



SAP-C02^{Q&As}

AWS Certified Solutions Architect - Professional

Pass Amazon SAP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/sap-c02.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

A retail company runs a business-critical web service on an Amazon Elastic Container Service (Amazon ECS) cluster that runs on Amazon EC2 instances. The web service receives POST requests from end users and writes data to a MySQL database that runs on a separate EC2 instance. The company needs to ensure that data loss does not occur. The current code deployment process includes manual updates of the ECS service. During a recent deployment, end users encountered intermittent 502 Bad Gateway errors in response to valid web requests.

The company wants to implement a reliable solution to prevent this issue from recurring. The company also wants to automate code deployments. The solution must be highly available and must optimize cost-effectiveness.

Which combination of steps will meet these requirements? (Select THREE.)

- A. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a blue/green deployment with validation testing to update the ECS service.
- B. Migrate the MySQL database to run on an Amazon RDS for MySQL Multi-AZ DB instance that uses Provisioned IOPS SSD (io2) storage.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an event source to receive the POST requests from the web service. Configure an AWS Lambda function to poll the queue. Write the data to the database.
- D. Run the web service on an ECS cluster that has a Fargate launch type. Use AWS CodePipeline and AWS CodeDeploy to perform a canary deployment to update the ECS service.

Correct Answer: CD

QUESTION 2

A company has a legacy monolithic application that is critical to the company's business. The company hosts the application on an Amazon EC2 instance that runs Amazon Linux 2. The company's application team receives a directive from the legal department to back up the data from the instance's encrypted Amazon Elastic Block Store (Amazon EBS) volume to an Amazon S3 bucket. The application team does not have the administrative SSH key pair for the instance. The application must continue to serve the users.

Which solution will meet these requirements?

- A. Attach a role to the instance with permission to write to Amazon S3. Use the AWS Systems Manager Session Manager option to gain access to the instance and run commands to copy data into Amazon S3.
- B. Create an image of the instance with the reboot option turned on. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.
- C. Take a snapshot of the EBS volume by using Amazon Data Lifecycle Manager (Amazon DLM). Copy the data to Amazon S3.
- D. Create an image of the instance. Launch a new EC2 instance from the image. Attach a role to the new instance with permission to write to Amazon S3. Run a command to copy data into Amazon S3.

Correct Answer: C

Taking a snapshot of the EBS volume using Amazon Data Lifecycle Manager (DLM) will meet the requirements because it allows you to create a backup of the volume without the need to access the instance or its SSH key pair.



Additionally, DLM allows you to schedule the backups to occur at specific intervals and also enables you to copy the snapshots to an S3 bucket. This approach will not impact the running application as the backup is performed on the EBS volume level.

QUESTION 3

A company is deploying a new cluster for big data analytics on AWS. The cluster will run across many Linux Amazon EC2 instances that are spread across multiple Availability Zones.

All of the nodes in the cluster must have read and write access to common underlying file storage. The file storage must be highly available, must be resilient, must be compatible with the Portable Operating System Interface (POSIX), and must accommodate high levels of throughput.

Which storage solution will meet these requirements?

- A. Provision an AWS Storage Gateway file gateway NFS file share that is attached to an Amazon S3 bucket. Mount the NFS file share on each EC2 instance in the cluster.
- B. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses General Purpose performance mode. Mount the EFS file system on each EC2 instance in the cluster.
- C. Provision a new Amazon Elastic Block Store (Amazon EBS) volume that uses the io2 volume type. Attach the EBS volume to all of the EC2 instances in the cluster.
- D. Provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode. Mount the EFS file system on each EC2 instance in the cluster.

Correct Answer: D

The best solution is to provision a new Amazon Elastic File System (Amazon EFS) file system that uses Max I/O performance mode and mount the EFS file system on each EC2 instance in the cluster. Amazon EFS is a fully managed, scalable, and elastic file storage service that supports the POSIX standard and can be accessed by multiple EC2 instances concurrently. Amazon EFS offers two performance modes: General Purpose and Max I/O. Max I/O mode is designed for highly parallelized workloads that can tolerate higher latencies than the General Purpose mode. Max I/O mode provides higher levels of aggregate throughput and operations per second, which are suitable for big data analytics applications. This solution meets all the requirements of the company. References: Amazon EFS Documentation, Amazon EFS performance modes

QUESTION 4

A company has migrated an application from on premises to AWS. The application frontend is a static website that runs on two Amazon EC2 instances behind an Application Load Balancer (ALB). The application backend is a Python application that runs on three EC2 instances behind another ALB. The EC2 instances are large, general purpose On-Demand Instances that were sized to meet the on-premises specifications for peak usage of the application.

The application averages hundreds of thousands of requests each month. However, the application is used mainly during lunchtime and receives minimal traffic during the rest of the day.

A solutions architect needs to optimize the infrastructure cost of the application without negatively affecting the application availability.

Which combination of steps will meet these requirements? (Choose two.)



- A. Change all the EC2 instances to compute optimized instances that have the same number of cores as the existing EC2 instances.
- B. Move the application frontend to a static website that is hosted on Amazon S3.
- C. Deploy the application frontend by using AWS Elastic Beanstalk. Use the same instance type for the nodes.
- D. Change all the backend EC2 instances to Spot Instances.
- E. Deploy the backend Python application to general purpose burstable EC2 instances that have the same number of cores as the existing EC2 instances.

Correct Answer: BD

Moving the application frontend to a static website that is hosted on Amazon S3 will save cost as S3 is cheaper than running EC2 instances. Using Spot instances for the backend EC2 instances will also save cost, as they are significantly

cheaper than On-Demand instances. This will be suitable for the application, as it has minimal traffic during the rest of the day, and the availability of spot instances will not negatively affect the application's availability.

Reference:

Amazon S3 pricing:<https://aws.amazon.com/s3/pricing/>

Amazon EC2 Spot Instances documentation:<https://aws.amazon.com/ec2/spot/>

AWS Elastic Beanstalk documentation:<https://aws.amazon.com/elasticbeanstalk/>

Amazon Elastic Compute Cloud (EC2) pricing:<https://aws.amazon.com/ec2/pricing/>

QUESTION 5

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. 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.
- 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. 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

D. 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.

Correct Answer: B

According to <https://aws.amazon.com/data-exchange/why-aws-data-exchange/redshift-data-tables/>

[SAP-C02 Practice Test](#)

[SAP-C02 Exam Questions](#)

[SAP-C02 Braindumps](#)