



CertsMania

Free Questions for **SAA-C03**

Shared by **Reagan** on **Jan 2, 2026**

For More Free Questions and Preparation Resources

Check the Links on Last Page



CertsMania

Questions # 1:

A company is building a data processing application that uses AWS Lambda functions. The Lambda functions need to communicate with an Amazon RDS DB instance deployed inside a VPC in the same AWS account.

Which solution meets these requirements in the most secure way?

Options:

A.

Configure the DB instance for public access. Allow Lambda public address space.

B.

Deploy Lambda inside the VPC. Attach a network ACL allowing outbound access to the VPC CIDR. Update the DB security group to allow traffic from 0.0.0.0/0.

C.

Deploy Lambda inside the VPC. Attach a security group to the Lambda functions. Allow outbound access only to the VPC CIDR. Update the DB instance security group to allow traffic from the Lambda security group.

D.

Peer the Lambda default VPC with the DB VPC and avoid security groups.

Answer

C

Explanation

For secure communication between Lambda and an RDS DB instance, AWS documentation recommends placing the Lambda functions inside the same VPC and controlling traffic using security groups.

Lambda should have a dedicated security group with outbound access to the VPC CIDR range, and the DB instance's security group should explicitly allow inbound connections from the Lambda security group. This follows the "least privilege" principle while avoiding public exposure.

Options A and B expose the DB to unnecessary risk. Option D is insecure because it bypasses security groups.

Questions # 2:

A company is developing an application in the AWS Cloud. The application's HTTP API contains critical information that is published in Amazon API Gateway. The critical information must be accessible from only a limited set of trusted IP addresses that belong to the company's internal network.

Which solution will meet these requirements?

Options:

A.

Set up an API Gateway private integration to restrict access to a predefined set of IP addresses.

B.

Create a resource policy for the API that denies access to any IP address that is not specifically allowed.

C.

Directly deploy the API in a private subnet. Create a network ACL. Set up rules to allow the traffic from specific IP addresses.

D.

Modify the security group that is attached to API Gateway to allow inbound traffic from only the trusted IP addresses.

Answer

B

Explanation

Amazon API Gateway supports resource policies, which allow you to control access to your API by specifying the IP addresses or ranges that can access the API. By creating a resource policy that explicitly denies access to any IP address outside the allowed set, you can ensure that only trusted IP addresses (such as those from your internal network) can access the critical information in your API. This approach provides fine-grained access control without the need for additional infrastructure or complex configurations.


Option A (Private integration): API Gateway private integrations are for creating private APIs that are only accessible within a VPC, but this solution is about restricting access to certain IP addresses.

Option C (Private subnet and ACLs): Deploying the API in a private subnet and using network ACLs adds unnecessary complexity and isn't the best fit for HTTP APIs.

Option D (Security group): API Gateway doesn't have a security group because it isn't a resource inside a VPC. Instead, resource policies are the correct mechanism for controlling IP-based access.

AWS References:

Controlling Access to API Gateway with Resource Policies



Questions # 3:

A company wants to design a microservices architecture for an application. Each microservice must perform operations that can be completed within 30 seconds.

The microservices need to expose RESTful APIs and must automatically scale in response to varying loads. The APIs must also provide client access control and rate limiting to maintain equitable usage and service availability.

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

Options:

A.

Use Amazon Elastic Container Service (Amazon ECS) on Amazon EC2 to host each microservice. Use Amazon API Gateway to manage the RESTful API requests.

B.

Deploy each microservice as a set of AWS Lambda functions. Use Amazon API Gateway to manage the RESTful API requests.

C.

Host each microservice on Amazon EC2 instances in Auto Scaling groups behind an Elastic Load Balancing (ELB) load balancer. Use the ELB to manage the RESTful API requests.

D.

Deploy each microservice on Amazon Elastic Beanstalk. Use Amazon CloudFront to manage the RESTful API requests.



Answer

C

Questions # 4:

A company is running a media store across multiple Amazon EC2 instances distributed across multiple Availability Zones in a single VPC. The company wants a high-performing solution to share data between all the EC2 instances, and prefers to keep the data within the VPC only.

What should a solutions architect recommend?

Options:

A.

Create an Amazon S3 bucket and call the service APIs from each instance's application.

B.

Create an Amazon S3 bucket and configure all instances to access it as a mounted volume.

C.

Configure an Amazon Elastic Block Store (Amazon EBS) volume and mount it across all instances.

D.

Configure an Amazon Elastic File System (Amazon EFS) file system and mount it across all instances.

Answer

D

Explanation

Amazon Elastic File System (EFS) is a managed file storage service that can be mounted across multiple EC2 instances. It provides a scalable and high-performing solution to share data among instances within a VPC.

High Performance: EFS provides scalable performance for workloads that require high throughput and IOPS. It is particularly well-suited for applications that need to share data across multiple instances.

Ease of Use: EFS can be easily mounted on multiple instances across different Availability

Zones, providing a shared file system accessible to all the instances within the VPC.

Security: EFS can be configured to ensure that data remains within the VPC, and it supports encryption at rest and in transit.

Why Not Other Options?:

Option A (Amazon S3 bucket with APIs): While S3 is excellent for object storage, it is not a file system and does not provide the low-latency access required for shared data between instances.

Option B (S3 bucket as a mounted volume): S3 is not designed to be mounted as a file system, and this approach would introduce unnecessary complexity and latency.

Option C (EBS volume shared across instances): EBS volumes cannot be attached to multiple instances simultaneously. It is not designed to be shared across instances like EFS.

AWS References:

Amazon EFS- Overview of Amazon EFS and its features.

Best Practices for Amazon EFS- Recommendations for using EFS with multiple instances.

Questions # 5:

A company stores data for multiple business units in a single Amazon S3 bucket that is in the company's payer AWS account. To maintain data isolation, the business units store data in separate prefixes in the S3 bucket by using an S3 bucket policy.

The company plans to add a large number of dynamic prefixes. The company does not want to rely on a single S3 bucket policy to manage data access at scale. The company wants to develop a secure access management solution in addition to the bucket policy to enforce prefix-level data isolation.

Options:

A.

Configure the S3 bucket policy to deny `s3:GetObject` permissions for all users. Configure the bucket policy to allow `s3:*` access to individual business units.

B.

Enable default encryption on the S3 bucket by using server-side encryption with Amazon S3 managed keys (SSE-S3).

C.

Configure resource-based permissions on the S3 bucket by creating an S3 access point for each business unit.

D.

Use pre-signed URLs to provide access to the S3 bucket.

Answer

C

Explanation



CertsMania

Why Option C is Correct:

S3 Access Points: Provide scalable management of access to large datasets with specific permissions for individual prefixes.

Dynamic Prefixes: Access points simplify managing access to a growing number of prefixes without relying solely on a single bucket policy.

Fine-Grained Control: Resource-based permissions on access points enforce prefix-level isolation effectively.

Why Other Options Are Not Ideal:

Option A: Using deny/allow bucket policies introduces complexity and is less scalable for dynamic prefixes.

Option B: Encryption ensures data security but does not address access management.

Option D: Pre-signed URLs are temporary and not suitable for managing access at scale.

AWS References:

Amazon S3 Access Points: [AWS Documentation - S3 Access Points](#)



CertsMania

Questions # 6:

A company needs to migrate a MySQL database from an on-premises data center to AWS within 2 weeks. The database is 180 TB in size. The company cannot partition the database.

The company wants to minimize downtime during the migration. The company's internet connection speed is 100 Mbps.

Which solution will meet these requirements?

Options:

A.

Order an AWS Snowball Edge Storage Optimized device. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes. Send the Snowball Edge device back to AWS to finish the migration. Continue to replicate ongoing changes.

B.

Establish an AWS Site-to-Site VPN connection between the data center and AWS. Use AWS Database Migration Service (AWS DMS) and the AWS Schema Conversion Tool (AWS SCT) to migrate the database to Amazon RDS for MySQL and replicate ongoing changes.

C.

Establish a 10 Gbps dedicated AWS Direct Connect connection between the data center and AWS. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

D.

Use the company's existing internet connection. Use AWS DataSync to replicate the database to Amazon S3. Create a script to import the data from Amazon S3 to a new Amazon RDS for MySQL database instance.

Answer

A

Explanation

Given the large size (180 TB) of the database and the time constraint, AWS Snowball Edge Storage Optimized is the best solution. Snowball Edge allows for the physical transfer of large datasets to AWS efficiently without relying on slow internet connections. AWS DMS and SCT can be used to perform ongoing replication of any changes made during the migration, ensuring minimal downtime.

Option B (VPN): Using a 100 Mbps internet connection would take far too long to transfer 180 TB.

Option C (Direct Connect): Establishing a 10 Gbps Direct Connect link might not be feasible within the 2-week timeframe.

Option D (DataSync over internet): With the existing internet connection, DataSync would also take too long.

AWS References:

AWS Snowball Edge

AWS DMS

Questions # 7:

A solutions architect is designing a three-tier web application. The architecture consists of an internet-facing Application Load Balancer (ALB) and a web tier that is hosted on Amazon EC2 instances in private subnets. The application tier with the business logic runs on EC2 instances in private subnets. The database tier consists of Microsoft SQL Server that runs on EC2 instances in private subnets. Security is a high priority for the company. Which combination of security group configurations should the solutions architect use? (Select THREE.)

Options:

A.

Configure the security group for the web tier to allow inbound HTTPS traffic from the security group for the ALB.

B.

Configure the security group for the web tier to allow outbound HTTPS traffic to 0.0.0.0/0.

C.

Configure the security group for the database tier to allow inbound Microsoft SQL Server traffic from the security group for the application tier.

D.

Configure the security group for the database tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.

E.

Configure the security group for the application tier to allow inbound HTTPS traffic from the security group for the web tier.

F.

Configure the security group for the application tier to allow outbound HTTPS traffic and Microsoft SQL Server traffic to the security group for the web tier.

Answer

A, C, E

Explanation

According to AWS best practices, each tier's security group must restrict inbound traffic to only the upstream trusted source. For the web tier, the Application Load Balancer must be the only entity allowed to send traffic. AWS documentation specifies: "Restrict the security groups associated with your targets to accept traffic only from the load balancer." This confirms that the web tier security group should allow inbound HTTPS from the ALB security group (A).

For communication between the web and application tiers, AWS states: "You can specify a security group as the source or destination in a rule" and "Create rules only for the protocols and ports required by your application." Therefore, the application tier security group must allow inbound HTTPS traffic from the web tier security group (E).

For the database tier, AWS guidance says: "Allow only the necessary ports for database communication." Microsoft SQL Server listens on port 1433 by default, so the database tier security group must allow inbound SQL Server traffic from the application tier security group (C).

Outbound rules (options B, D, and F) are unnecessary because AWS specifies that "Security groups are stateful. Return traffic is automatically allowed." This means once inbound rules are defined, the return path is automatically permitted without extra outbound configurations.

This combination (A, C, E) applies the principle of least privilege, ensures end-to-end secure communication across tiers, and follows AWS recommendations for ALB-to-target security group setups.

[References: • Elastic Load Balancing User Guide — Application Load Balancers: Security groups for your load balancer, Target security groups • Amazon VPC User Guide: Security groups for your VPC, Security group rules • AWS Well-Architected Framework — Security Pillar: Apply the principle of least privilege, , ,]

Questions # 8:

A company is planning to deploy its application on an Amazon Aurora PostgreSQL Serverless v2 cluster. The application will receive large amounts of traffic. The company wants to optimize the storage performance of the cluster as the load on the application increases

Which solution will meet these requirements MOST cost-effectively?

Options:

A.

Configure the cluster to use the Aurora Standard storage configuration.

B.

Configure the cluster storage type as Provisioned IOPS.

C.

Configure the cluster storage type as General Purpose.

D.

Configure the cluster to use the Aurora I/O-Optimized storage configuration.

Answer

D

Explanation

Aurora I/O-Optimized: This storage configuration is designed to provide consistent high performance for Aurora databases. It automatically scales IOPS as the workload increases, without needing to provision IOPS separately.

Cost-Effectiveness: With Aurora I/O-Optimized, you only pay for the storage and I/O you use, making it a cost-effective solution for applications with varying and unpredictable I/O demands.

Implementation:

During the creation of the Aurora PostgreSQL Serverless v2 cluster, select the I/O-Optimized storage configuration.

The storage system will automatically handle scaling and performance optimization based on the application load.

Operational Efficiency: This configuration reduces the need for manual tuning and ensures optimal performance without additional administrative overhead.

[References:, Amazon Aurora I/O-Optimized, , , ,]

Questions # 9:

An insurance company is creating an application to record personal user data. The data includes users' names, ages, and health data. The company wants to run the application in a private subnet on AWS.

Because of data security requirements, the company must have access to the operating

system of the compute resources that run the application tier. The company must use a low-latency NoSQL database to store the data.

Which solution will meet these requirements?

Options:

A.

Use Amazon EC2 instances for the application tier. Use an Amazon DynamoDB table for the database tier. Create a VPC endpoint for DynamoDB. Assign the instances an instance profile that has permission to access DynamoDB.

B.

Use AWS Lambda functions for the application tier. Use an Amazon DynamoDB table for the database tier. Assign a Lambda function an appropriate IAM role to access the table.

C.

Use AWS Fargate for the application tier. Create an Amazon Aurora PostgreSQL instance inside a private subnet for the database tier.

D.

Use Amazon EC2 instances for the application tier. Use an Amazon S3 bucket to store the data in JSON format. Configure the application to use Amazon Athena to read and write the data to and from the S3 bucket.

Answer

A

Explanation

The requirement to “have access to the operating system” means the compute layer must be Amazon EC2 (or containers on EC2). Managed runtimes such as Lambda and Fargate do not provide OS-level access.

The requirement for a “low-latency NoSQL database” maps directly to Amazon DynamoDB, which is a fully managed NoSQL key-value and document database that provides single-digit millisecond latency at any scale.

Because the application runs in a private subnet, AWS best practice is to access DynamoDB privately via a VPC endpoint (gateway endpoint for DynamoDB). This avoids traversing the public internet and simplifies security.

An instance profile (EC2 role) is the recommended method to grant EC2 instances permission to access DynamoDB without hardcoding credentials.

Why the other options are not correct:

B: Lambda does not provide OS access, which violates the security requirement.

C: Fargate does not provide OS access, and Aurora PostgreSQL is a relational database, not NoSQL.

D: S3 + Athena is an analytics pattern, not a low-latency NoSQL database solution; query latency is much higher and not suitable for OLTP-style app storage.

Questions # 10:

A company is designing a microservice-based architecture for a new application on AWS. Each microservice will run on its own set of Amazon EC2 instances. Each microservice will need to interact with multiple AWS services such as Amazon S3 and Amazon Simple Queue Service (Amazon SQS).

The company wants to manage permissions for each EC2 instance based on the principle of least privilege.

Which solution will meet this requirement?

Options:

A.

Assign an IAM user to each micro-service. Use access keys stored within the application code to authenticate AWS service requests.

B.

Create a single IAM role that has permission to access all AWS services. Associate the IAM role with all EC2 instances that run the microservices.

C.

Use AWS Organizations to create a separate account for each microservice. Manage permissions at the account level.

D.

Create individual IAM roles based on the specific needs of each microservice. Associate the IAM roles with the appropriate EC2 instances.

Answer

D

Explanation

When designing a microservice architecture where each microservice interacts with different AWS services, it's essential to follow the principle of least privilege. This means granting each microservice only the permissions it needs to perform its tasks, reducing the risk of unauthorized access or accidental actions.

The recommended approach is to create individual IAM roles with policies that grant each microservice the specific permissions it requires. Then, these roles should be associated with the EC2 instances that run the corresponding microservice. By doing so, each EC2 instance will assume its specific IAM role, and permissions will be automatically managed by AWS.

IAM roles provide temporary credentials via the instance metadata service, eliminating the need to hard-code credentials in your application code, which enhances security.

AWS References:

[IAM Roles for Amazon EC2](#) explains how EC2 instances can use IAM roles to securely access AWS services without managing long-term credentials.

[Best Practices for IAM](#) includes recommendations for implementing the least privilege principle and using IAM roles effectively.

Why the other options are incorrect:

A. Assign an IAM user to each microservice: This requires managing long-term credentials (access keys), which should be avoided. Storing keys in application code is insecure and creates a maintenance burden.

B. Create a single IAM role: This violates the principle of least privilege, as a single role with broad permissions across all services is less secure.

C. Use AWS Organizations: This approach adds unnecessary complexity. Managing permissions at the account level for each microservice is excessive for this use case and doesn't adhere to the principle of least privilege.

Questions # 11:

A company is developing a platform to process large volumes of data for complex analytics and machine learning (ML) tasks. The platform must handle compute-intensive workloads. The workloads currently require 20 to 30 minutes for each data processing step.

The company wants a solution to accelerate data processing.

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

Options:

A.

Deploy three Amazon EC2 instances. Distribute the EC2 instances across three Availability Zones. Use traditional batch processing techniques for data processing.

B.

Create an Amazon EMR cluster. Use managed scaling. Install Apache Spark to assist with data processing.

C.

Create an AWS Lambda function for each data processing step. Deploy an Amazon Simple Queue Service (Amazon SQS) queue to relay data between Lambda functions.

D.

Create a series of AWS Lambda functions to process the data. Use AWS Step Functions to orchestrate the Lambda functions into data processing steps.

Answer

B

Explanation

Amazon EMR provides a managed big data framework that supports Apache Spark, which is ideal for distributed and compute-intensive data transformations. Managed scaling dynamically adjusts cluster resources, ensuring high performance with minimal management.

From AWS Documentation:

“Amazon EMR provides a managed environment for big data frameworks such as Apache Spark and Hadoop. With managed scaling, EMR automatically resizes clusters to meet workload demands.”

(Source: Amazon EMR Developer Guide)

Why B is correct:

Provides distributed parallel processing for large datasets.

Reduces operational overhead with managed scaling and auto-termination.

Integrates easily with S3, Glue, and ML pipelines.

Optimized for heavy ETL and analytics workloads.

Why others are incorrect:

A: Manual scaling and limited processing capacity.

C & D: Lambda has execution time and memory limits unsuitable for 30-minute compute-intensive tasks.

[References:, Amazon EMR Developer Guide - "Using Managed Scaling", AWS Well-Architected Framework - Performance Efficiency Pillar, , ,]

Questions # 12:

A company is planning to deploy a data processing platform on AWS. The data processing platform is based on PostgreSQL. The company stores the data that the platform must process on premises.

To comply with regulations, the company must not migrate the data to the cloud. However, the company wants to use AWS managed data analytics solutions.

Which solution will meet these requirements?

Options:

A.

Create an Amazon RDS for PostgreSQL database in a VPC. Create an interface VPC endpoint to connect the on-premises PostgreSQL database to the RDS for PostgreSQL database.

B.

Create Amazon EC2 instances in an Auto Scaling group on AWS Outposts. Install PostgreSQL data analytics software on the instances.

C.

Create an Amazon EMR cluster on AWS Outposts. Connect the EMR cluster to the on-premises PostgreSQL database to perform data processing locally.

D.

Create an Amazon EMR cluster in a VPC. Connect the EMR cluster to Amazon RDS for SQL Server with a linked server to connect to the company's data processing platform.

Answer

C

Explanation

AWS Outposts extends AWS infrastructure and services to on-premises locations. Running Amazon EMR on Outposts allows for processing data that resides locally while benefiting from the managed services of EMR. This enables compliance with data residency requirements and provides scalability and manageability for analytics.

[Reference: AWS Documentation - Amazon EMR on Outposts, =====, ,]

Questions # 13:

A company hosts an application on AWS. The application has generated approximately 2.5 TB of data over the previous 12 years. The company currently stores the data on Amazon EBS volumes.

The company wants a cost-effective backup solution for long-term storage. The company must be able to retrieve the data within minutes when required for audits.

Which solution will meet these requirements?

Options:

A.

Create EBS snapshots to back up the data.

B.

Create an Amazon S3 bucket. Use the S3 Glacier Deep Archive storage class to back up the data.

C.

Create an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class to back up the data.

D.

Create an Amazon Elastic File System (Amazon EFS) file system to back up the data.

Answer

C

Explanation

Amazon S3 Glacier Flexible Retrieval is a low-cost archival storage class that supports retrieval of data within minutes (expedited access ~1-5 minutes), making it ideal for audit scenarios where occasional, quick access to archived data is required. In contrast, Glacier Deep Archive takes hours to retrieve.

[Reference: AWS Documentation – Amazon S3 Glacier Flexible Retrieval,
=====, ,]

Questions # 14:

A company runs an ecommerce platform with a monolithic architecture on Amazon EC2 instances. The platform runs web and API services. The company wants to decouple the architecture and enhance scalability. The company also wants the ability to track orders and reprocess any failed orders.

Which solution will meet these requirements?

Options:

A.

Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Implement an SQS dead-letter queue.

B.

Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure Amazon Elastic Container Service (Amazon ECS) tasks to consume the queue. Implement SQS visibility timeout.

C.

Use Amazon Kinesis Data Streams to queue orders. Use AWS Lambda functions to consume the data stream. Configure Amazon S3 to track and reprocess failed orders.

D.

Send orders to an Amazon Simple Queue Service (Amazon SQS) queue. Configure AWS Lambda functions to consume the queue and process orders. Configure the Lambda functions to use SQS long polling.

Answer

A

Explanation

To decouple the monolith and enhance scalability, AWS best practice is to introduce an asynchronous message queue, such as Amazon SQS, between the web/API tier and the order-processing logic.

AWS Lambda functions consuming from the SQS queue provide serverless, auto-scaling processing without managing servers.

To track and reprocess failed orders, SQS supports dead-letter queues (DLQs). Messages that cannot be processed successfully after a configurable number of attempts are automatically moved to the DLQ, where operations teams or automated processes can inspect and reprocess them.

Why others are not correct:

B: ECS tasks can consume an SQS queue, but this requires managing container infrastructure and does not inherently provide as simple reprocessing/visibility as combining Lambda with a DLQ. Visibility timeout is not a tracking or archival mechanism.

C: Kinesis is a streaming service designed for ordered event streams, not primarily for order-queue semantics and DLQs; SQS is simpler and purpose-built for this pattern.

D: Long polling reduces empty responses and API calls but does nothing for tracking or reprocessing failed messages; without a DLQ, failed orders are harder to manage

Questions # 15:

A company is running a highly sensitive application on Amazon EC2 backed by an Amazon RDS database. Compliance regulations mandate that all personally identifiable information (PII) be encrypted at rest.

Which solution should a solutions architect recommend to meet this requirement with the LEAST amount of changes to the infrastructure?

Options:

A.

Deploy AWS Certificate Manager to generate certificates Use the certificates to encrypt the database volume

B.

Deploy AWS CloudHSM. generate encryption keys, and use the keys to encrypt database volumes.

C.

Configure SSL encryption using AWS Key Management Service (AWS KMS) keys to encrypt database volumes.

D.

Configure Amazon Elastic Block Store (Amazon EBS) encryption and Amazon RDS encryption with AWS Key Management Service (AWS KMS) keys to encrypt instance and database volumes.

Answer

D

Explanation

EBS Encryption:

Default EBS Encryption: Can be enabled for new EBS volumes.

Use of AWS KMS: Specify AWS KMS keys to handle encryption and decryption of data transparently.

Amazon RDS Encryption:

RDS Encryption: Encrypts the underlying storage for RDS instances using AWS KMS.

Configuration: Enable encryption when creating the RDS instance or modify an existing instance to enable encryption.

Least Amount of Changes:

Both EBS and RDS support seamless encryption with AWS KMS, requiring minimal changes to the existing infrastructure.

Enables compliance with regulatory requirements without modifying the application.

Operational Efficiency: Using AWS KMS for both EBS and RDS ensures a consistent,

managed approach to encryption, simplifying key management and enhancing security.

[References:, Amazon EBS Encryption, Amazon RDS Encryption, AWS Key Management Service, , , , ,]

Questions # 16:

A data science team requires storage for nightly log processing. The size and number of logs is unknown and the logs will persist for 24 hours only.

What is the MOST cost-effective solution?

Options:

A.

Amazon S3 Glacier Deep Archive

B.

Amazon S3 Standard

C.

Amazon S3 Intelligent-Tiering

D.

Amazon S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer

B

Explanation

For logs that are:

Written and processed within a short period (24 hours)

Accessed quickly for compute/analytics

With unknown object count and size

Amazon S3 Standard is the most appropriate and cost-effective. Intelligent-Tiering is designed for data stored for longer periods (typically 30+ days) with changing access

patterns and charges a per-object monitoring and automation fee that becomes inefficient for very short-lived objects.

S3 Glacier Deep Archive and S3 One Zone-IA are optimized for long-term archival or infrequently accessed data with retrieval time or availability constraints that are not suitable for nightly active processing.

Questions # 17:

A company recently migrated a data warehouse to AWS. The company has an AWS Direct Connect connection to AWS. Company users query the data warehouse by using a visualization tool. The average size of the queries that the data warehouse returns is 50 MB. The average visualization that the visualization tool produces is 500 KB in size. The result sets that the data warehouse returns are not cached.

The company wants to optimize costs for data transfers between the data warehouse and the company.

Which solution will meet this requirement?

Options:

A.

Host the visualization tool on premises. Connect to the data warehouse directly through the internet.

B.

Host the visualization tool in the same AWS Region as the data warehouse. Access the visualization tool through the internet.

C.

Host the visualization tool on premises. Connect to the data warehouse through the Direct Connect connection.

D.

Host the visualization tool in the same AWS Region as the data warehouse. Access the visualization tool through the Direct Connect connection.

Answer

D

Explanation

A. On-premises tool via internet: Incurs high costs due to large data transfers over the internet.

B. AWS Region tool via internet: Does not utilize Direct Connect, leading to potential latency and higher costs.

C. On-premises tool via Direct Connect: Adds latency for querying and visualization.

D. AWS Region tool via Direct Connect: Reduces latency and leverages Direct Connect for optimized data transfer costs.

[References: AWS Direct Connect, , ,]

Questions # 18:

A company has multiple Amazon RDS DB instances that run in a development AWS account. All the instances have tags to identify them as development resources. The company needs the development DB instances to run on a schedule only during business hours.

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

Options:

A.

Create an Amazon CloudWatch alarm to identify RDS instances that need to be stopped
Create an AWS Lambda function to start and stop the RDS instances.

B.

Create an AWS Trusted Advisor report to identify RDS instances to be started and stopped.
Create an AWS Lambda function to start and stop the RDS instances.

C.

Create AWS Systems Manager State Manager associations to start and stop the RDS instances.

D.

Create an Amazon EventBridge rule that invokes AWS Lambda functions to start and stop the RDS instances.

Answer

D

Explanation

To run RDS instances only during business hours with the least operational overhead, you can use Amazon EventBridge to schedule events that invoke AWS Lambda functions. The Lambda functions can be configured to start and stop the RDS instances based on the specified schedule (business hours). EventBridge rules allow you to define recurring events easily, and Lambda functions provide a serverless way to manage RDS instance start and stop operations, reducing administrative overhead.

Option A: While CloudWatch alarms could be used, they are more suited for monitoring, and using Lambda with EventBridge is simpler.

Option B (Trusted Advisor): Trusted Advisor is not ideal for scheduling tasks.

Option C (Systems Manager): Systems Manager could also work, but EventBridge and Lambda offer a more streamlined and lower-overhead solution.

AWS References:

Amazon EventBridge Scheduler

AWS Lambda

Questions # 19:

A company collects data from sensors. The company needs a cloud-based solution to store and transform the sensor data to make critical decisions. The solution must store the data for up to 2 days. After 2 days, the solution must delete the data. The company needs to use the transformed data in an automated workflow that has manual approval steps.

Which solution will meet these requirements?

Options:

A.

Load the data into an Amazon Simple Queue Service (Amazon SQS) queue that has a retention period of 2 days. Use an Amazon EventBridge pipe to retrieve data from the queue, transform the data, and pass the data to an AWS Step Functions workflow.

B.

Load the data into AWS DataSync. Delete the DataSync task after 2 days. Invoke an AWS

Lambda function to retrieve the data, transform the data, and invoke a second Lambda function that performs the remaining workflow steps.

C.

Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic, transform the data, and send the data to Amazon EC2 instances to perform the remaining workflow steps.

D.

Load the data into an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge pipe to retrieve the data from the topic and transform the data into an appropriate format for an Amazon SQS queue. Use an AWS Lambda function to poll the queue to perform the remaining workflow steps.

Answer

A

Explanation

Amazon SQS with a 2-day retention ensures the data lives just as long as needed. EventBridge Pipes allow direct integration between event producers and consumers, with optional filtering and transformation. AWS Step Functions supports manual approval steps, which fits the workflow requirement perfectly.

[Reference: AWS Documentation – Amazon EventBridge Pipes, AWS Step Functions, =====, ,]

Questions # 20:

A company recently migrated a monolithic application to an Amazon EC2 instance and Amazon RDS. The application has tightly coupled modules. The existing design of the application gives the application the ability to run on only a single EC2 instance.

The company has noticed high CPU utilization on the EC2 instance during peak usage times. The high CPU utilization corresponds to degraded performance on Amazon RDS for read requests. The company wants to reduce the high CPU utilization and improve read request performance.

Which solution will meet these requirements?

Options:

A.

Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Configure an RDS read replica for read requests.

B.

Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Add an RDS read replica and redirect all read/write traffic to the replica.

C.

Configure an Auto Scaling group with a minimum size of 1 and maximum size of 2. Resize the RDS DB instance to an instance type that has more CPU capacity.

D.

Resize the EC2 instance to an EC2 instance type that has more CPU capacity. Configure an Auto Scaling group with a minimum and maximum size of 1. Resize the RDS DB instance to an instance type that has more CPU capacity.

Answer

A

Explanation

To address the high CPU utilization on the EC2 instance and the degraded performance of Amazon RDS for read requests, the solution involves two key actions: resizing the EC2 instance and leveraging Amazon RDS read replicas.

Resizing the EC2 Instance: The first step is to resize the EC2 instance to a type with more CPU capacity to handle the higher computational demands during peak usage times. This helps to alleviate the immediate pressure on the CPU.

Auto Scaling Group with a Size of 1: Although the application can only run on a single EC2 instance due to its monolithic nature, creating an Auto Scaling group with a minimum and maximum size of 1 ensures that the instance is automatically restarted or replaced in case of failure, maintaining high availability.

RDS Read Replica: Configuring an RDS read replica allows the application to offload read requests to a separate instance, thus reducing the load on the primary RDS instance. This improves the performance of read operations, which were previously bottlenecked due to the high CPU usage on the EC2 instance.

Why Not Other Options?:

Option B: Redirecting all traffic to the RDS read replica is not recommended because replicas are meant for read traffic only, not for write operations. This could lead to data consistency issues.

Option C: Increasing the RDS instance type capacity helps, but it doesn't address the high CPU usage on the EC2 instance, nor does it provide a solution for scaling reads.

Option D: While resizing both the EC2 and RDS instances increases their capacities, it doesn't address the specific need to offload read traffic from the primary RDS instance.

AWS References:

Amazon RDS Read Replicas- Explains how to create and use read replicas to offload read traffic from the primary database instance.

Resizing Your EC2 Instance- Guidance on resizing EC2 instances to meet workload demands.

Questions # 21:

A company uses AWS Cost Explorer to monitor its AWS costs. The company notices that Amazon Elastic Block Store (Amazon EBS) storage and snapshot costs increase every month. However, the company does not purchase additional EBS storage every month. The company wants to optimize monthly costs for its current storage usage.

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

Options:

A.

Use logs in Amazon CloudWatch Logs to monitor the storage utilization of Amazon EBS. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.

B.

Use a custom script to monitor space usage. Use Amazon EBS Elastic Volumes to reduce the size of the EBS volumes.

C.

Delete all expired and unused snapshots to reduce snapshot costs.

D.

Delete all nonessential snapshots. Use Amazon Data Lifecycle Manager to create and manage the snapshots according to the company's snapshot policy requirements.

Answer

D

Explanation

Amazon Data Lifecycle Manager (DLM) automates the creation, retention, and deletion of EBS snapshots. This allows organizations to define policies that ensure snapshots are only kept as long as needed, reducing costs automatically and minimizing manual effort. AWS recommends using DLM for optimizing storage and managing backup lifecycle with minimal overhead.

[Reference: AWS Documentation - Amazon Data Lifecycle Manager, =====, ,]

Questions # 22:

A company has an application that uses a MySQL database that runs on an Amazon EC2 instance. The instance currently runs in a single Availability Zone. The company requires a fault-tolerant database solution that provides a recovery time objective (RTO) and a recovery point objective (RPO) of 2 minutes or less. Which solution will meet these requirements?

Options:

A.

Migrate the MySQL database to Amazon RDS. Create a read replica in a second Availability Zone. Create a script that detects availability interruptions and promotes the read replica when needed.

B.

Migrate the MySQL database to Amazon RDS for MySQL. Configure the new RDS for MySQL database to use a Multi-AZ deployment.

C.

Create a second MySQL database in a second Availability Zone. Use native MySQL commands to sync the two databases every 2 minutes. Create a script that detects availability interruptions and promotes the second MySQL database when needed.

D.

Create a copy of the EC2 instance that runs the MySQL database. Deploy the copy in a second Availability Zone. Create a Network Load Balancer. Add both instances as targets.

Answer

B

Explanation

Amazon RDS Multi-AZ deployments provide automatic failover for relational databases such as MySQL, ensuring high availability and durability. The feature maintains synchronous replication between a primary DB instance and a standby in a separate Availability Zone. AWS guarantees that failover typically completes within minutes, ensuring an RTO and RPO of less than 2 minutes. Option A requires manual promotion of replicas, which cannot meet the strict RTO/RPO requirement. Option C depends on custom scripts and manual synchronization, introducing operational risk. Option D creates active-active EC2-based databases, which do not provide synchronous replication or automated failover. Therefore, Multi-AZ RDS (B) is the managed, resilient, and operationally efficient solution that meets the business requirements.

[References: • Amazon RDS User Guide — Multi-AZ deployments • AWS Well-Architected Framework — Reliability Pillar: High availability and disaster recovery, , , ,]

Questions # 23:

A company is designing a new Amazon Elastic Kubernetes Service (Amazon EKS) deployment to host multi-tenant applications that use a single cluster. The company wants to ensure that each pod has its own hosted environment. The environments must not share CPU, memory, storage, or elastic network interfaces.

Which solution will meet these requirements?

Options:

A.

Use Amazon EC2 instances to host self-managed Kubernetes clusters. Use taints and tolerations to enforce isolation boundaries.

B.

Use Amazon EKS with AWS Fargate. Use Fargate to manage resources and to enforce isolation boundaries.

C.

Use Amazon EKS and self-managed node groups. Use taints and tolerations to enforce isolation boundaries.

D.

Use Amazon EKS and managed node groups. Use taints and tolerations to enforce isolation boundaries.

Answer

B

Explanation

AWS Fargate provides per-pod isolation for CPU, memory, storage, and networking, making it ideal for multi-tenant use cases.

AWS Documentation References:

EKS with Fargate

Questions # 24:

A company is using AWS Identity and Access Management (IAM) Access Analyzer to refine IAM permissions for employee users. The company uses an organization in AWS Organizations and AWS Control Tower to manage its AWS accounts. The company has designated a specific member account as an audit account.

A solutions architect needs to set up IAM Access Analyzer to aggregate findings from all member accounts in the audit account.

What is the first step the solutions architect should take?

Options:

A.

Use AWS CloudTrail to configure one trail for all accounts. Create an Amazon S3 bucket in the audit account. Configure the trail to send logs related to access activity to the new S3 bucket in the audit account.

B.

Configure a delegated administrator account for IAM Access Analyzer in the AWS Control Tower management account. In the delegated administrator account for IAM Access Analyzer, specify the AWS account ID of the audit account.

C.

Create an Amazon S3 bucket in the audit account. Generate a new permissions policy, and add a service role to the policy to give IAM Access Analyzer access to AWS CloudTrail and the

S3 bucket in the audit account.

D.

Add a new trust policy that includes permissions to allow IAM Access Analyzer to perform sts:AssumeRole actions. Modify the permissions policy to allow IAM Access Analyzer to generate policies.

Answer

B

Explanation

The first step is to configure a delegated administrator account for IAM Access Analyzer at the organization level. Only after delegating the administrator account can you aggregate Access Analyzer findings from all member accounts into a designated audit account. This must be set up in the AWS Organizations management account.

AWS Documentation Extract:

“You must designate a delegated administrator for IAM Access Analyzer at the organization level. The delegated administrator account aggregates findings from all member accounts.”

(Source: IAM Access Analyzer documentation)

A, C, D: These steps do not establish the organization-wide aggregation required for Access Analyzer.

[Reference: AWS Certified Solutions Architect – Official Study Guide, Access Analyzer Delegation., ,]

Questions # 25:

A company uses Amazon EC2 instances and stores data on Amazon Elastic Block Store (Amazon EBS) volumes. The company must ensure that all data is encrypted at rest by using AWS Key Management Service (AWS KMS). The company must be able to control rotation of the encryption keys.

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

Options:

A.

Create a customer managed key Use the key to encrypt the EBS volumes.

B.

Use an AWS managed key to encrypt the EBS volumes. Use the key to configure automatic key rotation.

C.

Create an external KMS key with imported key material. Use the key to encrypt the EBS volumes.

D.

Use an AWS owned key to encrypt the EBS volumes.

Answer

A

Explanation

To meet the requirement of controlling key rotation with minimal operational overhead, creating a customer managed key (CMK) in AWS KMS is the optimal solution. With CMKs, you can define custom key rotation policies, ensuring that you retain control over the key lifecycle, including enabling automatic key rotation every year.

Key AWS features:

Custom Key Management: A customer managed key allows you to control the key policies, lifecycle, and enable key rotation for compliance.

Least Operational Overhead: Using a customer managed key simplifies encryption management while offering more flexibility than AWS managed or owned keys.

AWS Documentation: The AWS Well-Architected Framework recommends customer managed keys for environments where key control and flexibility are required.

Questions # 26:

A healthcare provider is planning to store patient data on AWS as PDF files. To comply with regulations, the company must encrypt the data and store the files in multiple locations. The data must be available for immediate access from any environment.

Options:

A.

Store the files in an Amazon S3 bucket. Use the Standard storage class. Enable server-side encryption with Amazon S3 managed keys (SSE-S3) on the bucket. Configure cross-Region replication on the bucket.

B.

Store the files in an Amazon Elastic File System (Amazon EFS) volume. Use an AWS KMS managed key to encrypt the EFS volume. Use AWS DataSync to replicate the EFS volume to a second AWS Region.

C.

Store the files in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Backup to back up the volume on a regular schedule. Use an AWS KMS key to encrypt the backups.

D.

Store the files in an Amazon S3 bucket. Use the S3 Glacier Flexible Retrieval storage class. Ensure that all PDF files are encrypted by using client-side encryption before the files are uploaded. Configure cross-Region replication on the bucket.

Answer

A

Explanation

AmazonS3 with the Standard storage class is the best solution:

Encryption: SSE-S3 ensures server-side encryption of the data, meeting compliance requirements.

Immediate access: The Standard storage class provides low-latency and high-throughput access to data.

Multi-location storage: Cross-Region replication ensures data is stored in multiple AWS Regions for redundancy.

Why Other Options Are Not Ideal:

Option B:

Amazon EFS is more costly and suited for file systems rather than object storage. Not cost-effective.

Option C:

Amazon EBS is block storage and not optimized for object storage like PDFs. Backup schedules do not ensure immediate availability. Not suitable.

Option D:


S3 Glacier Flexible Retrieval is designed for archival, not immediate access. Does not meet access requirements.

AWS References:

Amazon S3 Standard Storage: [AWS Documentation - S3 Storage Classes](#)

Amazon S3 Cross-Region Replication: [AWS Documentation - Cross-Region Replication](#)

AWS Encryption Options: [AWS Documentation - S3 Encryption](#)



To Get Premium Files for SAA-C03 Visit

<https://www.certsmania.com/amazon-web-services/saa-c03-practice>

For More Free Questions Visit

<https://www.certsmania.com/amazon-web-services/pdf/saa-c03>



CertsMania