

Neuester und gültiger SAP-C02 Test VCE Motoren-Dumps und SAP-C02 neueste Testfragen für die IT-Prüfungen



Laden Sie die neuesten DeutschPrüfung SAP-C02 PDF-Versionen von Prüfungsfragen kostenlos von Google Drive herunter: https://drive.google.com/open?id=1AbMdR1m-znIzGH_zeJGfdmBb-MYKqFi

Wenn Sie DeutschPrüfung wählen, kommt der Erfolg auf Sie zu. Die Examsfragen zur Amazon SAP-C02 Zertifizierungsprüfung wird Ihnen helfen, die Prüfung zu bestehen. Die Simulationsprüfung vor der Amazon SAP-C02 Zertifizierungsprüfung zu machen, ist ganz notwendig und effizient. Wenn Sie DeutschPrüfung wählen, können Sie 100% die Prüfung bestehen.

Als Anbieter des Amazon SAP-C02 (AWS Certified Solutions Architect - Professional (SAP-C02)) IT-Prüfungskompendium bieten IT-Experten von DeutschPrüfung ständig die Produkte von guter Qualität. Sie bieten den Kunden kostenlosen Online-Service rund um die Uhr und aktualisieren Amazon SAP-C02 (AWS Certified Solutions Architect - Professional (SAP-C02)) Prüfungsfragen und Antworten auch am schnellsten.

>> SAP-C02 Zertifizierungsantworten <<

SAP-C02 Übungsmaterialien & SAP-C02 Lernführung: AWS Certified Solutions Architect - Professional (SAP-C02) & SAP-C02 Lernguide

Die Amazon SAP-C02 Zertifizierungsprüfung ist eine der beliebten und wichtigen Prüfung in der IT-Branche. Wir haben die besten Lernhilfe und den besten Online-Service. Wir bieten den IT-Fachleuten eine Abkürzung. Die online Tests zur Amazon SAP-C02 Zertifizierungsprüfung von DeutschPrüfung enthalten viele Prüfungsinhalte und Antworten, die Sie wollen. Wenn Sie die Simulationsprüfung von DeutschPrüfung bestehen, dann finden Sie, dass DeutschPrüfung bietet genau was, was Sie wollen und dass Sie sich gut auf die Amazon SAP-C02 Prüfung vorbereiten können.

Die SAP-C02-Prüfung ist ein umfassender Test, der eine Vielzahl von Themen rund um die AWS-Architektur abdeckt, einschließlich der Gestaltung und Implementierung von hochverfügbaren, fehlertoleranten und skalierbaren Systemen, der Auswahl der geeigneten AWS-Services zur Erfüllung spezifischer Anforderungen von Anwendungen sowie der Gestaltung und Implementierung von unternehmensweiten skalierbaren Operationen auf AWS. Die Prüfung umfasst auch fortgeschrittene Sicherheitsthemen wie Datenverschlüsselung, Identitäts- und Zugriffsmanagement sowie Netzwerksicherheit.

Amazon AWS Certified Solutions Architect - Professional (SAP-C02) SAP-C02 Prüfungsfragen mit Lösungen (Q497-Q502):

497. Frage

A company has Linux-based Amazon EC2 instances. Users must access the instances by using SSH with EC2 SSH Key pairs. Each machine requires a unique EC2 Key pair.

The company wants to implement a key rotation policy that will, upon request, automatically rotate all the EC2 key pairs and keep the key in a securely encrypted place. The company will accept less than 1 minute of downtime during key rotation.

Which solution will meet these requirement?

- A. Store all the keys in Parameter Store, a capability of AWS Systems Manager, as a string. Define a Systems Manager maintenance window to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instance. Update the private keys in parameter.
- B. Add all the EC2 instances to Fleet Manager, a capability of AWS Systems Manager. Define a Systems Manager maintenance window to issue a Systems Manager Run Command document to generate new Key pairs and to rotate public keys to all the instances in Fleet Manager.
- **C. Store all the keys in AWS Secrets Manager. Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public Keys on EC2 instances. Update the private keys in Secrets Manager.**
- D. Import the EC2 key pairs into AWS Key Management Service (AWS KMS). Configure automatic key rotation for these key pairs. Create an Amazon EventBridge scheduled rule to invoke an AWS Lambda function to initiate the key rotation AWS KMS.

Antwort: C

Begründung:

To meet the requirements for automatic key rotation of EC2 SSH key pairs with minimal downtime, storing the keys in AWS Secrets Manager and defining a rotation schedule is the most suitable solution. AWS Secrets Manager supports automatic rotation of secrets, including SSH keys, by invoking a Lambda function that can handle the creation of new key pairs and the replacement of public keys on EC2 instances. Updating the corresponding private keys in Secrets Manager ensures secure and centralized management of SSH keys, complying with the key rotation policy and minimizing operational overhead.

:

AWS Secrets Manager Documentation: Describes how to store and rotate secrets, including SSH keys, using Secrets Manager and Lambda functions.

AWS Lambda Documentation: Provides information on creating Lambda functions for custom secret rotation logic.

AWS Best Practices for Security: Highlights the importance of key rotation and how AWS services like Secrets Manager can facilitate secure and automated key management.

498. Frage

A company recently wanted a web application from an on-premises data center to the AWS Cloud. The web application infrastructure consists of an Amazon CloudFront distribution that routes to an Application Load Balancer (ALB), with Amazon Elastic Container Service (Amazon ECS) to process requests. A recent security audit revealed that the web application is accessible by using both CloudFront and ALB endpoints. However, the company requires that the web application must be accessible only by using the CloudFront endpoint.

Which solution will meet this requirement with the LEAST amount of effort?

- **A. Update ALB security group ingress to allow access only from the CloudFront managed prefix list.**
- B. Create a new security group and attach it to the CloudFront distribution. Update the ALB security group ingress to allow access only from the CloudFront security group.

- C. Create a VPC interface endpoint for Elastic Load Balancing. Update the ALB scheme from internet-facing to internal
- D. Extract CloudFront IPS from the AWS provided ip-ranges.json document. Update ALB security group ingress to allow access only from CloudFront IPs.

Antwort: A

Begründung:

The CloudFront managed prefix list contains the IP ranges for all CloudFront edge locations. By updating the ALB security group ingress to allow access only from this prefix list, the web application will be accessible only by using the CloudFront endpoint. This solution requires the least amount of effort compared to the other options, which involve creating new resources or updating existing ones. This solution also avoids hard-coding IP addresses, which can change over time.

499. Frage

A company wants to containerize a multi-tier web application and move the application from an on-premises data center to AWS. The application includes web, application, and database tiers. The company needs to make the application fault tolerant and scalable. Some frequently accessed data must always be available across application servers. Frontend web servers need session persistence and must scale to meet increases in traffic.

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

- A. Run the application on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Use Amazon Elastic File System (Amazon EFS) for data that is frequently accessed between the web and application tiers. Store the frontend web server session data in Amazon Simple Queue Service (Amazon SQS).
- B. Run the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node groups. Use ReplicaSets to run the web servers and applications. Create an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system across all EKS pods to store frontend web server session data.
- C. Deploy the application on Amazon Elastic Kubernetes Service (Amazon EKS). Configure Amazon EKS to use managed node groups. Run the web servers and application as Kubernetes deployments in the EKS cluster. Store the frontend web server session data in an Amazon DynamoDB table. Create an Amazon Elastic File System (Amazon EFS) volume that all applications will mount at the time of deployment.
- D. Run the application on Amazon Elastic Container Service (Amazon ECS) on Amazon EC2. Use Amazon ElastiCache for Redis to cache frontend web server session data. Use Amazon Elastic Block Store (Amazon EBS) with Multi-Attach on EC2 instances that are distributed across multiple Availability Zones.

Antwort: C

Begründung:

Deploying the application on Amazon EKS with managed node groups simplifies the operational overhead of managing the Kubernetes cluster. Running the web servers and application as Kubernetes deployments ensures that the desired number of pods are always running and can scale up or down as needed. Storing the frontend web server session data in an Amazon DynamoDB table provides a fast, scalable, and durable storage option that can be accessed across multiple Availability Zones. Creating an Amazon EFS volume that all applications will mount at the time of deployment allows the application to share data that is frequently accessed between the web and application tiers. Reference:

<https://docs.aws.amazon.com/eks/latest/userguide/managed-node-groups.html>

<https://docs.aws.amazon.com/eks/latest/userguide/deployments.html>

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

<https://docs.aws.amazon.com/efs/latest/ug/mounting-fs.html>

500. Frage

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 ETL 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 IAM 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. 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.
- D. 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.

Antwort: B

Begründung:

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/>

501. Frage

A company has an application that runs as a ReplicaSet of multiple pods in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster has nodes in multiple Availability Zones. The application generates many small files that must be accessible across all running instances of the application. The company needs to back up the files and retain the backups for 1 year.

Which solution will meet these requirements while providing the FASTEST storage performance?

- A. Create an Amazon S3 bucket. Configure the ReplicaSet to mount the S3 bucket. Direct the application to store files in the S3 bucket. Configure S3 Versioning to retain copies of the data. Configure an S3 Lifecycle policy to delete objects after 1 year.
- **B. Create an Amazon Elastic File System (Amazon EFS) file system and a mount target for each subnet that contains nodes in the EKS cluster. Configure the ReplicaSet to mount the file system. Direct the application to store files in the file system.**

https://drive.google.com/open?id=1AbMdR1m-znlzGH_zleJGfdmBb-MYKqFi