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 ongoingreplication 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 CPUconsumption.

Answer: A

NEW QUESTION 2
A company has deployed an e-commerce web application in a new AWS account. An Amazon RDS for MySQL Multi-AZ DB instance is part of this deployment with a
database-1.xxxxxxxxxxxx.us-east-1.rds.amazonaws.com endpoint listening on port 3306. The company's Database Specialist is able to log in to MySQL and run queries from the bastion host using these details.
When users try to utilize the application hosted in the AWS account, they are presented with a generic error message. The application servers are logging a "could not connect to server: Connection times out" error message to Amazon CloudWatch Logs.
What is the cause of this error?

A. The user name and password the application is using are incorrect.
B. The security group assigned to the application servers does not have the necessary rules to allow inbound connections from the DB instance.
C. The security group assigned to the DB instance does not have the necessary rules to allow inbound connections from the application servers.
D. The user name and password are correct, but the user is not authorized to use the DB instance.

Answer: C

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 will the appropriate settings

Answer: C

NEW QUESTION 4
A Database Specialist needs to define a database migration strategy to migrate an on-premises Oracle database to an Amazon Aurora MySQL DB cluster. The company requires near-zero downtime for the data migration. The solution must also be cost-effective.
Which approach should the Database Specialist take?

A. Dump all the tables from the Oracle database into an Amazon S3 bucket using datapump (expdp).Rundata transformations in AWS Glu
B. Load the data from the S3 bucket to the Aurora DB cluster.
C. Order an AWS Snowball appliance and copy the Oracle backup to the Snowball applianc
D. Once theSnowball data is delivered to Amazon S3, create a new Aurora DB cluste
E. Enable the S3 integration tomigrate the data directly from Amazon S3 to Amazon RDS.
F. Use the AWS Schema Conversion Tool (AWS SCT) to help rewrite database objects to MySQL during theschema migratio
G. Use AWS DMS to perform the full load and change data capture (CDC) tasks.
H. Use AWS Server Migration Service (AWS SMS) to import the Oracle virtual machine image as an AmazonEC2 instanc
I. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to anAurora DB cluster.

Answer: D

NEW QUESTION 5
A company has migrated a single MySQL database to Amazon Aurora. The production data is hosted in a DB cluster in VPC_PROD, and 12 testing environments are hosted in VPC_TEST using the same AWS account. Testing results in minimal changes to the test data. The Development team wants each environment refreshed nightly so each test database contains fresh production data every day.
Which migration approach will be the fastest and most cost-effective to implement?

A. Run the master in Amazon Aurora MySQ
B. Create 12 clones in VPC_TEST, and script the clones to bedeleted and re-created nightly.
C. Run the master in Amazon Aurora MySQ
D. Take a nightly snapshot, and restore it into 12 databases inVPC_TEST using Aurora Serverless.
E. Run the master in Amazon Aurora MySQ
F. Create 12 Aurora Replicas in VPC_TEST, and script thereplicas to be deleted and re-created nightly.
G. Run the master in Amazon Aurora MySQL using Aurora Serverles
H. Create 12 clones in VPC_TEST, andscript the clones to be deleted and re-created nightly.

Answer: A

NEW QUESTION 6
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 7
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 bucke
B. Set a lifecycle policy to expire the objects after 90 days.
C. Modify the RDS databases to publish log to Amazon CloudWatch Log
D. Change the log retention policy for each log group to expire the events after 90 days.
E. Write a stored procedure in each RDS database to download the logs and consolidate the log files in an Amazon S3 bucke
F. Set a lifecycle policy to expire the objects after 90 days.
G. Create an AWS Lambda function to download the logs from the RDS databases and publish the logs to Amazon CloudWatch Log
H. Change the log retention policy for the log group to expire the events after 90 days.

Answer: A

NEW QUESTION 8
A Database Specialist is creating a new Amazon Neptune DB cluster, and is attempting to load fata 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 9
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

Answer: C

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
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 templat
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 templat
E. Configure DynamoDBto provision throughput capacity using the stack outputs.
F. Add values for the rcuCount and wcuCount parameters to the Mappings section of the template.Replacethe hard-coded values with calls to the Ref intrinsic function, referencing the new parameters.

Answer: B

NEW QUESTION 12
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 operationand store the logs in Amazon S3

Answer: A

NEW QUESTION 16
A company is planning to close for several days. A Database Specialist needs to stop all applications alongwith the DB instances to ensure employees do not have access to the systems during this time. All databasesare running on Amazon RDS for MySQL.
The Database Specialist wrote and executed a script to stop all the DB instances. When reviewing the logs,the Database Specialist found that Amazon RDS DB instances with read replicas did not stop.
How should the Database Specialist edit the script to fix this issue?

A. Stop the source instances before stopping their read replicas
B. Delete each read replica before stopping its corresponding source instance
C. Stop the read replicas before stopping their source instances
D. Use the AWS CLI to stop each read replica and source instance at the same

Answer: D

NEW QUESTION 20
A company is looking to move an on-premises IBM Db2 database running AIX on an IBM POWER7 server. Due to escalating support and maintenance costs, the company is exploring the option of moving the workload to an Amazon Aurora PostgreSQL DB cluster.
What is the quickest way for the company to gather data on the migration compatibility?

A. Perform a logical dump from the Db2 database and restore it to an Aurora DB cluste
B. Identify the gaps andcompatibility of the objects migrated by comparing row counts from source and target tables.
C. Run AWS DMS from the Db2 database to an Aurora DB cluste
D. Identify the gaps and compatibility of theobjects migrated by comparing the row counts from source and target tables.
E. Run native PostgreSQL logical replication from the Db2 database to an Aurora DB cluster to evaluate themigration compatibility.
F. Run the AWS Schema Conversion Tool (AWS SCT) from the Db2 database to an Aurora DB cluster.Create a migration assessment report to evaluate the migration compatibility.

Answer: D

NEW QUESTION 24
A company is deploying a solution in Amazon Aurora by migrating from an on-premises system. The IT department has established an AWS Direct Connect link from the company's data center. The company's Database Specialist has selected the option to require SSL/TLS for connectivity to prevent plaintext data from being set over the network. The migration appears to be working successfully, and the data can be queried from a desktop machine.
Two Data Analysts have been asked to query and validate the data in the new Aurora DB cluster. Both Analysts are unable to connect to Aurora. Their user names and passwords have been verified as valid and the Database Specialist can connect to the DB cluster using their accounts. The Database Specialist also verified that the security group configuration allows network from all corporate IP addresses.
What should the Database Specialist do to correct the Data Analysts' inability to connect?

A. Restart the DB cluster to apply the SSL change.
B. Instruct the Data Analysts to download the root certificate and use the SSL certificate on the connection string to connect.
C. Add explicit mappings between the Data Analysts' IP addresses and the instance in the security groupassigned to the DB cluster.
D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer: D

NEW QUESTION 29
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 apnortheast-1 Regio
B. Create an Amazon ElastiCache cluster in the ap-northeast-1 Region to cacheapplication 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-1Regio
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 replicainstance in the ap-northeast-1 Regio
F. Have the dashboard application read from the read replica.

G. Use an Amazon Aurora global databas
H. Deploy the writer instance in the us-east-1 Region and the replicain the ap-northeast-1 Regio
I. Have the dashboard application read from the replica ap-northeast-1 Region.

Answer: D

NEW QUESTION 30
A company is using Amazon RDS for PostgreSQL. The Security team wants all database connection requests to be logged and retained for 180 days. The RDS for PostgreSQL DB instance is currently using the default parameter group. A Database Specialist has identified that setting the log_connections parameter to 1 will enable connections logging.
Which combination of steps should the Database Specialist take to meet the logging and retention requirements? (Choose two.)

A. Update the log_connections parameter in the default parameter group
B. Create a custom parameter group, update the log_connections parameter, and associate the parameterwith the DB instance
C. Enable publishing of database engine logs to Amazon CloudWatch Logs and set the event expiration to180 days
D. Enable publishing of database engine logs to an Amazon S3 bucket and set the lifecycle policy to 180 days
E. Connect to the RDS PostgreSQL host and update the log_connections parameter in the postgresql.conf file

Answer: AE

NEW QUESTION 31
A company wants to migrate its existing on-premises Oracle database to Amazon Aurora PostgreSQL. The migration must be completed with minimal downtime using AWS DMS. A Database Specialist must validate that the data was migrated accurately from the source to the target before the cutover. The migration must have minimal impact on the performance of the source database.
Which approach will MOST effectively meet these requirements?

A. Use the AWS Schema Conversion Tool (AWS SCT) to convert source Oracle database schemas to the target Aurora DB cluste
B. Verify the datatype of the columns.
C. Use the table metrics of the AWS DMS task created for migrating the data to verify the statistics for the tables being migrated and to verify that the data definition language (DDL) statements are completed.
D. Enable the AWS Schema Conversion Tool (AWS SCT) premigration validation and review the premigrationchecklist to make sure there are no issues with the conversion.
E. Enable AWS DMS data validation on the task so the AWS DMS task compares the source and targetrecords, and reports any mismatches.

Answer: D

NEW QUESTION 35
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 39
A Database Specialist must create a read replica to isolate read-only queries for an Amazon RDS for MySQLDB 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 40
A company has an Amazon RDS Multi-AZ DB instances that is 200 GB in size with an RPO of 6 hours. To meet the company's disaster recovery policies, the database backup needs to be copied into another Region. The company requires the solution to be cost-effective and operationally efficient.
What should a Database Specialist do to copy the database backup into a different Region?

A. Use Amazon RDS automated snapshots and use AWS Lambda to copy the snapshot into another Region
B. Use Amazon RDS automated snapshots every 6 hours and use Amazon S3 cross-Region replication tocopy the snapshot into another Region
C. Create an AWS Lambda function to take an Amazon RDS snapshot every 6 hours and use a secondLambda function to copy the snapshot into another Region
D. Create a cross-Region read replica for Amazon RDS in another Region and take an automated snapshot ofthe read replica

Answer: D

NEW QUESTION 42
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 47
A Database Specialist has migrated an on-premises Oracle database to Amazon Aurora PostgreSQL. The schema and the data have been migrated successfully. The on-premises database server was also being used to run database maintenance cron jobs written in Python to perform tasks including data purging and generating data exports. The logs for these jobs show that, most of the time, the jobs completed within 5 minutes, but a few jobs took up to 10 minutes to complete. These maintenance jobs need to be set up for Aurora PostgreSQL.
How can the Database Specialist schedule these jobs so the setup requires minimal maintenance and provides high availability?

A. Create cron jobs on an Amazon EC2 instance to run the maintenance jobs following the required schedule.
B. Connect to the Aurora host and create cron jobs to run the maintenance jobs following the requiredschedule.
C. Create AWS Lambda functions to run the maintenance jobs and schedule them with Amazon CloudWatchEvents.
D. Create the maintenance job using the Amazon CloudWatch job scheduling plugin.

Answer: D

NEW QUESTION 50
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.

Answer: D

NEW QUESTION 51
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 then 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

Answer: C

NEW QUESTION 52
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 54
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.Enablecache coherence on the DB cluster, set the primary DB instance failover priority to tier-0, and assign afailover priority of tier-1 to the replicas.
B. Deploy an AWS Lambda function that calls the DescribeDBInstances action to establish which instancehas failed, and then use the PromoteReadReplica operation to promote one Aurora Replica to be theprimary DB instanc
C. Configure an Amazon RDS event subscription to send a notification to an AmazonSNS topic to which the Lambda function is subscribed.
D. Configure one Aurora Replica to have the same instance class as the primary DB instance.ImplementAurora PostgreSQL DB cluster cache managemen
E. Set the failover priority to tier-0 for the primary DBinstance and one replica with the same instance clas
F. Set the failover priority to tier-1 for the otherreplicas.
G. Configure both Aurora Replicas to have the same instance class as the primary DB instance.ImplementAurora PostgreSQL DB cluster cache managemen
H. Set the failover priority to tier-0 for the primary DBinstance and to tier-1 for the replicas.

Answer: D

NEW QUESTION 59
A gaming company wants to deploy a game in multiple Regions. The company plans to save local high scores in Amazon DynamoDB tables in each Region. A Database Specialist needs to design a solution to automate the deployment of the database with identical configurations in additional Regions, as needed. The solution should also automate configuration changes across all Regions.
Which solution would meet these requirements and deploy the DynamoDB tables?

A. Create an AWS CLI command to deploy the DynamoDB table to all the Regions and save it for future deployments.
B. Create an AWS CloudFormation template and deploy the template to all the Regions.
C. Create an AWS CloudFormation template and use a stack set to deploy the template to all the Regions.
D. Create DynamoDB tables using the AWS Management Console in all the Regions and create a step-bystep guide for future deployments.

Answer: B

NEW QUESTION 63
A Database Specialist is creating Amazon DynamoDB tables, Amazon CloudWatch alarms, and associated infrastructure for an Application team using a development AWS account. The team wants a deployment method that will standardize the core solution components while managing environment-specific settings separately, and wants to minimize rework due to configuration errors.
Which process should the Database Specialist recommend to meet these requirements?

A. Organize common and environmental-specific parameters hierarchically in the AWS Systems ManagerParameter Store, then reference the parameters dynamically from an AWS CloudFormation template.Deploy the CloudFormation stack using the environment name as a parameter.
B. Create a parameterized AWS CloudFormation template that builds the required object
C. Keep separateenvironment parameter files in separate Amazon S3 bucket
D. Provide an AWS CLI command that deploysthe CloudFormation stack directly referencing the appropriate parameter bucket.
E. Create a parameterized AWS CloudFormation template that builds the required object
F. Import thetemplate into the CloudFormation interface in the AWS Management Consol
G. Make the required changesto the parameters and deploy the CloudFormation stack.
H. Create an AWS Lambda function that builds the required objects using an AWS SD
I. Set the requiredparameter values in a test event in the Lambda console for each environment that the Application team canmodify, as neede
J. Deploy the infrastructure by triggering the test event in the console.

Answer: C

NEW QUESTION 67
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-onlyapplication 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-onlyapplication.
D. Use custom endpoints for the two read-only applications.

Answer: B

NEW QUESTION 68
......