



CertsMania

Free Questions for SAA-C03

Shared by **Helen** on **Dec 6, 2025**

For More Free Questions and Preparation Resources

Check the Links on Last Page



CertsMania

Questions # 1:

A company needs to run a critical data processing workload that uses a Python script every night. The workload takes 1 hour to finish.

Which solution will meet these requirements MOST cost-effectively?

Options:

A.

Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the AWS Fargate launch type. Use the Fargate Spot capacity provider. Schedule the job to run once every night.

B.

Deploy an Amazon Elastic Container Service (Amazon ECS) cluster with the Amazon EC2 launch type. Schedule the job to run once every night.

C.

Create an AWS Lambda function that uses the existing Python code. Configure Amazon EventBridge to invoke the function once every night.

D.

Create an Amazon EC2 On-Demand Instance that runs Amazon Linux. Migrate the Python script to the instance. Use a cron job to schedule the script. Create an AWS Lambda function to start and stop the instance once every night.

Answer

A

Explanation

AWS Fargate with Spot capacity is the most cost-effective and serverless container option for short-duration jobs that are flexible on start time. Since the job is not latency-critical and runs nightly, it is a good candidate for Fargate Spot, which can offer up to 70% cost savings over On-Demand pricing.

The job runs for 1 hour, which exceeds the AWS Lambda maximum execution time (15 minutes). Therefore, Lambda is not suitable. EC2-based solutions involve higher operational overhead and cost, making Fargate Spot the best low-cost, low-maintenance solution.

Questions # 2:

A company runs a latency-sensitive gaming service in the AWS Cloud. The gaming service runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). An Amazon DynamoDB table stores the gaming data. All the infrastructure is in a single AWS Region. The main user base is in that same Region.

A solutions architect needs to update the architecture to support a global expansion of the gaming service. The gaming service must operate with the least possible latency.

Which solution will meet these requirements?

Options:

A.

Create an Amazon CloudFront distribution in front of the ALB.

B.

Deploy an Amazon API Gateway regional API endpoint. Integrate the API endpoint with the ALB.

C.

Create an accelerator in AWS Global Accelerator. Add a listener. Configure the endpoint to point to the ALB.

D.

Deploy the ALB and the fleet of EC2 instances to another Region. Use Amazon Route 53 with geolocation routing.

Answer

C

Explanation

For latency-sensitive, global, non-cacheable workloads such as online gaming, AWS recommends AWS Global Accelerator. Global Accelerator:

Provides static anycast IP addresses at AWS edge locations worldwide.

Routes traffic over the AWS global network to the application's ALB endpoint in the optimal Region, reducing latency and jitter versus internet-based routing.

Supports TCP and UDP and is ideal for real-time gaming traffic.

CloudFront (Option A) is optimized for caching HTTP/HTTPS content and does not significantly improve latency for highly dynamic, interactive gaming traffic.

API Gateway (Option B) is not needed in front of an ALB for this gaming pattern and introduces extra hops and cost.

Geolocation routing (Option D) with multiple Regions adds complexity (including data replication for DynamoDB) and still relies on DNS.

Questions # 3:

A company is building a stock trading application in the AWS Cloud. The company requires a highly available solution that provides low-latency access to block storage across multiple Availability Zones.

Options:

A.

Use an Amazon S3 bucket and an S3 File Gateway as shared storage for the application.

B.

Create an Amazon EC2 instance in each Availability Zone. Attach a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume to each EC2 instance. Create a Bash script to sync data between volumes.

C.

Use an Amazon FSx for NetApp ONTAP Multi-AZ file system to access data by using the iSCSI protocol.

D.

Create an Amazon EC2 instance in each Availability Zone. Attach a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume to each EC2 instance. Create a Python script to sync data between volumes.

Answer

C

Explanation

Amazon FSx for NetApp ONTAP supports Multi-AZ, providing automatic failover between Availability Zones for high availability. It exposes ONTAP LUNs over the iSCSI protocol, delivering shared block storage semantics with low latency and consistent performance to EC2 clients across AZs. This meets the requirement for “highly available” and “low-latency” block access across multiple AZs. S3/S3 File Gateway (A) is object/file, not block storage. EBS (B, D) provides block storage to a single instance in a single AZ; EBS volumes cannot be shared across instances/AZs, and host-side sync scripts add latency, complexity, and do not provide true HA. FSx for ONTAP natively provides synchronous HA pair replication, fast failover, and supports iSCSI multipathing for resilient, performant access suited to latency-sensitive trading workloads.

[References: Amazon FSx for NetApp ONTAP — Multi-AZ file systems; iSCSI LUNs and host connectivity; High availability and failover behavior; Performance and client access guidance., , ,]

Questions # 4:

A company uses AWS to host a public website. The load on the web servers recently increased. The company wants to learn more about the traffic flow and traffic sources. The company also wants to increase the overall security of the website.

Which solution will meet these requirements?

Options:

A.

Deploy AWS WAF and set up logging. Use Amazon Data Firehose to deliver the log files to an Amazon S3 bucket for analysis.

B.

Deploy Amazon API Gateway and set up logging. Use Amazon Kinesis Data Streams to deliver the log files to an Amazon S3 bucket for analysis.

C.

Deploy a Network Load Balancer and set up logging. Use Amazon Data Firehose to deliver the log files to an Amazon S3 bucket for analysis.

D.

Deploy an Application Load Balancer and set up logging. Use Amazon Kinesis Data Streams to deliver the log files to an Amazon S3 bucket for analysis.

Answer

A

Explanation

AWS WAF (Web Application Firewall) is designed to protect public-facing web applications from common web exploits and allows for deep inspection of HTTP and HTTPS requests. By enabling logging on AWS WAF, you can gain insights into traffic flow, request sources, and blocked requests, which improves both security visibility and posture. Integrating AWS WAF logging with Amazon Kinesis Data Firehose allows automatic, near-real-time delivery of log data to Amazon S3 for analysis, reporting, or integration with analytics tools. This solution not only secures the website but also enables comprehensive traffic analysis, satisfying both requirements efficiently.

Reference Extract from AWS Documentation / Study Guide:

"AWS WAF provides detailed logs of web requests, which can be delivered to Amazon S3 via Amazon Kinesis Data Firehose. This logging enables analysis of web traffic and sources, supporting both security and operational monitoring."

Source: AWS Certified Solutions Architect – Official Study Guide, Security and Compliance section; AWS WAF Developer Guide (Logging and Monitoring).

Questions # 5:

A company is planning to migrate customer records to an Amazon S3 bucket. The company needs to ensure that customer records are protected against unauthorized access and are encrypted in transit and at rest. The company must monitor all access to the S3 bucket.

Options:

A.

Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an S3 bucket policy that includes the `aws:SecureTransport` condition. Use an IAM policy to control access to the records. Use AWS CloudTrail to monitor access to the records.

B.

Use AWS Nitro Enclaves to encrypt customer records at rest. Use AWS Key Management Service (AWS KMS) to encrypt the records in transit. Use an IAM policy to control access to the records. Use AWS CloudTrail and AWS Security Hub to monitor access to the records.

C.

Use AWS Key Management Service (AWS KMS) to encrypt customer records at rest. Create an Amazon Cognito user pool to control access to the records. Use AWS CloudTrail to monitor access to the records. Use Amazon GuardDuty to detect threats.

○ D.

Use server-side encryption with Amazon S3 managed keys (SSE-S3) with default settings to encrypt the records at rest. Access the records by using an Amazon CloudFront distribution that uses the S3 bucket as the origin. Use IAM roles to control access to the records. Use Amazon CloudWatch to monitor access to the records.

Answer

A

Explanation

Encryption at Rest: AWS Key Management Service (AWS KMS) provides centralized control over the cryptographic keys used to protect data. By using AWS KMS with Amazon S3, you can manage encryption keys and define policies to control access to them.

Encryption in Transit: By enforcing the `aws:SecureTransport` condition in the S3 bucket policy, you ensure that all data is transmitted over HTTPS, protecting data in transit.

Access Control: IAM policies allow you to define fine-grained permissions for users and roles, ensuring that only authorized entities can access the customer records.

Monitoring: AWS CloudTrail provides a record of actions taken by a user, role, or AWS service in Amazon S3, enabling you to monitor access to the records.

[References:, Protecting data using encryption, Using IAM policies to control access to Amazon S3 resources, Logging Amazon S3 API calls using AWS CloudTrail, ,]

Questions # 6:

The customers of a finance company request appointments with financial advisors by sending text messages. A web application that runs on Amazon EC2 instances accepts the appointment requests. The text messages are published to an Amazon Simple Queue Service (Amazon SQS) queue through the web application. Another application that runs on EC2 instances then sends meeting invitations and meeting confirmation email messages to the customers. After successful scheduling, this application stores the meeting information in an Amazon DynamoDB database.

As the company expands, customers report that their meeting invitations are taking longer to arrive.

What should a solutions architect recommend to resolve this issue?

Options:

A.

Add a DynamoDB Accelerator (DAX) cluster in front of the DynamoDB database.

B.

Add an Amazon API Gateway API in front of the web application that accepts the appointment requests.

C.

Add an Amazon CloudFront distribution. Set the origin as the web application that accepts the appointment requests.

D.

Add an Auto Scaling group for the application that sends meeting invitations. Configure the Auto Scaling group to scale based on the depth of the SQS queue.

Answer

D

Explanation

For SQS-backed worker architectures, AWS recommends scaling consumers based on queue metrics (e.g., `ApproximateNumberOfMessagesVisible`, `AgeOfOldestMessage`). EC2 Auto Scaling can “scale out based on Amazon CloudWatch alarms” tied to SQS backlog, increasing worker capacity to reduce latency as demand grows. DAX (A) accelerates DynamoDB reads, not message processing. API Gateway (B) and CloudFront (C) optimize request ingress and content delivery, not the asynchronous email-sending application. The bottleneck is consumer throughput; scaling workers with an SQS-driven policy restores timely processing without downtime and follows Well-Architected patterns for decoupled, elastic systems.

[References: Amazon SQS Developer Guide — “Scaling consumers,” “Queue metrics”; EC2 Auto Scaling — “Target tracking and step scaling with CloudWatch metrics”; Well-Architected — Performance Efficiency (queue-based load leveling).,]

Questions # 7:

A company uses a set of Amazon EC2 instances to host a website. The website uses an Amazon S3 bucket to store images and media files.

The company wants to automate website infrastructure creation to deploy the website to

multiple AWS Regions. The company also wants to provide the EC2 instances access to the S3 bucket so the instances can store and access data by using AWS Identity and Access Management (IAM).

Which solution will meet these requirements MOST securely?

Options:

A.

Create an AWS CloudFormation template for the web server EC2 instances. Save an IAM access key in the UserData section of the AWS::EC2::Instance entity in the CloudFormation template.

B.

Create a file that contains an IAM secret access key and access key ID. Store the file in a new S3 bucket. Create an AWS CloudFormation template. In the template, create a parameter to specify the location of the S3 object that contains the access key and access key ID.

C.

Create an IAM role and an IAM access policy that allows the web server EC2 instances to access the S3 bucket. Create an AWS CloudFormation template for the web server EC2 instances that contains an IAM instance profile entity that references the IAM role and the IAM access policy.

D.

Create a script that retrieves an IAM secret access key and access key ID from IAM and stores them on the web server EC2 instances. Include the script in the UserData section of the AWS::EC2::Instance entity in an AWS CloudFormation template.

Answer

C

Explanation

The most secure solution for allowing EC2 instances to access an S3 bucket is by using IAM roles. An IAM role can be created with an access policy that grants the required permissions (e.g., to read and write to the S3 bucket). The IAM role is then associated with the EC2 instances through an IAM instance profile.

By associating the role with the instances, the EC2 instances can securely assume the role and receive temporary credentials via the instance metadata service. This avoids the need to store credentials (such as access keys) on the instances or within the application, enhancing security and reducing the risk of credentials being exposed.

AWS CloudFormation can be used to automate the creation of the entire infrastructure, including EC2 instances, IAM roles, and associated policies.

AWS References:

IAM Roles for EC2 Instances outlines the use of IAM roles for secure access to AWS services.

AWS CloudFormation User Guide details how to create and manage resources using CloudFormation templates.

Why the other options are incorrect:

- A. Save IAM access key in UserData: This is insecure because it involves storing long-term credentials in the instance user data, which can be exposed.
- B. Store access keys in S3: This is also insecure, as it involves managing and distributing long-term credentials, which should be avoided.
- D. Retrieve access keys via a script: This approach is unnecessarily complex and less secure than using IAM roles, which provide temporary credentials automatically.

Questions # 8:

A company wants to standardize its Amazon Elastic Block Store (Amazon EBS) volume encryption strategy. The company also wants to minimize the cost and configuration effort required to operate the volume encryption check.

Which solution will meet these requirements?

Options:

A.

Write API calls to describe the EBS volumes and to confirm the EBS volumes are encrypted. Use Amazon EventBridge to schedule an AWS Lambda function to run the API calls.

B.

Write API calls to describe the EBS volumes and to confirm the EBS volumes are encrypted. Run the API calls on an AWS Fargate task.

C.

Create an AWS Identity and Access Management (IAM) policy that requires the use of tags on EBS volumes. Use AWS Cost Explorer to display resources that are not properly tagged. Encrypt the untagged resources manually.

○ D.

Create an AWS Config rule for Amazon EBS to evaluate if a volume is encrypted and to flag the volume if it is not encrypted.

Answer

D

Explanation



CertsMania

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a Config rule, you can automatically check whether your Amazon EBS volumes are encrypted and flag those that are not, with minimal cost and configuration effort.

AWS Config Rule: AWS Config provides managed rules that you can use to automatically check the compliance of your resources against predefined or custom criteria. In this case, you would create a rule to evaluate EBS volumes and determine if they are encrypted. If a volume is not encrypted, the rule will flag it, allowing you to take corrective action.

Operational Overhead: This approach significantly reduces operational overhead because once the rule is in place, it continuously monitors your EBS volumes for compliance, and there's no need for manual checks or custom scripting.

Why Not Other Options?:

Option A (Lambda with API calls and EventBridge): While this can work, it involves writing and maintaining custom code, which increases operational overhead compared to using a managed AWS Config rule.

Option B (API calls on Fargate): Running API calls on Fargate is more complex and costly compared to using AWS Config, which provides a simpler, managed solution.

Option C (IAM policy with Cost Explorer): This option does not directly enforce encryption compliance and involves manual intervention, making it less efficient and more prone to errors.

AWS References:

AWS Config Rules- Overview of AWS Config rules and how they can be used to evaluate resource configurations.

Amazon EBS Encryption- Information on how to manage and enforce encryption for EBS volumes.

A media company uses an Amazon CloudFront distribution to deliver content over the internet. The company wants only premium customers to have access to the media streams and file content. The company stores all content in an Amazon S3 bucket. The company also delivers content on demand to customers for a specific purpose, such as movie rentals or music downloads.

Which solution will meet these requirements?

Options:

A.

Generate and provide S3 signed cookies to premium customers

B.

Generate and provide CloudFront signed URLs to premium customers.

C.

Use origin access control (OAC) to limit the access of non-premium customers

D.

Generate and activate field-level encryption to block non-premium customers.

Answer

B

Explanation

CloudFront Signed URLs: These URLs allow you to provide limited access to content that is being served through an Amazon CloudFront distribution. Signed URLs can be generated to grant time-limited access to premium customers.

Content Restriction:

By using CloudFront signed URLs, you can control access to your media streams and file content stored in S3.

These URLs can be customized with an expiration time, ensuring that access is only available for a specific period, which is useful for scenarios like movie rentals or music downloads.

Security and Flexibility:

Signed URLs ensure that only authenticated users (premium customers) can access the

restricted content.

This approach integrates seamlessly with CloudFront and S3, providing an efficient way to manage access controls without additional overhead.

Operational Efficiency: Using CloudFront signed URLs leverages AWS managed services to handle the complexity of access control, reducing the need for custom implementation and maintenance.

[References:, Serving Private Content with Signed URLs and Signed Cookies, , , ,]

Questions # 10:

A company needs to create an AWS Lambda function that will run in a VPC in the company's primary AWS account. The Lambda function needs to access files that the company stores in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system, the solution must scale to meet the demand.

Which solution will meet these requirements MOST cost-effectively?

Options:

A.

Create a new EFS file system in the primary account. Use AWS DataSync to copy the contents of the original EFS file system to the new EFS file system.

B.

Create a VPC peering connection between the VPCs that are in the primary account and the secondary account.

C.

Create a second Lambda function in the secondary account that has a mount that is configured for the file system. Use the primary account's Lambda function to invoke the secondary account's Lambda function.

D.

Move the contents of the file system to a Lambda layer. Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

Answer

B

Explanation

Amazon EFS is a regional, elastic file system that “scales to petabytes” and can be accessed from “thousands of compute instances” concurrently. You can mount EFS across VPCs and across AWS accounts by providing network connectivity (e.g., VPC peering) to the EFS mount targets and allowing NFS (TCP 2049) in the mount target security group, optionally using EFS access points and a file system policy for cross-account access control. VPC peering has no hourly charge and minimal operational overhead, and EFS automatically scales as files are added—meeting the scalability and cost goals.

Option A duplicates data, incurs DataSync and extra storage costs, and adds sync lag.

Option C adds an extra Lambda hop and complexity without exposing a shared filesystem to the primary function.

Option D is impractical: Lambda layers are immutable artifacts with tight size limits (up to 50 MB compressed/250 MB uncompressed) and are not suited for dynamic, growing file sets.

[References: Amazon EFS User Guide — “Accessing EFS across VPCs and accounts,” “Mount targets and security groups,” “EFS access points,” and “EFS automatically scales to petabytes.”, ,]

Questions # 11:

A company has an on-premises SFTP file transfer solution. The company is migrating to the AWS Cloud to scale the file transfer solution and to optimize costs by using Amazon S3. The company's employees will use their credentials for the on-premises Microsoft Active Directory (AD) to access the new solution. The company wants to keep the current authentication and file access mechanisms.

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

Options:

A.

Configure an S3 File Gateway. Create SMB file shares on the file gateway that use the existing Active Directory to authenticate

B.

Configure an Auto Scaling group with Amazon EC2 instances to run an SFTP solution. Configure the group to scale up at 60% CPU utilization.

C.

Create an AWS Transfer Family server with SFTP endpoints Choose the AWS Directory Service option as the identity provider Use AD Connector to connect the on-premises Active Directory.

D.

Create an AWS Transfer Family SFTP endpoint. Configure the endpoint to use the AWS Directory Service option as the identity provider to connect to the existing Active Directory.

Answer

C

Explanation

AWS Transfer Family: This service provides fully managed support for file transfers directly into and out of Amazon S3 using the SFTP, FTPS, and FTP protocols.

SFTP Endpoints:

Set up an AWS Transfer Family server and configure SFTP endpoints to handle the file transfers.

This service is scalable and managed, reducing operational overhead compared to running an SFTP solution on EC2 instances.

Integration with Active Directory:

Choose the AWS Directory Service option as the identity provider for the Transfer Family server.

Use AD Connector to link the on-premises Active Directory with AWS, allowing employees to use their existing AD credentials to access the SFTP service.

Operational Efficiency: This solution leverages managed services for both file transfer and identity management, ensuring minimal changes to the current authentication mechanisms and reducing operational overhead.

[References:, AWS Transfer Family, AWS Directory Service and AD Connector, , , ,]

Questions # 12:

A company wants to create an Amazon EMR cluster that multiple teams will use. The company wants to ensure that each team's big data workloads can access only the AWS services that each team needs to interact with. The company does not want the workloads to have access to Instance Metadata Service Version 2 (IMDSv2) on the cluster's underlying EC2 instances.

Which solution will meet these requirements?

Options:

A.

Configure interface VPC endpoints for each AWS service that the teams need. Use the required interface VPC endpoints to submit the big data workloads.

B.

Create EMR runtime roles. Configure the cluster to use the runtime roles. Use the runtime roles to submit the big data workloads.

C.

Create an EC2 IAM instance profile that has the required permissions for each team. Use the instance profile to submit the big data workloads.

D.

Create an EMR security configuration that has the EnableApplicationScoped IAM Role option set to false. Use the security configuration to submit the big data workloads.

Answer

B

Explanation

EMR runtime roles allow fine-grained permissions per job, letting each team access only the services they are authorized to use. This isolates IAM permissions per workload and avoids exposing instance-level credentials through IMDSv2. Runtime roles improve security posture in multi-tenant EMR environments.

[Reference: AWS Documentation - EMR Runtime Roles and Access Isolation,
=====, ,]

Questions # 13:

A company hosts its applications in multiple private and public subnets in a VPC. The applications in the private subnets need to access an API. The API is available on the internet and is hosted in the company's on-premises data center. A solutions architect needs to establish connectivity for applications in the private subnets.

Which solution will meet these requirements MOST cost-effectively?

Options:

A.

Create a transit gateway to connect the VPC to the on-premises network. Use the transit gateway to route API calls from the private subnets to the on-premises data center.

B.

Create a NAT gateway in the public subnet of the VPC. Use the NAT gateway to allow the private subnets to access the API over the internet.

C.

Establish an AWS PrivateLink connection to connect the VPC to the on-premises network. Use PrivateLink to make API calls from the private subnets to the on-premises data center.

D.

Implement an AWS Site-to-Site VPN connection between the VPC and the on-premises data center. Use the VPN connection to make API calls from the private subnets to the on-premises data center.

Answer

D

Explanation

AWS Site-to-Site VPN is a cost-effective way to securely connect your on-premises data center with AWS resources. In this scenario:

Applications in private subnets require access to the API hosted in the on-premises data center.

A Site-to-Site VPN connection is a secure and cost-efficient option to route traffic between the VPC and on-premises resources.

Transit Gateway and PrivateLink are not cost-effective for this use case.

NAT Gateway only provides internet access for private subnets, which is not suitable for reaching an on-premises resource.

AWS Documentation Reference:

AWS Site-to-Site VPN

A company wants to implement new security compliance requirements for its development team to limit the use of approved Amazon Machine Images (AMIs).

The company wants to provide access to only the approved operating system and software for all its Amazon EC2 instances. The company wants the solution to have the least amount of lead time for launching EC2 instances.

Which solution will meet these requirements?

Options:

A.

Create a portfolio by using AWS Service Catalog that includes only EC2 instances launched with approved AMIs. Ensure that all required software is preinstalled on the AMIs. Create the necessary permissions for developers to use the portfolio.

B.

Create an AMI that contains the approved operating system and software by using EC2 Image Builder. Give developers access to that AMI to launch the EC2 instances.

C.

Create an AMI that contains the approved operating system Tell the developers to use the approved AMI Create an Amazon EventBridge rule to run an AWS Systems Manager script when a new EC2 instance is launched. Configure the script to install the required software from a repository.

D.

Create an AWS Config rule to detect the launch of EC2 instances with an AMI that is not approved. Associate a remediation rule to terminate those instances and launch the instances again with the approved AMI. Use AWS Systems Manager to automatically install the approved software on the launch of an EC2 instance.

Answer

A

Explanation

AWS Service Catalog is designed to allow organizations to manage a catalog of approved products (including AMIs) that users can deploy. By creating a portfolio that contains only EC2 instances launched with preapproved AMIs, the company can enforce compliance with the approved operating system and software for all EC2 instances. Service Catalog also streamlines the process of launching EC2 instances, reducing the lead time while ensuring that developers use only the approved configurations.

Option B (EC2 Image Builder): While EC2 Image Builder helps in creating and managing AMIs, it doesn't provide the enforcement mechanism that Service Catalog does.

Option C (EventBridge rule and Systems Manager script): This solution is reactive and involves more operational complexity compared to Service Catalog.

Option D (AWS Config rule): This option is reactive (it terminates non-compliant instances after launch) and introduces additional operational overhead.

AWS References:

AWS Service Catalog



CertsMania

Questions # 15:

A company runs a workload in an AWS Region. Users connect to the workload by using an Amazon API Gateway REST API.

The company uses Amazon Route 53 as its DNS provider and has created a Route 53 Hosted Zone. The company wants to provide unique and secure URLs for all workload users.

Which combination of steps will meet these requirements with the MOST operational efficiency? (Select THREE.)

Options:

A.

Create a wildcard custom domain name in the Route 53 hosted zone as an alias for the API Gateway endpoint.

B.

Use AWS Certificate Manager (ACM) to request a wildcard certificate that matches the custom domain in a second Region.

C.

Create a hosted zone for each user in Route 53. Create zone records that point to the API Gateway endpoint.

D.

Use AWS Certificate Manager (ACM) to request a wildcard certificate that matches the custom domain name in the same Region.

E.

Use API Gateway to create multiple API endpoints for each user.

F.

Create a custom domain name in API Gateway for the REST API. Import the certificate from AWS Certificate Manager (ACM).

Answer

A, D, F

Explanation



CertsMania

To provide unique, secure URLs efficiently:

A: Create a wildcard custom domain in Route 53 as an alias for the API Gateway endpoint.

D: Request a wildcard certificate in ACM in the same Region as API Gateway (certificates must be in the same Region as the API).

F: Create a custom domain name in API Gateway and attach the certificate.

“You can configure a custom domain name for your API Gateway API. To use a wildcard certificate, request it from ACM in the same Region as your API.”

— API Gateway Custom Domain Names

This combination provides secure wildcard URLs without creating separate endpoints or hosted zones per user.

Questions # 16:

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 # 17:

A company currently runs an on-premises stock trading application by using Microsoft Windows Server. The company wants to migrate the application to the AWS Cloud. The company needs to design a highly available solution that provides low-latency access to block storage across multiple Availability Zones. Which solution will meet these requirements with the LEAST implementation effort?

Options:

A.

Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon FSx for Windows File Server as shared storage between the two cluster nodes.

B.

Configure a Windows Server cluster that spans two Availability Zones on Amazon EC2 instances. Install the application on both cluster nodes. Use Amazon Elastic Block Store (Amazon EBS) General Purpose SSD (gp3) volumes as storage attached to the EC2 instances. Set up application-level replication to sync data from one EBS volume in one Availability Zone to another EBS volume in the second Availability Zone.

C.

Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use an Amazon FSx for NetApp ONTAP Multi-AZ file system to access the data by using Internet Small Computer Systems Interface (iSCSI) protocol.

D.

Deploy the application on Amazon EC2 instances in two Availability Zones. Configure one EC2 instance as active and the second EC2 instance in standby mode. Use Amazon Elastic Block Store (Amazon EBS) Provisioned IOPS SSD (io2) volumes as storage attached to the EC2 instances. Set up Amazon EBS level replication to sync data from one io2 volume in one Availability Zone to another io2 volume in the second Availability Zone.

Answer

A

Explanation

This solution is designed to provide high availability and low-latency access to block storage across multiple Availability Zones with minimal implementation effort.

Windows Server Cluster Across AZs: Configuring a Windows Server Failover Cluster (WSFC) that spans two Availability Zones ensures that the application can failover from one instance to another in case of a failure, meeting the high availability requirement.

Amazon FSx for Windows File Server: FSx for Windows File Server provides fully managed Windows file storage that is accessible via the SMB protocol, which is suitable for Windows-based applications. It offers high availability and can be used as shared storage between the cluster nodes, ensuring that both nodes have access to the same data with low latency.

Why Not Other Options?:

Option B (EBS with application-level replication): This requires complex configuration and management, as EBS volumes cannot be directly shared across AZs. Application-level replication is more complex and prone to errors.

Option C (FSx for NetApp ONTAP with iSCSI): While this is a viable option, it introduces additional complexity with iSCSI and requires more specialized knowledge for setup and management.

Option D (EBS with EBS-level replication): EBS-level replication is not natively supported across AZs, and setting up a custom replication solution would increase the implementation effort.

AWS References:

Amazon FSx for Windows File Server- Overview and benefits of using FSx for Windows File Server.

Windows Server Failover Clustering on AWS- Guide on setting up a Windows Server cluster on AWS.

Questions # 18:

Question:

A company operates an online photo-sharing service and stores data in AWS Account A in a centralized Amazon S3 bucket. The company wants to grant a second AWS account named Account B access to the centralized S3 bucket. The company owns Account B.

Options:

Options:

A.

Enable S3 Transfer Acceleration to provide Account B access to the centralized S3 bucket in Account A.

B.

Enable cross-Region replication between Account A and Account B to share the S3 bucket data.

C.

Use Amazon CloudFront to distribute the S3 bucket contents. Grant Account B access to the bucket contents through a signed URL.

D.

Create a bucket policy that grants Account B permission to access the centralized S3 bucket in Account A.

Answer

D

Explanation



CertsMania

To grant cross-account access to an Amazon S3 bucket, you can use a bucket policy that specifies the AWS account ID of the account you want to grant access to. This method allows Account B to access the S3 bucket in Account A without the need for additional services or configurations.

[Reference: Example 2: Bucket owner granting cross-account bucket permissions, ,]

Questions # 19:

A solutions architect has created an AWS Lambda function that makes queries to an Amazon Aurora MySQL DB instance. When the solutions architect performs a test, the DB instance shows an error for too many connections.

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

Options:

A.

Create a read replica for the DB instance. Query the replica DB instance instead of the primary DB instance.

B.

Migrate the data to an Amazon DynamoDB database.

C.

Configure the Amazon Aurora MySQL DB instance for Multi-AZ deployment.

D.

Create a proxy in Amazon RDS Proxy. Query the proxy instead of the DB instance.

Answer

D

Explanation

AWS RDS Proxy is a fully managed, highly available database proxy that allows applications to pool and share database connections efficiently.

In serverless architectures like Lambda, rapid invocations can open numerous concurrent connections to Aurora, potentially overwhelming the database and causing “too many connections” errors.

By using Amazon RDS Proxy, the solution:

Pools database connections.

Maintains warm connections that can be reused.

Supports IAM authentication and Secrets Manager integration.

Requires minimal application change and low operational effort.

This directly supports the Performance Efficiency pillar of the AWS Well-Architected Framework, ensuring the application scales without overloading the DB.

???? Reference:

[Amazon RDS Proxy Documentation](#)

[Lambda + RDS Best Practices](#)

Questions # 20:

A company is planning to migrate an on-premises online transaction processing (OLTP) database that uses MySQL to an AWS managed database management system. Several reporting and analytics applications use the on-premises database heavily on weekends and at the end of each month. The cloud-based solution must be able to handle read-heavy surges during weekends and at the end of each month.

Which solution will meet these requirements?

Options:

A.

Migrate the database to an Amazon Aurora MySQL cluster. Configure Aurora Auto Scaling to

use replicas to handle surges.

B.

Migrate the database to an Amazon EC2 instance that runs MySQL. Use an EC2 instance type that has ephemeral storage. Attach Amazon EBS Provisioned IOPS SSD (io2) volumes to the instance.

C.

Migrate the database to an Amazon RDS for MySQL database. Configure the RDS for MySQL database for a Multi-AZ deployment, and set up auto scaling.

D.

Migrate from the database to Amazon Redshift. Use Amazon Redshift as the database for both OLTP and analytics applications.

Answer

A

Explanation

A. Aurora MySQL:Handles OLTP workloads efficiently with built-in replication and auto-scaling capabilities.

B. EC2 with MySQL:Requires heavy manual maintenance and does not scale seamlessly.

C. RDS for MySQL:Limited in auto-scaling compared to Aurora.

D. Redshift:Primarily for OLAP, not suitable for OLTP workloads.

[References:Amazon Aurora, , , , ,]

Questions # 21:

A company's packaged application dynamically creates and returns single-use text files in response to user requests. The company is using Amazon CloudFront for distribution, but wants to further reduce data transfer costs. The company cannot modify the application's source code.

What should a solutions architect do to reduce costs?

Options:

A.

Use Lambda@Edge to compress the files as they are sent to users.

B.

Enable Amazon S3 Transfer Acceleration to reduce the response times.

C.

Enable caching on the CloudFront distribution to store generated files at the edge.

D.

Use Amazon S3 multipart uploads to move the files to Amazon S3 before returning them to users.

Answer

A

Explanation

Lambda@Edge allows you to run functions at CloudFront edge locations, enabling modifications to content before it's delivered to users, including compression — which directly reduces data transfer cost.

“Lambda@Edge can be used to compress or modify your content before delivering it to end users, which reduces the amount of data transferred.”

— Lambda@Edge Use Cases

Since code modifications aren't allowed, using Lambda@Edge is a non-invasive way to:

Compress responses.

Reduce transfer size = lower CloudFront cost.

Incorrect Options:

B: S3 Transfer Acceleration improves speed, not cost.

C: Caching doesn't help if content is always unique (single-use).

D: Multipart uploads help with large file uploads, not transfers.

[Reference:, Lambda@Edge for Content Compression, CloudFront Pricing and Cost Optimization, , ,]

A solutions architect needs to secure an Amazon API Gateway REST API. Users need to be able to log in to the API by using common external social identity providers (IdPs). The social IdPs must use standard authentication protocols such as SAML or OpenID Connect (OIDC). The solutions architect needs to protect the API against attempts to exploit application vulnerabilities.

Which combination of steps will meet these security requirements? (Select TWO.)

Options:

A.

Create an AWS WAF web ACL that is associated with the REST API. Add the appropriate managed rules to the ACL.

B.

Subscribe to AWS Shield Advanced. Enable DDoS protection. Associate Shield Advanced with the REST API.

C.

Create an Amazon Cognito user pool with a federation to the social IdPs. Integrate the user pool with the REST API.

D.

Create an API key in API Gateway. Associate the API key with the REST API.

E.

Create an IP address filter in AWS WAF that allows only the social IdPs. Associate the filter with the web ACL and the API.

Answer

A, C

Explanation

Step A: AWS WAF with managed rules protects the API against application-layer attacks, such as SQL injection and cross-site scripting (XSS).

Step C: Amazon Cognito provides secure authentication and supports federation with social IdPs using OIDC or SAML. It integrates seamlessly with API Gateway.

Option B: AWS Shield Advanced provides DDoS protection, which is not explicitly required in this scenario.

Option D:API keys provide identification, not authentication, and are insufficient for this use case.

Option E:IP filters in WAF are overly restrictive for federated authentication scenarios.

AWS Documentation References:

Amazon Cognito Federation

AWS WAF Managed Rules



CertsMania

Questions # 23:

A company runs a critical Amazon RDS for MySQL DB instance in a single Availability Zone. The company must improve the availability of the DB instance.

Which solution will meet this requirement?

Options:

A.

Configure the DB instance to use a multi-Region DB instance deployment.

B.

Create an Amazon Simple Queue Service (Amazon SQS) queue in the AWS Region where the company hosts the DB instance to manage writes to the DB instance.

C.

Configure the DB instance to use a Multi-AZ DB instance deployment.

D.

Create an Amazon Simple Queue Service (Amazon SQS) queue in a different AWS Region than the Region where the company hosts the DB instance to manage writes to the DB instance.

Answer

C

Explanation

To improve availability and fault tolerance of an Amazon RDS instance, the recommended approach is to configure a Multi-AZ deployment.

Multi-AZ deployments for RDS automatically replicate data to a standby instance in a different Availability Zone (AZ).

If a failure occurs in the primary AZ (due to hardware, network, or power), RDS will automatically failover to the standby instance with minimal downtime, without administrative intervention.

This is an AWS-managed feature and does not require application modification.

It does not provide scalability or load balancing; it's designed for high availability and resiliency.

Options A, B, and D are incorrect:

A refers to cross-Region, which is used for disaster recovery, not high availability.

B and D with SQS do not address high availability directly for the RDS instance; queues help decouple systems but do not make a database more resilient.

???? Reference:

Amazon RDS Multi-AZ Deployments

Questions # 24:

A company uses an AWS Transfer for SFTP public server endpoint and Amazon S3 storage to host large datasets for its customers. The company provides customers SSH private keys to authenticate and download their datasets. The Transfer for SFTP server is configured with structured logging that is saved to an S3 bucket. The company wants to charge customers based on their monthly data download usage. Which solution will meet these requirements?

Options:

A.

Configure VPC Flow Logs to write to a new S3 bucket. Run monthly queries on the flow logs to identify customer usage and calculate cost. Add the charges to the customers' monthly bills.

B.

Each month, use AWS Cost Explorer to examine the costs for Transfer for SFTP and obtain a breakdown by customer. Add the charges to the customers' monthly bills.

C.

Enable requester pays on the S3 bucket that hosts the software. Allocate the charges to each customer based on the customer's requests.

○ D.

Run Amazon Athena queries on the logging S3 bucket monthly to identify customer usage and calculate costs. Add the charges to the customers' monthly bills.

Answer

D

Explanation



CertsMania

Comprehensive and Detailed Step-by-Step Explanation:

To accurately charge customers based on their monthly data download usage, the following solution is recommended:

Structured Logging Configuration:

Action: Ensure that the AWS Transfer for SFTP server is configured to log user activity, including details about file downloads, to Amazon S3 in a structured format.

Implementation: Utilize AWS Transfer Family's structured logging feature to capture detailed information about user sessions, including actions performed and data transferred.

docs.aws.amazon.com

Justification: Structured logs provide comprehensive data necessary for analyzing customer-specific download activities.

Data Analysis with Amazon Athena:

Action: Use Amazon Athena to run SQL queries on the structured log data stored in the S3 bucket to calculate the amount of data each customer has downloaded.

Implementation:

a. Define a Schema: Create a table in Athena that maps to the structure of your log files. This involves specifying the format of the logs and the location in S3.

b. Query Data: Write SQL queries to sum the total bytes downloaded by each customer over the billing period. This can be achieved by filtering logs based on user identifiers and summing the data transfer amounts.

Justification: Athena allows for efficient querying

A company runs a critical three-tier web application that consists of multiple virtual machines (VMs) and virtual databases in an on-premises environment. The company wants to set up a disaster recovery (DR) environment in AWS.

The company requires a 15-minute recovery time objective (RTO). The company must be able to test the failover solution to validate the recovery. The solution must provide an automated failover mechanism.

Which solution will meet these requirements?

Options:

A.

Use AWS Backup to create backups of the on-premises VMs and to restore the backups in AWS. Configure recovery to Amazon EC2 instances to meet the RTO requirement.

B.

Use AWS Database Migration Service (AWS DMS) to replicate the on-premises databases to Amazon RDS. Set up AWS Storage Gateway for baseline and incremental data replication to AWS to meet the RTO requirement.

C.

Use AWS DataSync and AWS Storage Gateway to migrate the baseline and incremental data to AWS. Use Amazon EC2, Amazon S3, and an Application Load Balancer to set up the DR environment.

D.

Use AWS Elastic Disaster Recovery to replicate the VMs incrementally to AWS. Configure Elastic Disaster Recovery to automate the DR process.

Answer

D

Explanation

AWS Elastic Disaster Recovery (AWS DRS) enables fast, reliable, and cost-effective disaster recovery. It replicates on-premises machines to AWS using continuous block-level replication. It supports automated testing and failover, meeting aggressive RTO targets such as 15 minutes.

[Reference: AWS Documentation – Elastic Disaster Recovery Overview, [https://docs.aws.amazon.com/elasticdisasterrecovery/latest/ug/elastic-disaster-recovery-overview.html](#), ,]

Questions # 26:

A company stores petabytes of historical medical information on premises. The company has a process to manage encryption of the data to comply with regulations. The company needs a cloud-based solution for data backup, recovery, and archiving. The company must retain control over the encryption key material. Which combination of solutions will meet these requirements? (Select TWO.)

Options:

A.

Create an AWS Key Management Service (AWS KMS) key without key material. Import the company's key material into the KMS key.

B.

Create an AWS Key Management Service (AWS KMS) encryption key that contains key material generated by AWS KMS.

C.

Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage. Use S3 Bucket Keys with AWS Key Management Service (AWS KMS) keys.

D.

Store the data in an Amazon S3 Glacier storage class. Use server-side encryption with customer-provided keys (SSE-C).

E.

Store the data in AWS Snowball devices. Use server-side encryption with AWS KMS keys (SSE-KMS).

Answer

A, D

Explanation

Option A: Importing customer-managed keys into AWS KMS ensures that encryption key material remains under the company's control.

Option D: S3 Glacier with server-side encryption using customer-provided keys (SSE-C) complies with the need for controlled encryption and provides cost-effective storage for backups.

Questions # 27:

An ecommerce company hosts a three-tier web application in a VPC. The web tier runs on Amazon EC2 instances in two Availability Zones. The company stores a product catalog and customer sales information in Amazon DynamoDB.

The company's finance team uses a reporting application to generate reports of daily product sales. When the finance team runs the daily reports, a sudden performance decrease affects website customers.

The company wants to improve the performance of the system.

Which solution will meet these requirements with MINIMAL changes to the current architecture?

Options:

A.

Migrate the application to larger EC2 instances. Migrate the database to Amazon RDS for MySQL. Configure a read replica of the database in a second Availability Zone.

B.

Increase the compute capacity of the EC2 instances. Migrate the database to Amazon ElastiCache (Memcached).

C.

Implement DynamoDB Accelerator (DAX).

D.

Configure DynamoDB streams.

Answer

C

Explanation

DynamoDB Accelerator (DAX) is a fully managed, in-memory cache for DynamoDB that delivers up to a 10x performance improvement—even at millions of requests per second.

DAX requires minimal changes to existing applications and offloads the read workload from DynamoDB during report generation.

Reference Extract:

"DAX provides in-memory acceleration for DynamoDB tables, reducing response times from milliseconds to microseconds with minimal application changes."

Source: AWS Certified Solutions Architect - Official Study Guide, DynamoDB and Performance section.



CertsMania

Questions # 28:

A global company runs its workloads on AWS. The company's application uses Amazon S3 buckets across AWS Regions for sensitive data storage and analysis. The company stores millions of objects in multiple S3 buckets daily. The company wants to identify all S3 buckets that are not versioning-enabled.

Which solution will meet these requirements?

Options:

A.

Set up an AWS CloudTrail event that has a rule to identify all S3 buckets that are not versioning-enabled across Regions.

B.

Use Amazon S3 Storage Lens to identify all S3 buckets that are not versioning-enabled across Regions.

C.

Enable IAM Access Analyzer for S3 to identify all S3 buckets that are not versioning-enabled across Regions.

D.

Create an S3 Multi-Region Access Point to identify all S3 buckets that are not versioning-enabled across Regions.

Answer

B

Explanation

Amazon S3 Storage Lens:

S3 Storage Lens provides organization-wide visibility into object storage usage and activity trends.

It can generate metrics and insights about your S3 buckets, including versioning status.

Configuration:

Enable S3 Storage Lens at the organization level.

Configure the dashboard to include the versioning status metric.

Identify Non-Versioned Buckets:

Use the S3 Storage Lens dashboard to filter and identify buckets that do not have versioning enabled.

Storage Lens provides detailed insights and reports which can be used to enforce compliance and manage storage effectively.

Operational Efficiency: Using S3 Storage Lens provides a centralized, easy-to-use interface for monitoring bucket configurations across multiple Regions and accounts, reducing the need for custom scripts or manual checks.

[References:, Amazon S3 Storage Lens, S3 Storage Lens Metrics, , , ,]

Questions # 29:

A company is developing an application that uses an Amazon Aurora MySQL database. The company plans to regularly make changes to the MySQL database schema to test new features. The tests must not affect the existing production database.

When the company finishes testing, a developer needs to replicate the changes to the production database. The solution must cause minimal downtime.

Which solution will meet these requirements?

Options:

A.

Create a new staging Aurora MySQL database cluster based on the existing database. Make the schema changes to the new staging database cluster to test the new features.

B.

Create a read replica based on the existing Aurora MySQL database. Make the schema changes to the read replica. Promote the read replica to primary after successful testing.

C.

Create a blue/green deployment of the Aurora MySQL database. Make schema changes in the staging environment to test new features. Direct traffic from the green environment to the blue environment when testing is complete.

D.

Replicate the Aurora MySQL database to an Amazon DynamoDB table. Make the schema changes to the DynamoDB table to test the new features. Configure the application to use the DynamoDB table when testing is complete.

Answer

C

Explanation

Aurora blue/green deployments are specifically designed for:

Safely testing schema changes and database upgrades in an isolated environment (green) without impacting production (blue).

Performing a fast, low-downtime switchover once testing is complete and validated.

In this pattern:

The blue environment is the current production database.

The green environment is a synchronized copy used for testing changes.

You apply schema changes to the green environment, run tests, and when ready, perform a managed switchover that minimizes downtime and risk.

Why others are not ideal:

A: A separate staging cluster allows testing but does not provide automated, low-downtime synchronization and switchover.

B: Aurora read replicas are for read scaling; schema changes on replicas are not supported in the way described, and promotion alone doesn't solve controlled testing and replication of changes.

D: Moving to DynamoDB changes database engines and data models entirely, and does not match the requirement to keep using Aurora MySQL.

Questions # 30:

A company runs multiple web applications on Amazon EC2 instances behind a single Application Load Balancer (ALB). The application experiences unpredictable traffic spikes throughout each day. The traffic spikes cause high latency. The unpredictable spikes last less than 3 hours. The company needs a solution to resolve the latency issue caused by traffic spikes.

Options:

A.

Use EC2 instances in an Auto Scaling group. Configure the ALB and Auto Scaling group to use a target tracking scaling policy.

B.

Use EC2 Reserved Instances in an Auto Scaling group. Configure the Auto Scaling group to use a scheduled scaling policy based on peak traffic hours.

C.

Use EC2 Spot Instances in an Auto Scaling group. Configure the Auto Scaling group to use a scheduled scaling policy based on peak traffic hours.

D.

Use EC2 Reserved Instances in an Auto Scaling group. Replace the ALB with a Network Load Balancer (NLB).

Answer

A

Explanation

AWS recommends Auto Scaling with dynamic scaling policies to handle unpredictable workload spikes. Target tracking scaling policies automatically adjust capacity based on defined metrics such as average CPU utilization or request count per target. This approach ensures applications maintain responsiveness without overprovisioning. Scheduled scaling (options B and C) is only effective when traffic patterns are predictable, which is not the case here. Reserved Instances or Spot Instances also do not address sudden demand

changes effectively. Replacing the ALB with an NLB (D) does not solve latency caused by EC2 instance capacity. Therefore, using EC2 instances in an Auto Scaling group with a target tracking scaling policy is the most effective and cost-efficient solution for unpredictable, short-lived traffic spikes.

[References: • Amazon EC2 Auto Scaling User Guide — Target tracking scaling policies • AWS Well-Architected Framework — Performance Efficiency Pillar: Elasticity and scaling, ,]

Questions # 31:

A company runs an HPC workload that uses a 200-TB file system on premises. The company needs to migrate this data to Amazon FSx for Lustre. Internet capacity is 10 Mbps, and all data must be migrated within 30 days.

Which solution will meet this requirement?

Options:

A.

Use AWS DMS to transfer data into S3 and link FSx for Lustre to the bucket.

B.

Deploy AWS DataSync on premises and transfer directly into FSx for Lustre.

C.

Use AWS Storage Gateway Volume Gateway to move data into FSx for Lustre.

D.

Use an AWS Snowball Edge storage-optimized device to transfer data into S3 and link FSx for Lustre to the bucket.

Answer

D

Explanation

At 10 Mbps, the maximum transferable data in 30 days is far below 200 TB, making any online transfer (Options A, B, C) impossible within the time window.

AWS Snowball Edge storage-optimized devices support high-speed, offline bulk data migration of large datasets. Once the data is delivered to S3, FSx for Lustre can be linked to the S3 bucket to populate the Lustre filesystem.

DataSync cannot meet the time constraint over 10 Mbps. Storage Gateway is not designed for large-scale migrations.

=====

Questions # 32:

A company is creating a mobile financial app that gives users the ability to sign up and store personal information. The app uses an Amazon DynamoDB table to store user details and preferences.

The app generates a credit score report by using the data that is stored in DynamoDB. The app sends credit score reports to users once every month.

The company needs to provide users with an option to remove their data and preferences. The app must delete customer data within one month of receiving a request to delete the data.

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

Options:

A.

Create an AWS Lambda function to delete user information. Create an Amazon EventBridge rule that runs when a specified TTL expires. Configure the EventBridge rule to invoke the Lambda function.

B.

Create a DynamoDB stream. Create an AWS Lambda function to delete user information. When a specified TTL expires, write user information to the DynamoDB stream from the DynamoDB table. Configure the DynamoDB stream to invoke the Lambda function to delete user information.

C.

Enable TTL in DynamoDB. Set the expiration date as an attribute. Create an AWS Lambda function to set the TTL based on the expiration date value. Invoke the Lambda function when a user requests to delete personal data.

D.

Enable TTL in DynamoDB. Create an AWS Lambda function to delete user information. Configure AWS Config to detect the DynamoDB state change when TTL expires and to invoke the Lambda function.

Answer

C

Explanation

Amazon DynamoDB supports Time to Live (TTL), which automatically and asynchronously deletes expired items based on a timestamp attribute. TTL is ideal for automatic data expiration with very low operational overhead.

In this scenario:

When a user requests deletion, the system can calculate an expiration timestamp one month in the future.

A Lambda function is used once per request to write or update the item's TTL attribute to that timestamp (Option C).

DynamoDB TTL then automatically removes the item after the expiration time. Deletion typically occurs within 48 hours of the TTL timestamp, which satisfies the "within one month" requirement.

This approach offloads the actual deletion to DynamoDB and avoids building complex orchestration or periodic cleanup jobs.

Options A, B, and D add unnecessary components (EventBridge, streams, AWS Config, and Lambda-based deletion workflows) on top of TTL or custom logic, increasing operational overhead without providing additional value for the given requirement.

Questions # 33:

A multinational company operates in multiple AWS Regions. The company must ensure that its developers and administrators have secure, role-based access to AWS resources.

The roles must be specific to each user's geographic location and job responsibilities.

The company wants to implement a solution to ensure that each team can access only resources within the team's Region. The company wants to use its existing directory service to manage user access. The existing directory service organizes users into roles based on location. The system must be capable of integrating seamlessly with multi-factor authentication (MFA).

Which solution will meet these requirements?

Options:

A.

Use AWS Security Token Service (AWS STS) to generate temporary access tokens. Integrate STS with the directory service. Assign Region-specific roles.

B.

Configure AWS IAM Identity Center with federated access. Integrate IAM Identity Center with the directory service to set up Region-specific IAM roles.

C.

Create IAM managed policies that restrict access by location. Apply policies based on group membership in the directory.

D.

Use custom Lambda functions to dynamically assign IAM policies based on login location and job function.

Answer

B

Explanation

IAM Identity Center (formerly AWS SSO) is designed for:

Federated access from external directories (e.g., Active Directory, Okta)

Centralized permission management

Support for MFA

Granular control via Attribute-based access control (ABAC)

“IAM Identity Center allows you to manage SSO access to AWS accounts and business applications centrally. You can assign users and groups permissions based on directory attributes such as Region and job role.”

— IAM Identity Center Docs

This option ensures:

Federated, centralized access

Region-specific permissions

MFA and role mapping via existing directory service

[References:, IAM Identity Center (SSO) Overview, Set Up Attribute-Based Access Control, ,]

Questions # 34:

A company runs several applications on Amazon EC2 instances. The company stores configuration files in an Amazon S3 bucket.

A solutions architect must provide the company's applications with access to the configuration files. The solutions architect must follow AWS best practices for security.

Which solution will meet these requirements?

Options:

A.

Use the AWS account root user access keys.

B.

Use the AWS access key ID and the EC2 secret access key.

C.

Use an IAM role to grant the necessary permissions to the applications.

D.

Activate multi-factor authentication (MFA) and versioning on the S3 bucket.

Answer

C

Explanation

The best security practice when providing EC2 instances access to AWS services (like S3) is to use an IAM role with an instance profile. This avoids hardcoding secrets and enables automatic credential rotation.

“We strongly recommend that you use IAM roles for applications that run on Amazon EC2 instances to securely access AWS services.”

— IAM Roles for Amazon EC2

Benefits:

No manual credentials

Temporary and automatically rotated keys

Least privilege access via IAM policies

Incorrect Options:

A: Root user access is not to be used for programmatic access.

B: Storing secret keys is insecure and discouraged.

D: MFA/versioning improves object protection, not access control.

[References:, Best Practices for IAM, Using IAM Roles with EC2, , ,]

Questions # 35:

A marketing company receives a large amount of new clickstream data in Amazon S3 from a marketing campaign. The company needs to analyze the clickstream data in Amazon S3 quickly. Then the company needs to determine whether to process the data further in the data pipeline.

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

Options:

A.

Create external tables in a Spark catalog. Configure jobs in AWS Glue to query the data.

B.

Configure an AWS Glue crawler to crawl the data. Configure Amazon Athena to query the data.

C.

Create external tables in a Hive metastore. Configure Spark jobs in Amazon EMR to query the data.

D.

Configure an AWS Glue crawler to crawl the data. Configure Amazon Kinesis Data Analytics to use SQL to query the data.

Answer

B

Explanation

AWS Glue Crawler: AWS Glue is a fully managed ETL (Extract, Transform, Load) service that makes it easy to prepare and load data for analytics. A Glue crawler can automatically discover new data and schema in Amazon S3, making it easy to keep the data catalog up-to-date.

Crawling the Data:

Set up an AWS Glue crawler to scan the S3 bucket containing the clickstream data.

The crawler will automatically detect the schema and create/update the tables in the AWS Glue Data Catalog.

Amazon Athena:

Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL.

Once the data catalog is updated by the Glue crawler, use Athena to query the clickstream data directly in S3.

Operational Efficiency: This solution leverages fully managed services, reducing operational overhead. Glue crawlers automate data cataloging, and Athena provides a serverless, pay-per-query model for quick data analysis without the need to set up or manage infrastructure.

[References:, AWS Glue, Amazon Athena, , , ,]

Questions # 36:

A company wants to use an API to translate text from one language to another. The API must receive an HTTP header value and pass the value to an embedded library. The API translates documents in 6 minutes. The API requires a custom authorization mechanism.

Options:

A.

Configure an Amazon API Gateway REST API with AWS_PROXY integration to synchronously call an AWS Lambda function to perform translations.

B.

Configure an AWS Lambda function with a Lambda function URL to synchronously call a second function to perform translations.

C.

Configure an Amazon API Gateway REST API with AWS_PROXY integration to asynchronously call an AWS Lambda function to perform translations.

D.

Configure an Amazon API Gateway REST API with HTTP_PROXY integration to synchronously call a web endpoint that is hosted on an EC2 instance.

Answer

A

Explanation

The AWS_PROXY integration with Amazon API Gateway allows the API to invoke a Lambda function synchronously, making it a suitable solution for the custom authorization mechanism and text translation use case.

Synchronous Invocation: The API Gateway REST API with AWS_PROXY integration enables synchronous processing of HTTP requests and responses, which is required for document translation.

Custom Authorization: API Gateway supports custom authorizers for fine-grained access control.

Lambda Function Execution: Although Lambda's execution time limit is 15 minutes, this is sufficient for the 6-minute document translation requirement.

Why Other Options Are Not Ideal:

Option B:

Introducing a Lambda function URL to invoke another Lambda function unnecessarily complicates the architecture. Not efficient.

Option C:

Asynchronous invocation cannot guarantee real-time response delivery for document translation tasks. Not suitable.

Option D:

Hosting the API on an EC2 instance increases operational overhead. HTTP_PROXY integration does not add significant benefits here. Not cost-effective or efficient.

AWS References:

API Gateway Lambda Proxy Integration:AWS Documentation - Proxy Integration

Custom Authorization in API Gateway:AWS Documentation - Custom Authorization



CertsMania



CertsMania

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