



# DBS-C01<sup>Q&As</sup>

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.geekcert.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

An ecommerce company has tasked a Database Specialist with creating a reporting dashboard that visualizes critical business metrics that will be pulled from the core production database running on Amazon Aurora. Data that is read by the dashboard should be available within 100 milliseconds of an update.

The Database Specialist needs to review the current configuration of the Aurora DB cluster and develop a cost-effective solution. The solution needs to accommodate the unpredictable read workload from the reporting dashboard without any impact on the write availability and performance of the DB cluster.

Which solution meets these requirements?

- A. Turn on the serverless option in the DB cluster so it can automatically scale based on demand.
- B. Provision a clone of the existing DB cluster for the new Application team.
- C. Create a separate DB cluster for the new workload, refresh from the source DB cluster, and set up ongoing replication using AWS DMS change data capture (CDC).
- D. Add an automatic scaling policy to the DB cluster to add Aurora Replicas to the cluster based on CPU consumption.

Correct Answer: A

---

### QUESTION 2

A company wants to improve its ecommerce website on AWS. A database specialist decides to add Amazon ElastiCache for Redis in the implementation stack to ease the workload off the database and shorten the website response times. The database specialist must also ensure the ecommerce website is highly available within the company's AWS Region.

How should the database specialist deploy ElastiCache to meet this requirement?

- A. Launch an ElastiCache for Redis cluster using the AWS CLI with the `-cluster-enabled` switch.
- B. Launch an ElastiCache for Redis cluster and select read replicas in different Availability Zones.
- C. Launch two ElastiCache for Redis clusters in two different Availability Zones. Configure Redis streams to replicate the cache from the primary cluster to another.
- D. Launch an ElastiCache cluster in the primary Availability Zone and restore the cluster's snapshot to a different Availability Zone during disaster recovery.

Correct Answer: B

Explanation: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/AutoFailover.html>

You can enable Multi-AZ only on Redis (cluster mode disabled) clusters that have at least one available read replica. Clusters without read replicas do not provide high availability or fault tolerance.

---



### QUESTION 3

A company hosts a 2 TB Oracle database in its on-premises data center. A database specialist is migrating the database from on premises to an Amazon Aurora PostgreSQL database on AWS.

The database specialist identifies a problem that relates to compatibility Oracle stores metadata in its data dictionary in uppercase, but PostgreSQL stores the metadata in lowercase. The database specialist must resolve this problem to complete the migration.

What is the MOST operationally efficient solution that meets these requirements?

- A. Override the default uppercase format of Oracle schema by encasing object names in quotation marks during creation.
- B. Use AWS Database Migration Service (AWS DMS) mapping rules with rule-action as convert-lowercase.
- C. Use the AWS Schema Conversion Tool conversion agent to convert the metadata from uppercase to lowercase.
- D. Use an AWS Glue job that is attached to an AWS Database Migration Service (AWS DMS) replication task to convert the metadata from uppercase to lowercase.

Correct Answer: B

Explanation: <https://aws.amazon.com/premiumsupport/knowledge-center/dms-mapping-oracle-postgresql/>

### QUESTION 4

A financial institution uses AWS to host its online application. Amazon RDS for MySQL is used to host the application's database, which includes automatic backups.

The program has corrupted the database logically, resulting in the application being unresponsive. The exact moment the corruption occurred has been determined, and it occurred within the backup retention period.

How should a database professional restore a database to its previous state prior to corruption?

- A. Use the point-in-time restore capability to restore the DB instance to the specified time. No changes to the application connection string are required.
- B. Use the point-in-time restore capability to restore the DB instance to the specified time. Change the application connection string to the new, restored DB instance.
- C. Restore using the latest automated backup. Change the application connection string to the new, restored DB instance.
- D. Restore using the appropriate automated backup. No changes to the application connection string are required.

Correct Answer: B

Explanation: When you perform a restore operation to a point in time or from a DB Snapshot, a new DB Instance is created with a new endpoint (the old DB Instance can be deleted if so desired). This is done to enable you to create multiple DB Instances from a specific DB Snapshot or point in time."

**QUESTION 5**

A company just migrated to Amazon Aurora PostgreSQL from an on-premises Oracle database. After the migration, the company discovered there is a period of time every day around 3:00 PM where the response time of the application is noticeably slower. The company has narrowed down the cause of this issue to the database and not the application.

Which set of steps should the Database Specialist take to most efficiently find the problematic PostgreSQL query?

- A. Create an Amazon CloudWatch dashboard to show the number of connections, CPU usage, and disk space consumption. Watch these dashboards during the next slow period.
- B. Launch an Amazon EC2 instance, and install and configure an open-source PostgreSQL monitoring tool that will run reports based on the output error logs.
- C. Modify the logging database parameter to log all the queries related to locking in the database and then check the logs after the next slow period for this information.
- D. Enable Amazon RDS Performance Insights on the PostgreSQL database. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

Correct Answer: D

[DBS-C01 VCE Dumps](#)

[DBS-C01 Practice Test](#)

[DBS-C01 Braindumps](#)