

Sure SAP-C02 Pass | SAP-C02 Latest Demo



DOWNLOAD the newest Pass4Leader SAP-C02 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1-PB9emVx31knudoERGf0zUHqxSde1HcB>

On the other hand, those who do not score well can again try reading all the AWS Certified Solutions Architect - Professional (SAP-C02) (SAP-C02) dumps questions and then give the SAP-C02 exam. This will help them polish their skills and clear all their doubts. Also, you must note down your AWS Certified Solutions Architect - Professional (SAP-C02) (SAP-C02) practice test score every time you try the Amazon Exam Questions. It will help you keep a record of your study and how well you are doing in them.

The SAP-C02 Certification is recognized by employers worldwide as a validation of an individual's expertise in AWS. It is a valuable credential for IT professionals who want to advance their careers and increase their earning potential. AWS Certified Solutions Architect - Professional (SAP-C02) certification is valid for three years, after which individuals must recertify to maintain their status as an AWS Certified Solutions Architect - Professional.

>> Sure SAP-C02 Pass <<

Buy Pass4Leader SAP-C02 Exam Dumps Today and Get Free Updates for 1 year

Generally speaking, every candidate wants to pass the exam just one time. SAP-C02 learning materials of us can do that for you. Since we have a professional team to collect and research the latest information for the exam, and therefore the quality can be guaranteed. We offer you free demo for SAP-C02 Exam Materials to have a try, so that you can know what the complete version is like. Besides, we also pass guarantee and money back guarantee, and if you fail to pass the exam after using SAP-C02 exam

materials of us, we will give you refund.

Amazon AWS Certified Solutions Architect - Professional (SAP-C02) Sample Questions (Q539-Q544):

NEW QUESTION # 539

A solutions architect is importing a VM from an on-premises environment by using the Amazon EC2 VM Import feature of AWS Import/Export. The solutions architect has created an AMI and has provisioned an Amazon EC2 instance that is based on that AMI. The EC2 instance runs inside a public subnet in a VPC and has a public IP address assigned. The EC2 instance does not appear as a managed instance in the AWS Systems Manager console. Which combination of steps should the solutions architect take to troubleshoot this issue"? (Select TWO)

- A. Verify that the instance is assigned an appropriate IAM role for Systems Manager
- B. Verify that Systems Manager Agent is installed on the instance and is running
- C. Verify that the AWS Application Discovery Agent is configured
- D. Verify the existence of a VPC endpoint on the VPC
- E. Verify the correct configuration of service-linked roles for Systems Manager

Answer: A,B,C

NEW QUESTION # 540

A company has an organization in AWS Organizations. The organization consists of a large number of AWS accounts that belong to separate business units. The company requires all Amazon EC2 instances to be provisioned with custom, hardened AMIs. The company wants a solution that provides each AWS account access to the AMIs. Which solution will meet these requirements with the MOST operational efficiency?

- A. Create and share the AMIs with EC2 Image Builder. Use AWS Service Catalog to configure a product that provides access to the AMIs across all AWS accounts.
- B. Deploy Jenkins on an EC2 instance. Create jobs to create and share the AMIs across all AWS accounts.
- C. Create the AMIs with EC2 Image Builder. Create an AWS CodePipeline pipeline to share the AMIs across all AWS accounts.
- D. Create the AMIs with EC2 Image Builder. Create an AWS Lambda function to share the AMIs across all AWS accounts.

Answer: A

Explanation:

Explanation

<https://docs.aws.amazon.com/servicecatalog/latest/adminguide/introduction.html>

<https://docs.aws.amazon.com/iot/latest/developerguide/kinesis-rule-action.html>

NEW QUESTION # 541

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. 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.
- 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. 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.

Answer: D

NEW QUESTION # 542

A company needs to build a disaster recovery (DR) solution for its ecommerce website. The web application is hosted on a fleet of t3.large Amazon EC2 instances and uses an Amazon RDS for MySQL DB instance.

The EC2 instances are in an Auto Scaling group that extends across multiple Availability Zones.

In the event of a disaster, the web application must fail over to the secondary environment with an RPO of 30 seconds and an R TO of 10 minutes.

Which solution will meet these requirements MOST cost-effectively?

- A. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Manually restore the backed-up data on new instances. Use an Amazon Route 53 simple routing policy to automatically fail over to the DR Region in the event of a disaster.
- B. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create an Amazon Aurora global database. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the Auto Scaling group of EC2 instances at full capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster.
- C. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. Run the EC2 instances at the minimum capacity in the DR Region. Use an Amazon Route 53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. Increase the desired capacity of the Auto Scaling group.
- D. Use infrastructure as code (IaC) to provision the new infrastructure in the DR Region. Create a cross-Region read replica for the DB instance. Set up a backup plan in AWS Backup to create cross-Region backups for the EC2 instances and the DB instance. Create a cron expression to back up the EC2 instances and the DB instance every 30 seconds to the DR Region. Recover the EC2 instances from the latest EC2 backup. Use an Amazon Route 53 geolocation routing policy to automatically fail over to the DR Region in the event of a disaster.

Answer: C

Explanation:

The company should use infrastructure as code (IaC) to provision the new infrastructure in the DR Region.

The company should create a cross-Region read replica for the DB instance. The company should set up AWS Elastic Disaster Recovery to continuously replicate the EC2 instances to the DR Region. The company should run the EC2 instances at the minimum capacity in the DR Region. The company should use an Amazon Route

53 failover routing policy to automatically fail over to the DR Region in the event of a disaster. The company should increase the desired capacity of the Auto Scaling group. This solution will meet the requirements most cost-effectively because AWS Elastic Disaster Recovery (AWS DRS) is a service that minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. AWS DRS enables RPOs of seconds and RTOs of minutes¹. AWS DRS continuously replicates data from the source servers to a staging area subnet in the DR Region, where it uses low-cost storage and minimal compute resources to maintain ongoing replication. In the event of a disaster, AWS DRS automatically converts the servers to boot and run natively on AWS and launches recovery instances on AWS within minutes². By using AWS DRS, the company can save costs by removing idle recovery site resources and paying for the full disaster recovery site only when needed. By creating a cross-Region read replica for the DB instance, the company can have a standby copy of its primary database in a different AWS Region³. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way⁴. By using an Amazon Route 53 failover routing policy, the company can route traffic to a resource that is healthy or to another resource when the first resource becomes unavailable.

The other options are not correct because:

* Using AWS Backup to create cross-Region backups for the EC2 instances and the DB instance would not meet the RPO and RTO requirements. AWS Backup is a service that enables you to centralize and automate data protection across AWS services. You can use AWS Backup to back up your application data across AWS services in your account and across accounts. However, AWS Backup does not provide continuous replication or fast recovery; it creates backups at scheduled intervals and requires manual restoration. Creating backups every 30 seconds would also incur high costs and network bandwidth.

* Creating an Amazon API Gateway Data API service integration with Amazon Redshift would not help with disaster recovery. 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 client. It is useful for building applications that interact with Amazon Redshift, but not for replicating or recovering data.

* Creating an AWS Data Exchange datashare by connecting AWS Data Exchange to the Redshift cluster would not help with disaster recovery. AWS Data Exchange is a service that makes it easy for AWS customers to exchange data in the cloud. You can use AWS Data Exchange to subscribe to a diverse selection of third-party data products or offer your own data products to other AWS customers. 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 data. It is useful for sharing query results and views with other users, but not for replicating or recovering data.

References:

- * <https://aws.amazon.com/disaster-recovery/>
- * <https://docs.aws.amazon.com/drs/latest/userguide/what-is-drs.html>
- * https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html#USER_ReadRepl

XRgn

- * <https://aws.amazon.com/cloudformation/>
- * <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover.html>
- * <https://aws.amazon.com/backup/>
- * <https://docs.aws.amazon.com/redshift/latest/mgmt/data-api.html>
- * <https://aws.amazon.com/data-exchange/>
- * <https://docs.aws.amazon.com/redshift/latest/dg/datashare-overview.html>

NEW QUESTION # 543

A company has a Windows-based desktop application that is packaged and deployed to the users' Windows machines. The company recently acquired another company that has employees who primarily use machines with a Linux operating system. The acquiring company has decided to migrate and rehost the Windows-based desktop application to AWS.

All employees must be authenticated before they use the application. The acquiring company uses Active Directory on premises but wants a simplified way to manage access to the application on AWS for all the employees.

Which solution will rehost the application on AWS with the LEAST development effort?

- A. Set up and provision an Amazon Workspaces virtual desktop for every employee. Implement authentication by using Amazon Cognito identity pools. Instruct employees to run the application from their provisioned Workspaces virtual desktops.
- B. Create an Auto Scaling group of Windows-based Amazon EC2 instances. Join each EC2 instance to the company's Active Directory domain. Implement authentication by using the Active Directory that is running on premises. Instruct employees to run the application by using a Windows remote desktop.
- C. Use an Amazon AppStream 2.0 image builder to create an image that includes the application and the required configurations. Provision an AppStream 2.0 On-Demand fleet with dynamic Fleet Auto Scaling policies for running the image. Implement authentication by using AppStream 2.0 user pools. Instruct the employees to access the application by starting browser-based AppStream 2.0 streaming sessions.
- D. Refactor and containerize the application to run as a web-based application. Run the application in Amazon Elastic Container Service (Amazon ECS) on AWS Fargate with step scaling policies. Implement authentication by using Amazon Cognito user pools. Instruct the employees to run the application from their browsers.

Answer: C

Explanation:

Option C leverages Amazon AppStream 2.0, a fully managed application streaming service. With AppStream 2.0, you can create an image that includes the Windows-based desktop application and the required configurations.

NEW QUESTION # 544

.....

The Pass4Leader SAP-C02 exam practice test questions provide a way to assess your understanding of the material, identify areas for improvement, and build confidence and test-taking skills. The Pass4Leader SAP-C02 exam practice test questions are real and verified by AWS Certified Solutions Architect - Professional (SAP-C02) (SAP-C02) exam trainers. They work collectively and strive hard to ensure the top standard of AWS Certified Solutions Architect - Professional (SAP-C02) (SAP-C02) exam practice questions all the time.

SAP-C02 Latest Demo: <https://www.pass4leader.com/Amazon/SAP-C02-exam.html>

BONUS!!! Download part of Pass4Leader SAP-C02 dumps for free: <https://drive.google.com/open?id=1-PB9emVx31knudoERGf0zUHqxSde1HcB>