

# SAP-C02 Real Testing Environment & SAP-C02 New Practice Materials



2026 Latest SureTorrent SAP-C02 PDF Dumps and SAP-C02 Exam Engine Free Share: <https://drive.google.com/open?id=16IWMVOSYtTwqztSuUagZRX9Ii5C1E0yK>

If you are willing to purchase valid Amazon SAP-C02 reliable vce exam simulator, you should be eagle-eyed since there are so much information on the internet. Valid products are hard to tell, once you find them, you will fell as if you'd found a priceless treasure. Our SAP-C02 reliable vce exam simulator will be your priceless products. Our passing rate is 100% recent two years. We can assure you that No Pass Full Refund. Our materials are valid and the best absolutely.

The SAP-C02 Exam covers a wide range of topics, including designing and deploying scalable, highly available, and fault-tolerant systems, selecting appropriate AWS services for a given scenario, migrating complex multi-tier applications on AWS, designing and deploying secure and compliant applications and infrastructures, and troubleshooting common issues with AWS deployments.

[\*\*>> SAP-C02 Real Testing Environment <<\*\*](#)

## Get Valid Amazon SAP-C02 Exam Questions and Answer

You can easily download these formats of Amazon SAP-C02 actual dumps and use them to prepare for the Amazon SAP-C02 certification test. You don't need to enroll yourself in expensive SAP-C02 Exam Training classes. With the Amazon SAP-C02 valid dumps, you can easily prepare well for the actual Amazon SAP-C02 exam at home.

## Amazon AWS Certified Solutions Architect - Professional (SAP-C02) Sample Questions (Q608-Q613):

### NEW QUESTION # 608

A company runs an application on Amazon EC2 and AWS Lambda.

a. The application stores temporary data in Amazon S3. The S3 objects are deleted after 24 hours.

The company deploys new versions of the application by launching AWS CloudFormation stacks. The stacks create the required resources. After validating a new version, the company deletes the old stack. The deletion of an old development stack recently failed.

A solutions architect needs to resolve this issue without major architecture changes.

Which solution will meet these requirements?

- A. Update the CloudFormation stack to add a `DeletionPolicy` attribute with a value of `Snapshot` for the S3 bucket resource.
- B. Modify the CloudFormation stack to attach a `DeletionPolicy` attribute with a value of `Delete` to the S3 bucket.
- C. Update the CloudFormation template to create an Amazon EFS file system to store temporary files instead of Amazon S3. Configure the Lambda functions to run in the same VPC as the EFS file system.
- D. Create a Lambda function to delete objects from the S3 bucket. Add the Lambda function as a custom resource in the CloudFormation stack with a `DependsOn` attribute that points to the S3 bucket resource.

### Answer: D

Explanation:

CloudFormation cannot delete non-empty S3 buckets. Option A allows you to create a custom Lambda resource that deletes all objects in the S3 bucket before the stack deletes it. The `DependsOn` ensures the bucket deletion occurs only after the Lambda has completed.

B: Adding `DeletionPolicy: Delete` does not resolve the issue if the bucket still contains objects.

C: `Snapshot` doesn't apply to S3 and won't help here.

D: Changing to Amazon EFS would require architectural changes, which are not allowed per requirements.

### NEW QUESTION # 609

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. 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.
- 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. 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.
- 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: A

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 minutes<sup>1</sup>. 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 minutes<sup>2</sup>. 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 Region<sup>3</sup>. By using infrastructure as code (IaC), the company can provision the new infrastructure in the DR Region in an automated and consistent way<sup>4</sup>. 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](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 # 610

A company is hosting a monolithic REST-based API for a mobile app on five Amazon EC2 instances in public subnets of a VPC. Mobile clients connect to the API by using a domain name that is hosted on Amazon Route 53. The company has created a Route 53 multivalue answer routing policy with the IP addresses of all the EC2 instances. Recently, the app has been overwhelmed by large and sudden increases to traffic. The app has not been able to keep up with the traffic.

A solutions architect needs to implement a solution so that the app can handle the new and varying load.

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

- A. Separate the API into individual AWS Lambda functions. Configure an Amazon API Gateway REST API with Lambda integration for the backend. Update the Route 53 record to point to the API Gateway API.
- B. Create an Auto Scaling group. Place all the EC2 instances in the Auto Scaling group. Configure the Auto Scaling group to perform scaling actions that are based on CPU utilization. Create an AWS Lambda function that reacts to Auto Scaling group changes and updates the Route 53 record.
- C. Create an Application Load Balancer (ALB) in front of the API. Move the EC2 instances to private subnets in the VPC. Add the EC2 instances as targets for the ALB. Update the Route 53 record to point to the ALB.
- D. Containerize the API logic. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. Run the containers in the cluster by using Amazon EC2. Create a Kubernetes ingress. Update the Route 53 record to point to the Kubernetes

ingress.

#### Answer: A

Explanation:

By breaking down the monolithic API into individual Lambda functions and using API Gateway to handle the incoming requests, the solution can automatically scale to handle the new and varying load without the need for manual scaling actions. Additionally, this option will automatically handle the traffic without the need of having EC2 instances running all the time and only pay for the number of requests and the duration of the execution of the Lambda function.

By updating the Route 53 record to point to the API Gateway, the solution can handle the traffic and also it will direct the traffic to the correct endpoint.

### NEW QUESTION # 611

A global ecommerce company has many data centers worldwide. The company needs scalable cloud storage for legacy file applications. Requirements:

Must support iSCSI access from on-premises servers.

Must support point-in-time snapshots via AWS Backup.

Must retain low-latency access to frequently accessed data. Which solution will meet these requirements?

- A. Provision an AWS Storage Gateway volume gateway in cache mode. Back up the volumes using AWS Backup.
- B. Provision an AWS Storage Gateway file gateway in cache mode. Use AWS Backup.
- C. Use Amazon FSx File Gateway and S3 File Gateway. Use AWS Backup.
- D. Provision an AWS Storage Gateway tape gateway with S3 and AWS Backup.

#### Answer: A

Explanation:

C is correct because the Storage Gateway Volume Gateway in cache mode supports iSCSI block storage, which is exactly what legacy applications need. It also caches frequently accessed data locally for low-latency performance and integrates with AWS Backup for point-in-time copies.

A is incorrect - tape gateway is designed for archival backup only.

B and D are incorrect because FSx and file gateway offer file (NFS/SMB) storage, not block/iSCSI support.

Reference:

Volume Gateway Overview

AWS Backup Integration

### NEW QUESTION # 612

A company wants to deploy an API to AWS. The company plans to run the API on AWS Fargate behind a load balancer. The API requires the use of header-based routing and must be accessible from on-premises networks through an AWS Direct Connect connection and a private VIF.

The company needs to add the client IP addresses that connect to the API to an allow list in AWS. The company also needs to add the IP addresses of the API to the allow list. The company's security team will allow /27 CIDR ranges to be added to the allow list. The solution must minimize complexity and operational overhead.

Which solution will meet these requirements?

- A. Create a new Network Load Balancer (NLB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the API. Attach the new security group to the Fargate tasks. Provide the security team with the NLB's IP addresses for the allow list.
- B. Create a new Application Load Balancer (ALB) in the same subnets as the Fargate task deployments. Create a security group that includes only the client IP addresses that need access to the API. Attach the security group to the ALB. Provide the security team with the ALB's IP addresses for the allow list.
- C. Create two new /27 subnets. Create a new Application Load Balancer (ALB) that extends across the new subnets. Create a security group that includes only the client IP addresses that need access to the API. Attach the security group to the ALB. Provide the security team with the new subnet IP ranges for the allow list.
- D. Create two new /27 subnets. Create a new Network Load Balancer (NLB) that extends across the new subnets. Create a new Application Load Balancer (ALB) within the new subnets. Create a security group that includes only the client IP addresses that need access to the API. Attach the security group to the ALB. Add the ALB's IP addresses as targets behind the NLB. Provide the security team with the NLB's IP addresses for the allow list.

**Answer: A**

## NEW QUESTION # 613

A minor mistake may result you to lose chance even losing out on your SAP-C02 Exam. So we hold responsible tents when compiling the SAP-C02 learning guide. The principles of our SAP-C02 practice materials can be expressed in words like clarity, correction and completeness. Experts expressed their meaning with clarity by knowledgeable and understandable words which cannot be misunderstood.

**SAP-C02 New Practice Materials:** <https://www.suretorrent.com/SAP-C02-exam-guide-torrent.html>

2026 Latest SureTorrent SAP-C02 PDF Dumps and SAP-C02 Exam Engine Free Share: <https://drive.google.com/open?id=16IWMVOSYtTwqz1SuUagZRX9i5C1E0yK>