

ACA Database Certification

Alibaba Cloud ACA-Database

Version Demo

Total Demo Questions: 11

Total Premium Questions: 111

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

Which of the following are database architectures which have emerged during the evolution of database systems? (Number of correct answers: 3)

- A. Single Instance
- B. Shared Nothing
- C. Shared Everything
- D. Dedicated

ANSWER: A B C

Explanation:

Database-system evolution is commonly described through single-instance, shared-everything, and shared-nothing architectures. A single-instance architecture runs the database engine on one server, with processing, memory, and storage resources managed by that instance. It is the foundational deployment model and remains appropriate for many workloads that do not require distributed scale-out capability.

In a shared-everything architecture, multiple processors or nodes access common resources, such as shared memory or storage. This model supports coordinated processing against a common data set and is an important historical and practical database architecture. In a shared-nothing architecture, each node owns its processing, memory, and typically its local data partition; nodes coordinate through a network. This design enables horizontal scaling and fault isolation, making it a core model for distributed databases and data warehouses.

Alibaba Cloud's managed database portfolio reflects this progression from conventional instance-based services toward distributed and decoupled compute-and-storage designs. For example, ApsaraDB RDS provides managed database instances, while PolarDB uses a distributed architecture with separated compute and shared storage. See the [ApsaraDB RDS overview](#) and the [PolarDB overview](#).

QUESTION NO: 2

When high availability switchovers occur in a database, which of the following features provided by ApsaraDB RDS for PostgreSQL can prevent the incremental data loss of logical subscriptions?

- A. Ganos
- B. Pase
- C. Slot failover
- D. High concurrency optimization
- E. Point-in-time restore

ANSWER: C

Explanation:

Slot failover is the ApsaraDB RDS for PostgreSQL feature designed to preserve logical replication continuity during a high-availability switchover. Logical subscriptions consume changes from a publisher through a logical replication slot. That slot maintains the subscriber's confirmed consumption position, which identifies the point from which incremental changes must continue to be delivered. During an HA role change, preserving and synchronizing this slot state to the new primary is essential; otherwise, the subscriber can lose its incremental-change position and may need to resynchronize data.

When slot failover is enabled, RDS synchronizes logical replication slot metadata between the primary and standby instances. After the standby is promoted during a switchover or failover, the corresponding slot is available on the new

primary with the required progress information. The logical subscription can therefore resume consuming WAL changes from the appropriate LSN, avoiding incremental data loss caused by the primary-role transition. This feature is specifically intended for logical replication scenarios that require availability as well as continuity of subscription data. See the Alibaba Cloud documentation for [replication slot failover](#) and PostgreSQL's overview of [logical replication](#).

QUESTION NO: 3

Which of the following features are provided by PolarDB-O's proxy? (Number of correct answers: 4)

- A. Read/write splitting
- B. Load balancing
- C. Health checks for nodes
- D. Failover strategies for high availability
- E. Data encryption

ANSWER: A B C D

Explanation:

PolarDB-O's database proxy provides the traffic-management functions needed to use a cluster's primary and read-only nodes efficiently and reliably. **Read/write splitting** directs write requests to the primary node while routing eligible read requests to read-only nodes, reducing read pressure on the primary. **Load balancing** distributes read traffic among available read-only nodes according to the proxy's routing and balancing policies, helping the cluster use its read capacity effectively.

The proxy also performs **health checks for nodes** so it can identify whether cluster nodes are available for connections and request processing. This node-awareness is essential to safe routing because the proxy must avoid sending traffic to an unavailable or unsuitable node. In addition, its **failover strategies for high availability** allow client access to continue through the proxy when a primary-node switchover or failure occurs. The proxy tracks topology changes and updates routing so that applications can reconnect or continue with minimal dependence on manually changing database endpoints. These capabilities centralize connection routing and high-availability handling at the database access layer. See Alibaba Cloud's [Database Proxy documentation](#) and [read/write splitting documentation](#).

QUESTION NO: 4

Which of the following architectures is not supported by ApsaraDB RDS for MySQL?

- A. The single-node Basic Edition architecture, which can be used in test and non-core production environments
- B. The high-availability architecture, in which primary and read-only nodes are deployed to implement automatic failovers
- C. The three-node Enterprise Edition architecture, which ensures zero data loss
- D. The serverless architecture, in which computing resources are charged based on the Pay-As-You-Go (PAYG) billing method

ANSWER: D

Explanation:

The serverless architecture, in which computing resources are charged based on the Pay-As-You-Go (PAYG) billing method, is the correct choice because it is not an ApsaraDB RDS for MySQL deployment architecture in the edition-based architecture classification tested here. For RDS for MySQL, Alibaba Cloud documents the supported instance architectures as Basic Edition, High-availability Edition, and Cluster Edition. Basic Edition uses a single database node. High-availability Edition provides a primary and secondary-node design with automatic failover capabilities. Cluster Edition uses three nodes to provide stronger availability and data-protection capabilities, including protection against data loss caused by a node failure.

Serverless describes an elastic service and consumption model rather than one of those underlying Basic, high-availability, or cluster deployment architectures. PAYG billing similarly describes how resources may be charged; it does not define a database-node architecture. Therefore, when the question asks specifically for an unsupported architecture among the listed RDS for MySQL architecture models, the serverless/PAYG description is the appropriate answer. Alibaba Cloud's architecture overview and edition comparison are available in the [ApsaraDB RDS for MySQL instance types documentation](#) and the [RDS for MySQL product editions documentation](#).

QUESTION NO: 5

If a database exception occurs in a PolarDB cluster because important data is deleted by mistake, how can data be restored?

- A. Level-2 backup + Redo log
- B. Level 1 backup + Redo log
- C. Level-2 backup
- D. Level-1 backup

ANSWER: A

Explanation:

Level-2 backup + Redo log is the appropriate recovery method because an accidental deletion normally requires point-in-time recovery, rather than restoration only to the time at which a backup was created. The Level-2 backup provides the baseline copy of the data. Redo logs record subsequent committed changes, allowing PolarDB to replay changes from that backup forward to a selected recovery time immediately before the unintended deletion occurred. This lets the administrator recover the required historical state and retrieve the deleted data with minimal data loss relative to the chosen recovery point.

In practice, the recovery target should be selected just before the DELETE operation was committed. The restored data can then be validated and, where appropriate, exported or copied back into the production cluster. This recovery approach depends on having retained both the relevant Level-2 backup and the redo-log data for the required time window. Alibaba Cloud documents PolarDB backup and restoration concepts, including backup retention and restoration capabilities, in its [Backup and restoration overview](#). For the operational process of recovering data to a historical point, see the PolarDB [Restore data documentation](#).

QUESTION NO: 6 - (SIMULATION)

Which of the following are features supported by Alibaba Cloud RDS for MySQL? (The number of

ANSWER: Check the explanation for the answer

Explanation:

Supported RDS for MySQL features include:

Automated backups and point-in-time recovery (PITR).

High availability, using a primary/secondary architecture with automatic failover for HA editions.

Read-only instances, which can be used to scale read workloads.

Monitoring and alerting for instance and database performance metrics.

Database proxy / read-write splitting (where supported by the selected edition/version).

SQL audit and performance optimization capabilities, such as SQL Explorer/SQL Insights where enabled.

Therefore, select the options corresponding to managed backup and recovery, HA/failover, read-only scaling, and the applicable monitoring/audit or proxy capabilities. RDS for MySQL is a managed database service; Alibaba Cloud handles core operations such as backup, recovery, and high-availability maintenance.

QUESTION NO: 7

From which of the following data sources or databases can Data Lake Analytics (DLA) read data for analysis? (Number of correct answers: 6)

- A. MySQL
- B. Oracle
- C. OTS
- D. MongoDB
- E. AnalyticDB
- F. HOLO
- G. OSS

ANSWER: A B C D E G

Explanation:

Data Lake Analytics (DLA) is designed for federated, serverless data analysis. It can define external tables over data held in Alibaba Cloud storage and in supported relational or NoSQL systems, allowing SQL queries to access data without first copying all of it into a separate analytics repository. Therefore, DLA can read MySQL and Oracle data through its database connectors, use OTS (Table Store) as a supported NoSQL data source, and access MongoDB data for SQL-based analysis. It also supports AnalyticDB as an external data source and can directly query files and datasets stored in OSS, which is a core DLA data-lake storage integration.

In practice, the connection is configured by creating a catalog or external table definition containing the endpoint and access credentials, then querying the resulting table using DLA SQL. This architecture lets teams combine operational database data with object-storage data for exploration, reporting, and ETL-style processing while retaining the data in its original service. Alibaba Cloud's DLA documentation describes its support for querying OSS and connecting to multiple external database and NoSQL data sources; see the [DLA overview](#) and the [DLA data sources documentation](#).

QUESTION NO: 8

Which of these are NOT good use cases for DTS (Data Transmission Service)? (Number of correct answers: 2)

- A. Incremental synchronization for disaster recovery
- B. Full database synchronization for database migration
- C. Undoing changes made to a database by mistake
- D. Saving database logs

ANSWER: C D

Explanation:

Undoing changes made to a database by mistake and **Saving database logs** are not appropriate DTS use cases. Alibaba Cloud Data Transmission Service is designed to move and continuously replicate database data. Its core capabilities include data migration, change tracking, and data synchronization between supported source and destination databases. DTS can perform an initial full-data migration or synchronization and then replicate incremental changes, which makes it suitable for migration projects, data integration, and maintaining a synchronized standby or downstream database.

Undoing an accidental database change requires a recovery capability based on backups, point-in-time recovery, transaction controls, or database-native rollback mechanisms. DTS replicates committed source changes rather than serving as a general-purpose mechanism for reversing them. Likewise, DTS consumes database logs or change records as part of its

replication process, but it is not a log-retention or log-archiving service. Persisting and managing database logs should use the database's backup, logging, audit, or dedicated storage and observability features. Alibaba Cloud describes DTS as a service for migration, synchronization, and change tracking in its [DTS overview](#). Its [data synchronization documentation](#) also explains the full and incremental synchronization model that underpins its intended replication scenarios.

QUESTION NO: 9

Which of the following statements about high availability of ApsaraDB RDS are correct? (Number of correct answers: 3)

- A. The primary instance processes write requests.
- B. A secondary instance can take over when the primary instance fails.
- C. A multi-zone deployment improves resilience against zone-level faults.
- D. A read-only instance and a secondary instance have the same purpose.
- E. Regular backups remove the need for a high-availability architecture.

ANSWER: A B C

Explanation:

The primary instance processes write requests in a high-availability RDS architecture. **A secondary instance can take over when the primary instance fails**, which helps restore database availability through an automatic failover mechanism. **A multi-zone deployment improves resilience against zone-level faults** by distributing high-availability resources across zones within a region.

A read-only instance is used to scale read workloads and is not the same as the secondary instance that provides high availability. Backups are essential for recovery from data loss, accidental deletion, and corruption, but they do not replace a high-availability architecture because backup restoration generally requires more time than failover. See the [ApsaraDB RDS high availability documentation](#).

QUESTION NO: 10

What is the purpose of a foreign key in a relational database?

- A. To enforce a relationship between tables
- B. To encrypt sensitive column values
- C. To create a full database backup
- D. To automatically partition a table across database nodes

ANSWER: A

Explanation:

To enforce a relationship between tables is correct. A foreign key is a column or group of columns in one table that refers to a primary key or unique key in another table. It helps enforce referential integrity by preventing child-table records from referencing parent records that do not exist.

For example, an orders table can contain a customer ID foreign key that references the customer ID primary key in a customers table. The database can then ensure that each order is associated with a valid customer. A foreign key does not itself provide encryption, backup, or automatic query optimization.

QUESTION NO: 11

When a production database table requires a controlled schema modification with approval and auditing, which Data Management (DMS) capability is most appropriate?

- A. Change Management
- B. Data Transmission Service
- C. Database Backup Service
- D. Change Tracking

ANSWER: A

Explanation:

Change Management is correct. DMS Change Management provides a controlled workflow for database changes, including schema modifications and data changes. It can apply approval processes, execution controls, and audit records so that production changes are reviewed and traceable.

This capability helps database development and operations teams reduce the risk of unauthorized or incorrectly executed changes. Data migration is performed by DTS, backup and restoration are handled by DBS or database backup features, and change tracking is a DTS capability for capturing database changes. See the [Data Management documentation](#).