

ACA Developer Certification

Alibaba Cloud ACA-Developer

Version Demo

Total Demo Questions: 13

Total Premium Questions: 133

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Alibaba Cloud Product Knowledge	119
Topic 2, Application Development Knowledge	14
Total	133

QUESTION NO: 1

A developer accesses logs in a Log Service Logstore via the API. The error code returned by the server is 404. Which of these could be the cause?

- A. The log project (Project) does not exist
- B. The requested digital signature does not match
- C. Server internal error
- D. Server is busy: please try later

ANSWER: A

Explanation:

The log project (Project) does not exist is correct. In Alibaba Cloud Simple Log Service (SLS), a Logstore belongs to a Project, and API requests identify the target Project as part of the request path or endpoint context. Before SLS can locate and return logs from a Logstore, it must first resolve that Project. If the specified Project name is absent, has been deleted, is in the wrong region or endpoint context, or is otherwise not found for the request, SLS returns an HTTP 404 response to indicate that the requested resource cannot be found.

When investigating this response, confirm that the Project name in the API call is spelled exactly as created, that the request is sent to the endpoint for the Project's region, and that the Project still exists. Then verify the Logstore name and resource configuration. Alibaba Cloud documents SLS API error codes and their HTTP status mappings in the [SLS API error codes reference](#). The SLS documentation also describes the Project as the resource container used to manage Logstores and related SLS resources; see [Manage a Project](#).

QUESTION NO: 2

A Developer is creating a Function Compute function. This function requires around 1 GB of temporary storage for files while executing. These files will not be needed after the function is completed for a fixed period. How can this developer MOST efficiently handle the temporary files?

- A. Store the files in Cloud Disk and delete the files when Function Compute function call completes.
- B. Copy the files to NFS and delete the files when the Function Compute function call completes.
- C. Store the files in the /tmp directory on a temporary ECS instance, and delete the files at when the Function Compute function call completes.
- D. Copy the files to an OSS bucket with a lifecycle policy configured to delete the files automatically.
- E. Configure the Function Compute temporary disk to at least 1 GB and store the files in the function runtime's /tmp directory.

ANSWER: E

Explanation:

Configure the Function Compute function with at least 1 GB of temporary disk space and write the working files to the function runtime's /tmp directory. This is the most efficient approach because the files are needed only during execution and therefore do not require a separately provisioned or managed persistent storage service. Function Compute provides a writable temporary disk for each function instance; its size is configurable, including values large enough for a 1 GB workload. The temporary disk is intended for transient artifacts such as downloaded source files, decompression output, intermediate processing data, and generated files that are consumed before the invocation completes.

Using the runtime-local temporary directory also avoids the network transfer, request overhead, and operational work associated with external storage for data that has no value once processing is done. Data on the temporary disk can remain available when the same instance is reused, but it must not be treated as durable storage because the instance can be reclaimed at any time. The function should therefore clean up sensitive or unnecessarily large files when practical, while relying on the ephemeral runtime disk for normal execution-time workspace. Alibaba Cloud documents temporary disk sizing

and its use in Function Compute at [Configure the disk size](#) and provides runtime filesystem guidance in the [Function Compute documentation](#).

QUESTION NO: 3

Alibaba Cloud's ECS disk snapshotting tool can create crash-consistent snapshots for all types of Cloud Disks. It is a convenient and efficient disaster recovery method. It is often used for data backup, custom mirroring, application disaster recovery, etc. Which of the following is NOT a correct description of disk snapshots?

- A. A snapshot is a backup file of the cloud disk data state at a certain point in time. The first snapshot of the cloud disk is a full snapshot of the ac usage. Empty data blocks are not backed up. All subsequent snapshots are incremental snapshots, and only changed data blocks are stored.
- B. Snapshots can be created manually through the web console or command line tools, or automatically according to a schedule.
- C. In the event of a system failure, you can use a snapshot to roll back the Cloud Disk to an earlier point in time.
- D. Snapshot rollbacks can be performed while the ECS instance is running.

ANSWER: D

Explanation:

“Snapshot rollbacks can be performed while the ECS instance is running.” is the statement that is not correct. Alibaba Cloud requires the associated ECS instance to be stopped before a disk can be rolled back from a snapshot. This prerequisite helps ensure that no active operating-system, application, or file-system writes conflict with restoration of the disk's blocks to their earlier snapshot state. The rollback operation replaces the current disk data with the data captured at the selected snapshot recovery point, so it is intended to be performed during a planned interruption of the affected instance. After the rollback finishes, the instance can be started again and will use the restored disk contents. This operational requirement is particularly important when rolling back a system disk because the system disk contains the operating system and boot environment, but Alibaba Cloud's rollback procedure also specifies stopping the relevant ECS instance before restoring a data disk. Administrators should therefore plan snapshot rollback as a recovery action that includes instance downtime, rather than expecting it to occur transparently on a running workload. See the Alibaba Cloud [Roll back a disk](#) documentation and the [Snapshots overview](#) for the applicable operation and snapshot guidance.

QUESTION NO: 4

An application uploads objects to OSS and requires the uploaded data to be encrypted by OSS by using a customer-managed key in Key Management Service (KMS). Which server-side encryption method should be selected?

- A. SSE-KMS
- B. SSE-OSS
- C. Client-side encryption only
- D. HTTPS encryption

ANSWER: A

Explanation:

SSE-KMS is correct. Server-Side Encryption with KMS causes OSS to encrypt object data at rest by using a data key protected by a KMS key. A customer-managed KMS key can be selected when the organization needs centralized key lifecycle management, authorization control, auditing, and key-rotation capabilities.

With SSE-KMS, encryption and decryption are performed transparently by OSS for authorized requests. The application does not need to implement encryption logic or manage plaintext data keys itself. In contrast, SSE-OSS uses OSS-managed

keys, while client-side encryption requires the application to encrypt data before upload. See [OSS server-side encryption](#) and the [Key Management Service overview](#).

QUESTION NO: 5

Which of the following data volume types is not supported by Elastic Container Instance (ECI)?

- A. EmptyDir
- B. NFS
- C. ConfigFile
- D. FixedVolume

ANSWER: D

Explanation:

FixedVolume is the correct answer because it is not an Elastic Container Instance data-volume type. ECI defines volume types through its container-group configuration and API schema, where storage is specified using supported volume structures such as temporary empty-directory storage, NFS-backed storage, configuration-file volumes, and extensible FlexVolume-based storage. "FixedVolume" is not a recognized volume definition or configuration field in the ECI volume model, so it cannot be used to attach or mount data storage to an ECI container group.

When provisioning ECI programmatically, each volume must be declared using a supported volume category and its associated parameters, such as a mount path and the relevant backend configuration. This allows ECI to validate the storage request and correctly prepare the volume before containers start. A name that is not part of the supported ECI volume schema, including FixedVolume, cannot represent a valid data-volume configuration. The ECI documentation describes the supported data-volume capabilities, and the CreateContainerGroup API reference shows the volume configuration structures accepted when creating container groups. See the [ECI data volumes overview](#) and the [CreateContainerGroup API reference](#).

QUESTION NO: 6

The event monitoring feature in CloudMonitor provides custom event monitoring functions. Which of the following data reporting methods are included? (Number of correct answers: 3)

- A. Report event monitoring through the Python SDK
- B. Report event monitoring through the Java SDK
- C. Report event monitoring data via HTTP
- D. Report event monitoring data through the command line (CU)

ANSWER: A B C

Explanation:

CloudMonitor custom event monitoring lets an application submit business or operational events to CloudMonitor, where they can be used to create event-triggered alert rules. The supported reporting approaches include calling the reporting capability through the Python SDK, which allows Python applications to authenticate with Alibaba Cloud credentials and submit custom event payloads programmatically. The Java SDK provides the equivalent integration path for Java services, enabling event reporting directly from application code. Custom event data can also be reported over HTTP by making an API request with the required event fields and authenticated Alibaba Cloud API parameters. These mechanisms support automated reporting from different application stacks while using the same CloudMonitor custom-event service and event model. Alibaba Cloud documents custom event monitoring and its reporting workflow in the [CloudMonitor Custom Event Monitoring guide](#). The underlying API operation for submitting custom events is also documented in the [PutCustomEvent API reference](#), which is accessible through HTTP-based API calls and supported SDKs.

QUESTION NO: 7

A developer runs a web application on a physical machine in an IDC. The machine is running Ubuntu Linux. Which of the following is the recommended method if he wants to migrate this web application onto an Alibaba Cloud ECS instance?

- A.** Use a third-party image creation tool to create an image from the physical server. Upload the image to the ECS console to turn it into a custom image. Then, use the custom image to create ECS instances.
- B.** Use SMC(Server Migration Center) to collect data from the origin physical server. Create an intermediate ECS instance on Alibaba Cloud and wait for the data transfer to complete. Then generate a new image directly from that intermediate ECS instance.
- C.** Use a third-party image creation tool to create an image from the physical server. The image is then uploaded to Alibaba Cloud Object Storage Service (OSS) for storage. In the ECS console, import that image from OSS, then, create an ECS instance from the custom image.

ANSWER: B

Explanation:

Use SMC(Server Migration Center) to collect data from the origin physical server. Create an intermediate ECS instance on Alibaba Cloud and wait for the data transfer to complete. Then generate a new image directly from that intermediate ECS instance. This is the recommended approach because Server Migration Center (SMC) is Alibaba Cloud's managed migration service for moving physical servers, virtual machines, and cloud-hosted servers to Elastic Compute Service (ECS). An SMC client is installed on the Ubuntu source server and transfers the operating system, application files, configurations, and disks through a controlled migration job. For a migration whose destination is an ECS image, SMC uses an intermediate ECS instance during the conversion and synchronization process, then produces a custom image that can be used to create the target ECS instance. This workflow is designed to preserve the source environment as consistently as possible while avoiding the operational complexity of manually creating, validating, uploading, and importing a disk image. Before migration, the developer should ensure that the source server can reach Alibaba Cloud migration endpoints, has sufficient available disk space, and that the target region has appropriate ECS quotas and permissions. After the custom image is created, an ECS instance can be launched from it and the web application can be validated before cutover. See [Alibaba Cloud: What is Server Migration Center?](#) and [Alibaba Cloud: Migrate a server to an ECS image](#).

QUESTION NO: 8

Which of these are suitable use cases for Alibaba Cloud's MongoDB service? (Number of correct answers: 3)

- A.** Applications requiring high OPS, such as logging apps
- B.** Scenarios where complex transaction support is required
- C.** Applications where JSON is a good choice of storage format
- D.** Applications requiring geo-location queries, such as logistics or order status apps

ANSWER: A C D

Explanation:

Alibaba Cloud's MongoDB service is well suited to applications that need scalable, high-throughput document storage, including logging applications that generate large volumes of write operations. MongoDB's flexible document model allows records with changing or varied fields to be stored efficiently, which is useful for operational logs, event data, and similar workloads. Alibaba Cloud positions ApsaraDB for MongoDB as a fully managed document database service designed for high availability and elastic scalability; see the [Alibaba Cloud MongoDB product page](#).

Applications where JSON is a good choice of storage format are also an appropriate fit because MongoDB stores data as BSON documents, a binary representation closely aligned with JSON-style objects. This model supports nested structures and evolving schemas without requiring a fixed relational-table design. In addition, applications requiring geo-location queries are suitable because MongoDB provides geospatial indexes and query operators for storing coordinates and finding data by proximity, containment, or intersection. These capabilities are directly applicable to logistics, delivery tracking,

location-aware order status, and nearby-resource searches. MongoDB's geospatial query capabilities are documented in the [MongoDB geospatial queries documentation](#).

QUESTION NO: 9

Which of the following cloud services can be integrated with Alibaba Cloud KMS (Key Managed Service)? (Number of correct answers: 3)

- A. Elastic Compute Service (ECS)
- B. Object Storage Service (OSS)
- C. Relational Database Service (RDS)
- D. Server Load Balancer (SLB)

ANSWER: A B C

Explanation:

Elastic Compute Service (ECS), Object Storage Service (OSS), and Relational Database Service (RDS) can integrate with Alibaba Cloud Key Management Service to use customer master keys for encryption operations. ECS uses KMS when encrypting Elastic Block Storage disks, allowing disk data to be protected at rest with an encryption key managed through KMS. OSS supports server-side encryption with KMS-managed keys, so objects written to a bucket can be encrypted transparently and decrypted when authorized users or applications access them. RDS integrates with KMS for database encryption, enabling encryption of database storage and related backup data while the key lifecycle, authorization, rotation, and auditing are centrally managed in KMS. These integrations let an organization retain centralized control over cryptographic keys instead of placing key-management responsibilities in individual applications or service configurations. Access to the keys is governed through Alibaba Cloud identity and resource permissions, supporting separation of duties between key administrators and service users. See the [Key Management Service overview](#) and the [OSS server-side encryption documentation](#) for implementation details.

QUESTION NO: 10

You have an application with the following requirements:

7. There will be a lot of video content that need to be uploaded.
2. Each uploaded video needs to be preprocessed to provide different video resolutions and quality levels to accommodate different devices and network conditions.
3. It should be possible to process video in parallel at scale, to ensure that videos are preprocessed in a reasonable amount of time.

Which of the following Alibaba Cloud applications can be combined to provide a low operation and maintenance burden and relatively low operating costs? (Number of correct answers: 3)

- A. Serverless Workflow
- B. Function Compute
- C. Log Service
- D. Elastic Compute Service (ECS)
- E. Server Load Balancer (SLB)

ANSWER: A B C

Explanation:

Serverless Workflow, Function Compute, and Log Service provide a serverless, event-driven approach to video-preprocessing pipelines. Serverless Workflow can coordinate the stages of a transcoding process, including invoking processing tasks, handling retries, tracking execution state, and managing parallel branches. This avoids building and operating custom workflow-control services. Function Compute supplies elastic execution for processing jobs. Functions can be invoked for individual uploaded-video events and scale out concurrently as workload increases, so many videos and their resolution variants can be processed in parallel without provisioning or maintaining a fleet of servers. Its pay-as-you-go model also aligns cost with actual processing activity. Log Service supplies centralized collection, querying, and analysis of function and workflow logs, making execution status, errors, and performance observable without operating a separate logging platform. Together, these services reduce infrastructure administration while supporting scalable asynchronous processing and operational monitoring. Alibaba Cloud describes Function Compute as a fully managed, event-driven compute service that automatically allocates compute resources; see [Function Compute documentation](#). Workflow orchestration capabilities are documented at [Serverless Workflow documentation](#).

QUESTION NO: 11

Which of the following web applications are suitable for deployment on Simple Application Server (SAS), as compared with LCS? (Number of correct answers: 3)

- A. Quickly build a personal blog website
- B. Quickly build a small e-commerce website
- C. Quickly build serverless site using Function Compute
- D. Set up a NoSQL database service
- E. Quickly build a personal, learning, or test environment

ANSWER: A B E

Explanation:

Simple Application Server is intended to reduce the operational overhead of hosting lightweight applications by providing straightforward server provisioning, application images, networking, and basic management capabilities. A personal blog is a core use case because it can be deployed rapidly on a small instance using common web stacks and blog-oriented application images. A small e-commerce website is also appropriate when its expected traffic and architecture fit a lightweight single-server deployment; SAS provides a practical starting point for hosting the web application and its supporting runtime. Personal learning and test environments are likewise well suited because developers can quickly create an isolated Linux or Windows server, install tools, test code, and remove the environment when it is no longer required. These workloads benefit from SAS's simplified configuration and predictable, entry-level server resources. Alibaba Cloud describes Simple Application Server as a service for building applications, websites, and development or test environments with simplified management. For workload details and available application-image deployment paths, see [Alibaba Cloud: What is Simple Application Server?](#) and [Alibaba Cloud Simple Application Server product page](#).

QUESTION NO: 12

By default Function Compute cannot access resources within a VPC network group. To enable Function Compute to access resources located in a VPC, which of the following parameters needs to be configured? (Number of correct answers: 3)

- A. vpcId
- B. vswitchId
- C. securityGroupId
- D. ENIID

ANSWER: A B C

Explanation:

Function Compute requires VPC configuration that identifies the network environment in which its elastic network interface is created. The `vpcId` parameter specifies the target Virtual Private Cloud. This establishes the logical private network that contains the resources the function must reach, such as an Elastic Compute Service instance, database, or internal service endpoint.

The `vswitchIds` parameter supplies one or more vSwitches in that VPC. Function Compute uses the selected vSwitches to allocate network interfaces and private IP addresses for function instances. Selecting suitable vSwitches also determines the zones and available IP address ranges used for this connectivity.

The `securityGroupId` parameter associates the function's network interface with a security group. Security-group rules control allowed inbound and outbound traffic, so this association is necessary to permit the function's required communication with target VPC resources. Together, these three settings provide the VPC identity, subnet placement, and network-access policy required for Function Compute VPC access. See the Alibaba Cloud documentation on [configuring a network for Function Compute](#) and the [CreateFunction API VPC configuration](#).

QUESTION NO: 13

You can use Alibaba Cloud Data transmission service (DTS) or the open source tool Redis-shake to migrate and synchronize data from a local Redis instance to an Alibaba Cloud Redis instance. What is one feature which DTS has but Redis-shake does not?

- A. Full data migration
- B. Incremental data migration
- C. One-way data synchronization
- D. Two-way data synchronization

ANSWER: D

Explanation:

Two-way data synchronization is the distinguishing DTS capability. Alibaba Cloud DTS provides managed synchronization architectures that include bidirectional synchronization, allowing changes made in either participating data source to be synchronized to the other. This is valuable in migration, hybrid-cloud, and active-active-style scenarios where applications may continue writing to both sides and data consistency must be maintained during the synchronization period. DTS manages the synchronization task, connectivity, progress monitoring, and operational controls through the Alibaba Cloud service, rather than requiring the user to build bidirectional replication orchestration independently.

Redis-shake is an open-source Redis migration and synchronization utility designed around a source-to-target replication workflow. It can perform an initial transfer and continue forwarding subsequent source changes, but it does not supply a managed, bidirectional synchronization topology in which updates from each endpoint are replicated back to the other. Thus, among the listed capabilities, bidirectional synchronization is the feature provided by DTS but not Redis-shake. Alibaba Cloud documents the data synchronization topologies, including two-way synchronization, in its [DTS data synchronization overview](#). Redis-shake's source-to-target synchronization model is described in its [official project repository](#).