Free Amazon SAP-C02 Practice Questions

Shared by Battle on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is running applications on AWS in a multi-account environment. The company's sales team and marketing team use separate AWS accounts in AWS Organizations.

The sales team stores petabytes of data in an Amazon S3 bucket. The marketing team uses Amazon QuickSight for data visualizations. The marketing team needs access to data that the sates team stores in the S3 bucket. The company has encrypted the S3 bucket with an AWS Key Management Service (AWS KMS) key. The marketing team has already created the IAM service role for QuickSight to provide QuickSight access in the marketing AWS account. The company needs a solution that will provide secure access to the data in the S3 bucket across AWS accounts.

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

Options:

A- Create a new S3 bucket in the marketing account. Create an S3 replication rule in the sales account to copy the objects to the new S3 bucket in the marketing account. Update the QuickSight permissions in the marketing account to grant access to the new S3 bucket.
B- Create an SCP to grant access to the S3 bucket to the marketing account. Use AWS Resource Access Manager (AWS RAM) to share the KMS key from the sates account with the marketing account. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
C- Update the S3 bucket policy in the marketing account to grant access to the QuickSight role. Create a KMS grant for the encryption key that is used in the S3 bucket. Grant decrypt access to the QuickSight role. Update the QuickSight permissions in the marketing account to grant access to the S3 bucket.
D- Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight rote, to create a trust relationship with the new IAM role in the sales account.

Answer:

D

Explanation:

Create an IAM role in the sales account and grant access to the S3 bucket. From the marketing account, assume the IAM role in the sales account to access the S3 bucket. Update the QuickSight role, to create a trust relationship with the new IAM role in the sales account.

This approach is the most secure way to grant cross-account access to the data in the S3 bucket while minimizing operational overhead. By creating an IAM role in the sales account, the

marketing team can assume the role in their own account, and have access to the S3 bucket. And updating the QuickSight role, to create a trust relationship with the new IAM role in the sales account will grant the marketing team to access the data in the S3 bucket and use it for data visualization using QuickSight.

AWS Resource Access Manager (AWS RAM) also allows sharing of resources between accounts, but it would require additional management and configuration to set up the sharing, which would increase operational overhead.

Using S3 replication would also replicate the data to the marketing account, but it would not provide the marketing team access to the original data, and also it would increase operational overhead with managing the replication process.

IAM roles and policies, KMS grants and trust relationships are a powerful combination for managing cross-account access in a secure and efficient manner.

AWS IAM Roles

AWS KMS - Key Grants

AWS RAM

Question 2

Question Type: MultipleChoice

A company is running a critical application that uses an Amazon RDS for MySQL database to store data. The RDS DB instance is deployed in Multi-AZ mode.

A recent RDS database failover test caused a 40-second outage to the application A solutions architect needs to design a solution to reduce the outage time to less than 20 seconds.

Which combination of steps should the solutions architect take to meet these requirements? (Select THREE.)

Options:

A- Use Amazon ElastiCache for Memcached in front of the database
B- Use Amazon ElastiCache for Redis in front of the database.
C- Use RDS Proxy in front of the database
D- Migrate the database to Amazon Aurora MySQL
E- Create an Amazon Aurora Replica
F- Create an RDS for MySQL read replica

Answer:

C, D, E

Explanation:

Migrate the database to Amazon Aurora MySQL. - Create an Amazon Aurora Replica. - Use RDS Proxy in front of the database. - These options are correct because they address the requirement of reducing the failover time to less than 20 seconds. Migrating to Amazon Aurora MySQL and creating an Aurora replica can reduce the failover time to less than 20 seconds. Aurora has a built-in, fault-tolerant storage system that can automatically detect and repair failures. Additionally, Aurora has a feature called 'Aurora Global Database' which allows you to create read-only replicas across multiple AWS regions which can further help to reduce the failover time. Creating an Aurora replica can also help to reduce the failover time as it can take over as the primary DB instance in case of a failure. Using RDS proxy can also help to reduce the failover time as it can route the queries to the healthy DB instance, it also helps to balance the load across multiple DB instances.

Question 3

Question Type: MultipleChoice

A global company runs an analytics application on Amazon EC2 for computing. The company uses Amazon EBS as primary storage for raw and processed data. Users manually upload raw data daily to Amazon EC2 by using SSH from a local on-premises storage computer. The analytics application processes the data and a user manually uploads the data to Amazon S3 for long-term storage.

The company wants to containerize the processing logic and migrate the processing logic to Amazon EKS. The company needs an automated solution to upload and move the processed data. The solution must have multiprotocol support and be usable from the EKS cluster.

Which solution meets these requirements with the LEAST operational effort?

Options:

A- Use AWS DataSync to copy raw data to Amazon EFS. Mount Amazon EFS on Amazon EKS as a volume. Use AWS Transfer for SFTP to copy processed data from Amazon EFS to Amazon S3.
B- Use AWS DataSync to copy raw data to Amazon FSx for Lustre. Mount FSx for Lustre on Amazon EKS as a volume. Use DataSync to copy processed data from FSx for Lustre to Amazon S3.
C- Use AWS DataSync to copy raw data to Amazon FSx for NetApp ONTAP. Mount FSx for NetApp ONTAP on Amazon EKS as a volume. Use DataSync to copy processed data from FSx for NetApp ONTAP to Amazon S3.

D- Use AWS DataSync to copy raw data to Amazon FSx for NetApp ONTAP. Mount FSx for NetApp ONTAP on Amazon EKS as a volume. Use AWS Transfer for SFTP to copy processed data from FSx for NetApp ONTAP to Amazon S3.

Answer:

C

Explanation:

This explanation is based on AWS documentation and best practices but is paraphrased, not a literal extract.

The company wants to move from a manual EC2 and EBS-based workflow to a containerized application on Amazon EKS and automate data movement. The solution must:

Support automated transfer of raw and processed data.

Offer multiprotocol support.

Be directly usable from the EKS cluster as a mounted volume.

Minimize operational effort by using managed services where possible.

AWS DataSync is a managed service designed to move data between on-premises storage and AWS storage services or between AWS storage services. It can perform scheduled or continuous transfers with minimal operational overhead. For storage accessible from Amazon EKS, a shared file system that supports mounting as a volume is appropriate.

Amazon FSx for NetApp ONTAP provides a fully managed file system with multiprotocol support, including NFS and SMB, and supports features such as snapshots and storage efficiencies. Because it supports multiple protocols, it satisfies the requirement for multiprotocol access and can be mounted by applications running in Amazon EKS using standard Kubernetes persistent volume mechanisms.

In the correct solution (option C), DataSync is used to copy raw data from the on-premises environment to FSx for NetApp ONTAP. The FSx for NetApp ONTAP file system is then mounted as a volume in the EKS cluster, allowing the containerized analytics processing logic to read and write data directly. After processing, DataSync is again used to copy processed data from FSx for NetApp ONTAP to Amazon S3 for long-term storage. This leverages DataSync's native integration with both FSx for NetApp ONTAP and Amazon S3, and avoids the need to run or manage custom upload tooling.

Option A uses Amazon EFS, which supports NFS but does not provide multiprotocol support (for example, SMB), so it does not fully meet the multiprotocol requirement. It also introduces AWS Transfer for SFTP for the processed data upload, which adds an additional managed endpoint and SFTP-based flow, increasing complexity relative to using DataSync end-to-end.

Option B uses Amazon FSx for Lustre, which is optimized for high-performance compute workloads and integrates well with S3, but it is not a multiprotocol file system and is typically

accessed via NFS. It does not meet the stated multiprotocol requirement.

Option D uses FSx for NetApp ONTAP (which supports multiprotocol) but relies on AWS Transfer for SFTP to move processed data to S3. While this can work, it adds another managed input endpoint and requires SFTP client configuration and management. Using DataSync directly from FSx for NetApp ONTAP to Amazon S3 (as in option C) is more straightforward, better suited for automated large-scale transfers, and involves less operational overhead.

Therefore, option C meets all the requirements with the least operational effort by using DataSync with FSx for NetApp ONTAP and S3.

Question 4

Question Type: MultipleChoice

A company has multiple AWS accounts and manages these accounts with AWS Organizations. A developer was given IAM user credentials to access AWS resources. The developer should have read-only access to all Amazon S3 buckets in the account. However, when the developer tries to access the S3 buckets from the console, they receive an access denied error message with no buckets listed.

A solutions architect reviews the permissions and finds that the developer's IAM user is listed as having read-only access to all S3 buckets in the account.

Which additional steps should the solutions architect take to troubleshoot the issue? (Choose TWO.)

Options:

A- Check the bucket policies for all S3 buckets.
B- Check the ACLs for all S3 buckets.
C- Check the SCPs set at the organizational units (OUs).
D- Check for the permissions boundaries set for the IAM user.
E- Check if an appropriate IAM role is attached to the IAM user.

Answer:

C, D

Explanation:

When an IAM user appears to have the correct permissions but still receives Access Denied errors, especially in an AWS Organizations environment, the effective permissions must be

evaluated across all permission layers. In AWS, permissions are the intersection of all applicable controls, and an explicit deny at any layer overrides any allow.

First, because the company uses AWS Organizations, service control policies (SCPs) must be evaluated. SCPs define the maximum permissions that accounts or organizational units can have. Even if an IAM user or IAM policy allows an action, an SCP attached to the account or OU can explicitly or implicitly deny that action. If an SCP does not allow s3:ListAllMyBuckets or related S3 read actions, the IAM user will not be able to list buckets in the S3 console and will see no buckets at all. Therefore, checking the SCPs applied to the OU or account is a required troubleshooting step.

Second, IAM permissions boundaries can further restrict the effective permissions of an IAM user. A permissions boundary defines the maximum permissions that an IAM user can exercise, regardless of what their attached policies allow. If the permissions boundary does not include the required Amazon S3 read-only permissions (such as s3:ListAllMyBuckets or s3:GetBucketLocation), the user will receive access denied errors even though the user's IAM policy appears to allow read-only access. Reviewing whether a permissions boundary is attached, and whether it allows the necessary S3 actions, is essential.

Option A (checking bucket policies) can be relevant for access to specific buckets or objects, but a user seeing no buckets at all in the S3 console is more commonly caused by missing or denied account-level list permissions rather than individual bucket policies. Bucket policies generally do not control the ability to list all buckets in an account unless they contain explicit denies, which is less common as a first troubleshooting step in an Organizations setup.

Option B (checking ACLs) is less relevant because ACLs primarily control object-level and bucket-level access and are not typically used to manage console-level visibility of all buckets in an account. ACLs also do not commonly block the ListAllMyBuckets action.

Option E is incorrect because IAM users do not require IAM roles to access AWS services. Roles are assumed by users or services, but the presence or absence of a role attached to an IAM user does not affect the user's direct permissions.

Therefore, the most appropriate troubleshooting steps are to check the SCPs applied through AWS Organizations and to verify whether an IAM permissions boundary is restricting the user's effective permissions.

Question 5

Question Type: MultipleChoice

A company runs an ecommerce website on Amazon ECS behind an Application Load Balancer (ALB). The company stores the container images in Amazon ECR. The website stores data in an Amazon Aurora MySQL DB cluster. The company uses an Amazon S3 bucket to store backup data.

The company needs to prevent data tampering. The website domain is registered with Amazon

Route 53. The company wants to recreate the setup in a second AWS Region with an RPO of 5 minutes and an RTO of 15 minutes. The company has created an ALB in the second Region.

Which solution will meet these requirements?

Options:

A- Create a new ECS deployment that uses the Fargate launch type. Use the ECR repository in the current Region to store and pull container images. Set up a cross-Region read replica in Amazon RDS. Create a backup vault in compliance mode and a backup plan in AWS Backup. Set up a Route 53 primary record in the main Region and a secondary record with a multivalue answer routing policy.
B- Create a new ECS deployment that uses the Fargate launch type. Use the ECR repository in the current Region to store and pull container images. Set up a cross-Region read replica in Amazon RDS. Set up a Route 53 primary record in the main Region and a secondary record with a failover routing policy.
C- Set up ECR cross-Region replication. Create a new ECS deployment that uses the Fargate launch type. Migrate the DB cluster to an Aurora global database. Create a backup vault in compliance mode and a backup plan in AWS Backup. Enable point-in-time recovery and cross-Region replication for Amazon S3. Set up a Route 53 primary record in the main Region and a secondaryrecord with a failover routing policy.
D- Set up ECR cross-Region replication. Create a new ECS deployment that uses the Fargate launch type. Migrate the DB cluster to an Aurora global database. Create a backup vault in governance mode and a backup plan in AWS Backup. Set up a Route 53 primary record in the main Region and a secondary record with a geolocation routing policy.

Answer:

C

Question 6

Question Type: MultipleChoice

A company's inventory application stores data in an Amazon Aurora PostgreSQL DB cluster in a single AWS Region. The company wants to improve resiliency by extending the database infrastructure to a secondary Region. The company wants an RTO of 15 minutes and an RPO of 5 minutes. The solution must not run Aurora DB instances in the secondary Region when the application is operational in the main Region. Which solution meets these requirements?

Options:

A- Configure AWS DMS to copy the Aurora DB cluster in the primary Region to the secondary Region. Use AWS DMS to synchronize the primary DB cluster with the secondary DB cluster.

B- Create a new Aurora PostgreSQL DB cluster in the secondary Region. Use AWS Backup to synchronize the primary DB cluster with the secondary DB cluster.
C- Create a headless Aurora DB cluster in the second Region that is part of the same global DB cluster as the primary Region's DB cluster.
D- Create an AWS Backup job to back up the DB cluster and copy the DB cluster to the secondary Region every 5 minutes.

Answer:

C

Explanation:

Creating a headless Aurora DB cluster in the secondary Region as part of an Aurora global database ensures continuous, low-latency replication without running active instances in the secondary Region.

This configuration meets the RTO and RPO requirements by allowing rapid promotion of the secondary cluster in a failover scenario.

It also reduces operational costs compared to continuously running standby instances in the secondary Region.

This aligns with AWS best practices for DR and cross-Region resiliency.

Question 7

Question Type: MultipleChoice

A company needs to implement disaster recovery for a critical application that runs in a single AWS Region. The application's users interact with a web frontend that is hosted on Amazon EC2 Instances behind an Application Load Balancer (ALB). The application writes to an Amazon RD5 tor MySQL DB instance. The application also outputs processed documents that are stored in an Amazon S3 bucket

The company's finance team directly queries the database to run reports. During busy periods, these queries consume resources and negatively affect application performance.

A solutions architect must design a solution that will provide resiliency during a disaster. The solution must minimize data loss and must resolve the performance problems that result from the finance team's queries.

Which solution will meet these requirements?

Options:

A- Migrate the database to Amazon DynamoDB and use DynamoDB global tables. Instruct the finance team to query a global table in a separate Region. Create an AWS Lambda function to periodically synchronize the contents of the original S3 bucket to a new S3 bucket in the separate Region. Launch EC2 instances and create an ALB in the separate Region. Configure the application to point to the new S3 bucket.
B- Launch additional EC2 instances that host the application in a separate Region. Add theadditional instances to the existing ALB. In the separate Region, create a read replica of the RDS DB instance. Instruct the finance team to run queries ageist the read replica. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 Docket in the separate Region. During a disaster, promote the read replace to a standalone DB instance. Configure the application to point to the new S3 bucket and to the newly project read replica.
C- Create a read replica of the RDS DB instance in a separate Region. Instruct the finance team to run queries against the read replica. Create AMIs of the EC2 instances mat host the application frontend- Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, promote the read replica to a standalone DB instance. Launch EC2 instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket.
D- Create hourly snapshots of the RDS DB instance. Copy the snapshots to a separate Region. Add an Amazon Elastic ache cluster m front of the existing RDS database. Create AMIs of the EC2 instances that host the application frontend Copy the AMIs to the separate Region. Use S3 Cross-Region Replication (CRR) from the original S3 bucket to a new S3 bucket in the separate Region. During a disaster, restore The database from the latest RDS snapshot. Launch EC2 Instances from the AMIs and create an ALB to present the application to end users. Configure the application to point to the new S3 bucket

Answer:

C

Explanation:

Implementing a disaster recovery strategy that minimizes data loss and addresses performance issues involves creating a read replica of the RDS DB instance in a separate region and directing the finance team's queries to this replica. This solution alleviates the performance impact on the primary database. Using Amazon S3 Cross-Region Replication (CRR) ensures that processed documents are available in the disaster recovery region. In the event of a disaster, the read replica can be promoted to a standalone DB instance, and EC2 instances can be launched from pre-created AMIs to serve the web frontend, thereby ensuring resiliency and minimal data loss.

AWS Documentation on Amazon RDS Read Replicas, Amazon S3 Cross-Region Replication, and Amazon EC2 AMIs provides comprehensive guidance on implementing a robust disaster recovery solution. This approach is in line with AWS best practices for high availability and disaster recovery planning.

Question 8

Question Type: MultipleChoice

A company recently completed the migration from an on-premises data center to the AWS Cloud by using a replatforming strategy. One of the migrated servers is running a legacy Simple Mail Transfer Protocol (SMTP) service that a critical application relies upon. The application sends outbound email messages to the company's customers. The legacy SMTP server does not support TLS encryption and uses TCP port 25. The application can use SMTP only.

The company decides to use Amazon Simple Email Service (Amazon SES) and to decommission the legacy SMTP server. The company has created and validated the SES domain. The company has lifted the SES limits.

What should the company do to modify the application to send email messages from Amazon SES?

Options:

A- Configure the application to connect to Amazon SES by using TLS Wrapper. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Attach the IAM role to an Amazon EC2 instance.
B- Configure the application to connect to Amazon SES by using STARTTLS. Obtain Amazon SES SMTP credentials. Use the credentials to authenticate with Amazon SES.
C- Configure the application to use the SES API to send email messages. Create an IAM role that has ses:SendEmail and ses:SendRawEmail permissions. Use the IAM role as a service role for Amazon SES.
D- Configure the application to use AWS SDKs to send email messages. Create an IAM user for Amazon SES. Generate API access keys. Use the access keys to authenticate with Amazon SES.

Answer:

B

Explanation:

To set up a STARTTLS connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 25, 587, or 2587, issues an EHLO command, and waits for the server toannounce that it supports the STARTTLS SMTP extension. The client then issues the STARTTLS command, initiating TLS negotiation. When negotiation is complete, the client issues an EHLO command over the new encrypted connection, and the SMTP session proceeds normally To set up a TLS Wrapper connection, the SMTP client connects to the Amazon SES SMTP endpoint on port 465 or 2465. The server presents its certificate, the client issues an EHLO command, and the SMTP session proceeds normally.

https://docs.aws.amazon.com/ses/latest/dg/smtp-connect.html

Question 9

Question Type: MultipleChoice

A company has multiple lines of business (LOBs) that toll up to the parent company. The company has asked its solutions architect to develop a solution with the following requirements

* Produce a single AWS invoice for all of the AWS accounts used by its LOBs.

* The costs for each LOB account should be broken out on the invoice

* Provide the ability to restrict services and features in the LOB accounts, as defined by the company's governance policy

* Each LOB account should be delegated full administrator permissions regardless of the governance policy

Which combination of steps should the solutions architect take to meet these requirements'? (Choose TWO.)

Options:

A- Use AWS Organizations to create an organization in the parent account for each LOB Then invite each LOB account to the appropriate organization
B- Use AWS Organizations to create a single organization in the parent account Then, invite each LOB's AWS account lo join the organization.
C- Implement service quotas to define the services and features that are permitted and apply the quotas to each LOB. as appropriate
D- Create an SCP that allows only approved services and features then apply the policy to the LOB accounts
E- Enable consolidated billing in the parent account's billing console and link the LOB accounts

Answer:

B, E

Explanation:

Create AWS Organization:

In the AWS Management Console, navigate to AWS Organizations and create a new organization in the parent account.

Invite LOB Accounts:

Invite each Line of Business (LOB) account to join the organization. This allows centralized management and governance of all accounts.

Enable Consolidated Billing:

Enable consolidated billing in the billing console of the parent account. Link all LOB accounts to ensure a single consolidated invoice that breaks down costs per account.

Apply Service Control Policies (SCPs):

Implement Service Control Policies (SCPs) to define the services and features permitted for each LOB account as per the governance policy, while still delegating full administrative permissions to the LOB accounts.

By consolidating billing and using AWS Organizations, the company can achieve centralized billing and governance while maintaining independent administrative control for each LOB account

Question 10

Question Type: MultipleChoice

A company operates a proxy server on a fleet of Amazon EC2 instances. Partners in different countries use the proxy server to test the company's functionality. The EC2 instances are running in a VPC. and the instances have access to the internet.

The company's security policy requires that partners can access resources only from domains that the company owns.

Which solution will meet these requirements?

Options:

A- Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. Configure a rule that has a low numeric value that allows requests for domains in the allowed list. Associate the rule group with the VPC.
B- Create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. Configure a Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Associate the domain list with the outbound endpoint.
C- Create an Amazon Route 53 traffic flow policy to match the allowed domains. Configure the traffic flow policy to forward requests that match to the Route 53 Resolver. Associate the traffic flow policy with the VPC.
D- Create an Amazon Route 53 outbound endpoint. Associate the outbound endpoint with the VPC. Configure a Route 53 traffic flow policy to forward requests for allowed domains to the outbound endpoint. Associate the traffic flow policy with the VPC.

Answer:

A

Explanation:

The company should create an Amazon Route 53 Resolver DNS Firewall domain list that contains the allowed domains. The company should configure a DNS Firewall rule group with a rule that has a high numeric value that blocks all requests. The company should configure a rule that has a low numeric value that allows requests for domains in the allowed list. The company should associate the rule group with the VPC. This solution will meet the requirements because Amazon Route 53 Resolver DNS Firewall is a feature that enables you to filter and regulate outbound DNS traffic for your VPC. You can create reusable collections of filtering rules in DNS Firewall rule groups and associate them with your VPCs.You can specify lists of domain names to allow or block, and you can customize theresponses for the DNS queries that you block1. By creating a domain list with the allowed domains and a rule group with rules to allow or block requests based on the domain list, the company can enforce its security policy and control access to sites.

The other options are not correct because:

Configuring a Route 53 outbound endpoint and associating it with the VPC would not help with filtering outbound DNS traffic.A Route 53 outbound endpoint is a resource that enables you to forward DNS queries from your VPC to your networkover AWS Direct Connect or VPN connections2.It does not provide any filtering capabilities.

Creating a Route 53 traffic flow policy to match the allowed domains would not help with filtering outbound DNS traffic.A Route 53 traffic flow policy is a resource that enables you to route traffic based on multiple criteria, such as endpoint health, geographic location, and latency3.It does not provide any filtering capabilities.

Creating a Gateway Load Balancer (GWLB) would not help with filtering outbound DNS traffic.A GWLB is a service that enables you to deploy, scale, and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems, and deep packet inspection systems in the cloud4.It does not provide any filtering capabilities.

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-dns-firewall.html

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-outbound-endpoints.html

https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/traffic-flow.html

https://docs.aws.amazon.com/elasticloadbalancing/latest/gateway/introduction.html

Question 11

Question Type: MultipleChoice

A company wants to refactor its retail ordering web application that currently has a load-balanced Amazon EC2 instance fleet for web hosting, database API services, and business logic. The company needs to create a decoupled, scalable architecture with a mechanism for retaining failed orders while also minimizing operational costs.

Which solution will meet these requirements?

Options:

A- Use Amazon S3 for web hosting with Amazon API Gateway for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use Amazon Elastic Container Service (Amazon ECS) for business logic with Amazon SQS long polling for retaining failed orders.
B- Use AWS Elastic Beanstalk for web hosting with Amazon API Gateway for database API services. Use Amazon MQ for order queuing. Use AWS Step Functionsfor business logic with Amazon S3 Glacier Deep Archive for retaining failed orders.
C- Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.
D- Use Amazon Lightsail for web hosting with AWS AppSync for database API services. Use Amazon Simple Email Service (Amazon SES) for order queuing. UseAmazon Elastic Kubernetes Service (Amazon EKS) for business logic with Amazon OpenSearch Service for retaining failed orders.

Answer:

C

Explanation:

*Use Amazon S3 for web hosting with AWS AppSync for database API services. Use Amazon Simple Queue Service (Amazon SQS) for order queuing. Use AWS Lambda for business logic with an Amazon SQS dead-letter queue for retaining failed orders.

This solution will allow you to:

*Host a static website on Amazon S3 without provisioning or managing servers1.

*Use AWS AppSync to create a scalable GraphQL API that connects to your database and other data sources1.

*Use Amazon SQS to decouple and scale your order processing microservices1.

*Use AWS Lambda to run code for your business logic without provisioning or managing servers1.

*Use an Amazon SQS dead-letter queue to retain messages that can't be processed by your Lambda function1.

P2P exams