



Free Questions for SAA-C03 by certsdeals

Shared by Cole on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A social media company has workloads that collect and process data. The workloads store the data in on-premises NFS storage. The data store cannot scale fast enough to meet the company's expanding business needs. The company wants to migrate the current data store to AWS.

Which solution will meet these requirements MOST cost-effectively?

Options:

- A-** Set up an AWS Storage Gateway Volume Gateway. Use an Amazon S3 Lifecycle policy to transition the data to the appropriate storage class.
- B-** Set up an AWS Storage Gateway Amazon S3 File Gateway. Use an Amazon S3 Lifecycle policy to transition the data to the appropriate storage class.
- C-** Use the Amazon Elastic File System (Amazon EFS) Standard-Infrequent Access (Standard-IA) storage class. Activate the infrequent access lifecycle policy.
- D-** Use the Amazon Elastic File System (Amazon EFS) One Zone-Infrequent Access (One Zone-IA) storage class. Activate the infrequent access lifecycle policy.

Answer:

B

Explanation:

This solution meets the requirements most cost-effectively because it enables the company to migrate its on-premises NFS data store to AWS without changing the existing applications or workflows. AWS Storage Gateway is a hybrid cloud storage service that provides seamless and secure integration between on-premises and AWS storage. Amazon S3 File Gateway is a type of AWS Storage Gateway that provides a file interface to Amazon S3, with local caching for low-latency access. By setting up an Amazon S3 File Gateway, the company can store and retrieve files as objects in Amazon S3 using standard file protocols such as NFS. The company can also use an Amazon S3 Lifecycle policy to automatically transition the data to the appropriate storage class based on the frequency of access and the cost of storage. For example, the company can use S3 Standard for frequently accessed data, S3 Standard-Infrequent Access (S3 Standard-IA) or S3 One Zone-Infrequent Access (S3 One Zone-IA) for less frequently accessed data, and S3 Glacier or S3 Glacier Deep Archive for long-term archival data.

Option A is not a valid solution because AWS Storage Gateway Volume Gateway is a type of AWS Storage Gateway that provides a block interface to Amazon S3, with local caching for low-latency access. Volume Gateway is not suitable for migrating an NFS data store, as it requires attaching the volumes to EC2 instances or on-premises servers using the iSCSI protocol. Option C is not a valid solution because Amazon Elastic File System (Amazon EFS) is a fully managed elastic NFS file system that is designed for workloads that require high availability, scalability, and performance. Amazon EFS Standard-Infrequent Access (Standard-IA) is a storage class within Amazon EFS that is optimized for infrequently accessed files, with a lower price per GB and a higher price per access. Using Amazon EFS Standard-IA for migrating an NFS data store would not be cost-effective, as it would incur higher access charges and require additional configuration to enable lifecycle management. Option D is not a valid solution because Amazon EFS One Zone-Infrequent Access (One Zone-IA) is a storage class within Amazon EFS that is optimized for infrequently accessed files that do not require the availability and durability of Amazon EFS Standard or Standard-IA. Amazon EFS One Zone-IA stores data in a single

Availability Zone, which reduces the cost by 47% compared to Amazon EFS Standard-IA, but also increases the risk of data loss in the event of an Availability Zone failure. Using Amazon EFS One Zone-IA for migrating an NFS data store would not be cost-effective, as it would incur higher access charges and require additional configuration to enable lifecycle management. It would also compromise the availability and durability of the data.

[AWS Storage Gateway - Amazon Web Services](#)

[Amazon S3 File Gateway - AWS Storage Gateway](#)

[Object Lifecycle Management - Amazon Simple Storage Service](#)

[AWS Storage Gateway Volume Gateway - AWS Storage Gateway]

[Amazon Elastic File System - Amazon Web Services]

[Using EFS storage classes - Amazon Elastic File System]

Question 2

Question Type: MultipleChoice

A company runs containers in a Kubernetes environment in the company's local data center. The company wants to use Amazon Elastic Kubernetes Service (Amazon EKS) and other AWS managed services. Data must remain locally in the company's data center and cannot be stored in any remote site or cloud to maintain compliance.

Which solution will meet these requirements?

Options:

- A- Deploy AWS Local Zones in the company's data center
- B- Use an AWS Snowmobile in the company's data center
- C- Install an AWS Outposts rack in the company's data center
- D- Install an AWS Snowball Edge Storage Optimized node in the data center

Answer:

C

Explanation:

AWS Outposts is a fully managed service that delivers AWS infrastructure and services to virtually any on-premises or edge location for a consistent hybrid experience. AWS Outposts supports Amazon EKS, which is a managed service that makes it easy to run Kubernetes on AWS and on-premises. By installing an AWS Outposts rack in the company's data center, the company can run containers in a Kubernetes environment using Amazon EKS and other AWS managed services, while keeping the data locally in the company's data center and meeting the compliance requirements. AWS Outposts also provides a seamless connection to the local AWS Region for access to a broad range of AWS services.

Option A is not a valid solution because AWS Local Zones are not deployed in the company's data center, but in large metropolitan areas closer to end users. AWS Local Zones are owned, managed, and operated by AWS, and they provide low-latency access to the public internet and the local AWS Region. Option B is not a valid solution because AWS Snowmobile is a service that transports exabytes of data to AWS using a 45-foot long ruggedized shipping container pulled by a semi-trailer truck. AWS Snowmobile is not designed for running containers or AWS managed services on-premises, but for large-scale data migration. Option D is not a valid solution because AWS Snowball Edge Storage Optimized is a device that provides 80 TB of HDD or 210 TB of NVMe storage capacity for data transfer and edge computing. AWS Snowball Edge Storage Optimized does not support Amazon EKS or other AWS managed services, and it is not suitable for running containers in a Kubernetes environment.

[AWS Outposts - Amazon Web Services](#)

[Amazon EKS on AWS Outposts - Amazon EKS](#)

[AWS Local Zones - Amazon Web Services](#)

[AWS Snowmobile - Amazon Web Services](#)

[[AWS Snowball Edge Storage Optimized - Amazon Web Services](#)]

Question 3

Question Type: MultipleChoice

A company has an Amazon S3 data lake. The company needs a solution that transforms the data from the data lake and loads the data into a data warehouse every day. The data warehouse must have massively parallel processing (MPP) capabilities.

Data analysts then need to create and train machine learning (ML) models by using SQL commands on the data. The solution must use serverless AWS services wherever possible.

Which solution will meet these requirements?

Options:

- A-** Run a daily Amazon EMR job to transform the data and load the data into Amazon Redshift. Use Amazon Redshift ML to create and train the ML models.
- B-** Run a daily Amazon EMR job to transform the data and load the data into Amazon Aurora Serverless. Use Amazon Aurora ML to create and train the ML models.
- C-** Run a daily AWS Glue job to transform the data and load the data into Amazon Redshift Serverless. Use Amazon Redshift ML to create and train the ML models.
- D-** Run a daily AWS Glue job to transform the data and load the data into Amazon Athena tables. Use Amazon Athena ML to create and train the ML models.

Answer:

C

Explanation:

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. AWS Glue can automatically discover your data in Amazon S3 and catalog it, so you can query and search the data using SQL. AWS Glue can also run serverless ETL jobs using Apache Spark and Python to transform and load your data into various destinations, such as Amazon Redshift, Amazon Athena, or Amazon Aurora. AWS Glue is a serverless service, so you only pay for the resources consumed by the jobs, and you don't need to provision or manage any infrastructure.

Amazon Redshift is a fully managed, petabyte-scale data warehouse service that enables you to use standard SQL and your existing business intelligence (BI) tools to analyze your data. Amazon Redshift also supports massively parallel processing (MPP), which means it can distribute and execute queries across multiple nodes in parallel, delivering fast performance and scalability. Amazon Redshift Serverless is a new option that automatically scales query compute capacity based on the queries being run, so you don't need to manage clusters or capacity. You only pay for the query processing time and the storage consumed by your data.

Amazon Redshift ML is a feature that enables you to create, train, and deploy machine learning (ML) models using familiar SQL commands. Amazon Redshift ML can automatically discover the best model and hyperparameters for your data, and store the model in Amazon SageMaker, a fully managed service that provides a comprehensive set of tools for building, training, and deploying ML models. You can then use SQL functions to apply the model to your data in Amazon Redshift and generate predictions.

The combination of AWS Glue, Amazon Redshift Serverless, and Amazon Redshift ML meets the requirements of the question, as it provides a serverless, scalable, and SQL-based solution to transform, load, and analyze the data from the Amazon S3 data lake, and to create and train ML models on the data.

Option A is not correct, because Amazon EMR is not a serverless service. Amazon EMR is a managed service that simplifies running Apache Spark, Apache Hadoop, and other big data frameworks on AWS. Amazon EMR requires you to launch and configure clusters of EC2 instances to run your ETL jobs, which adds complexity and cost compared to AWS Glue.

Option B is not correct, because Amazon Aurora Serverless is not a data warehouse service, and it does not support MPP. Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora, a relational database service that is compatible with MySQL and PostgreSQL. Amazon Aurora Serverless can automatically adjust the database capacity based on the traffic, but it does not distribute the data and queries across multiple nodes like Amazon Redshift does. Amazon Aurora Serverless is more suitable for transactional workloads than analytical workloads.

Option D is not correct, because Amazon Athena is not a data warehouse service, and it does not support MPP. Amazon Athena is an interactive query service that enables you to analyze data in Amazon S3 using standard SQL. Amazon Athena is serverless, so you only pay for the queries you run, and you don't need to load the data into a database. However, Amazon Athena does not store the data in a columnar format, compress the data, or optimize the query execution plan like Amazon Redshift does. Amazon Athena is more suitable for ad-hoc queries than complex analytics and ML.

[AWS Glue](#)

[Amazon Redshift](#)

[Amazon Redshift Serverless](#)

[Amazon Redshift ML](#)

[Amazon EMR](#)

[Amazon Aurora Serverless](#)

[Amazon Athena](#)

Question 4

Question Type: MultipleChoice

A company has a mobile app for customers. The app's data is sensitive and must be encrypted at rest. The company uses AWS Key Management Service (AWS KMS).

The company needs a solution that prevents the accidental deletion of KMS keys. The solution must use Amazon Simple Notification Service (Amazon SNS) to send an email notification to administrators when a user attempts to delete a KMS key.

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

Options:

- A-** Create an Amazon EventBridge rule that reacts when a user tries to delete a KMS key. Configure an AWS Config rule that cancels any deletion of a KMS key. Add the AWS Config rule as a target of the EventBridge rule. Create an SNS topic that notifies the administrators.
- B-** Create an AWS Lambda function that has custom logic to prevent KMS key deletion. Create an Amazon CloudWatch alarm that is activated when a user tries to delete a KMS key. Create an Amazon EventBridge rule that invokes the Lambda function when the DeleteKey operation is performed. Create an SNS topic. Configure the EventBridge rule to publish an SNS message that notifies the administrators.
- C-** Create an Amazon EventBridge rule that reacts when the KMS DeleteKey operation is performed. Configure the rule to initiate an AWS Systems Manager Automation.

runbook Configure the runbook to cancel the deletion of the KMS key Create an SNS topic Configure the EventBndge rule to publish an SNS message that notifies the administrators.

D- Create an AWS CloudTrail trail Configure the trail to delrver logs to a new Amazon CloudWatch log group Create a CloudWatch alarm based on the metric filter for the CloudWatch log group Configure the alarm to use Amazon SNS to notify the administrators when the KMS DeleteKey operation is performed

Answer:

C

Explanation:

This solution meets the requirements with the least operational overhead because it uses AWS services that are fully managed and scalable. The EventBridge rule can detect the DeleteKey operation from the AWS KMS API and trigger the Systems Manager Automation runbook, which can execute a predefined workflow to cancel the key deletion. The EventBridge rule can also publish an SNS message to the topic that sends an email notification to the administrators. This way, the company can prevent the accidental deletion of KMS keys and notify the administrators of any attempts to delete them.

Option A is not a valid solution because AWS Config rules are used to evaluate the configuration of AWS resources, not to cancel the deletion of KMS keys. Option B is not a valid solution because it requires creating and maintaining a custom Lambda function that has logic to prevent KMS key deletion, which adds operational overhead. Option D is not a valid solution because it only notifies the administrators of the DeleteKey operation, but does not cancel it.

[Using Amazon EventBridge rules to trigger Systems Manager Automation workflows - AWS Systems Manager](#)

Using Amazon SNS for system-to-administrator communications - Amazon Simple Notification Service

Deleting AWS KMS keys - AWS Key Management Service

Question 5

Question Type: MultipleChoice

A company needs a solution to prevent AWS CloudFormation stacks from deploying AWS Identity and Access Management (IAM) resources that include an inline policy or "*" in the statement. The solution must also prohibit deployment of Amazon EC2 instances with public IP addresses. The company has AWS Control Tower enabled in its organization in AWS Organizations.

Which solution will meet these requirements?

Options:

- A-** Use AWS Control Tower proactive controls to block deployment of EC2 instances with public IP addresses and inline policies with elevated access or '*'
- B-** Use AWS Control Tower detective controls to block deployment of EC2 instances with public IP addresses and inline policies with elevated access or ''
- C-** Use AWS Config to create rules for EC2 and IAM compliance. Configure the rules to run an AWS Systems Manager Session Manager automation to delete a resource when it is not compliant.

D- Use a service control policy (SCP) to block actions for the EC2 instances and IAM resources if the actions lead to noncompliance

Answer:

D

Explanation:

A service control policy (SCP) is a type of policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization, allowing you to ensure your accounts stay within your organization's access control guidelines. SCPs are available only in an organization that has all features enabled. SCPs do not grant permissions; instead, they specify the maximum permissions for an organization or organizational unit (OU). SCPs limit permissions that identity-based policies or resource-based policies grant to entities (users or roles) within the account, but do not grant permissions to entities. You can use SCPs to restrict the actions that the root user in an account can perform. You can also use SCPs to prevent users or roles in any account from creating or modifying certain AWS resources, such as EC2 instances with public IP addresses or IAM resources with inline policies or `inline`. For example, you can create an SCP that denies the `ec2:RunInstances` action if the request includes the `AssociatePublicIpAddress` parameter set to `true`. You can also create an SCP that denies the `iam:PutUserPolicy` and `iam:PutRolePolicy` actions if the request includes a policy document that contains `inline`. By attaching these SCPs to your organization or OUs, you can prevent the deployment of AWS CloudFormation stacks that violate these rules.

AWS Control Tower proactive controls are guardrails that enforce preventive policies on your accounts and resources. Proactive guardrails are implemented as AWS Organizations service control policies (SCPs) and AWS Config rules. However, AWS Control Tower does not provide a built-in proactive guardrail to block EC2 instances with public IP addresses or IAM resources with inline policies or `inline`. You would have to create your own custom guardrails using AWS CloudFormation templates and SCPs, which is essentially the

same as option D. Therefore, option A is not correct.

AWS Control Tower detective controls are guardrails that detect and alert on policy violations in your accounts and resources. Detective guardrails are implemented as AWS Config rules and Amazon CloudWatch alarms. Detective guardrails do not block or remediate noncompliant resources; they only notify you of the issues. Therefore, option B is not correct.

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. AWS Config rules are customizable, AWS Lambda functions that AWS Config invokes to evaluate your AWS resource configurations. You can use AWS Config rules to check for compliance with your policies, such as ensuring that EC2 instances have public IP addresses disabled or IAM resources do not have inline policies or "*" . However, AWS Config rules alone cannot prevent the deployment of AWS CloudFormation stacks that violate these policies; they can only report the compliance status. You would need to use another service, such as AWS Systems Manager Session Manager, to run automation scripts to delete or modify the noncompliant resources. This would require additional configuration and permissions, and may not be the most efficient or secure way to enforce your policies. Therefore, option C is not correct.

[Service Control Policies](#)

[AWS Control Tower Guardrails](#)

[AWS Config](#)

Question 6

Question Type: MultipleChoice

A company wants to use NAT gateways in its AWS environment. The company's Amazon EC2 instances in private subnets must be able to connect to the public internet through the NAT gateways.

Which solution will meet these requirements'?

Options:

- A-** Create public NAT gateways in the same private subnets as the EC2 instances
- B-** Create private NAT gateways in the same private subnets as the EC2 instances
- C-** Create public NAT gateways in public subnets in the same VPCs as the EC2 instances
- D-** Create private NAT gateways in public subnets in the same VPCs as the EC2 instances

Answer:

C

Explanation:

A public NAT gateway enables instances in a private subnet to send outbound traffic to the internet, while preventing the internet from initiating connections with the instances. A public NAT gateway requires an elastic IP address and a route to the internet gateway for the

VPC. A private NAT gateway enables instances in a private subnet to connect to other VPCs or on-premises networks through a transit gateway or a virtual private gateway. A private NAT gateway does not require an elastic IP address or an internet gateway. Both private and public NAT gateways map the source private IPv4 address of the instances to the private IPv4 address of the NAT gateway, but in the case of a public NAT gateway, the internet gateway then maps the private IPv4 address of the public NAT gateway to the elastic IP address associated with the NAT gateway. When sending response traffic to the instances, whether it's a public or private NAT gateway, the NAT gateway translates the address back to the original source IP address.

Creating public NAT gateways in the same private subnets as the EC2 instances (option A) is not a valid solution, as the NAT gateways would not have a route to the internet gateway. Creating private NAT gateways in the same private subnets as the EC2 instances (option B) is also not a valid solution, as the instances would not be able to access the internet through the private NAT gateways. Creating private NAT gateways in public subnets in the same VPCs as the EC2 instances (option D) is not a valid solution either, as the internet gateway would drop the traffic from the private NAT gateways.

Therefore, the only valid solution is to create public NAT gateways in public subnets in the same VPCs as the EC2 instances (option C), as this would allow the instances to access the internet through the public NAT gateways and the internet gateway. Reference:

[NAT gateways - Amazon Virtual Private Cloud](#)

[NAT gateway use cases - Amazon Virtual Private Cloud](#)

[Amazon Web Services -- Introduction to NAT Gateways](#)

[What is AWS NAT Gateway? - KnowledgeHut](#)

To Get Premium Files for SAA-C03 Visit

<https://www.p2pexams.com/products/saa-c03>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/saa-c03>

