DUMPSBOSS.

HCIA-Cloud Service V3.0

Huawei H13-811_V3.0

Version Demo

Total Demo Questions: 46

Total Premium Questions: 467

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing Basics	12
Topic 2, Huawei Cloud Introduction	9
Topic 3, Huawei Cloud Computing Services	82
Topic 4, Huawei Cloud Storage Services	107
Topic 5, Huawei Cloud Networking Services	110
Topic 6, Huawei Cloud Security Services	48
Topic 7, Huawei Cloud Database Services	32
Topic 8, Huawei Cloud Application Services	43
Topic 9, Huawei Cloud AI Services	14
Topic 10, Huawei Cloud Big Data Services	10
Total	467

QUESTION NO: 1

After mounting the newly purchased cloud hard disk to the cloud server, it can be used directly.

A. True

B. False

ANSWER: B

Explanation:

False is correct. Attaching a newly purchased cloud hard disk to an ECS only makes the block storage device available to the server operating system; it does not automatically make it ready for application data. Before use, the operating system must recognize the new device, and the administrator normally needs to initialize it by creating a partition table and one or more partitions when appropriate. A file system, such as ext4 or XFS on Linux, must then be created on the usable partition or disk. Finally, the file system must be mounted to a directory so that applications and users can access it through the server's file system hierarchy. For persistent use after a reboot, the mount configuration should also be added to the operating system's startup configuration, typically by using the disk UUID in `/etc/fstab` on Linux. These preparation steps protect data usability and ensure that the attached EVS disk has a defined file-system structure before it stores files. Huawei Cloud's EVS guidance describes disk initialization as the required post-attachment procedure: Initializing Disks. Additional ECS disk-attachment guidance is available in the Huawei Cloud ECS User Guide.

QUESTION NO: 2

Description of the secret key, Which of the following is correct?

A. The key pair information used by the elastic cloud server can be queried through the console

B. It is not possible to create a key pair to log in to the Windows single-sex cloud server

C. You can change the key pair to log in to the elastic cloud server to log in with a password

D. You can use the key to authenticate when logging in with VNC

ANSWER: A B C

Explanation:

Key-pair authentication is an ECS login mechanism in which Huawei Cloud retains the public key and the user securely keeps the private key. The console provides key-pair management capabilities, including viewing the key-pair information associated with ECS resources, so administrators can identify the authentication method configured for their instances. This is especially useful when managing multiple Linux ECSs and controlling access through centrally managed keys.

Key pairs are intended for Linux ECS login authentication. Windows ECSs use password-based access mechanisms rather than SSH private-key authentication, so a key pair cannot be used as the login credential for a Windows ECS. For a Linux ECS originally created with key-pair authentication, Huawei Cloud supports setting or resetting an OS password, allowing password-based login to be used when operational requirements change. The applicable reset operation requires the instance and its guest-agent configuration to support password injection. Huawei Cloud documents key-pair use and lifecycle management in its ECS Key Pair documentation, and documents the password reset process in the ECS Resetting a Password guide.

QUESTION NO: 3

How many types of VPNs can be classified according to business purposes?

A. Lay2 and Lay3 tunnels VPN

B. Access VPN

C. Extranet VPN

D. Intranet VPN

ANSWER: B C D

Explanation:

When VPNs are classified by business purpose, there are three categories: Access VPN, Extranet VPN, and Intranet VPN. An Access VPN is used to provide secure remote access for individual users, such as employees connecting from a home office or while traveling, to enterprise resources. An Intranet VPN is intended for secure connectivity between sites or networks that belong to the same organization, enabling geographically distributed branches to communicate as though they are on one private network. An Extranet VPN securely extends selected enterprise resources to external parties, such as suppliers, customers, or business partners, while supporting controlled interorganization access. This business-oriented classification focuses on who is connecting and the organizational relationship between the connected parties. Huawei Cloud VPN provides encrypted connectivity for remote access and connections between enterprise networks, which aligns with these common business-use scenarios. See the Huawei Cloud VPN product page and the IETF's VPN terminology reference for background on VPN use and private network connectivity.

QUESTION NO: 4

The virtual private cloud supports the interconnection and intercommunication of VPCs and traditional data centers through dedicated lines/VPNs and other connection methods, which flexibly integrate resources.

A. True

B. False

ANSWER: A

Explanation:

True is correct. Huawei Cloud Virtual Private Cloud provides an isolated, logically defined network in which customers can create and manage cloud-side network resources such as subnets, route tables, security groups, and elastic network interfaces. It is also designed for hybrid-cloud connectivity rather than functioning only as an isolated cloud network.

A VPC can communicate with an on-premises or traditional data center by using services such as Direct Connect, which provides dedicated private connectivity, or Virtual Private Network, which establishes encrypted connectivity over the Internet. These connection methods enable enterprise workloads, applications, and data to be distributed between cloud resources and existing data-center environments while maintaining planned routing and access-control policies. Huawei Cloud also provides connectivity mechanisms for networking between cloud virtual private networks when an architecture requires it. This flexibility allows organizations to extend existing networks to the cloud, migrate services in stages, and integrate cloud and on-premises resources into a consistent hybrid environment. Huawei Cloud's VPC product overview describes VPC networking and hybrid connectivity capabilities; its Direct Connect documentation details dedicated connections between on-premises networks and Huawei Cloud. See Huawei Cloud Virtual Private Cloud and Huawei Cloud Direct Connect.

QUESTION NO: 5

API Which of the following calling methods are there?

A. GET

B. OPTIONS

C. ANY

D. POST

ANSWER: A B C D

Explanation:

API calling methods describe the HTTP request method that an API Gateway route accepts and forwards to its backend. **GET** is a standard HTTP method used to retrieve a representation of a resource. **POST** is a standard HTTP method used to submit data for processing or create subordinate resources. **OPTIONS** is also a defined HTTP method and is commonly used to obtain communication options; it is particularly relevant to browser cross-origin request handling (CORS). Huawei Cloud API Gateway supports these methods when defining an API.

ANY is a supported API Gateway route method selection rather than a separate HTTP method defined by the HTTP standard. When an API is configured with ANY, the gateway can match requests using any HTTP request method, which is useful when the backend service is designed to handle multiple methods through one API route. Therefore, GET, OPTIONS, ANY, and POST are all valid API calling-method selections in this context. Huawei Cloud documents API request-method configuration in its API Gateway user documentation, and the standardized semantics of GET, POST, and OPTIONS are defined by the IETF HTTP Semantics specification. See Huawei Cloud API Gateway User Guide and RFC 9110: HTTP Semantics.

QUESTION NO: 6

Which of the following container exceptions can the Container Security Service monitor?

A. High-risk system calls

B. Escape vulnerability attack

C. Escape file access

D. Container abnormal process

ANSWER: A B C D

Explanation:

Container Security Service provides runtime protection and monitoring for workloads running in containers. Its container runtime detection capability identifies suspicious behavior by observing process activity, system-call behavior, and file operations from within protected containers. Therefore, **High-risk system calls** are monitorable because dangerous syscall activity can indicate exploitation or an attempt to perform privileged operations. **Escape vulnerability attack** is monitorable because runtime protection detects behaviors associated with attempting to break isolation between a container and its host environment. **Escape file access** is also a relevant monitored exception: access attempts involving sensitive files outside the expected container scope can indicate an escape-related risk. Finally, **Container abnormal process** is monitored because unexpected or malicious processes are a primary indicator of runtime compromise, such as unauthorized command execution or malicious tooling launched in a container. Together, these event types reflect the service's purpose of detecting container runtime threats and producing security alerts so that administrators can investigate and respond promptly. Huawei documents container security as part of its Host Security Service capabilities; see the Huawei Cloud Host Security Service product page and the Huawei Cloud HSS documentation center.

QUESTION NO: 7

Which of the following are the functions of relational database services?

A. Create a database cluster

B. Expand the disk

C. Restore the backup to the specified time node

D. Add read replica

ANSWER: A B C D

Explanation:

Huawei Cloud Relational Database Service (RDS) provides managed lifecycle, capacity, availability, and data-protection capabilities for supported relational database engines. “Create a database cluster” represents the service’s ability to provision managed database resources, including high-availability primary/standby deployments where applicable. The platform handles core infrastructure provisioning and supplies a database endpoint for application access.

“Expand the disk” is a supported capacity-management function. As application data grows, RDS storage can be scaled to provide additional disk capacity without requiring users to build and manage new database servers themselves. “Restore the backup to the specified time node” reflects point-in-time restoration: RDS uses automated backups and transaction logs to restore data to a requested point within the available backup retention period. This is an essential managed data-recovery capability.

“Add read replica” is also a standard RDS scalability function. Read replicas asynchronously replicate data from a primary instance and can serve read-intensive workloads, helping separate read traffic from write processing. Huawei documents RDS as a managed service with instance provisioning, backup and restoration, scaling, and read-replica capabilities. See the Huawei Cloud RDS Product Description and Huawei Cloud RDS User Guide.

QUESTION NO: 8

Elastic scaling service is based on the user's business needs, A service that automatically adjusts its business resources through strategies.

A. True

B. False

ANSWER: A

Explanation:

“True” is correct because Elastic Scaling is designed to automatically adjust the amount of compute capacity available to an application in response to business demand and administrator-defined scaling policies. A user can create scaling configurations that specify the desired instance template and scaling policies that determine when capacity should increase or decrease. These policies can be scheduled, triggered by monitoring metrics such as CPU utilization, or invoked by alarms and other events. This automation helps maintain application availability during demand increases while reducing unnecessary resource consumption when demand falls.

In Huawei Cloud, Auto Scaling manages Elastic Cloud Server capacity in an Auto Scaling group according to configured rules and the group’s minimum and maximum instance limits. The business therefore retains control over the boundaries and strategy, while the platform performs the resource adjustments automatically. This directly matches the statement that elastic scaling adjusts business resources according to user business requirements through strategies. Huawei’s service overview describes Auto Scaling as automatically adjusting cloud resources based on policies, and its scaling-policy documentation explains the mechanisms used to trigger those adjustments. See Huawei Cloud Auto Scaling product description and Huawei Cloud Auto Scaling policy documentation.

QUESTION NO: 9

In which of the following situations, the network delay between the deployed elastic cloud servers must be the shortest?

A. In the same security group

B. In the same VPC

C. In the same district

D. Same zone

ANSWER: D

Explanation:

Same zone is correct because an availability zone is a physically isolated infrastructure area within a Huawei Cloud region, designed with independent power, cooling, and network-fault isolation. Elastic Cloud Servers deployed in the same availability zone are located in the closest cloud infrastructure scope available to one another. Their traffic therefore avoids the additional inter-zone network path required when instances reside in different zones. This makes same-zone deployment the appropriate choice when the primary requirement is the lowest possible network latency between ECS instances.

Huawei Cloud recommends selecting the same availability zone for workloads that require close communication and low network delay, such as application tiers with frequent internal calls, tightly coupled distributed components, or latency-sensitive services. A multi-zone architecture is valuable for high availability and resilience, but it deliberately places resources in separate fault-isolation domains and can introduce additional network latency compared with communication within one zone. The decision should therefore balance latency requirements against availability objectives. For the shortest delay specifically, place the related ECS instances in the same zone. See Huawei Cloud's ECS product description and its guidance for purchasing ECSs.

QUESTION NO: 10

Cannot transfer the elastic cloud server under one account to another account.

A. True

B. False

ANSWER: A

Explanation:

"True" is correct. An Elastic Cloud Server (ECS) is created as a resource of the Huawei Cloud account that purchases and manages it. Its ownership, billing relationship, IAM permissions, associated network resources, and lifecycle management remain bound to that account; Huawei Cloud does not provide a direct operation to transfer an existing ECS resource to a different account. This account-level isolation is fundamental to Huawei Cloud resource management and access control.

When a workload must be moved to a different account, the supported approach is to migrate or recreate it rather than change the ownership of the original ECS. For example, an administrator can create an image or back up the server data, share or copy the required image as permitted, and then create a new ECS in the destination account. Data and configuration must be migrated according to the relevant service capabilities and security requirements. The original ECS remains a resource of its original account until it is released.

Huawei Cloud's ECS documentation describes ECS as an account-managed compute resource, while IAM documentation explains that resources and permissions are governed within an account scope. See the Huawei Cloud ECS documentation and the Huawei Cloud IAM User Guide.

QUESTION NO: 11

When creating a scaling configuration, What are the options for configuration templates?

A. Use existing cloud server specifications as a template

B. Use an existing image as a template

C. Use shared templates

D. Use new template

ANSWER: A B D

Explanation:

When an Auto Scaling configuration is created in Huawei Cloud, the configuration-template selection supports three practical starting points: copying the specifications of an existing cloud server, basing the configuration on an existing image, or defining a new template manually. Using existing cloud server specifications captures the relevant instance settings from a

selected ECS so that newly scaled-out servers can be created consistently with that reference server. This is useful when a known-good ECS already has the intended flavor, disks, network configuration, and related settings.

Using an existing image is also a valid template approach because the image supplies the operating system and preinstalled application environment for instances added by the scaling group. This enables repeatable deployment of a prepared workload baseline. A new template is the manual alternative: the administrator specifies the cloud-server parameters directly while creating the scaling configuration, including the desired image and other instance attributes. Huawei Cloud's Auto Scaling documentation describes AS configurations as the source of the specifications used when the service automatically creates ECS instances. See the Huawei Cloud Auto Scaling configuration documentation and the Huawei Cloud Auto Scaling service overview.

QUESTION NO: 12

can access file systems across VPCs.

A. True

B. mistake

ANSWER: A

Explanation:

True is correct. Huawei Cloud Scalable File Service (SFS), including SFS Turbo where applicable, can be accessed by compute resources located in another VPC when network connectivity between the VPCs has been established. This cross-VPC access is not automatic merely because the file system exists; the client network must have a valid route to the file system's mount target. In practice, this can be provided through an approved VPC interconnection method, such as VPC peering or Cloud Connect/Enterprise Router, depending on the network architecture and regional design. The client must also be able to resolve and reach the mount address and satisfy the file system's access-control requirements, including the applicable security-group and mount permissions. Once these connectivity and access prerequisites are configured, instances in the connected VPC can mount and use the same shared file system. This enables workloads in separated VPC environments to use centralized shared storage without copying data into each VPC. Huawei Cloud describes SFS as a shared file-storage service for cloud servers and provides configuration guidance in its SFS product documentation and the SFS Help Center.

QUESTION NO: 13

IP The Sec framework contains several protocols, which are the important ones?

A. AH

B. IP

C. ESP

D. IKE

ANSWER: A C D

Explanation:

IPsec's principal protocols include AH, ESP, and IKE. AH (Authentication Header) supplies connectionless integrity, data-origin authentication, and optional anti-replay protection for IP packets. ESP (Encapsulating Security Payload) is the IPsec protocol most commonly used to protect VPN traffic because it can provide confidentiality through encryption as well as integrity, authentication, and anti-replay services. IKE (Internet Key Exchange) provides the control-plane functions needed to negotiate security associations, authenticate peers, select IPsec protection parameters, and derive or establish the keys used by AH and ESP. In practical Huawei VPN deployments, IKE negotiates the IPsec tunnel parameters and AH or ESP applies the negotiated packet protection; ESP is generally selected where encrypted communication is required. Together, these protocols represent the core IPsec architecture: packet-protection protocols plus a secure key-management and negotiation protocol. The IETF IPsec architecture defines AH and ESP as the security protocols used for IP packet

protection, while IKEv2 defines the exchange mechanism for establishing security associations and keying material. See RFC 4301: Security Architecture for IP and RFC 7296: Internet Key Exchange Protocol Version 2.

QUESTION NO: 14

What are the lifecycle management of elastic cloud servers?

A. Shut down

B. delete

C. Boot up

D. Unmount disk

ANSWER: A B C

Explanation:

Elastic Cloud Server lifecycle management covers the core actions that control an ECS instance's running state and its existence. **Boot up** starts an ECS that is stopped, allocating compute resources so that its operating system and applications can run. **Shut down** stops the instance and ends its active running state; this is a standard lifecycle operation used for maintenance, cost control, or planned service suspension. **delete** terminates the ECS as a managed resource, removing the server instance after the appropriate deletion choices are confirmed. These operations correspond to the normal operational flow of an ECS: create and start it, stop it when required, and delete it when it is no longer needed. Huawei Cloud documents ECS instance operations, including starting, stopping, and deleting servers, within the ECS User Guide. Lifecycle actions apply to the ECS instance itself, whereas storage attachment management is handled separately through disk operations. For operational details and prerequisites, consult the Huawei Cloud ECS User Guide and the Elastic Cloud Server product documentation.

QUESTION NO: 15

The object is the basic unit of data storage in OBS. Which of the following parts does it include?

A. AK/SK

B. Metadata

C. Data

D. Key

ANSWER: B C D

Explanation:

In Object Storage Service (OBS), an object is the fundamental storage entity held in a bucket. Each object includes a **key**, **data**, and **metadata**. The key is the object's unique name or identifier within its bucket. It can include prefixes that create a folder-like representation in OBS's flat object namespace, allowing objects to be organized and addressed through their bucket and key.

The data is the actual content stored in the object, such as a document, image, video, backup archive, or application file. Metadata is descriptive information associated with that stored content. OBS maintains system-defined metadata, including properties such as content length and content type, and supports user-defined metadata for application-specific identification and management needs. Together, these elements enable OBS to uniquely locate an object, retain its content, and manage properties needed for storage, retrieval, and HTTP-based access.

Huawei's OBS documentation describes objects as consisting of a key, data, and metadata. See the Huawei Cloud OBS object overview and the Huawei Cloud OBS metadata documentation.

QUESTION NO: 16

API What is included in the gateway's multiple security protection?

A. Threat protection

B. Access security

C. Log audit

D. Authentication service extension

ANSWER: A B C D

Explanation:

API Gateway’s multiple-security-protection capability spans the full API exposure lifecycle. **Threat protection** protects backend services by applying controls such as request throttling and safeguards against malicious or abnormal traffic patterns. **Access security** controls which clients and network sources can invoke an API, including restrictions based on client identity and source access policies. **Log audit** records API invocation and operational activity, providing traceability that supports security investigations, compliance review, and ongoing monitoring. **Authentication service extension** is also part of this protection model because API Gateway supports authentication and authorization mechanisms and can integrate authentication processing with external or customized services where required. Together, these controls help ensure that APIs are exposed only to authorized callers, protected while processing requests, and auditable after requests occur. Huawei Cloud presents API Gateway as a managed API publishing and management service with integrated security controls, while its API Gateway documentation describes the available authentication, access-control, throttling, and monitoring capabilities. See the Huawei Cloud API Gateway product page and the API Gateway API reference.

QUESTION NO: 17

Cloud service operation log can be queried in which service?

A. Security Expert Service

B. Cloud Audit Service

C. Cloud Monitoring Service

D. Cloud Analysis Service

ANSWER: B

Explanation:

Cloud Audit Service is the correct service for querying cloud service operation logs. It centrally records and manages audit trails generated by operations on cloud resources, such as creating, modifying, deleting, or configuring resources. These records provide essential traceability by identifying details such as the time of an operation, the user or agency that initiated it, the affected resource, the request parameters, and the operation result. This makes Cloud Audit Service a core capability for security auditing, operational troubleshooting, compliance evidence collection, and investigating unexpected changes in a cloud environment.

In Huawei Cloud, Cloud Audit Service collects operation records from supported cloud services and allows users to search, filter, and review those records through the console or APIs. Audit trails can also be delivered to Object Storage Service for long-term retention and further analysis, helping organizations retain an auditable history of management activity. Huawei’s documentation describes Cloud Audit Service as the service used to record users’ operations on cloud resources and query those traces. See the Cloud Audit Service overview and the Cloud Trace Service trace query documentation.

QUESTION NO: 18

BS Browser supports batch upload function.

A. True

B. mistake

ANSWER: A

Explanation:

“True” is correct. OBS Browser is Huawei Cloud Object Storage Service’s browser-based graphical tool for managing buckets and objects. A key capability of this interface is uploading multiple local files to an OBS bucket in one operation, commonly described as batch upload. The user selects the destination bucket and folder, chooses files from the local computer, and OBS Browser creates upload tasks for the selected files. This makes the tool appropriate for routine object-management work where users need to transfer a collection of files without using APIs, command-line tools, or writing code.

Batch upload is particularly useful because OBS is an object-storage service designed to store large numbers of independent objects. The browser tool provides a visual workflow for selecting files, monitoring task status, and handling the upload process from a web browser. Huawei’s OBS documentation identifies OBS Browser as a graphical client and documents its object upload operations and task-based management behavior. For product and tool details, see the Huawei Cloud OBS documentation and the OBS management console.

QUESTION NO: 19

What are the ways to connect to the public network (Internet)?

A. Elastic public network ip

B. Cloud monitoring

C. NAT gateway

D. Elastic load balancing

ANSWER: A C D

Explanation:

Huawei Cloud provides several public-network connectivity patterns depending on whether the requirement is direct inbound/outbound access, shared outbound access, or publishing an application service. An Elastic Public IP (EIP) supplies a publicly routable address that can be associated with an eligible cloud resource, enabling Internet communication. A NAT gateway provides controlled, scalable Internet egress for resources in a VPC that use private addresses; through its SNAT capability, multiple instances can access the Internet without each requiring its own public address. Elastic Load Balance can be configured as an Internet-facing load balancer by assigning a public IP address. It accepts client traffic from the Internet and distributes that traffic to backend servers, allowing an application to be exposed publicly while backend instances can remain private. These services therefore address the principal Internet-access and Internet-service-publication models in Huawei Cloud. Huawei documents EIPs as public IP resources that can be bound to cloud resources and describes NAT gateways as providing Internet connectivity for VPC resources. It also documents public Elastic Load Balance instances for receiving requests from Internet clients. See Huawei Cloud Elastic IP documentation and Huawei Cloud Elastic Load Balance documentation.

QUESTION NO: 20

does not need to delete the virtual gateway and virtual interface bound to the physical leased line when unsubscribing from the physical leased line.

A. True

B. mistake

ANSWER: B

Explanation:

mistake is correct because Huawei Cloud Direct Connect resources have dependency relationships that must be cleared before a physical leased line connection can be unsubscribed from or deleted. A virtual interface is the logical connection that carries traffic between the customer's on-premises network and a virtual gateway in the virtual private cloud. Since these logical resources are bound to the physical connection, they must first be removed to release that dependency.

This cleanup sequence helps ensure that the Direct Connect configuration is consistently removed and that no active logical connectivity remains associated with a physical leased line that is being terminated. In operational terms, administrators should identify the virtual interfaces associated with the physical connection, delete those interfaces, and remove the related virtual gateway configuration as required before proceeding with the physical leased-line unsubscription. Therefore, the claim that deleting the bound virtual gateway and virtual interface is unnecessary is a mistake.

See Huawei Cloud's Direct Connect product documentation and the Direct Connect User Guide for service concepts and lifecycle operations.

QUESTION NO: 21

VBS Supports simultaneous backup of all cloud hard drives in the server.

A. True

B. mistake

ANSWER: A

Explanation:

True is correct. Huawei Cloud Volume Backup Service (VBS) supports server-level backup, which protects all cloud disks attached to a server as a consistent group rather than requiring each disk to be backed up independently. When a server backup is created, VBS coordinates backup of the system disk and its associated data disks, helping preserve the relationship between operating-system files, applications, and data that may span multiple disks. This is particularly important for workloads that write to several volumes, because recovering only one volume could leave the workload in an inconsistent state. The service uses cloud backup capabilities and incremental backup mechanisms to provide protection and enable restoration when needed. In practical administration, the server-level backup capability simplifies policy management: a backup policy can be associated with the server and applied to its disks together. Huawei Cloud documentation describes VBS/server backup as protecting ECS server disks collectively and explains the relevant backup and restoration workflow. See the Huawei Cloud VBS User Guide and the Huawei Cloud Volume Backup Service overview for service capabilities and usage details.

QUESTION NO: 22

Which of the following descriptions about an OBS bucket are correct?

A. A bucket is a container for storing objects

B. A bucket is created in a specified region

C. Object names must be unique within the same bucket

D. A bucket can be attached to an ECS as a block disk

ANSWER: A B C

Explanation:

An OBS bucket is a container for storing objects. Buckets are created in a specified region, and object names must be unique within the same bucket. Bucket-level configurations can be used to manage access control and other storage behaviors for the objects in the bucket.

A bucket is not an individual object and cannot be attached to an ECS as a block disk. See the Huawei Cloud OBS Product Description.

QUESTION NO: 23

management console can be used without logging in

A. True

B. False

ANSWER: B

Explanation:

The statement is false: access to Huawei Cloud management-console functions requires a user to authenticate first. The console is the portal through which users manage cloud resources, view service information, configure permissions, create resources, and perform operational tasks. Huawei Cloud's documentation describes logging in with a Huawei Cloud account, an IAM user, or another supported identity method as the starting point for using the console. Authentication establishes the identity under which requests are made and allows Huawei Cloud IAM to apply the appropriate permissions, account boundaries, and security controls.

Unauthenticated visitors may be able to reach public Huawei Cloud webpages, documentation, pricing pages, or the sign-in screen itself, but that is not equivalent to using the management console to access or manage tenant resources. Consequently, the appropriate response is "False." This reflects the fundamental cloud-security practice that management-plane access is authenticated and authorized before resource operations are permitted. Huawei's login guidance is available in the Huawei Cloud Account User Guide, and the role of identity and permission control is described in the IAM Product Overview.

QUESTION NO: 24

Call AP IWhich of the following security authentication methods are there?

A. Huawei IAM authentication

B. No authentication

C. APP authentication

D. Micro authentication

ANSWER: A B C

Explanation:

Huawei Cloud API Gateway supports several API security authentication modes, including Huawei IAM authentication, APP authentication, and no authentication. Huawei IAM authentication is used when an API caller must be identified and authorized through Huawei Cloud IAM. The caller obtains IAM credentials, such as a token or an access key and secret key, and API Gateway can apply the associated IAM permissions to control access. This is suitable for secure calls between Huawei Cloud users, services, or applications that use IAM-based authorization.

APP authentication is intended for application clients that invoke published APIs. An API provider creates an application and credentials for it, allowing API consumers to authenticate as that application when accessing the API. This approach supports management of calling applications and their access to APIs. No authentication is also an available API Gateway mode when an API is deliberately made publicly callable and the backend does not require Gateway-level identity verification. The selected mode is configured per API according to its security and access requirements. Huawei documents these API Gateway authentication mechanisms in its API Gateway User Guide; IAM credential and authorization concepts are described in the IAM API documentation.

QUESTION NO: 25

What are the advantages of the shared service system architecture?

A. Services must be developed using the same programming language and tools

B. Allows to frequently publish different services at the same time,Maintain the availability and stability of other parts of the system

C. Low complexity

D. Each service can be independently developed by different teams, Do not affect each other, Accelerate time to market

ANSWER: B D

Explanation:

A shared-service system architecture, in the sense of independently deployable services, enables teams to deliver and operate individual services without requiring a coordinated release of the entire application. This makes it practical to publish different services frequently and in parallel while keeping unaffected parts of the system available and stable. Service boundaries also contain the impact of a deployment or change, supporting continuous delivery and more reliable operations.

It also supports independent development ownership. Different teams can design, build, test, deploy, and evolve the services for which they are responsible with limited dependency on other teams' delivery schedules. That autonomy reduces cross-team release coordination, permits work to proceed concurrently, and can significantly shorten the time required to bring new functions to market. Huawei Cloud's Cloud Service Engine documentation describes microservice-oriented service governance and independent service management capabilities, while Huawei's cloud-native guidance discusses building applications from independently managed services. See Huawei Cloud Service Engine User Guide and Huawei Cloud Container Engine.

QUESTION NO: 26

Which of the following are typical application scenarios of web application firewalls?

A. Anti-tampering

B. Key detection latest CVE Loopholes

C. 0Day vulnerability outbreak protection

D. E-commerce snapped up spike protection

ANSWER: A C D

Explanation:

Web application firewalls are typically deployed in front of web-facing applications to protect the HTTP/HTTPS layer from attacks and abusive traffic. **Anti-tampering** is a standard WAF scenario: WAF can protect designated web pages from unauthorized modification and help ensure that visitors receive trusted page content. This is especially useful for public-facing portals, brand sites, and pages containing important notices.

0Day vulnerability outbreak protection is also a core WAF use case. WAF protection rules and emergency defenses can be used as a virtual patching layer, filtering malicious requests that attempt to exploit newly disclosed or actively exploited web vulnerabilities before an application patch can be deployed. This reduces exposure during the period between vulnerability discovery and remediation.

E-commerce snapped up spike protection describes the protection required during flash-sales, limited-quantity promotions, and similar high-interest events. WAF can identify and mitigate malicious high-frequency requests, bots, and application-layer denial-of-service or CC-style traffic, helping preserve availability for legitimate shoppers. Huawei Cloud describes WAF capabilities including web attack protection, web tamper prevention, and protection against malicious traffic in its Web Application Firewall product overview and WAF service description.

QUESTION NO: 27

Shared path of the file system, What is the correct format of the NFS type?

A. \\File system domain name:/path

B. File system domain name:/path

C. \\File system domain name:,path

D. File system domain name:, path

ANSWER: B

Explanation:

The correct NFS shared-path format is **File system domain name:/path**. NFS clients identify a remote export using a server hostname or domain name followed by a colon and the exported directory path. For example, a mount target can be expressed as `example.sfs.myhuaweicloud.com:/share`. This syntax supplies the two pieces of information required by an NFS mount operation: the network endpoint of the file system and the specific exported path to mount.

In Huawei Cloud Scalable File Service documentation, NFS mounting uses the format `<mount address>:/<share name>`. The mount address is the file system's domain name or IP address, and the portion after the colon begins with a forward slash because it is the directory/export path on the NFS server. This shared path can then be passed to the Linux `mount` command along with a local mount directory. The format is independent of the local directory selected on the client; it solely describes the remote NFS resource. See Huawei Cloud's SFS mounting documentation and the mount.nfs manual for the standard server-and-export-path notation.

QUESTION NO: 28

The use of the cloud hard disk is the same as that of the transfer hard disk. The cloud hard disk mounted on the cloud server can be formatted and file system created to realize the persistent storage of data.

A. True

B. False

ANSWER: A

Explanation:

True is correct. Huawei Cloud Elastic Volume Service (EVS) provides block storage volumes that are used much like traditional physical hard disks after they are attached to an Elastic Cloud Server. The operating system on the cloud server detects the attached volume as a disk device. An administrator can then partition the disk if needed, format it with an appropriate file system, create a mount point, and mount it for application or user data.

This process enables persistent storage because the data resides on the EVS disk rather than only in the cloud server's temporary operating-system environment. The disk can retain its data independently of normal server operations, and it can be detached and attached according to supported EVS procedures. This makes EVS appropriate for data such as databases, application files, logs, and other workloads that require durable block storage. Huawei's EVS documentation describes EVS disks as scalable block storage devices that can be attached to cloud servers and managed similarly to physical disks. See the Huawei Cloud EVS product description and the EVS user guide for disk attachment and usage information.

QUESTION NO: 29

What service is CloudTable based on?

A. HDFS

B. OBS

C. MySQL

D. HBase

ANSWER: D

Explanation:

CloudTable is based on Apache HBase. Huawei Cloud describes CloudTable as a distributed, scalable, fully managed NoSQL database service that is compatible with Apache HBase. HBase provides the underlying wide-column data model and distributed storage architecture that allow CloudTable to support massive tables, high-throughput reads and writes, and low-latency access to large volumes of structured or semi-structured data. In practical use, applications can use standard HBase APIs and ecosystem tools to connect to CloudTable, reducing the effort required to migrate HBase workloads to Huawei Cloud. CloudTable also provides managed capabilities around the HBase foundation, such as cluster deployment, scaling, monitoring, and maintenance, so users do not need to operate the underlying distributed database infrastructure themselves. This HBase basis is particularly suitable for scenarios such as time-series data, IoT data collection, log storage, user profiles, and other workloads requiring horizontal scalability and random real-time access. Huawei Cloud's product documentation explicitly identifies CloudTable as an HBase-compatible NoSQL database service. See the Huawei CloudTable product page and the CloudTable product description.

QUESTION NO: 30

Which statement about Anti-DDos traffic cleaning is wrong?

A. Provides public IP protection against elastic cloud hosts and load balancing devices

B. Provide monitoring capability to view a single public network IP,Including current protection status, current protection configuration parameters, traffic conditions from 24 hours ago to the present, 2 hours of abnormal events

C. Possess Web vulnerability security detection capabilities

D. Provide configuration and modification for public network IP Anti-DDosAbility to correlate parameters

ANSWER: C

Explanation:

"Possess Web vulnerability security detection capabilities" is the wrong statement. Anti-DDoS traffic cleaning is intended to maintain availability during volumetric and protocol-based denial-of-service attacks. It monitors traffic to protected public IP addresses, detects attack traffic using traffic models and thresholds, and diverts or filters malicious traffic so that legitimate traffic can still reach the cloud resource. Its core protection objective is therefore attack detection and traffic scrubbing, rather than discovering weaknesses in web applications.

Web-vulnerability detection involves inspecting an application's HTTP/HTTPS behavior and identifying or blocking application-layer threats, such as SQL injection, cross-site scripting, malicious web shells, and other exploit attempts. Those functions belong to web application security products, particularly a Web Application Firewall, which applies web-specific rules and application-layer inspection. Huawei Cloud separately presents Anti-DDoS as a DDoS-mitigation service and WAF as a service for protecting web applications against web attacks. This functional separation is why web vulnerability security detection should not be attributed to Anti-DDoS traffic cleaning. See Huawei Cloud's Anti-DDoS service overview and Web Application Firewall overview.

QUESTION NO: 31

Description of cloud server group, Which of the following is correct?

A. There is no association between the cloud server group and the cloud server

B. Elastic cloud servers in the same cloud server group will try to be created on the same host

C. Elastic cloud servers in the same cloud server group will be created on different hosts as dispersed as possible

D. The elastic cloud servers in the cloud server group comply with the same policy

ANSWER: C D

Explanation:

An ECS cloud server group is a logical grouping used to apply a consistent placement policy to its member elastic cloud servers. With the anti-affinity policy commonly used for high availability, Huawei Cloud attempts to place ECSs in the same server group on separate physical hosts and distribute them as widely as possible. This reduces the likelihood that a single physical-host failure will affect all application instances at once. The wording “as dispersed as possible” is important: placement is attempted according to available infrastructure resources and scheduling conditions, rather than being an unconditional guarantee in every situation.

All ECSs associated with one cloud server group follow that group’s configured policy. Consequently, a server group provides a shared placement rule for its members, allowing an application with multiple ECS instances to be deployed in a manner that supports resilience and service continuity. This is especially useful for clustered applications, load-balanced service tiers, and redundant application nodes. Huawei Cloud documents ECS server groups and their policy-based host placement in its ECS user guide: Huawei Cloud ECS User Guide and Huawei Cloud ECS API Reference.

QUESTION NO: 32

After deactivating the scaling group, Which of the following operations or activities will still be performed?

A. Health check continues to check the health status of the instance

B. The scaling activity will still proceed after manually adjusting the expected number of instances

C. Stretching activities that have started will continue

D. Instances with abnormal health checks will be removed

ANSWER: A C D

Explanation:

When a Huawei Cloud Auto Scaling group is deactivated, Auto Scaling stops initiating new scaling activities that would normally adjust capacity in response to policies, alarms, schedules, or desired-capacity changes. Deactivation does not turn the group into an unmanaged collection of servers, however. Health checking remains active: Auto Scaling continues to monitor the health state of instances belonging to the group. If an instance is determined to be unhealthy, Auto Scaling removes that unhealthy instance according to the group’s health-check handling behavior. This preserves the group’s ability to identify and dispose of failed instances even while ordinary capacity scaling is suspended.

In addition, a scaling activity that was already started before the group was deactivated is allowed to finish. Deactivation prevents subsequent scaling operations from being started; it does not abruptly cancel an in-progress operation, which could otherwise leave resources in an inconsistent transitional state. These behaviors are consistent with Huawei Cloud Auto Scaling’s scaling-group lifecycle and health-check design. See the Huawei Cloud Auto Scaling User Guide and the Huawei Cloud Auto Scaling documentation portal.

QUESTION NO: 33

public cloud provides multiple types of elastic cloud servers, Including but not limited to what?

A. Bare Metal Server Type

B. GPU accelerated

C. Disk enhanced

D. General calculation type

ANSWER: B C D

Explanation:

Huawei Cloud Elastic Cloud Server (ECS) provides multiple instance families so workloads can be matched to the appropriate balance of vCPU performance, memory, storage I/O, and specialized acceleration. **General calculation type** is an ECS category intended for common enterprise applications, web services, development environments, and other workloads needing balanced compute and memory resources. **Disk enhanced** represents storage-optimized or disk-intensive ECS capacity, which is intended for scenarios that need substantial local storage capability and high disk I/O, such as big-data processing and data-intensive applications. **GPU accelerated** is also an ECS category, supplying GPU resources for workloads that benefit from massively parallel processing, including graphics rendering, video processing, scientific computing, and AI-related tasks. These categories demonstrate the on-demand nature of public-cloud ECS: customers select a flavor appropriate to the workload rather than deploying a single fixed server specification. Huawei's ECS documentation describes the available ECS flavor families and their workload-oriented characteristics, while the ECS product page summarizes its configurable compute, storage, network, and acceleration capabilities. See Huawei Cloud ECS instance type documentation and Huawei Cloud Elastic Cloud Server product page.

QUESTION NO: 34

Which of the following is not a restriction on the use of relational database service RDS?

A. The RDS read replica instance must be created in the same subnet as the master instance

B. Only allow access to applications that are in the same VPC as the instance

C. The RDS instance must be created in the VPC subnet

D. Users can self-create and manage various database engine instances in the RDS system

ANSWER: D

Explanation:

Users can self-create and manage various database engine instances in the RDS system is the correct statement because it describes a core capability of Huawei Cloud Relational Database Service rather than a usage restriction. RDS is a managed database service that lets customers provision database instances through the management console, APIs, or infrastructure automation tools. During creation, users select the supported database engine and version, compute flavor, storage configuration, availability option, VPC, subnet, and security group according to the choices available for the selected region and engine. After provisioning, customers retain operational control over service-level tasks such as instance configuration, account and database management, backups, monitoring, scaling where supported, and connection settings. Huawei Cloud operates and maintains the underlying database infrastructure and managed-service components, while the customer uses the self-service interface to create and administer their own RDS resources. This customer self-service provisioning model is a defining cloud-service characteristic, not a limitation. Huawei's RDS product documentation describes RDS as an online relational database service supporting creation and management of DB instances, and its user guide documents the instance-creation workflow: Huawei Cloud RDS Product Overview and Huawei Cloud RDS User Guide.

QUESTION NO: 35

IoT Internet of Vehicles solution, All used HUAWEI CLOUD EI Which services?

A. Map Reduce service MRS

B. Data Access Service DIS

C. Real-time streaming computing service Cloud Stream

D. Graph Computing Service GES

ANSWER: A B C D

Explanation:

An Internet of Vehicles solution needs a connected data pipeline that can ingest high-volume vehicle telemetry, process events with low latency, retain and analyze historical information, and discover relationships among vehicles, roads, drivers, and other entities. Data Access Service DIS supports the data-ingestion layer by collecting and transmitting continuous streams of data from connected vehicles and related devices. Real-time streaming computing service Cloud Stream supports real-time stream processing, enabling immediate analysis of events such as vehicle location updates, abnormal driving behavior, and traffic conditions.

Map Reduce service MRS provides the big-data processing and analytics capabilities required for large-scale, historical vehicle data, including batch analysis and data-mining workloads. Graph Computing Service GES is applicable where the solution needs to model and query complex relationships, such as associations between vehicles, routes, charging stations, users, and traffic networks. Together, these services provide the ingestion, streaming, big-data analytics, and graph-analysis capabilities commonly used in an IoT Internet of Vehicles architecture. Huawei Cloud describes DIS as a managed service for ingesting streaming data and MRS as a managed big-data analytics platform; see Huawei Cloud DIS and Huawei Cloud MRS.

QUESTION NO: 36

Which of the following components are required when an ECS without an EIP accesses the Internet through a public NAT gateway?

A. NAT gateway

B. SNAT rule

C. DNAT rule

D. VPN connection

ANSWER: A B

Explanation:

For an ECS with only a private IP address to access the Internet, a **NAT gateway** provides the translation function and an **SNAT rule** defines the private subnet and the public egress EIP used for source address translation. The VPC route table must also direct Internet-bound traffic to the NAT gateway.

An inbound DNAT rule is used when external users need to access a private resource through an EIP and port mapping. It is not required for outbound Internet access initiated by the ECS. See Huawei Cloud NAT Gateway User Guide.

QUESTION NO: 37

What are the advantages of using HUAWEI CLOUD E1?

A. Elastic resource scaling

B. Free operation and maintenance, users only need to focus on business

C. Have AI Basic platform that provides general AI capabilities

D. ELB

E. One-stop big data basic service

ANSWER: A B C E

Explanation:

HUAWEI CLOUD EI, commonly written as Enterprise Intelligence, is designed to make artificial-intelligence capabilities available as managed cloud services. **Elastic resource scaling** is an advantage because cloud-based intelligence

workloads can use computing and storage resources according to changing business demand, without requiring customers to pre-purchase and operate all underlying infrastructure. **Free operation and maintenance, users only need to focus on business** accurately reflects the managed-service model: Huawei Cloud operates the service platform and its underlying infrastructure so that customers can concentrate on applying intelligence to their own scenarios.

Have AI Basic platform that provides general AI capabilities is also correct because EI provides broadly applicable intelligence capabilities, including services and tools that enable organizations to incorporate AI into applications. In addition, **One-stop big data basic service** is an EI advantage because data processing and analysis are foundational to AI use cases; integrating these capabilities in the cloud reduces the effort needed to prepare, manage, and use data for intelligent applications. Huawei Cloud describes EI as a platform that provides AI capabilities and related cloud services. See the Huawei Cloud Enterprise Intelligence product page and the Enterprise Intelligence product documentation.

QUESTION NO: 38

GPU is suitable for application scenarios such as graphics rendering, engineering drawing and desktop cloud scenarios running in a virtualized environment?

A. True

B. mistake

ANSWER: A

Explanation:

True is correct. GPUs provide highly parallel processing capability and are designed to accelerate workloads that must process large volumes of graphical or image-related calculations simultaneously. This makes them well suited to graphics rendering and engineering design or drawing applications, where responsive display performance and efficient rendering of complex models are important.

In a virtualized desktop-cloud deployment, GPU resources can be assigned or shared with virtual machines so users of virtual desktops can run graphics-intensive applications with a user experience closer to that of a physical workstation. This is particularly useful for remote visualization, computer-aided design, three-dimensional modeling, and other professional graphics workloads. Huawei Cloud Elastic Cloud Server offerings include GPU-accelerated instance families intended for graphics acceleration and high-performance computing scenarios. The Huawei Cloud documentation describes GPU-accelerated instances as providing GPU computing capability for applicable workloads, while its desktop-cloud services support virtual desktop delivery for enterprise users.

Using GPU acceleration in these scenarios shifts suitable parallel graphics-processing tasks away from general-purpose CPU resources, helping provide the rendering throughput and display responsiveness required by virtualized graphical desktops. See Huawei Cloud Elastic Cloud Server and Huawei Cloud Workspace.

QUESTION NO: 39

The cloud dedicated line is a Layer 2 network connection.

A. True

B. mistake

ANSWER: A

Explanation:

The statement is true. A cloud dedicated line, provided through Huawei Cloud Direct Connect, establishes a dedicated physical network path between an enterprise’s on-premises network and Huawei Cloud. The dedicated connection operates at Layer 2 of the OSI model: it delivers Ethernet connectivity rather than, by itself, defining end-to-end IP routing. This Layer 2 foundation allows customers to create virtual interfaces over the dedicated connection and configure Layer 3 communication, such as BGP routing, for access to cloud virtual private clouds and associated services.

This distinction is important in cloud networking design. The dedicated line supplies predictable, private, high-bandwidth connectivity that avoids using the public Internet, while routing policies, IP addressing, and route exchange are configured through the virtual interface and the connected networks. Huawei Cloud documentation describes Direct Connect as connecting on-premises data centers to Huawei Cloud through dedicated connections and virtual interfaces, with virtual interfaces used to establish communication with VPCs. Therefore, characterizing the cloud dedicated line itself as a Layer 2 network connection is accurate. See the Huawei Cloud Direct Connect overview and the Huawei Cloud Direct Connect product page.

QUESTION NO: 40

Regarding the scaling configuration, which of the following is correct?

A. you An infinite number of scaling configurations can be created.

B. The scaling configuration defines the specification data of the cloud server to be added in the scaling group.

C. You can modify the created scaling configuration as needed.

D. After the scaling group is enabled, the scaling configuration cannot be changed.

ANSWER: B

Explanation:

The scaling configuration defines the specification data of the cloud server to be added in the scaling group. In Huawei Cloud Auto Scaling, a scaling configuration is the reusable template Auto Scaling uses whenever it needs to create new ECS instances for a scaling group. It specifies the essential instance deployment information, including the ECS flavor, image, disks, login method, VPC and subnet settings, security groups, elastic IP configuration where applicable, and related server parameters. This ensures that instances automatically added during scale-out activities are created consistently and meet the workload requirements established for the scaling group.

Creating this configuration separately from the scaling group is important operationally: administrators can prepare a standardized server template and then associate it with a scaling group so that scheduled, alarm-triggered, or manual scaling actions deploy the intended server specification. Huawei Cloud documents scaling configurations as the resource used to define the specifications for instances added by an Auto Scaling group. See the Huawei Cloud Auto Scaling User Guide: Creating a Scaling Configuration and the Huawei Cloud Auto Scaling API reference for scaling configurations.

QUESTION NO: 41

Cloud Monitoring, What alarm status does Cloud Monitoring currently support?

A. Warning

B. Back to normal

C. Insufficient data

D. All of the above

ANSWER: D

Explanation:

All of the above is correct. Huawei Cloud Monitoring Service (CES) evaluates configured alarm rules continuously or at their specified monitoring intervals and presents a state that reflects the evaluation result. A metric that reaches the configured threshold is reported as a warning/alarm condition. When subsequent metric values no longer meet the alarm condition, the alarm returns to a normal state, commonly presented as “Back to normal” in alarm notifications and history. CES can also show an insufficient-data state when it does not have enough valid metric datapoints to determine whether the configured threshold has been breached.

These states are important because they distinguish an active event requiring attention, a recovered condition, and a situation where monitoring data is not currently adequate for a conclusive evaluation. They also support meaningful alarm notifications and operational handling: teams can investigate warning events, recognize recovery notifications, and check metric collection or reporting when data is insufficient. Huawei's CES alarm API documents the underlying alarm-state values, including `alarm`, `ok`, and `insufficient_data`; the console wording may use user-oriented labels such as warning and back to normal. See the Huawei Cloud CES ShowAlarm API reference and the Huawei Cloud CES alarm documentation.

QUESTION NO: 42

Virtual Private Cloud (VirtualPrivateCloud, Abbreviated as VPC)It is an isolated and private virtual network environment that users apply for on HUAWEI CLOUD.

A. True

B. False

ANSWER: A

Explanation:

True is correct. Huawei Cloud Virtual Private Cloud (VPC) provides a logically isolated virtual networking environment in Huawei Cloud. A user creates a VPC and then defines its network range, subnets, route tables, security groups, and related connectivity resources according to the application's requirements. This lets cloud resources such as Elastic Cloud Servers operate in a private, user-defined network rather than directly sharing an unmanaged flat network with other tenants.

The isolation described in the statement is a core VPC characteristic. Each VPC is designed to provide network separation and administrative control, allowing users to plan IP address ranges and establish controlled communication among cloud resources. Connectivity to external networks, on-premises environments, or other VPCs can then be configured through appropriate Huawei Cloud networking services and rules. Therefore, describing VPC as an isolated, private virtual network environment that users apply for on Huawei Cloud accurately reflects its purpose and standard service definition. Huawei's VPC documentation describes the service as a logically isolated virtual network where users have control over network configuration. See the Huawei Cloud VPC product description and the VPC User Guide.

QUESTION NO: 43

Which of the following parameters need to be configured when purchasing a VPN connection?

A. VPN gateway

B. Remote gateway

C. This terminal network

D. Remote subnet

ANSWER: A B C D

Explanation:

When creating a Huawei Cloud VPN connection, the VPN gateway, remote gateway, local network, and remote subnet are all required configuration elements because together they define the two VPN endpoints and the protected traffic routes. The VPN gateway is the Huawei Cloud-side gateway through which the VPC accesses the IPsec VPN tunnel. The remote gateway identifies the peer device or peer gateway on the customer side, including the public IP address used to establish the tunnel. "This terminal network" represents the local subnet, meaning the VPC subnet or CIDR range whose traffic is to be protected by the VPN connection. The remote subnet specifies the on-premises or peer-network CIDR range reachable through that tunnel. These local and remote protected subnet definitions are essential for determining which traffic is encrypted and routed across the site-to-site VPN connection. Huawei Cloud documentation describes VPN connections as being created between a VPN gateway and a customer gateway, with local and remote subnets configured as protected data flows. See the Huawei Cloud Creating a VPN Connection guide and the VPN service overview.

QUESTION NO: 44

The traditional IT infrastructure architecture of enterprises is no longer capable of supporting the efficient operation of enterprises, Urgently calling for the arrival of structural changes, Which is not the main reason?

A. Slow business launch,Complex life cycle management

B. TCO remains high

C. Critical application performance is limited byIObottleneck

D. Internet speed is too slow

ANSWER: D

Explanation:

Internet speed is too slow is not a principal architectural reason that enterprises need to transform traditional IT infrastructure. Enterprise infrastructure transformation is driven primarily by internal operational and technical needs: accelerating service delivery, simplifying the management of increasingly complex systems, improving resource utilization, reducing lifecycle cost, and ensuring that applications receive scalable and reliable compute, storage, and network resources. These needs concern how enterprise IT is designed, provisioned, operated, and evolved.

Internet access speed can affect users' experience when they use externally hosted services, but it is generally an external connectivity condition rather than an inherent limitation of the traditional enterprise infrastructure architecture. It can be addressed through connectivity planning, bandwidth upgrades, dedicated connections, content delivery, or network optimization without necessarily requiring a structural redesign of the entire IT environment. Huawei Cloud describes cloud infrastructure as enabling flexible, on-demand resources and supporting enterprise digital transformation; these capabilities target the core operational pressures behind infrastructure modernization. See Huawei Cloud and Huawei Cloud ECS User Guide.

QUESTION NO: 45

cloud hard drives need to be used with elastic cloud servers or bare metal servers.

A. True

B. False

ANSWER: A

Explanation:

True is correct. In Huawei Cloud, Elastic Volume Service (EVS) disks, commonly called cloud hard drives or cloud disks, provide persistent block storage rather than an independently usable compute environment. To store, read, and write application data, an EVS disk is attached to a server that supplies the operating system, file system, and application access. Huawei Cloud supports attaching EVS disks to Elastic Cloud Servers (ECSs) and Bare Metal Servers (BMSs), allowing those compute instances to use the disk as an additional data disk or, where applicable, as a system disk.

This attachment-based model separates compute from storage. A disk can remain available independently of an individual server's lifecycle and can be detached and attached as operational needs require, subject to disk and attachment constraints. However, the disk's block devices are consumed through the attached ECS or BMS; it is not a standalone service for directly running workloads. This is why cloud hard drives need to be used with elastic cloud servers or bare metal servers. Huawei describes EVS as a block storage service for ECSs and BMSs in its EVS product description and documents the server attachment workflow in the EVS User Guide.

QUESTION NO: 46

What are the following restrictions on cloud server instances manually moved into a scaling group? (Multiple choice)

A. This instance is the same as the scaling group VPC

B. The instance has been bound to an elastic public IP

C. The scaling group AZ contains the AZ where the instance is located

D. This instance is not used by other scaling groups

ANSWER: A C D

Explanation:

When an existing cloud server is manually added to a Huawei Cloud Auto Scaling group, its networking and placement attributes must be compatible with the group. The instance must use the same VPC as the scaling group because a scaling group manages instances within its configured network environment. This ensures that the instance can use the group's subnets, security configuration, load-balancing integration, and other network-related resources consistently.

The instance's availability zone must also be included in the availability zones configured for the scaling group. Auto Scaling uses the group's enabled AZs to control where capacity is managed and distributed, so an instance outside that scope cannot be incorporated into the group. In addition, an instance can belong to only one scaling group at a time. This prevents conflicting lifecycle operations, including health checks, scaling actions, instance removal, and policy-driven configuration changes, from being applied by multiple groups. These requirements are part of the compatibility checks performed when an existing ECS is added to an Auto Scaling group. See Huawei Cloud's Auto Scaling user documentation and the Huawei Cloud Auto Scaling service overview.