

AWS DAS-C01 Study Guide PDF

GRAB THE AWS DATA ANALYTICS SPECIALTY CERTIFICATION PDF
QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

DAS-C01

AWS Certified Data Analytics - Specialty

65 Questions Exam – 750 / 1000 Cut Score – Duration of 180 minutes

www.VMExam.com

Table of Contents

Get an Overview of the DAS-C01 Certification:.....	3
Why Should You Earn the AWS DAS-C01 Certification?	3
What Is the AWS DAS-C01 Data Analytics Specialty Certification Exam Structure?	3
Enhance Knowledge with DAS-C01 Sample Questions:	4
What Study Guide Works Best in Acing the AWS DAS- C01 Data Analytics Specialty Certification?.....	9
Explore the Syllabus Topics and Learn from the Core:	9
Make Your Schedule:	9
Get Expert Advice from the Training:	9
Get Access to the PDF Sample Questions:	9
Avoid Dumps and Utilize the AWS DAS-C01 Practice Test:	9

Get an Overview of the DAS-C01 Certification:

Who should take the [DAS-C01 exam](#)? This is the first question that comes to a candidate's mind when preparing for the Data Analytics Specialty certification. The DAS-C01 certification is suitable for candidates who are keen to earn knowledge on the Specialty and grab their AWS Certified Data Analytics - Specialty. When it is about starting the preparation, most candidates get confused regarding the study materials and study approach. But DAS-C01 study guide PDF is here to solve the problem. DAS-C01 PDF combines some effective sample questions and offers valuable tips to pass the exam with ease.

Why Should You Earn the AWS DAS-C01 Certification?

There are several reasons why one should grab the DAS-C01 certification.

- The Data Analytics Specialty certification proves to be one of the most recognized certifications.
- The certification badge proves the knowledge of the candidate regarding subject matters and makes his resume presentable to potential candidates.
- Thus earning the [AWS Certified Data Analytics - Specialty](#) is a powerful qualification for a prosperous career.

What Is the AWS DAS-C01 Data Analytics Specialty Certification Exam Structure?

Exam Name	AWS Certified Data Analytics - Specialty (Data Analytics Specialty)
Exam Code	DAS-C01
Exam Price	\$300 USD
Duration	180 minutes
Number of Questions	65
Passing Score	750 / 1000
Recommended Training / Books	Data Analytics Fundamentals Big Data on AWS
Schedule Exam	PEARSON VUE
Sample Questions	AWS DAS-C01 Sample Questions
Recommended Practice	AWS Certified Data Analytics - Specialty Practice Test

Enhance Knowledge with DAS-C01 Sample Questions:

Question: 1

An online retail company wants to perform analytics on data in large Amazon S3 objects using Amazon EMR.

An Apache Spark job repeatedly queries the same data to populate an analytics dashboard. The analytics team wants to minimize the time to load the data and create the dashboard.

Which approaches could improve the performance?

(Select TWO.)

- a) Copy the source data into Amazon Redshift and rewrite the Apache Spark code to create analytical reports by querying Amazon Redshift.
- b) Copy the source data from Amazon S3 into Hadoop Distributed File System (HDFS) using s3distcp.
- c) Load the data into Spark DataFrames.
- d) Stream the data into Amazon Kinesis and use the Kinesis Connector Library (KCL) in multiple Spark jobs to perform analytical jobs.
- e) Use Amazon S3 Select to retrieve the data necessary for the dashboards from the S3 objects.

Answer: c, e

Question: 2

A company needs to implement a near-real-time fraud prevention feature for its ecommerce site.

User and order details need to be delivered to an Amazon SageMaker endpoint to flag suspected fraud. The amount of input data needed for the inference could be as much as 1.5 MB.

Which solution meets the requirements with the LOWEST overall latency?

- a) Create an Amazon Managed Streaming for Kafka cluster and ingest the data for each order into a topic. Use a Kafka consumer running on Amazon EC2 instances to read these messages and invoke the Amazon SageMaker endpoint.
- b) Create an Amazon Kinesis Data Streams stream and ingest the data for each order into the stream. Create an AWS Lambda function to read these messages and invoke the Amazon SageMaker endpoint.
- c) Create an Amazon Kinesis Data Firehose delivery stream and ingest the data for each order into the stream. Configure Kinesis Data Firehose to deliver the data to an Amazon S3 bucket. Trigger an AWS Lambda function with an S3 event notification to read the data and invoke the Amazon SageMaker endpoint.
- d) Create an Amazon SNS topic and publish the data for each order to the topic. Subscribe the Amazon SageMaker endpoint to the SNS topic.

Answer: a

Question: 3

A publisher website captures user activity and sends clickstream data to Amazon Kinesis Data Streams.

The publisher wants to design a cost-effective solution to process the data to create a timeline of user activity within a session. The solution must be able to scale depending on the number of active sessions.

Which solution meets these requirements?

- a) Include a variable in the clickstream data from the publisher website to maintain a counter for the number of active user sessions. Use a timestamp for the partition key for the stream. Configure the consumer application to read the data from the stream and change the number of processor threads based upon the counter. Deploy the consumer application on Amazon EC2 instances in an EC2 Auto Scaling group.
- b) Include a variable in the clickstream to maintain a counter for each user action during their session. Use the action type as the partition key for the stream. Use the Kinesis Client Library (KCL) in the consumer application to retrieve the data from the stream and perform the processing. Configure the consumer application to read the data from the stream and change the number of processor threads based upon the counter. Deploy the consumer application on AWS Lambda.
- c) Include a session identifier in the clickstream data from the publisher website and use as the partition key for the stream. Use the Kinesis Client Library (KCL) in the consumer application to retrieve the data from the stream and perform the processing. Deploy the consumer application on Amazon EC2 instances in an EC2 Auto Scaling group. Use an AWS Lambda function to reshard the stream based upon Amazon CloudWatch alarms.
- d) Include a variable in the clickstream data from the publisher website to maintain a counter for the number of active user sessions. Use a timestamp for the partition key for the stream. Configure the consumer application to read the data from the stream and change the number of processor threads based upon the counter. Deploy the consumer application on AWS Lambda.

Answer: c

Question: 4

A real estate company is receiving new property listing data from its agents through .csv files every day and storing these files in Amazon S3.

The data analytics team created an Amazon QuickSight visualization report that uses a dataset imported from the S3 files. The data analytics team wants the visualization report to reflect the current data up to the previous day.

How can a data analyst meet these requirements?

- a) Schedule an AWS Lambda function to drop and re-create the dataset daily.
- b) Configure the visualization to query the data in Amazon S3 directly without loading the data into SPICE.
- c) Schedule the dataset to refresh daily.
- d) Close and open the Amazon QuickSight visualization.

Answer: c

Question: 5

A data engineer needs to create a dashboard to display social media trends during the last hour of a large company event. The dashboard needs to display the associated metrics with a consistent latency of less than 2 minutes.

Which solution meets these requirements?

- a) Publish the raw social media data to an Amazon Kinesis Data Firehose delivery stream. Use Kinesis Data Analytics for SQL Applications to perform a sliding window analysis to compute the metrics and output the results to a Kinesis Data Streams data stream. Configure an AWS Lambda function to save the stream data to an Amazon DynamoDB table. Deploy a real-time dashboard hosted in an Amazon S3 bucket to read and display the metrics data stored in the DynamoDB table.
- b) Publish the raw social media data to an Amazon Kinesis Data Firehose delivery stream. Configure the stream to deliver the data to an Amazon Elasticsearch Service cluster with a buffer interval of 0 seconds. Use Kibana to perform the analysis and display the results.
- c) Publish the raw social media data to an Amazon Kinesis Data Streams data stream. Configure an AWS Lambda function to compute the metrics on the stream data and save the results in an Amazon S3 bucket. Configure a dashboard in Amazon QuickSight to query the data using Amazon Athena and display the results.
- d) Publish the raw social media data to an Amazon SNS topic. Subscribe an Amazon SQS queue to the topic. Configure Amazon EC2 instances as workers to poll the queue, compute the metrics, and save the results to an Amazon Aurora MySQL database. Configure a dashboard in Amazon QuickSight to query the data in Aurora and display the results.

Answer: a

Question: 6

A media company is migrating its on-premises legacy Hadoop cluster with its associated data processing scripts and workflow to an Amazon EMR environment running the latest Hadoop release. The developers want to reuse the Java code that was written for data processing jobs for the on-premises cluster.

Which approach meets these requirements?

- a) Deploy the existing Oracle Java Archive as a custom bootstrap action and run the job on the EMR cluster.
- b) Compile the Java program for the desired Hadoop version and run it using a CUSTOM_JAR step on the EMR cluster.
- c) Submit the Java program as an Apache Hive or Apache Spark step for the EMR cluster.
- d) Use SSH to connect the master node of the EMR cluster and submit the Java program using the AWS CLI.

Answer: b

Question: 7

A company is providing analytics services to its marketing and human resources (HR) departments. The departments can only access the data through their business intelligence (BI) tools, which run Presto queries on an Amazon EMR cluster that uses the EMR File System (EMRFS).

The marketing data analyst must be granted access to the advertising table only. The HR data analyst must be granted access to the personnel table only.

Which approach will satisfy these requirements?

- a) Create separate IAM roles for the marketing and HR users. Assign the roles with AWS Glue resourcebased policies to access their corresponding tables in the AWS Glue Data Catalog. Configure Presto to use the AWS Glue Data Catalog as the Apache Hive metastore.
- b) Create the marketing and HR users in Apache Ranger. Create separate policies that allow access to the user's corresponding table only. Configure Presto to use Apache Ranger and an external Apache Hive metastore running in Amazon RDS.
- c) Create separate IAM roles for the marketing and HR users. Configure EMR to use IAM roles for EMRFS access. Create a separate bucket for the HR and marketing data. Assign appropriate permissions so the users will only see their corresponding datasets.
- d) Create the marketing and HR users in Apache Ranger. Create separate policies that allows access to the user's corresponding table only. Configure Presto to use Apache Ranger and the AWS Glue Data Catalog as the Apache Hive metastore.

Answer: a

Question: 8

A company ingests a large set of clickstream data in nested JSON format from different sources and stores it in Amazon S3.

Data analysts need to analyze this data in combination with data stored in an Amazon Redshift cluster. Data analysts want to build a cost-effective and automated solution for this need.

Which solution meets these requirements?

- a) Use Apache Spark SQL on Amazon EMR to convert the clickstream data to a tabular format. Use the Amazon Redshift COPY command to load the data into the Amazon Redshift cluster.
- b) Use AWS Lambda to convert the data to a tabular format and write it to Amazon S3. Use the Amazon Redshift COPY command to load the data into the Amazon Redshift cluster.
- c) Use the Relationalize class in an AWS Glue ETL job to transform the data and write the data back to Amazon S3. Use Amazon Redshift Spectrum to create external tables and join with the internal tables.
- d) Use the Amazon Redshift COPY command to move the clickstream data directly into new tables in the Amazon Redshift cluster.

Answer: c

Question: 9

A financial company uses Amazon EMR for its analytics workloads. During the company's annual security audit, the security team determined that none of the EMR clusters' root volumes are encrypted. The security team recommends the company encrypt its EMR clusters' root volume as soon as possible.

Which solution would meet these requirements?

- a) Enable at-rest encryption for EMR File System (EMRFS) data in Amazon S3 in a security configuration. Re-create the cluster using the newly created security configuration.
- b) Specify local disk encryption in a security configuration. Re-create the cluster using the newly created security configuration.
- c) Detach the Amazon EBS volumes from the master node. Encrypt the EBS volume and attach it back to the master node.
- d) Re-create the EMR cluster with LZO encryption enabled on all volumes.

Answer: b

Question: 10

A company is currently using Amazon DynamoDB as the database for a user support application.

The company is developing a new version of the application that will store a PDF file for each support case ranging in size from 1–10 MB. The file should be retrievable whenever the case is accessed in the application.

How can the company store the file in the MOST cost-effective manner?

- a) Store the file in Amazon DocumentDB and the document ID as an attribute in the DynamoDB table.
- b) Store the file in Amazon S3 and the object key as an attribute in the DynamoDB table.
- c) Split the file into smaller parts and store the parts as multiple items in a separate DynamoDB table.
- d) Store the file as an attribute in the DynamoDB table using Base64 encoding.

Answer: b

What Study Guide Works Best in Acing the AWS DAS-C01 Data Analytics Specialty Certification?

The DAS-C01 study guide is a combination of some proven study tips and the combination of all valuable study materials like sample questions, syllabus and practice tests in one place.

Explore the Syllabus Topics and Learn from the Core:

If you are determined to earn success in the Data Analytics Specialty exam, getting in full touch of the [syllabus](#) is mandatory. During preparation, you might not like all syllabus sections or topics, but try to get at least the fundamental knowledge from the sections you don't like. The more you possess knowledge on all syllabus sections, the more is the chance to attempt maximum number of questions during the actual exam.

Make Your Schedule:

Studying and completing the syllabus becomes easier, if you work on the syllabus topics after making a schedule. Your syllabus must mention what areas you want to cover and within what time. Once you make a schedule and follow it regularly, syllabus completion becomes easier and preparation becomes smoother.

Get Expert Advice from the Training:

Do not forget to join the AWS DAS-C01 training if it is providing any. Training enhances the practical knowledge of a candidate, which helps them to work well in the practical field during projects.

Get Access to the PDF Sample Questions:

If your study material is in a [PDF format](#) or the materials are mobile-friendly, what could be better than that? Get access to the free sample questions and keep enhancing your knowledge beyond the syllabus.

Avoid Dumps and Utilize the AWS DAS-C01 Practice Test:

Why should you rely on practice tests? The reason is simple: you must get familiar with the exam pattern before reaching the exam hall. An aspirant aware of the exam structure and time management during the exam preparation can perform well in the actual exam and attempt the maximum number of questions during the exam.

Many aspirants prefer to read from dumps, but they miss out on the self assessment method. Therefore, DAS-C01 practice tests always stand out to be the better choice than dumps PDF.

Avail the Proven DAS-C01 Practice Test for Success!!!

Do you want to pass the DAS-C01 exam on your first attempt? Stop worrying; we, VMExam.com are here to provide you the best experience during your AWS Certified Data Analytics - Specialty preparation. Try out our free mock tests to get a glimpse of our quality study materials, and build your confidence with the premium [DAS-C01 practice tests](#). Our expert-designed questions help you to improve performance and pass the exam on your first attempt.