



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.lead4pass.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 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 2

A Database Specialist is planning to create a read replica of an existing Amazon RDS for MySQL Multi-AZ DB instance. When using the AWS Management Console to conduct this task, the Database Specialist discovers that the source RDS

DB instance does not appear in the read replica source selection box, so the read replica cannot be created.

What is the most likely reason for this?

- A. The source DB instance has to be converted to Single-AZ first to create a read replica from it.
- B. Enhanced Monitoring is not enabled on the source DB instance.
- C. The minor MySQL version in the source DB instance does not support read replicas.
- D. Automated backups are not enabled on the source DB instance.

Correct Answer: D

Reference: <https://aws.amazon.com/rds/features/read-replicas/>

QUESTION 3

A large company has a variety of Amazon DB clusters. Each of these clusters has various configurations that adhere to various requirements. Depending on the team and use case, these configurations can be organized into broader categories.

A database administrator wants to make the process of storing and modifying these parameters more systematic. The



database administrator also wants to ensure that changes to individual categories of configurations are automatically applied to all instances when required.

Which AWS service or feature will help automate and achieve this objective?

- A. AWS Systems Manager Parameter Store
- B. DB parameter group
- C. AWS Config
- D. AWS Secrets Manager

Correct Answer: B

Reference: https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/USER_WorkingWithParamGroups.html

QUESTION 4

A company maintains several databases using Amazon RDS for MySQL and PostgreSQL. Each RDS database generates log files with retention periods set to their default values. The company has now mandated that database logs be maintained for up to 90 days in a centralized repository to facilitate real-time and after-the-fact analyses.

What should a Database Specialist do to meet these requirements with minimal effort?

- A. Create an AWS Lambda function to pull logs from the RDS databases and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- B. Modify the RDS databases to publish log to Amazon CloudWatch Logs. Change the log retention policy for each log group to expire the events after 90 days.
- C. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucket. Set a lifecycle policy to expire the objects after 90 days.
- D. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Logs. Change the log retention policy for the log group to expire the events after 90 days.

Correct Answer: A

QUESTION 5

A ride-hailing application uses an Amazon RDS for MySQL DB instance as persistent storage for bookings. This application is very popular and the company expects a tenfold increase in the user base in next few months. The application experiences more traffic during the morning and evening hours.

This application has two parts:

1.

An in-house booking component that accepts online bookings that directly correspond to simultaneous requests from users.



2.

A third-party customer relationship management (CRM) component used by customer care representatives. The CRM uses queries to access booking data.

A database specialist needs to design a cost-effective database solution to handle this workload.

Which solution meets these requirements?

A. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to the RDS for MySQL DB instance used by the CRM.

B. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to an Amazon SQS queue. This triggers another Lambda function that pulls data from Amazon SQS and writes it to the RDS for MySQL DB instance used by the CRM.

C. Use Amazon ElastiCache for Redis to accept the bookings. Associate an AWS Lambda function to capture changes and push the booking data to an Amazon Redshift database used by the CRM.

D. Use Amazon DynamoDB to accept the bookings. Enable DynamoDB Streams and associate an AWS Lambda function to capture changes and push the booking data to Amazon Athena, which is used by the CRM.

Correct Answer: A

QUESTION 6

A company is running Amazon RDS for MySQL for its workloads. There is downtime when AWS operating system patches are applied during the Amazon RDS-specified maintenance window.

What is the MOST cost-effective action that should be taken to avoid downtime?

A. Migrate the workloads from Amazon RDS for MySQL to Amazon DynamoDB

B. Enable cross-Region read replicas and direct read traffic to them when Amazon RDS is down

C. Enable a read replicas and direct read traffic to it when Amazon RDS is down

D. Enable an Amazon RDS for MySQL Multi-AZ configuration

Correct Answer: C

QUESTION 7

A company is running its line of business application on AWS, which uses Amazon RDS for MySQL at the persistent data store. The company wants to minimize downtime when it migrates the database to Amazon Aurora. Which migration method should a Database Specialist use?

A. Take a snapshot of the RDS for MySQL DB instance and create a new Aurora DB cluster with the option to migrate snapshots.

B. Make a backup of the RDS for MySQL DB instance using the mysqldump utility, create a new Aurora DB cluster, and restore the backup.



- C. Create an Aurora Replica from the RDS for MySQL DB instance and promote the Aurora DB cluster.
- D. Create a clone of the RDS for MySQL DB instance and promote the Aurora DB cluster.

Correct Answer: A

Reference: <https://d1.awsstatic.com/whitepapers/RDS/Migrating%20your%20databases%20to%20Amazon%20Aurora.pdf> (10)

QUESTION 8

A Database Specialist is working with a company to launch a new website built on Amazon Aurora with several Aurora Replicas. This new website will replace an on-premises website connected to a legacy relational database. Due to stability

issues in the legacy database, the company would like to test the resiliency of Aurora.

Which action can the Database Specialist take to test the resiliency of the Aurora DB cluster?

- A. Stop the DB cluster and analyze how the website responds
- B. Use Aurora fault injection to crash the master DB instance
- C. Remove the DB cluster endpoint to simulate a master DB instance failure
- D. Use Aurora Backtrack to crash the DB cluster

Correct Answer: B

Reference: <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Managing.FaultInjectionQueries.html>

QUESTION 9

A database specialist at a large multi-national financial company is in charge of designing the disaster recovery strategy for a highly available application that is in development. The application uses an Amazon DynamoDB table as its data

store. The application requires a recovery time objective (RTO) of 1 minute and a recovery point objective (RPO) of 2 minutes.

Which operationally efficient disaster recovery strategy should the database specialist recommend for the DynamoDB table?

- A. Create a DynamoDB stream that is processed by an AWS Lambda function that copies the data to a DynamoDB table in another Region.
- B. Use a DynamoDB global table replica in another Region. Enable point-in-time recovery for both tables.
- C. Use a DynamoDB Accelerator table in another Region. Enable point-in-time recovery for the table.
- D. Create an AWS Backup plan and assign the DynamoDB table as a resource.

Correct Answer: C



QUESTION 10

A database specialist is building a system that uses a static vendor dataset of postal codes and related territory information that is less than 1 GB in size. The dataset is loaded into the application's cache at start up. The company needs to store this data in a way that provides the lowest cost with a low application startup time.

Which approach will meet these requirements?

- A. Use an Amazon RDS DB instance. Shut down the instance once the data has been read.
- B. Use Amazon Aurora Serverless. Allow the service to spin resources up and down, as needed.
- C. Use Amazon DynamoDB in on-demand capacity mode.
- D. Use Amazon S3 and load the data from flat files.

Correct Answer: A

QUESTION 11

A company uses the Amazon DynamoDB table contractDB in us-east-1 for its contract system with the following schema:

1.
orderID (primary key)
2.
timestamp (sort key)
3.
contract (map)
4.
createdBy (string)
5.
customerEmail (string)

After a problem in production, the operations team has asked a database specialist to provide an IAM policy to read items from the database to debug the application. In addition, the developer is not allowed to access the value of the customerEmail field to stay compliant.

Which IAM policy should the database specialist use to achieve these requirements?



```
A. {
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Allow",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderID",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```



```
B. {  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "IAMPolicy",  
      "Effect": "Allow",  
      "Action": [  
        "dynamodb: Query"  
      ],  
      "Resource": [  
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"  
      ],  
      "Condition": {  
        "ForAllValues:StringLike": {  
          "dynamodb:Attributes": [  
            "customerEmail"  
          ]  
        },  
        "StringEquals": {  
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"  
        }  
      }  
    }  
  ]  
}
```




```
C. {
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb: Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "customerEmail"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```



```
D. {
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "IAMPolicy",
      "Effect": "Deny",
      "Action": [
        "dynamodb:Query"
      ],
      "Resource": [
        "arn:aws:dynamodb:us-east-1:123456789012:table/contractDB"
      ],
      "Condition": {
        "ForAllValues:StringLike": {
          "dynamodb:Attributes": [
            "orderID",
            "timestamp",
            "contract",
            "createdBy"
          ]
        },
        "StringEquals": {
          "dynamodb:Select": "SPECIFIC_ATTRIBUTES"
        }
      }
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: A

QUESTION 12



An electric utility company wants to store power plant sensor data in an Amazon DynamoDB table. The utility company has over 100 power plants and each power plant has over 200 sensors that send data every 2 seconds. The sensor data includes time with milliseconds precision, a value, and a fault attribute if the sensor is malfunctioning. Power plants are identified by a globally unique identifier. Sensors are identified by a unique identifier within each power plant. A database specialist needs to design the table to support an efficient method of finding all faulty sensors within a given power plant.

Which schema should the database specialist use when creating the DynamoDB table to achieve the fastest query time when looking for faulty sensors?

- A. Use the plant identifier as the partition key and the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- B. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a local secondary index (LSI) on the fault attribute.
- C. Create a composite of the plant identifier and sensor identifier as the partition key. Use the measurement time as the sort key. Create a global secondary index (GSI) with the plant identifier as the partition key and the fault attribute as the sort key.
- D. Use the plant identifier as the partition key and the sensor identifier as the sort key. Create a local secondary index (LSI) on the fault attribute.

Correct Answer: B

[Latest DBS-C01 Dumps](#)

[DBS-C01 VCE Dumps](#)

[DBS-C01 Practice Test](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

- 100% Guaranteed Success
- 100% Money Back Guarantee
- 365 Days Free Update
- Instant Download After Purchase
- 24x7 Customer Support
- Average 99.9% Success Rate
- More than 800,000 Satisfied Customers Worldwide
- Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.lead4pass.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket:



	<p>One Year Free Update</p> <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>		<p>Money Back Guarantee</p> <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>		<p>Security & Privacy</p> <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
--	--	--	--	--	--

Any charges made through this site will appear as Global Simulators Limited.
All trademarks are the property of their respective owners.
Copyright © lead4pass, All Rights Reserved.