



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.passapply.com/das-c01.html>

100% Passing Guarantee
100% Money Back Assurance

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

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

A data analytics specialist needs to encrypt the storage for an Amazon EMR cluster that processes data for a company's monthly financial report. The EMR cluster uses a security configuration template. The local disks of the EMR cluster are not encrypted.

What should the data analytics specialist do to encrypt the local disks?

- A. Update the existing security configuration to enable encryption for the local disks. Restart the existing cluster for the configuration to take effect.
- B. Create a new security configuration to enable encryption for the local disks. Create a new cluster that uses the new security configuration.
- C. Create a new security configuration to enable encryption for the local disks. Restart the existing cluster with the new security configuration.
- D. Update the existing security configuration to enable encryption for the local disks. Create a new cluster that uses the updated security configuration.

Correct Answer: A

QUESTION 2

A banking company plans to build a data warehouse solution on AWS to run join queries on 20 TB of data. These queries will be complex and analytical. About 10% of the data is from the past 3 months. Data older than 3 months needs to be accessed occasionally to run queries.

Which solution MOST cost-effectively meets these requirements?

- A. Use Amazon S3 as the data store and use Amazon Athena for the queries. Use Amazon S3 Glacier Flexible Retrieval for storing data older than 3 months by using S3 lifecycle policies.
- B. Use Amazon Redshift to build a data warehouse solution. Create an AWS Lambda function that is orchestrated by AWS Step Functions to run the UNLOAD command on data older than 3 months from the Redshift database to Amazon S3. Use Amazon Redshift Spectrum to query the data in Amazon S3.
- C. Use Amazon Redshift to build a data warehouse solution. Use RA3 instances for the Redshift cluster so that data requested for a query is stored in a solid state drive (SSD) for fast local storage and Amazon S3 for longer-term durable storage.
- D. Use Amazon Elastic File System (Amazon EFS) to build a data warehouse solution for data storage. Use Amazon EFS lifecycle management to retire data older than 3 months to the S3 Standard-Infrequent Access (S3 Standard-IA) class. Use Apache Presto on an Amazon EMR cluster to query the data interactively.

Correct Answer: C

QUESTION 3



A financial company hosts a data lake in Amazon S3 and a data warehouse on an Amazon Redshift cluster. The company uses Amazon QuickSight to build dashboards and wants to secure access from its on-premises Active Directory to Amazon QuickSight.

How should the data be secured?

- A. Use an Active Directory connector and single sign-on (SSO) in a corporate network environment.
- B. Use a VPC endpoint to connect to Amazon S3 from Amazon QuickSight and an IAM role to authenticate Amazon Redshift.
- C. Establish a secure connection by creating an S3 endpoint to connect Amazon QuickSight and a VPC endpoint to connect to Amazon Redshift.
- D. Place Amazon QuickSight and Amazon Redshift in the security group and use an Amazon S3 endpoint to connect Amazon QuickSight to Amazon S3.

Correct Answer: B

QUESTION 4

A company is planning to create a data lake in Amazon S3. The company wants to create tiered storage based on access patterns and cost objectives. The solution must include support for JDBC connections from legacy clients, metadata management that allows federation for access control, and batch-based ETL using PySpark and Scala. Operational management should be limited.

Which combination of components can meet these requirements? (Choose three.)

- A. AWS Glue Data Catalog for metadata management
- B. Amazon EMR with Apache Spark for ETL
- C. AWS Glue for Scala-based ETL
- D. Amazon EMR with Apache Hive for JDBC clients
- E. Amazon Athena for querying data in Amazon S3 using JDBC drivers
- F. Amazon EMR with Apache Hive, using an Amazon RDS with MySQL-compatible backed metastore

Correct Answer: ACE

Reference: <https://d1.awsstatic.com/whitepapers/Storage/data-lake-on-aws.pdf>

QUESTION 5

A large retailer has successfully migrated to an Amazon S3 data lake architecture. The company's marketing team is using Amazon Redshift and Amazon QuickSight to analyze data, and derive and visualize insights. To ensure the marketing team has the most up-to-date actionable information, a data analyst implements nightly refreshes of Amazon Redshift using terabytes of updates from the previous day.

After the first nightly refresh, users report that half of the most popular dashboards that had been running correctly before the refresh are now running much slower. Amazon CloudWatch does not show any alerts.



What is the MOST likely cause for the performance degradation?

- A. The dashboards are suffering from inefficient SQL queries.
- B. The cluster is undersized for the queries being run by the dashboards.
- C. The nightly data refreshes are causing a lingering transaction that cannot be automatically closed by Amazon Redshift due to ongoing user workloads.
- D. The nightly data refreshes left the dashboard tables in need of a vacuum operation that could not be automatically performed by Amazon Redshift due to ongoing user workloads.

Correct Answer: B

[Latest DAS-C01 Dumps](#)

[DAS-C01 Practice Test](#)

[DAS-C01 Braindumps](#)