

DBS-C01 Dumps

AWS Certified Database - Specialty

<https://www.certleader.com/DBS-C01-dumps.html>



NEW QUESTION 1

A company is writing a new survey application to be used with a weekly televised game show. The application will be available for 2 hours each week. The company expects to receive over 500,000 entries every week, with each survey asking 2-3 multiple choice questions of each user. A Database Specialist needs to select a platform that is highly scalable for a large number of concurrent writes to handle the anticipated volume.

Which AWS services should the Database Specialist consider? (Choose two.)

- A. Amazon DynamoDB
- B. Amazon Redshift
- C. Amazon Neptune
- D. Amazon Elasticsearch Service
- E. Amazon ElastiCache

Answer: AE

NEW QUESTION 2

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 3

A retail company is about to migrate its online and mobile store to AWS. The company's CEO has strategic plans to grow the brand globally. A Database Specialist has been challenged to provide predictable read and write database performance with minimal operational overhead.

What should the Database Specialist do to meet these requirements?

- A. Use Amazon DynamoDB global tables to synchronize transactions
- B. Use Amazon EMR to copy the orders table data across Regions
- C. Use Amazon Aurora Global Database to synchronize all transactions
- D. Use Amazon DynamoDB Streams to replicate all DynamoDB transactions and sync them

Answer: A

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). Run data transformations in AWS Glue
- 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 appliance
- D. Once the Snowball data is delivered to Amazon S3, create a new Aurora DB cluster
- E. Enable the S3 integration to migrate 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 the schema migration
- 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 Amazon EC2 instance
- I. Use the Oracle Logical Dump utility to migrate the Oracle data from Amazon EC2 to an Aurora 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 MySQL
- B. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.
- C. Run the master in Amazon Aurora MySQL
- D. Take a nightly snapshot, and restore it into 12 databases in VPC_TEST using Aurora Serverless.
- E. Run the master in Amazon Aurora MySQL
- F. Create 12 Aurora Replicas in VPC_TEST, and script the replicas to be deleted and re-created nightly.
- G. Run the master in Amazon Aurora MySQL using Aurora Serverless
- H. Create 12 clones in VPC_TEST, and script the clones to be deleted and re-created nightly.

Answer: A

NEW QUESTION 6

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
- 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 bucket
- 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 7

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 8

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 9

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 AWS
- B. Leverage AWS SCT and apply the converted schema to Amazon Redshift
- 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 Redshift
- 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 Redshift
- 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 AWS
- 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 10

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 RDS
- 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 10

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 group assigned to the DB cluster.
- D. Modify the Data Analysts' local client firewall to allow network traffic to AWS.

Answer: D

NEW QUESTION 12

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 cluster.
- 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 premigration checklist 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 target records, and reports any mismatches.

Answer: D

NEW QUESTION 16

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 21

An ecommerce company is using Amazon DynamoDB as the backend for its order-processing application. The steady increase in the number of orders is resulting in increased DynamoDB costs. Order verification and reporting perform many repeated GetItem functions that pull similar datasets, and this read activity is contributing to the increased costs. The company wants to control these costs without significant development efforts. How should a Database Specialist address these requirements?

- A. Use AWS DMS to migrate data from DynamoDB to Amazon DocumentDB
- B. Use Amazon DynamoDB Streams and Amazon Kinesis Data Firehose to push the data into Amazon Redshift
- C. Use an Amazon ElastiCache for Redis in front of DynamoDB to boost read performance
- D. Use DynamoDB Accelerator to offload the reads

Answer: B

NEW QUESTION 24

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 28

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 29

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 separate environment parameter files in separate Amazon S3 bucket
- D. Provide an AWS CLI command that deploys the CloudFormation stack directly referencing the appropriate parameter bucket.
- E. Create a parameterized AWS CloudFormation template that builds the required object
- F. Import the template into the CloudFormation interface in the AWS Management Console
- G. Make the required changes to 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 required parameter values in a test event in the Lambda console for each environment that the Application team can modify, as needed
- J. Deploy the infrastructure by triggering the test event in the console.

Answer: C

NEW QUESTION 32

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your DBS-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/DBS-C01-dumps.html>