DAS-C01 Q&As

AWS Certified Data Analytics - Specialty (DAS-C01)

Pass Amazon DAS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/das-c01.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

QUESTION 1

A data analyst is designing an Amazon QuickSight dashboard using centralized sales data that resides in Amazon Redshift. The dashboard must be restricted so that a salesperson in Sydney, Australia, can see only the Australia view and that a salesperson in New York can see only United States (US) data.

What should the data analyst do to ensure the appropriate data security is in place?

A. Place the data sources for Australia and the US into separate SPICE capacity pools.

B. Set up an Amazon Redshift VPC security group for Australia and the US.

C. Deploy QuickSight Enterprise edition to implement row-level security (RLS) to the sales table.

D. Deploy QuickSight Enterprise edition and set up different VPC security groups for Australia and the US.

Correct Answer: D

Reference: https://docs.aws.amazon.com/quicksight/latest/user/working-with-aws-vpc.html

QUESTION 2

A company has a production AWS account that runs production workloads. The company created a new security AWS account to store and analyze security logs from the production AWS account. The security logs in the production AWS account are stored in Amazon CloudWatch Logs. The company needs to use Amazon Kinesis Data Streams to deliver the security logs to the security AWS account.

Which solution will meet these requirements?

A. Create a destination data stream in the production AWS account. In the security AWS account, create an IAM role that has cross-account permissions to Kinesis Data Streams in the production AWS account.

B. Create a destination data stream in the security AWS account. Create an IAM role and a trust policy to grant CloudWatch Logs the permission to put data into the stream. Create a subscription filter in the security AWS account.

C. Create a destination data stream in the production AWS account. In the production AWS account, create an IAM role that has cross-account permissions to Kinesis Data Streams in the security AWS account.

D. Create a destination data stream in the security AWS account. Create an IAM role and a trust policy to grant CloudWatch Logs the permission to put data into the stream. Create a subscription filter in the production AWS account.

Correct Answer: D

QUESTION 3

A company has an application that ingests streaming data. The company needs to analyze this stream over a 5-minute timeframe to evaluate the stream for anomalies with Random Cut Forest (RCF) and summarize the current count of status codes. The source and summarized data should be persisted for future use.

Which approach would enable the desired outcome while keeping data persistence costs low?

A. Ingest the data stream with Amazon Kinesis Data Streams. Have an AWS Lambda consumer evaluate the stream, collect the number status codes, and evaluate the data against a previously trained RCF model. Persist the source and results as a time series to Amazon DynamoDB.

B. Ingest the data stream with Amazon Kinesis Data Streams. Have a Kinesis Data Analytics application evaluate the stream over a 5-minute window using the RCF function and summarize the count of status codes. Persist the source and results to Amazon S3 through output delivery to Kinesis Data Firehouse.

C. Ingest the data stream with Amazon Kinesis Data Firehose with a delivery frequency of 1 minute or 1 MB in Amazon S3. Ensure Amazon S3 triggers an event to invoke an AWS Lambda consumer that evaluates the batch data, collects the number status codes, and evaluates the data against a previously trained RCF model. Persist the source and results as a time series to Amazon DynamoDB.

D. Ingest the data stream with Amazon Kinesis Data Firehose with a delivery frequency of 5 minutes or 1 MB into Amazon S3. Have a Kinesis Data Analytics application evaluate the stream over a 1-minute window using the RCF function and summarize the count of status codes. Persist the results to Amazon S3 through a Kinesis Data Analytics output to an AWS Lambda integration.

Correct Answer: B

QUESTION 4

A real estate company maintains data about all properties listed in a market. The company receives data about new property listings from vendors who upload the data daily as compressed files into Amazon S3. The company\\'s leadership team wants to see the most up-to-date listings as soon as the data is uploaded to Amazon S3. The data analytics team must automate and orchestrate the data processing workflow of the listings to feed a dashboard. The team also must provide the ability to perform one-time queries and analytical reporting in a scalable manner.

Which solution meets these requirements MOST cost-effectively?

A. Use Amazon EMR for processing incoming data. Use AWS Step Functions for workflow orchestration. Use Apache Hive for one-time queries and analytical reporting. Bulk ingest the data in Amazon OpenSearch Service (Amazon Elasticsearch Service). Use OpenSearch Dashboards (Kibana) on Amazon OpenSearch Service (Amazon Elasticsearch Service) for the dashboard.

B. Use Amazon EMR for processing incoming data. Use AWS Step Functions for workflow orchestration. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.

C. Use AWS Glue for processing incoming data. Use AWS Step Functions for workflow orchestration. Use Amazon Redshift Spectrum for one-time queries and analytical reporting. Use OpenSearch Dashboards (Kibana) on Amazon OpenSearch Service (Amazon Elasticsearch Service) for the dashboard.

D. Use AWS Glue for processing incoming data. Use AWS Lambda and S3 Event Notifications for workflow orchestration. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.

Correct Answer: B

Reference: https://aws.amazon.com/blogs/compute/visualizing-aws-step-functions-workflows-from-the-amazon-athena-console/

QUESTION 5

A company hosts an on-premises PostgreSQL database that contains historical data. An internal legacy application uses the database for read-only activities. The company\\'s business team wants to move the data to a data lake in Amazon S3 as soon as possible and enrich the data for analytics.

The company has set up an AWS Direct Connect connection between its VPC and its on-premises network. A data analytics specialist must design a solution that achieves the business team\\'s goals with the least operational overhead.

Which solution meets these requirements?

A. Upload the data from the on-premises PostgreSQL database to Amazon S3 by using a customized batch upload process. Use the AWS Glue crawler to catalog the data in Amazon S3. Use an AWS Glue job to enrich and store the result in a separate S3 bucket in Apache Parquet format. Use Amazon Athena to query the data.

B. Create an Amazon RDS for PostgreSQL database and use AWS Database Migration Service (AWS DMS) to migrate the data into Amazon RDS. Use AWS Data Pipeline to copy and enrich the data from the Amazon RDS for PostgreSQL table and move the data to Amazon S3. Use Amazon Athena to query the data.

C. Configure an AWS Glue crawler to use a JDBC connection to catalog the data in the on-premises database. Use an AWS Glue job to enrich the data and save the result to Amazon S3 in Apache Parquet format. Create an Amazon Redshift cluster and use Amazon Redshift Spectrum to query the data.

D. Configure an AWS Glue crawler to use a JDBC connection to catalog the data in the on-premises database. Use an AWS Glue job to enrich the data and save the result to Amazon S3 in Apache Parquet format. Use Amazon Athena to query the data.

Correct Answer: B