DUMPSBOSS.

ACA Cloud Computing Associate

Alibaba Cloud ACA-Cloud1

Version Demo

Total Demo Questions: 10

Total Premium Questions: 102

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	3
Topic 2, Alibaba Cloud Core Products	67
Topic 3, Cloud Computing Security	19
Topic 4, Cloud Computing Architecture	13
Total	102

QUESTION NO: 1

You are developing a highly available web application using stateless web servers. Which services are suitable for storing session state data? (Choose two.)

A. RDS

B. Table Store

C. MaxCompute

D. OSS

ANSWER: A B

Explanation:

RDS and Table Store are suitable choices for holding session state outside stateless web servers. Persisting session records in a shared managed data service allows any healthy web-server instance to retrieve a user's session after load balancing, instance replacement, scaling, or failover. This design prevents session affinity from becoming a dependency and supports horizontal scaling of the application tier.

RDS provides managed relational databases and is appropriate when session information benefits from relational storage, transactional updates, SQL queries, or integration with an application's existing database workflow. Session identifiers can be used as keys for rows containing serialized session attributes and expiration information. Alibaba Cloud documents its managed database capabilities at Alibaba Cloud RDS documentation.

Table Store is a managed, highly scalable NoSQL service designed for low-latency reads and writes. Its primary-key-based data model fits session access patterns particularly well: the session ID can be the primary key, while session fields and expiry-related data are stored as attributes. This makes it a strong shared session store for web applications that must scale across many stateless servers. See Alibaba Cloud Table Store documentation.

QUESTION NO: 2

Auto Scaling is a management service that can automatically adjust elastic computing resources based on your business needs and policies. Which of the following application scenarios are suitable for Auto Scaling? (Choose three.)

A. You need to flexibly adjust elastic computing resources based on your business needs, and ECS instances will be added seamlessly to address traffic peaks and surges.

B. You need to flexibly adjust elastic computing resources based on your business needs, and ECS instances will be reduced automatically to save costs when business needs drop.

C. When an ECS instance reacts abnormally in a scaling group, Auto Scaling will create a new ECS instance to take over the computing tasks conducted by the faulty instance.

D. You want to flexibly adjust elastic computing resources based on your business needs, and ECS instances will be automatically upgraded (with more Memory, more powerful CPU) when needed.

ANSWER: A B C

Explanation:

Auto Scaling is suitable when an application's required compute capacity varies over time or when the application needs automated instance-level fault recovery. "You need to flexibly adjust elastic computing resources based on your business needs, and ECS instances will be added seamlessly to address traffic peaks and surges." is a core use case: scaling rules can increase the number of Elastic Compute Service instances in a scaling group when demand rises, helping the application maintain capacity during predictable or unexpected traffic growth. "You need to flexibly adjust elastic computing resources based on your business needs, and ECS instances will be reduced automatically to save costs when business needs drop." is also appropriate because scale-in rules remove excess instances when the configured conditions are met, so users pay only for capacity that remains necessary. "When an ECS instance reacts abnormally in a scaling group, Auto Scaling will create a new ECS instance to take over the computing tasks conducted by the faulty instance." is suitable because health checks allow Auto Scaling to identify unhealthy instances and replace them according to the scaling group

configuration. This supports service availability without requiring manual intervention for each failed instance. Alibaba Cloud documents scaling-out, scaling-in, and health-check-based replacement as key Auto Scaling capabilities. See the Auto Scaling overview and health check documentation.

QUESTION NO: 3

Your website has oscillating traffic peaks that are difficult to predict in advance. In this scenario, it is recommended to use SLB andtogether with ECS.

A. RDS

B. Auto Scaling

C. VPC

D. MaxCompute

ANSWER: B

Explanation:

Auto Scaling is the appropriate service to use with Server Load Balancer (SLB) and Elastic Compute Service (ECS) when website traffic has unpredictable peaks. Auto Scaling monitors demand through configured metrics or schedules and automatically adjusts the number of ECS instances in a scaling group. When traffic increases, it can add ECS instances; when demand subsides, it can remove instances, helping the application maintain capacity while avoiding payment for unneeded compute resources.

SLB distributes incoming requests across the available ECS backend servers. When Auto Scaling and SLB are associated, instances added by a scaling activity can automatically be attached to the relevant SLB server group, allowing them to begin receiving traffic. Instances removed during scale-in are detached appropriately. This combination provides elastic capacity, load distribution, and improved availability without requiring the website operator to accurately forecast every traffic surge in advance. Alibaba Cloud documents the integration of scaling groups with load balancers and the automatic addition or removal of backend servers during scaling activities. See the Alibaba Cloud Auto Scaling documentation for attaching a load balancer and the Server Load Balancer documentation.

QUESTION NO: 4

Your friend purchases some Alibaba Cloud resources and uses them to build an online training website. Before the site goes online, your friend wants to perform some performance tests but doesn't know how to do it. In this case, which of the following services can help?

A. After-Sales Support Service

B. Support Ticket Service

C. Load and Performance Testing Service

D. None Of The Above

ANSWER: C

Explanation:

Load and Performance Testing Service is the appropriate service because it is designed to evaluate how a website, application, or server behaves when subjected to simulated user traffic and varying workload levels before production launch. The service enables users to create performance tests that model concurrent access and requests, then observe key indicators such as response time, throughput, and error rates. These results help the website owner identify capacity limits and performance bottlenecks, validate that the deployed cloud resources can support expected training-site traffic, and make scaling or application-tuning decisions before real users are affected. Alibaba Cloud documents this capability under its performance-testing offerings, describing how users can generate load and analyze the resulting performance metrics for

target services. Running these tests prior to launch is a practical reliability step for an online training platform, where many learners may access videos, course pages, and authentication functions concurrently. See the Alibaba Cloud Performance Testing PTS product page and the Performance Testing PTS documentation for service capabilities and usage guidance.

QUESTION NO: 5

Which of the following practices help improve the security of an Alibaba Cloud account? (Choose two.)

A. Enable MFA for the Alibaba Cloud account.

B. Use RAM users or RAM roles for routine operations.

C. Share the Alibaba Cloud account password among project members.

D. Embed AccessKey secrets in public source code.

ANSWER: A B

Explanation:

Enabling multi-factor authentication (MFA) for the Alibaba Cloud account and avoiding routine use of that account are both recommended security practices. The Alibaba Cloud account has full control over resources and permissions. MFA adds an additional verification factor to the account sign-in process and helps reduce the risk of unauthorized access if the password is compromised.

For daily operations, administrators should create RAM users or use RAM roles and grant only the permissions required for the job. This follows the principle of least privilege and limits the impact of an accidental or unauthorized action. AccessKey secrets should also be protected and should never be exposed in source code repositories or shared publicly. See the Alibaba Cloud account security best practices and MFA documentation.

QUESTION NO: 6

Company B builds a music download website based on OSS and ECS, and users can download mp3 files after registering for the website. Recently, the public network traffic to the OSS has doubled but the increase of registered users is less than 10%. After in-depth analysis, engineers find that many user download requests are raised from search engine but not the website itself. In this case, which of the following measures is thought to be able to address this issue.

A. Migrates music files from the OSS to the data disks for ECS instances.

B. Changes the website access mode to HTTPS.

C. Limits access sources by configuring “Anti-leech settings” of OSS bucket attributes.

D. Introduce SLB for extra protection.

ANSWER: C

Explanation:

Limits access sources by configuring “Anti-leech settings” of OSS bucket attributes is the appropriate measure because the unexpected traffic is caused by direct links to objects being discovered and used through search engines rather than by visitors navigating through the registered website. This is commonly called hotlinking or leeching: a third-party page, search result, or other external source embeds or exposes an OSS object URL, and downloads consume the bucket owner’s outbound bandwidth.

OSS hotlink protection controls access using the HTTP `Referer` header. The bucket owner can configure a whitelist containing the company’s legitimate website domains and, where required, define behavior for requests with an empty Referer. Requests originating from approved website pages can continue to download the MP3 objects, while requests referred by unapproved external sources are denied. This directly reduces unauthorized public download traffic while preserving the existing OSS-based content-delivery architecture. The configuration can be applied at the bucket level and

should be tested with the site's real download flows, including mobile applications or clients that may not send a Referer header. See Alibaba Cloud's OSS hotlink protection documentation and the OSS security overview.

QUESTION NO: 7

Which of the following functions can be provided by Alibaba Cloud API Gateway? (Choose two.)

A. Publish APIs and route requests to backend services.

B. Configure authentication for API callers.

C. Store large media files as objects.

D. Create relational database tables for application data.

ANSWER: A B

Explanation:

API Gateway can publish APIs and route API requests to configured backend services. Backends can include ECS-hosted applications, Function Compute functions, or other supported services. This enables an organization to expose application capabilities through managed API endpoints.

API Gateway also supports API authentication and authorization mechanisms, such as AppCode and signature-based authentication, depending on the API configuration. These controls help restrict API access to authorized clients. API Gateway is not a replacement for object storage and does not provide persistent relational database storage. See What is API Gateway?.

QUESTION NO: 8

Anti-DDoS is one of the major features in Alibaba Cloud Security Service. Many websites have experienced different types of DDoS attacks, an accurate understanding of DDoS is crucial to protecting websites. Which of the following descriptions about DDoS attack is the most accurate one?

A. The main target for DDoS attack is the database.

B. DDoS attack will try to crack the password through a large number of attempts.

C. The main purpose of DDoS attack is to make the target server unable to provide service normally, which is one of the network attacks that are most difficult to defend.

D. The purpose of DDoS attack is to steal confidential information.

ANSWER: C

Explanation:

The main purpose of DDoS attack is to make the target server unable to provide service normally, which is one of the network attacks that are most difficult to defend. A distributed denial-of-service attack uses traffic or requests from numerous distributed sources to consume a target's available network bandwidth, connection capacity, or computing resources. The intended outcome is degraded performance or an outage that prevents legitimate users from accessing the application, website, or other online service. This focus on disrupting availability is the defining characteristic of DDoS attacks and is why layered detection, traffic scrubbing, rate limiting, and scalable mitigation capacity are important. Alibaba Cloud Anti-DDoS services are designed to detect malicious volumetric and application-layer traffic, scrub attack traffic, and forward legitimate traffic to protected assets so that services can remain available during an attack. Because attack sources can be widely distributed and attack patterns can vary, effective protection requires specialized, continuously updated mitigation capabilities. See the Alibaba Cloud Anti-DDoS Pro and Premium overview and the Cloudflare explanation of DDoS attacks.

QUESTION NO: 9

When using Alibaba Cloud OSS, you can set access control in 3 different levels, which are _____ (Choose three.)

A. Bucket level

B. Object level

C. RAM account level

D. Directory level

ANSWER: A B D

Explanation:

Alibaba Cloud Object Storage Service (OSS) supports access control list (ACL) configuration at the bucket, object, and directory levels. **Bucket level** access control defines the default public or private access posture for the entire bucket, making it the foundational scope for protecting all content stored in that bucket. **Object level** access control permits an individual object to have an ACL that differs from the bucket's ACL, which is useful when a specific file needs more restricted or more broadly accessible permissions than the bucket default. **Directory level** access control applies to objects grouped under a specified prefix (a directory-like path in OSS). This provides a practical way to manage access consistently for a logical set of objects without configuring every object separately. OSS directories are implemented as prefixes rather than physical folders, but OSS supports directory ACL settings for this purpose. These three scopes allow permissions to be applied broadly, precisely, or to a logical collection of objects, depending on the required access-management granularity. See the Alibaba Cloud documentation for OSS ACL overview and managing OSS ACLs.

QUESTION NO: 10

Alibaba Cloud's object storage (Object Storage Service, OSS for short) can be used to store massive amount of files, but there is still a maximum number of files in a single bucket.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because Alibaba Cloud OSS does not impose a maximum number of objects (files) that can be stored in one bucket. A bucket is a logical container for objects, and OSS is designed for massive-scale object storage; users can continue adding objects to a bucket without encountering a fixed object-count quota. OSS documentation lists limits for relevant properties such as bucket naming, object key length, object size, multipart-upload parts, and request behavior, but not a maximum number of objects per bucket.

In practice, applications should use a sensible object-key naming convention, often including prefixes such as dates, tenant identifiers, or business categories. Prefixes make large collections easier to organize and can support efficient listing and lifecycle-management designs. However, these prefixes are part of object names and do not create a maximum file-count limit for the bucket. Storage consumption and applicable billing are determined by the amount of data stored and the operations performed, rather than by a capped number of files. See Alibaba Cloud's OSS limits documentation and bucket documentation for the bucket and object model.