Exam Amazon DOP-C02 Answers | DOP-C02 Premium Files

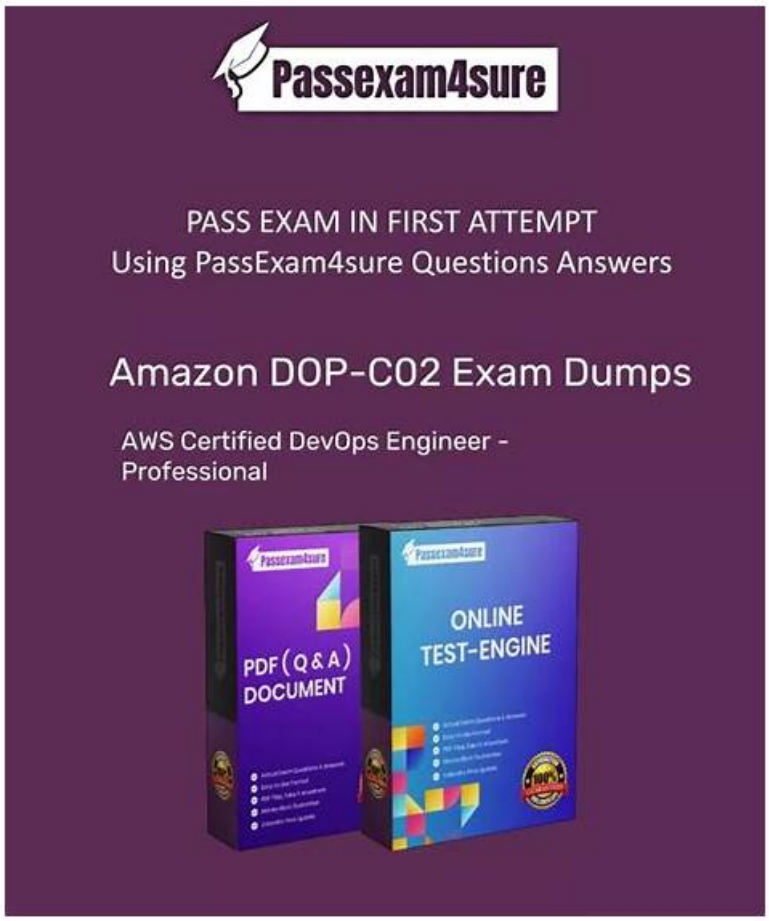

P.S. Free & New DOP-C02 dumps are available on Google Drive shared by ExamCost: https://drive.google.com/open?id=1FvFPCpxTcuSWeOoNBCde6KD5aHMhB-IP

For candidates who want to get the certificate of the exam, choosing a proper DOP-C02 learning material is important. We will provide you the DOP-C02 learning with high accuracy and high quality. If you fail to pass the exam, money back guarantee and it will returning to your account, and if you have any questions about the DOP-C02 Exam Dumps, our online service staff will help to solve any problem you have, just contact us without any hesitation.

To prepare for the Amazon DOP-C02 Certification Exam, candidates can take advantage of various resources provided by AWS, such as training courses, practice exams, and sample questions. Candidates can also leverage their experience with AWS services and DevOps methodologies to prepare for the exam. It is recommended that candidates have at least two years of experience with AWS services and one year of experience with DevOps practices before attempting the certification exam.

>> Exam Amazon DOP-C02 Answers <<

DOP-C02 Premium Files | Test DOP-C02 Centres

Our DOP-C02 guide questions boost many advantages and varied functions. You can have a free download and tryout of our DOP-C02 exam questions before the purchase and our purchase procedures are easy and fast. You can receive our DOP-C02 exam questions in a few minutes and we provide 3 versions for you to choose. You need little time to learn the DOP-C02 Exam Torrent and prepare the exam. Our passing rate and the hit rate is very high. After you pass the DOP-C02 exam you will gain a lot of benefits such as enter in the big company and double your wage.

Amazon DOP-C02 Certification is an excellent way for experienced DevOps professionals to validate their skills and knowledge, enhance their career prospects, and make a valuable contribution to their organizations. If you are interested in this certification, you

can find more information on the AWS website, including study materials, exam details, and registration information.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q61-Q66):

NEW QUESTION # 61
A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.
The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.
Which solution is the MOST cost-effective way to reduce the application startup time?

- A. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- B. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.
- C. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
- D. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Stopped state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.

Answer: D

Explanation:
Option A is the most cost-effective solution. By configuring a warm pool of EC2 instances in the Stopped state, the company can reduce the time it takes for new instances to be ready to serve requests. When the Auto Scaling group launches a new instance, it can attach the stopped EC2 instance from the warm pool. The instance can then be started up immediately, rather than having to wait for the data to be downloaded and processed. This reduces the overall startup time for the application.

NEW QUESTION # 62
A company has a workflow that generates a file for each of the company ' s products and stores the files in a production environment Amazon S3 bucket. The company ' s users can access the S3 bucket.
Each file contains a product ID. Product IDs for products that have not been publicly announced are prefixed with a specific UUID. Product IDs are 12 characters long. IDs for products that have not been publicly announced begin with the letter P.
The company does not want information about products that have not been publicly announced to be available in the production environment S3 bucket.
Which solution will meet these requirements?

- A. Create a new staging S3 bucket. Generate all files in the new staging bucket. Create an Amazon Macie custom data identifier to identify product IDs in the new bucket that begin with the specific UUID.
  Launch an Amazon Macie sensitive data discovery job with the custom data identifier. Copy all files that do not have a Macie finding to the production S3 bucket.
- B. Create an Amazon Macie sensitive data discovery job with a managed data identifier. Remove all files that have a Macie finding from the production S3 bucket.
- C. Create a new staging S3 bucket. Generate all files in the new staging bucket. Launch an Amazon Macie sensitive data discovery job with a managed data identifier. Copy all files that do not have a Macie finding to the production S3 bucket.
- D. Create an Amazon Macie custom data identifier to identify product IDs in the production bucket that begin with the specific UUID. Launch an Amazon Macie sensitive data discovery job with the custom data identifier. Remove all files that have a Macie finding from the production S3 bucket.

Answer: A

Explanation:
The requirement is to prevent any product files containing unannounced product IDs (prefixed with a specific UUID) from being stored in the production S3 bucket that users can access.

To achieve this, a best practice is to use a staging bucket as a control point before files go to production, combined with Amazon Macie's data classification capabilities.
* Creating a custom data identifier in Amazon Macie allows precise detection of product IDs starting with the specific UUID, which default managed identifiers will not detect.
* By running a Macie sensitive data discovery job on the staging bucket, you can identify files containing these sensitive product IDs.
* Only files without findings (i.e., files that do not contain unannounced product IDs) are copied to the production bucket, ensuring no sensitive information is exposed. This approach aligns with AWS best practices for data classification and staged deployment workflows, maximizing control and reducing risk. Using Macie on the production bucket directly (options B and D) risks exposing sensitive data before detection and deletion. Option C uses managed data identifiers, which will likely not detect the custom UUID prefix pattern.
Reference from AWS Official Documentation and Study Guide:
* Amazon Macie Custom Data Identifiers: " You can create custom data identifiers in Amazon Macie to find sensitive data that is unique to your organization. " (Amazon Macie User Guide)
* Data Security Best Practices: " Use staging environments to inspect and sanitize data before moving it to production to reduce exposure risks. " (AWS Security Best Practices)

NEW QUESTION # 63
A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.
The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.
A DevOps engineer must implement a solution to identify in near real time any AWS service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.
Which solution will meet these requirements with the LEAST development overhead?

- A. Use CloudFormation drift detection to identify noncompliant resources. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediation. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log group. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
- B. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resources. Enable AWS Security Hub with the ~no-enable-default-standards option in all the AWS accounts. Set up AWS Config managed rules and custom rules. Set up automatic remediation by using AWS Config conformance packs. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
- C. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resources. Use CloudWatch Logs filters for drift detection. Use Amazon EventBridge to invoke the Lambda function for remediation. Stream filtered CloudWatch logs to Amazon OpenSearch Service. Set up a dashboard on OpenSearch Service for tracking.
- D. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resources. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediation. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.

Answer: B

Explanation:
The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.
The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.
Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function

has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.
:
AWS Config conformance packs
Introducing AWS Config conformance packs
Managing conformance packs across all accounts in your organization

NEW QUESTION # 64
A video-sharing company stores its videos in an Amazon S3 bucket. The company needs to analyze user access patterns such as the number of users who access a specific video each month.
Which solution will meet these requirements with the LEAST development effort?

- A. Record a log message in Amazon CloudWatch Logs for every S3 object access event. Configure a log stream in CloudWatch Logs to write the file access information, including user ID, S3 bucket ID, and file key, to an Amazon Managed Service for Apache Flink application. Perform a sliding window analysis on the user access patterns.
- B. Enable Amazon S3 server access logging. Load the access logs into an Amazon Aurora database. Run SQL queries on the Aurora database to analyze the user access patterns.
- C. Invoke an AWS Lambda function for every S3 object access event. Configure the Lambda function to write the file access information, including user ID, S3 bucket ID, and file key, to an Amazon Aurora database. Run SQL queries on the Aurora database to analyze the user access patterns.
- D. Enable Amazon S3 server access logging. Use Amazon Athena to create an external table that contains the access logs. Run SQL queries on the Athena table to analyze the user access patterns.

Answer: D

Explanation:
Amazon S3 can generate server access logs that record detailed information about each request, including requester, bucket, key, operation, time, and status. These logs are written as objects to an S3 bucket. To analyze access patterns, the simplest and most serverless approach is to use Amazon Athena directly on those logs without building ingestion pipelines or databases.
Option B enables S3 server access logging and then creates an Athena external table over the log bucket.
AWS provides standard log formats and even example schemas for S3 access logs. The analytics team can run ad hoc SQL queries to count the number of accesses per object per time period, filter by user, and perform aggregations, all without provisioning compute or managing databases.
Option A requires ingesting logs into Aurora, which adds ETL complexity and ongoing database management. Option C requires a Lambda function for every access event plus DB writes, which is more complex and potentially expensive at scale. Option D uses CloudWatch Logs and Managed Flink, which is more suited for streaming analytics and is significantly more complex than necessary for monthly summary reports.
Therefore, Option B provides the required analysis with the least development and operational effort.

NEW QUESTION # 65
A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.
Which solution will accomplish this?

- A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoints. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
- B. Store the Aurora endpoint in AWS Systems Manager Parameter Store. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store. Code the application to reload the endpoint from Parameter Store if a database connection fails.

- C. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
- D. Create an Aurora custom endpoint to point to the primary database instance. Configure the application to use this endpoint. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.

Answer: B

NEW QUESTION # 66

.....

DOP-C02 Premium Files: https://www.examcost.com/DOP-C02-practice-exam.html

- Fast-Download Exam DOP-C02 Answers - Pass DOP-C02 Once - First-Grade DOP-C02 Premium Files □ Search for ➠ DOP-C02 □ and obtain a free download on ➤ www.troytecdumps.com □ □DOP-C02 Latest Dumps Free
- High Hit-Rate Amazon - DOP-C02 - Exam AWS Certified DevOps Engineer - Professional Answers □ Open 【 www.pdfvce.com 】 and search for ➽ DOP-C02 □ to download exam materials for free □DOP-C02 Certification Sample Questions
- DOP-C02 Exam Material □ DOP-C02 New Dumps Free □ DOP-C02 New Dumps Ebook □ Search for ➥ DOP-C02 □ on { www.prepawaypdf.com } immediately to obtain a free download □New DOP-C02 Test Cram
- Latest updated Exam DOP-C02 Answers | Easy To Study and Pass Exam at first attempt - Hot Amazon AWS Certified DevOps Engineer - Professional □ Open ➤ www.pdfvce.com □ and search for ▹ DOP-C02 ◃ to download exam materials for free □DOP-C02 Latest Dumps Free
- Real DOP-C02 Dumps □ Practice DOP-C02 Test Engine □ DOP-C02 Latest Exam Test □ Search for ⇛ DOP-C02 ⇚ and download it for free immediately on ☀ www.easy4engine.com □☀□ □Real DOP-C02 Dumps
- DOP-C02 New Dumps Free □ DOP-C02 Latest Exam Test □ Formal DOP-C02 Test □ Open [www.pdfvce.com] enter ➠ DOP-C02 □ and obtain a free download □Updated DOP-C02 CBT
- Real DOP-C02 Dumps □ DOP-C02 Valid Exam Question □ Valid DOP-C02 Exam Prep □ Search for □ DOP-C02 □ and download it for free on ✔ www.prep4away.com □✔□ website □DOP-C02 New Dumps Ebook
- Using Amazon DOP-C02 Dumps, Improve Your Exam Skills □ Open website ✔ www.pdfvce.com □✔□ and search for ✔ DOP-C02 □✔□ for free download □Reliable DOP-C02 Exam Tips
- DOP-C02 Valid Exam Question □ DOP-C02 New Dumps Ebook □ VCE DOP-C02 Exam Simulator ↪ Download □ DOP-C02 □ for free by simply entering □ www.testkingpass.com □ website □DOP-C02 New Dumps Ebook
- Latest updated Exam DOP-C02 Answers | Easy To Study and Pass Exam at first attempt - Hot Amazon AWS Certified DevOps Engineer - Professional □ Search for " DOP-C02 " and download it for free on [www.pdfvce.com] website □ □VCE DOP-C02 Exam Simulator
- www.troytecdumps.com Latest DOP-C02 Dumps Will Help You Build A SuccessFul Career □ Search for 《 DOP-C02 》 and download exam materials for free through □ www.troytecdumps.com □ □DOP-C02 Latest Dumps Free
- jimmgdy451372.blogars.com, lawsonqnjy417179.activablog.com, safatxfg460928.blog-gold.com, ronaldbbko397061.blogdosaga.com, bookmarkrange.com, directoryarmy.com, bookmark-media.com, nanaulqg476821.blgwiki.com, jasondhjz176493.ttblogs.com, mariahrwnj583476.get-blogging.com, Disposable vapes

BONUS!!! Download part of ExamCost DOP-C02 dumps for free: https://drive.google.com/open?id=1FvFPCpxTcuSWeOoNBCde6KD5aHMhB-IP