SAP-C02 Reliable Exam Voucher | SAP-C02 Test Valid

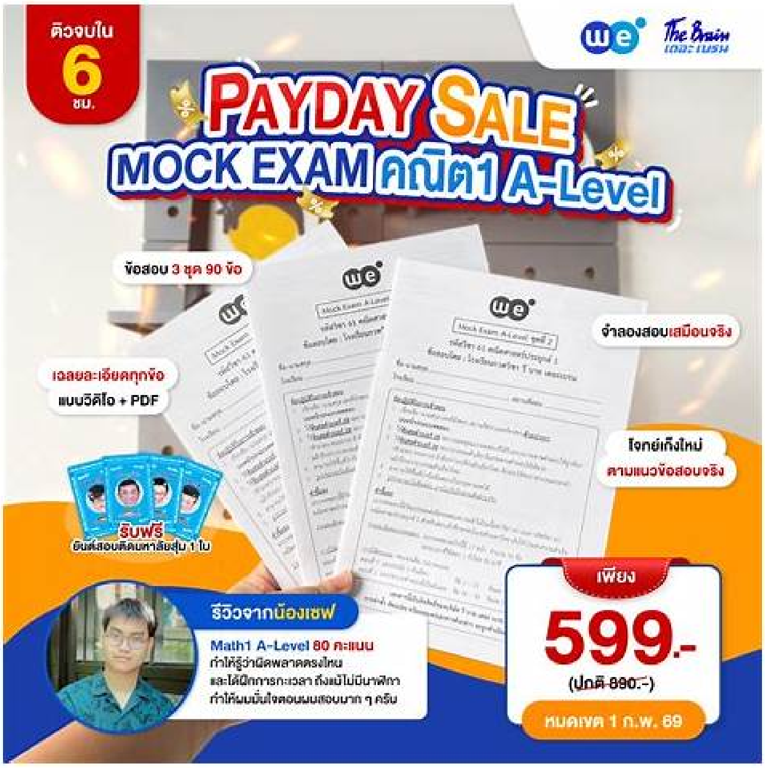

DOWNLOAD the newest TestKingIT SAP-C02 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=10PqNyDsHedn7bpLRMbI-KtfOqnrxlLdK

The study material is available in three formats, i.e. PDF format, web-based practice exam, and desktop practice test software. The PDF format is easy for those who always have their smart devices and love to study from them. Users can also make notes of printed PDF Amazon AWS Certified Solutions Architect - Professional (SAP-C02) certification exam so they can study them anywhere to pass Amazon SAP-C02 Certification test with a good score.

Amazon SAP-C02 Exam is a certification exam aimed at IT professionals who wish to validate their skills and knowledge in designing and deploying AWS solutions. It is the second version of the AWS Certified Solutions Architect - Professional exam and is designed to test an individual's ability to design and deploy scalable, cost-effective, and fault-tolerant systems on AWS.

>> SAP-C02 Reliable Exam Voucher <<

Amazon SAP-C02 Test Valid | Test SAP-C02 Online

It is really a tough work to getting SAP-C02 certification in their spare time because preparing actual exam dumps needs plenty time and energy. As the one of certification exam dumps provider, TestKingIT enjoys a high popularity for its profession of SAP-C02 Exam Dumps and training materials. You will get high passing score in test with the help of our SAP-C02 braindumps torrent.

Passing the SAP-C02 Exam can be a valuable asset for professionals looking to advance their career in the field of AWS solutions architecture. It demonstrates to employers and clients that the individual has the knowledge and skills required to design and deploy complex AWS solutions using best practices and architectural principles. Additionally, it can increase the individual's earning potential and open up new career opportunities.

Amazon AWS Certified Solutions Architect - Professional (SAP-C02) Sample Questions (Q308-Q313):

NEW QUESTION # 308
A media company has a 30-TB repository of digital news videos These videos are stored on tape in an on-premises tape library and referenced by a Media Asset Management (MAM) system. The company wants to enrich the metadata for these videos in an

automated fashion and put them into a searchable catalog by using a MAM feature. The company must be able to search based on information in the video such as objects scenery items or people's faces A catalog is available that contains faces of people who have appeared in the videos that include an image of each person. The company would like to migrate these videos to AWS
The company has a high-speed AWS Direct Connect connection with AWS and would like to move the MAM solution video content directly from its current file system
How can these requirements be met by using the LEAST amount of ongoing management overhead and causing MINIMAL disruption to the existing system"

- A. Configure a video ingestion stream by using Amazon Kinesis Video Streams Use the catalog of faces to build a collection in Amazon Rekognition Stream the videos from the MAM solution into Kinesis Video Streams Configure Amazon Rekognition to process the streamed videos Then, use a stream consumer to retrieve the required metadata and push the metadata into the MAM solution Configure the stream to store the videos in Amazon S3
- B. Set up an AWS Storage Gateway file gateway appliance on-premises. Use the MAM solution to extract the videos from the current archive and push them into the file gateway Use the catalog of faces to build a collection in Amazon Rekognition Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Rekognition pull the video from the Amazon S3 files backing the file gateway, retrieve the required metadata and push the metadata into the MAM solution
- C. Set up an Amazon EC2 instance that runs the OpenCV libranes Copy the videos, images, and face catalog from the on-premises library into an Amazon EBS volume mounted on this EC2 instance Process the videos to retrieve the required metadata, and push the metadata into the MAM solution, while also copying the video files to an Amazon S3 bucket
- D. Set up an AWS Storage Gateway tape gateway appliance on-premises Use the MAM solution to extract the videos from the current archive and push them into the tape gateway Use the catalog of faces to build a collection in Amazon Rekognition Build an AWS Lambda function that invokes the Rekognition Javascript SDK to have Amazon Rekognition process the video in the tape gateway retrieve the required metadata, and push the metadata into the MAM solution

Answer: A

NEW QUESTION # 309
A solutions architect has deployed a web application that serves users across two AWS Regions under a custom domain The application uses Amazon Route 53 latency-based routing The solutions architect has associated weighted record sets with a pair of web servers in separate Availability Zones for each Region The solutions architect runs a disaster recovery scenario When all the web servers in one Region are stopped Route 53 does not automatically redirect users to the other Region Which of the following are possible root causes of this issue? (Select TWO.)

- A. One of the web servers in the secondary Region did not pass its HTTP health check
- B. Latency resource record sets cannot be used in combination with weighted resource record sets
- C. An HTTP health check has not been set up for one or more of the weighted resource record sets associated with the stopped web servers
- D. The weight for the Region where the web servers were stopped is higher than the weight for the other Region
- E. The setting to evaluate target health is not turned on for the latency alias resource record set that is associated with the domain in the Region where the web servers were stopped

Answer: B,E

Explanation:
Explanation
Latency resource record sets cannot be used in combination with weighted resource record sets in Route 53, so this is the likely root cause of the issue. Additionally, Route 53 needs to be configured to evaluate target health for the latency alias resource record set that is associated with the domain in the Region where the web servers were stopped. If this setting is not turned on, Route 53 will not automatically redirect users to the other Region. Reference: AWS Certified Solutions Architect Professional Official Text Book, Chapter 5, Section
5.2.2.

NEW QUESTION # 310
A company has application services that have been containerized and deployed on multiple Amazon EC2 instances with public IPs. An Apache Kafka cluster has been deployed to the EC2 instances. A PostgreSQL database has been migrated to Amazon RDS for PostgreSQL. The company expects a significant increase of orders on its platform when a new version of its flagship product is released.
What changes to the current architecture will reduce operational overhead and support the product release?

- A. Create an EC2 Auto Scaling group behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- B. Deploy the application on a Kubernetes cluster created on the EC2 instances behind an Application Load Balancer. Deploy the DB instance in Multi-AZ mode and enable storage auto scaling. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
- C. Create an EC2 Auto Scaling group behind an Application Load Balancer. Create additional read replicas for the DB instance. Create Amazon Kinesis data streams and configure the application services to use the data streams. Store and serve static content directly from Amazon S3.
- D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.

Answer: D

Explanation:
The correct answer is D. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS) with AWS Fargate and enable auto scaling behind an Application Load Balancer. Create additional read replicas for the DB instance. Create an Amazon Managed Streaming for Apache Kafka cluster and configure the application services to use the cluster. Store static content in Amazon S3 behind an Amazon CloudFront distribution.
Option D meets the requirements of the scenario because it allows you to reduce operational overhead and support the product release by using the following AWS services and features:
Amazon Elastic Kubernetes Service (Amazon EKS) is a fully managed service that allows you to run Kubernetes applications on AWS without needing to install, operate, or maintain your own Kubernetes control plane. You can use Amazon EKS to deploy your containerized application services on a Kubernetes cluster that is compatible with your existing tools and processes.
AWS Fargate is a serverless compute engine that eliminates the need to provision and manage servers for your containers. You can use AWS Fargate as the launch type for your Amazon EKS pods, which are the smallest deployable units of computing in Kubernetes. You can also enable auto scaling for your pods, which allows you to automatically adjust the number of pods based on the demand or custom metrics.
An Application Load Balancer (ALB) is a load balancer that distributes traffic across multiple targets in multiple Availability Zones using HTTP or HTTPS protocols. You can use an ALB to balance the load across your Amazon EKS pods and provide high availability and fault tolerance for your application.
Amazon RDS for PostgreSQL is a fully managed relational database service that supports the PostgreSQL open source database engine. You can create additional read replicas for your DB instance, which are copies of your primary DB instance that can handle read-only queries and improve performance. You can also use read replicas to scale out beyond the capacity of a single DB instance for read-heavy workloads.
Amazon Managed Streaming for Apache Kafka (Amazon MSK) is a fully managed service that makes it easy to build and run applications that use Apache Kafka to process streaming data. Apache Kafka is an open source platform for building real-time data pipelines and streaming applications. You can use Amazon MSK to create and manage a Kafka cluster that is highly available, secure, and compatible with your existing Kafka applications. You can also configure your application services to use the Amazon MSK cluster as a source or destination of streaming data.
Amazon S3 is an object storage service that offers high durability, availability, and scalability. You can store static content such as images, videos, or documents in Amazon S3 buckets, which are containers for objects. You can also serve static content directly from Amazon S3 using public URLs or presigned URLs.
Amazon CloudFront is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds. You can use Amazon CloudFront to create a distribution that caches static content from your Amazon S3 bucket at edge locations closer to your users. This can improve the performance and user experience of your application.
Option A is incorrect because creating an EC2 Auto Scaling group behind an ALB would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances for your containers. Creating additional read replicas for the DB instance would not provide high availability or fault tolerance in case of a failure of the primary DB instance, unlike deploying the DB instance in Multi-AZ mode. Creating Amazon Kinesis data streams would not be compatible with your existing Apache Kafka applications, unlike using Amazon MSK.
Option B is incorrect because creating an EC2 Auto Scaling group behind an ALB would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances for your containers. Creating Amazon Kinesis data streams would not be compatible with your existing Apache Kafka applications, unlike using Amazon MSK. Storing and serving static content directly from Amazon S3 would not provide optimal performance and user experience, unlike using Amazon CloudFront.
Option C is incorrect because deploying the application on a Kubernetes cluster created on the EC2 instances behind an ALB

would not reduce operational overhead as much as using AWS Fargate with Amazon EKS, as you would still need to manage EC2 instances and Kubernetes control plane for your containers. Using Amazon API Gateway to interact with the application would add an unnecessary layer of complexity and cost to your architecture, as you would need to create and maintain an API gateway that proxies requests to your ALB.

NEW QUESTION # 311

A company runs an application on AWS. The company curates data from several different sources. The company uses proprietary algorithms to perform data transformations and aggregations. After the company performs E TL processes, the company stores the results in Amazon Redshift tables. The company sells this data to other companies. The company downloads the data as files from the Amazon Redshift tables and transmits the files to several data customers by using FTP. The number of data customers has grown significantly. Management of the data customers has become difficult.

The company will use AWS Data Exchange to create a data product that the company can use to share data with customers. The company wants to confirm the identities of the customers before the company shares data.

The customers also need access to the most recent data when the company publishes the data.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Publish the Amazon Redshift data to an Open Data on AWS Data Exchange. Require the customers to subscribe to the data product in AWS Data Exchange. In the AWS account of the company that produces the data, attach 1AM resource-based policies to the Amazon Redshift tables to allow access only to verified AWS accounts.
- B. Download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically. Use AWS Data Exchange for S3 to share data with customers. Configure subscription verification. Require the data customers to subscribe to the data product.
- C. Use AWS Data Exchange for APIs to share data with customers. Configure subscription verification. In the AWS account of the company that produces the data, create an Amazon API Gateway Data API service integration with Amazon Redshift. Require the data customers to subscribe to the data product.
- D. In the AWS account of the company that produces the data, create an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster. Configure subscription verification. Require the data customers to subscribe to the data product.

Answer: B

Explanation:

The company should download the data from the Amazon Redshift tables to an Amazon S3 bucket periodically and use AWS Data Exchange for S3 to share data with customers. The company should configure subscription verification and require the data customers to subscribe to the data product. This solution will meet the requirements with the least operational overhead because AWS Data Exchange for S3 is a feature that enables data subscribers to access third-party data files directly from data providers' Amazon S3 buckets.

Subscribers can easily use these files for their data analysis with AWS services without needing to create or manage data copies. Data providers can easily set up AWS Data Exchange for S3 on top of their existing S3 buckets to share direct access to an entire S3 bucket or specific prefixes and S3 objects. AWS Data Exchange automatically manages subscriptions, entitlements, billing, and payment1.

The other options are not correct because:

Using AWS Data Exchange for APIs to share data with customers would not work because AWS Data Exchange for APIs is a feature that enables data subscribers to access third-party APIs directly from data providers' AWS accounts. Subscribers can easily use these APIs for their data analysis with AWS services without needing to manage API keys or tokens. Data providers can easily set up AWS Data Exchange for APIs on top of their existing API Gateway resources to share direct access to an entire API or specific routes and stages2. However, this feature is not suitable for sharing data from Amazon Redshift tables, which are not exposed as APIs.

Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not work because the Data API is a feature that enables you to query your Amazon Redshift cluster using HTTP requests, without needing a persistent connection or a SQL client3. It is useful for building applications that interact with Amazon Redshift, but not for sharing data files with customers.

Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not work because AWS Data Exchange does not support datashares for Amazon Redshift clusters. A datashare is a feature that enables you to share live and secure access to your Amazon Redshift data across your accounts or with third parties without copying or moving the underlying data4.

It is useful for sharing query results and views with other users, but not for sharing data files with customers.

Publishing the Amazon Redshift data to an Open Data on AWS Data Exchange would not work because Open Data on AWS Data Exchange is a feature that enables you to find and use free and public datasets from AWS customers and partners. It is useful for accessing open and free data, but not for confirming the identities of the customers or charging them for the data.

References:

https://aws.amazon.com/data-exchange/why-aws-data-exchange/s3/
https://aws.amazon.com/data-exchange/why-aws-data-exchange/api/
https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html
https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html
https://aws.amazon.com/data-exchange/open-data/

NEW QUESTION # 312
A solutions architect is auditing the security setup of an AWS Lambda function for a company.
The Lambda function retrieves the latest changes from an Amazon Aurora database. The Lambda function and the database run in the same VPC. Lambda environment variables are providing the database credentials to the Lambda function. The Lambda function aggregates data and makes the data available in an Amazon S3 bucket that is configured for server-side encryption with AWS KMS managed encryption keys (SSE-KMS). The data must not travel across the internet. If any database credentials become compromised, the company needs a solution that minimizes the impact of the compromise.
What should the solutions architect recommend to meet these requirements?

- A. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Enforce HTTPS on the connection to Amazon S3 during data transfers.
- B. Save the database credentials in AWS Secrets Manager. Set up password rotation on the credentials in Secrets Manager. Change the IAM role for the Lambda function to allow the function to access Secrets Manager. Modify the Lambda function to retrieve the credentials Om Secrets Manager. Enforce HTTPS on the connection to Amazon S3 during data transfers.
- C. Enable IAM database authentication on the Aurora DB cluster. Change the IAM role for the Lambda function to allow the function to access the database by using IAM database authentication. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.
- D. Save the database credentials in AWS Systems Manager Parameter Store. Set up password rotation on the credentials in Parameter Store. Change the IAM role for the Lambda function to allow the function to access Parameter Store. Modify the Lambda function to retrieve the credentials from Parameter Store. Deploy a gateway VPC endpoint for Amazon S3 in the VPC.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/UsingWithRDS.IAMDBAuth.
html

NEW QUESTION # 313
......

- VCE SAP-C02 Dumps Reliable SAP-C02 Guide Files SAP-C02 Reliable Braindumps Pdf Search for ⇒ SAP-C02 ⇐ and download it for free on www.pdfvce.com website Exam Cram SAP-C02 Pdf
- Amazon SAP-C02 Exam Dumps - Secret Hacks To Crack SAP-C02 Exam ❤ Search for { SAP-C02 } and download exam materials for free through ➤ www.vce4dumps.com Latest SAP-C02 Test Guide
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

BONUS!!! Download part of TestKingIT SAP-C02 dumps for free: https://drive.google.com/open?id=10PqNyDsHedn7bpLRMbI-KtfOqnrxlLdK