



Free Questions for [SAP-C02](#) by [actualtestdumps](#)

Shared by [Jacobson](#) on [29-01-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

A company provides a software as a service (SaaS) application that runs in the AWS Cloud. The application runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The instances are in an Auto Scaling group and are distributed across three Availability Zones in a single AWS Region.

The company is deploying the application into additional Regions. The company must provide static IP addresses for the application to customers so that the customers can add the IP addresses to allow lists.

The solution must automatically route customers to the Region that is geographically closest to them.

Which solution will meet these requirements?

Options:

A- Create an Amazon CloudFront distribution. Create a CloudFront origin group. Add the NLB for each additional Region to the origin group. Provide customers with the IP address ranges of the distribution's edge locations.

B- Create an AWS Global Accelerator standard accelerator. Create a standard accelerator endpoint for the NLB in each additional Region. Provide customers with the Global Accelerator IP address.

C- Create an Amazon CloudFront distribution. Create a custom origin for the NLB in each additional Region. Provide customers with the IP address ranges of the distribution's edge locations.

D- Create an AWS Global Accelerator custom routing accelerator. Create a listener for the custom routing accelerator. Add the IP address and ports for the NLB in each additional Region. Provide customers with the Global Accelerator IP address.

Answer:

B

Explanation:

[What is AWS Global Accelerator?](#)

[Standard accelerator endpoints](#)

[AWS Global Accelerator IP addresses](#)

Question 2

Question Type: MultipleChoice

A company built an ecommerce website on AWS using a three-tier web architecture. The application is Java-based and composed of an Amazon CloudFront distribution, an Apache web server layer of Amazon EC2 instances in an Auto Scaling group, and a backend Amazon Aurora MySQL database.

Last month, during a promotional sales event, users reported errors and timeouts while adding items to their shopping carts. The operations team recovered the logs created by the web servers and reviewed Aurora DB cluster performance metrics. Some of the web servers were terminated before logs could be collected and the Aurora metrics were not sufficient for query performance analysis.

Which combination of steps must the solutions architect take to improve application performance visibility during peak traffic events?
(Choose three.)

Options:

- A-** Configure the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs.
- B-** Implement the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java.
- C-** Configure the Aurora MySQL DB cluster to stream slow query and error logs to Amazon Kinesis
- D-** Install and configure an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs.
- E-** Enable and configure AWS CloudTrail to collect and analyze application activity from Amazon EC2 and Aurora.
- F-** Enable Aurora MySQL DB cluster performance benchmarking and publish the stream to AWS X-Ray.

Answer:

A, B, D

Explanation:

Configuring the Aurora MySQL DB cluster to publish slow query and error logs to Amazon CloudWatch Logs will allow the solutions architect to monitor and troubleshoot the database performance by identifying slow or problematic queries¹. CloudWatch Logs also provides features such as metric filters, alarms, and dashboards to analyze and visualize the log data².

Implementing the AWS X-Ray SDK to trace incoming HTTP requests on the EC2 instances and implement tracing of SQL queries with the X-Ray SDK for Java will allow the solutions architect to measure and map the end-to-end latency and performance of the web application³. X-Ray traces show how requests travel through the application components, such as web servers, load balancers, microservices, and databases⁴. X-Ray also provides features such as service maps, annotations, histograms, and error rates to analyze and optimize the application performance.

Installing and configuring an Amazon CloudWatch Logs agent on the EC2 instances to send the Apache logs to CloudWatch Logs will allow the solutions architect to monitor and troubleshoot the web server performance by collecting and storing the Apache access and error logs. CloudWatch Logs also provides features such as metric filters, alarms, and dashboards to analyze and visualize the log data².

[Publishing Aurora MySQL logs to Amazon CloudWatch Logs](#)

[Working with log data in CloudWatch Logs](#)

[Instrumenting your application with the X-Ray SDK for Java](#)

[Tracing requests with AWS X-Ray](#)

[\[Analyzing application performance with AWS X-Ray\]](#)

[\[Using CloudWatch Logs with your Apache web server\]](#)

Question 3

Question Type: MultipleChoice

A company has a solution that analyzes weather data from thousands of weather stations. The weather stations send the data over an Amazon API Gateway REST API that has an AWS Lambda function integration. The Lambda function calls a third-party service for data pre-processing. The third-party service gets overloaded and fails the pre-processing, causing a loss of data.

A solutions architect must improve the resiliency of the solution. The solutions architect must ensure that no data is lost and that data can be processed later if failures occur.

What should the solutions architect do to meet these requirements?

Options:

- A-** Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure the queue as the dead-letter queue for the API.
- B-** Create two Amazon Simple Queue Service (Amazon SQS) queues: a primary queue and a secondary queue. Configure the secondary queue as the dead-letter queue for the primary queue. Update the API to use a new integration to the primary queue. Configure the Lambda function as the invocation target for the primary queue.
- C-** Create two Amazon EventBridge event buses: a primary event bus and a secondary event bus. Update the API to use a new integration to the primary event bus. Configure an EventBridge rule to react to all events on the primary event bus. Specify the Lambda function as the target of the rule. Configure the secondary event bus as the failure destination for the Lambda function.

D- Create a custom Amazon EventBridge event bus. Configure the event bus as the failure destination for the Lambda function.

Answer:

C

Explanation:

Using Amazon EventBridge with AWS Lambda

Using multiple event buses

Using failure destinations

[Using dead-letter queues]

=====

Question 4

Question Type: MultipleChoice

A company has a web application that securely uploads pictures and videos to an Amazon S3 bucket. The company requires that only authenticated users are allowed to post content. The application generates a presigned URL that is used to upload objects through a browser interface. Most users are reporting slow upload times for objects larger than 100 MB.

What can a Solutions Architect do to improve the performance of these uploads while ensuring only authenticated users are allowed to post content?

Options:

- A-** Set up an Amazon API Gateway with an edge-optimized API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using a COGNITO_USER_POOLS authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload objects.
- B-** Set up an Amazon API Gateway with a regional API endpoint that has a resource as an S3 service proxy. Configure the PUT method for this resource to expose the S3 PutObject operation. Secure the API Gateway using an AWS Lambda authorizer. Have the browser interface use API Gateway instead of the presigned URL to upload API objects.
- C-** Enable an S3 Transfer Acceleration endpoint on the S3 bucket. Use the endpoint when generating the presigned URL. Have the browser interface upload the objects to this URL using the S3 multipart upload API.
- D-** Configure an Amazon CloudFront distribution for the destination S3 bucket. Enable PUT and POST methods for the CloudFront cache behavior. Update the CloudFront origin to use an origin access identity (OAI). Give the OAI user s3:PutObject permissions in the bucket policy. Have the browser interface upload objects using the CloudFront distribution.

Answer:

C

Explanation:

S3 Transfer Acceleration

Using Transfer Acceleration with presigned URLs

Uploading objects using multipart upload API

Question 5

Question Type: MultipleChoice

A company is deploying a third-party web application on AWS. The application is packaged as a Docker image. The company has deployed the Docker image as an AWS

Fargate service in Amazon Elastic Container Service (Amazon ECS). An Application Load Balancer (ALB) directs traffic to the application.

The company needs to give only a specific list of users the ability to access the application from the internet. The company cannot change the application and cannot integrate the application with an identity provider. All users must be authenticated through multi-factor authentication (MFA).

Which solution will meet these requirements?

Options:

A- Create a user pool in Amazon Cognito. Configure the pool for the application. Populate the pool with the required users. Configure the pool to require MFA.

Configure a listener rule on the ALB to require authentication through the Amazon Cognito hosted UI.

B- Configure the users in AWS Identity and Access Management (IAM). Attach a resource policy to the Fargate service to require users to use MFA. Configure a

listener rule on the ALB to require authentication through IAM.

C- Configure the users in AWS Identity and Access Management (IAM). Enable AWS IAM Identity Center (AWS Single Sign-On).

Configure resource protection for

the ALB. Create a resource protection rule to require users to use MFA.

D- Create a user pool in AWS Amplify. Configure the pool for the application. Populate the pool with the required users. Configure the pool to require MFA.

Configure a listener rule on the ALB to require authentication through the Amplify hosted UI.

Answer:

A

Explanation:

Creating a user pool in Amazon Cognito and configuring it for the application will meet the requirement of giving only a specific list of users the ability to access the application from the internet. A user pool is a directory of users that can sign in to an application with a username and password¹. The company can populate the user pool with the required users and configure the pool to require MFA for additional security². Configuring a listener rule on the ALB to require authentication through the Amazon Cognito hosted UI will meet the requirement of not changing the application and not integrating it with an identity provider. The ALB can use Amazon Cognito as an authentication action to authenticate users before forwarding requests to the Fargate service³. The Amazon Cognito hosted UI is a customizable web page that provides sign-in and sign-up functionality for users⁴.

Question 6

Question Type: MultipleChoice

A company is planning to migrate its on-premises transaction-processing application to AWS. The application runs inside Docker containers that are hosted on VMS in the company's data center. The Docker containers have shared storage where the application records transaction data.

The transactions are time sensitive. The volume of transactions inside the application is unpredictable. The company must implement a low-latency storage solution that will automatically scale throughput to meet increased demand. The company cannot develop the application further and cannot continue to administer the Docker hosting environment.

How should the company migrate the application to AWS to meet these requirements?

Options:

- A-** Migrate the containers that run the application to Amazon Elastic Kubernetes Service (Amazon EKS). Use Amazon S3 to store the transaction data that the containers share.
- B-** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic File System (Amazon EFS) file system. Create a Fargate task definition. Add a volume to the task definition to point to the EFS file system
- C-** Migrate the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS). Create an Amazon Elastic Block Store (Amazon EBS) volume. Create a Fargate task definition. Attach the EBS volume to each running task.
- D-** Launch Amazon EC2 instances. Install Docker on the EC2 instances. Migrate the containers to the EC2 instances. Create an Amazon Elastic File System (Amazon EFS) file system. Add a mount point to the EC2 instances for the EFS file system.

Answer:

B

Explanation:

Migrating the containers that run the application to AWS Fargate for Amazon Elastic Container Service (Amazon ECS) will meet the requirement of not administering the Docker hosting environment. AWS Fargate is a serverless compute engine that runs containers without requiring any infrastructure management³. Creating an Amazon Elastic File System (Amazon EFS) file system and adding a volume to the Fargate task definition to point to the EFS file system will meet the requirement of low-latency storage that will automatically scale throughput to meet increased demand. Amazon EFS is a fully managed file system service that provides shared access to data from multiple containers, supports NFSv4 protocol, and offers consistent performance and high availability⁴. Amazon EFS

also supports automatic scaling of throughput based on the amount of data stored in the file system⁵.

Question 7

Question Type: MultipleChoice

A company has a legacy application that runs on multiple .NET Framework components. The components share the same Microsoft SQL Server database and

communicate with each other asynchronously by using Microsoft Message Queueing (MSMQ).

The company is starting a migration to containerized .NET Core components and wants to refactor the application to run on AWS. The .NET Core components require complex orchestration. The company must have full control over networking and host configuration. The application's database model is strongly relational.

Which solution will meet these requirements?

Options:

A- Host the .NET Core components on AWS App Runner. Host the database on Amazon RDS for SQL Server. Use Amazon EventBridge for asynchronous messaging.

- B-** Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the AWS Fargate launch type. Host the database on Amazon DynamoDB. Use Amazon Simple Notification Service (Amazon SNS) for asynchronous messaging.
- C-** Host the .NET Core components on AWS Elastic Beanstalk. Host the database on Amazon Aurora PostgreSQL Serverless v2. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) for asynchronous messaging.
- D-** Host the .NET Core components on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. Host the database on Amazon Aurora MySQL Serverless v2. Use Amazon Simple Queue Service (Amazon SQS) for asynchronous messaging.

Answer:

D

Explanation:

Hosting the .NET Core components on Amazon ECS with the Amazon EC2 launch type will meet the requirements of having complex orchestration and full control over networking and host configuration. Amazon ECS is a fully managed container orchestration service that supports both AWS Fargate and Amazon EC2 as launch types. The Amazon EC2 launch type allows users to choose their own EC2 instances, configure their own networking settings, and access their own host operating systems. Hosting the database on Amazon Aurora MySQL Serverless v2 will meet the requirements of having a strongly relational database model and using the same database engine as SQL Server. MySQL is a compatible relational database engine with SQL Server, and it can support most of the legacy application's database model. Amazon Aurora MySQL Serverless v2 is a serverless version of Amazon Aurora MySQL that can scale up and down automatically based on demand. Using Amazon SQS for asynchronous messaging will meet the requirements of providing a compatible replacement for MSMQ, which is a queue-based messaging system³. Amazon SQS is a fully managed message queuing service that enables decoupled and scalable microservices, distributed systems, and serverless applications.

Question 8

Question Type: MultipleChoice

A company is preparing to deploy an Amazon Elastic Kubernetes Service (Amazon EKS) cluster for a workload. The company expects the cluster to support an

unpredictable number of stateless pods. Many of the pods will be created during a short time period as the workload automatically scales the number of replicas that the workload uses.

Which solution will MAXIMIZE node resilience?

Options:

- A-** Use a separate launch template to deploy the EKS control plane into a second cluster that is separate from the workload node groups.
- B-** Update the workload node groups. Use a smaller number of node groups and larger instances in the node groups.
- C-** Configure the Kubernetes Cluster Autoscaler to ensure that the compute capacity of the workload node groups stays under provisioned.
- D-** Configure the workload to use topology spread constraints that are based on Availability Zone.

Answer:

D

Explanation:

Configuring the workload to use topology spread constraints that are based on Availability Zone will maximize the node resilience of the workload node groups. This will ensure that the pods are evenly distributed across different Availability Zones, reducing the impact of failures or disruptions in one Availability Zone². This will also improve the availability and scalability of the workload node groups, as they can leverage the low-latency, high-throughput, and highly redundant networking between Availability Zones¹.

Question 9

Question Type: MultipleChoice

An online gaming company needs to optimize the cost of its workloads on AWS. The company uses a dedicated account to host the production environment for its online gaming application and an analytics application.

Amazon EC2 instances host the gaming application and must always be available. The EC2 instances run all year. The analytics application uses data that is stored in Amazon S3. The analytics application can be interrupted and resumed without issue.

Which solution will meet these requirements MOST cost-effectively?

Options:

- A-** Purchase an EC2 Instance Savings Plan for the online gaming application instances. Use On-Demand Instances for the analytics application.
- B-** Purchase an EC2 Instance Savings Plan for the online gaming application instances. Use Spot Instances for the analytics application.
- C-** Use Spot Instances for the online gaming application and the analytics application. Set up a catalog in AWS Service Catalog to provision services at a discount.
- D-** Use On-Demand Instances for the online gaming application. Use Spot Instances for the analytics application. Set up a catalog in AWS Service Catalog to provision services at a discount.

Answer:

B

Explanation:

The correct answer is B.

B) This solution is the most cost-effective because it uses an EC2 Instance Savings Plan for the online gaming application instances, which provides the lowest prices and savings up to 72% compared to On-Demand prices. The EC2 Instance Savings Plan applies to any instance size within the same family and region, regardless of availability zone, operating system, or tenancy. The online gaming application instances run all year and must always be available, so they are not suitable for Spot Instances, which can be interrupted with a two-minute notice. This solution also uses Spot Instances for the analytics application, which can reduce the cost by up to 90% compared to On-Demand prices. The analytics application can be interrupted and resumed without issue, so it is a good fit for Spot

Instances, which use spare EC2 capacity. This solution does not require AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use on AWS, but does not provide any discounts¹²³

A) This solution is incorrect because it uses On-Demand Instances for the analytics application, which are more expensive than Spot Instances. The analytics application can be interrupted and resumed without issue, so it can benefit from the lower cost of Spot Instances, which use spare EC2 capacity.

C) This solution is incorrect because it uses Spot Instances for the online gaming application, which can be interrupted with a two-minute notice. The online gaming application instances must always be available, so they are not suitable for Spot Instances, which use spare EC2 capacity. This solution also uses AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use on AWS, but does not provide any discounts.

D) This solution is incorrect because it uses On-Demand Instances for the online gaming application, which are more expensive than an EC2 Instance Savings Plan. The online gaming application instances run all year and must always be available, so they are suitable for an EC2 Instance Savings Plan, which provides the lowest prices and savings up to 72% compared to On-Demand prices. This solution also uses AWS Service Catalog, which is a service that helps to create and manage catalogs of IT services that are approved for use on AWS, but does not provide any discounts.

[1: EC2 Instance Savings Plans -- Amazon Web Services](#) [2: Amazon EC2 Spot Instances](#) [3: Cloud Management and Governance -- AWS Service Catalog -- Amazon Web Services](#)

Question 10

Question Type: MultipleChoice

A company has AWS accounts that are in an organization in AWS rganizations. The company wants to track Amazon EC2 usage as a metric.

The company's architecture team must receive a daily alert if the EC2 usage is more than 10% higher than the average EC2 usage from the last 30 days.

Which solution will meet these requirements?

Options:

- A-** Configure AWS Budgets in the organization's management account. Specify a usage type of EC2 running hours. Specify a daily period. Set the budget amount to be 10% more than the reported average usage for the last 30 days from AWS Cost Explorer.
- B-** Configure an alert to notify the architecture team if the usage threshold is met. Configure AWS Cost Anomaly Detection in the organization's management account. Configure a monitor type of AWS Service. Apply a filter of Amazon EC2. Configure an alert subscription to notify the architecture team if the usage is 10% more than the average usage for the last 30 days.
- C-** Enable AWS Trusted Advisor in the organization's management account. Configure a cost optimization advisory alert to notify the architecture team if the EC2 usage is 10% more than the reported average usage for the last 30 days.
- D-** Configure Amazon Detective in the organization's management account. Configure an EC2 usage anomaly alert to notify the architecture team if Detective identifies a usage anomaly of more than 10%.

Answer:

B

Explanation:

The correct answer is B.

B) This solution meets the requirements because it uses AWS Cost Anomaly Detection, which is a feature of AWS Cost Management that uses machine learning to identify and alert on anomalous spend and usage patterns. By configuring a monitor type of AWS Service and applying a filter of Amazon EC2, the solution can track the EC2 usage as a metric across the organization's accounts. By configuring an alert subscription with a threshold of 10%, the solution can notify the architecture team via email or Amazon SNS if the EC2 usage is more than 10% higher than the average usage for the last 30 days¹²

A) This solution is incorrect because it uses AWS Budgets, which is a feature of AWS Cost Management that helps to plan and track costs and usage. However, AWS Budgets does not support usage type of EC2 running hours as a budget type. The only supported usage types are Amazon S3 storage, Amazon EC2 RI utilization, and Amazon EC2 RI coverage. Moreover, AWS Budgets does not support setting the budget amount based on the reported average usage from AWS Cost Explorer. The budget amount has to be a fixed or variable value³⁴

C) This solution is incorrect because it uses AWS Trusted Advisor, which is a feature of AWS Premium Support that provides recommendations to follow best practices for cost optimization, security, performance, and fault tolerance. However, AWS Trusted Advisor does not support configuring custom alerts based on EC2 usage or average usage for the last 30 days. The only supported alerts are based on predefined checks and thresholds that are applied to all services and resources in the account⁵⁶

D) This solution is incorrect because it uses Amazon Detective, which is a service that helps to analyze and visualize security data to investigate potential security issues. However, Amazon Detective does not support configuring EC2 usage anomaly alerts based on average usage for the last 30 days. The only supported alerts are based on GuardDuty findings and other security-related events that are detected by machine learning models⁷⁸

1: [AWS Cost Anomaly Detection - Amazon Web Services](#) 2: [Getting started with AWS Cost Anomaly Detection](#) 3: [Set Custom Cost and Usage Budgets -- AWS Budgets -- Amazon Web Services](#) 4: [Creating a budget - AWS Cost Management](#) 5: [AWS Trusted Advisor](#) 6: [AWS Trusted Advisor - AWS Support](#) 7: [Security Investigation Visualization - Amazon Detective - AWS](#) 8: [What is Amazon Detective? - Amazon Detective](#)

To Get Premium Files for SAP-C02 Visit

<https://www.p2pexams.com/products/sap-c02>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/sap-c02>

