



# 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

A company runs online transaction processing (OLTP) workloads on an Amazon RDS for PostgreSQL Multi-AZ DB instance. Tests were run on the database after work hours, which generated additional database logs. The free storage of the RDS DB instance is low due to these additional logs.

What should the company do to address this space constraint issue?

- A. Log in to the host and run the `rm $PGDATA/pg_logs/*` command
- B. Modify the `rds.log_retention_period` parameter to 1440 and wait up to 24 hours for database logs to be deleted
- C. Create a ticket with AWS Support to have the logs deleted
- D. Run the `SELECT rds_rotate_error_log()` stored procedure to rotate the logs

Correct Answer: B

Explanation: To set the retention period for system logs, use the `rds.log_retention_period` parameter. You can find `rds.log_retention_period` in the DB parameter group associated with your DB instance. The unit for this parameter is minutes.

For example, a setting of 1,440 retains logs for one day. The default value is 4,320 (three days). The maximum value is 10,080 (seven days).

[https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER\\_LogAccess.Concepts.PostgreSQL.html](https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_LogAccess.Concepts.PostgreSQL.html)

## QUESTION 2

A gaming company is evaluating Amazon ElastiCache as a solution to manage player leaderboards. Millions of players around the world will compete in annual tournaments. The company wants to implement an architecture that is highly available. The company also wants to ensure that maintenance activities have minimal impact on the availability of the gaming platform.

Which combination of steps should the company take to meet these requirements? (Choose two.)

- A. Deploy an ElastiCache for Redis cluster with read replicas and Multi-AZ enabled.
- B. Deploy an ElastiCache for Memcached global datastore.
- C. Deploy a single-node ElastiCache for Redis cluster with automatic backups enabled. In the event of a failure, create a new cluster and restore data from the most recent backup.
- D. Use the default maintenance window to apply any required system changes and mandatory updates as soon as they are available.
- E. Choose a preferred maintenance window at the time of lowest usage to apply any required changes and mandatory updates.

Correct Answer: AE

Explanation: <https://aws.amazon.com/blogs/database/configuring-amazon-elasticache-for-redis-for-higher-availability/>



### QUESTION 3

A company migrated an on-premises Oracle database to Amazon RDS for Oracle. A database specialist needs to monitor the latency of the database.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Publish RDS Performance insights metrics to Amazon CloudWatch. Add AWS CloudTrail filters to monitor database performance
- B. Install Oracle Statspack. Enable the performance statistics feature to collect, store, and display performance data to monitor database performance.
- C. Enable RDS Performance Insights to visualize the database load. Enable Enhanced Monitoring to view how different threads use the CPU
- D. Create a new DB parameter group that includes the AllocatedStorage, DBInstanceClassMemory, and DBInstanceVCPU variables. Enable RDS Performance Insights

Correct Answer: C

Explanation:

### QUESTION 4

A company is moving its fraud detection application from on premises to the AWS Cloud and is using Amazon Neptune for data storage. The company has set up a 1 Gbps AWS Direct Connect connection to migrate 25 TB of fraud detection data from the on-premises data center to a Neptune DB instance. The company already has an Amazon S3 bucket and an S3 VPC endpoint, and 80% of the company's network bandwidth is available.

How should the company perform this data load?

- A. Use an AWS SDK with a multipart upload to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- B. Use AWS Database Migration Service (AWS DMS) to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- C. Use AWS DataSync to transfer the data from on premises to the S3 bucket. Use the Loader command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.
- D. Use the AWS CLI to transfer the data from on premises to the S3 bucket. Use the Copy command for Neptune to move the data in bulk from the S3 bucket to the Neptune DB instance.

Correct Answer: C

Explanation: "AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services."  
<https://docs.aws.amazon.com/neptune/latest/userguide/bulk-load.html>

### QUESTION 5



A software-as-a-service (SaaS) company is using an Amazon Aurora Serverless DB cluster for its production MySQL database. The DB cluster has general logs and slow query logs enabled. A database engineer must use the most operationally efficient solution with minimal resource utilization to retain the logs and facilitate interactive search and analysis.

Which solution meets these requirements?

- A. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- B. Download the logs from the DB cluster and store them in Amazon S3 by using manual scripts. Use Amazon Athena and Amazon QuickSight to search and analyze the logs.
- C. Use an AWS Lambda function to ship database logs to an Amazon S3 bucket. Use Amazon Elasticsearch Service (Amazon ES) and Kibana to search and analyze the logs.
- D. Use Amazon CloudWatch Logs Insights to search and analyze the logs when the logs are automatically uploaded by the DB cluster.

Correct Answer: D

Explanation: <https://aws.amazon.com/premiumsupport/knowledge-center/aurora-serverless-logs-enable-view/>  
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/AnalyzingLogData.html>

[DBS-C01 VCE Dumps](#)

[DBS-C01 Study Guide](#)

[DBS-C01 Braindumps](#)