



# Amazon-Web-Services

## Exam Questions DBS-C01

AWS Certified Database - Specialty

#### NEW 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.

**Answer: A**

#### NEW QUESTION 2

A clothing company uses a custom ecommerce application and a PostgreSQL database to sell clothes to thousands of users from multiple countries. The company is migrating its application and database from its on premises data center to the AWS Cloud. The company has selected Amazon EC2 for the application and Amazon RDS for PostgreSQL for the database. The company requires database passwords to be changed every 60 days. A Database Specialist needs to ensure that the credentials used by the web application to connect to the database are managed securely. Which approach should the Database Specialist take to securely manage the database credentials?

- A. Store the credentials in a text file in an Amazon S3 bucket
- B. Restrict permissions on the bucket to the IAM role associated with the instance profile on the EC2 instance
- C. Modify the application to download the text file and retrieve the credentials on start up
- D. Update the text file every 60 days.
- E. Configure IAM database authentication for the application to connect to the database
- F. Create an IAM user and map it to a separate database user for each ecommerce user
- G. Require users to update their passwords every 60 days.
- H. Store the credentials in AWS Secrets Manager
- I. Restrict permissions on the secret to only the IAM role associated with the instance profile on the EC2 instance
- J. Modify the application to retrieve the credentials from Secrets Manager on start up
- K. Configure the rotation interval to 60 days.
- L. Store the credentials in an encrypted text file in the application AMI
- M. Use AWS KMS to store the key for decrypting the text file
- N. Modify the application to decrypt the text file and retrieve the credentials on start up
- O. Update the text file and publish a new AMI every 60 days.

**Answer: B**

#### NEW QUESTION 3

A team of Database Specialists is currently investigating performance issues on an Amazon RDS for MySQL DB instance and is reviewing related metrics. The team wants to narrow the possibilities down to specific database wait events to better understand the situation. How can the Database Specialists accomplish this?

- A. Enable the option to push all database logs to Amazon CloudWatch for advanced analysis
- B. Create appropriate Amazon CloudWatch dashboards to contain specific periods of time
- C. Enable Amazon RDS Performance Insights and review the appropriate dashboard
- D. Enable Enhanced Monitoring with the appropriate settings

**Answer: C**

#### NEW QUESTION 4

A company is building a new web platform where user requests trigger an AWS Lambda function that performs an insert into an Amazon Aurora MySQL DB cluster. Initial tests with less than 10 users on the new platform yielded successful execution and fast response times. However, upon more extensive tests with the actual target of 3,000 concurrent users, Lambda functions are unable to connect to the DB cluster and receive too many connections errors. Which of the following will resolve this issue?

- A. Edit the my.cnf file for the DB cluster to increase max\_connections
- B. Increase the instance size of the DB cluster
- C. Change the DB cluster to Multi-AZ
- D. Increase the number of Aurora Replicas

**Answer: B**

#### NEW QUESTION 5

A company is hosting critical business data in an Amazon Redshift cluster. Due to the sensitive nature of the data, the cluster is encrypted at rest using AWS KMS. As a part of disaster recovery requirements, the company needs to copy the Amazon Redshift snapshots to another Region. Which steps should be taken in the AWS Management Console to meet the disaster recovery requirements?

- A. Create a new KMS customer master key in the source Region
- B. Switch to the destination Region, enable Amazon Redshift cross-Region snapshots, and use the KMS key of the source Region.
- C. Create a new IAM role with access to the KMS key in the source Region
- D. Enable Amazon Redshift cross-Region replication using the new IAM role, and use the KMS key of the source Region.
- E. Enable Amazon Redshift cross-Region snapshots in the source Region, and create a snapshot copy grant and use a KMS key in the destination Region.
- F. Create a new KMS customer master key in the destination Region and create a new IAM role with access to the new KMS key
- G. Enable Amazon Redshift cross-Region replication in the source Region and use the KMS key of the destination Region.

Answer: A

#### NEW QUESTION 6

A Database Specialist is migrating a 2 TB Amazon RDS for Oracle DB instance to an RDS for PostgreSQL DB instance using AWS DMS. The source RDS Oracle DB instance is in a VPC in the us-east-1 Region. The target RDS for PostgreSQL DB instance is in a VPC in the use-west-2 Region. Where should the AWS DMS replication instance be placed for the MOST optimal performance?

- A. In the same Region and VPC of the source DB instance
- B. In the same Region and VPC as the target DB instance
- C. In the same VPC and Availability Zone as the target DB instance
- D. In the same VPC and Availability Zone as the source DB instance

Answer: D

#### NEW QUESTION 7

A Database Specialist is migrating an on-premises Microsoft SQL Server application database to Amazon RDS for PostgreSQL using AWS DMS. The application requires minimal downtime when the RDS DB instance goes live. What change should the Database Specialist make to enable the migration?

- A. Configure the on-premises application database to act as a source for an AWS DMS full load with ongoing change data capture (CDC)
- B. Configure the AWS DMS replication instance to allow both full load and ongoing change data capture(CDC)
- C. Configure the AWS DMS task to generate full logs to allow for ongoing change data capture (CDC)
- D. Configure the AWS DMS connections to allow two-way communication to allow for ongoing change datacapture (CDC)

Answer: A

#### NEW QUESTION 8

A gaming company is designing a mobile gaming app that will be accessed by many users across the globe. The company wants to have replication and full support for multi-master writes. The company also wants to ensure low latency and consistent performance for app users. Which solution meets these requirements?

- A. Use Amazon DynamoDB global tables for storage and enable DynamoDB automatic scaling
- B. Use Amazon Aurora for storage and enable cross-Region Aurora Replicas
- C. Use Amazon Aurora for storage and cache the user content with Amazon ElastiCache
- D. Use Amazon Neptune for storage

Answer: A

#### NEW QUESTION 9

A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load data from Amazon S3 into the Neptune DB cluster using the Neptune bulk loader API. The Database Specialist receives the following error:

“Unable to connect to s3 endpoint. Provided source = s3://mybucket/graphdata/ and region = us-east-1. Please verify your S3 configuration.”

Which combination of actions should the Database Specialist take to troubleshoot the problem? (Choose two.)

- A. Check that Amazon S3 has an IAM role granting read access to Neptune
- B. Check that an Amazon S3 VPC endpoint exists
- C. Check that a Neptune VPC endpoint exists
- D. Check that Amazon EC2 has an IAM role granting read access to Amazon S3
- E. Check that Neptune has an IAM role granting read access to Amazon S3

Answer: BD

#### NEW QUESTION 10

A gaming company has recently acquired a successful iOS game, which is particularly popular during the holiday season. The company has decided to add a leaderboard to the game that uses Amazon DynamoDB. The application load is expected to ramp up over the holiday season.

Which solution will meet these requirements at the lowest cost?

- A. DynamoDB Streams
- B. DynamoDB with DynamoDB Accelerator
- C. DynamoDB with on-demand capacity mode
- D. DynamoDB with provisioned capacity mode with Auto Scaling

Answer: C

#### NEW QUESTION 10

An online gaming company is planning to launch a new game with Amazon DynamoDB as its data store. The database should be designated to support the following use cases:

Update scores in real time whenever a player is playing the game.

Retrieve a player's score details for a specific game session.

A Database Specialist decides to implement a DynamoDB table. Each player has a unique user\_id and each game has a unique game\_id.

Which choice of keys is recommended for the DynamoDB table?

- A. Create a global secondary index with game\_id as the partition key
- B. Create a global secondary index with user\_id as the partition key
- C. Create a composite primary key with game\_id as the partition key and user\_id as the sort key
- D. Create a composite primary key with user\_id as the partition key and game\_id as the sort key

**Answer: B**

#### NEW QUESTION 13

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 diskspace consumption
- B. Watch these dashboards during the next slow period.
- C. 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.
- D. 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.
- E. Enable Amazon RDS Performance Insights on the PostgreSQL databases
- F. Use the metrics to identify any queries that are related to spikes in the graph during the next slow period.

**Answer: D**

#### NEW QUESTION 18

A large financial services company requires that all data be encrypted in transit. A Developer is attempting to connect to an Amazon RDS DB instance using the company VPC for the first time with credentials provided by a Database Specialist. Other members of the Development team can connect, but this user is consistently receiving an error indicating a communications link failure. The Developer asked the Database Specialist to reset the password a number of times, but the error persists.

Which step should be taken to troubleshoot this issue?

- A. Ensure that the database option group for the RDS DB instance allows ingress from the Developer machine's IP address
- B. Ensure that the RDS DB instance's subnet group includes a public subnet to allow the Developer to connect
- C. Ensure that the RDS DB instance has not reached its maximum connections limit
- D. Ensure that the connection is using SSL and is addressing the port where the RDS DB instance is listening for encrypted connections

**Answer: B**

#### NEW QUESTION 20

A company developed an AWS CloudFormation template used to create all new Amazon DynamoDB tables in its AWS account. The template configures provisioned throughput capacity using hard-coded values. The company wants to change the template so that the tables it creates in the future have independently configurable read and write capacity units assigned.

Which solution will enable this change?

- A. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Configure DynamoDB to provision throughput capacity using the stack's mappings.
- B. Add values for two Number parameters, rcuCount and wcuCount, to the template
- C. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.
- D. Add values for the rcuCount and wcuCount parameters as outputs of the template
- E. Configure DynamoDB to provision throughput capacity using the stack outputs.
- F. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template. Replace the hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

**Answer: B**

#### NEW QUESTION 24

A marketing company is using Amazon DocumentDB and requires that database audit logs be enabled. A Database Specialist needs to configure monitoring so that all data definition language (DDL) statements performed are visible to the Administrator. The Database Specialist has set the audit\_logs parameter to enabled in the cluster parameter group.

What should the Database Specialist do to automatically collect the database logs for the Administrator?

- A. Enable DocumentDB to export the logs to Amazon CloudWatch Logs
- B. Enable DocumentDB to export the logs to AWS CloudTrail
- C. Enable DocumentDB Events to export the logs to Amazon CloudWatch Logs
- D. Configure an AWS Lambda function to download the logs using the download-db-log-file-portion operation and store the logs in Amazon S3

**Answer: A**

#### NEW QUESTION 27

A Database Specialist modified an existing parameter group currently associated with a production Amazon RDS for SQL Server Multi-AZ DB instance. The change is associated with a static parameter type, which controls the number of user connections allowed on the most critical RDS SQL Server DB instance for the company. This change has been approved for a specific maintenance window to help minimize the impact on users.

How should the Database Specialist apply the parameter group change for the DB instance?

- A. Select the option to apply the change immediately
- B. Allow the preconfigured RDS maintenance window for the given DB instance to control when the change is applied
- C. Apply the change manually by rebooting the DB instance during the approved maintenance window
- D. Reboot the secondary Multi-AZ DB instance

**Answer: D**

#### NEW QUESTION 29

A company is closing one of its remote data centers. This site runs a 100 TB on-premises data warehouse solution. The company plans to use the AWS Schema

Conversion Tool (AWS SCT) and AWS DMS for the migration to AWS. The site network bandwidth is 500 Mbps. A Database Specialist wants to migrate the on-premises data using Amazon S3 as the data lake and Amazon Redshift as the data warehouse. This move must take place during a 2-week period when source systems are shut down for maintenance. The data should stay encrypted at rest and in transit.

Which approach has the least risk and the highest likelihood of a successful data transfer?

- A. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- B. Leverage AWSSCT and apply the converted schema to Amazon Redshif
- C. Once complete, start an AWS DMS task to move the data from the source to Amazon S3. Use AWS Glue to load the data from Amazon S3 to Amazon Redshift.
- D. Leverage AWS SCT and apply the converted schema to Amazon Redshif
- E. Start an AWS DMS task with two AWS Snowball Edge devices to copy data from on-premises to Amazon S3 with AWS KMS encryption. Use AWS DMS to finish copying data to Amazon Redshift.
- F. Leverage AWS SCT and apply the converted schema to Amazon Redshif
- G. Once complete, use a fleet of 10 TB dedicated encrypted drives using the AWS Import/Export feature to copy data from on-premises to Amazon S3 with AWS KMS encryption
- H. Use AWS Glue to load the data to Amazon redshift.
- I. Set up a VPN tunnel for encrypting data over the network from the data center to AW
- J. Leverage a native database export feature to export the data and compress the file
- K. Use the aws S3 cp multi-port upload command to upload these files to Amazon S3 with AWS KMS encryption
- L. Once complete, load the data to Amazon Redshift using AWS Glue.

**Answer: C**

#### NEW QUESTION 31

A Database Specialist migrated an existing production MySQL database from on-premises to an Amazon RDS for MySQL DB instance. However, after the migration, the database needed to be encrypted at rest using AWS KMS. Due to the size of the database, reloading, the data into an encrypted database would be too time-consuming, so it is not an option.

How should the Database Specialist satisfy this new requirement?

- A. Create a snapshot of the unencrypted RDS DB instance
- B. Create an encrypted copy of the unencrypted snapshot
- C. Restore the encrypted snapshot copy.
- D. Modify the RDS DB instance
- E. Enable the AWS KMS encryption option that leverages the AWS CLI.
- F. Restore an unencrypted snapshot into a MySQL RDS DB instance that is encrypted.
- G. Create an encrypted read replica of the RDS DB instance
- H. Promote it the master.

**Answer: A**

#### NEW QUESTION 34

A company has an on-premises system that tracks various database operations that occur over the lifetime of a database, including database shutdown, deletion, creation, and backup.

The company recently moved two databases to Amazon RDS and is looking at a solution that would satisfy these requirements. The data could be used by other systems within the company.

Which solution will meet these requirements with minimal effort?

- A. Create an Amazon Cloudwatch Events rule with the operations that need to be tracked on Amazon RD
- B. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.
- C. Create an AWS Lambda function to trigger on AWS CloudTrail API call
- D. Filter on specific RDS API calls and write the output to the tracking systems.
- E. Create RDS event subscription
- F. Have the tracking systems subscribe to specific RDS event system notifications.
- G. Write RDS logs to Amazon Kinesis Data Firehose
- H. Create an AWS Lambda function to act on these rules and write the output to the tracking systems.

**Answer: C**

#### NEW QUESTION 38

A large ecommerce company uses Amazon DynamoDB to handle the transactions on its web portal. Traffic patterns throughout the year are usually stable; however, a large event is planned. The company knows that traffic will increase by up to 10 times the normal load over the 3-day event. When sale prices are published during the event, traffic will spike rapidly.

How should a Database Specialist ensure DynamoDB can handle the increased traffic?

- A. Ensure the table is always provisioned to meet peak needs
- B. Allow burst capacity to handle the additional load
- C. Set an AWS Application Auto Scaling policy for the table to handle the increase in traffic
- D. Preprovision additional capacity for the known peaks and then reduce the capacity after the event

**Answer: B**

#### NEW QUESTION 42

A retail company with its main office in New York and another office in Tokyo plans to build a database solution on AWS. The company's main workload consists of a mission-critical application that updates its application data in a data store. The team at the Tokyo office is building dashboards with complex analytical queries using the application data. The dashboards will be used to make buying decisions, so they need to have access to the application data in less than 1 second.

Which solution meets these requirements?

- A. Use an Amazon RDS DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region
- B. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cache application data from the replica to generate the dashboards.
- C. Use an Amazon DynamoDB global table in the us-east-1 Region with replication into the ap-northeast-1 Region

- D. Use Amazon QuickSight for displaying dashboard results.
- E. Use an Amazon RDS for MySQL DB instance deployed in the us-east-1 Region with a read replica instance in the ap-northeast-1 Region.
- F. Have the dashboard application read from the read replica.
- G. Use an Amazon Aurora global database.
- H. Deploy the writer instance in the us-east-1 Region and the replica in the ap-northeast-1 Region.
- I. Have the dashboard application read from the replica in the ap-northeast-1 Region.

**Answer: D**

#### NEW QUESTION 45

A company wants to automate the creation of secure test databases with random credentials to be stored safely for later use. The credentials should have sufficient information about each test database to initiate a connection and perform automated credential rotations. The credentials should not be logged or stored anywhere in an unencrypted form.

Which steps should a Database Specialist take to meet these requirements using an AWS CloudFormation template?

- A. Create the database with the MasterUserName and MasterUserPassword properties set to the default value
- B. Then, create the secret with the user name and password set to the same default value
- C. Add a Secret Target Attachment resource with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database
- D. Finally, update the secret's password value with a randomly generated string set by the GenerateSecretString property.
- E. Add a Mapping property from the database Amazon Resource Name (ARN) to the secret ARN
- F. Then, create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property
- G. Add the database with the MasterUserName and MasterUserPassword properties set to the user name of the secret.
- H. Add a resource of type AWS::SecretsManager::Secret and specify the GenerateSecretString property. Then, define the database user name in the SecureStringTemplate template
- I. Create a resource for the database and reference the secret string for the MasterUserName and MasterUserPassword properties
- J. Then, add a resource of type AWS::SecretsManager::SecretTargetAttachment with the SecretId and TargetId properties set to the Amazon Resource Names (ARNs) of the secret and the database.
- K. Create the secret with a chosen user name and a randomly generated password set by the GenerateSecretString property
- L. Add a SecretTargetAttachment resource with the SecretId property set to the Amazon Resource Name (ARN) of the secret and the TargetId property set to a parameter value matching the desired database ARN
- M. Then, create a database with the MasterUserName and MasterUserPassword properties set to the previously created values in the secret.

**Answer: C**

#### NEW QUESTION 46

A Database Specialist is performing a proof of concept with Amazon Aurora using a small instance to confirm a simple database behavior. When loading a large dataset and creating the index, the Database Specialist encounters the following error message from Aurora:

ERROR: cloud not write block 7507718 of temporary file: No space left on device

What is the cause of this error and what should the Database Specialist do to resolve this issue?

- A. The scaling of Aurora storage cannot catch up with the data load
- B. The Database Specialist needs to modify the workload to load the data slowly.
- C. The scaling of Aurora storage cannot catch up with the data load
- D. The Database Specialist needs to enable Aurora storage scaling.
- E. The local storage used to store temporary tables is full
- F. The Database Specialist needs to scale up the instance.
- G. The local storage used to store temporary tables is full
- H. The Database Specialist needs to enable local storage scaling.

**Answer: C**

#### NEW QUESTION 49

A company is developing a multi-tier web application hosted on AWS using Amazon Aurora as the database.

The application needs to be deployed to production and other non-production environments. A Database Specialist needs to specify different MasterUsername and MasterUserPassword properties in the AWS CloudFormation templates used for automated deployment. The CloudFormation templates are version controlled in the company's code repository. The company also needs to meet compliance requirements by routinely rotating its database master password for production.

What is the most secure solution to store the master password?

- A. Store the master password in a parameter file in each environment
- B. Reference the environment-specific parameter file in the CloudFormation template.
- C. Encrypt the master password using an AWS KMS key
- D. Store the encrypted master password in the CloudFormation template.
- E. Use the secretsmanager dynamic reference to retrieve the master password stored in AWS SecretsManager and enable automatic rotation.
- F. Use the ssm dynamic reference to retrieve the master password stored in the AWS Systems Manager Parameter Store and enable automatic rotation.

**Answer: C**

#### NEW QUESTION 51

A Database Specialist is designing a new database infrastructure for a ride hailing application. The application data includes a ride tracking system that stores GPS coordinates for all rides. Real-time statistics and metadata lookups must be performed with high throughput and microsecond latency. The database should be fault tolerant with minimal operational overhead and development effort.

Which solution meets these requirements in the MOST efficient way?

- A. Use Amazon RDS for MySQL as the database and use Amazon ElastiCache
- B. Use Amazon DynamoDB as the database and use DynamoDB Accelerator
- C. Use Amazon Aurora MySQL as the database and use Aurora's buffer cache
- D. Use Amazon DynamoDB as the database and use Amazon API Gateway

**Answer: D**

#### NEW QUESTION 53

A company has a database monitoring solution that uses Amazon CloudWatch for its Amazon RDS for SQL Server environment. The cause of a recent spike in CPU utilization was not determined using the standard metrics that were collected. The CPU spike caused the application to perform poorly, impacting users. A Database Specialist needs to determine what caused the CPU spike.

Which combination of steps should be taken to provide more visibility into the processes and queries running during an increase in CPU load? (Choose two.)

- A. Enable Amazon CloudWatch Events and view the incoming T-SQL statements causing the CPU to spike.
- B. Enable Enhanced Monitoring metrics to view CPU utilization at the RDS SQL Server DB instance level.
- C. Implement a caching layer to help with repeated queries on the RDS SQL Server DB instance.
- D. Use Amazon QuickSight to view the SQL statement being run.
- E. Enable Amazon RDS Performance Insights to view the database load and filter the load by waits, SQL statements, hosts, or users.

**Answer:** BE

#### NEW QUESTION 55

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

**Answer:** C

#### NEW QUESTION 56

A company is going to use an Amazon Aurora PostgreSQL DB cluster for an application backend. The DB cluster contains some tables with sensitive data. A Database Specialist needs to control the access privileges at the table level.

How can the Database Specialist meet these requirements?

- A. Use AWS IAM database authentication and restrict access to the tables using an IAM policy.
- B. Configure the rules in a NACL to restrict outbound traffic from the Aurora DB cluster.
- C. Execute GRANT and REVOKE commands that restrict access to the tables containing sensitive data.
- D. Define access privileges to the tables containing sensitive data in the pg\_hba.conf file.

**Answer:** C

#### NEW QUESTION 57

A company is running a two-tier ecommerce application in one AWS account. The web server is deployed using an Amazon RDS for MySQL Multi-AZ DB instance. A Developer mistakenly deleted the database in the production environment. The database has been restored, but this resulted in hours of downtime and lost revenue.

Which combination of changes in existing IAM policies should a Database Specialist make to prevent an error like this from happening in the future? (Choose three.)

- A. Grant least privilege to groups, users, and roles
- B. Allow all users to restore a database from a backup that will reduce the overall downtime to restore the database
- C. Enable multi-factor authentication for sensitive operations to access sensitive resources and API operations
- D. Use policy conditions to restrict access to selective IP addresses
- E. Use AccessList Controls policy type to restrict users for database instance deletion
- F. Enable AWS CloudTrail logging and Enhanced Monitoring

**Answer:** ACD

#### NEW QUESTION 58

A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQL DB instance. Immediately after creating the read replica, users that query it report slow response times.

What could be causing these slow response times?

- A. New volumes created from snapshots load lazily in the background
- B. Long-running statements on the master
- C. Insufficient resources on the master
- D. Overload of a single replication thread by excessive writes on the master

**Answer:** B

#### NEW QUESTION 60

A company is about to launch a new product, and test databases must be re-created from production data. The company runs its production databases on an Amazon Aurora MySQL DB cluster. A Database Specialist needs to deploy a solution to create these test databases as quickly as possible with the least amount of administrative effort.

What should the Database Specialist do to meet these requirements?

- A. Restore a snapshot from the production cluster into test clusters
- B. Create logical dumps of the production cluster and restore them into new test clusters
- C. Use database cloning to create clones of the production cluster
- D. Add an additional read replica to the production cluster and use that node for testing

Answer: D

#### NEW QUESTION 64

A financial company has allocated an Amazon RDS MariaDB DB instance with large storage capacity to accommodate migration efforts. Post-migration, the company purged unwanted data from the instance. The company now want to downsize storage to save money. The solution must have the least impact on production and near-zero downtime.

Which solution would meet these requirements?

- A. Create a snapshot of the old databases and restore the snapshot with the required storage
- B. Create a new RDS DB instance with the required storage and move the databases from the old instance to the new instance using AWS DMS
- C. Create a new database using native backup and restore
- D. Create a new read replica and make it the primary by terminating the existing primary

Answer: A

#### NEW QUESTION 68

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.

Answer: A

#### NEW QUESTION 71

A Database Specialist needs to speed up any failover that might occur on an Amazon Aurora PostgreSQL DB cluster. The Aurora DB cluster currently includes the primary instance and three Aurora Replicas.

How can the Database Specialist ensure that failovers occur with the least amount of downtime for the application?

- A. Set the TCP keepalive parameters low
- B. Call the AWS CLI failover-db-cluster command
- C. Enable Enhanced Monitoring on the DB cluster
- D. Start a database activity stream on the DB cluster

Answer: B

#### NEW QUESTION 74

A company needs a data warehouse solution that keeps data in a consistent, highly structured format. The company requires fast responses for end-user queries when looking at data from the current year, and users must have access to the full 15-year dataset, when needed. This solution also needs to handle a fluctuating number incoming queries. Storage costs for the 100 TB of data must be kept low. Which solution meets these requirements?

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

Answer: C

#### NEW QUESTION 78

A manufacturing company's website uses an Amazon Aurora PostgreSQL DB cluster.

Which configurations will result in the LEAST application downtime during a failover? (Choose three.)

- A. Use the provided read and write Aurora endpoints to establish a connection to the Aurora DB cluster.
- B. Create an Amazon CloudWatch alert triggering a restore in another Availability Zone when the primary Aurora DB cluster is unreachable.
- C. Edit and enable Aurora DB cluster cache management in parameter groups.
- D. Set TCP keepalive parameters to a high value.
- E. Set JDBC connection string timeout variables to a low value.
- F. Set Java DNS caching timeouts to a high value.

Answer: ABC

#### NEW QUESTION 82

An Amazon RDS EBS-optimized instance with Provisioned IOPS (PIOPS) storage is using less than half of its allocated IOPS over the course of several hours under constant load. The RDS instance exhibits multi-second read and write latency, and uses all of its maximum bandwidth for read throughput, yet the instance uses less than half of its CPU and RAM resources.

What should a Database Specialist do in this situation to increase performance and return latency to sub-second levels?

- A. Increase the size of the DB instance storage
- B. Change the underlying EBS storage type to General Purpose SSD (gp2)
- C. Disable EBS optimization on the DB instance
- D. Change the DB instance to an instance class with a higher maximum bandwidth

**Answer:** B

#### NEW QUESTION 87

A company is load testing its three-tier production web application deployed with an AWS CloudFormation template on AWS. The Application team is making changes to deploy additional Amazon EC2 and AWS Lambda resources to expand the load testing capacity. A Database Specialist wants to ensure that the changes made by the Application team will not change the Amazon RDS database resources already deployed. Which combination of steps would allow the Database Specialist to accomplish this? (Choose two.)

- A. Review the stack drift before modifying the template
- B. Create and review a change set before applying it
- C. Export the database resources as stack outputs
- D. Define the database resources in a nested stack
- E. Set a stack policy for the database resources

**Answer:** AD

#### NEW QUESTION 90

A company is running an Amazon RDS for PostgreSQL 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
- B. Set up replication from the source RDS for PostgreSQL DB instance using AWS DMS to the target DB cluster.
- C. Use the pg\_dump and pg\_restore utilities to extract and restore the RDS for PostgreSQL DB instance to the Aurora PostgreSQL DB cluster.
- D. Create a database snapshot of the RDS for PostgreSQL DB instance and use this snapshot to create the Aurora PostgreSQL DB cluster.
- E. Migrate data from the RDS for PostgreSQL DB instance to an Aurora PostgreSQL DB cluster using an Aurora Replic
- F. Promote the replica during the cutover.

**Answer:** C

#### NEW QUESTION 91

A company has multiple applications serving data from a secure on-premises database. The company is migrating all applications and databases to the AWS Cloud. The IT Risk and Compliance department requires that auditing be enabled on all secure databases to capture all log ins, log outs, failed logins, permission changes, and database schema changes. A Database Specialist has recommended Amazon Aurora MySQL as the migration target, and leveraging the Advanced Auditing feature in Aurora.

Which events need to be specified in the Advanced Auditing configuration to satisfy the minimum auditing requirements? (Choose three.)

- A. CONNECT
- B. QUERY\_DCL
- C. QUERY\_DDL
- D. QUERY\_DML
- E. TABLE
- F. QUERY

**Answer:** ACE

#### NEW QUESTION 95

A company is using Amazon with Aurora Replicas for read-only workload scaling. A Database Specialist needs to split up two read-only applications so each application always connects to a dedicated replica. The Database Specialist wants to implement load balancing and high availability for the read-only applications. Which solution meets these requirements?

- A. Use a specific instance endpoint for each replica and add the instance endpoint to each read-only application connection string.
- B. Use reader endpoints for both the read-only workload applications.
- C. Use a reader endpoint for one read-only application and use an instance endpoint for the other read-only application.
- D. Use custom endpoints for the two read-only applications.

**Answer:** B

#### NEW QUESTION 97

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### DBS-C01 Practice Exam Features:

- \* DBS-C01 Questions and Answers Updated Frequently
- \* DBS-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* DBS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DBS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DBS-C01 Practice Test Here](#)**