

# ACP Cloud Computing Certification

Alibaba Cloud ACP-Cloud1

Version Demo

Total Demo Questions: 20

Total Premium Questions: 206

Buy Premium PDF

<https://passqueen.com>

[support@passqueen.com](mailto:support@passqueen.com)

## Topic Break Down

| Topic                                 | No. of Questions |
|---------------------------------------|------------------|
| Topic 1, Cloud Computing Basics       | 6                |
| Topic 2, Alibaba Cloud Products       | 127              |
| Topic 3, Cloud Computing Architecture | 21               |
| Topic 4, Cloud Computing Security     | 34               |
| Topic 5, Cloud Computing O&M          | 18               |
| Total                                 | 206              |

## QUESTION NO: 1

Which of the following statements about file uploading to OSS is NOT correct?

- A. All file uploading methods provided by Alibaba Cloud OSS service will support resuming from break point.
- B. If you are uploading some small files ( < 1MB), you can choose the PUT method to finish uploading in one HTTP request.
- C. If you are using OSS to store log files that are continuously being updated, it ' s better to choose the APPEND method instead of overwriting or generating new versions.
- D. If you need to upload very large files, it's better to use multi-part uploading to get the capability of resuming from a breakpoint.

**ANSWER: A**

### Explanation:

“All file uploading methods provided by Alibaba Cloud OSS service will support resuming from break point.” is the statement that is not correct. Breakpoint-resume capability depends on the upload mechanism and the client-side implementation; it is not an inherent feature of every OSS upload request. In OSS, multipart upload divides an object into independently uploaded parts. A client can persist the upload ID and the completed-part information, then resume an interrupted transfer by uploading only the parts that remain and completing the multipart upload after all parts are available. This makes multipart upload particularly appropriate where network interruptions are possible or where objects are large enough that retrying an entire upload would be inefficient. Alibaba Cloud documents multipart upload as a multi-step process consisting of initiating the upload, uploading parts, and completing the upload; the SDKs can use checkpoint data to implement resumable uploads. In contrast, a normal single-request upload is handled as one request and does not provide the same multipart checkpoint process. Therefore, OSS upload methods cannot accurately be described as universally supporting resumption from a breakpoint. See the [Alibaba Cloud OSS multipart upload documentation](#) and the [OSS resumable upload documentation](#).

## QUESTION NO: 2

An entrepreneur team builds a network disk application based on Alibaba Cloud Products. This application allows users to upload and download files online. Currently, all the files are stored on the data disks of ECS instances. As the service volume increases, the number of ECS instances and overall O&M costs soar. To reduce O&M costs, using the \_\_\_\_\_ product in conjunction with ECS instances provides an ideal solution.

- A. SLB
- B. OSS
- C. CDN
- D. ApsaraDB for RDS

**ANSWER: B**

### Explanation:

OSS is the appropriate product because a network disk application primarily needs durable, elastic storage for user-uploaded files. Alibaba Cloud Object Storage Service is designed to store massive quantities of unstructured objects, including documents, images, videos, and backups. The application can retain ECS instances for application logic, authentication, metadata processing, and upload/download request handling, while placing the actual file content in OSS buckets. This separation allows storage capacity to grow independently from compute capacity, so the team does not need to add or maintain ECS instances merely to obtain more data-disk space.

OSS also provides highly durable storage, scalable capacity without advance provisioning, lifecycle management for transitioning infrequently accessed objects to lower-cost storage classes, and access-control mechanisms such as RAM policies and signed URLs. These capabilities reduce operational work associated with disk expansion, file replication, and storage-server maintenance. For client transfers, the application can issue temporary signed URLs so users upload to or download from OSS directly, reducing load on the ECS application tier. Alibaba Cloud identifies OSS as an object storage

service for storing and accessing any amount of data, including unstructured data, with elastic scalability. See the [Alibaba Cloud OSS overview](#) and [OSS storage classes documentation](#).

### QUESTION NO: 3

A large enterprise wants to migrate the entire business system to Alibaba Cloud to save the overall IT procurement and O&M costs. From the security aspect, the company requires that

1. Must support secured remote O&M because the administrator often takes business trips.
  2. Networks between subsystems should be isolated because subsystems are independently used by different departments.
- Which of the followings should be used together to meet the company's requirements? (Number of correct answers: 3)
- A.** Configure a VPN connection to the VPC and use the bastion host for remote O&M. The administrator uses VPN-encrypted communication during O&M.
  - B.** Build an independent ECS instance as the bastion host for remote logon and O&M, and authorize the bastion host to access ECS instances running other subsystems.
  - C.** Use the security group function of the ECS instance, and respectively deploy ECS instances running different subsystems to independent security groups.
  - D.** Create multiple ECS instances in the VPC to install subsystems of different departments- Allocate only Intranet IP addresses to all ECS instances, and deploy them in the same security groups.

**ANSWER: A B C**

#### Explanation:

Secure remote O&M is achieved by establishing encrypted VPN connectivity between the administrator's remote environment and the cloud VPC, then using a controlled bastion host as the management entry point. This avoids exposing every workload instance directly to the public Internet and provides a centralized path for administrative access. Alibaba Cloud VPN Gateway supports encrypted connections between on-premises or remote networks and a VPC, making it suitable for protected remote administration. A dedicated bastion host further centralizes remote logon and O&M permissions: it can be the only authorized management source for the workload ECS instances, which supports controlled access and simplifies auditing of privileged operations.

Network separation between independently operated departmental subsystems should be enforced through distinct security groups. Security groups are stateful virtual firewalls for ECS instances; assigning each subsystem's ECS instances to its own security group allows inbound and outbound rules to be defined separately. Access between groups can then be explicitly permitted only where a real business dependency exists, maintaining least-privilege segmentation. Together, encrypted VPN access, a bastion-host management path, and separate security groups address both secure remote maintenance and isolation among departmental systems. See [Alibaba Cloud VPN Gateway documentation](#) and [Alibaba Cloud security group overview](#).

### QUESTION NO: 4

The Alibaba Cloud CDN can directly accelerate access to the files stored in OSS and reduce OSS traffic costs.

- A.** True
- B.** False

**ANSWER: A**

#### Explanation:

**True** is correct. Alibaba Cloud CDN can use an Object Storage Service (OSS) bucket as an origin server. After a domain name is added to CDN and configured with the OSS bucket as its origin, users request content through CDN edge nodes rather than retrieving every object directly from OSS. CDN caches eligible static content, such as images, videos, installation

packages, and web assets, at edge locations close to users. Cache hits are served directly by the edge node, which reduces latency and improves download and page-load performance.

This architecture also reduces traffic served directly from OSS because OSS is contacted primarily when content is not yet cached, has expired, or must be refreshed. Alibaba Cloud documents OSS as a supported CDN origin type and describes CDN as a service that caches origin content at globally distributed points of presence. This makes CDN an appropriate solution for accelerating public access to OSS-hosted files while reducing the OSS traffic associated with repeated client requests. CDN traffic and any applicable cache-miss or origin-fetch charges should still be evaluated under the relevant billing rules. See the [Alibaba Cloud CDN overview](#) and [documentation for accelerating OSS resources with CDN](#).

#### QUESTION NO: 5

Auto Scaling is a management service that can automatically adjust elastic computing resources based on your business needs and policies. To ensure ECS instances created by Auto Scaling can correctly respond to user requests without additional configuration, applications deployed on these ECS instances must be \_\_\_\_\_.

- A. Stateless and horizontally scalable
- B. Stateful
- C. Stateless and idempotent
- D. All data is locally stored on ECS instance disks

**ANSWER: A**

**Explanation:**

**Stateless and horizontally scalable** is correct because Auto Scaling must be able to add, remove, and replace ECS instances without requiring application-specific intervention. A stateless application keeps session state, user data, and other persistent information outside the individual compute node, such as in a database, distributed cache, or shared storage service. Therefore, a newly created ECS instance can begin serving requests using the same deployment configuration as the existing instances, while a removed instance does not take indispensable application state with it.

Horizontal scalability also enables the application tier to increase capacity by running additional equivalent instances behind a load-balancing endpoint. Auto Scaling can then respond to changing demand by adjusting the number of ECS instances, while requests continue to be distributed to healthy instances. This architecture supports the disposable-instance model used by scaling activities and ensures that replacement or expansion capacity can serve users immediately after initialization and health checking. Alibaba Cloud Auto Scaling is designed to automatically adjust ECS resources in response to demand and configured rules; application workloads should consequently be deployable as interchangeable instances. See the [Alibaba Cloud Auto Scaling overview](#) and [scaling group documentation](#).

#### QUESTION NO: 6

The backend server pool of an Alibaba Cloud SLB contains multiple ECS instances, which may have different service capacities. To exploit the different service capacities of backend ECS instances, which of the following statements is correct?

- A. Choose Weighted Round Robin mode to set higher weights to ECS instances with higher capacities. The higher the weight of the backend ECS instance, the higher the chance that the instance will receive requests.
- B. SLB cannot assign more requests to certain ECS instances.
- C. Choose Round-Robin model and ECS instances with higher capacities will surely be assigned with more requests.
- D. Choose Weighted Least-Connection mode, and ECS instances with higher capacities will surely be assigned with more requests.

**ANSWER: A**

**Explanation:**

Choose Weighted Round Robin mode to set higher weights to ECS instances with higher capacities. The higher the weight of the backend ECS instance, the higher chance that the instance will receive requests. This is the appropriate approach because a backend server's weight represents its relative share of traffic within the Server Load Balancer backend pool. After weights are configured, the load balancer distributes new requests in proportion to those values. For example, assigning a weight of 200 to a higher-capacity ECS instance and 100 to another instance gives the first instance approximately twice the request-distribution opportunity, assuming both are healthy and available. This lets an administrator align traffic allocation with practical capacity differences such as CPU and memory resources, application-worker limits, or measured throughput. Weights should be selected from capacity testing and monitored over time so that the traffic ratio continues to reflect actual application performance. Alibaba Cloud documents weighted round robin as a scheduling algorithm for distributing requests according to backend-server weights. See the [Classic Load Balancer scheduling-algorithm documentation](#) and the [backend-server configuration documentation](#).

#### QUESTION NO: 7

After creating Alibaba Cloud Elastic Compute Service (ECS) instances, Alibaba Cloud will provide free basic security services to help defend against attacks, including \_\_\_\_\_. (Choose four.)

- A. Unexpected remote login
- B. Key management service
- C. DDoS attacks
- D. Brute-force password cracking
- E. Sensitive information leakage

**ANSWER: A C D E**

#### Explanation:

Alibaba Cloud ECS instances can use free baseline protections supplied through Alibaba Cloud security services. **Unexpected remote login** is covered by Security Center's baseline detection capabilities, which identify abnormal or suspicious logon behavior and generate security alerts. **Brute-force password cracking** is also monitored as an attack behavior, helping administrators identify repeated password-guessing attempts against ECS instances. **Sensitive information leakage** is included in the intended baseline security scope through detection and alerting for security risks that can expose sensitive data or files. These capabilities give ECS users initial visibility into common host-security threats without requiring a separately purchased advanced edition.

**DDoS attacks** are covered by the default DDoS mitigation capability provided for eligible Alibaba Cloud public IP resources. DDoS Basic is automatically enabled and provides foundational mitigation against volumetric attacks, while higher protection requirements may require paid DDoS services. Together, host-threat detection and baseline DDoS mitigation are the free security protections described in the question. Alibaba Cloud documents the Security Center editions and baseline capabilities at [Security Center editions](#) and describes the default service at [DDoS Basic](#).

#### QUESTION NO: 8

Alibaba Cloud provides comprehensive security protection covering networks, hosts, applications, databases, and content for on-cloud clients. Users using Alibaba Cloud services only need to focus on their own services, without the need to pay any attention to security protection issues.

- A. True
- B. False

**ANSWER: B**

#### Explanation:

**False** is correct because cloud security operates under a shared responsibility model. Alibaba Cloud is responsible for securing the underlying cloud infrastructure, including facilities, physical hardware, and foundational cloud services. However, customers retain responsibility for the security of what they deploy and configure in the cloud. This includes managing identities and permissions through RAM, protecting account credentials, configuring network controls such as security groups, applying operating-system and application patches where applicable, encrypting and protecting data, and monitoring their own workloads for security events.

Alibaba Cloud offers a broad set of security capabilities, such as Security Center, Cloud Firewall, Anti-DDoS, WAF, and encryption services. These services can help customers build effective defenses, but they do not eliminate the customer's security obligations. The exact division of responsibility also varies according to the service model: managed services can reduce operational work, while ECS-based workloads require customers to take a more direct role in host and application security. Customers must therefore actively design, configure, operate, and continuously review security controls for their own services and data. Alibaba Cloud describes this allocation in its [Shared Responsibility Model](#) and provides customer-focused guidance in the [Security Center documentation](#).

#### QUESTION NO: 9

Which of the following statements about ECS snapshots are correct? (Number of correct answers: 3)

- A. A snapshot can be used to create a new disk.
- B. A snapshot can be used as a source for creating a custom image.
- C. Automatic snapshots can be created by using an automatic snapshot policy.
- D. A snapshot guarantees that all applications running on the instance are application-consistent without any application coordination.
- E. Snapshots can be used only for system disks and cannot be used for data disks.

#### ANSWER: A B C

##### Explanation:

An ECS snapshot is a point-in-time backup of a disk and can be used to create a new disk or roll back a disk when recovery is required. Snapshots are also commonly used as the source for creating custom images, which can then be used to create ECS instances with a standardized operating system and application environment.

Snapshots can be created manually or automatically by configuring an automatic snapshot policy. They are not a replacement for all application-consistent backup procedures, because an application may need to flush data or coordinate its own backup process before a snapshot is created. In addition, deleting an ECS instance does not automatically mean that snapshots are deleted; snapshots are managed as separate resources according to their retention and deletion configuration. See [ECS snapshots documentation](#).

#### QUESTION NO: 10

You are using Auto Scaling with one scaling group already created, then you want to execute a task at a specific time such as removing 1 ECS instance every night at 00:00. To achieve this, which of the following operations should be performed? (Number of correct answers; 2)

- A. Create a scaling rule
- B. Create a scheduled task.
- C. Create a new scaling group.
- D. Create an event-triggered task.
- E. Create a scaling rule that removes one ECS instance and a scheduled task that runs the rule every day at 00:00.

#### ANSWER: E

**Explanation:**

Create a scaling rule that removes one ECS instance and a scheduled task that runs it daily at 00:00 is correct because Alibaba Cloud Auto Scaling separates the scaling action from its trigger. The scaling rule defines the required capacity adjustment, such as reducing the number of instances by one. A scheduled task then specifies when Auto Scaling must execute that rule, including its recurrence and the daily 00:00 execution time. The task is associated with the existing scaling group through the scaling rule, so no additional scaling group is required. This arrangement provides predictable, time-based capacity reduction while allowing Auto Scaling to select an eligible instance according to the scaling group's configuration and instance-removal policy. When creating the scheduled task, configure its scheduled execution time, recurrence type, and the target scaling rule; ensure that the scaling group minimum-size setting permits the requested reduction. Alibaba Cloud documents scheduled tasks as the mechanism for triggering scaling rules at specified times, while scaling rules define the actual adjustment behavior. See the [Alibaba Cloud scheduled task documentation](#) and the [Alibaba Cloud scaling rule documentation](#).

**QUESTION NO: 11**

A company plans to migrate a self-managed MySQL database to ApsaraDB RDS for MySQL by using Data Transmission Service (DTS). Which of the following preparations are required or recommended before the migration task is started? (Number of correct answers: 3)

- A. Ensure that DTS can reach both the source database and the destination RDS instance.
- B. Configure database accounts with the permissions required by the migration task.
- C. Review and resolve failures reported by the DTS precheck before starting the task.
- D. Disable the destination RDS IP address allowlist before migration.
- E. Assign a public IP address to every ECS instance in the source environment.

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

The source MySQL database and destination RDS instance must be reachable from the DTS migration task. The required network configuration can use a private network, public network, or another supported connectivity method, depending on the source and destination deployment. The destination RDS instance must also permit the required connection source through its IP address allowlist or applicable network-access settings.

The source and destination database accounts must have the required permissions for the selected migration objects and migration type. DTS performs a precheck before the migration starts, and the administrator should review and resolve precheck failures before proceeding. The precheck verifies important items such as connectivity, account permissions, and configuration conditions. See the [DTS data migration overview](#) and [RDS IP allowlist documentation](#).

**QUESTION NO: 12**

A customer uses Alibaba Cloud Auto Scaling service and creates a scaling group. He/she sets the "Minimum number of instances" to 2 and "Maximum number of instances" to 7.

After correctly configuring scaling settings, the customer adds a scaling rule "Add 6 ECS instances", and creates a scheduled task based on this rule.

And then, the customer checks the scaling group and see it already contains 3 valid ECS instances, he/she then enables this scaling group immediately.

How many ECS instance will be added in the scaling group when the scheduled task is activated once'?

- A. 0
- B. 2
- C. 4

**ANSWER: C****Explanation:**

4 is correct. At the time the scheduled task triggers the scaling rule, the scaling group already has three valid ECS instances. Although the rule requests that six instances be added, Auto Scaling must enforce the scaling group's configured maximum size of seven instances. Therefore, it can increase the group from three to seven instances, which requires creation of only four additional ECS instances. The requested adjustment is constrained by the group capacity boundary rather than allowing the total number of instances to exceed the maximum. This behavior ensures that scheduled tasks and scaling rules remain subject to the scaling group's minimum and maximum instance limits, even when the rule itself specifies a larger adjustment value. Alibaba Cloud documents that a scaling group has minimum and maximum instance limits and that scaling activities are performed within those limits. See the [Alibaba Cloud Auto Scaling overview](#) and the [scaling rule documentation](#).

**QUESTION NO: 13**

As a concept stem from cloud computing, cloud storage service provides on-demand storage and data services to users through a unified interface Which of the following characteristics about cloud storage are correct? (Number of correct answers: 4)

- A. Management and maintenance: To prevent emergent service termination, users need to know the underlying detail like how the data is archived in the resource pool.
- B. Service on demand Cloud storage has a huge resources pool, users can apply for and use the resources on demand.
- C. High reliability and availability: Provides very high reliability and availability through technologies like 3 redundant copies, automatic failure recovery, etc.
- D. Ultra-large scale: Data storage can reach TB or even PB level
- E. Transparent service: Provides the service through a single SDK
- F. High scalability Storage capacity can be dynamically scaled to meet the needs.

**ANSWER: B C D F****Explanation:**

Cloud storage is designed to expose storage as an elastic, managed service rather than as fixed, user-maintained physical infrastructure. "Service on demand Cloud storage has a huge resources pool, users can apply for and use the resources on demand." is correct because users can provision and consume storage as needed through cloud service interfaces. "High reliability and availability: Provides very high reliability and availability through technologies like 3 redundant copies, automatic failure recovery, etc." accurately reflects the use of data redundancy and automated recovery mechanisms to protect stored data and maintain service continuity. "Ultra-large scale: Data storage can reach TB or even PB level" is also a core cloud-storage characteristic: cloud platforms pool distributed infrastructure to support very large data volumes, from terabytes to petabytes and beyond. Finally, "High scalability Storage capacity can be dynamically scaled to meet the needs." is correct because capacity can be expanded without the traditional process of purchasing, installing, and managing additional storage hardware. Alibaba Cloud Object Storage Service is an example of this model, offering scalable object storage and data-redundancy capabilities. See [Alibaba Cloud: What is OSS?](#) and [Alibaba Cloud: Data redundancy types](#).

**QUESTION NO: 14**

Alibaba Cloud Content Delivery Network (CDN) is a distributed network built on top of a base network with edge node server clusters distributed across various regions. If a user finds that the CDN cache hit rate is low, the possible reasons are \_\_\_\_\_. (Choose three.)

- A. When websites have low access traffic and the files are not popular enough, the CDN will receive fewer requests and result in a lower cache hit rate.

- B.** If the cache configuration is inappropriate (e.g., overly short Cache Expiry Time for popular files), this may result in frequent back-to-source operations.
- C.** Insufficient bandwidth was selected when enabling the CDN service, making the service capacity reach its limit.
- D.** If the origin site has a large number of dynamic resources or mostly non-cacheable content, this can also lead to frequent back-to-source operations.
- E.** The HTTP Header setting may cause the cache to fail. In this case, users must check the origin site's "Cache-Control" and "Expires" settings.

**ANSWER: A B D E**

**Explanation:**

Technically, four of the listed statements describe valid causes of a low CDN cache hit rate, so the "Choose three" instruction appears inconsistent with the options. When traffic is low or objects are requested infrequently, edge nodes have fewer repeat requests from which to serve cached copies; consequently, more requests may require retrieval from the origin. Cache-expiration configuration also directly affects the hit rate: short TTL values cause cached objects to expire quickly and increase back-to-origin requests, including for frequently requested files. A workload containing many dynamic or otherwise non-cacheable resources similarly requires origin processing rather than edge-cache delivery. Finally, origin HTTP response headers can override or prevent CDN caching. In particular, the origin's `Cache-Control` and `Expires` headers must allow an object to be cached for the intended period; restrictive directives or unsuitable expiration values can cause cache misses and back-to-origin traffic. Alibaba Cloud documents cache-expiration rules and their role in controlling how long resources remain cached at edge points of presence in the [CDN cache expiration rules documentation](#). The broader behavior of serving cached static resources from edge nodes is described in the [Alibaba Cloud CDN overview](#).

**QUESTION NO: 15**

When using Alibaba Cloud SLB, a user configures the Health Check as follows: response timeout value set to 5 seconds, Health Check interval set to 2 seconds, unhealthy threshold value set to 3, and healthy threshold value set to 3. That is to say, to confirm a backend ECS instance is healthy, three successive normal responses are required.

If the current status of a backend ECS instance is unhealthy, and you finish the repair work immediately, it will take at least \_\_\_\_\_ for SLB Health Check to confirm this instance is healthy again.

- A.** 11 seconds
- B.** 21 seconds
- C.** 5 seconds
- D.** 6 seconds

**ANSWER: A**

**Explanation:**

**11 seconds** is correct. A backend ECS instance that is currently marked unhealthy must meet the configured healthy threshold before Classic Load Balancer (SLB) changes its state back to healthy and resumes normal traffic distribution. Here, the healthy threshold is three, so SLB must receive three consecutive successful health-check results after the repair is complete.

For health-check timing, the 2-second interval determines when consecutive checks are initiated, while the 5-second response timeout defines the maximum period SLB allows for the final required health-check response. Using the health-check convergence calculation applied in this scenario, the three required check opportunities account for  $3 \times 2$  seconds, and the final response may take up to the configured 5 seconds. Therefore, the minimum confirmation period to allow for is  $6 + 5 = 11$  seconds. This is a control-plane health-state transition time; completing the repair does not immediately alter the load balancer's recorded health status.

Alibaba Cloud documents the health-check interval, response timeout, and healthy threshold as configurable parameters that govern backend health determination. See the [Classic Load Balancer health-check configuration documentation](#) and the [Classic Load Balancer health-check overview](#).

#### QUESTION NO: 16

A Virtual Private Cloud (VPC) is an isolated network environment that is completely isolated from each other. The following statements about VPC are correct \_\_\_\_\_. (Number of correct answers: 2)

- A. Each VPC has an independent tunnel ID, and a tunnel ID corresponds to a virtualized network.
- B. Subnets can be divided like the traditional network environment. Different cloud servers within each subnet are connected by the same router and different subnets are interconnected by switches.
- C. ECS instances in a VPC use security group firewalls for Layer 2 network access control.
- D. The internal networks of different VPCs are completely isolated and can be interconnected through IP addresses mapped to the outside.

**ANSWER: A D**

#### Explanation:

Each VPC has an independent tunnel ID, and a tunnel ID corresponds to a virtualized network. Alibaba Cloud VPC is implemented as an isolated, software-defined virtual network over the shared physical cloud infrastructure. The tunnel identifier distinguishes one tenant's virtual network traffic from another's, allowing the platform to maintain logical network separation even when traffic traverses common underlying infrastructure. This isolation model is a core characteristic of a VPC and enables customers to define their own IP address ranges, vSwitches, route tables, and security policies independently.

The internal networks of different VPCs are completely isolated and can be interconnected through IP addresses mapped to the outside. By default, private IP address spaces in separate VPCs cannot directly reach one another. Workloads can communicate through externally reachable, mapped public IP addresses when appropriate public-network configuration and security rules are in place. Alibaba Cloud also provides private connectivity products such as VPC peering and Cloud Enterprise Network for architectures that require private inter-VPC communication. The relevant point is that VPC isolation does not prevent connectivity when a supported communication path is explicitly configured. See [Alibaba Cloud: What is a VPC?](#) and [Alibaba Cloud: Public IP addresses](#).

#### QUESTION NO: 17

Alibaba Cloud Content Delivery Network (CDN) is a distributed network that is built and overlaid on the bearer network. Moreover, it is composed of edge node server clusters distributed across different regions. It replaces the traditional data transmission mode, which is centered on Web servers. When using Alibaba Cloud CDN, a user's request will first reach the edge node, and then receive data from the origin site by means of back-to-source. Moreover, the admin can obtain visitor's real IP on the origin site. Which of the following descriptions relate to "obtaining visitor's real IP" are correct? (Number of correct answers: 2)

- A. "Visitor's real IP" is saved in "X-Forwarded-For" header in HTTP protocol. It can be directly obtained in the user-defined LOG of Apache and Nginx.
- B. "Visitor's real IP" can only be obtained by modifying the application.
- C. You can one-step activate the "recording visitor's real IP" function in Alibaba Cloud CDN console to directly view the visitor's real IP in the access log.
- D. In Windows, if IIS is used: after installing "F5XForwardedFor" extension module, "Visitor's real IP" can then be seen in the log.

**ANSWER: A D**

**Explanation:**

"Visitor's real IP" is saved in "X-Forwarded-For" header in HTTP protocol. It can be directly obtained in the user-defined LOG of Apache and Nginx. is correct because CDN edge nodes forward requests to the origin and add the client address to the standard X-Forwarded-For HTTP request header. The origin server therefore sees the CDN node as the direct network peer, but its web-server logging configuration can record the original client address from this forwarded header. Apache supports defining a custom log format that uses the header, and NGINX supports logging it through its request-header variables. This lets origin access logs retain meaningful client-IP information while traffic is accelerated by CDN. Alibaba Cloud documents this mechanism and provides origin-server-specific configuration guidance in its [Retrieve client IP addresses](#) documentation.

In Windows, if IIS is used: after installing "F5XForwardedFor" extension module. "Visitor's real IP" can then be seen in the log. is also correct. IIS does not natively replace its logged remote address with the value supplied in X-Forwarded-For. The F5XForwardedFor extension is designed to read that header and make the original client IP available to IIS logging, enabling logs to record visitors rather than CDN edge-node addresses. This is the IIS-specific approach described by Alibaba Cloud for CDN-origin deployments. For the relevant IIS extension and setup background, see the [F5 X-Forwarded-For module repository](#).

**QUESTION NO: 18**

Alibaba Cloud Object Storage Service (OSS) provides multiple file uploading modes for different scenarios, which are\_\_\_\_\_.

(Number of correct answers: 2)

- A. Direct upload from the OSS management console no matter how large the file is
- B. Uploading using the API
- C. Directly copying to/from OSS using portable disks
- D. Upload using the SDK

**ANSWER: B D****Explanation:**

Uploading using the API and Upload using the SDK are supported OSS upload modes. OSS exposes RESTful API operations for object upload, including PutObject for direct object uploads and multipart-upload operations for dividing large objects into parts, uploading those parts independently, and then completing the upload. This API approach is suitable when an application needs direct control over HTTP requests, authentication, headers, object metadata, and upload workflow.

SDKs provide the same OSS capabilities through language-specific libraries, making application integration more convenient. Alibaba Cloud OSS SDKs support common development languages and encapsulate request signing, retries, multipart upload, resumable upload, and other operational details. Therefore, SDK uploads are especially appropriate for developers building upload functionality into services, desktop software, or mobile applications. For large objects or networks that may be interrupted, multipart and resumable-upload capabilities allow uploads to be continued efficiently rather than restarted from the beginning. OSS documentation describes both the underlying object upload APIs and SDK-based upload procedures: [PutObject API reference](#) and [OSS SDK overview](#).

**QUESTION NO: 19**

For ECS and RDS instances under different Alibaba Cloud accounts but in the same region, which of the following statements is NOT correct for migrating self-built MySQL databases (running on ECS) to RDS?

- A. The data can be imported via the Intranet
- B. The data cannot be migrated.
- C. The data can be imported via the public network.
- D. The data can be imported by running mysqldump.

**ANSWER: B**

**Explanation:**

**The data cannot be migrated.** is the statement that is not correct. Alibaba Cloud does support migration from a self-managed MySQL database hosted on ECS to an ApsaraDB RDS for MySQL instance when the resources belong to separate accounts. The accounts do not create an inherent migration restriction; instead, the migration configuration must provide network reachability and the required database permissions. A public endpoint can be used where appropriate, while private connectivity can also be used when the relevant VPC networking and cross-account access arrangements are in place. The destination RDS account must authorize the source connection, such as by configuring the RDS IP whitelist, and the migration user needs sufficient privileges for the selected migration method. Alibaba Cloud Database Transmission Service can migrate supported self-managed MySQL sources to RDS and supports full-data migration, with optional ongoing change migration when the source configuration meets the prerequisites. For smaller or planned one-time migrations, a logical export and subsequent import can also be used. Therefore, separate Alibaba Cloud accounts in the same region do not mean that the MySQL data is impossible to migrate. See the [DTS MySQL-to-RDS migration guide](#) and [RDS IP allowlist documentation](#).

**QUESTION NO: 20**

Alibaba Cloud OSS is a cloud storage service that features massive capacity, outstanding security, low cost, and high reliability. OSS provides multiple types of monitoring reports to help you monitor OSS resource consumption in real time. Which of the following monitoring functions are supported by OSS? (Choose three.)

- A. Storage space monitoring
- B. User device type (PC end or mobile end) monitoring for accessing OSS files
- C. Monitoring for OSS service requests (including the number of PUT and GET requests)
- D. UV (Unique Visitor) monitoring for OSS service requests

**ANSWER: A C D**

**Explanation:**

OSS supports **Storage space monitoring** through monitoring metrics and reports that show the amount of data stored in buckets. This enables administrators to observe storage consumption trends, understand capacity growth, and make appropriate lifecycle, archiving, or capacity-planning decisions. OSS also supports **Monitoring for OSS service requests (including the number of PUT and GET requests)**. Request metrics provide operational visibility into how buckets are being used, including request volume and request types, which helps teams evaluate application activity and estimate request-related costs.

**UV (Unique Visitor) monitoring for OSS service requests** is also part of OSS access-analysis capabilities. UV statistics help identify the number of distinct visitors accessing OSS resources during a selected period. Used together, storage, request, and visitor metrics provide a practical picture of OSS consumption: how much data is retained, how actively objects are requested, and how broadly the content is accessed. Alibaba Cloud describes OSS as an object-storage service with monitoring and management capabilities for stored data and access activity. See the [Alibaba Cloud OSS product page](#) and the [OSS monitoring documentation](#).