

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 company is looking to migrate a 1 TB Oracle database from on-premises to an Amazon Aurora PostgreSQL DB cluster. The company\\'s Database Specialist discovered that the Oracle database is storing 100 GB of large binary objects (LOBs) across multiple tables. The Oracle database has a maximum LOB size of 500 MB with an average LOB size of 350 MB. The Database Specialist has chosen AWS DMS to migrate the data with the largest replication instances. How should the Database Specialist optimize the database migration using AWS DMS?

A. Create a single task using full LOB mode with a LOB chunk size of 500 MB to migrate the data and LOBs together

B. Create two tasks: task1 with LOB tables using full LOB mode with a LOB chunk size of 500 MB and task2 without LOBs

C. Create two tasks: task1 with LOB tables using limited LOB mode with a maximum LOB size of 500 MB and task 2 without LOBs

D. Create a single task using limited LOB mode with a maximum LOB size of 500 MB to migrate data and LOBs together

Correct Answer: C

QUESTION 2

A business\\'s production database is hosted on a single-node Amazon RDS for MySQL DB instance. The database instance is hosted in a United States AWS Region.

A week before a significant sales event, a fresh database maintenance update is released. The maintenance update has been designated as necessary. The firm want to minimize the database instance\\'s downtime and requests that a

database expert make the database instance highly accessible until the sales event concludes.

Which solution will satisfy these criteria?

A. Defer the maintenance update until the sales event is over.

B. Create a read replica with the latest update. Initiate a failover before the sales event.

C. Create a read replica with the latest update. Transfer all read-only traffic to the read replica during the sales event.

D. Convert the DB instance into a Multi-AZ deployment. Apply the maintenance update.

Correct Answer: D

Explanation: https://aws.amazon.com/premiumsupport/knowledge-center/rds-required- maintenance/

QUESTION 3

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

QUESTION 4

A bike rental company operates an application to track its bikes. The application receives location and condition data from bike sensors. The application also receives rental transaction data from the associated mobile app.

The application uses Amazon DynamoDB as its database layer. The company has configured DynamoDB with provisioned capacity set to 20% above the expected peak load of the application. On an average day, DynamoDB used 22 billion read capacity units (RCUs) and 60 billion write capacity units (WCUs). The application is running well. Usage changes smoothly over the course of the day and is generally shaped like a bell curve. The timing and magnitude of peaks vary based on the weather and season, but the general shape is consistent.

Which solution will provide the MOST cost optimization of the DynamoDB database layer?

- A. Change the DynamoDB tables to use on-demand capacity.
- B. Use AWS Auto Scaling and configure time-based scaling.

C. Enable DynamoDB capacity-based auto scaling.

D. Enable DynamoDB Accelerator (DAX).

Correct Answer: C

https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

QUESTION 5

A company is migrating its on-premises database workloads to the AWS Cloud. A database specialist performing the move has chosen AWS DMS to migrate an Oracle database with a large table to Amazon RDS. The database specialist notices that AWS DMS is taking significant time to migrate the data. Which actions would improve the data migration speed? (Choose three.)

A. Create multiple AWS DMS tasks to migrate the large table.

- B. Configure the AWS DMS replication instance with Multi-AZ.
- C. Increase the capacity of the AWS DMS replication server.
- D. Establish an AWS Direct Connect connection between the on-premises data center and AWS.
- E. Enable an Amazon RDS Multi-AZ configuration.
- F. Enable full large binary object (LOB) mode to migrate all LOB data for all large tables.

Correct Answer: CDE

QUESTION 6

A company is using an Amazon Aurora PostgreSQL DB cluster with an xlarge primary instance master and two large Aurora Replicas for high availability and read-only workload scaling. A failover event occurs and application performance is poor for several minutes. During this time, application servers in all Availability Zones are healthy and responding normally.

What should the company do to eliminate this application performance issue?

A. Configure both of the Aurora Replicas to the same instance class as the primary DB instance. Enable cache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign a failover priority of tier-1 to the replicas.

B. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instance has failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be the primary DB instance. Configure an Amazon RDS event subscription to send a notification to an Amazon SNS topic to which the Lambda function is subscribed.

C. Configure one Aurora Replica to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and one replica with the same instance class. Set the failover priority to tier-1 for the other replicas.

D. Configure both Aurora Replicas to have the same instance class as the primary DB instance. Implement Aurora PostgreSQL DB cluster cache management. Set the failover priority to tier-0 for the primary DB instance and to tier-1 for the replicas.

Correct Answer: C

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.clu ster-cache-mgmt.html

https://aws.amazon.com/blogs/database/introduction-to-aurora-postgresql-cluster-cache-management/

"You can customize the order in which your Aurora Replicas are promoted to the primary instance after a failure by assigning each replica a priority. Priorities range from 0 for the first priority to 15 for the last priority. If the primary instance

fails, Amazon RDS promotes the Aurora Replica with the better priority to the new primary instance. You can modify the priority of an Aurora Replica at any time. Modifying the priority doesn\\'t trigger a failover. More than one Aurora Replica

can share the same priority, resulting in promotion tiers. If two or more Aurora Replicas share the same priority, then Amazon RDS promotes the replica that is largest in size. If two or more Aurora Replicas share the same priority and size,



then Amazon RDS promotes an arbitrary replica in the same promotion tier. "

Amazon Aurora with PostgreSQL compatibility now supports cluster cache management, providing a faster path to full performance if there\\'s a failover. With cluster cache management, you designate a specific reader DB instance in your

Aurora PostgreSQL cluster as the failover target. Cluster cache management keeps the data in the designated reader\\'s cache synchronized with the data in the read-write instance\\'s cache. If a failover occurs, the designated reader is

promoted to be the new read-write instance, and workloads benefit immediately from the data in its cache.

QUESTION 7

A company with 500,000 employees needs to supply its employee list to an application used by human resources. Every 30 minutes, the data is exported using the LDAP service to load into a new Amazon DynamoDB table. The data model has a base table with Employee ID for the partition key and a global secondary index with Organization ID as the partition key.

While importing the data, a database specialist receives ProvisionedThroughputExceededException errors. After increasing the provisioned write capacity units (WCUs) to 50,000, the specialist receives the same errors. Amazon CloudWatch metrics show a consumption of 1,500 WCUs.

What should the database specialist do to address the issue?

- A. Change the data model to avoid hot partitions in the global secondary index.
- B. Enable auto scaling for the table to automatically increase write capacity during bulk imports.
- C. Modify the table to use on-demand capacity instead of provisioned capacity.
- D. Increase the number of retries on the bulk loading application.

Correct Answer: A

Explanation: https://aws.amazon.com/premiumsupport/knowledge-center/dynamodb- table-throttled/

QUESTION 8

An AWS CloudFormation stack that included an Amazon RDS DB instance was accidentally deleted and recent data was lost. A Database Specialist needs to add RDS settings to the CloudFormation template to reduce the chance of accidental instance data loss in the future.

Which settings will meet this requirement? (Choose three.)

- A. Set DeletionProtection to True
- B. Set MultiAZ to True
- C. Set TerminationProtection to True
- D. Set DeleteAutomatedBackups to False
- E. Set DeletionPolicy to Delete



F. Set DeletionPolicy to Retain

Correct Answer: ACF

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws- attribute- deletionpolicy.html https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-accidental- updates/

QUESTION 9

A company is running an Amazon RDS for PostgeSQL DB instance and wants to migrate it to an Amazon Aurora PostgreSQL DB cluster. The current database is 1 TB in size. The migration needs to have minimal downtime. What is the FASTEST way to accomplish this?

A. Create an Aurora PostgreSQL DB cluster. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.

B. Use the pg_dump and pg_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.

C. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.

D. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replica. Promote the replica during the cutover.

Correct Answer: D

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Mi grating.html Migrating data from an RDS PostgreSQL DB instance to an Aurora PostgreSQL DB cluster by using an Aurora read replica. https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Mi grating.html#AuroraPostgreSQL.Migrating.RDSPostgreSQ https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraPostgreSQL.Mi grating.html#AuroraPostgreSQL.Migrating.RDSPostgreSQL.Replica

QUESTION 10

A company has applications running on Amazon EC2 instances in a private subnet with no internet connectivity. The company deployed a new application that uses Amazon DynamoDB, but the application cannot connect to the DynamoDB tables. A developer already checked that all permissions are set correctly.

What should a database specialist do to resolve this issue while minimizing access to external resources?

A. Add a route to an internet gateway in the subnet///s route table.

- B. Add a route to a NAT gateway in the subnet///s route table.
- C. Assign a new security group to the EC2 instances with an outbound rule to ports 80 and 443.
- D. Create a VPC endpoint for DynamoDB and add a route to the endpoint in the subnet///s route table.

Correct Answer: D

Explanation: https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/vpc- endpoints-dynamodb.html



QUESTION 11

A company is developing an application that performs intensive in-memory operations on advanced data structures such as sorted sets. The application requires sub-millisecond latency for reads and writes. The application occasionally must run a group of commands as an ACID-compliant operation. A database specialist is setting up the database for this application. The database specialist needs the ability to create a new database cluster from the latest backup of the production cluster.

Which type of cluster should the database specialist create to meet these requirements?

- A. Amazon ElastiCache for Memcached
- B. Amazon Neptune
- C. Amazon ElastiCache for Redis
- D. Amazon DynamoDB Accelerator (DAX)
- Correct Answer: C

Explanation: https://aws.amazon.com/elasticache/redis-vs-memcached/ https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/elasticache-use- cases.html #elasticache-for-redis-usecases-gaming

QUESTION 12

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 13



A web-based application uses Amazon DocumentDB (with MongoDB compatibility) as its underlying data store. Sufficient access control IS in place, but a database specialist wants to be able to review logs if the primary DocumentDB database is deleted.

Which combination of steps Should the database specialist take to meet this requirement? (Select TWO_)

- A. Set the audit_logs cluster parameter to enabled
- B. Enable DocumentDB log export to Amazon CloudWatch Logs.
- C. Enable Enhanced Monitoring tor DocumentDB.
- D. Enable AWS CloudTrail for DocumentDB.
- E. use AWS Config to monitor the state of DocumentDB.

Correct Answer: AB

Option A is correct because it sets the audit_logs cluster parameter to enabled. This enables auditing on the DocumentDB cluster, which records events that were performed in the cluster, such as successful and failed authentication

attempts, dropping a collection in a database, or creating an index1. By enabling auditing, the database specialist can review the logs to see who and when deleted the primary DocumentDB database, and what other actions were taken on

the cluster.

Option B is correct because it enables DocumentDB log export to Amazon CloudWatch Logs. This allows the DocumentDB cluster to export its auditing records (JSON documents) to Amazon CloudWatch Logs, where they can be analyzed,

monitored, and archived1. By enabling log export, the database specialist can access the logs even if the primary DocumentDB database is deleted, as they are stored in a separate service.

QUESTION 14

A database specialist is managing an application in the us-west-1 Region and wants to set up disaster recovery in the us-east-1 Region. The Amazon Aurora MySQL DB cluster needs an RPO of 1 minute and an RTO of 2 minutes. Which approach meets these requirements with no negative performance impact?

- A. Enable synchronous replication.
- B. Enable asynchronous binlog replication.
- C. Create an Aurora Global Database.

D. Copy Aurora incremental snapshots to the us-east-1 Region.

Correct Answer: C

Explanation: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora- global-database-disaster-recovery.html



QUESTION 15

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

Latest DBS-C01 Dumps

DBS-C01 Study Guide

DBS-C01 Exam Questions