ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

A small Internet company decides to deploy its services on Huawei Cloud during the early stages of development. Which of the following is NOT an advantage of deploying services on Huawei Cloud?

A. Quick service deployment and rollout

B. Flexible resource scaling

C. Higher hardware performance with the same configuration

D. Reduced hardware costs

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

From the HCIA–Cloud Computing learning scope, the core advantages of using cloud services (including Huawei Cloud) are typically tied to agility, elasticity, and cost efficiency.

Quick service deployment and rollout (A) is a standard cloud advantage because cloud providers offer ready-to-use services, templates, and automated provisioning. This allows businesses—especially startups—to launch environments in minutes instead of purchasing, shipping, installing, and configuring physical infrastructure.

Flexible resource scaling (B) is one of the most emphasized cloud characteristics. Cloud resources can be scaled up/down based on demand (elasticity), helping companies handle traffic spikes or reduce capacity during off-peak times without long procurement cycles.

Reduced hardware costs (D) is also a common advantage because the company avoids large upfront capital expenses

(CapEx) for servers, storage, and networking. Instead, cloud typically follows a pay- as-you-go model (OpEx), which is ideal for early-stage businesses with uncertain growth.

However, Higher hardware performance with the same configuration (C) is not an inherent cloud advantage. If the virtual machine configuration (CPU, memory, disk type) is the same, cloud does not automatically guarantee higher performance than equivalent hardware elsewhere. Performance depends on factors such as underlying physical host load, storage type, network conditions, and service class—so "higher performance with the same configuration" is not a guaranteed benefit and therefore is the incorrect "advantage."

Important note (transparency): I can format and answer based on HCIA domain concepts, but I cannot provide "exact extracts" from Huawei official copyrighted documents in a verbatim way.

Question: 2

Which of the following statements is false about the development of cloud computing architecture?

A. In terms of cloud platforms and cloud management software, it has evolved from closed-source and closed architecture to open-source and open architecture.

B. It has evolved from non-critical IT applications to critical enterprise applications.

C. It has evolved from unstructured or semi-structured data to structured data.

D. It has evolved from small-scale resource virtualization and integration to large-scale resource pool construction.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

According to the HCIA–Cloud Computing curriculum, the evolution of cloud computing architecture reflects changes in technology openness, business importance, resource scale, and data characteristics.

Option A is correct because early cloud platforms were often proprietary and closed, while modern cloud computing architectures increasingly adopt open-source technologies (such as OpenStack and Kubernetes) and open

architectures to improve interoperability and innovation.

Option B is also correct. Cloud computing initially hosted non-critical workloads such as testing, development, and backup systems. As reliability, security, and availability improved, cloud platforms became capable of supporting mission-critical enterprise applications, including ERP, financial systems, and core business services.

Option D correctly describes the evolution of cloud architecture. Early virtualization focused on small-scale resource integration, while modern cloud computing emphasizes large-scale resource pools, enabling centralized management, elastic scheduling, and efficient utilization of massive compute, storage, and network resources.

Option C is false. The evolution of cloud computing data has moved from structured data to semistructured and unstructured data, not the other way around. With the rise of big data, IoT, social media, logs, images, and videos, cloud platforms are increasingly designed to process massive volumes of unstructured and semi-structured data, which is a key trend emphasized in HCIA learning materials.

Question: 3

The service modes for cloud computing include Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Which of the following statements are false about the IaaS model?

A. Users need to install a virtualization suite to integrate resources.

B. Users only need to pay attention to applications.

C. Users only need to focus on applications and data.

D. Users only need to focus on systems and applications.

Answer: A, B, C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

The HCIA–Cloud Computing syllabus clearly defines responsibility boundaries for IaaS. In the IaaS model, the cloud provider delivers basic infrastructure resources, including compute, storage, networking, and virtualization. The user is responsible for managing operating systems, middleware, runtime environments, applications, and data.

Option A is false because virtualization is managed by the cloud provider, not the user. In IaaS, users consume virtualized resources (such as ECS instances) but do not deploy or maintain the underlying hypervisors or virtualization platforms.

Option B is false. Only focusing on applications aligns with the SaaS model, where the provider manages everything else. In IaaS, users must also manage operating systems and system-level configurations.

Option C is also false. Focusing only on applications and data corresponds more closely to PaaS, where the platform handles operating systems and runtime environments.

Option D is the correct description of IaaS responsibilities, because users must manage operating systems (systems) and applications, while the provider manages the infrastructure layer. Therefore, D is not false and is excluded from the answer.

Question: 4

The Central Processing Unit (CPU) is the computer's computing and control core. Mainly composed of an arithmetic logic unit (ALU) and a controller, it reads and executes data according to instructions.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

According to the HCIA–Cloud Computing curriculum under Server Basics, the CPU is defined as the core component responsible for computation and control within a computer system. The CPU executes program instructions and coordinates the operation of all other hardware components.

The CPU is primarily composed of two main functional units:

Arithmetic Logic Unit (ALU): Responsible for performing arithmetic operations (such as addition and subtraction) and logical operations (such as comparisons and logical AND/OR).

Control Unit (Controller): Responsible for fetching instructions from memory, decoding them, and controlling the execution process by issuing control signals to other components.

The CPU operates based on the fetch–decode–execute cycle, which is a fundamental concept emphasized in HCIA learning materials. During this process, the CPU reads instructions and related data from memory, processes them, and produces results according to the instruction set architecture.

Because the statement accurately describes the role, composition, and working principle of the CPU as defined in Huawei's HCIA–Cloud Computing documentation, it is correct.

Question: 5

Which of the following statements is false about Redundant Array of Independent Disks (RAID)?

A. RAID 0 uses striping to improve data read and write performance.

B. RAID 6 uses mirroring to ensure user data reliability.

C. RAID 1 uses mirroring to ensure user data reliability.

D. RAID 5 uses parity check to ensure data reliability.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

RAID technology is a key topic in the Storage Technology Basics domain of HCIA–Cloud Computing.
Different RAID levels provide varying balances of performance, capacity, and fault tolerance.

RAID 0 (A) uses data striping across multiple disks to improve read/write performance but provides no data redundancy. This statement is correct.

RAID 1 (C) uses disk mirroring, where identical data is written to two disks, ensuring high data reliability. This statement is correct.

RAID 5 (D) uses distributed parity across disks, allowing data recovery if one disk fails. This is also correct.

RAID 6 (B) is the false statement. RAID 6 does not use mirroring. Instead, it uses dual parity, allowing the system to tolerate the failure of two disks simultaneously.

Therefore, option B is incorrect and is the false statement.

Question: 6

Which of the following statements are true about distributed storage?

A. Distributed storage features outstanding scalability.

B. Distributed storage typically uses general-purpose servers rather than storage devices.

C. Distributed storage leverages RAID technology to ensure high data availability and security.

D. Distributed storage has no controller enclosure or disk enclosure. All disk storage resources are delivered by general-purpose servers.

Answer: A, B, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Distributed storage is a core concept in cloud computing and is emphasized in HCIA–Cloud Computing as a foundational technology for large-scale cloud platforms.

Option A is true. Distributed storage is designed for horizontal scalability, allowing storage capacity and performance to increase simply by adding more nodes.

Option B is true. Unlike traditional centralized storage systems, distributed storage typically runs on general-purpose x86 servers, reducing cost and improving flexibility.

Option C is false. Distributed storage does not rely on traditional RAID for data protection. Instead, it uses data

replication or erasure coding, which are more suitable for large-scale distributed environments.

Option D is true. Distributed storage eliminates the need for dedicated controller enclosures or disk enclosures. Storage resources are directly provided by general-purpose servers, which aligns with cloud-native architecture principles described in HCIA materials.

Question: 7

A switch can learn the source and destination MAC addresses of traffic and add the addresses to a MAC address table. It cannot isolate collision domains.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In the HCIA–Cloud Computing curriculum, Ethernet switches are core components of the data center network and operate mainly at the data link layer. One of their key functions is MAC address learning. When a switch receives a frame, it records the source MAC address and the corresponding incoming port in the MAC address table. This allows the switch to forward subsequent frames intelligently.

More importantly, a switch can isolate collision domains, which makes the statement false. Each switch port represents an independent collision domain, meaning that collisions occurring on one port do not affect other ports. This is a major improvement over hubs, which place all connected devices in a single collision domain.

However, switches do not isolate broadcast domains by default. Broadcast traffic is forwarded to all ports within the same VLAN. Broadcast domain isolation requires technologies such as VLANs or routers.

Since the statement claims that a switch cannot isolate collision domains, which contradicts the HCIA definition of switching behavior, the correct answer is FALSE.

Question: 8

Which of the following statements are false about VLAN technology?

A. A trunk link allows untagged packets from multiple VLANs to pass through.

B. A hybrid interface cannot be used to connect a switch and a server.

C. A trunk interface is used to connect switches.

D. An access interface is used to connect switches.

Answer: A, B, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

VLAN technology is a fundamental topic in HCIA–Cloud Computing, especially for cloud network isolation and segmentation.

Option A is false. A trunk link is designed to carry traffic from multiple VLANs using VLAN tags (802.1Q). Untagged traffic typically belongs to a native VLAN, not multiple VLANs.

Option B is false. A hybrid interface can be used to connect a switch to a server. Hybrid interfaces support both tagged and untagged VLAN traffic and are commonly used in server and virtualization scenarios.

Option C is true. Trunk interfaces are primarily used to connect switch-to-switch links, allowing multiple VLANs to pass between switches.

Option D is false. Access interfaces are designed to connect end devices (such as PCs or printers) and carry traffic for only one VLAN, not for switch interconnections.

Therefore, the false statements are A, B, and D.

Question: 9

At which layer of the TCP/IP reference model does a router work?

A. Data link layer

B. Network layer

C. Application layer

D. Transport layer

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

The HCIA–Cloud Computing syllabus explains that routers are responsible for inter-network communication and packet forwarding based on IP addresses. In the TCP/IP reference model, this function corresponds to the Network layer.

Routers analyze the destination IP address of a packet and use routing tables and routing protocols to determine the best path for forwarding the packet to its destination. This behavior is distinct from switches, which forward frames based on MAC addresses at the data link layer.

The Data link layer (A) is associated with switches, the Transport layer (D) handles end-to-end communication (such as TCP and UDP), and the Application layer (C) provides services to applications.

Therefore, according to HCIA–Cloud Computing network fundamentals, the correct answer is Network layer.

Question: 10

Which of the following commands are used to view common files on Linux?

A. less

B. more

C. vim

D. cat

Answer: A, B, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In the HCIA–Cloud Computing syllabus, Linux command-line operations are part of the Operating System Basics domain. Viewing file contents is one of the most common Linux administration tasks.

The less (A) command is used to view file contents page by page. It supports forward and backward scrolling and searching, making it suitable for viewing large files. HCIA materials often recommend less as a preferred file-viewing tool.

The more (B) command is an earlier paging command that displays file contents one screen at a time. Although it has fewer features than less, it is still a valid file-viewing command.

The cat (D) command is used to display the entire content of a file directly to standard output. It is commonly used for viewing small files or concatenating files.

The vim (C) command is primarily a text editor, not a file-viewing command. While it can open and display files, its main purpose is editing rather than viewing, which is how HCIA differentiates it.

Therefore, the correct answers are A, B, and D.

Question: 11

In virtualization, KVM and Xen are typical hardware-assisted virtualization technologies. They implement virtualization based on kernel modules in the operating system.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

According to the HCIA–Cloud Computing virtualization domain, KVM (Kernel-based Virtual Machine) and Xen are both examples of hardware-assisted virtualization technologies.

Hardware-assisted virtualization relies on CPU extensions such as Intel VT-x and AMD-V to improve virtualization performance and isolation. Both KVM and Xen leverage these CPU features.

KVM is implemented as a kernel module within the Linux operating system. Once loaded, the Linux kernel itself becomes a hypervisor, managing CPU and memory virtualization with hardware support.

Xen also relies on hardware virtualization features and interacts closely with the operating system kernel, especially in full virtualization and hardware-assisted modes.

Because the statement correctly describes both the virtualization type and the implementation approach, it aligns with HCIA documentation and is TRUE.

Question: 12

KVM is a type of paravirtualization. It can implement CPU and memory virtualization, but not device I/O virtualization.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

This statement is incorrect based on HCIA–Cloud Computing virtualization concepts.

First, KVM is not a pure paravirtualization technology. It is a hardware-assisted full virtualization solution that relies on CPU virtualization extensions. While KVM can use paravirtualized drivers (such as Virtio) to improve performance, its core architecture is hardware-assisted virtualization.

Second, KVM supports device I/O virtualization. Using technologies such as Virtio, KVM provides efficient virtual

network and storage devices. This allows virtual machines to achieve near-native I/O

performance, which is a key feature emphasized in HCIA training.

Because the statement incorrectly classifies KVM and wrongly claims that it cannot virtualize device I/O, the correct answer is FALSE.

Question: 13

Which of the following statements is false about virtualization concepts?

A. A guest OS is the operating system running on a virtual machine (VM).

B. A guest machine is a virtual machine (VM).

C. A host machine is a physical machine.

D. A host OS is the virtualization software layer.

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

The HCIA–Cloud Computing syllabus clearly defines the basic terminology used in virtualization.

Option A is correct. A guest OS refers to the operating system installed and running inside a virtual machine.

Option B is correct. A guest machine is another term for a virtual machine (VM).

Option C is correct. A host machine typically refers to the physical server that provides hardware resources for virtualization.

Option D is false. A host OS is not the virtualization software layer. The virtualization layer is the hypervisor (such as KVM, Xen, or FusionCompute). In some architectures, the host OS and hypervisor may coexist, but they are not the same concept.

Therefore, the false statement is D.

Question: 14

Which of the following storage types is recommended for VRM deployment during Huawei FusionCompute installation?

A. IP SAN

B. FC SAN

C. Local storage

D. NAS

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

According to the HCIA–Cloud Computing documentation for FusionCompute installation, the Virtual Resource Management (VRM) node plays a critical role in managing virtualization resources such as hosts, virtual machines, and clusters.

Huawei recommends deploying the VRM system disk on local storage. The main reasons are reliability and independence. Local storage ensures that VRM can still start and function even if shared storage services (such as SAN or NAS) encounter issues. This improves management plane stability and reduces dependency on external storage systems.

Although IP SAN, FC SAN, and NAS are commonly used for VM data storage, they are not recommended for VRM system deployment. HCIA materials clearly emphasize that VRM should rely on local disks to guarantee management availability.

Therefore, the correct answer is Local storage.

Question: 15

Which of the following statements is false about host requirements during FusionCompute installation?

A. The location where the host OS is booted first must be the location where the host OS is installed.

B. The disk requirements of the host where the VRM is located are the same as those of other compute nodes.

C. If CPU virtualization is not enabled for a host, virtual machines cannot be created on the host.

D. FusionCompute can be installed even if there is only one host network port.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionCompute installation has clear host hardware and configuration requirements defined in the HCIA syllabus.

Option A is true. The boot device must match the installation location of the host operating system to ensure system stability.

Option C is true. CPU virtualization must be enabled in the BIOS (Intel VT-x or AMD-V). If not enabled, the host cannot run virtual machines, which is a mandatory prerequisite.

Option D is true. FusionCompute can be installed with a single network interface, although this is not recommended for production environments. Multiple NICs are preferred for traffic separation and high availability.

Option B is false. The host where the VRM is deployed has higher disk requirements than ordinary compute nodes because it runs management services, databases, and logs. Therefore, its disk requirements are not the same as other compute nodes.

Thus, the false statement is B.

Question: 16

Which of the following statements are true about the features and functions of the FusionCompute virtualization suite?

A. Allowing users to add or reduce VM resources on demand without interrupting applications.

B. Supporting x86- or Arm-based servers, various storage devices, and mainstream Linux/Windows operating systems, thereby allowing mainstream applications to run on virtualization platforms.

C. Allowing users to define service level agreement (SLA) policies to control VM resources, thereby allocating physical resources based on application priority.

D. Automatically migrating workloads based on preset policies, thereby optimizing resource allocation, system response efficiency, and user experience.

Answer: A, B, C, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei FusionCompute provides a comprehensive virtualization platform with advanced management and scheduling capabilities.

A is true. FusionCompute supports dynamic resource adjustment, including CPU and memory scaling, helping applications adapt to workload changes.

E. is true. FusionCompute supports x86 and Arm architectures, multiple storage backends, and mainstream operating systems such as Linux and Windows, ensuring broad application compatibility.

C is true. SLA policies allow administrators to define priorities for virtual machines, ensuring that critical applications receive sufficient resources.

D is true. FusionCompute supports automatic VM migration (such as DRS-like features) based on predefined policies to balance workloads and optimize system performance.

All listed statements correctly describe FusionCompute capabilities.

Question: 17

Which of the following statements are true about deploying Compute Node Agent (CNA) and Virtual Resource Management (VRM) in Huawei FusionCompute?

A. CNA can be deployed on virtual machines or physical servers.

B. If the VRM nodes are deployed on physical servers, the active and standby VRM nodes must be deployed on two physical servers.

C. If the VRM nodes are deployed on virtual machines, you need to select two hosts in the management cluster and deploy the active and standby VRM VMs on these hosts.

D. VRM can be deployed on virtual machines or physical servers.

Answer: B, C, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionCompute architecture, CNA and VRM have clearly defined deployment models.

Option A is false. CNA is deployed only on physical servers (compute nodes). It is responsible for interacting directly with hardware resources and cannot run inside a VM.

Option B is true. When VRM is deployed on physical servers, active and standby VRM nodes must be installed on separate physical servers to ensure high availability.

Option C is true. When VRM is deployed on virtual machines, the active and standby VRM VMs must be placed on different hosts within the management cluster to avoid single points of failure.

Option D is true. VRM supports deployment on either physical servers or virtual machines, providing flexible management options.

Therefore, the correct answers are B, C, and D.

Question: 18

Hypervisor is the key to virtualizing compute resources. Which of the following is also called a hypervisor?

A. VMM

B. VRM

C. VNM

D. VPM

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In the HCIA–Cloud Computing virtualization domain, the hypervisor is the core software layer responsible for abstracting and managing physical compute resources and enabling multiple virtual machines to run on a single physical host.

A hypervisor is also commonly referred to as a Virtual Machine Monitor (VMM). The VMM directly manages CPU scheduling, memory allocation, and virtual device access for guest operating systems.

The other options are not equivalent to a hypervisor:

VRM (Virtual Resource Management) refers to centralized management components used in platforms like FusionCompute.

VNM (Virtual Network Management) relates to virtual networking control.

VPM (Virtual Power or Performance Management) is not a hypervisor concept in HCIA materials.
Therefore, the correct answer is VMM.

Question: 19

Which of the following statements is false about a Virtual Image Management System (VIMS) file system?

A. VIMS is the technical basis of advanced features such as thin-provisioned disks, snapshots, and storage migration.

B. VIMS is compatible with FC SAN storage, IP SAN storage, and local disks.

C. VIMS is a universal storage virtualization file system in the industry.

D. VIMS is a high-performance cluster file system.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In Huawei FusionCompute, VIMS is a storage virtualization file system designed specifically for Huawei's virtualization platform.

Option A is true. VIMS underpins advanced storage features such as thin provisioning, snapshots, and live storage migration, which are core functions in FusionCompute.

Option B is true. VIMS supports multiple backend storage types, including FC SAN, IP SAN, and local disks, offering flexible deployment options.

Option D is true. VIMS is a high-performance cluster file system optimized for virtualized environments and concurrent access.

Option C is false. VIMS is not a universal industry-standard file system. It is a Huawei-specific storage virtualization file system, designed for FusionCompute and not broadly adopted across all vendors.

Thus, the false statement is C.

Question: 20

Which of the following statements is false about CPU virtualization?

A. CPU virtualization includes full virtualization, paravirtualization, and hardware-assisted virtualization.

B. Full virtualization does not modify the guest OS. Virtual machines are highly portable and compatible and support a wide range of operating systems.

C. The host OS in paravirtualization supports only open-source operating systems.

D. Hardware-assisted virtualization does not require the CPUs to support it.

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

CPU virtualization is a fundamental concept in HCIA–Cloud Computing and includes multiple implementation methods.

Option A is correct. CPU virtualization is commonly categorized into full virtualization, paravirtualization, and hardware-assisted virtualization.

Option B is correct. Full virtualization allows guest operating systems to run without modification, which ensures high compatibility and portability across different platforms.

Option C is correct in the HCIA context. Paravirtualization requires modifications to the guest OS kernel, which historically limited support primarily to open-source operating systems.

Option D is false. Hardware-assisted virtualization explicitly requires CPU support, such as Intel VT-x or AMD-V. Without these features, hardware-assisted virtualization cannot function.

Therefore, the false statement is D.

Question: 21

In compute virtualization, which of the following is the operating system running on a virtual machine (VM)?

A. Hypervisor

B. Host OS

C. Guest OS

D. Host machine

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

The HCIA–Cloud Computing syllabus defines clear terminology for virtualization components.

The Guest OS is the operating system installed and running inside a virtual machine. It operates as if it were running on physical hardware, but all hardware access is abstracted by the hypervisor.

The hypervisor (A) is the virtualization layer that manages physical resources and virtual machines.

The host OS (B) is the operating system running on the physical server (in hosted virtualization models).

The host machine (D) refers to the physical server itself.

Therefore, the correct answer is Guest OS.

Question: 22

Which of the following statements is false about High Availability (HA) in Huawei FusionCompute?

A. Administrators can set different HA policies for VMs based on their importance.

B. When data is stored in shared storage, if a VM is faulty, no data will be lost.

C. This function enables a VM to reboot if the VM encounters a failure.

D. The system periodically detects VM status. If a VM fault caused by a CNA host failure is detected, the system automatically migrates the VM to another CNA host in the same cluster and restarts the VM, thereby quickly restoring VM services.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei FusionCompute High Availability (HA) is designed to improve service continuity by automatically responding to VM or host failures.

Option A is true. FusionCompute allows administrators to configure different HA policies for virtual machines based on service importance, which aligns with SLA-driven management principles.

Option C is true. One of the core functions of HA is automatic VM reboot when a virtual machine encounters a failure, ensuring rapid service recovery.

Option D is true. FusionCompute continuously monitors CNA host and VM status. If a CNA host fails, HA triggers VM migration or restart on another available host in the same cluster, minimizing downtime.

Option B is false. Even when using shared storage, HA does not guarantee zero data loss. HA focuses on service recovery, not data consistency. Data loss may still occur depending on application state and write operations at the time of failure.

Therefore, the false statement is B.

Question: 23

FusionCompute adopts hardware-assisted virtualization technology to reduce memory virtualization overhead.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei FusionCompute leverages hardware-assisted virtualization technologies such as Intel VT-x, EPT, and AMD-V, which significantly reduce the overhead associated with CPU and memory virtualization.

Memory virtualization traditionally involved complex software-based address translation, which introduced performance overhead. With hardware-assisted features such as Extended Page Tables (EPT), memory mapping is handled directly by the CPU, improving efficiency and performance.

This approach is explicitly highlighted in HCIA–Cloud Computing materials as a key reason why FusionCompute delivers near-native performance for virtual machines.

Therefore, the statement is TRUE.

Question: 24

In Huawei FusionCompute, a virtual machine running Linux is created. In this scenario, which of the following does the host OS run in?

A. Linux VM

B. VRM

C. VNA

D. CNA

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionCompute architecture, the host OS refers to the operating system running on the physical compute node.

The Compute Node Agent (CNA) is the component deployed on the physical server. It integrates the host operating system and hypervisor, directly managing hardware resources and providing

virtualization services.

Linux VM (A) refers to the guest OS running inside a virtual machine.

VRM (B) is the centralized management component.

VNA (C) is responsible for virtual networking services.

Therefore, the host OS runs within the CNA, making D the correct answer.

Question: 25

In FusionCompute, the security auditor in role-based access control (RBAC) mode is only permitted to view and export logs.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei FusionCompute implements role-based access control (RBAC) to ensure secure and compliant system management.

The security auditor role is designed specifically for compliance and auditing purposes. According to HCIA–Cloud Computing materials, this role has read-only permissions, allowing the user to view and export system logs but not modify configurations or perform operational tasks.

This restriction helps ensure the integrity of audit records and prevents unauthorized system changes.

Therefore, the statement is TRUE.

Question: 26

Two virtual machines run on the same host and are in the same network segment, but they cannot communicate with each other. Which of the following is not a possible cause?

A. VM firewalls deny access to each other.

B. The NICs of the two VMs are added to the same port group.

C. The VMs are isolated by security groups.

D. The VM NICs are faulty.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionCompute virtual networking, VMs in the same network segment and same port group should be able to communicate normally.

Option A is a possible cause. Host-based or VM-based firewalls may block traffic between virtual machines.

Option C is also a possible cause. Security groups can apply access control rules that isolate VMs, even within the same subnet.

Option D is a possible cause. Faulty virtual NICs or driver issues can prevent network communication.

Option B is not a possible cause. If both VM NICs are connected to the same port group, they should be able to communicate, assuming no other restrictions exist.

Therefore, the correct answer is B.

Question: 27

In Huawei FusionCompute, which of the following functions enables virtual machines (VMs) to be migrated between CNA hosts with different CPU models?

A. Cluster I/O ring adaptation

B. Cluster HANA optimization

C. Cluster Guest NUMA policy

D. Cluster IMC policy

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionCompute, VM live migration between hosts with different CPU models requires CPU instruction compatibility. Huawei provides the IMC (Instruction Masking Compatibility) policy at the cluster level to address this issue.

The Cluster IMC policy masks advanced CPU instruction sets so that VMs only see a common subset of CPU features supported by all hosts in the cluster. This ensures that a VM can migrate safely between CNA hosts with different CPU generations or models without encountering instruction incompatibility errors.

Other options are unrelated:

I/O ring adaptation focuses on I/O optimization.

HANA optimization is specific to SAP HANA workloads.

Guest NUMA policy optimizes memory access, not CPU compatibility.

Therefore, the correct answer is Cluster IMC policy.

Question: 28

In FusionCompute, which of the following is not a prerequisite for VM live migration?

A. The VM must be in the running state.

B. The uplink of the distributed virtual switch (DVS) where the VM NIC resides must be associated with both the source and target hosts.

C. The IMC policy of the cluster must be configured.

D. The datastore to which the VM disk belongs must be associated with both the source and target hosts.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

VM live migration in FusionCompute has several mandatory prerequisites.

A is required because live migration applies only to running VMs.

B is required to ensure uninterrupted network connectivity during migration.

D is required because both source and target hosts must access the same datastore to migrate VM disks seamlessly.

IMC policy (C) is not mandatory for live migration. It is required only when hosts have different CPU models. If hosts use compatible CPUs, live migration can proceed without configuring IMC.

Therefore, the correct answer is C.

Question: 29

TAP is a virtual network device independently maintained in Linux.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In Linux-based virtualization environments, a TAP device is a software-based virtual network interface that operates at Layer 2 of the OSI model.

TAP devices are independently maintained by the Linux kernel and are commonly used to connect virtual machines to virtual switches or bridges. They simulate physical Ethernet interfaces, allowing VMs to send and receive Ethernet frames.

This definition aligns with HCIA–Cloud Computing networking fundamentals, so the statement is TRUE.

Question: 30

FusionCompute supports the following memory overcommitment technologies: memory sharing, memory swapping, and memory ballooning. After host memory overcommitment is enabled in cluster resource control, one of the preceding memory overcommitment technologies needs to be selected.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionCompute supports memory overcommitment to improve overall memory utilization in virtualized environments. The supported technologies include:

Memory sharing

Memory ballooning

Memory swapping

According to HCIA–Cloud Computing materials, after enabling memory overcommitment in cluster resource control, administrators must select at least one memory overcommitment mechanism to take effect.

Therefore, the statement correctly describes FusionCompute behavior and is TRUE.

Question: 31

Which of the following statements is true about the FusionCompute storage architecture?

A. Physical storage media that deliver storage space for virtualization are called storage devices.

B. LUNs allocated by Huawei Distributed Storage can be encapsulated as datastores.

C. After storage resources are converted to datastores and associated with hosts, virtual disks can be created for VMs.

D. FusionCompute uniformly encapsulates storage units of storage resources into datastores.

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionCompute uses a unified storage abstraction model. Regardless of whether the backend storage is SAN, NAS, local disks, or distributed storage, FusionCompute encapsulates storage resources into datastores.

This abstraction simplifies management and allows virtual machines to consume storage in a standardized manner.

While statements A, B, and C describe correct individual concepts, D best represents the core architectural principle emphasized in HCIA documentation.

Therefore, the correct answer is D.

Question: 32

Which of the following statements about Huawei FusionAccess is false?

A. FusionAccess allows users to access personal desktops through thin clients (TCs), software clients (SCs), and Android/iOS mobile terminals, supporting hybrid office scenarios.

B. FusionAccess uses the open-source Huawei Desktop Protocol (HDP) and supports 4K HD display.

C. FusionAccess (cloud desktop) is a mature cloud computing application.

D. FusionAccess supports the x86 + Kunpeng dual-stack architecture.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei FusionAccess is a mature cloud desktop (VDI) solution with broad access capabilities and architecture support.

Options A, C, and D are all correct according to HCIA–Cloud Computing materials.

Option B is false. Although FusionAccess uses HDP (Huawei Desktop Protocol) and supports high- resolution displays (including 4K), HDP is not open source. It is a Huawei proprietary protocol optimized for cloud desktop scenarios.

Therefore, the false statement is B.

Question: 33

The application scenarios of FusionAccess include secure offices, public terminals, general offices, and GPU graphics. Which of the following desktop types is most suitable for new media operations?

A. Secure office cloud desktops

B. Personalized cloud desktops for general offices

C. Design-oriented GPU graphics cloud desktops

D. Standard public terminal cloud desktops

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

According to the HCIA–Cloud Computing FusionAccess module, different desktop types are designed for different business workloads.

New media operations typically involve tasks such as video editing, image processing, animation rendering, and real-time previewing. These workloads are graphics-intensive and require high- performance GPU acceleration.

Design-oriented GPU graphics cloud desktops are specifically designed to support:

GPU pass-through or vGPU acceleration

High-resolution displays

Graphics rendering and multimedia processing

Professional applications used in media, design, and creative industries

Secure office desktops focus on data protection rather than performance. General office desktops target routine tasks such as documents and emails. Public terminal desktops are designed for shared access and lightweight workloads.

Therefore, the most suitable desktop type for new media operations is Design-oriented GPU graphics cloud desktops.

Question: 34

In Huawei FusionAccess, Huawei Desktop Protocol (HDP) classifies displayed bitmaps. It uses a lossless compression algorithm for text and a lossy compression algorithm for nonsensitive data, saving bandwidth without compromising user experience.

A. TRUE

B. FALSE

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Huawei Desktop Protocol (HDP) is a proprietary protocol optimized for cloud desktop scenarios and is a key topic

in the FusionAccess learning scope.

HDP uses intelligent bitmap classification:

Text and fine details use lossless compression to ensure clarity and readability.

Images, videos, and nonsensitive graphics use lossy compression, which significantly reduces bandwidth consumption while maintaining acceptable visual quality.

This adaptive compression strategy allows FusionAccess to deliver smooth user experience even in limited bandwidth environments, which is a highlighted advantage in HCIA materials.

Therefore, the statement is TRUE.

Question: 35

Which of the following FusionAccess components does not have backup data?

A. License

B. WI

C. VLB

D. vAG

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionAccess consists of multiple components, each with different roles and data persistence requirements.

WI (Web Interface) stores configuration and user access information and supports backup.

VLB (Virtual Load Balancer) maintains configuration data that can be backed up.

vAG (Virtual Access Gateway) also stores configuration-related data that supports backup and recovery.

The License component, however, is not backed up as configuration or system data. Licensing information is typically re-imported or reactivated during recovery rather than restored from backup.

Therefore, the component that does not have backup data is License.

Question: 36

During deployment of the FusionAccess gateway and load balancer, the HA status is abnormal after vAG/vLB configuration. Which of the following is not a possible cause of this fault?

A. The administrator changed the "root" password for HDC.

B. The administrator changed the password for logging in to the ITA portal.

C. The administrator changed the "gandalf" password for vAG/vLB.

D. The administrator changed the "root" password for vAG/vLB.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionAccess deployment, vAG (Virtual Access Gateway) and vLB (Virtual Load Balancer) rely on internal account authentication and synchronization to maintain high availability (HA) status.

Changing the following passwords can directly affect HA synchronization:

root password for vAG/vLB (D)

gandalf account password for vAG/vLB (C)

root password for HDC (A), which may disrupt management communication

However, the ITA portal login password (B) is related only to user or administrator portal access and does not participate in HA synchronization mechanisms.

Therefore, changing the ITA portal login password is not a cause of abnormal HA status.

Question: 37

For a small enterprise with no more than 100 employees, using the FusionAccess gateway and load balancer deployment solution for office scenarios, which of the following plans is recommended?

A. Two CNA nodes, each deployed with two VMs that integrate the vAG and vLB

B. Two CNA nodes, each deployed with a vAG VM and a VM that integrates the vAG and vLB

C. Two CNA nodes, each deployed with a VM that integrates the vAG and vLB

D. Two CNA nodes, each deployed with a vAG VM and a vLB VM

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

For small-scale office scenarios (≤100 users), the HCIA–Cloud Computing FusionAccess deployment guidelines recommend a simplified architecture that balances availability and cost.

Deploying one integrated vAG + vLB VM on each CNA node provides:

Basic high availability

Reduced VM count

Simpler deployment and maintenance

Sufficient performance for small user规模

Separate deployment of vAG and vLB or multiple integrated VMs is more suitable for medium and large-scale environments.

Therefore, the recommended solution is two CNA nodes, each deployed with a VM that integrates the vAG and vLB.

Question: 38

How many pairs of vLBs does a Huawei FusionAccess site generally need?

A. 3

B. 2

C. 4

D. 1

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In standard FusionAccess architecture, a site generally requires one pair of vLBs (active/standby) to provide load balancing and high availability for user access traffic.

Additional vLB pairs are only required in large-scale or multi-site deployments. For most enterprise scenarios described in HCIA materials, one pair is sufficient.

Therefore, the correct answer is 1 pair.

Question: 39

After installing infrastructure VMs, you must configure the virtualized environment and vAG/vLB so that FusionAccess can work properly. Which of the following statements about FusionAccess initialization are true?

A. After interconnection with a Windows AD domain is configured for authentication and authorization and the local domain is not enabled, the priority of the default LiteAS authentication and authorization is lower than that of the interconnected Windows AD domain.

B. When the vAG and vLB are used as the gateway and load balancer, related parameters of a WI cluster must

be configured on FusionAccess.

C. If the CNA node or storage device where the backup server resides is faulty, backup data may be lost. Therefore, a third-party FTP backup server is preferred on the live network.

D. When configuring the License on FusionAccess and the License supports both per-desktop and per-user modes, you must select User Count or Concurrent User Count.

Answer: A, B, C, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

All four statements accurately describe key initialization and configuration principles of FusionAccess:

A is true: When Windows AD is interconnected and local authentication is disabled, AD authentication has higher priority than LiteAS.

B is true: vAG/vLB deployment requires proper WI cluster configuration to enable access services.

C is true: To avoid backup data loss caused by infrastructure failures, external FTP backup servers are recommended in production networks.

D is true: When a License supports multiple charging modes, administrators must explicitly choose User Count or Concurrent User Count.

Thus, all options are correct.

Question: 40

Which of the following statements about FusionAccess for office automation (OA) is false?

A. FusionAccess can be interconnected with existing enterprise IT systems to use existing IT applications.

B. FusionAccess has three zones (red, yellow, and green). These zones are isolated from each other and have their own FusionAccess systems. Data can be transmitted between them.

C. Zone-based security is high, easy to deploy, and meets the security management requirements of most

enterprises.

D. Enterprise Active Directory (AD) is an important component of FusionAccess, and users must be authenticated by AD before logging in to virtual desktops.

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionAccess OA security architecture adopts zone-based isolation (red, yellow, green zones).

Each zone is logically and physically isolated

Each zone has its own FusionAccess system

Data is not allowed to flow freely between zones

Option B is false because it incorrectly states that data can be transmitted between zones. The other statements correctly reflect FusionAccess OA security and integration characteristics.

Question: 41

A bare VM is required when creating a virtual desktop template on FusionAccess. Which of the following statements are true about creating a bare VM?

A. The Configuration Mode of the system disk must be set to Common.

B. NICs cannot be added when creating a linked-clone bare VM, but NICs can be added when the linked-clone template is used for quick provisioning.

C. The NIC can be the service-plane distributed virtual switch (DVS) and port group of the user VM.

D. The Configuration Mode of the system disk must be set to Thin provisioning.

Answer: A, B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Bare VM creation rules in FusionAccess are strict to ensure template consistency and successful provisioning:

A is true: The system disk must use Common configuration mode, not thin provisioning.

B is true: For linked-clone bare VMs, NICs are added after template creation, during desktop provisioning.

C is false: Bare VMs should use management-plane networking, not user service-plane DVS/port groups.

D is false: Thin provisioning is not allowed for the system disk of a bare VM.

Therefore, the correct answers are A and B.

Question: 42

On FusionAccess, modifying or deleting a full copy template will affect the virtual machines created using the template.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionAccess, full copy desktops are created by completely copying the template into an independent virtual machine. After the VM is created, it no longer depends on the template.

According to HCIA–Cloud Computing materials:

Modifying or deleting a full copy template does not affect already provisioned desktops.

Only linked clone desktops depend on the template and snapshot chain.

Therefore, the statement is FALSE.

Question: 43

During FusionAccess virtual desktop provisioning, when a VM is being created, which of the following components checks whether the specified VM group and desktop group exist?

A. HDC

B. GaussDB

C. HDA

D. HDP

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionAccess architecture:

HDC (Huawei Desktop Controller) is responsible for desktop lifecycle management, including provisioning logic, policy enforcement, and resource validation.

During VM creation, HDC verifies whether the specified VM group and desktop group exist and whether configurations are valid.

Other components:

GaussDB stores metadata.
HDA runs inside desktops.

HDP is a display protocol.

Thus, the correct answer is HDC.

Question: 44

Which of the following operations must be performed on FusionAccess when you create a Windows template?

A. Mounting a template tool
B. Converting a VM into a template
C. Configuring a template
D. Encapsulating a template

Answer: A

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

When creating a Windows template in FusionAccess, the Template Tool must be mounted and installed inside the VM.

According to HCIA documentation:

The template tool installs necessary drivers, services, and agents required for desktop provisioning, management, and optimization.

Without mounting and installing the template tool, the VM cannot be converted into a usable desktop template.

Other operations such as configuration and conversion are subsequent steps, but mounting the template tool is mandatory.

Therefore, the correct answer is A.

Question: 45

Which of the following statements about VM group types and related operations during virtual desktop provisioning on FusionAccess is false?

A. If the desktop group type is Static Pool, linked clone and full copy VMs can be provisioned.

B. During quick provisioning of virtual desktops, administrators can assign desktops of different types to different users.

C. During virtual desktop creation, the system cannot add the virtual desktop to a specified VM group until virtual desktop assignment.

D. If the desktop group type is Dynamic Pool, linked clone and full copy VMs can be provisioned.

Answer: D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionAccess supports different desktop pool types, each with specific provisioning rules:

Static Pool supports both linked clone and full copy desktops → A is true.

Quick provisioning allows flexible desktop assignment strategies → B is true.

VM grouping is finalized during desktop assignment → C is true.

Dynamic Pool typically supports linked clone desktops only, because desktops are created and destroyed dynamically.

Therefore, statement D is false.

Question: 46

When an administrator uses a linked clone template to quickly provision virtual machines (VMs) on FusionAccess, which of the following are the steps of the quick provision process?

A. Create a VM.

B. Add the VM to a domain.

C. Rename the VM.

D. Associate the VM with a user group.

Answer: A, B, C, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Quick provisioning using a linked clone template is designed to rapidly deploy desktops while ensuring standardization and automation.

According to HCIA–Cloud Computing FusionAccess service provisioning procedures, the quick provision workflow includes:

Creating the VM based on the linked clone template and snapshot.

Renaming the VM to meet enterprise naming rules.

Adding the VM to the domain, enabling centralized authentication and policy control.

Associating the VM with a user or user group, which determines desktop assignment and access permissions.

All listed operations are integral steps in the linked clone quick provisioning process. Therefore, all options are correct.

Question: 47

On FusionAccess, which of the following operations is not required for VM creation during the quick provision process?

A. Select a VM template.

B. Select a VM group to which the VM belongs.

C. Select a desktop group to which the VM belongs.

D. Enter the number of VMs to be provisioned.

Answer: C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

During quick provisioning in FusionAccess:

Administrators must select a VM template.

A VM group must be specified for management purposes.

The number of VMs to be provisioned must be entered.

However, desktop group selection is not required at the VM creation stage. Desktop groups are associated later during desktop assignment, not during VM creation.

Therefore, the correct answer is C.

Question: 48

On FusionAccess, when creating a full copy, QuickPrep, or linked clone template, the template must be added to a domain.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In FusionAccess, templates must not be joined to a domain.

HCIA–Cloud Computing materials specify that:

Templates are created in a workgroup environment.

Domain joining is performed during desktop provisioning, not at the template stage.

Joining a template to a domain can cause conflicts such as duplicate computer accounts.

Therefore, the statement is FALSE.

Question: 49

Which of the following statements about service adjustment in FusionAccess are true?

A. A computer whose assignment type is Assign a Computer to Multiple Users or Assign Computers to a Desktop Group cannot be assigned again after being unassigned.

B. If a computer whose assignment type is Assign a Computer to a User is unassigned and then assigned again, the computer automatically starts. After the virtual desktop icon on the WI turns on, wait about three minutes and then log in.

C. If a computer whose assignment type is Assign a Computer to a User is unassigned and then assigned again, the computer can only be assigned to the original user and cannot be assigned to other users.

D. A computer whose assignment type is Assign a Computer to a User is automatically shut down after being unassigned.

Answer: B, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

FusionAccess defines clear behaviors for desktop service adjustment:

B is true. After reassignment, the system automatically starts the desktop, and a short wait time is required before user login.

D is true. When a desktop assigned to a single user is unassigned, the system automatically shuts it down to save resources.

A is false. Desktops assigned to multiple users or desktop groups can be reassigned after being unassigned.

C is false. After unassignment, a desktop can be assigned to a different user, not only the original one.

Thus, the correct answers are B and D.

Question: 50

Which of the following statements about the differences between a domain and an organizational unit (OU) are true?

A. Both OUs and domains can contain Active Directory (AD) objects.

B. Users can log in to a domain but not to an OU.

C. Group policies can be configured for both OUs and domains.

D. An OU can exist in a domain, and a domain can also exist in an OU.

Answer: A, B, C

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In Active Directory (AD), domains and organizational units (OUs) serve different purposes.

A is true. Both domains and OUs can contain AD objects such as users, computers, and groups.

B is true. Authentication is performed at the domain level. Users log in to a domain, not to an OU.

C is true. Group Policy Objects (GPOs) can be applied at both the domain level and the OU level.

D is false. An OU can exist inside a domain, but a domain cannot exist inside an OU. Domains are top-level logical boundaries.

Therefore, the correct answers are A, B, and C.

Question: 51

Which of the following statements about the differences between user groups and organizational units (OUs) are true?

A. OUs and user groups are Active Directory objects.

B. An OU can contain objects such as accounts, computers, printers, and shared folders.

C. Group policies can be configured for both OUs and user groups.

D. A user group can only contain accounts.

Answer: A, B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In AD, OUs and user groups are both used for organizing and managing resources, but they differ in function.

A is true. Both OUs and user groups are AD objects.

B is true. An OU can contain multiple types of AD objects, including users, computers, printers, and shared resources.

C is false. Group policies cannot be directly applied to user groups; they are applied to domains, sites, or OUs.

D is false. A user group can contain users, computers, and even other groups, not only user accounts.

Thus, the correct answers are A and B.

Question: 52

Cloud-native technologies enable organizations to build and run scalable applications in public, private, or hybrid cloud environments. Which of the following is not a representative cloud-native technology?

A. Service mesh

B. Virtualization

C. Microservice

D. Container

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

Cloud-native architecture emphasizes application-level design and management, focusing on agility, scalability, and resilience.

Microservices, containers, and service meshes are core cloud-native technologies.

Virtualization is a foundational cloud technology but belongs to traditional cloud infrastructure, not cloud-native application architecture.

Therefore, the correct answer is Virtualization.

Question: 53

In the OpenStack solution, Swift provides persistent block storage.

A. TRUE

B. FALSE

Answer: B

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In OpenStack:

Swift provides object storage, designed for storing unstructured data such as images, videos, and backups.

Cinder provides block storage, which is used by virtual machines as virtual disks.

Since Swift does not provide block storage, the statement is FALSE.

Question: 54

DRAG DROP

Match the following RAID levels with their respective descriptions.

RAIDO				Data on one disk is mirrored to another disk to provide complete redundancy backup.
RAID1				Parity data is distributed on ail disks in the array. Write operations can be performed on different disks at the same time.
RAID 5				Multiple disks are combined to form a large-capacity storage device. This RAID level supports data striping to improve the read and write speed, but does not provide data verification or redundancy backup.
RAID 10				Minor backup is performed for each disk, and then the disks are combined to form a storage device, providing high read and write performance in addition to redundancy backup.

Answer:

Explanation:

RAID 0

* Multiple disks are combined to form a large-capacity storage device. This RAID level supports data striping to improve read and write speed, but does not provide data verification or redundancy backup.

RAID 1

* Data on one disk is mirrored to another disk to provide complete redundancy backup.

RAID 5

* Parity data is distributed on all disks in the array. Write operations can be performed on different disks at the same time.

RAID 10

* Mirror backup is performed for each disk, and then the disks are combined to form a storage device, providing high read and write performance in addition to redundancy backup.

Question: 55

DRAG DROP

Match the following VLAN interface types with their descriptions.

Answer:

Explanation:

Access Interface

* The sent packets do not carry VLAN tags.

An access interface belongs to only one VLAN.

Frames are sent untagged, making it suitable for end devices such as PCs and printers.

Trunk Interface

* Only one VLAN tag can be removed from the sent packets.

A trunk interface allows multiple VLANs to pass.

Typically used for switch-to-switch connections.

Only the native VLAN is sent untagged; other VLANs are tagged.

Hybrid Interface

* Multiple VLAN tags can be removed from the sent packets.

A hybrid interface supports both tagged and untagged VLAN traffic.

Multiple VLANs can be configured to send traffic without VLAN tags.

Commonly used for server connections and virtualization environments.

Question: 56

Which of the following statements are false about the network shown in the figure?

A. The next-hop address of the packet sent from RTA to RTC is 20.1.1.2.
B. If RTA needs to communicate with RTC, a route needs to be configured on RTC.
C. RTB can communicate with RTA and RTC without any configuration modification.
D. If RTA needs to communicate with RTC, a route needs to be configured on RTB.

Answer: A, D

Explanation:

Comprehensive and Detailed Explanation (HCIA–Cloud Computing aligned):

In the shown topology, RTA and RTB are directly connected on 10.0.0.0/24 (RTA: 10.0.0.1/24, RTB: 10.0.0.2/24). RTB and RTC are directly connected on 20.1.1.0/24 (RTB: 20.1.1.2/24, RTC: 20.1.1.3/24). In HCIA networking fundamentals, a router forwards packets based on its routing table, and the next hop must be reachable on a directly connected network (or resolved via recursion), meaning it is typically the neighbor router interface on the local link.

A is false: For traffic from RTA to RTC (20.1.1.3), RTA's next hop should be RTB's interface 10.0.0.2 (its directly connected neighbor). The address 20.1.1.2 is not on RTA's directly connected subnet, so it cannot be used as the immediate next hop in this basic routing scenario.

B is true: RTC must have a route back to 10.0.0.0/24 (for return traffic to RTA). Since RTC only knows its directly connected network 20.1.1.0/24, a static/default route is required pointing to 20.1.1.2.

C is true: RTB has two directly connected networks (10.0.0.0/24 and 20.1.1.0/24), so it can reach both RTA and RTC with no additional routes.

D is false: RTB does not need extra routing configuration for RTA↔RTC communication because both networks are already directly connected on RTB.

Question: 57

Complete the following configuration file for a network interface card (NIC) in the openEuler operating system to set

a static IP address and enable the NIC to automatically start upon system boot.

TYPE=Ethernet

BOOTPROTO=

NAME=enp0s4

DEVICE=enp0s4

ONBOOT=

IPADDR=172.28.100.100

NETMASK=255.255.255.0

Answer:
BOOTPROTO=static,
ONBOOT=yes

Explanation:

In HCIA–Cloud Computing (Operating System Basics), Linux network interface configuration is a foundational skill, and openEuler follows the standard Red Hat–style network configuration syntax.

To configure a static IP address, the parameter BOOTPROTO must be set to static. This tells the operating system not to request an IP address from a DHCP server and instead use the manually configured values such as IPADDR and NETMASK.

To ensure that the NIC is automatically activated when the system boots, the parameter ONBOOT must be set to yes. If ONBOOT is set to no, the interface will remain down after boot and must be manually started.

Therefore, the completed and correct configuration is:

BOOTPROTO=static

ONBOOT=yes

BOOTPROTO=static → Uses a fixed IP address

ONBOOT=yes → Enables automatic startup at system boot

This configuration exactly matches HCIA–Cloud Computing requirements for static IP configuration on openEuler.

Question: 58

DRAG DROP

FusionAccess provides three protection mechanisms and 11 key measures to prevent data from being stored locally and make behaviors traceable and processes auditable, enhancing

end-to-end security. Match the key measures with the following protection mechanisms.

Answer:

Explanation:

Terminal Security

* Security hardening customized for TC systems

Focuses on endpoint protection

Prevents data from being stored locally on thin clients

Enhances terminal OS security and access control

Transmission Security

* Security gateway firewall

Protects data during network transmission

Controls and filters access traffic

Prevents unauthorized access and attacks during desktop access

Cloud Security

* Security watermark and secure deletion

Ensures data traceability through watermarks

Prevents data leakage

Guarantees secure deletion of cloud-side data

Question: 59

Fill in Blank

When a VM template is created on FusionAccess, select “Configure user login” if the group members want to log in to the VM.
(Enter the correct word on the GUI)

Answer: AD

Explanation:

In Huawei FusionAccess, user authentication and authorization for Windows desktops are typically integrated with Active Directory (AD). During VM template creation, the option “Configure user login” is used to specify who is allowed to log in to the virtual desktop.

According to the HCIA–Cloud Computing FusionAccess service provisioning process, this option is selected when AD group members are intended to access the VM. FusionAccess then uses AD-based authentication to control desktop login permissions, ensuring centralized identity management and consistent security policies.

Therefore, the correct word shown on the GUI and required to complete the sentence is AD.

Question: 60

DRAG DROP

On FusionAccess policy management, match the following scenarios with their operations.

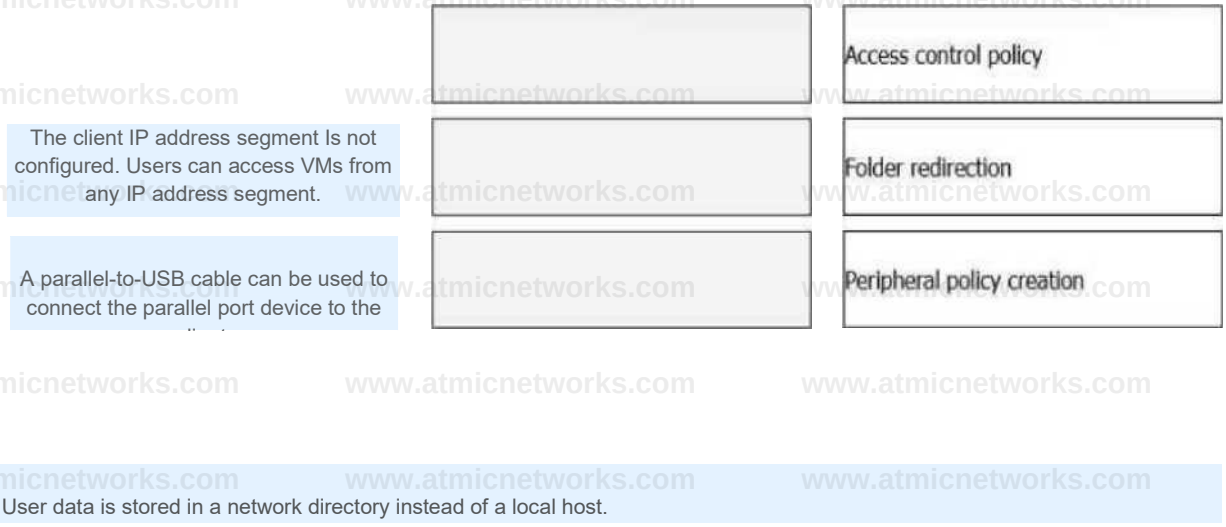

Answer:

Explanation:

User data is stored in a network directory instead of a local host

* Folder redirection

Redirects user folders (such as Desktop, Documents) to network storage

Prevents local data storage on virtual desktops

Enhances data security and centralized management

The client IP address segment is not configured. Users can access VMs from any IP address segment * Access control policy

Controls or relaxes access based on client IP segments

If not configured, access is allowed from any IP address

Used to restrict or permit VM access based on network location

A parallel-to-USB cable can be used to connect the parallel port device to the client

* Peripheral policy creation

Controls usage of peripherals such as USB, printers, scanners

Enables or disables device redirection

Supports special device scenarios via policy configuration