



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

Shared by [Holden](#) on [24-05-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

A company needs to migrate an on-premises SFTP site to AWS. The SFTP site currently runs on a Linux VM. Uploaded files are made available to downstream applications through an NFS share.

As part of the migration to AWS, a solutions architect must implement high availability. The solution must provide external vendors with a set of static public IP addresses that the vendors can allow. The company has set up an AWS Direct Connect connection between its on-premises data center and its VPC.

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

Options:

A- Create an AWS Transfer Family server, configure an internet-facing VPC endpoint for the Transfer Family server, specify an Elastic IP address for each subnet, configure the Transfer Family server to place files into an Amazon Elastic File System (Amazon EFS) file system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.

B- Create an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family server to place files into an Amazon Elastic File System [Amazon EFS} the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the its endpoint instead.

C- Use AWS Application Migration service to migrate the existing Linux VM to an Amazon EC2 instance. Assign an Elastic IP address to the EC2 instance. Mount an Amazon Elastic File System (Amazon EFS) the system to the EC2 instance. Configure the SFTP server to place files in. the EFS file system. Modify the configuration on the downstream applications that access the existing NFS share to mount the EFS endpoint instead.

D- Use AWS Application Migration Service to migrate the existing Linux VM to an AWS Transfer Family server. Configure a publicly accessible endpoint for the Transfer Family server. Configure the Transfer Family sever to place files into an Amazon FSx for Luster the system that is deployed across multiple Availability Zones. Modify the configuration on the downstream applications that access the existing NFS share to mount the FSx for Luster endpoint instead.

Answer:

A

Explanation:

To migrate an on-premises SFTP site to AWS with high availability and a set of static public IP addresses for external vendors, the best solution is to create an AWS Transfer Family server with an internet-facing VPC endpoint. Assigning Elastic IP addresses to each subnet and configuring the server to store files in an Amazon Elastic File System (EFS) that spans multiple Availability Zones ensures high availability and consistent access. This approach minimizes operational overhead by leveraging AWS managed services and eliminates the need to manage underlying infrastructure.

Question 2

Question Type: MultipleChoice

A company uses AWS Organizations AWS account. A solutions architect must design a solution in which only administrator roles are allowed to use IAM actions. However the solutions architect does not have access to all the AWS accounts throughout the company.

Which solution meets these requirements with the LEAST operational overhead?

Options:

- A-** Create an SCP that applies to all the AWS accounts to allow IAM actions only for administrator roles. Apply the SCP to the root OLI.
- B-** Configure AWS CloudTrail to invoke an AWS Lambda function for each event that is related to IAM actions. Configure the function to deny the action. If the user who invoked the action is not an administrator.
- C-** Create an SCP that applies to all the AWS accounts to deny IAM actions for all users except for those with administrator roles. Apply the SCP to the root OU.
- D-** Set an IAM permissions boundary that allows IAM actions. Attach the permissions boundary to every administrator role across all the AWS accounts.

Answer:

A

Explanation:

To restrict IAM actions to only administrator roles across all AWS accounts in an organization, the most operationally efficient solution is to create a Service Control Policy (SCP) that allows IAM actions exclusively for administrator roles and apply this SCP to the root Organizational Unit (OU) of AWS Organizations. This method ensures a centralized governance mechanism that uniformly applies the policy across all accounts, thereby minimizing the need for individual account-level configurations and reducing operational complexity.

Question 3

Question Type: MultipleChoice

A company use an organization in AWS Organizations to manage multiple AWS accounts. The company hosts some applications in a VPC in the company's snared services account. The company has attached a transit gateway to the VPC in the Shared services account.

The company is developing a new capability and has created a development environment that requires access to the applications that are in the snared services account. The company intends to delete and recreate resources frequently in the development account. The company also wants to give a development team the ability to recreate the team's connection to the shared services account as required.

Which solution will meet these requirements?

Options:

- A-** Create a transit gateway in the development account. Create a transit gateway peering request to the shared services account. Configure the shared services transit gateway to automatically accept peering connections.
- B-** Turn on automate acceptance for the transit gateway in the shared services account. Use AWS Resource Access Manager (AWS RAM) to share the transit gateway resource in the shared services account with the development account. Accept the resource in the development account. Create a transit gateway attachment in the development account.
- C-** Turn on automate acceptance for the transit gateway in the shared services account. Create a VPC endpoint. Use the endpoint policy to grant permissions on the VPC endpoint for the development account. Configure the endpoint service to automatically accept connection requests. Provide the endpoint details to the development team.
- D-** Create an Amazon EventBridge rule to invoke an AWS Lambda function that accepts the transit gateway attachment value the development account makes an attachment request. Use AWS Network Manager to store the transit gateway in the shared services account with the development account. Accept the transit gateway in the development account.

Answer:

B

Explanation:

For a development environment that requires frequent resource recreation and connectivity to applications hosted in a shared services account, the most efficient solution involves using AWS Resource Access Manager (RAM) and the transit gateway in the shared services account. By turning on automatic acceptance for the transit gateway in the shared services account and sharing it with the development account through AWS RAM, the development team can easily recreate their connection as needed without manual intervention. This setup allows for scalable, flexible connectivity between accounts while minimizing operational overhead and ensuring consistent access to shared services.

Question 4

Question Type: MultipleChoice

A company that is developing a mobile game is making game assets available in two AWS Regions. Game assets are served from a set of Amazon EC2 instances behind an Application Load Balancer (ALB) in each Region. The company requires game assets to be fetched from the closest Region. If game assets become unavailable in the closest Region, they should be fetched from the other Region.

What should a solutions architect do to meet these requirements?

Options:

- A-** Create an Amazon CloudFront distribution. Create an origin group with one origin for each ALB. Set one of the origins as primary.
- B-** Create an Amazon Route 53 health check for each ALB. Create a Route 53 failover routing record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.
- C-** Create two Amazon CloudFront distributions, each with one ALB as the origin. Create an Amazon Route 53 failover routing record pointing to the two CloudFront distributions. Set the Evaluate Target Health value to Yes.
- D-** Create an Amazon Route 53 health check for each ALB. Create a Route 53 latency alias record pointing to the two ALBs. Set the Evaluate Target Health value to Yes.

Answer:

A

Explanation:

To ensure that game assets are fetched from the closest region and have a fallback option in case the assets become unavailable in the closest region, a solution architect should leverage Amazon CloudFront, a global content delivery network (CDN) service. By creating an Amazon CloudFront distribution and setting up origin groups, the architect can specify multiple origins (in this case, the Application Load Balancers in each region). The primary origin will serve content under normal circumstances, and if the content becomes unavailable, CloudFront will automatically switch to the secondary origin. This approach not only meets the requirement of regional proximity and redundancy but also optimizes latency and enhances the gaming experience by serving assets from the nearest geographical location to the end-user.

Question 5

Question Type: MultipleChoice

A company is developing an application that will display financial reports. The company needs a solution that can store financial information that comes from multiple systems. The solution must provide the reports through a web interface and must serve the data with less than 500 milliseconds of latency to end users. The solution also must be highly available and must have an RTO of 30 seconds.

Which solution will meet these requirements?

Options:

- A-** Use an Amazon Redshift cluster to store the data. Use a state website that is hosted on Amazon S3 with backend APIs that are served by an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to provide the reports to the application.
- B-** Use Amazon S3 to store the data. Use Amazon Athena to provide the reports to the application. Use AWS App Runner to serve the application to view the reports.
- C-** Use Amazon DynamoDB to store the data, use an embedded Amazon QuickSight dashboard with direct Query datasets to provide the reports to the application.
- D-** Use Amazon Keyspaces (for Apache Cassandra) to store the data, use AWS Elastic Beanstalk to provide the reports to the application.

Answer:

C

Explanation:

For an application requiring low-latency access to financial information and high availability with a Recovery Time Objective (RTO) of 30 seconds, using Amazon DynamoDB for data storage and Amazon QuickSight for reporting is the most suitable solution. DynamoDB offers fast, consistent, and single-digit millisecond latency for data retrieval, meeting the latency requirements. QuickSight's ability to

directly query DynamoDB datasets and provide embedded dashboards for reporting enables real-time financial report generation. This combination ensures high availability and meets the RTO requirement, providing a robust solution for the application's needs.

Amazon DynamoDB Documentation: Describes the features and benefits of DynamoDB, emphasizing its performance and scalability for applications requiring low-latency access to data.

Amazon QuickSight Documentation: Provides information on using QuickSight for creating and embedding interactive dashboards, including direct querying of DynamoDB datasets for real-time data visualization.

Question 6

Question Type: MultipleChoice

A company has many services running in its on-premises data center. The data center is connected to AWS using AWS Direct Connect (DX) and an IPsec VPN. The service data is sensitive and connectivity cannot traverse the internet. The company wants to expand to a new market segment and begin offering its services to other companies that are using AWS.

Which solution will meet these requirements?

Options:

A- Create a VPC Endpoint Service that accepts TCP traffic, host it behind a Network Load Balancer, and make the service available

over DX.

B- Create a VPC Endpoint Service that accepts HTTP or HTTPS traffic, host it behind an Application Load Balancer, and make the service available over DX.

C- Attach an internet gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

D- Attach a NAT gateway to the VPC. and ensure that network access control and security group rules allow the relevant inbound and outbound traffic.

Answer:

B

Explanation:

To offer services to other companies using AWS without traversing the internet, creating a VPC Endpoint Service hosted behind an Application Load Balancer (ALB) and making it available over AWS Direct Connect (DX) is the most suitable solution. This approach ensures that the service traffic remains within the AWS network, adhering to the requirement that connectivity must not traverse the internet. An ALB is capable of handling HTTP/HTTPS traffic, making it appropriate for web-based services. Utilizing DX for connectivity between the on-premises data center and AWS further secures and optimizes the network path.

AWS Direct Connect Documentation: Explains how to set up DX for private connectivity between AWS and an on-premises network.

Amazon VPC Endpoint Services (AWS PrivateLink) Documentation: Provides details on creating and configuring endpoint services for private, secure access to services hosted in AWS.

AWS Application Load Balancer Documentation: Offers guidance on configuring ALBs to distribute HTTP/HTTPS traffic efficiently.

Question 7

Question Type: MultipleChoice

A company has Linux-based Amazon EC2 instances. Users must access the instances by using SSH with EC2 SSH Key pairs. Each machine requires a unique EC2 Key pair.

The company wants to implement a key rotation policy that will, upon request, automatically rotate all the EC2 key pairs and keep the key in a securely encrypted place. The company will accept less than 1 minute of downtime during key rotation.

Which solution will meet these requirements?

Options:

A- Store all the keys in AWS Secrets Manager. Define a Secrets Manager rotation schedule to invoke an AWS Lambda function to generate new key pairs. Replace public Keys on EC2 instances. Update the private keys in Secrets Manager.

- B-** Store all the keys in Parameter. Store, a capability of AWS Systems Manager, as a string. Define a Systems Manager maintenance window to invoke an AWS Lambda function to generate new key pairs. Replace public keys on EC2 instance. Update the private keys in parameter.
- C-** Import the EC2 key pairs into AWS Key Management Service (AWS KMS). Configure automatic key rotation for these key pairs. Create an Amazon EventBridge scheduled rule to invoke an AWS Lambda function to initiate the key rotation AWS KMS.
- D-** Add all the EC2 instances to Fleet Manager, a capability of AWS Systems Manager. Define a Systems Manager maintenance window to issue a Systems Manager Run Command document to generate new Key pairs and to rotate public keys to all the instances in Fleet Manager.

Answer:

A

Explanation:

To meet the requirements for automatic key rotation of EC2 SSH key pairs with minimal downtime, storing the keys in AWS Secrets Manager and defining a rotation schedule is the most suitable solution. AWS Secrets Manager supports automatic rotation of secrets, including SSH keys, by invoking a Lambda function that can handle the creation of new key pairs and the replacement of public keys on EC2 instances. Updating the corresponding private keys in Secrets Manager ensures secure and centralized management of SSH keys, complying with the key rotation policy and minimizing operational overhead.

AWS Secrets Manager Documentation: Describes how to store and rotate secrets, including SSH keys, using Secrets Manager and Lambda functions.

AWS Lambda Documentation: Provides information on creating Lambda functions for custom secret rotation logic.

AWS Best Practices for Security: Highlights the importance of key rotation and how AWS services like Secrets Manager can facilitate secure and automated key management.

Question 8

Question Type: MultipleChoice

A company is using an organization in AWS organization to manage AWS accounts. For each new project the company creates a new linked account. After the creation of a new account, the root user signs in to the new account and creates a service request to increase the service quota for Amazon EC2 instances. A solutions architect needs to automate this process.

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

Options:

A- Create an Amazon EventBridge rule to detect creation of a new account Send the event to an Amazon Simple Notification Service (Amazon SNS) topic that invokes an AWS Lambda function. Configure the Lambda function to run the request-service-quota-increase command to request a service quota increase for EC2 instances.

B- Create a Service Quotas request template in the management account. Configure the desired service quota increases for EC2 instances.

C- Create an AWS Config rule in the management account to set the service quota for EC2 instances.

D- Create an Amazon EventBridge rule to detect creation of a new account. Send the event to an Amazon simple Notification service (Amazon SNS) topic that involves an AWS Lambda function. Configure the Lambda function to run the create-case command to request a service quota increase for EC2 instances.

Answer:

A

Explanation:

Automating the process of increasing service quotas for Amazon EC2 instances in new AWS accounts with minimal operational overhead can be effectively achieved by using Amazon EventBridge, Amazon SNS, and AWS Lambda. An EventBridge rule can detect the creation of a new account and trigger an SNS topic, which in turn invokes a Lambda function. This function can then programmatically request a service quota increase for EC2 instances using the AWS Service Quotas API. This approach streamlines the process, reduces manual intervention, and ensures that new accounts are automatically configured with the desired service quotas.

Amazon EventBridge Documentation: Provides guidance on setting up event rules for detecting AWS account creation.

AWS Lambda Documentation: Details how to create and configure Lambda functions to perform automated tasks, such as requesting service quota increases.

AWS Service Quotas Documentation: Offers information on managing and requesting increases for AWS service quotas programmatically.

Question 9

Question Type: MultipleChoice

A company's factory and automaton applications are running in a single VPC. More than 23 applications run on a combination of Amazon EC2, Amazon Elastic Container Service (Amazon ECS), and Amazon RDS.

The company has software engineers spread across three teams. One of the three teams owns each application, and each team is responsible for the cost and performance of all of its applications. Team resources have tags that represent their application and team. The teams use IAM access for daily activities.

The company needs to determine which costs on the monthly AWS bill are attributable to each application or team. The company also must be able to create reports to compare costs over the last 12 months and to help forecast costs for the next 12 months. A solution architect must recommend an AWS Billing and Cost Management solution that provides these cost reports.

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

Options:

- A-** Activate the user-defined cost allocation tags that represent the application and the team.
- B-** Activate the AWS generated cost allocation tags that represent the application and the team.

- C-** Create a cost category for each application in Billing and Cost Management
- D-** Activate IAM access to Billing and Cost Management.
- E-** Create a cost budget
- F-** Enable Cost Explorer.

Answer:

A, C, F

Explanation:

To attribute AWS costs to specific applications or teams and enable detailed cost analysis and forecasting, the solution architect should recommend the following actions: A. Activating user-defined cost allocation tags for resources associated with each application and team allows for detailed tracking of costs by these identifiers. C. Creating a cost category for each application within AWS Billing and Cost Management enables the organization to group costs according to application, facilitating detailed reporting and analysis. F. Enabling Cost Explorer is essential for analyzing and visualizing AWS spending over time. It provides the capability to view historical costs and forecast future expenses, supporting the company's requirement for cost comparison and forecasting.

AWS Billing and Cost Management Documentation: Covers the activation of cost allocation tags, creation of cost categories, and the use of Cost Explorer for cost management.

AWS Tagging Strategies: Provides best practices for implementing tagging strategies that support cost allocation and reporting.

AWS Cost Explorer Documentation: Details how to use Cost Explorer to analyze and forecast AWS costs.

Question 10

Question Type: MultipleChoice

A company is planning a migration from an on-premises data center to the AWS cloud. The company plans to use multiple AWS accounts that are managed in an organization in AWS organizations. The company will cost a small number of accounts initially and will add accounts as needed. A solution architect must design a solution that turns on AWS accounts.

What is the MOST operationally efficient solution that meets these requirements.

Options:

- A-** Create an AWS Lambda function that creates a new cloudTrail trail in all AWS account in the organization. Invoke the Lambda function dally by using a scheduled action in Amazon EventBridge.
- B-** Create a new CloudTrail trail in the organizations management account. Configure the trail to log all events for all AYYs accounts in the organization.
- C-** Create a new CloudTrail trail in all AWS accounts in the organization. Create new trails whenever a new account is created.
- D-** Create an AWS systems Manager Automaton runbook that creates a cloud trail in all AWS accounts in the organization. Invoke the automation by using Systems Manager State Manager.

Answer:

B

Explanation:

The most operationally efficient solution for turning on AWS CloudTrail across multiple AWS accounts managed within an AWS Organization is to create a single CloudTrail trail in the organization's management account and configure it to log events for all accounts within the organization. This approach leverages CloudTrail's ability to consolidate logs from all accounts in an organization, thereby simplifying management, reducing overhead, and ensuring consistent logging across accounts. This method eliminates the need for manual intervention in each account, making it an operationally efficient choice for organizations planning to scale their AWS usage.

AWS CloudTrail Documentation: Provides detailed instructions on setting up CloudTrail, including how to configure it for an organization.

AWS Organizations Documentation: Offers insights into best practices for managing multiple AWS accounts and how services like CloudTrail integrate with AWS Organizations.

AWS Best Practices for Security and Governance: Guides on how to effectively use AWS services to maintain a secure and well-governed AWS environment, with a focus on centralized logging and monitoring.

Question 11

Question Type: MultipleChoice

A company is planning to migrate an on-premises data center to AWS. The company currently hosts the data center on Linux-based VMware VMs. A solutions architect must collect information about network dependencies between the VMs. The information must be in the form of a diagram that details host IP addresses, hostnames, and network connection information.

Which solution will meet these requirements?

Options:

- A-** Use AWS Application Discovery Service. Select an AWS Migration Hub home AWS Region. Install the AWS Application Discovery Agent on the on-premises servers for data collection. Grant permissions to Application Discovery Service to use the Migration Hub network diagrams.
- B-** Use the AWS Application Discovery Service Agentless Collector for server data collection. Export the network diagrams from the AWS Migration Hub in .png format.
- C-** Install the AWS Application Migration Service agent on the on-premises servers for data collection. Use AWS Migration Hub data in Workload Discovery on AWS to generate network diagrams.
- D-** Install the AWS Application Migration Service agent on the on-premises servers for data collection. Export data from AWS Migration Hub in .csv format into an Amazon CloudWatch dashboard to generate network diagrams.

Answer:

B

Explanation:

To effectively gather information about network dependencies between VMs in an on-premises data center for migration to AWS, it's crucial to use tools that can capture detailed application and server dependencies. The AWS Application Discovery Service is designed for this purpose, particularly when migrating from environments like Linux-based VMware VMs. By installing the AWS Application Discovery Agent on the on-premises servers, the service can collect necessary data such as host IP addresses, hostnames, and network connection information. This data is crucial for creating a comprehensive network diagram that outlines the interactions and dependencies between various components of the on-premises infrastructure. The integration with AWS Migration Hub enhances this process by allowing the visualization of these dependencies in a network diagram format, aiding in the planning and execution of the migration process. This approach ensures a thorough understanding of the on-premises environment, which is essential for a successful migration to AWS.

AWS Documentation on Application Discovery Service: This provides detailed guidance on how to use the Application Discovery Service, including the installation and configuration of the Discovery Agent.

AWS Migration Hub User Guide: Offers insights on how to integrate Application Discovery Service data with Migration Hub for comprehensive migration planning and tracking.

AWS Solutions Architect Professional Learning Path: Contains advanced topics and best practices for migrating complex on-premises environments to AWS, emphasizing the use of AWS services and tools for effective migration planning and execution.

Question 12

Question Type: MultipleChoice

A company wants to migrate an Amazon Aurora MySQL DB cluster from an existing AWS account to a new AWS account in the same AWS Region. Both accounts are members of the same organization in AWS Organizations.

The company must minimize database service interruption before the company performs DNS cutover to the new database.

Which migration strategy will meet this requirement?

Options:

- A-** Take a snapshot of the existing Aurora database. Share the snapshot with the new AWS account. Create an Aurora DB cluster in the new account from the snapshot.
- B-** Create an Aurora DB cluster in the new AWS account. Use AWS Database Migration Service (AWS DMS) to migrate data between the two Aurora DB clusters.
- C-** Use AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account. Create an Aurora DB cluster in the new AWS account from the snapshot.
- D-** Create an Aurora DB cluster in the new AWS account. Use AWS Application Migration Service to migrate data between the two Aurora DB clusters.

Answer:

B

Explanation:

The best migration strategy to meet the requirement of minimizing database service interruption before the DNS cutover is to use AWS DMS to migrate data between the two Aurora DB clusters. AWS DMS can perform continuous replication of data with high availability and consolidate databases into a petabyte-scale data warehouse by streaming data to Amazon Redshift and Amazon S3¹. AWS DMS supports homogeneous migrations such as migrating from one Aurora MySQL DB cluster to another, as well as heterogeneous migrations between different database platforms². AWS DMS also supports cross-account migrations, as long as the source and target databases are in the same AWS Region³.

The other options are not optimal for the following reasons:

Option A: Taking a snapshot of the existing Aurora database and restoring it in the new account would require a downtime during the snapshot and restore process, which could be significant for large databases. Moreover, any changes made to the source database after the snapshot would not be replicated to the target database, resulting in data inconsistency⁴.

Option C: Using AWS Backup to share an Aurora database backup from the existing AWS account to the new AWS account would have the same drawbacks as option A, as AWS Backup uses snapshots to create backups of Aurora databases.

Option D: Using AWS Application Migration Service to migrate data between the two Aurora DB clusters is not a valid option, as AWS Application Migration Service is designed to migrate applications, not databases, to AWS. AWS Application Migration Service can migrate applications from on-premises or other cloud environments to AWS, using agentless or agent-based methods.

1:What Is AWS Database Migration Service? - AWS Database Migration Service

2:Sources for Data Migration - AWS Database Migration Service

3:AWS Database Migration Service FAQs

4:Working with DB Cluster Snapshots - Amazon Aurora

: [Backing Up and Restoring an Amazon Aurora DB Cluster - Amazon Aurora]

: [What is AWS Application Migration Service? - AWS Application Migration Service]

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>

