



DBS-C01^{Q&As}

AWS Certified Database - Specialty (DBS-C01)

Pass Amazon DBS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/aws-certified-database-specialty.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

A business needs a data warehouse system that stores data consistently and in a highly organized fashion. The organization demands rapid response times for end-user inquiries including current-year data, and users must have access to the whole 15-year dataset when necessary. Additionally, this solution must be able to manage a variable volume of incoming inquiries. Costs associated with storing the 100 TB of data must be maintained to a minimum.

Which solution satisfies these criteria?

- A. Leverage an Amazon Redshift data warehouse solution using a dense storage instance type while keeping all the data on local Amazon Redshift storage. Provision enough instances to support high demand.
- B. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Provision enough instances to support high demand.
- C. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Enable Amazon Redshift Concurrency Scaling.
- D. Leverage an Amazon Redshift data warehouse solution using a dense storage instance to store the most recent data. Keep historical data on Amazon S3 and access it using the Amazon Redshift Spectrum layer. Leverage Amazon Redshift elastic resize.

Correct Answer: C

Explanation: <https://docs.aws.amazon.com/redshift/latest/dg/concurrency-scaling.html> "With the Concurrency Scaling feature, you can support virtually unlimited concurrent users and concurrent queries, with consistently fast query performance. When concurrency scaling is enabled, Amazon Redshift automatically adds additional cluster capacity when you need it to process an increase in concurrent read queries. Write operations continue as normal on your main cluster. Users always see the most current data, whether the queries run on the main cluster or on a concurrency scaling cluster. You're charged for concurrency scaling clusters only for the time they're in use. For more information about pricing, see Amazon Redshift pricing. You manage which queries are sent to the concurrency scaling cluster by configuring WLM queues. When you enable concurrency scaling for a queue, eligible queries are sent to the concurrency scaling cluster instead of waiting in line."

QUESTION 2

A company is running an on-premises application comprised of a web tier, an application tier, and a MySQL database tier. The database is used primarily during business hours with random activity peaks throughout the day. A database specialist needs to improve the availability and reduce the cost of the MySQL database tier as part of the company's migration to AWS.

Which MySQL database option would meet these requirements?

- A. Amazon RDS for MySQL with Multi-AZ
- B. Amazon Aurora Serverless MySQL cluster
- C. Amazon Aurora MySQL cluster
- D. Amazon RDS for MySQL with read replica



Correct Answer: C

QUESTION 3

The Development team recently executed a database script containing several data definition language (DDL) and data manipulation language (DML) statements on an Amazon Aurora MySQL DB cluster. The release accidentally deleted thousands of rows from an important table and broke some application functionality. This was discovered 4 hours after the release. Upon investigation, a Database Specialist tracked the issue to a DELETE command in the script with an incorrect WHERE clause filtering the wrong set of rows.

The Aurora DB cluster has Backtrack enabled with an 8-hour backtrack window. The Database Administrator also took a manual snapshot of the DB cluster before the release started. The database needs to be returned to the correct state as quickly as possible to resume full application functionality. Data loss must be minimal. How can the Database Specialist accomplish this?

- A. Quickly rewind the DB cluster to a point in time before the release using Backtrack.
- B. Perform a point-in-time recovery (PITR) of the DB cluster to a time before the release and copy the deleted rows from the restored database to the original database.
- C. Restore the DB cluster using the manual backup snapshot created before the release and change the application configuration settings to point to the new DB cluster.
- D. Create a clone of the DB cluster with Backtrack enabled. Rewind the cloned cluster to a point in time before the release. Copy deleted rows from the clone to the original database.

Correct Answer: A

QUESTION 4

A database specialist was alerted that a production Amazon RDS MariaDB instance with 100 GB of storage was out of space. In response, the database specialist modified the DB instance and added 50 GB of storage capacity. Three hours later, a new alert is generated due to a lack of free space on the same DB instance. The database specialist decides to modify the instance immediately to increase its storage capacity by 20 GB.

What will happen when the modification is submitted?

- A. The request will fail because this storage capacity is too large.
- B. The request will succeed only if the primary instance is in active status.
- C. The request will succeed only if CPU utilization is less than 10%.
- D. The request will fail as the most recent modification was too soon.

Correct Answer: D

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_PIOPS.StorageTypes.html

QUESTION 5



A user has a non-relational key-value database. The user is looking for a fully managed AWS service that will offload the administrative burdens of operating and scaling distributed databases. The solution must be cost-effective and able to handle unpredictable application traffic.

What should a Database Specialist recommend for this user?

- A. Create an Amazon DynamoDB table with provisioned capacity mode
- B. Create an Amazon DocumentDB cluster
- C. Create an Amazon DynamoDB table with on-demand capacity mode
- D. Create an Amazon Aurora Serverless DB cluster

Correct Answer: C

Reference: <https://aws.amazon.com/dynamodb/>

Key-value database -> DynamoDB Capable of dealing with unexpected application traffic -> on-demand capacity mode

A key-value database is a type of nonrelational database that uses a simple key-value method to store data. A key-value database stores data as a collection of key-value pairs in which a key serves as a unique identifier. On-demand mode is a good option to create new tables with unknown workloads.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.html#HowItWorks.OnDemand>

[DBS-C01 Study Guide](#)

[DBS-C01 Exam Questions](#)

[DBS-C01 Braindumps](#)