, and detach the disk through the cloud-management interface. Once the detach operation has completed and the disk is no longer linked to a cloud host, it can be deleted subject to the tenant's permissions and any applicable protection policies. This behavior is consistent with the general cloud-volume lifecycle model, where volume attachment is a managed resource state and deletion is performed only after the attachment has been removed. Huawei's cloud documentation resources are available through the [FusionSphere OpenStack product documentation portal](#); the comparable OpenStack volume lifecycle and detach workflow is documented in the [OpenStack Horizon volume management guide](#).

QUESTION NO: 66

(Multiple choice) Which of the following isFusionCompute. What are the necessary conditions for a virtual machine to communicate properly with other virtual machines?

- A. Configure the correct virtual machineIPaddress

- B. Configure the correct configuration for this virtual machineMA, Address
- C. Configure pinning for this virtual machineIPAddress
- D. Configure the gateway address for this virtual machine

ANSWER: A B

Explanation:

For virtual machines on a FusionCompute virtual network to communicate correctly, each VM needs a valid IP address that belongs to the intended network and does not conflict with another endpoint. The IP address provides the Layer 3 identity used by the guest operating system and peer systems for packet delivery. The VM must also have a valid and unique MAC address on its virtual NIC. FusionCompute and the connected virtual switch use the MAC address for Layer 2 frame forwarding, including learning the source location and delivering frames to the appropriate virtual port. In normal deployments, FusionCompute assigns the MAC address when the virtual NIC is created, but the configured value must remain valid and unique if it is customized or restored from a template. Together, correct IP and MAC addressing allow the VM to participate in local subnet communication and enable the virtual switching and IP networking functions required for VM-to-VM connectivity. Huawei's enterprise support portal provides FusionCompute product documentation and networking configuration guidance at [Huawei Enterprise Support](#). Additional Huawei Cloud infrastructure networking concepts are available through [Huawei Cloud](#).

QUESTION NO: 67

FusionCare has two functions: health inspection and information collection. Which of the following information collection modes are available some items? (many)

- A. Collection by scene
- B. Collect by component
- C. Collection by account
- D. to collect by time period

ANSWER: A B

Explanation:

FusionCare information collection supports task templates organized around the operational context in which troubleshooting or maintenance is being performed. **Collection by scene** is therefore available: an administrator selects a predefined service or fault-handling scenario, and FusionCare gathers the relevant diagnostic information for that scenario. This approach makes it easier to obtain a consistent data set for common maintenance and support workflows.

Collect by component is also available. It allows collection to be targeted at a particular infrastructure or software component, so that logs, configuration data, alarms, and related status information can be gathered for the selected component. Component-oriented collection is useful when analysis has already narrowed the issue to a specific product area or system service. Together, scenario-based and component-based collection provide both workflow-oriented and granular collection entry points, supporting efficient health inspection and fault diagnosis. Huawei's enterprise support portal provides product documentation and service resources for FusionCare and related cloud-computing products: [Huawei Enterprise Support](#). Huawei also describes its enterprise services and support offerings at [Huawei Enterprise Services](#).

QUESTION NO: 68

(Multiple Choice) AboutFusionSphereOf the four supported disk configuration modes, which of the following descriptions are correct?

- A. Virtualized storage supports normal disks, thin disks, and normal delayed-zeroed disks
- B. Raw device disks willSANphysical storageLUNDirectly as a disk, bound to the business virtual machine for use

C. Non-virtualized storage of block storage only supports ordinary disks and raw device disks

D. NASStorage supports normal, thin disks and raw device disks

ANSWER: A B C

Explanation:

FusionSphere/FusionCompute distinguishes virtual-disk provisioning modes according to the underlying storage architecture. For virtualized storage, the supported virtual-disk choices are normal disks, thin disks, and normal delayed-zeroed disks. These modes allow the virtualization platform and its storage virtualization layer to create and manage the virtual disk while applying the selected allocation behavior. A raw device disk is different: it maps a physical SAN storage LUN directly to a service virtual machine, providing the VM with direct use of that LUN rather than creating a conventionally provisioned virtual disk on the datastore. For non-virtualized block storage, the applicable choices are ordinary disks and raw device disks. This reflects the direct relationship between the host, the block-storage LUN, and the VM disk mapping in a non-virtualized storage configuration. These distinctions are important when selecting disk provisioning because the available modes, capacity-allocation behavior, and LUN presentation model depend on the storage type configured for the cluster. Huawei's FusionSphere documentation and product-support resources provide the relevant storage and virtualization configuration guidance; see [Huawei FusionSphere Virtualization support](#) and [Huawei FusionSphere product information](#).

QUESTION NO: 69

The ManageOne system can realize the connection of third-party systems by configuring the 3A authentication protocol, and the system no longer passes the

The user management function performs user authentication, and the user authentication is implemented by the 3A system when logging in. Where 3A stands for which of the following content ?

A. Accounting

B. Accessing

C. Authentication

D. Authorization

ANSWER: A C D

Explanation:

The correct content is **Authentication**, **Authorization**, and **Accounting**. Together, these are the three functions conventionally represented by 3A in network and identity-access management systems. Authentication establishes the identity of the person or system attempting to log in, such as by validating a username/password, certificate, token, or an external identity-provider assertion. Authorization then determines the permissions available to that authenticated identity, including the resources, services, projects, or operations it is allowed to access. Accounting records access and usage activity, providing traceability for activities such as login events, session duration, resource consumption, and administrative actions.

In a ManageOne integration with a third-party 3A platform, this division enables the external platform to centrally control identity verification and access policy while retaining auditable records of user activity. This supports centralized identity governance, consistent role and privilege enforcement, and operational auditing across connected systems. The Internet Engineering Task Force defines this model as authentication, authorization, and accounting in its AAA architecture terminology; see [RFC 3539: AAA Architecture](#). Huawei's enterprise support portal also provides ManageOne product documentation and integration guidance at [Huawei ManageOne Support](#).

QUESTION NO: 70

(Multiple choice) The following about FusionAccess middle TCM Description of the component, which options are correct?

- A. TCM can upgrade HDA.
- B. TCM can manage TC upgrades.
- C. TCM can deliver policies to virtual desktops.
- D. TCM can view the run logs of managed TCs.

ANSWER: B D

Explanation:

TCM (Terminal Configuration Manager) is the FusionAccess component used for centralized management of TC thin-client terminals. A core TCM lifecycle function is software version management: administrators can centrally organize and perform upgrades for managed TCs, avoiding the need to upgrade each endpoint individually. TCM also provides operational maintenance capabilities for these managed terminals, including collecting and viewing terminal run logs. These logs help administrators investigate endpoint faults, assess operating status, and support routine maintenance of the virtual-desktop access environment. Therefore, the statements that TCM can manage TC upgrades and view the run logs of managed TCs accurately describe its endpoint-management role. Huawei's FusionAccess documentation and product resources describe the platform's centralized desktop and endpoint management architecture; see the [Huawei FusionAccess support page](#) and the [Huawei FusionAccess product page](#).

QUESTION NO: 71

(Single choice) When backing up cloud computing business, support multiple backup types, which one of the following backup Type is not supported?

- A. Cumulative full backup
- B. Differential Incremental Backup
- C. Cumulative incremental backup D. Full backup

ANSWER: A

Explanation:

Cumulative full backup is the unsupported backup type. In Huawei cloud backup terminology, a full backup is a complete baseline copy of protected data, while subsequent backups are normally managed as incremental backups that record changed blocks or files relative to the backup chain. Huawei Cloud Cloud Backup and Recovery documents describe this backup-chain approach: the first backup is full and later backups are incremental, with the system managing the underlying data and restoration dependencies. "Cumulative full backup" is not a defined Huawei backup mode because it combines two distinct concepts into a nonstandard type. A full backup is independently complete at the time it is created; "cumulative" is terminology more commonly associated with incremental or differential change accumulation, rather than a separate form of full backup. Therefore, when identifying a backup type unsupported by the platform's documented backup model, cumulative full backup is the correct selection. See Huawei Cloud's [CBR backup concepts documentation](#) and the [Cloud Backup and Recovery product documentation](#) for the service's backup mechanisms and capabilities.

QUESTION NO: 72

The following aboutFusionComputeThe description of virtual machine snapshot restore, which option is incorrect?

- A. If the port group to which the virtual machine NIC is connected is deleted after the snapshot is created, the system will automatically restore the port group and NIC when the snapshot is used to restore the virtual machine.
- B. If a certain disk attribute of the virtual machine is independent and persistent, when the virtual machine is restored using a snapshot, the data on the disk will not be restored.
- C. If a new disk is bound to the virtual machine after the snapshot is created, the new disk will be automatically unbound from the virtual machine when the snapshot is used to restore the virtual machine.

D. When a virtual machine fails or data needs to be restored, the existing snapshot of the virtual machine can be used to restore the data of the virtual machine to the state when the snapshot was created.

ANSWER: A

Explanation:

The statement that a deleted port group and the virtual machine NIC will be automatically restored is incorrect. A FusionCompute virtual-machine snapshot is designed to preserve and roll back the virtual machine's state and snapshot-supported virtual-machine configuration; it is not a backup or recovery mechanism for independently managed virtual-network resources. A port group is a network-side object managed in the virtualization environment. Consequently, if that port group has been deleted after the snapshot was taken, restoring the virtual-machine snapshot cannot recreate the deleted port group automatically. The administrator must first recreate or otherwise provide an appropriate port group, then ensure that the virtual machine NIC is connected to the required network. This distinction is important during recovery planning: snapshots can return eligible VM disks and VM configuration to their recorded state, but they do not replace protection of external infrastructure dependencies such as network configuration. Huawei's enterprise support portal provides the FusionCompute product documentation and version-specific virtualization administration guidance at [Huawei Enterprise Support](#). Huawei also describes FusionCompute as a virtualization platform whose resources, including virtual networking, are administered as platform infrastructure at [Huawei FusionCompute](#).

QUESTION NO: 73

(Multiple choice) In HuaweiFusionsphere Openstack, which of the following services needs to interact with all other services?

- A Cinder
- B. Neutron
- C. keystone
- D Nowa

Answer: C

- A. Cinder
- B. Neutron

ANSWER:

Explanation:

Keystone is the OpenStack Identity service and is the common authentication and authorization point used throughout a FusionSphere OpenStack deployment. It maintains users, projects, roles, domains, service definitions, and API endpoints. Before a user, tenant application, or OpenStack component invokes another service API, it obtains an authentication token from Keystone or presents a token for validation. The receiving service relies on Keystone identity information and the service catalog to determine whether the request is authorized and where the requested service endpoint is located.

This central identity role means Keystone must integrate broadly with the OpenStack control plane. Compute, networking, block storage, image, orchestration, dashboard, and other services use its token-based access-control model to provide a consistent multi-tenant security boundary. The OpenStack Keystone documentation describes Identity as the service that provides authentication, authorization, and a service catalog: [OpenStack Keystone documentation](#). Its architecture and token-validation behavior are also described in the [OpenStack Identity concepts guide](#). Therefore, Keystone is the service that needs to interact with all other OpenStack services.

QUESTION NO: 74

Users can either FusionAccess Modify the specifications of the virtual machine on the FusionCompute Modify the specifications of the virtual

machine above.

A. True

B. False

ANSWER: A

Explanation:

True is correct. FusionAccess desktop virtual machines are implemented on the underlying FusionCompute virtualization platform, so their compute specifications—such as vCPU count, memory capacity, and related virtual hardware resources—are managed in coordination with both systems. FusionAccess provides desktop-oriented management functions for administrators, while FusionCompute provides the infrastructure-level virtual-machine management capabilities. This allows an administrator to adjust the specifications through FusionAccess where that operation is supported by the desktop management workflow, or to perform the virtual-machine specification adjustment on FusionCompute when managing the underlying VM directly.

The key operational requirement is that the desktop VM remains correctly associated with its FusionAccess desktop resources and that any applicable synchronization, status checks, and power-state prerequisites are observed after a change. This coordinated model reflects FusionAccess's role as a VDI management solution built on FusionCompute's virtualization and resource-management foundation. Huawei describes FusionAccess as a desktop-cloud solution using virtualization resources and documents its deployment and management relationship with FusionCompute in its product documentation. See the [Huawei FusionAccess documentation portal](#) and the [Huawei FusionCompute documentation portal](#).

QUESTION NO: 75

Backup supports two backup methods. CBT does not keep snapshots on the production side, and occupies its storage space.

Smaller, snapshots are stored on the production side compared to snapshots, which take up more storage space.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In the CBT (Changed Block Tracking) backup method, the backup system identifies blocks that have changed since the preceding backup and transfers only the required changed data. This method is designed to avoid retaining a long-lived production-side snapshot as the backup recovery point. Consequently, it imposes less production-storage consumption than a backup approach that maintains snapshots on the production environment. Snapshot-based backup uses a point-in-time snapshot to preserve the source data state while backup processing occurs. Since that snapshot is retained on production storage for the relevant period, it consumes storage capacity there; its actual consumption can grow as blocks change after the snapshot is created. Therefore, the statement accurately distinguishes CBT's lower production-side snapshot storage impact from the snapshot backup method's need to retain snapshots on the production side. This behavior supports selecting CBT when minimizing storage pressure on production infrastructure is an important operational requirement, while still enabling incremental protection through changed-block identification. Huawei's data-protection documentation and backup-service guidance describe snapshot and incremental-backup concepts relevant to this distinction: [Huawei Enterprise Cloud Computing Support](#) and [Huawei Cloud Volume Backup documentation](#).

QUESTION NO: 76

Backup is the process of copying all or part of the data set from the hard disk or array applied to the host to other storage media in order to prevent data loss due to system operating errors or system failures. Which of the following options are backup media?

- A. CD library
- B. cloud storage
- C. Disk Array
- D. cassette library

ANSWER: A B C D

Explanation:

All listed items can serve as backup media because each provides a destination separate from the production data source where backup copies can be retained and recovered. A CD library uses optical discs to store backup data and is an example of removable-media-based backup storage. Cloud storage provides offsite, scalable capacity for backup copies and is commonly used to improve resilience against local hardware or site failures. A disk array can be used as a disk-based backup target, supporting relatively fast backup and restoration operations; enterprise backup systems often use disk storage as a staging or primary backup repository. A cassette library, more commonly called a tape library, uses removable tape cartridges and remains a widely used medium for long-term retention, archival, and large-volume backup workloads. The essential characteristic is not the underlying technology—optical, cloud, disk, or tape—but the use of that storage medium to hold a recoverable copy of data. Huawei backup solutions support backup repositories based on disk storage and cloud-oriented architectures, while tape remains a standard backup medium in enterprise data-protection designs. See [Huawei OceanProtect documentation](#) and the [NIST storage guidance](#).

QUESTION NO: 77

(multiple choice) The company plans to build a cloud data center and switch the existing office system to run in a cloud computing environment. When planning the total storage capacity, which of the following factors need to be considered?

- A. LUN
- B. RAID
- C. IOPS
- D. Disk capacity

ANSWER: B D

Explanation:

RAID and **Disk capacity** must be considered when sizing total storage capacity. Raw disk capacity is the starting point: planners must count the number of disks and their individual usable capacities to establish the physical capacity available to the storage system. RAID configuration then determines how much of that raw capacity remains usable for service data. RAID groups reserve capacity for redundancy information, such as parity or mirrored copies, and may also include hot-spare capacity according to the selected protection design. Consequently, two systems with identical disk counts and disk sizes can provide materially different usable capacities when configured with different RAID levels. Capacity planning for a cloud data center should therefore begin with the required service capacity, allow for growth and operational reserve, and calculate sufficient raw disk capacity after accounting for the RAID protection overhead. This approach ensures that the storage pool can meet application data requirements while retaining the reliability and fault-tolerance expected from enterprise storage. Huawei's storage documentation and technical-support portal provide product-specific RAID and capacity-calculation guidance: [Huawei Enterprise Storage Support](#) and [Huawei Storage Products](#).

QUESTION NO: 78

In FusionCloud, the cloud fortress can provide terminal security protection services and database security protection. Protection capabilities, such as providing intrusion detection and defense, data desensitization, data auditing, fault repair and other security protection capabilities to ensure the security of related services.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In FusionCloud security architecture, the cloud fortress function is intended to provide centrally controlled, auditable protection for access to important assets, including terminal and database-related resources. It helps enforce controlled operations and security monitoring around administrator access paths, which is a key requirement for protecting cloud-hosted services. The stated capabilities align with the broader security protections available in the FusionCloud environment: intrusion detection and defense help identify and respond to malicious activity; data desensitization limits unnecessary exposure of sensitive values; data auditing creates traceable records of database and operational activity; and repair capabilities support remediation when faults or security issues are identified. Together, these controls protect service availability, confidentiality, and accountability. Huawei's Cloud Bastion Host documentation describes the bastion-host role in centralized management, controlled access, and audit of operations on cloud resources. Huawei's database security documentation also describes database-focused protections such as auditing and sensitive-data safeguards. Therefore, the statement accurately summarizes the terminal and database security protection capabilities associated with the FusionCloud cloud fortress solution. See [Huawei Cloud Bastion Host documentation](#) and [Huawei Database Security Service documentation](#).