

New SAP-C02 Dumps Book | Updated SAP-C02 CBT



SAP-C02
AWS Certified Solutions
Architect
Professional

Certification Questions
& Exams Dumps

www.edurely.com

BONUS!!! Download part of PassLeaderVCE SAP-C02 dumps for free: <https://drive.google.com/open?id=1KZiR3OFs2ICZz6mRLNgKccyAh6PcQxV>

The Amazon SAP-C02 web-based practice test software is very user-friendly and simple to use. It is accessible on all browsers. It will save your progress and give a report of your mistakes which will surely be beneficial for your overall exam preparation. A useful certification will bring you much outstanding advantage when you apply for any jobs about Amazon company or products.

The AWS Certified Solutions Architect - Professional (SAP-C02) exam covers a wide range of topics, including AWS services and their integration, AWS architecture best practices, disaster recovery and business continuity, migration strategies, and security best practices. SAP-C02 Exam also tests the candidate's knowledge of AWS cost optimization, deployment and management, and troubleshooting.

>> [New SAP-C02 Dumps Book](#) <<

Pass Guaranteed Quiz 2026 Amazon The Best SAP-C02: New AWS Certified Solutions Architect - Professional (SAP-C02) Dumps Book

If you buy our Software version of the SAP-C02 study questions, you can enjoy the similar real exam environment for that this version has the advantage of simulating the real exam. In addition, the software version of our SAP-C02 learning guide is not limited to the number of the computer. As long as you use it on the Windows system, then you can enjoy the convenience of this version brings. So do not hesitate and buy our Software version of SAP-C02 Preparation exam, you will benefit a lot from it.

The SAP-C02 Certification is recognized by companies around the world as a measure of a candidate's advanced skills in AWS architecture. Obtaining this certification can help IT professionals advance their careers by demonstrating their expertise in designing and deploying complex cloud-based solutions on AWS. Additionally, the certification can help candidates stand out in a competitive job market by demonstrating their commitment to continuous learning and professional development.

Amazon AWS Certified Solutions Architect - Professional (SAP-C02) Sample Questions (Q402-Q407):

NEW QUESTION # 402

A retail company is operating its ecommerce application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company uses an Amazon RDS DB instance as the database backend. Amazon CloudFront is configured with one origin that points to the ALB. Static content is cached. Amazon Route 53 is used to host all public zones. After an update of the application, the ALB occasionally returns a 502 status code (Bad Gateway) error. The root cause is malformed HTTP headers that are returned to the ALB. The webpage returns successfully when a solutions architect reloads the webpage immediately after the error occurs.

While the company is working on the problem, the solutions architect needs to provide a custom error page instead of the standard ALB error page to visitors.

Which combination of steps will meet this requirement with the LEAST amount of operational overhead?

(Choose two.)

- A. **Modify the existing Amazon Route 53 records by adding health checks. Configure a fallback target if the health check fails. Modify DNS records to point to a publicly accessible webpage.**
- B. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Target.FailedHealthChecks is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a publicly accessible web server.
- C. Create an Amazon S3 bucket. Configure the S3 bucket to host a static webpage. Upload the custom error pages to Amazon S3.
- D. Create an Amazon CloudWatch alarm to invoke an AWS Lambda function if the ALB health check response Elb.InternalError is greater than 0. Configure the Lambda function to modify the forwarding rule at the ALB to point to a public accessible web server.
- E. **Add a custom error response by configuring a CloudFront custom error page. Modify DNS records to point to a publicly accessible web page.**

Answer: A,E

Explanation:

"Save your custom error pages in a location that is accessible to CloudFront. We recommend that you store them in an Amazon S3 bucket, and that you don't store them in the same place as the rest of your website or application's content. If you store the custom error pages on the same origin as your website or application, and the origin starts to return 5xx errors, CloudFront can't get the custom error pages because the origin server is unavailable."

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/GeneratingCustomErrorResponses.html>

NEW QUESTION # 403

A company is running a traditional web application on Amazon EC2 instances. The company needs to refactor the application as microservices that run on containers. Separate versions of the application exist in two distinct environments: production and testing. Load for the application is variable, but the minimum load and the maximum load are known. A solutions architect needs to design the updated application with a serverless architecture that minimizes operational complexity.

Which solution will meet these requirements MOST cost-effectively?

- A. Upload the container images to AWS Elastic Beanstalk. In Elastic Beanstalk, create separate environments and deployments for production and testing. Configure two separate Application Load Balancers to direct traffic to the Elastic Beanstalk deployments.
- B. Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Kubernetes Service (Amazon EKS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the EKS clusters.
- C. **Upload the container images to Amazon Elastic Container Registry (Amazon ECR). Configure two auto scaled Amazon Elastic Container Service (Amazon ECS) clusters with the Fargate launch type to handle the expected load. Deploy tasks from the ECR images. Configure two separate Application Load Balancers to direct traffic to the ECS clusters.**
- D. Upload the container images to AWS Lambda as functions. Configure a concurrency limit for the associated Lambda functions to handle the expected peak load. Configure two separate Lambda integrations within Amazon API Gateway: one for production and one for testing.

Answer: C

Explanation:

minimizes operational + microservices that run on containers = AWS Elastic Beanstalk

NEW QUESTION # 404

A solutions architect has been assigned to migrate a 50 TB Oracle data warehouse that contains sales data from on-premises to Amazon Redshift. Major updates to the sales data occur on the final calendar day of the month. For the remainder of the month, the data warehouse only receives minor daily updates and is primarily used for reading and reporting. Because of this, the migration process must start on the first day of the month and must be complete before the next set of updates occur. This provides approximately 30 days to complete the migration and ensure that the minor daily changes have been synchronized with the Amazon Redshift data warehouse.

Because the migration cannot impact normal business network operations, the bandwidth allocated to the migration for moving data over the internet is 50 Mbps. The company wants to keep data migration costs low.

Which steps will allow the solutions architect to perform the migration within the specified timeline?

- A. Install Oracle database software on an Amazon EC2 instance.

To minimize the migration time, configure VPN connectivity between AWS and the company's data center by provisioning a 1 Gbps AWS Direct Connect connection.

Configure the Oracle database running on Amazon EC2 to be a read replica of the data center Oracle database.

Start the synchronization process between the company's on-premises data center and the Oracle database on Amazon EC2.

When the Oracle database on Amazon EC2 is synchronized with the on-premises database, create an AWS DMS ongoing replication task to migrate the data from the Oracle database read replica that is running on Amazon EC2 to Amazon Redshift.

Verify the data migration is complete and perform the cut over to Amazon Redshift.

- B. Create an AWS Snowball import job. Export a backup of the Oracle data warehouse.

Copy the exported data to the Snowball device. Return the Snowball device to AWS.

Create an Amazon RDS for Oracle database and restore the backup file to that RDS instance.

Create an AWS DMS task to migrate the data from the RDS for Oracle database to Amazon Redshift.

Copy daily incremental backups from Oracle in the data center to the RDS for Oracle database over the internet.

Verify the data migration is complete and perform the cut over to Amazon Redshift.

- C. Install Oracle database software on an Amazon EC2 instance.

Configure VPN connectivity between AWS and the company's data center.

Configure the Oracle database running on Amazon EC2 to join the Oracle Real Application Clusters (RAC).

When the Oracle database on Amazon EC2 finishes synchronizing, create an AWS DMS ongoing replication task to migrate the data from the Oracle database on Amazon EC2 to Amazon Redshift.

Verify the data migration is complete and perform the cut over to Amazon Redshift.

- D. Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent.

Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema.

Create a new project in AWS SCT using the registered data extraction agent.

Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes.

Copy data to the Snowball device and return the Snowball device to AWS.

Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift.

Verify that the data migration is complete and perform the cut over to Amazon Redshift.

Answer: D

Explanation:

Create an AWS Snowball import job. Configure a server in the company's data center with an extraction agent. Use AWS SCT to manage the extraction agent and convert the Oracle schema to an Amazon Redshift schema. Create a new project in AWS SCT using the registered data extraction agent. Create a local task and an AWS DMS task in AWS SCT with replication of ongoing changes. Copy data to the Snowball device and return the Snowball device to AWS.

Allow AWS DMS to copy data from Amazon S3 to Amazon Redshift. Verify that the data migration is complete and perform the cut over to Amazon Redshift.

<https://aws.amazon.com/getting-started/hands-on/migrate-oracle-to-amazon-redshift/>

NEW QUESTION # 405

A company is providing weather data over a REST-based API to several customers. The API is hosted by Amazon API Gateway and is integrated with different AWS Lambda functions for each API operation. The company uses Amazon Route 53 for DNS and has created a resource record of weather.example.com. The company stores data for the API in Amazon DynamoDB tables. The company needs a solution that will give the API the ability to fail over to a different AWS Region.

Which solution will meet these requirements?

- A. Deploy a new API Gateway API in a new Region. Change the Lambda functions to global functions.

Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.

- B. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a multivalue answer. Add both API Gateway APIs to the answer. Enable target health monitoring. Convert the DynamoDB tables to global tables.
- C. Deploy a new API Gateway API and Lambda functions in another Region. Change the Route 53 DNS record to a failover record. Enable target health monitoring. Convert the DynamoDB tables to global tables.
- D. Deploy a new set of Lambda functions in a new Region. Update the API Gateway API to use an edge-optimized API endpoint with Lambda functions from both Regions as targets. Convert the DynamoDB tables to global tables.

Answer: C

Explanation:

Explanation

<https://docs.aws.amazon.com/apigateway/latest/developerguide/dns-failover.html>

NEW QUESTION # 406

A solutions architect is preparing to deploy a new security tool into several previously unused AWS Regions. The solutions architect will deploy the tool by using an AWS CloudFormation stack set. The stack set's template contains an IAM role that has a custom name. Upon creation of the stack set, no stack instances are created successfully.

What should the solutions architect do to deploy the stacks successfully?

- A. Specify the CAPABILITY_NAMED_IAM capability and the SELF_MANAGED permissions model during the creation of the stack set.
- B. Use the Service Quotas console to request a quota increase for the number of CloudFormation stacks in each new Region in all relevant accounts. Specify the CAPABILITY_IAM capability during the creation of the stack set.
- C. Enable the new Regions in all relevant accounts. Specify the CAPABILITY_NAMED_IAM capability during the creation of the stack set.
- D. Specify an administration role ARN and the CAPABILITY_IAM capability during the creation of the stack set.

Answer: C

Explanation:

The CAPABILITY_NAMED_IAM capability is required when creating or updating CloudFormation stacks that contain IAM resources with custom names. This capability acknowledges that the template might create IAM resources that have broad permissions or affect other resources in the AWS account. The stack set's template contains an IAM role that has a custom name, so this capability is needed. Enabling the new Regions in all relevant accounts is also necessary to deploy the stack set across multiple Regions and accounts.

Option B is incorrect because the Service Quotas console is used to view and manage the quotas for AWS services, not for CloudFormation stacks. The number of stacks per Region per account is not a service quota that can be increased.

Option C is incorrect because the SELF_MANAGED permissions model is used when the administrator wants to retain full permissions to manage stack sets and stack instances. This model does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.

Option D is incorrect because an administration role ARN is optional when creating a stack set. It is used to specify a role that CloudFormation assumes to create stack instances in the target accounts. It does not affect the creation of the stack set or the requirement for the CAPABILITY_NAMED_IAM capability.

1: AWS CloudFormation stack sets

2: Acknowledging IAM resources in AWS CloudFormation templates

3: AWS CloudFormation stack set permissions

NEW QUESTION # 407

.....

Updated SAP-C02 CBT: <https://www.passleadervce.com/AWS-Certified-Solutions-Architect/reliable-SAP-C02-exam-learning-guide.html>

- SAP-C02 Real Dumps Free ↗ Guide SAP-C02 Torrent □ SAP-C02 Real Dumps Free □ Open ► www.verifieddumps.com ▲ and search for ⇒ SAP-C02 ⇍ to download exam materials for free □ SAP-C02 Test Questions Vce

2026 Latest PassLeaderVCE SAP-C02 PDF Dumps and SAP-C02 Exam Engine Free Share: <https://drive.google.com/open?id=1KZiR3OFs2ICZz6mRLNgKccyAh6PcQxV>