SAP-C02 Valid Test Fee & Reliable SAP-C02 Test Sims

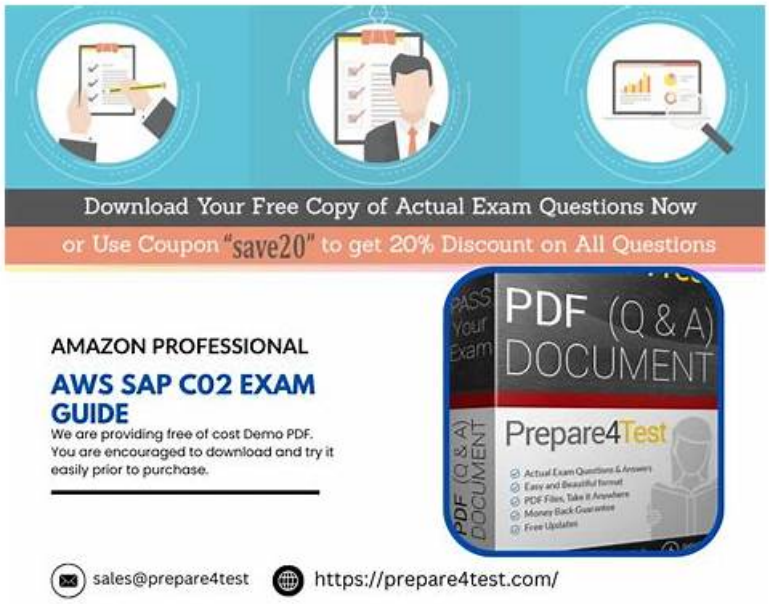

DOWNLOAD the newest Actual4Labs SAP-C02 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1Umi2Yn8V9dnub-gofjEG0lWWDWI1fMMp

In today's society, there are increasingly thousands of people put a priority to acquire certificates to enhance their abilities. With a total new perspective, SAP-C02 study materials have been designed to serve most of the office workers who aim at getting a SAP-C02 certification. Our SAP-C02 Test Guide keep pace with contemporary talent development and makes every learner fit in the needs of the society. There is no doubt that our SAP-C02 latest question can be your first choice for your relevant knowledge accumulation and ability enhancement.

The AWS Certified Solutions Architect - Professional (SAP-C02) certification exam is designed to validate the advanced technical skills and experience of individuals in designing and deploying AWS solutions. AWS Certified Solutions Architect - Professional (SAP-C02) certification is intended for professionals who have already earned the AWS Certified Solutions Architect - Associate certification and have hands-on experience in designing and deploying scalable, fault-tolerant, and highly available AWS systems.

>> SAP-C02 Valid Test Fee <<

Reliable Amazon SAP-C02 Test Sims, Dumps SAP-C02 Free Download

Our SAP-C02 exam guide have also set a series of explanation about the complicated parts certificated by the syllabus and are based on the actual situation to stimulate exam circumstance in order to provide you a high-quality and high-efficiency user experience. In addition, the SAP-C02 exam guide function as a time-counter, and you can set fixed time to fulfill your task, so that promote your efficiency in real test. The key strong-point of our SAP-C02 Test Guide is that we impart more important knowledge with fewer questions and answers, with those easily understandable SAP-C02 study braindumps, you will find more interests in them and experience an easy learning process.

To prepare for the SAP-C02 Exam, you will need to have a deep understanding of AWS services and how they work together. You will also need to be familiar with AWS tools, such as CloudFormation, Elastic Beanstalk, and OpsWorks, as well as other third-party tools that integrate with AWS. Additionally, you should have experience with designing and deploying highly available and fault-tolerant systems on AWS.

Amazon AWS Certified Solutions Architect - Professional (SAP-C02) Sample Questions (Q200-Q205):

NEW QUESTION # 200
A company wants to create a single Amazon S3 bucket for its data scientists to store work-related documents.

The company uses AWS 1AM Identity Center to authenticate all users. A group for the data scientists was created.
The company wants to give the data scientists access lo only their own work. The company also wants to create monthly reports that show which documents each user accessed.
Which combination of steps will meet these requirements? (Select TWO.)

- A. Configure AWS CloudTrail to log S3 data events and deliver the logs to an S3 bucket. Use Amazon Athena to run queries on the CloudTrail logs in Amazon S3 and generate reports.
- B. Create a custom 1AM Identity Center permission set to grant the data scientists access to an S3 bucket prefix that matches their username tag. Use a policy to limit access to paths with the S{aws:PrincipalTag /userName}/* condition.
- C. Enable S3 access logging to EMR File System (EMRFS). Use Amazon S3 Select to query logs and generate reports.
- D. Configure AWS CloudTrail to log S3 management events to CloudWatch. Use Amazon Athena's CloudWatch connector to query the logs and generate reports.
- E. Create an 1AM Identity Center role for the data scientists group that has Amazon S3 read access and write access. Add an S3 bucket policy that allows access to the 1AM Identity Center role.

Answer: A,B

NEW QUESTION # 201
A software company has deployed an application that consumes a REST API by using Amazon API Gateway.
AWS Lambda functions, and an Amazon DynamoDB table. The application is showing an increase in the number of errors during PUT requests. Most of the PUT calls come from a small number of clients that are authenticated with specific API keys.
A solutions architect has identified that a large number of the PUT requests originate from one client. The API is noncritical, and clients can tolerate retries of unsuccessful calls. However, the errors are displayed to customers and are causing damage to the API's reputation.
What should the solutions architect recommend to improve the customer experience?

- A. Implement retry logic with exponential backoff and irregular variation in the client application. Ensure that the errors are caught and handled with descriptive error messages.
- B. Implement API throttling through a usage plan at the API Gateway level. Ensure that the client application handles code 429 replies without error.
- C. Implement reserved concurrency at the Lambda function level to provide the resources that are needed during sudden increases in traffic.
- D. Turn on API caching to enhance responsiveness for the production stage. Run 10-minute load tests. Verify that the cache capacity is appropriate for the workload.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/aws-batch-requests-error/
https://aws.amazon.com/premiumsupport/knowledge-center/api-gateway-429-limit/ The main problem is that one client is generating a large number of PUT requests, resulting in increased errors. Because the clients are authenticated by API keys and the problematic traffic is concentrated among a small set of clients, API Gateway usage plans and throttling are a direct control mechanism to limit the request rate per client (or per API key). The API is noncritical and clients can tolerate retries, so returning throttling responses is acceptable as long as clients handle them correctly. The goal is to reduce visible customer-facing errors and protect the API from being overwhelmed by a single client.
API Gateway usage plans can set request throttling limits (rate and burst) per API key. When a client exceeds the configured throttle, API Gateway returns HTTP 429 (Too Many Requests). If the client application is designed to treat 429 as a normal, retryable condition rather than a fatal error, the user experience improves because the client can back off and retry instead of surfacing raw failures. This approach prevents one noisy client from degrading service for other users and reduces the overall error rate.
Option B addresses the root cause (excess requests from a particular client) by enforcing throttling at the API Gateway layer and improving client behavior when throttled.
Option A focuses on client-side retries and error handling but does not enforce fairness or prevent a single client from consuming disproportionate API capacity. Without throttling, the same client can continue to overload downstream Lambda and DynamoDB capacity and still cause errors for everyone.
Option C (API caching) does not help for PUT requests because PUT operations are writes and are generally not cacheable in a way that reduces write load. Caching is mainly beneficial for GET responses and read- heavy workloads.
Option D (reserved concurrency) can protect Lambda capacity for critical functions, but it can also cause throttling at the Lambda layer and does not directly address controlling the request rate per API key. It also does not fix the issue at the API front door

where the noisy client should be managed.
Therefore, implementing API Gateway throttling via a usage plan and ensuring clients handle 429 responses properly is the best recommendation.
References:AWS documentation on Amazon API Gateway usage plans and API key throttling, including rate and burst limits and 429 responses.AWS best practices for handling throttling responses with retries and backoff for improved client experience.

NEW QUESTION # 202
A company has an application that uses Amazon EC2 instances in an Auto Scaling group. The quality assurance (QA) department needs to launch a large number of short-lived environments to test the application. The application environments are currently launched by the manager of the department using an AWS CloudFormation template. To launch the stack, the manager uses a role with permission to use CloudFormation, EC2, and Auto Scaling APIs. The manager wants to allow testers to launch their own environments, but does not want to grant broad permissions to each user.
Which set up would achieve these goals?

- A. Create an AWS Service Catalog product from the environment template. Add a launch constraint to the product with the existing role. Give users in the QA department permission to use AWS Service Catalog APIs only. Train users to launch the template from the AWS Service Catalog console.
- B. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to use CloudFormation and S3 APIs, with conditions that restrict the permissions to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
- C. Upload the AWS CloudFormation template to Amazon S3. Give users in the QA department permission to assume the manager's role and add a policy that restricts the permissions to the template and the resources it creates. Train users to launch the template from the CloudFormation console.
- D. Create an AWS Elastic Beanstalk application from the environment template. Give users in the QA department permission to use Elastic Beanstalk permissions only. Train users to launch Elastic Beanstalk environments with the Elastic Beanstalk CLI, passing the existing role to the environment as a service role.

Answer: A

NEW QUESTION # 203
A company has multiple AWS accounts as part of an organization created with AWS Organizations. Each account has a VPC in the us-east-2 Region and is used for either production or development workloads.
Amazon EC2 instances across production accounts need to communicate with each other, and EC2 instances across development accounts need to communicate with each other, but production and development instances should not be able to communicate with each other.
To facilitate connectivity, the company created a common network account. The company used AWS Transit Gateway to create a transit gateway in the us-east-2 Region in the network account and shared the transit gateway with the entire organization by using AWS Resource Access Manager. Network administrators then attached VPCs in each account to the transit gateway, after which the EC2 instances were able to communicate across accounts. However, production and development accounts were also able to communicate with one another.
Which set of steps should a solutions architect take to ensure production traffic and development traffic are completely isolated?

- A. Modify the security groups assigned to development EC2 instances to block traffic from production EC2 instances. Modify the security groups assigned to production EC2 instances to block traffic from development EC2 instances.
- B. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Modify the AWS Transit Gateway routing table to route production tagged attachments to one another and development tagged attachments to one another.
- C. Create a tag on each VPC attachment with a value of either production or development, according to the type of account being attached. Using the Network Manager feature of AWS Transit Gateway, create policies that restrict traffic between VPCs based on the value of this tag.
- D. Create separate route tables for production and development traffic. Delete each account's association and route propagation to the default AWS Transit Gateway route table. Attach development VPCs to the development AWS Transit Gateway route table and production VPCs to the production route table, and enable automatic route propagation on each attachment.

Answer: D

Explanation:
Explanation

https://docs.aws.amazon.com/vpc/latest/tgw/vpc-tgw.pdf

NEW QUESTION # 204
A company is processing videos in the AWS Cloud by using Amazon EC2 instances in an Auto Scaling group. It takes 30 minutes to process a video. Several EC2 instances scale in and out depending on the number of videos in an Amazon Simple Queue Service (Amazon SQS) queue.
The company has configured the SQS queue with a redrive policy that specifies a target dead- letter queue and a maxReceiveCount of 1. The company has set the visibility timeout for the SQS queue to 1 hour. The company has set up an Amazon CloudWatch alarm to notify the development team when there are messages in the dead-letter queue.
Several times during the day, the development team receives notification that messages are in the dead-letter queue and that videos have not been processed properly. An investigation finds no errors in the application logs.
How can the company solve this problem?

- A. Turn on termination protection for the EC2 instances.
- B. Configure scale-in protection for the instances during processing.
- C. Update the redrive policy and set maxReceiveCount to 0.
- D. Update the visibility timeout for the SOS queue to 3 hours.

Answer: B

NEW QUESTION # 205
.....